

Christie, Iain (2026) *Balancing media immersion and reality awareness for passengers through augmented reality TV ornamentations.* PhD thesis.

https://theses.gla.ac.uk/86374/

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses
https://theses.gla.ac.uk/
research-enlighten@glasgow.ac.uk

Balancing Media Immersion and Reality Awareness for Passengers through Augmented Reality TV Ornamentations

Iain Christie

Submitted in fulfilment of the requirements for the
Degree of Doctor of Philosophy

School of Computer Science
College of Science and Engineering
University of Glasgow



June 2026

Abstract

Passengers on public transport frequently use travel time for reading, working, relaxing, or watching video. Emerging extended reality (XR) technologies offer new opportunities to enhance these activities through immersion and engagement. However, highly immersive XR can reduce awareness of physical and social surroundings, creating challenges for safety, social interaction, and journey management. This thesis addresses the resulting design challenge: how to support immersive media consumption while preserving environmental awareness.

To investigate this tension, the thesis introduces augmented reality television (ARTV) Ornamentation — a design paradigm that extends media content spatially into the surrounding environment, rather than replacing it. Four studies inform this work: a survey of commuter activities and a series of participatory design workshops, two controlled experiments in simulated XR environments, and a situated evaluation using pass-through augmented reality (AR) in a physical subway carriage. The studies examine how device characteristics, augmentation intensity, and content design affect immersion and awareness. The findings show that spatially integrated augmentation can significantly increase media immersion relative to unaugmented viewing, while preserving substantially greater self-reported environmental awareness than fully immersive alternatives.

Synthesising these results, the thesis develops the Passenger Immersion–Awareness Tradeoff (PIAT) as a design framing for understanding how immersion and environmental awareness are dynamically negotiated in shared mobile environments, treating them as distinct outcomes of design decisions. From this framing it derives an ARTV design framework with five dimensions that shape the balance between immersion and awareness: ornamentation intensity, spatial integration, content coupling, adaptivity and awareness preservation. The framework also gives practical guidance for designing immersive media for passengers.

This work advances XR interaction design beyond stationary settings, demonstrating that immersion is achievable through environmental integration rather than replacement. The thesis contributes an empirical characterisation of passenger media experience, evidence on how ARTV Ornamentation shapes immersion and awareness, and a design framing and framework for XR systems supporting engaging, contextually appropriate media consumption on public transport.

Contents

Abstract **i**

List of Acronyms **xii**

Acknowledgements **xiv**

Declaration **xvi**

1 Introduction **1**

1.1 Motivation 1
1.2 Terminology and Scope 3
1.3 Augmented Reality Television in Passenger Contexts 5
1.4 Thesis Statement 6
1.5 Research Questions 6
1.6 Thesis Contributions 7
1.7 Thesis Structure 8

2 Review of the Literature **10**

2.1 Framing The Problem 10
2.2 Understanding Commuting and Passenger Behaviour 13
2.2.1 Commuting as a Use of Time and Health Problem 13
2.2.2 Activities Undertaken by Commuters 15
2.3 Media Immersion 17
2.3.1 Presence and Immersion in XR Research 17
2.3.2 Media Immersion 19
2.3.3 Display Size, Enclosure, and Field of View 19
2.3.4 Environmental Theming, Immersive Cinema and the Limits of Physical Augmentation 20
2.3.5 Implications for AR-Based Media Immersion 22
2.3.6 Measuring Immersion and Engagement 23
2.3.7 Summary 26

2.4 Supporting Media Immersion through ARTV 26
2.5 XR for Passengers: Opportunities and Barriers 30
2.5.1 Barriers to XR Use in Passenger Transit: Physical, Cognitive, Perceptual, and Social-Normative Constraints 32
2.6 AR Hardware Constraints for ARTV: Consumer Reality (2026) and Plausible Advances (2030–2035) 40
2.6.1 The Consumer Reality of AR in 2026: Device Families and Hard Limits 40
2.6.2 Plausible Technology Advances 2030–2035 43
2.6.3 Implications for Passenger ARTV 45
2.7 Evaluating AR using VR 45
2.7.1 Two Uses of VR Simulation in AR Research 45
2.7.2 Validity and Limitations of VR-Simulated AR 46
2.7.3 Practical and Ethical Constraints of Real-World AR Evaluation 48
2.8 Summary and Research Gap 49

3 Current and Desired Commuter Activities: A Survey and Participatory Design Workshops on Design Opportunities for Passenger XR 50
3.1 Introduction 50
3.2 Desired Commuter Activity Survey 51
3.2.1 Recruitment and Scope 52
3.2.2 Respondents 54
3.2.3 Existing Journey Experience 57
3.2.4 Desired Journey Experience 58
3.2.5 Discussion 62
3.3 Participatory Design Workshops: Exploring XR Design Opportunities for Commuters 66
3.3.1 Participants 67
3.3.2 Procedure 67
3.3.3 Analysis 68
3.3.4 Findings 68
3.3.5 Implications for XR Media Design 71
3.4 Conclusions 72

4 Using ARTV Ornamentations to Balance Media Immersion and Environmental Awareness for Passengers on Public Transport 74
4.1 Introduction 74
4.2 Study Overview 75
4.3 Experiment 76
4.3.1 Design 76

4.3.2 Measures . . . 78
4.3.3 Participants . . . 80
4.3.4 Procedure . . . 81
4.3.5 Analysis . . . 81
4.4 Results . . . 82
4.4.1 Effect of FoV . . . 83
4.4.2 Effect of Degree of ARTV Ornamentation . . . 85
4.4.3 Interaction between FoV and Degree of ARTV Ornamentation . . . 86
4.4.4 Comparisons with the AR and VR Baselines . . . 86
4.4.5 Condition Preferences . . . 87
4.4.6 Qualitative Results . . . 89
4.5 Discussion . . . 92
4.5.1 FoV Moderated the Benefit of Ornamentation . . . 92
4.5.2 Comparison with VR . . . 92
4.5.3 Balancing Awareness against Media Immersion . . . 93
4.5.4 Design Implications . . . 93
4.5.5 Limitations . . . 94
4.6 Conclusions . . . 96

5 Ornamentation Specificity and Adaptivity for Multi-Scene Narrative Media 98
5.1 Introduction . . . 98
5.2 Study: ARTV Objects and Biomes . . . 99
5.2.1 Design . . . 99
5.2.2 Media Content . . . 100
5.2.3 Conditions . . . 100
5.2.4 Measures . . . 104
5.2.5 Participants . . . 105
5.2.6 Procedure . . . 105
5.2.7 Analysis . . . 105
5.3 Results . . . 106
5.3.1 Comparisons with the Baseline . . . 106
5.3.2 Genre-Based and Show-Specific Augmentations . . . 106
5.3.3 Biome-Specific Augmentations . . . 109
5.3.4 Environmental Awareness . . . 109
5.3.5 Condition Preferences . . . 109
5.3.6 Qualitative Results . . . 111
5.4 Discussion . . . 112
5.4.1 RQ3-2: Should Ornamentation Be Show-Specific, or Is Genre Sufficient? 112
5.4.2 RQ3-3: Should Ornamentation Adapt to Changes in Scene Biome? . . . 113

5.4.3 RQ2-3: Would Passengers Choose Ornamented ARTV? 114
5.4.4 Applicability across Genres 114
5.4.5 Design Guidelines 115
5.4.6 Limitations 115
5.5 Conclusion 116

6 ARTV Ornamentations In Situ with Realistic Awareness Task 118
6.1 Introduction 118
6.2 Methodology 119
6.2.1 Study Design 119
6.2.2 Conditions 119
6.2.3 Apparatus and Environment 121
6.2.4 Participants 123
6.2.5 Tasks 123
6.2.6 Procedure 127
6.2.7 Measures 127
6.2.8 Analysis 128
6.3 Results 130
6.3.1 Quantitative Results 130
6.3.2 Preferences 134
6.3.3 Interviews 135
6.4 Discussion 140
6.4.1 RQ4-1: Comparing Simulated and Physically Situated AR 140
6.4.2 RQ4-2: Immersion–Awareness Trade-offs Across XR Viewing Modalities 144
6.4.3 Limitations 147
6.5 Conclusions 149

7 General Discussion 150
7.1 Introduction 150
7.2 The Passenger Immersion–Awareness Trade-off 151
7.2.1 Origins of the PIAT 151
7.2.2 The Relationship Observed Across the Research Programme 152
7.2.3 How Immersion and Awareness Vary Together 153
7.2.4 The Composition of Environmental Awareness 154
7.2.5 Contextual Variation in Awareness Demands 155
7.2.6 Relation to Established Models 156
7.2.7 The Status and Contribution of the PIAT 157
7.2.8 Applying the PIAT to Unstudied Contexts 159

7.3 A Design Framework for ARTV in Passenger Contexts . . . 159
7.3.1 ARTV as a Design Space for Responding to the PIAT . . . 160
7.3.2 Dimension 1: Ornamentation Intensity . . . 161
7.3.3 Dimension 2: Spatial Integration . . . 162
7.3.4 Dimension 3: Content Coupling . . . 163
7.3.5 Dimension 4: Adaptivity . . . 164
7.3.6 Dimension 5: Awareness Preservation . . . 165
7.3.7 Implications of the Framework . . . 166
7.3.8 ARTV Ornamentation as the Primary Framework Instantiation . . . 167
7.4 Coherence, Plausibility, and the Authoring of Augmentations . . . 167
7.4.1 Coherence and Plausibility in Passenger ARTV . . . 168
7.4.2 Authoring Augmentations at Scale . . . 168
7.4.3 Social Effects and Adoption Over Time . . . 169
7.5 Design Implications for Passenger XR Systems . . . 170
7.5.1 Design for Graduated Immersion . . . 170
7.5.2 Support Dynamic and Context-Aware Augmentation . . . 171
7.5.3 Preserve Environmental Reference Through Spatial Integration . . . 171
7.5.4 Design XR Media as Socially Situated Experiences . . . 172
7.5.5 Prioritise Perceptual Stability and Environmental Consistency . . . 172
7.5.6 Support Flexible Attention Management . . . 172
7.6 Limitations and Future Research Directions . . . 173
7.6.1 Limitations of the Research . . . 173
7.6.2 Future Research Directions . . . 177

8 Conclusions 186
8.1 Answering the Research Questions . . . 186
8.1.1 Research Question 1 . . . 186
8.1.2 Research Question 2 . . . 187
8.1.3 Research Question 3 . . . 187
8.1.4 Research Question 4 . . . 188
8.1.5 Summary of Research Question Outcomes . . . 189
8.2 Contributions of the Thesis . . . 189
8.2.1 C1: Opportunities for Immersive XR Media in Public Transport . . . 189
8.2.2 C2: Introduction and Evaluation of ARTV Ornamentation . . . 189
8.2.3 C3: How Design and Technical Parameters Shape the Trade-off . . . 190
8.2.4 C4: VR-Simulated AR as an Evaluation Tool . . . 190
8.2.5 C5: The PIAT as a Design Concept . . . 191
8.3 Conclusion . . . 191

A Raw Datasets **192**

B Survey Instruments **193**

B.1 Study 1: Commuter Behaviour Survey 193

B.1.1 Activity Engagement During a Journey 193

B.1.2 Follow-Up Questions (Conditional) 194

B.1.3 Current Commute Version 195

B.1.4 Notes on Survey Structure 195

B.2 Study 2: ARTV Ornamentation Evaluation 195

B.2.1 Participant Information Sheet and Consent Form 195

B.2.2 Demographic Questionnaire 196

B.2.3 Per-Condition Questionnaire Battery 197

B.2.4 End-of-Study Questionnaire 200

B.3 Study 3: ARTV Ornamentations for a New Genre 201

B.3.1 Participant Information Sheet and Consent Form 201

B.3.2 Demographic Questionnaire 203

B.3.3 Per-Condition Questionnaire Battery 204

B.3.4 End-of-Study Questionnaire 204

B.4 Study 4: ARTV In Situ with Realistic Awareness Task 205

B.4.1 Participant Information Sheet and Consent Form 205

B.4.2 Demographic Questionnaire 207

B.4.3 Per-Condition Questionnaire Battery 208

B.4.4 End-of-Study Questionnaire 209

C Interview Guides **211**

C.1 Study 1: ARTV Ornamentation Evaluation (Chapter 4) 211

C.1.1 Opening 211

C.1.2 Plant Density: Low vs. Medium 212

C.1.3 Augmentation Level: Medium vs. High 212

C.1.4 Passenger Awareness 212

C.1.5 Closing 212

C.2 Study 2: ARTV Ornamentations for a New Genre (Chapter 5) 213

C.2.1 Opening: Condition Preference Ranking 213

C.2.2 Biome Transitions: Scene Changes 214

C.2.3 Particle Effects: Snow and Sand 214

C.2.4 Object Specificity: Show-Specific vs. Genre-Based 214

C.2.5 Object Adaptation: Static vs. Scene-Responsive 215

C.2.6 Real-World Use and Generalisation 215

C.2.7 Repeated Clip 215

C.2.8 Closing 215
C.3 Study 3: ARTV In Situ with Realistic Awareness Task (Chapter 6) 216
C.3.1 Simulated AR Conditions (Lab) 216
C.3.2 Physically Situated AR Conditions (Subway Carriage) 217
C.3.3 VR Cinema Condition 217
C.3.4 VR Videosphere Condition 217
C.3.5 Comparison: Simulated AR vs. Physically Situated AR 217
C.3.6 Comparison: Train Conditions vs. Cinema and Videosphere 218
C.3.7 Stop Awareness Task 218
C.3.8 Real-World Use 218
C.3.9 Condition Preference Ranking 219
C.3.10 Closing 219

D Complete Statistical Results: Chapter 4 220

E Complete Statistical Results: Chapter 5 240

F Complete Statistical Results: Chapter 6 247

List of Tables

3.1 Commuting activities in the survey and their sources 53
3.2 Activity categorisation counts 59

4.1 Fields of view and augmentation for each condition 78
4.2 Internal consistency of the questionnaires 82
4.3 Factorial analysis of the six ARTV Ornamentation conditions 83
4.4 Comparisons of the ARTV conditions with the AR and VR Baselines 85

5.1 Internal consistency of the questionnaires 106
5.2 ART ANOVAs and contrasts for the four conditions 108

6.1 Median scores per condition for UEQ-S subscales (range -3 to $+3$), UES-SF subscales (range 1–7), and awareness items (range 1–5). Awareness-P = awareness of other passengers; Awareness-A = awareness of augmentations, analysed only for the three conditions containing added visual content. 130
6.2 Two-way ART ANOVA results for the four AR and simulated-AR conditions, crossing Medium (pass-through, simulated) with Ornamentation (baseline, ornamented). All effects have df $= (1,54)$. Significant effects are shown in bold. . 131
6.3 ART ANOVA results and significant pairwise contrasts for the six-condition analysis. F statistics have df $= (5,90)$ throughout. Contrasts are ART-C with Tukey correction; g_z is Hedges' g_z calculated from raw paired differences; negative g_z indicates the first-named condition scored lower. Non-significant measures (UEQ-S Overall, UEQ-S Pragmatic, UES-SF Perceived Usability) are omitted; see appendix F for the complete results. 132

7.1 Relationship between the framework dimensions (section 7.3) and the design implications below, with the primary source of evidence for each. 170

D.1 Descriptive statistics (Median, IQR, Mean, *SD*) for all measures by condition ($n = 24$ per condition). 221
D.2 Chapter 4: omnibus tests for the factorial design 225
D.3 Chapter 4: post hoc comparisons for the factorial design 227

D.4 Chapter 4: comparisons with the AR Baseline 230
D.5 Chapter 4: comparisons with the VR Baseline 234
D.6 Chapter 4: equivalence tests against the VR Baseline 238
D.7 Chapter 4: exit questionnaire descriptives 238
D.8 Chapter 4: pairwise comparisons of the exit questionnaire 239

E.1 Descriptive statistics (Median, IQR, Mean, SD) for all Chapter 5 measures by condition ($n = 20$ per condition). Film-IEQ subscale scores are summed item totals; UEQ-S subscales scored -3 to $+3$; UES-SF items were answered on a 7-point scale and scored 1–7, with Perceived Usability items reversed; awareness items scored 1–5. 241
E.2 Omnibus ART ANOVA results for all Chapter 5 measures. F statistics have $\mathrm{df} = (3, 57)$ throughout; η_p^2 is partial eta squared. 242
E.3 Complete ART-C pairwise contrasts for all Chapter 5 measures. Contrasts are Tukey-corrected; $\mathrm{df} = 57$ throughout. Estimates, standard errors, and t are on the aligned-rank scale; g_z is Hedges' paired effect size computed from the raw scores. Negative values indicate the first-named condition scored lower. Significant contrasts are marked in the final column. Conditions: BL = Baseline; SSA = Show-Specific Augmentations; GBA = Genre-Based Augmentations; BSA = Biome-Specific Augmentations. 243

F.1 Median scores per condition for Chapter 6 measures. UEQ-S subscales are scored from -3 to $+3$, UES-SF subscales from 1–7, and awareness items from 1–5. Awareness of the augmentations was only meaningful in the three conditions containing added visual content. 247
F.2 Omnibus ART ANOVA results for the six-condition analysis of Chapter 6. F statistics have $\mathrm{df} = (5, 90)$ throughout; η_p^2 is partial eta squared. 248
F.3 Supplementary tests reported in Chapter 6. Awareness of the augmentations was compared across the three conditions containing added visual content; the stop task was compared across all six; the bystander model covers the four physical-carriage conditions; and the exit-survey measures cover all six conditions for the 19 analysed participants. 248
F.4 Complete ART-C pairwise contrasts for the six-condition analysis of Chapter 6. Contrasts are Tukey-corrected. Estimates, standard errors and t are on the aligned-rank scale; g_z is Hedges' paired effect size computed from the raw scores. Negative values indicate that the first-named condition scored lower. Significant contrasts are marked in the final column. 249

List of Figures

1.1 Examples of ARTV from prior work 9

2.1 IllumiRoom peripheral projected illusions 27
2.2 The ARTV Continuum 28
2.3 Virtual living room used in ARTV evaluation 30
2.4 Interaction space in commercial VR compared with transport seating 35
2.5 Reality Anchors in a simulated passenger environment 36
2.6 HoloLens 2 Mixed Reality Capture compared with a photograph through the optics 42

3.1 Transport modes used by commuters 55
3.2 Commute duration by primary transport mode 56
3.3 Primary commute activity 57
3.4 Activity categorisation heatmap 60

4.1 The eight viewing conditions 79
4.2 Film-IEQ, UEQ-S, and UES-SF scores by condition 84
4.3 Prospective use of each condition on a train 87
4.4 Participant rankings of the conditions 88

5.1 The viewing conditions 101
5.2 Questionnaire scores by condition 107
5.3 Interest in using each condition on a train 110
5.4 Participant rankings of the conditions 110

6.1 The four AR and simulated AR conditions 124
6.2 The two fully immersive VR conditions 125
6.3 The Glasgow Subway carriage at the Riverside Museum 126
6.4 Condition preference rankings 135
6.5 Interest in using each condition instead of a phone or laptop 136

7.1 Focused attention against environmental awareness across the three studies . . 153

List of Acronyms

ANOVA	analysis of variance
AR	augmented reality
ARI	Augmented Reality Immersion questionnaire
ART	Aligned Rank Transform
ART-C	Aligned Rank Transform Contrast
ARTV	augmented reality television
AV	autonomous vehicle
BSA	biome-specific augmentation
DoA	degree of ARTV ornamentation
DR	diminished reality
EMG	electromyography
Film-IEQ	Film Immersive Experience Questionnaire
FoV	field of view
GBA	genre-based augmentation
HCI	human–computer interaction
HMD	head-mounted display
IEQ	Immersive Experience Questionnaire
LiDAR	light detection and ranging
MR	mixed reality
MRUK	Mixed Reality Utility Kit
NDRT	non-driving-related task
NRPS	National Rail Passenger Survey
OST	optical see-through
PIAT	Passenger Immersion–Awareness Trade-off

SSA show-specific augmentation

SWB subjective wellbeing

TOST two one-sided tests

UEQ User Experience Questionnaire

UEQ-S User Experience Questionnaire–Short

UES User Engagement Scale

UES-SF User Engagement Scale–Short Form

VOT value of travel time savings

VR virtual reality

VST video see-through

XR extended reality

Acknowledgements

After agonising over the order of these acknowledgments for much too long I have decided that I will go with the best order - no particular order!

I will first thank my supervisors Mark McGill and Stephen Brewster. Mark, thank you for providing such an interesting Master's project proposal that caught my attention and captured my imagination, thus sending me down the path of XR research and culminating in this PhD. Ever since that Master's project proposal you have been unendingly supportive and helpful, especially when it comes to paper writing against a deadline. The speed of replies and after work hours feedback and redrafts sometimes made me worry for your work life balance. Not even the birth of two children prevented you from providing support to your PhD students! Steve, thank you for welcoming me to the Multimodal Interaction Group and the supportive research environment it provides. The Wednesday group meetings and associated lunches were always a highlight - both for picking people's brains about particular problems and for the levity the lunches provided. Likewise the MIG summer outings were brilliant for some recreation and to mark another year passing during this blur of a PhD.

Continuing the MIG thread, I have to thank all of my wonderful 'MIG mates'. Firstly my office compatriots who were in the trenches with me, Jake - the Sonomancer himself, Diego - the always friendly and always fashionable Italian, and Kieran - a brilliant PhD student who sets a truly high standard for what a PhD student can be, do, and achieve. To complete the office section, Graham. Thank you for always letting me pick your brain every time I had a, not infrequent, question. You were always a friendly face on a Monday morning and somehow always had the answer - even before the morning coffee! Thanks to Rune for his brief but memorable visit to Glasgow from Aalborg and the many pints and games of pool that we shared. Thanks to James for the shared fantasy of our Dungeons & Dragons sessions and for slowly teaching me to skateboard through his own PhD studies. Thanks to my two unofficial mentors - Melvin & Shaun; you both always gave me much needed advice and helped me get to grips with this whole PhD thing. Lastly but not least, Paddy. Thank you for your unrelenting support and the encouragement that you provided, to not just me but to everyone in the group. Your enthusiasm is infectious and brightened up countless Wednesday meetings. As I can't thank the entire department by name I will have to settle for broadly thanking everyone else in MIG and GIST who made the experience of being part of those groups as friendly and enjoyable as it was.

Thank you to my examiners, Florian Weidner and Anthony Steed. You both set me at ease at the start of the viva without even trying, which allowed us to have a very interesting and insightful discussion without pre-viva nerves getting in the way. It was a very enjoyable experience, insofar as a viva — which I had been stressing about for weeks — can be enjoyable.

I need to thank both my editors (who only edited the previous draft so any mistakes or typos herein are wholly my responsibility), Anna - my wonderful girlfriend, and my ever patient Mum (Shirley).

Anna, the past three years of working on this PhD has been made infinitely more enjoyable with you in my corner. You kept me sane during this whole process and you never doubted my ability, even at the times when I doubted my own. As you come to write your own PhD thesis I just hope that I can provide even a fraction of the support that you've given me. Unfortunately, I don't have your knack for editing but I'll try my best!

Thank you Mum, for always supporting me in everything that I do. I would have thought you would be tired of that by now 28 years on but you still support and encourage me each and every day and it feels just as good to know that now as it did on that first day of primary school all those years ago when I had to go in by myself for the first time.

Thank you Dad, although you weren't able to provide any editing help (Mum had to edit your own thesis after all), you're always encouraging and able to make me laugh even on a bad day. Thank you to both of you, Mum and Dad, you have raised my spirits when it was needed and reminded me that I was never alone. You always believed in me and you have done more for me than I could ever repay - but hopefully having the bragging rights that your son has a PhD somewhat helps. And finally my brother, David. You were probably the smart one stopping after your Masters degree. Thanks for always pushing me to be the best version of myself (I had to say something nice). You are a champion sword fighter after all who owns several swords - and knows how to use them!

Thank you to all my friends but especially to Adam and Selina. You both made sure that I actually dedicated time to living my life during this PhD, with our many holidays and adventures keeping me sane. You've both helped me much more than you realised. I promise I'll be a more attentive friend now that this sword of Damocles has been lifted.

Finally, thank you to everyone else who has touched my life, friends, family, teachers. You have all had an impact on me and helped me get to this point, the point that I can scroll past *Mr* on a drop down menu online and instead select *Dr*.

Declaration

Unless otherwise explicitly stated, the research presented in this thesis is entirely the author's own work.

Chapter 4 has been published as a paper at the ACM International Conference on Interactive Media Experiences 2024: Iain Christie, Mark McGill, and Stephen Brewster. "Welcome to the Jungle: Using ARTV to Balance Media Immersion and Reality Awareness for Passengers on Public Transport". In: Proceedings of the 2024 ACM International Conference on Interactive Media Experiences. Stockholm, Sweden: ACM, June 2024, pp. 219–230. DOI: 10.1145/3639701.3656321

Chapter 5 has been published as a paper at the ACM International Conference on Interactive Media Experiences 2025: Iain Christie, Graham Wilson, Mark McGill, and Stephen Anthony Brewster. "Winter is Coming… To Your Commute: Adapting Passenger ARTV Ornamentations to Match On-screen Visuals For Enhanced Immersion in Fantasy Drama TV Content". In: Proceedings of the 2025 ACM International Conference on Interactive Media Experiences. Niterói, Brazil: ACM, June 2025, pp. 275–293. DOI: 10.1145/3706370.3727854

Chapter 6 has been accepted for publication at the ACM International Conference on Mobile Human-Computer Interaction 2026: Iain Christie, Graham Wilson, Mark McGill, and Stephen Brewster. "Watch the Thrones, But Mind the Gap: Attention Management and Ecological Validity in In-Transit AR Media". In: Proceedings of the 26th International Conference on Mobile Human-Computer Interaction. Swansea, Wales: ACM, August 2026.

Chapter 1

Introduction

Advances in extended reality (XR), an umbrella term encompassing virtual reality (VR), augmented reality (AR), and mixed reality (MR) technologies [121], have made immersive media a realistic alternative to conventional screen-based viewing. Improvements in display fidelity, tracking, and wearable form factors mean that XR systems are no longer confined to laboratories or domestic environments, but are instead designed for use in everyday, mobile, and shared contexts. Commercial platforms now explicitly promote scenarios such as watching films or television on large virtual displays while commuting on trains, buses, or aircraft [198]. As XR technologies move into these real-world settings, they raise new challenges for human–computer interaction (HCI), particularly for attention and environmental awareness, which in turn affect user safety.

1.1 Motivation

Commuting occupies a substantial and recurring portion of daily life. In England, the average one-way commute is now 27 minutes, amounting to almost an hour of travel on each day a person commutes [48]. Aggregated across a working life this represents a considerable quantity of time, frequently characterised as unproductive or wasted [35, 68], and associated in the wider literature with measurable costs to wellbeing: longer commutes have been linked to elevated stress [47], poorer self-reported health [56, 84], and increased prevalence of depressive symptoms [91]. Ettema et al. [45] argue that travel contributes to subjective wellbeing (SWB) not only by enabling activities at the destination, but through the affect associated with the experience of travelling itself. If the quality of the journey matters independently of where it leads, then improving the experience of travelling, as opposed to shortening the trip or changing the mode, offers a plausible route to improving passenger wellbeing.

Passengers already attempt this for themselves. Analysis of over 80,000 responses to the National Rail Passenger Survey between 2004 and 2014 shows a marked rise in technology-mediated activities such as watching video, browsing, and messaging, with a corresponding decline in activities such as reading print, talking to other passengers, and window-gazing [103].

Singleton [161] found that passengers on *riding modes* (modes in which the traveller is carried and does not operate the vehicle) engage in a greater number and variety of activities than those operating their own mode of travel, and almost 70% of transit commuters reported at least one technology-mediated activity. He interprets much of this activity as coping with the journey rather than productive use of the time: engagement was associated with dissatisfaction with travel time and with longer journeys, which suggests that passengers reach for their devices to make an unwelcome journey tolerable, not to work through emails or documents. On this reading, entertainment is the principal use of travel time.

Yet the media experiences available to passengers remain constrained by the devices that deliver them. Handheld screens subtend a small visual angle at typical viewing distances, and both screen size and viewing angle are known to affect reported immersion and presence [148, 64]; attention is further fragmented by the surrounding environment and by the device itself. While mobile devices support convenient access to media, they offer limited opportunity for deeply engaging experiences that extend beyond the boundaries of a handheld display and integrate with the surrounding environment. A passenger watching video on a phone is therefore having a materially poorer experience than the same content would afford at home, precisely in the context where its compensatory value is greatest. XR headsets remove this constraint: content can be rendered at any apparent size, positioned anywhere in the surrounding space, and extended beyond the boundaries of a display altogether.

Commercial XR hardware now supports at least some of the capabilities relevant to passenger media: headsets that show the wearer's surroundings through outward-facing cameras are commercially available [5], display glasses that present a large virtual screen are sold by several vendors [198], glasses with an integrated in-lens display reached the consumer market in 2025 [115], and see-through waveguide optics (thin transparent lenses that guide a projected image to the eye) demonstrated in 2026 reach fields of view intended for video viewing [178]. Manufacturers are also positioning these capabilities for travel [198]. Passenger XR media is therefore no longer purely speculative, though its design for the specific demands of the transit context remains largely unaddressed.

There is good reason to think the setting matters. Public transport environments are socially dynamic and attention-critical. For journeys with several possible disembarkation points, passengers must remain aware of their surroundings, monitor journey progress, and attend to personal safety and social norms. Lee et al. [92] document this difficulty in existing smartphone use, where passengers negotiate between engagement and the risk of missing their stop, and in-the-wild deployment of immersive headsets on real train journeys confirms that comparable concerns persist outside the laboratory [18]. Chapter 2 reviews this literature in full. This introduces a tension between media immersion and environmental awareness that is particularly acute in passenger contexts, where users must monitor their surroundings and journey progress while engaged with media, and which any immersive media experience deployed in transit must

resolve. Designing poorly for this context therefore risks producing experiences that passengers are unwilling to use, or that compromise their safety if they do.

Existing XR research in passenger settings has largely pursued productivity and full immersion, or has addressed the safety and social-acceptability barriers that both raise; chapter 2 reviews this work in detail. Much of it either pursues full immersion, or transposes an existing activity into the vehicle using virtual displays positioned around the user, leaving the activity itself unchanged: Medeiros et al. [110], for example, reconstruct a multi-monitor office setup on a train. Less attention has been paid to enriching the everyday media experience passengers already have: making watching something on a commute better than it currently is, without replacing the journey or working through it. Little work has asked passengers directly what they would want to do during their journey if immersive technology were available. Without that grounding, both research and industry may miss opportunities to let passenger insight inform the development of passenger XR.

1.2 Terminology and Scope

Several terms used throughout this thesis carry more than one meaning in the literature; this section fixes the senses used here, and chapter 2 discusses the underlying distinctions in full. This thesis concerns *passengers*: travellers who do not have primary responsibility for navigation or vehicle control. This excludes drivers, walkers, and cyclists, whose attention is necessarily committed to the task of travelling. The distinction is one of attentional availability, and differs from the active/passive distinction common in the travel-behaviour literature, where active travel denotes physically active modes such as walking and cycling and a car driver would be classified as a passive traveller. The thesis further concerns journeys with multiple possible disembarkation points, such as bus, subway, and rail services, where passengers must actively monitor progress in order to alight correctly. Single-destination journeys with enforced disembarkation, such as most flights and ferry crossings, impose weaker awareness demands and are not the focus here. The empirical studies concern seated passengers: chapters 4 and 5 frame the journey as a longer rail journey on which passengers need not track stops, whereas chapter 6 adds an explicit stop-monitoring task. Standing, crowded, and luggage-managing passengers are not studied.

Immersion is used in two distinct senses in the literatures this thesis draws on: as an objective property of a system, and as a user's perceptual, cognitive, and emotional absorption in mediated content. This thesis keeps the two apart. The properties it manipulates (such as device field of view (FoV), the extent and intensity of augmentation, and the degree to which the real environment is occluded) are referred to as a configuration's *immersive properties*; the user's absorption, which the thesis measures, is referred to as *media immersion*. *Environmental awareness* is likewise an outcome measured in the user, not a property of the system. It denotes

a passenger's awareness of their physical and social surroundings: other passengers, journey progress, personal belongings, and events in the carriage, and is distinguished where relevant from awareness of the virtual content itself. This thesis measures immersion and awareness separately and examines how a system's immersive properties shape both media immersion, and awareness of the real and the virtual. The term *environmental awareness* is used in preference to *situational awareness*, which carries specific theoretical commitments in the human factors literature that this thesis does not operationalise in full; where that literature is discussed, the term is used in its established sense and identified as such.

Because this thesis compares AR and VR media experiences throughout, it is worth stating briefly how each is delivered. A VR head-mounted display (HMD) occludes the wearer's view of the physical environment and renders a synthetic environment in its place. AR instead presents virtual content alongside the real environment, through one of two architectures. The first, video see-through (VST) or passthrough, occludes the eye as VR does, reconstructs the surroundings from outward-facing cameras, and composites virtual content into that reconstructed view. The second, optical see-through (OST), instead projects imagery onto a transparent combiner, so that light from the real world reaches the eye directly and is never replaced by a camera feed. The terms AR and MR are applied inconsistently across this space; this thesis reserves MR for systems in which virtual content is registered to, and can be occluded by, real-world geometry, and uses AR as the broader term. Chapter 2 returns to the question of whether these categories are best understood as points on a continuum, following recent re-examinations of that framing [163].

These architectures differ in ways that matter for video content specifically. Conventional OST displays combine light additively: virtual content is added to the light already reaching the eye and cannot replace it. Dark image regions therefore render as transparent, and contrast falls as ambient light rises. Both properties are poorly suited to video, which depends on controlled contrast and on dark tones appearing dark. These are, however, active engineering targets: occlusion-capable architectures that mask incoming light on a per-pixel basis are an established research direction [146, 203], and commercial devices already ship dimming mechanisms to raise effective contrast [27]. Section 2.6 examines this trajectory, and its remaining constraints, in full.

This thesis therefore treats OST as a plausible near-term platform for passenger augmented reality television (ARTV) (introduced in section 1.3), while conducting its empirical work under conditions that current hardware can deliver: chapters 4 and 5 use VR-simulated AR to manipulate FoV and augmentation opacity beyond what shipping optical devices allow, and chapter 6 deploys on passthrough hardware.

A related scope decision concerns how AR experiences are evaluated. Owing to practical and technical constraints, several of the studies in this thesis use VR-simulated AR, in which a virtual reconstruction of a real environment approximates an AR experience under controlled

conditions; chapter 2 reviews the validity of this approach in detail. This thesis does not treat simulation as an unquestioned proxy for real-world AR: the majority of the empirical work uses simulated AR-in-VR to establish foundational insights under controlled conditions, and is complemented by a final *in situ* investigation deploying ARTV Ornamentation on consumer AR hardware within a physical public transport setting, allowing the thesis to examine how findings derived from simulation behave under real-world conditions.

1.3 Augmented Reality Television in Passenger Contexts

The occlusion inherent to VR carries a direct cost for passenger use: high levels of immersion are gained at the expense of awareness of the real world. Attempts to reintroduce that awareness within VR have been investigated: the Reality Anchors of Bajorunaite et al. [17] selectively represent real-world elements such as other passengers and stop signage inside the virtual environment. Their findings show that this shifts the tension without resolving it: anchors salient enough to support awareness disrupted immersion, while less intrusive anchors failed to provide reassurance. Recovering awareness in VR also depends on the system first recognising and reconstructing the relevant elements of the environment. AR inverts this dependency: the surrounding environment remains directly visible, so awareness is not contingent on computational reconstruction, and the design problem becomes one of placing virtual content so that it does not occlude what the passenger needs to see. This is not to say AR is unconstrained. Limited device FoV, display luminance, and visual clutter all shape what is achievable, and visibility of the environment does not by itself guarantee that a passenger attends to it (see section 2.5). It is because these are design problems within a system whose awareness baseline is given, rather than dependencies on a recognition system that must be built and validated, that this thesis takes AR as its scope. The alternative, that VR supplemented by well-designed awareness cues could match or exceed AR, is not tested empirically in this thesis; the evidence bearing on it, including the Reality Anchors work, is reviewed in section 2.5.1.

Within HCI research, the concept of ARTV has been introduced as a design space in which video content extends beyond a physical screen through spatial augmentations embedded in the user's environment [183]. These augmentations may take multiple forms, including secondary virtual screens, informational panels, environmental effects, or narrative elements integrated into the surrounding space. Figure 1.1 illustrates this range: in a, a virtual animal leaves the screen and enters the room while a companion display provides supporting context, whereas the scenarios in b show augmentation used to present alternative perspectives on a scene and to place supplementary information alongside the screen. What these share is that the augmentation is anchored to the viewer's physical surroundings and extends beyond the display, so the room itself becomes part of the viewing experience.

ARTV therefore reimagines media viewing as a spatial experience in which digital content

extends beyond the primary screen into the viewer's environment. Existing work has developed this idea almost entirely for the living room, where the surrounding environment is private, familiar, static, and makes few competing demands on the viewer's attention. Prior work has not systematically examined how such augmentation can enhance immersion while preserving the environmental awareness required in attention-critical public contexts. This thesis addresses that gap through ARTV Ornamentation: rather than ARTV in general, which may extend narrative or informational content anywhere in the surrounding environment, Ornamentation specifically denotes augmentation that is thematically tied to the on-screen content and spatially integrated with the immediate viewing environment, distinguishing it from augmentation that is purely informational or otherwise unrelated to what is being watched. This tension between immersive engagement and the need for ongoing environmental awareness forms the central challenge addressed in this thesis.

1.4 Thesis Statement

Immersive media experiences on public transport must balance media engagement with the need for passengers to remain aware of their surroundings. This thesis investigates how AR can support immersive passenger media consumption through ARTV Ornamentation: content-themed AR augmentations that extend video content beyond the display into the surrounding environment, complementing on-screen media without fully occluding the real world. Through controlled experimentation and situated evaluation, the research demonstrates that environmental augmentation can increase reported media immersion relative to unaugmented viewing, and characterises how the parameters governing that augmentation shape passengers' reported awareness of their surroundings.

1.5 Research Questions

The research presented in this thesis is guided by the following research questions:

- **RQ1:** How could XR improve the current experience of immersive media on public transport?
- **RQ2:** How can ARTV with Ornamentation be used to balance media immersion and environmental awareness in passenger transportation?
- **RQ3:** What is the impact of device FoV, ornamentation density, and ornamentation type on the media immersion–environmental awareness trade-off?
- **RQ4:** How does the passenger immersion–awareness trade-off apply when moving from simulated AR-in-VR to AR experiences situated in the real world?

Chapter 3 addresses RQ1, chapters 4 and 5 address RQ2 and RQ3, and chapter 6 addresses RQ4. Each empirical chapter divides its research questions into sub-questions, numbered RQ2-1, RQ3-1, and so on.

1.6 Thesis Contributions

This thesis makes the following contributions to research in HCI and immersive media design:

- **C1: An empirical characterisation of opportunities for immersive XR media in public transport contexts.** Through a survey of current and desired commuter activities and a series of participatory design workshops, the thesis identifies limitations of existing media consumption practices on public transport and characterises the design space in which XR technologies might address unmet passenger needs, thereby addressing **RQ1**.

- **C2: The introduction and systematic evaluation of ARTV Ornamentation as a design mechanism for balancing immersion and awareness.** The thesis introduces ARTV Ornamentation as a class of content-themed AR augmentations that extend on-screen media into the surrounding environment, and provides empirical evidence of how such ornamentations influence user immersion and environmental awareness, addressing **RQ2**.

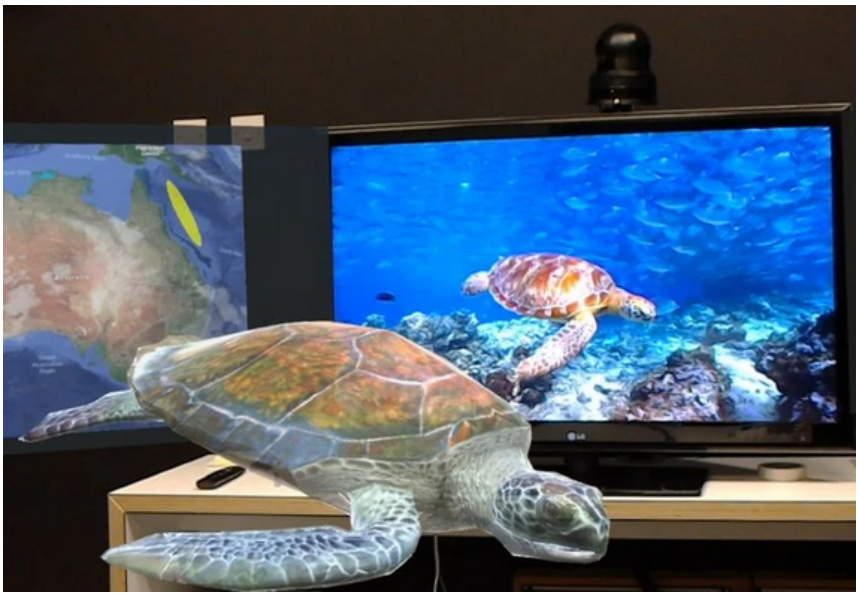
- **C3: An analysis of how key design and technical parameters shape the immersion–awareness trade-off.** By examining the effects of device FoV, ornamentation density, and ornamentation type, the thesis demonstrates how specific configuration choices affect reported media immersion, engagement, and environmental awareness, addressing **RQ3**.

- **C4: A methodological case study of VR-simulated AR as a tool for evaluating AR media experiences.** Through a paired comparison of simulated and situated deployments of ARTV Ornamentation ($N = 19$), the thesis identifies the specific dimensions along which simulation did and did not transfer: equivalence between the two media held for some kinds of finding but not others. This addresses **RQ4**.

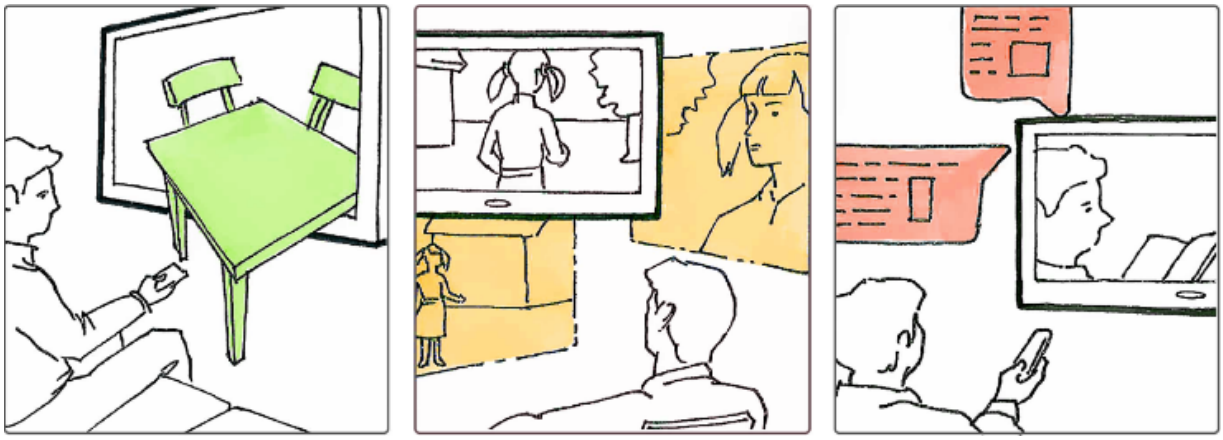
- **C5: The development of the Passenger Immersion–Awareness Trade-off (PIAT) as a design concept for passenger XR.** Synthesising evidence across the empirical studies, the thesis develops an intermediate-level design concept of how the immersive properties of a passenger media experience shape media immersion and environmental awareness. It distinguishes what a design leaves available to the passenger from how immersive it is, and sets out the conditions under which the account would not hold.

1.7 Thesis Structure

The remainder of this thesis is organised as follows. Chapter 2 reviews related work on commuter behaviour, immersive media consumption, ARTV, and the evaluation of AR using VR-based simulation, identifying gaps in both design knowledge and empirical validation. Chapter 3 reports a survey of current and desired commuter activities together with participatory design workshops, establishing the design space for immersive XR media in passenger contexts. Chapter 4 examines how ARTV Ornamentation can be used to balance media immersion and environmental awareness through a controlled study of device FoV and ornamentation intensity. Chapter 5 extends this investigation to ornamentation transitions and genre-specific designs, examining how different presentation strategies influence the trade-off. Chapter 6 presents an *in situ* study deploying ARTV Ornamentation on consumer AR hardware within a physical subway carriage, incorporating a realistic awareness task, and examines how findings from the simulated studies behave in a situated setting. Chapter 7 synthesises these findings, develops the PIAT as a design concept for reasoning about immersive media in passenger contexts, and sets out the resulting design framework and its implications for the evaluation of XR experiences in public contexts. Chapter 8 returns to the research questions and summarises the contributions of the thesis.



(a) ARTV content added to a nature documentary. A virtual sea turtle crosses in front of the television and a virtual side screen displays the location of its habitat.



(b) Artist illustrations of three (of many possible) ARTV scenarios: virtual objects emerging from the television screen (left), simultaneous alternative perspectives of the same scene (centre), and supplementary content such as character names displayed beside the screen (right).

Figure 1.1: Examples of ARTV from prior work, in which augmentation extends beyond the screen boundary. (a) adapted from Saeghe et al. [150]; (b) reproduced from Popovici and Vatavu [140].

Chapter 2

Review of the Literature

2.1 Framing The Problem

The daily commute is a reality for hundreds of millions of people worldwide and is widely recognised as an activity that can impact wellbeing, productivity, and perceived quality of life [35]. Today, advances in XR technologies and, in particular, AR stand to create new possibilities for transforming how passengers experience travel time [107]. With the ability to change how the user sees the world, XR affords new opportunities for immersion, entertainment, and support for meaningful activities during transit [110, 95].

A substantial body of work [103, 161] indicates that media consumption is one of the most common uses of travel time for passengers, often taking precedence over productivity-oriented activities. Passengers frequently engage in activities such as watching video, listening to music, playing games, or browsing content, reflecting both the constraints of the travel environment and preferences for lightweight, interruptible forms of engagement [161]. While prior research has also explored the use of travel time for work or productivity, this thesis focuses specifically on media consumption during self-managed journeys [15]: a primary and widely adopted commuter activity [161, 103] in which individuals are not responsible for vehicle control but must nevertheless remain perceptually and socially engaged with their surroundings. This context is especially relevant because the absorption and attentional focus associated with media immersion can intensify the tension between immersive engagement and environmental awareness. These constraints on immersive media consumption differ substantially from both domestic and fully private environments.

In this thesis, media immersion refers to the degree to which a passenger becomes perceptually, cognitively, and emotionally absorbed in mediated audiovisual content, such that attention is primarily directed toward the media experience rather than the surrounding environment [69]. This differs from the use of immersion in some XR and VR literature, where immersion is often treated as a property of the technological system or as a precursor to presence within a simulated environment [166]. Here, the concern is not whether passengers feel transported into a wholly

virtual world, but whether XR can deepen engagement with media while the passenger remains physically and socially situated in public transport. Situational awareness, in contrast, refers to the user's ability to perceive, comprehend, and act upon relevant cues in the surrounding environment [44]. In passenger contexts these cues include information about upcoming stops, nearby passengers, and the location of personal belongings. A central challenge in designing XR media experiences for passengers therefore lies in balancing two competing requirements: sustaining media immersion while preserving environmental awareness. These two requirements are inherently in tension. Increasing a system's immersive properties typically reduces awareness of the surrounding environment, whilst preserving awareness can limit the depth of media immersion (section 1.2).

Media consumption is an especially important context for examining passenger XR because absorption, focus, and attentional withdrawal are not incidental to the experience — they are central to its quality. Passengers already watch video, listen to audio, and browse content during their journeys, but these activities are constrained by small displays, fragmented attention, and the need to remain aware of the surrounding transport environment. XR could enhance these experiences by enlarging, spatialising, or environmentally integrating audiovisual content — yet this same capacity to deepen media immersion risks reducing awareness of journey progress, nearby passengers, and personal belongings. This thesis therefore addresses the immersion–awareness trade-off specifically as it manifests in passenger media consumption, through a passenger-focused implementation of ARTV [183]. This implementation blends media content with supportive spatial augmentations — *ornamentations* — in the surrounding environment, aiming to enhance immersion without fully occluding the real world.

There exist several key challenges to integrating XR experiences into public transport settings, most notably those relating to (i) maintaining environmental awareness, (ii) ensuring social acceptability in shared environments, and (iii) managing perceptual and physical constraints such as motion and display limitations. Prior work has highlighted issues such as motion sickness and social acceptability in shared environments [78], which reflect these broader perceptual, physical, and social constraints. For XR to be viable in commuting contexts, all of these issues must be addressed. This thesis addresses the requirement to balance the degree of immersion in media experiences with the need for environmental awareness. A passenger wants to be immersed enough that distractions such as other passengers fade from their perception, but needs to maintain enough awareness to ensure their personal safety, avoid their belongings being stolen, and to remember to disembark at their destination. This immersion–awareness trade-off must be considered carefully. VR provides extremely immersive experiences but this comes at a cost to the user's environmental awareness; VR, by design, blocks the user's view of their surroundings which provides an immersive experience; this same feature, however, creates the aforementioned safety and awareness risks [18]. AR, by contrast, preserves the user's view of the surrounding environment, but is more limited in its immersive potential by the constraints of current display

technologies. Section 1.2 introduced the two display architectures used throughout this thesis. Their differences have direct consequences for passenger media. Because VST systems present the environment as a rendered image, everything the passenger perceives of their surroundings is subject to the camera pipeline: resolution limits, colour reproduction, latency between real motion and displayed image, and geometric distortion arising from camera placement offset from the eyes. Environmental awareness in a VST system is therefore mediated awareness, bounded by reconstruction fidelity. OST systems avoid this, since light from the environment reaches the eye unmodified and awareness is limited only by what virtual content occludes. The cost is borne on the virtual side. Because the combination is additive, the system can add light but not remove it: black regions of an image are not rendered black but left transparent, and the passenger sees the environment through them. For video content this is a substantial constraint — dark scenes become windows onto the carriage, and because the background behind them is uncontrolled and in motion, what is visible varies with the vehicle's surroundings rather than with the media. Section 2.6 returns to this problem and to the occlusion technologies that may eventually address it.

Several strands of prior research motivate this thesis's focus on the tension between media immersion and environmental awareness, even though none has articulated it as a unified passenger-specific framing. Work on media immersion characterises immersive viewing as a state of perceptual and cognitive absorption, in which attention is directed toward mediated content and away from competing environmental stimuli [69]. Research on passenger behaviour and transit multitasking shows, in parallel, that travel environments require ongoing monitoring of journey progress, nearby passengers, social cues, and personal belongings; lapses in this awareness can lead to missed stops, discomfort, or reduced perceived safety [92]. XR technologies intensify this tension because they alter the relationship between mediated content and the physical environment: fully occlusive VR supports strong immersion by removing environmental cues, whereas AR preserves environmental visibility but typically provides weaker immersive affordances. This problem has concrete illustrations in existing passenger XR work. Williamson et al. [188] identified loss of environmental awareness as a major concern for passengers considering VR use on a plane. Bajorunaite et al. [17] addressed this issue through Reality Anchors—selectively reintroducing real-world elements into VR—and found that while awareness cues improved perceived safety, they could also disrupt immersion when continuously present. This dynamic, in which each goal constrains the other, is precisely the tension this thesis frames and investigates. These strands of work have not, however, been brought together into a passenger-specific account of how media immersion and environmental awareness should be balanced in XR. This thesis investigates that relationship empirically across four studies, and chapter 7 develops the resulting account.

This thesis argues that passenger media XR should prioritise selective environmental augmentation rather than full visual occlusion. Specifically, the thesis investigates whether ARTV

ornamentation—spatially aligned visual elements that extend media content into the surrounding environment—can enhance media immersion while preserving the environmental awareness that passengers report needing for safety and social acceptability in public transit. To evaluate such approaches, it is necessary to understand commuter behaviours, XR affordances and limitations, media immersion theory, environmental awareness requirements, and the methodological challenges of studying AR—often through VR simulations.

To support this investigation, this chapter reviews literature across six areas. First, it examines how commuters currently use their travel time and the constraints that shape their behaviour. Second, it considers how immersion is conceptualised and measured across XR and media research, and establishes the sense used in this thesis. Third, it introduces ARTV as an approach to supporting media immersion through spatial augmentation. Fourth, it reviews the opportunities and barriers for XR use in transit. Fifth, it considers the state of consumer AR hardware and its implications for in-transit use. Finally, it discusses methodological approaches to evaluating AR using VR simulations, and identifies the gaps that motivate the empirical studies presented in this thesis.

2.2 Understanding Commuting and Passenger Behaviour

In order to understand how XR technologies might improve the commuting experience, it is first necessary to examine how commuters currently use their travel time as well as the conditions under which these activities occur.

2.2.1 Commuting as a Use of Time and Health Problem

According to the United Kingdom's Department for Transport's *Transport Statistics Great Britain: 2024 Domestic Travel* (the most recent published statistics), the average person spends 28 minutes each day commuting to work, for a total of 56 minutes spent commuting each day [179]. This figure increases significantly for rail passengers, who spend 64 minutes getting to work - 128 minutes commuting daily. South Korea has a significantly longer average commute than the UK at 101 minutes spent commuting each day (approximately 50 minutes each way). A study comparing commuters with long (>60 min) commutes in South Korea found that they had a significant association with depressive symptoms for both men and women when compared to those with short (<30 min) commutes, even when accounting for social factors [91].

It is not just in South Korea that commuting was found to negatively impact commuter health. Studies in Sweden [56] and the UK [84] found that longer commutes correlated to increased health issues such as stress, fatigue, and poor sleep, and an increase in visits to the doctor. Longer passenger trips raise stress and reduce mood. Evans and Wener. [47] showed that, among New York suburban rail commuters, each additional commuting minute was linked

to higher salivary cortisol and greater perceived stress. Likewise, Kahneman et al. [74] and Kahneman and Krueger [73] found, through time-use surveys, that commuting ranked very low on *instant enjoyment* and high on *stress* compared to other daily activities. In short, longer commutes consistently correlate with worse mood, higher stress levels, and more mental health complaints.

However, in certain cases the commute can have some benefits. Jain and Lyons [68] argue that travel time, on an individual level, can be thought of as a *gift*. It provides a *Transition Time* as individuals transfer from different roles and settings, such as work and home. People can use time in transit to mentally prepare for the demands of the destination activity or to unwind and de-stress after work before arriving home. It additionally affords time for solitude and relaxation, so-called *Time Out* in which the travel space becomes a protected, cocooned environment. Finally, the commute can also allow for engaging in *Equipped Time*; in which by bringing objects such as mobile phones, laptops, books, and music along with them, the traveller can use this time productively or for entertainment.

Cornet et al.'s concept of worthwhile travel time defines and expands the notion of beneficial commuting. Cornet et al. [35] argues that travel is worthwhile if it is in some form one or more of enjoyable, productive, or improving the individual's health (such as an active commute - walking or cycling). One of the questions in the survey they administered asked if participants would teleport to their given destination had they the option. They found that, if passengers deem their commute or travel worthwhile, then they do not wish to skip their commute by teleporting to the destination. They concluded that if the commute was affecting the individual negatively the commuter would opt to skip the commute entirely.

We can also examine the impact of commuting on individuals through the lens of SWB, which refers to a person's self-evaluated quality of life, encompassing both affective states (e.g., stress, enjoyment) and cognitive evaluations (e.g., satisfaction with daily activities). Ettema et al. [45] argue that the welfare effects of travel cannot be captured by objective measures such as travel time or distance alone, and propose a framework in which travel contributes to SWB through three routes: by enabling participation in goal-directed activities, by imposing travel-related stress that diminishes wellbeing, and through the affect associated with the experience of travelling in itself. They propose that the impact of commuting on SWB depends on the conditions under which travel occurs. Factors such as the purpose of the journey, anticipated activities at the destination, and the experiential quality of the journey itself all contribute to how commuting is subjectively evaluated. In this sense, SWB acts as a holistic measure that captures how these contextual and experiential factors are integrated into an individual's overall assessment of their commuting experience. This perspective is particularly relevant to this thesis, as it frames commuting as an experiential context in which factors such as engagement, comfort, and attentional demand contribute to overall wellbeing. Consequently, interventions such as immersive media and ARTV Ornamentation must be understood not only in terms of

functionality, but in terms of how they shape the subjective quality of the passenger experience.

These findings suggest that improving the quality of the commuting experience, rather than eliminating travel time altogether, represents a viable pathway for enhancing commuter wellbeing. If passengers are able to engage in enjoyable or meaningful activities during their journeys, travel time may be transformed from a burdensome necessity into a more worthwhile and rewarding experience. This observation has important implications for emerging technologies such as XR, which have the potential to support such engagement by transforming travel time into immersive media experiences.

2.2.2 Activities Undertaken by Commuters

Identifying the activities passengers already undertake during transit provides insight into which experiences might be meaningfully enhanced through technological intervention. This thesis limits itself to passengers: those travellers who do not have primary responsibility for navigation or vehicle control, a category that excludes drivers, walkers, and cyclists. This scope was chosen because the challenges of developing XR technology for passengers differ substantially from those faced by travellers operating their own mode of transport, whose attention is necessarily committed to the travelling task.

The National Travel Survey [48] seeks to understand commuting habits such as transportation mode and duration. The United Kingdom's National Rail Passenger Survey (NRPS) [103] looked at commuting habits for rail passengers specifically, and the activities undertaken during the commute. Transport Focus, the body that administers the NRPS, removed the passenger activity question after 2014, meaning data are available only from 2004 (n=25,596), 2010 (n=27,556), and 2014 (n=27,812). Lyons et al. [103] analysed these datasets and found an increase in the prevalence of 'ICT-enabled activities' such as watching videos, checking emails, internet browsing and using social media, with a corresponding decrease in *low-tech* activities such as sleeping, reading for leisure, talking to other passengers, window gazing or people watching, and being bored. Given continued growth in smartphone ownership and capability since 2014, the prevalence of ICT-enabled commuting activities has likely increased further [131].

Singleton [161] surveyed 656 commuters in Portland, Oregon in 2016, examining activity participation across automobile, nonmotorized, and transit modes. Of the sample, 48.7% were car drivers, 16.5% cyclists, 12.9% rail users, 12.5% bus users, 5.1% car passengers, and 4.3% walkers, with non-auto commuters deliberately oversampled. Exploratory factor analysis across 23 activities produced two interpretable groups: an ICT factor (texting, emailing or messaging; reading electronically; using social websites or apps) and a passive factor (viewing scenery or people watching; thinking or daydreaming) [161]. Singleton [161] frames the modal differences as a distinction between operating modes, where attention is committed to the travelling task, and riding modes, where it is not. For operating modes – auto drivers and cyclists – the mental

and physical resources required for the driving task permitted little beyond the most passive activities. Riding modes, comprising transit riders and auto passengers, supported a greater number and variety of activities, including more resource-intensive ICT activities: participation in at least one ICT activity was reported by almost 70% of transit commuters, compared with 11% across automobile modes and 6% among nonmotorized commuters. Considering texting, emailing or messaging alone, 60% of rail passengers and 52% of bus passengers reported the activity, though Singleton [161] found no significant difference in activity participation between rail and bus commuters.

Notably, Singleton [161] interprets much of this participation as coping rather than production. ICT engagement among transit riders was associated with dissatisfaction with travel time and longer commutes, leading him to conclude that these commuters may be using electronic devices to pass the time rather than to make productive use of it. This aligns with his later argument that the value of travel time savings (VOT)—the willingness of an individual to pay for a small reduction in travel time—may have been overestimated for autonomous vehicles [160], since the riding modes that most resemble autonomous vehicle (AV) travel already show multitasking oriented toward coping and entertainment rather than productivity. Motion sickness may further constrain productivity-oriented activities in AVs, with around two thirds of the population potentially experiencing motion sickness while riding in a car [39].

Journey characteristics, such as duration and environment, are likely to influence the likelihood that passengers engage in productive activities while travelling. A passenger undertaking a thirty-minute uninterrupted journey may be more inclined to take out a laptop or begin a task than a commuter making a short ten-minute trip. In addition to journey duration, the physical motion environment experienced by passengers may also affect the feasibility of such activities. Passengers subjected to frequent changes in acceleration or direction may therefore be less inclined to look down at a device: head-down, visually demanding tasks reduce access to external motion cues, can increase visual–vestibular conflict, and are associated with greater discomfort and increased motion-sickness risk in moving vehicles [82, 117].

Across these studies, passengers consistently prioritise media consumption and lightweight entertainment over productivity-focused tasks. However, current media consumption during transit is constrained by small displays, fragmented attention, and environmental distractions inherent to shared transport environments. XR technologies offer new affordances for enlarging displays, spatialising content, and integrating media into the surrounding environment, potentially transforming how passengers experience in-transit entertainment. Despite this potential, relatively little work has examined how XR might enhance everyday commuter media experiences in realistic transit settings. This gap motivates the first research question addressed in this thesis:

RQ1: *How could XR improve the current experience of immersive media on public transport?*

2.3 Media Immersion

Immersion is a widely used concept across both media studies and XR research, but it is not used in the same way across these literatures; Nilsson et al. [128] review existing definitions and find the term applied inconsistently both within and between fields, dividing extant usage into immersion as a property of a system, as a response to narrative content, and as a response to challenge. In XR research, immersion is commonly discussed in relation to the technological properties of a system and the psychological experience of presence that such systems may produce [166]. In film, television, and interactive media research, by contrast, immersion is more often used to describe perceptual, cognitive, and emotional absorption in mediated content [69, 149]. Clarifying this distinction is important for the present thesis, as the work sits at the intersection of these two traditions. Specifically, this research examines how XR display technologies can be used to support immersion in audiovisual media, rather than to create a fully simulated virtual environment in which the user is perceptually relocated. This thesis keeps the two senses apart (section 1.2). System-level properties such as spatial presentation, FoV, and environmental integration are treated as *immersive properties* that the studies manipulate, alongside content-level factors such as narrative, aesthetics, and thematic coherence; the user's absorption in the content is treated as *media immersion*, which the studies measure. In the context of passenger travel, where environmental distractions are numerous and attention is frequently fragmented, understanding how immersion is conceptualised, how it is produced, and how it is measured is essential for designing effective XR media experiences.

2.3.1 Presence and Immersion in XR Research

In this literature, immersion is often treated as a property of a technological system, while presence refers to the psychological experience produced by that system. Slater and Wilbur [166] describe immersion as the objective characteristics of a display system that determine the extent to which it can deliver surrounding, inclusive, and sensorimotor-consistent stimuli to the user. These characteristics include factors such as field of view, stereoscopic rendering, display resolution, tracking fidelity, and system latency. Systems that provide wider visual coverage, more accurate tracking, and stronger correspondence between user actions and system responses are therefore considered more immersive because they more effectively surround the user perceptually and reduce the influence of stimuli from the physical world.

The experiential outcome of immersive systems is typically discussed in terms of presence. Witmer and Singer [192] define presence as the subjective sensation of 'being there' within a mediated environment rather than in the user's physical surroundings. Presence arises when the perceptual cues and interaction affordances provided by a system are sufficiently convincing that users respond to the mediated environment as if it were real. Subsequent work has emphasised that presence does not emerge from technological immersion alone, but from the interaction

between immersive system properties and user cognition, including attention, involvement, and task engagement [164, 155]. As a result, presence is commonly understood as a psychological state produced when immersive technological conditions align with the user's perceptual and cognitive processes.

Much of the XR presence literature has focused on immersive virtual environments designed to elicit a sense of presence [166, 192, 155, 37]—with a substantial applied emphasis on simulation-based training in domains such as defence and medicine [126, 159, 4, 125]. In these contexts, design efforts typically aim to maximise perceptual replacement of the physical environment so that the virtual world becomes the user's dominant frame of reference. Related formulations within the VR literature, such as place illusion and plausibility illusion, likewise concern the extent to which users respond to a virtual environment as if it were the place in which they are located and the events occurring within it were really happening [165]. More recent syntheses have sought to consolidate these constructs. Skarbez et al. [162] survey presence and its related concepts, distinguishing place illusion from plausibility illusion and identifying coherence—the internal consistency of the scenario presented—as the system property from which plausibility arises. This distinction is useful for the present work: where place illusion concerns whether a user feels located in a virtual space, coherence concerns whether what they perceive behaves consistently, which remains relevant even when no virtual space is being substituted for the real one. Coherence is also relevant to ornamentation: augmentation that does not fit the content on screen risks undermining plausibility, a point that the findings of chapter 5 bear out and chapter 7 returns to.

A third illusion, embodiment, concerns the sense that a virtual body is one's own, and decomposes into self-location, agency, and body ownership [76, 162]. It is not engaged by the work presented within this thesis: participants are seated throughout, are presented with no virtual body or avatar, and the experiences are deliberately non-interactive, so neither agency nor body ownership is manipulated. Self-location remains relevant in a weaker sense, since a passenger's sense of where their body is situated bears on the confined-space and proximity constraints discussed in section 2.5.1, but it is not a construct this thesis measures.

Measuring presence in these contexts has proved challenging and contested. The most widely adopted approach uses self-report questionnaires such as the Presence Questionnaire developed by Witmer and Singer [192], which asks participants to rate their subjective sense of being in the virtual environment across dimensions including control, sensory fidelity, and distraction. Slater [165] has argued, however, that self-report measures are insufficient because participants know retrospectively that they were not actually present in the environment; his preferred indicator is behavioural response—whether participants act toward virtual events as if they were real, including physiological reactions such as postural adjustment or stress responses. This disagreement reflects a deeper tension in the literature between treating presence as a reportable subjective experience and treating it as an implicit perceptual and behavioural state.

Neither approach has achieved consensus, and the two do not always converge.

2.3.2 Media Immersion

Many forms of media consumption, particularly film, television, and video streaming, do not attempt to relocate the user into a simulated environment, but instead aim to capture attention and sustain engagement with audiovisual content presented within a bounded display. In these contexts, immersion is typically conceptualised differently. Rather than referring to the illusion of spatial presence within a virtual world, immersion describes the degree to which viewers become perceptually, cognitively, and emotionally absorbed in media content. Research on immersion in games, from which the Immersive Experience Questionnaire (IEQ) originates, and in audiovisual media has characterised immersion as a state in which attentional resources become focused on the mediated experience, reducing the salience of stimuli in the surrounding environment [69]. This form of immersion emphasises narrative engagement, attentional focus, and the temporary dissociation from the immediate physical context of viewing.

It is because the work presented in this thesis investigates engagement with audiovisual media content presented through XR displays, rather than navigation within fully simulated virtual environments, that the concept of *media immersion* provides the most appropriate analytical lens. Accordingly, this thesis does not operationalise immersion primarily in terms of classic VR constructs such as presence, place illusion, or plausibility illusion. Instead, it examines immersion as media-related absorption and engagement, and measures it using instruments designed for audiovisual and interactive media experiences, rather than presence questionnaires developed for virtual environments. This terminological commitment holds throughout the thesis: where immersion is discussed or measured in subsequent chapters, it refers to this media-specific sense unless otherwise stated, and the system-level properties that the studies manipulate are referred to as immersive properties (section 1.2).

2.3.3 Display Size, Enclosure, and Field of View

One of the most direct drivers of media immersion is display size relative to the viewer's field of view. Larger displays occupy a greater portion of the visual field, reduce the salience of peripheral distractions, and promote attentional focus on the media content [64]. Experimental work has generally found that larger displays increase presence and related responses [102, 64], although the effect on immersion may level off beyond a certain size: Rigby et al. [148] found lower immersion on a phone than on a laptop or monitor, but no difference between the latter two.

An alternative approach to immersion is found in immersive media formats such as 360° video and volumetric capture. Rather than enlarging a statically positioned display, these formats place the viewer inside a captured or reconstructed scene, allowing them to explore the

environment freely by looking around. They are usually viewed in an HMD, although surround domes and cylindrical projection screens are also used. Prior research has shown that such formats can produce strong feelings of spatial presence because the viewer becomes perceptually situated within the media environment rather than observing it through a bounded screen [106]. However, these formats also impose different interaction and attentional demands, as viewers must actively explore the scene to follow narrative events. Because the focus of this thesis is on enhancing conventional audiovisual media—such as films, television, and streaming video—through spatial augmentation, the work presented here concentrates on planar media experiences rather than fully immersive 360° or volumetric formats.

For planar media, increases in display size and enclosure have historically been used to enhance immersion by expanding the perceptual footprint of the media experience. This relationship between display size and the viewer's field of view has been exploited for decades in cinema design, from early widescreen formats to contemporary large-format theatre systems such as IMAX theatres, where large-scale projection, controlled lighting, and reduced visual distractions are used to heighten presence and narrative absorption [65].

Beyond traditional cinemas, planetariums and dome theatres further extend this principle by surrounding the viewer with content, so that much or all of the visual field is filled by the display [89]. In such settings, immersion is thought to arise less from resolution or image quality than from the viewer's reduced opportunity to attend to stimuli outside the mediated experience. Recent large-scale venues, such as the Las Vegas Sphere, exemplify this approach by combining extreme display scale with environmental enclosure to produce highly immersive media experiences [81, 167]. These strategies however, depend on fixed infrastructure, dedicated spaces, and controlled viewing conditions. They are therefore fundamentally incompatible with mobile and constrained environments such as public transport, where passengers are limited to small personal devices and must share space with others. Most commuters consume media on smartphones or tablets [161], where limited display size, narrow field of view, and constant environmental interruption restrict immersion. This contrast highlights a key design challenge for passenger XR: how can immersive media experiences be recreated in constrained mobile environments without the large displays or controlled conditions traditionally used to support immersion?

2.3.4 Environmental Theming, Immersive Cinema and the Limits of Physical Augmentation

In addition to increasing display size and enclosure, media immersion has also been achieved by extending media content into the surrounding physical environment through theming and multisensory augmentation. Recent immersive cinema formats demonstrate a growing audience appetite for media experiences that extend beyond passive viewing. Productions such as Secret

Cinema and other themed screening events integrate narrative settings, live performance, and environmental design to transform film viewing into a participatory experience [12, 51]. Atkinson and Kennedy [11] consolidate research on this growing field of *live cinema*, in which screenings are accompanied by events that take place around and in tandem with the film. These experiences combine film screenings with physical sets, props, live performances, themed food, sound design, and even scent, transforming spectators into participants within a constructed narrative world. Rather than simply watching a film, audiences are placed inside an environment that spatially, aesthetically, and sensorially aligns with the media being presented. The continued popularity, scale, and commercial success of these events provide clear evidence that viewers value environmental augmentation—such as props, theming, and contextual additions—as a means of increasing immersion [10]. At the extreme end of this design spectrum are purpose-built immersive venues such as the Las Vegas Sphere, which combines a 160,000 sq. ft. 16K wraparound LED display with spatial audio, haptic seating, and environmental effects to create large-scale immersive audiovisual experiences [168]. Productions such as the Sphere's 2025 presentation of *The Wizard of Oz* illustrate how traditional planar films can be adapted to such environments through a combination of visual effects and artificial intelligence. The original 1939 film was shot in a conventional 4:3 aspect ratio intended for rectangular cinema screens; in order to fill the Sphere's surrounding display, engineers used AI-assisted techniques including super-resolution to increase the image fidelity, outpainting to extend backgrounds and characters beyond the original frame, and performance generation to synthesise plausible motion for off-screen elements based on archival footage and production materials [169]. These methods effectively expand the visual field of the film beyond its original boundaries, transforming a traditionally framed cinematic work into an enveloping spatial experience that surrounds the audience and integrates with environmental effects such as wind, fog, and spatialised sound. While such approaches demonstrate the extreme possibilities of large-scale immersive cinema, they rely on specialised infrastructure and computational pipelines that are far removed from the constraints of everyday viewing contexts such as commuting. Comparable effects have been pursued at domestic scale through ambient lighting systems that extend on-screen colour into the surrounding room, and through projector-camera research systems such as Beamatron, which uses a steerable projector and depth camera to place responsive content on real surfaces [189]. These approaches share with immersive cinema the principle that the viewing environment can be recruited into the media experience, but differ in operating within an ordinary room rather than a purpose-built venue.

Importantly, many of these immersive cinema experiences are built around traditional 2D films rather than fully virtual or 360° media. Immersion is achieved not by reconstructing the media content itself in three dimensions, but by extending its themes, aesthetics, and narrative cues into the surrounding environment. This distinction is significant because it demonstrates that immersion can be enhanced through environmental augmentation even when the underlying

media content itself remains unchanged.

Despite their effectiveness, immersive cinema and physically themed environments are subject to substantial limitations. They are location-specific, resource-intensive, and reliant on physical infrastructure, bespoke props, and human performers. As such, they cannot be scaled to everyday contexts such as commuting, where space is constrained, environments are shared, and users are mobile.

Furthermore, physical theming is inherently inflexible. It cannot adapt dynamically to different media content, individual preferences, or changing environmental conditions. While immersive cinema demonstrates a clear desire for environmental augmentation, its reliance on fixed, physical installations makes it unsuitable for routine media consumption in passenger contexts.

These constraints highlight a clear gap: although audiences value additions and environmental extensions that enhance immersion, existing physical approaches cannot deliver such experiences in mobile, everyday settings. These observations suggest that the benefits of environmental augmentation may need to be reproduced digitally rather than physically in order to support immersive media in everyday environments. AR offers one potential approach, enabling media content to extend into the surrounding environment without requiring physical modification of space. This idea forms the basis of ARTV, which is explored in the following section as a potential mechanism for enhancing media immersion in passenger contexts.

2.3.5 Implications for AR-Based Media Immersion

AR offers a means of addressing this gap by enabling digital environmental augmentation without requiring physical modification of space. Through AR, elements commonly associated with immersive cinema—such as themed surroundings, ambient visual effects, and spatial extensions of the screen—can be recreated digitally and deployed anywhere, at any scale.

In passenger contexts, AR enables media experiences that move beyond a single floating display by augmenting the surrounding environment in ways that are thematically aligned with the content being consumed. This approach mirrors the principles underlying immersive cinema while avoiding its physical, spatial, and logistical constraints. Crucially, AR allows these augmentations to remain selective and adjustable, supporting immersion while preserving environmental awareness [107].

These observations motivate the exploration of ARTV ornamentation as a means of supporting media immersion in transit. By digitally extending 2D media into the passenger's environment—rather than recreating the media itself—ARTV has the potential to deliver aspects of immersive cinema in contexts where traditional approaches are infeasible. This, in turn, raises questions about how the amount, visibility, and type of ornamentation influence the balance between immersion and awareness, which are investigated empirically in the studies presented later in this thesis.

2.3.6 Measuring Immersion and Engagement

Throughout this thesis, immersion refers specifically to media immersion — the degree to which viewers become perceptually, cognitively, and emotionally absorbed in audiovisual content, as defined in section 2.3.2. To measure this, and related constructs of engagement and enjoyment, this thesis uses short-form versions of several validated questionnaires: the User Experience Questionnaire–Short (UEQ-S), the User Engagement Scale–Short Form (UES-SF) and the Film Immersive Experience Questionnaire (Film-IEQ). Combined, these instruments provide complementary coverage of cognitive, hedonic, and usability-related aspects of the experiences while keeping participant burden lower than it would be if the full version of each questionnaire were to be administered.

The selection of measurement tools was guided by three principal criteria. First, the instruments provide theoretical coverage across usability and pragmatic quality, affective and hedonic responses, attentional engagement, and media-specific immersion. As no single validated short measure spans all of these domains with adequate reliability, a combined approach is necessary. Second, all instruments used in this thesis have strong psychometric foundations, have undergone validation through factor analyses, and demonstrate evidence of internal consistency [157, 130, 149]. This ensures that the measured constructs map to established theoretical dimensions rather than ad-hoc or researcher-defined questionnaires. Third, the use of short forms ensures pragmatic suitability for experiments in which participants complete multiple tasks. Shorter questionnaires minimise fatigue effects, reduce dropout risk, and improve data quality while retaining validated factor structure.

It is also worth noting that several questionnaires have been proposed specifically for evaluating immersive AR experiences. One example is the Augmented Reality Immersion questionnaire (ARI), which was developed to measure users' perceived immersion when interacting with augmented environments [52]. The ARI focuses on dimensions such as engagement with virtual content, perceived realism, and the degree to which virtual elements are integrated with the surrounding physical environment. Similar AR-focused evaluation approaches often emphasise spatial presence and the perceptual coherence between virtual and real-world elements [23]. These instruments were developed for interactive, object-level augmentation, in which the user acts on virtual content placed in the physical environment. The experiences studied in this thesis are instead non-interactive viewing of bounded-display narrative media, the context for which the Film-IEQ was developed and validated. Although two of its subscales overlap conceptually with presence (discussed below), the Film-IEQ as a whole targets absorption in the content rather than in the augmented space. Consequently, the evaluation approach adopted here prioritises instruments that capture attentional engagement, narrative absorption, and experiential quality during media viewing.

The UEQ-S

The User Experience Questionnaire (UEQ) is a widely used instrument for capturing users' subjective impressions of interactive systems, designed to provide a rapid but comprehensive assessment of both pragmatic and hedonic qualities of interaction. The full UEQ comprises six dimensions—Attractiveness, Perspicuity, Efficiency, Dependability, Stimulation, and Novelty—which together reflect users' perceptions of usability, control, and experiential appeal. These dimensions are grounded in a distinction between pragmatic qualities (task-oriented aspects such as efficiency and clarity) and hedonic qualities (non-task-oriented aspects such as stimulation and novelty), both of which contribute to overall product attractiveness [90].

The UEQ-S is a validated short form of this instrument, consisting of eight items that capture these two higher-level dimensions—pragmatic and hedonic quality—while substantially reducing completion time. It was explicitly designed for scenarios in which the full UEQ is impractical, such as studies requiring repeated measurements across multiple conditions within a single session [157]. Empirical validation shows that the UEQ-S closely approximates the results of the full questionnaire, providing a reliable estimate of overall user experience while minimizing participant burden [157].

In this thesis, the UEQ-S is used to capture a concise measure of overall experiential quality across multiple XR viewing conditions. The experimental designs employed require participants to evaluate several conditions in sequence, making longer questionnaires impractical due to fatigue and reduced response quality. In this context, the UEQ-S provides an appropriate balance between measurement validity and efficiency: it preserves sensitivity to both usability-related (pragmatic) and affective (hedonic) aspects of experience, while remaining sufficiently lightweight for repeated within-subject comparisons. As such, it serves as a complementary measure to more domain-specific instruments (e.g., immersion or engagement scales), enabling consistent comparison of global user experience across conditions.

The UES-SF

The User Engagement Scale (UES) was developed as a multidimensional instrument for capturing users' engagement with interactive systems, conceptualising engagement as a combination of cognitive, emotional, and behavioural investment in an experience [129]. Subsequent refinement and validation of the instrument support a four-factor structure—Focused Attention, Perceived Usability, Aesthetic Appeal, and Reward—which is preserved in both the full and short forms [130]. These dimensions respectively capture sustained attentional absorption, usability-related effort, sensory and aesthetic appeal, and the intrinsic value or desirability of the experience.

The short form UES-SF was introduced to enable efficient measurement of engagement in contexts where multiple conditions must be evaluated without imposing excessive participant burden [130]. Empirical validation demonstrates that the UES-SF provides a reliable approx-

imation of the full scale while retaining sensitivity to key engagement constructs, making it suitable for repeated-measures experimental designs.

In this thesis, the UES-SF is used to assess engagement across multiple XR viewing conditions, complementing broader experience measures such as the UEQ-S and domain-specific immersion measures. In particular, it enables the differentiation between general experiential quality and deeper engagement-related constructs such as attentional absorption and perceived reward. This is especially important in XR media contexts, where conditions may be similarly usable but differ substantially in how strongly they capture attention or motivate continued interaction. The brevity of the UES-SF further makes it appropriate for the within-subject study designs employed here, where participants evaluate several conditions sequentially.

The Film-IEQ

The IEQ was originally developed to measure immersion in interactive digital environments, conceptualising immersion as a combination of attentional focus, emotional involvement, and reduced awareness of the real world [69]. The IEQ captures immersion as a multifaceted construct, including cognitive absorption, affective engagement, and dissociation from the physical environment.

Building on this foundation, the Film-IEQ was developed as a domain-specific adaptation for film and video contexts, where immersion arises from narrative engagement rather than interaction [149]. The Film-IEQ comprises four validated subscales—Captivation, Real-World Dissociation, Comprehension, and Transportation—which together capture key aspects of film-based immersion, including attentional capture, narrative understanding, perceptual transport, and disengagement from the surrounding environment.

Two of these subscales, Real-World Dissociation and Transportation, share conceptual territory with presence constructs such as place illusion, in that both concern a reduction in engagement with the physical viewing environment. The two are nonetheless conceptually distinct: place illusion refers to the sensation of spatial relocation into a navigable environment that the viewer can explore by looking around, whereas Real-World Dissociation and Transportation were developed and validated for bounded-display narrative media, where no such relocation is possible. Real-World Dissociation captures attentional withdrawal from the surrounding environment, while Transportation captures absorption into the narrative itself — a related but distinct process from spatial relocation. The overlap between these constructs is therefore conceptual rather than operational: both describe a viewer's disengagement from physical surroundings, but attribute it to different perceptual mechanisms—spatial relocation in the case of presence, and narrative absorption in the case of Film-IEQ. This is consistent with the terminological distinction set out in section 2.3.2.

Unlike more general user experience or engagement measures, the Film-IEQ is specifically sensitive to narrative absorption and attentional engagement in audiovisual media. This makes

it particularly appropriate for evaluating XR media experiences, where the primary goal is not only usability or engagement, but the extent to which users are drawn into the content while potentially disengaging from their physical surroundings.

In this thesis, the Film-IEQ is used to assess immersion across different XR viewing conditions, complementing broader measures such as the UEQ-S and UES-SF. This distinction is important in the context of ARTV and in-transit media consumption, where immersion must be considered alongside environmental awareness. The Film-IEQ enables direct measurement of this immersion component, allowing the relationship between immersive engagement and environmental awareness to be examined across conditions.

2.3.7 Summary

Across the three instruments, this thesis captures a multidimensional view of immersion and engagement grounded in validated psychometric frameworks. The UEQ-S offers a concise assessment of pragmatic and hedonic experience; the UES-SF provides a reliable measure of attention, usability, aesthetics, and motivational reward; and the Film-IEQ introduces specialised sensitivity to narrative and perceptual transport. Collectively, these instruments offer comprehensive coverage of usability, hedonic quality, attentional engagement, and narrative immersion.

As discussed above, the evaluation approach focuses on media-related immersion and engagement rather than spatial presence within augmented environments. The selected instruments therefore provide a theoretically grounded and practically feasible framework for assessing immersive media experiences in the passenger XR studies presented in this thesis.

2.4 Supporting Media Immersion through ARTV

ARTV is the idea that the experience of watching television can be enhanced using AR [182, 71, 70]. This can take the form of contextual information about what is being viewed displayed outside the bounds of the television, additional screens showing alternative camera angles at a sports event, additional information about an animal during a nature documentary, or even an animated model of that same animal in the room with the user—for example, a turtle swimming through the air in the living room whilst a nature documentary is playing [150]. ARTV experiences can vary widely, not only in the types of augmentation but also in how they are displayed. The television itself may be a physical set, viewed through or surrounded by augmentation, or a virtual display rendered by the headset; the work in this thesis takes the latter form. Some ARTV systems, such as Microsoft's IllumiRoom, which projected graphics onto the walls surrounding a television to extend game and film content beyond the screen (fig. 2.1) [71], assume a stationary home setup, but portable devices such as AR glasses are likewise compatible with ARTV. While originally conceived for domestic viewing contexts, ARTV is particularly well-suited to



Figure 2.1: IllumiRoom, a proof-of-concept system that augmented the physical environment surrounding a television to enhance interactive experiences. Using a 3D scan of the room (a), the system could (b) directly extend the field of view of the game, (c) selectively render scene elements, and (d) augment the appearance of the physical environment, here rendered as a cartoon. Reproduced from Jones et al. [71].

passenger environments where physical constraints prevent large displays or physical theming. However, most prior ARTV work has focused on domestic or stationary viewing contexts. Much less attention has been given to how ARTV might function in constrained mobile environments such as passenger transport, where users must balance media immersion with environmental awareness. This thesis therefore investigates ARTV ornamentation in passenger contexts, examining how the placement, density, and type of augmentations influence the balance between media immersion and environmental awareness.

Vatavu et al. [183] suggest the use of the ARTV continuum to better describe the variation of ARTV experiences and the relationships between them. On one axis is the *TV Real-Virtual Continuum* - from a normal television set (real), to a conventional TV with off-TV augmentation, to a fully virtual TV with no physical screen. On the other axis is the *World Real-Virtual Continuum*. This axis goes from a fully real living room environment, to an augmented one, to a fully virtual watching environment. This thesis proposes the use of themed augmentations, presented in an AR HMD, to enhance the entertainment experience of a commuting passenger. This work falls into the ARTV continuum category of a fully virtual TV with an augmented real-world environment, in which content themed to the media being watched augments the user's surroundings. Using Vatavu et al.'s ARTV Continuum, this proposal occupies position 7 — *Augmented World / Virtual TV* (fig. 2.2). This position is particularly consequential in passenger contexts. A virtual display can increase apparent screen size beyond the constraints of handheld devices, while world augmentation can extend media theming into the surrounding carriage. However, increasing the virtuality of the display or the environment also changes how passengers perceive and attend to the shared, dynamic physical setting around them. Passenger ARTV therefore sits at a point on the continuum where immersive potential and situated awareness demands are directly entangled: the same augmentations that deepen media engagement may also alter a passenger's perceptual relationship with co-present others, personal belongings, and journey cues.

Vatavu et al.'s two-axis continuum is itself a descendant of Milgram and Kishino's Reality–Virtuality continuum, which is widely cited but represents only part of what was originally proposed [121]. The same work sets out a three-dimensional taxonomy — extent of world knowledge, reproduction fidelity, and extent of presence metaphor — which, as Skarbez et al.

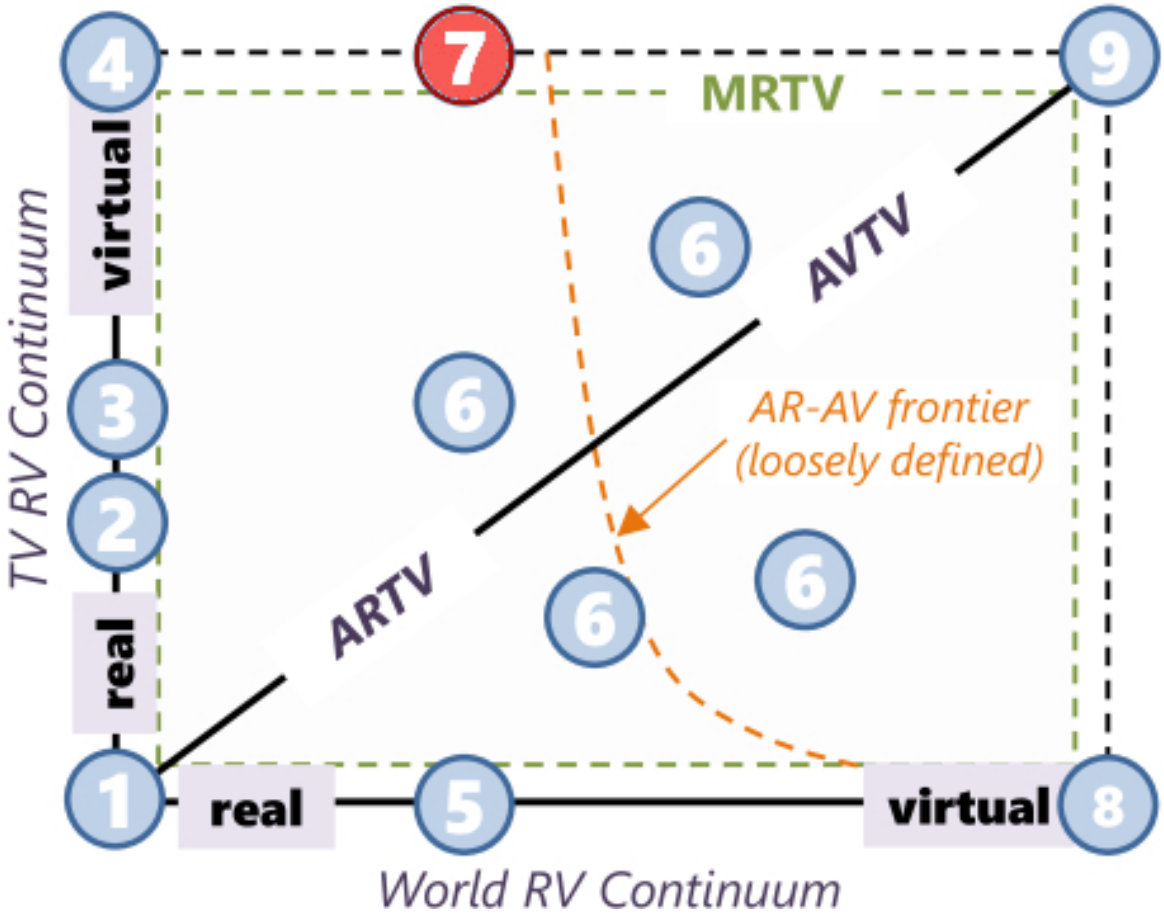


Figure 2.2: The ARTV Continuum, positioning ARTV experiences by the virtuality of the television (one axis) and of the viewing environment (the other). The passenger ARTV studied in this thesis occupies position 7, *Augmented World / Virtual TV*. Adapted from Vatavu et al. [183].

[163] observe, the field largely set aside in favour of the single axis. Skarbez et al. revisit the continuum directly, arguing that it is discontinuous and that conventional VR is itself a form of mixed reality, and propose a revised taxonomy of extent of world knowledge, immersion, and coherence. This matters for passenger ARTV because the dimension most relevant here — how much the system knows about, and preserves of, the real environment — is precisely the one the reduction to a single axis discards.

Users have expressed strong interest in immersive ARTV applications. Popovici and Vatavu [140] conducted an exploratory study with the aim of understanding user preferences and prioritisation of ARTV as a concept. They also investigated 20 distinct ARTV scenarios; such as virtual objects coming out of the screen, watching different perspectives of the same movie scene on different displays, and having additional content such as character names displayed next to the TV screen. The third, fourth, and fifth highest rated concepts – in terms of perceived value – revolved around increasing the field of view of content beyond the display of the TV and having objects come out of the TV [140].

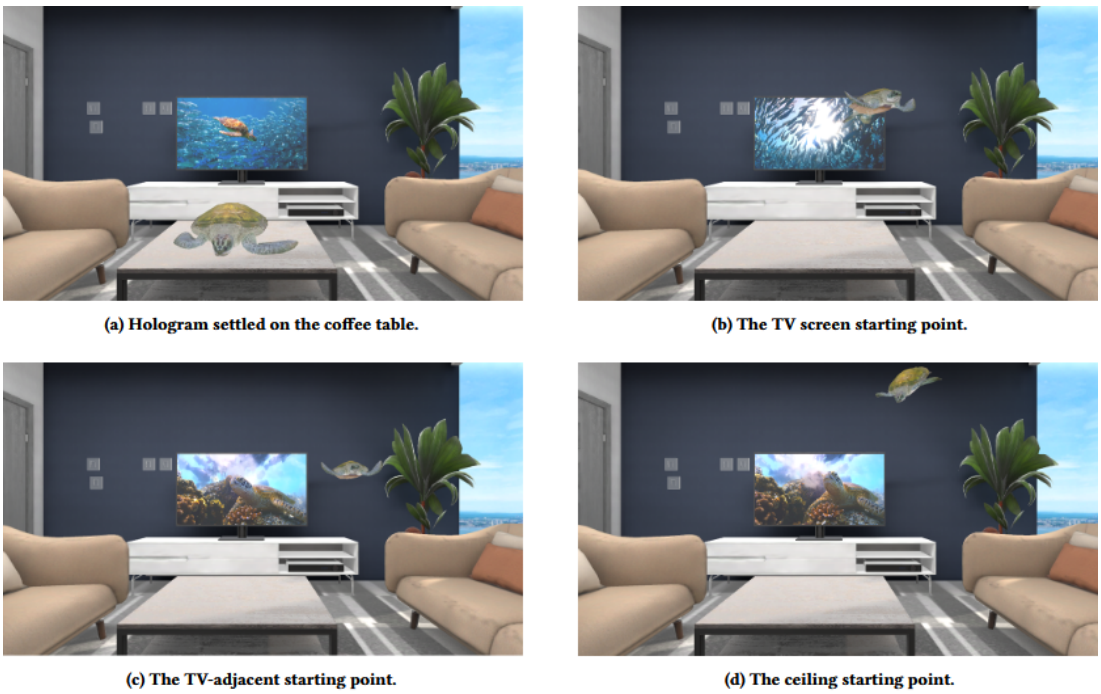
ARTV has additionally previously been categorised into distinct dimensions and themes

[151], which were then revised and updated [152]. Saeghe et al. [152] conducted a systematic review and proposed an ARTV design space with six dimensions (abstraction, interaction, time, display, context, editorial control). In the updated framework, the scenario that this thesis targets can be described as:

- Framed TV with Unframed AR Content
- Visual and Auditory Experience
- Pre-recorded Viewing
- Synchronised AR Content
- One Viewer Experience with Bystanders
- No Interaction
- Unfamiliar Viewing Location
- Public Indoor Location
- Content Augments the Space
- Content Affected by the Geometry of the Space
- Content Does Not Augment Objects

These categories (drawing on Saeghe et al.'s taxonomy) clarify the specific ARTV design scenario examined in this thesis: a single passenger viewing pre-recorded media in a shared public environment, where augmentations extend into the surrounding space but do not require interaction. This framing distinguishes passenger ARTV from domestic installations and highlights the importance of carefully controlling augmentation visibility and placement.

Saeghe et al. [150, 153] also performed user studies of ARTV in practice to understand the impact of environmental augmentations on viewer engagement. In one study using a Microsoft HoloLens, they augmented a nature documentary with a programme-related, lifelike hologram that appeared in the surrounding environment synchronised with events in the video (see fig. 2.3). The results indicated that such content-linked augmentations could increase viewer engagement with the media experience [150]. However, a subsequent study revealed a clear engagement–distraction trade-off: when the hologram (a life-sized turtle) was made more visually engaging or mobile, participants tended to pay less attention to the television content. As one participant noted, 'the more engagement you get with the hologram, the more distraction from the [TV] clip you get also' [153]. These findings suggest that while AR augmentations can enrich media experiences, their salience and behaviour must be carefully controlled to avoid diverting attention away from the primary content. For this reason, the ARTV ornamentations used in this thesis are deliberately designed to be supportive rather than disruptive. In passenger contexts, this engagement–distraction trade-off directly maps onto the broader immersion–awareness challenge, as overly salient augmentations risk drawing attention away from both the primary media and the surrounding environment.



(a) Hologram settled on the coffee table.

(b) The TV screen starting point.

(c) The TV-adjacent starting point.

(d) The ceiling starting point.

Figure 2.3: The virtual living room used to evaluate ARTV content, showing the hologram settled on the coffee table and the three other points that the hologram could start moving from. Reproduced from Saeghe et al. [153].

2.5 XR for Passengers: Opportunities and Barriers

It is likely that passengers who currently consume media on small-screened mobile devices will increasingly turn to XR devices for entertainment as the form factor becomes lighter, more socially acceptable, and more widely available. This transition is already visible in the consumer market: recent industry reports indicate that growth in XR device shipments is being driven less by traditional VR/MR headsets and more by lightweight smart glasses and display glasses, with vendors such as Meta, XREAL, Xiaomi, RayNeo, and Viture competing around everyday wearable form factors [66]. Some commercially available display glasses, in particular, foreground mobile media consumption and private large-screen viewing as central use cases — devices such as the XREAL One Pro [200], for example, offer a portable virtual-screen experience with a 57-degree field of view, and review coverage has explicitly framed such devices around replacing small handheld or seat-back screens during travel [57]. Alongside these display-oriented devices, AI smart glasses are moving toward mainstream eyewear designs, with Meta's Ray-Ban Display [115] introducing an in-lens display in a socially wearable frame, and Google working with Samsung, Warby Parker, and Gentle Monster on Android XR eyewear expected in 2026. These developments do not yet mean that immersive passenger XR is mainstream, but they indicate that the hardware ecosystem is moving toward lighter, more wearable devices that could plausibly be used during everyday travel. There has been significant interest in bringing passenger XR to the car. Hock et al. [61] developed an in-car VR prototype, and McGill et al.

[108] provide an open-source platform for passenger XR research. There have also been developments outside of academia from well known companies. Audi has partnered with Holoride [13, 62] to bring passenger XR to consumers and Meta has partnered with BMW to do the same [112]. There has even been work done on using headsets whilst driving: Ghiurău et al. [53] worked with Volvo to create an MR prototype for drivers and Cupra created a prototype of a track racing experience for a driver using a VR headset whilst driving a real car in a car park [186]. However, little work has been undertaken to investigate how best to specifically make the media consumption experience in private vehicles and on public transportation more immersive and engaging. Media consumption is a particularly compelling target for passenger XR because it aligns with existing commuter behaviour, tolerates intermittent attention, and does not require sustained interaction or productivity.

Prior work offers partial evidence about which activities passengers would want to undertake with XR; chapter 3 addresses this question directly. Li et al. [94] investigated VR usage for productivity in cars and found participants were reticent to use it, thinking it better suited to watching videos rather than the reading and writing tasks that they would usually do on commutes. Wilson et al. [190] used an online survey (n=1378) and in-depth interviews (n=18) to understand the tasks that people wanted to perform in AVs and the motivations for private AV ownership. 88% of the 566 that answered the question on AV activities agreed or strongly agreed that they would undertake leisure activities when the vehicle was driving autonomously. Preference for leisure activities outranked that for other activities included in the survey: the remaining options were, in descending order of agreement, resting and sleeping, socialising, being productive, remaining in the driving position, a morning routine, and working with a colleague.

Entertainment is preferred over productivity not just for passengers in cars. Bajorunaite et al. [16] administered two surveys (N=60, N=108). The first concerned the use of VR on aeroplanes and found strong interest in the technology (65% somewhat or very interested). Respondents ranked their interest across the three use cases of entertainment, communication, and work, with entertainment ranking the highest. Reasons for this preference included the perceived set-up cost associated with using a VR headset and the belief that the system did not replace traditional devices for work. This mirrors Li et al.'s finding that passengers saw VR as better suited to watching video than to work. In the same vein, Williamson et al. [188] looked at the usability and social acceptability of VR devices on a plane. The biggest issue identified in their survey was the loss of awareness of the environment and the concerns surrounding that. The need for reality awareness within the VR experience was also identified by O'Hagan et al. [132] and Bajorunaite et al. [17].

Participants across these studies identified the potential of VR for viewing videos, but without ways to support environmental awareness the issues of personal safety, comfort, and awareness remain unresolved. A VR cinema that simply provides a large screen also uses only part of what XR offers: it enlarges the display but leaves the surrounding environment unused.

These concerns are not specific to cars and aircraft. Any in-transit XR experience must account for awareness demands that are much higher than those of a living room: passengers must track the vehicle's progress and remain aware of other passengers and of their own belongings, particularly where pickpocketing is common. The occlusive properties of VR are therefore not always desirable in these environments, and the public are reticent to use VR on public transport [16]. If AR is used instead, passengers may be able to enjoy the benefits of XR without the loss of visual awareness associated with VR.

Collectively, prior work suggests that passengers are interested in XR primarily for entertainment, but that fully occlusive VR—while well-suited for immersive media viewing—undermines the environmental awareness and social acceptability that passengers report needing in transit [16]. This creates an immersion–awareness trade-off: increasing perceptual immersion tends to reduce awareness of stops, nearby passengers, and personal belongings. What remains unclear, however, is how to design passenger media XR experiences that increase immersion without reducing awareness below the level passengers need for safety and journey management. This thesis addresses this gap by investigating AR-based media viewing with ARTV ornamentation, leading to **RQ2:** *How can ARTV with Ornamentation be used to balance media immersion and environmental awareness in passenger transportation?*

2.5.1 Barriers to XR Use in Passenger Transit: Physical, Cognitive, Perceptual, and Social-Normative Constraints

The engagement with specific activities and varying productivity of passengers has been found to be a product of the transport modality [161]. These differences are born of the constraints to which passengers are subjected by each modality. The following sections detail four core types of constraints – physical, cognitive, perceptual, and social-normative.

Physical Constraints

Although train passengers were generally found to be the most productive, as compared to other types of commuters, it is doubtful that this would be the case for all rail passengers. For example, passengers on the London underground, if standing during a standing-room-only peak-time commute, are unlikely to get out their laptop to write an email. A cross-country rail service with a table and a power socket, however, is a space much more conducive to productivity.

The motion environment experienced during travel also plays an important role in passenger comfort and activity feasibility. Acceleration and jerk in particular can influence how comfortable passengers feel when attempting to read, use electronic devices, or perform other tasks during transit [176].

Road-based public transport often exposes passengers to relatively frequent acceleration and jerk events arising from braking, acceleration from stops, and turning manoeuvres. For ex-

ample, field measurements on urban buses report longitudinal events of approximately -0.32 g (≈ 3.14 m/s^2) during heavy braking and around 0.16 g (≈ 1.57 m/s^2) when accelerating away from signals, with typical operational accelerations often within the order of 1–2 m/s^2 [50]. Lateral motion during turning can also be significant; passenger discomfort thresholds have been reported at average lateral accelerations of approximately 1.5, 1.75 and 2.0 m/s^2 for uncomfortable, very uncomfortable and extremely uncomfortable conditions respectively [127]. Analyses of bus braking behaviour similarly suggest operational envelopes of approximately 1.1–1.5 m/s^2 deceleration with jerk values of up to around 3 m/s^3 in routine service [40].

By contrast, rail rapid transit systems typically operate with smoother commanded motion profiles. Representative operating conditions reported in rail capacity studies include service braking of approximately 1.0 m/s^2 and jerk rates around 0.75 m/s^3, with surveyed mean initial accelerations of roughly 1.12 m/s^2 across systems [134]. Operational specifications for some UK rolling stock similarly constrain jerk during normal traction and braking, for example limiting it to around 0.5 m/s^3 [176].

These differences suggest that rail journeys may provide a motion environment more conducive to activities such as reading or working on electronic devices. However, it remains unclear whether the greater prevalence of productive activity among rail passengers arises primarily from these smoother motion characteristics or from the generally longer and less interrupted nature of rail commutes. These motion characteristics may also influence the suitability of immersive technologies during travel, as HMDs must operate within dynamic motion environments that can exacerbate discomfort or cybersickness.

This physical motion from the journey may induce sickness in users of a VR HMD. Motion sickness is commonly explained by sensory conflict theory, which proposes that discomfort arises when the signals received by the visual system conflict with those from the vestibular system and proprioceptors, as originally described by Reason and Brand [147]. Subsequent empirical work has provided support for this explanation [184].

A distinction should be drawn between motion sickness, which arises from physical motion acting on the vestibular system [147], and cybersickness, which occurs in the absence of physical motion and has likewise been attributed to visual–vestibular conflict [174]. Passenger XR is unusual in that both may operate together, since the vehicle supplies physical motion while the display supplies visual motion that may or may not correspond to it [139]. More recent investigations of cybersickness in virtual environments similarly link discomfort to visual–vestibular mismatch. For example, Teixeira et al. [174] found that participants experiencing cybersickness in a commercial VR game were significantly more likely to report *unexpected vection*—illusory self-motion that differed from anticipated motion cues—further supporting sensory conflict as a causal mechanism.

These issues are particularly relevant in passenger contexts where users are already experiencing real-world vehicle motion. Studies of VR use for passenger tasks under simulated vehi-

cle motion suggest that visual–vestibular mismatches can contribute to motion sickness, and that adding congruent visual motion cues can reduce sickness [139]. This suggests that immersive displays must carefully account for the motion environment in which they are used. AR may mitigate some of these effects in principle because the outside world—and thus real visual motion cues—remains visible, potentially reducing visual–vestibular mismatch. No study in this thesis involved vehicle motion, so motion sickness was not measured; chapter 7 returns to this limitation.

Transport environments are also challenging due to the physical constraints placed upon the user's movement. In a confined space, such as a train or aircraft seat, users cannot swing their arms and make large gestures to interact with the virtual world as they would do at home without risking contact with the seat in front or neighbouring passengers [191]. This mismatch is stark when comparing the interaction volumes assumed by mainstream room-scale VR (typically on the order of 1–2 m^2) with the available free space in transport seating (tens of centimetres in width and depth), as illustrated in fig. 2.4. Wilson et al. [191] operationalised this constraint by evaluating interaction within a $50 \times 50 \times 60$ cm volume chosen to approximate economy seating, contrasted against a 1×1 m at-home baseline. They found that techniques which improved task performance under constraint could still produce substantial *unsafe* behaviour: Linear Gain yielded performance and user experience closest to the at-home condition, but at the cost of frequent play-boundary violations and large arm movements [191]. Conversely, AlphaCursor substantially reduced arm movement and boundary violations (reported at $\sim$10% of trials, compared to 22% for Linear Gain and 31% for the at-home mapping), but with poorer performance and subjective experience [191]. For passenger media XR, this highlights that confined seating does not merely reduce comfort; it creates a design tension between interaction effectiveness and physical/social safety. This thesis therefore avoids making arm-driven interaction central to the experience. Instead, ARTV ornamentation is designed to be predominantly passive and supportive—enhancing immersion through spatial augmentation while requiring minimal deliberate interaction beyond occasional adjustment or dismissal when necessary.

Combined, these constraints define a highly specific design space for XR use in transit. In this thesis, the focus is therefore placed on *seated passenger contexts with limited physical movement and relatively low perceived motion intensity*, such as typical rail journeys outside of peak standing-room-only conditions. This choice reflects both practical feasibility—where users have sufficient stability and available attention to engage with media—and methodological control, allowing the effects of ARTV ornamentation on immersion and awareness to be examined without the confounding influence of extreme motion or severe physical restriction.

Accordingly, the systems and evaluations presented here are designed for constrained, predominantly passive media consumption scenarios, rather than highly interactive or movement-intensive use. This aligns with the broader goal of supporting in-transit media experiences that balance immersion with environmental awareness under realistic, but tractable, physical condi-

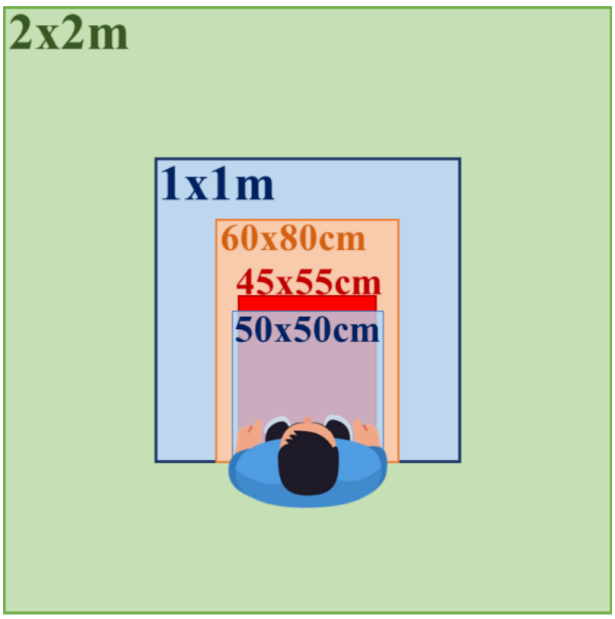


Figure 2.4: Interaction space assumed by commercial VR experiences (1×1 m and 2×2 m) compared with examples of constrained transport seating, and the spaces used by Wilson et al. (1×1 m at-home space and 50×50 cm constrained space, in blue). Reproduced from Wilson et al. [191].

tions.

Cognitive Constraints

Passengers on public transit operate under substantial cognitive and attentional constraints as they must divide their effort between monitoring their surroundings, tracking route progress, and managing personal safety while often-times simultaneously engaging in secondary mobile tasks such as checking messages, browsing social media, watching short videos, or listening to music. Lee et al. [92] investigated multitasking in the context of public transit passengers and found that the nature of travel-based multitasking is shaped by constant interference between the demands for cognition and attention and by the dynamic, interruption-prone nature of transit environments, which can easily lead to missed stops or reduced situational awareness. These findings underscore that technologies intended for use in transit must support—not compete with—passengers' need to remain perceptually open to environmental cues. In this context, AR presents a more suitable paradigm than fully occlusive VR, as it can layer digital content without blocking critical awareness of stops, crowd dynamics, and nearby passengers, and so aligns with the cognitive realities of transit multitasking rather than exacerbating them.

These constraints highlight that in-transit XR experiences must be designed for *divided attention* rather than full attentional capture. Mobile AR interaction compounds this challenge because virtual content must compete with real-world perceptual demands: prior work in mobile and in-the-wild AR contexts shows that environmental conditions — including traffic lev-

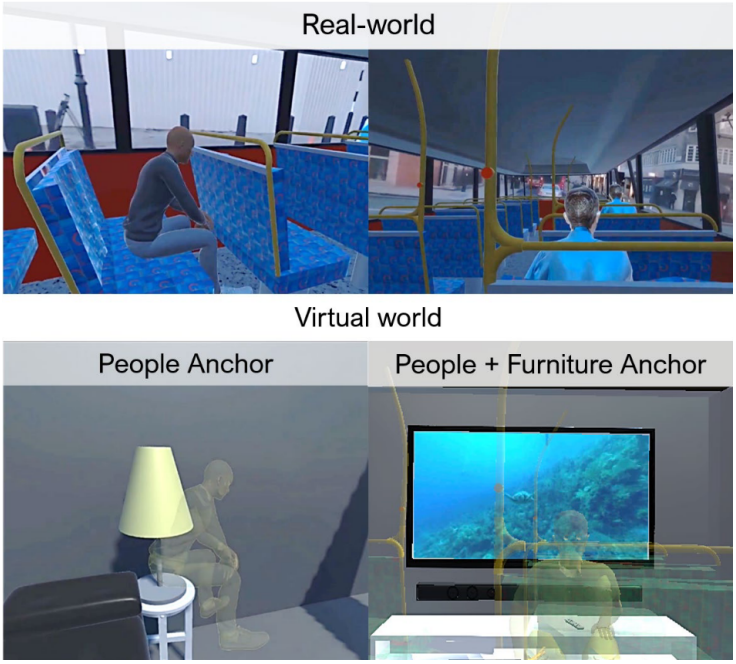


Figure 2.5: Reality Anchors, which selectively reintroduce real-world elements into an otherwise occlusive virtual environment. The upper images show the user's real surroundings (simulated virtually for the purposes of the study), the lower images show the anchor designs, featuring highlighted and transparent treatments. The original study depicted a bus; the same design logic applies to rail. Reproduced from Bajorunaite [15].

els, lighting variability, and the presence of bystanders — significantly shape how AR interfaces must be configured and used, suggesting that attention management rather than full attentional capture is the appropriate design goal [26, 28]. In this thesis, the focus is therefore placed on media consumption scenarios in which users engage with content while remaining partially attentive to their surroundings. This reflects a shift from traditional immersive design goals — maximising attentional absorption — towards supporting *attention management*, where experiences must be interruptible, glanceable, and compatible with ongoing situational monitoring. Accordingly, ARTV ornamentation is designed to enhance engagement with media content without demanding exclusive attention, enabling users to fluidly balance immersion and awareness within the cognitively dynamic conditions of transit environments.

Perceptual Awareness Constraints

Wearing HMDs on public transport brings with it an obvious reduction in the visual and aural perceptive abilities of the wearer. A passenger wearing headphones will have a reduction in their ability to hear the world around them and a passenger wearing a VR HMD will entirely forego their ability to see their surroundings.

Bajorunaite et al. [17] explored the concept of reintroducing awareness of the environment in VR through the integration of *Reality Anchors*—elements of the real world, such as other pas-

sengers and current stop signage—to keep the user aware (fig. 2.5). Two studies were conducted, investigating how Reality Anchors could alleviate concerns around safety, social acceptability, and comfort when using VR on public transport. Their findings revealed a clear tension between immersion and environmental awareness: anchors that provided stronger cues about the surrounding environment increased participants' perceived safety and comfort, but simultaneously disrupted immersion in the virtual experience. Participants reported that awareness of personal belongings and the location of other passengers was more important than general environmental awareness. In the second study [17], participants were able to toggle individual anchors on and off during the journey. While this allowed users to periodically check the real world, keeping anchors continuously visible was reported as distracting and reduced immersion. The study also found that preferences for which anchors should be visible varied depending on the content being consumed and the nature of the journey. The authors suggested that visually theming anchors to better match the media content might reduce this disruption, although this approach was not investigated in the study. These findings demonstrate that Reality Anchors do not resolve the immersion–awareness trade-off; rather, they highlight its limits. Anchors that are sufficiently salient to support awareness tend to disrupt immersion, while less intrusive anchors fail to provide reassurance or safety.

Since there was such variance in the reality anchors users wanted and as there was a desire expressed to make the visualisation of the reality anchors match the content being consumed [17], it makes sense to approach this problem from an AR perspective where the necessary reality anchors are included naturally by virtue of the experience being grounded in AR. The way that we provide the user with reality awareness differs drastically between the VR and AR experience. VR relies on sensing and reintroducing elements of reality, which entails the problem of identifying items such as personal belongings and subway information signs to represent them in VR. AR reality awareness on the other hand is the extent to which we can occlude the user's perception of reality before we diminish their awareness past a comfortable threshold. If the virtual content is not rendered over objects that the user needs to be aware of, then users will maintain awareness.

By utilising AR instead of VR, the challenge of explicitly reconstructing environmental awareness within the system can be reduced. VR-based approaches such as Reality Anchors require the system to recognise and selectively represent elements of the real world—such as other passengers, signage, or personal belongings—in order to maintain user awareness. In AR, however, the real environment remains directly visible to the user, meaning that awareness of these elements is not dependent on computational reconstruction in the way that VR requires. However, visual access to the real world should not be equated with situational awareness. A systematic review of AR systems designed to support situational awareness identifies information presentation design as a critical factor, suggesting that awareness cannot be assumed to arise automatically from the visibility of the environment afforded by AR devices [194]. In

mobile and public contexts this distinction is particularly consequential: maintaining situational awareness requires not merely perceiving environmental cues but interpreting and anticipating their relevance [177]. Empirical work examining AR's influence across Endsley's three levels of situational awareness—perception, comprehension, and projection—indicates that AR can affect these levels differently, and that higher-order awareness is not guaranteed even when the environment remains visible [177]. In passenger contexts, this means that even an AR experience which does not occlude the physical world may nonetheless impair a passenger's ability to comprehend journey progress or anticipate when to disembark, if virtual content draws sustained attentional resources away from environmental monitoring. AR experiences must therefore focus on carefully placing virtual content within the environment so that it does not occlude important real-world cues such as nearby passengers or personal belongings, while also being designed so that the attentional demands of the virtual content do not crowd out the higher-order interpretive processing that environmental awareness requires. If virtual content does occlude something important, the user can be given simple controls to reposition or remove it. A further benefit of AR is that it preserves real-world visual motion cues, which may help mitigate motion sickness by reducing visual–vestibular conflict; related evidence in passenger VR shows that adding congruent visual motion cues can reduce sickness [139].

The Immersion–Awareness Trade-off

The challenge is better described as a trade-off than as a continuum. The immersive properties of a configuration—in particular, how much of the real environment it occludes or replaces—act on two separate user outcomes: absorption in the media, and awareness of the surroundings (section 1.2). Fully occlusive VR favours the former by replacing the real world with virtual content, but in doing so removes visual access to the surroundings; a floating AR screen preserves that access but does little to extend the media into the environment. These are not two ends of a single scale. Awareness itself has more than one target: Reality Anchors attempt to restore awareness of the *real* environment inside a virtual one, whereas the engagement–distraction trade-off observed by Saeghe et al. [153] (section 2.4) concerns attention drawn towards the *virtual* layer at the expense of the primary content. Configurations can therefore differ in what they leave available to the passenger as well as in how immersive they are. Reality Anchors illustrate the difficulty of managing this within VR: anchors salient enough to support awareness tend to disrupt immersion, while less intrusive anchors fail to provide reassurance.

ARTV approaches the problem from the opposite direction. Rather than reintroducing elements of reality into a virtual world, AR-based media experiences begin with the real environment and selectively add virtual content, so that awareness depends on which parts of the environment the virtual content occludes and how strongly it draws attention. This connects to the extent-of-world-knowledge dimension discussed in section 2.4: the more a system knows about the passenger's surroundings, the more precisely augmentation can be placed so that it

does not obscure other passengers, belongings, or journey information.

Prior work therefore reveals a persistent tension between immersion and awareness in passenger XR, and suggests that AR with ornamentation may enhance media immersion without fully occluding the environment. This motivates an investigation into how such designs affect both outcomes (RQ2 - *How can ARTV with Ornamentation be used to balance media immersion and environmental awareness in passenger transportation?*).

Social-Normative Constraints

Aside from the demands on the passenger's awareness resulting from the unfamiliar, dynamic, public environment, passengers also experience constraints arising from social expectations. In shared transport environments, behaviour that attracts attention or appears socially unusual may lead to discomfort for both the user and surrounding passengers. For example, Koelle [78] highlight that wearable technologies with discreet, glasses-like form factors are perceived as more socially acceptable than bulky HMDs.

Empirical studies examining immersive technologies in public spaces further show that conspicuous XR devices can disrupt co-located social interactions and generate awkwardness for both users and bystanders. Bystanders can shape XR use by making users self-conscious about their appearance and behaviour, requiring social negotiation alongside engagement with the virtual content, and introducing attentional fragmentation as users divide awareness between the experience and co-present strangers [170]. Users may additionally adopt strategies such as concealing the device or altering their seating position in order to minimise social friction [133]. More broadly, research on XR adoption in shared environments identifies device visibility, perceived intrusiveness, and bystander reactions as key determinants of social acceptability [42, 188, 17, 133].

These concerns become particularly salient in transit settings where strangers must coordinate movement in confined spaces and rely on subtle social cues to maintain comfort and safety. Fully occlusive VR headsets—which block eye contact, conceal facial cues, and inhibit environmental monitoring—run counter to these expectations. In contrast, lightweight and glanceable AR modalities better align with the normative requirements of public travel by allowing passengers to remain visibly and socially attuned to those around them while still engaging with digital content. This alignment between social acceptability and environmental awareness positions AR as the more appropriate XR approach for transit contexts, where personal technology must coexist harmoniously with shared public space.

2.6 AR Hardware Constraints for ARTV: Consumer Reality (2026) and Plausible Advances (2030–2035)

ARTV is constrained as much by near-eye hardware as by rendering or tracking. Contemporary reviews of AR/VR display systems repeatedly identify limited FoV, insufficient optical efficiency/brightness, and the absence of convincing occlusion (virtual content correctly blocking the real world) as persistent bottlenecks that prevent current headsets from matching everyday visual conditions and human visual capabilities [197]. These constraints matter for ARTV because they directly bound (i) how much of the user's visual field is available for ornamentations in addition to the screen, (ii) how *solid* ornamentations can appear against real-world lighting, and (iii) whether ornamentations can be used for socially meaningful spatial effects (e.g. soft occlusion/shielding in passenger contexts).

2.6.1 The Consumer Reality of AR in 2026: Device Families and Hard Limits

A practical way to frame consumer AR hardware in 2026 is as three partially overlapping device families, each with different implications for ARTV experiments: lightweight 'display glasses' optimised around a floating personal screen; *optical see-through waveguide headsets*, primarily enterprise; and VST (pass-through) headsets. The first two differ mainly in tracking and registration. Display glasses such as the XREAL range are typically tethered to a phone or computer and present a screen that floats in front of the wearer, with limited or no anchoring to the surrounding environment. Waveguide headsets such as the HoloLens 2 track the wearer's position and can register content to the environment, but offer a narrow FoV.

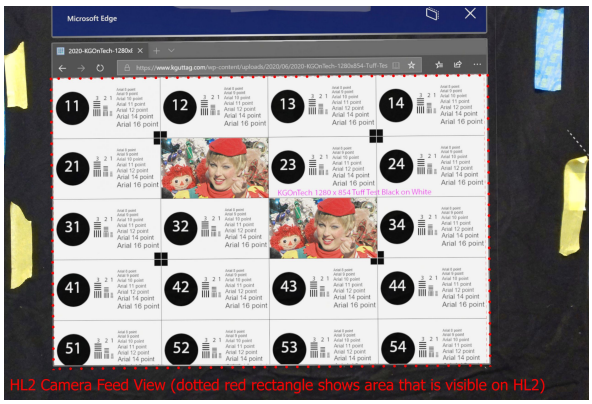
In the term *waveguide headset* a *waveguide* is a thin, transparent optical element (often laminated into a visor or lens) that guides injected light from a microdisplay into the eye while remaining see-through. A waveguide transports the image by total internal reflection and outcouples it toward the eye; the mixing of virtual and real imagery follows not from any combining function of the waveguide itself, but from the fact that the element remains largely optically transparent, so ambient light passes through it unaltered. This typically imposes tight constraints on FoV, eyebox, and optical efficiency [118, 197].

FoV constraints and why they matter for ornamentation-heavy ARTV. Recent AR devices such as the Xreal Air 2 Ultra (2024), and Microsoft Hololens 2 (2019) have a limited FoV (52° and 43° horizontal respectively) [199, 118]. By contrast, VR headsets have larger FoVs, e.g. the Meta Quest 2 has a Horizontal FoV of 104°, with the Quest 3 increasing to 110°, but still far less than human vision. If we compare these values to human vision, the gap becomes clearer: binocular horizontal visual field extent is commonly stated as approximately $\sim$214°,

with a single eye's FoV of approximately $\sim$107° [171]. This makes mainstream consumer VR headsets markedly narrower than the human peripheral envelope, and optical see-through AR headsets narrower still. Wider devices exist at both ends of the market. Among consumer hardware, the Pimax Crystal Super offers a modular optical system whose widest configuration provides a stated 140° horizontal FoV, traded against pixel density [137]. In the enterprise segment, the StarVR One reached 210° horizontal and 130° vertical, approaching the binocular extent of human vision [180, 34]. Such devices demonstrate that the narrow fields of view characteristic of mainstream headsets are not a fundamental limit, but they remain specialist products. The StarVR One was sold only to enterprise customers and neither device resembles the lightweight hardware a passenger would plausibly wear in transit. A distinction should also be drawn between field of view, the visual angle visible at any one moment, and field of regard, the total region around the user that can be viewed by moving the head [24]. For a fixed display the two effectively coincide, since turning the head reveals no further content. For a head-mounted display they diverge: content can be placed anywhere around the user, so the field of regard is the full surround, but at any moment the user sees only the device's field of view, which replaces the wider natural field of the eyes. In VST devices this limit applies to the view of the real environment as well as to virtual content; in OST devices it applies only to the region that can be augmented. The values discussed here are head-mounted fields of view, and therefore describe how much can be augmented at any one moment.

In ARTV experiences, device FoV does not merely constrain visual fidelity; it fundamentally shapes how ornamentations are perceived, how densely they can be presented, how much of the real environment remains visible, and consequently how much peripheral awareness a user has. One route to a wider augmentable FoV is VST AR on a VR headset such as the Meta Quest Pro (section 1.2). This provides a wider FoV within which ARTV ornamentations can be rendered, at the cost of lower visual quality, since the real world can only be captured and displayed at the quality of the device's cameras and screens. Therefore, FoV is not just a device spec; it determines the available peripheral canvas for ornamentation and likely moderates the immersion–awareness trade-off.

Contrast and occlusion: why optical see-through AR often looks 'ghostly'. A second, ARTV-critical constraint is contrast and occlusion: how 'solid' and visually authoritative ornamentations can appear against a real scene. Optical see-through displays add light to the user's view while real-world light still passes through to the eye, which makes dark tones and correct real-world occlusion difficult; hence virtual elements often appear semi-transparent or 'washed out' under bright ambient illumination [197]. This characteristic is not evident from most publicly available depictions of optical see-through output. Captures produced by the HoloLens 2, for instance, composite a separately rendered image over the headset's forward-facing camera feed rather than photographing the display itself, yielding a view that is both higher in quality



(a) Mixed Reality Capture, with the red rectangle marking the region actually visible through the display.



(b) Photograph taken through the optics, processed to approximate what the eye sees.

Figure 2.6: Two depictions of the same optical see-through display. (a) A Mixed Reality Capture composites a separately rendered image over the headset's forward-facing camera feed rather than photographing the display, producing a view both higher in quality and considerably wider in field of view than the device presents; the red rectangle marks the region actually visible. (b) A photograph taken through the optics with the camera and headset fixed, showing the reduced contrast, non-uniform brightness and colour variation characteristic of the display in use. Captures of the first kind constitute most publicly available depictions of optical see-through output. Both images reproduced from Guttag [55].

and considerably wider in field of view than the device presents to the wearer (fig. 2.6a) [55]. Consumer AR glasses frequently mitigate the contrast problem via global dimming approaches (e.g. tinted or electrochromic lenses) that reduce ambient light transmission, improving perceived contrast but not providing true, depth-correct per-pixel occlusion [27, 119].

What consumer ARTV looks like in practice. In current consumer devices, the most common form of ARTV is not rich augmentation of television content, but the rendering of the media display itself as a large virtual screen positioned within the user's environment [198, 5]. In terms of the ARTV design space proposed by Saeghe et al. [152] and the broader continuum articulated by Vatavu et al. [183], this corresponds to a relatively limited form of ARTV in which the primary contribution of AR is to replace or relocate the physical television display rather than to substantially augment the content or surrounding environment [151, 152, 183]. Put differently, AR is used here mainly to render the *TV* in space, while the media itself remains largely unchanged. Commercial viewing applications have begun to theme the viewing space—Disney+ on Apple Vision Pro, for example, offers environments drawn from its franchises [175], and the Apple TV app a virtual cinema [7]—but these replace the surroundings rather than augmenting them.

This configuration is nevertheless important, as it reflects the current consumer reality of AR media viewing. Previous work by Medeiros et al. [110] investigated the placement of AR

panels across a range of transport modes and found a trade-off between physical comfort (e.g. placing the display at head level in front of the user) and social comfort (e.g. not obscuring another passenger's face). Different users prioritised these factors differently, suggesting that such systems should ideally be customisable to accommodate user preferences for display placement. Medeiros et al. further suggested that additional AR modifications and diminished-reality techniques might help alleviate the social discomfort of looking towards another passenger while using a virtual screen. ARTV ornamentations could serve this role by placing supplementary virtual content in the environment so as to reduce this social discomfort while moving beyond the simple floating-screen paradigm.

Compared to VR, ARTV offers reality awareness by default and therefore has the potential to support immersive media experiences that remain compatible with passenger safety and environmental awareness requirements. However, designing such experiences presents a number of perceptual challenges. In AR, immersion cannot rely on fully replacing the visual environment; instead it must be achieved through carefully positioned augmentations that coexist with the real world. This means that the effectiveness of ARTV ornamentation depends not only on the media content itself, but also on perceptual design parameters such as how much augmentation is present, where it appears within the user's field of view, and how thematically aligned it is with the media being consumed. These factors are particularly important given the limited field of view and display constraints of current AR hardware, which restrict how much of the surrounding environment can be augmented at once. As a result, designers must carefully balance augmentation density, placement, and type in order to enhance immersion without obscuring important environmental cues.

Despite their importance, these perceptual parameters have not been systematically examined in passenger contexts. This gap motivates the following research question:

RQ3: *What is the impact of device FoV, ornamentation density, and ornamentation type on the media immersion–environmental awareness trade-off?*

2.6.2 Plausible Technology Advances 2030–2035

In time, OST headsets may be able to match the FoV of current VR HMDs without requiring solutions such as VST. Recent papers propose new AR display solutions that allow for wider FoVs than current devices [197, 202]. The effect of such increases on media viewing, and on immersive ARTV in particular, is not yet clear; at present, the limited FoV of these devices restricts the types of experiences they can provide.

Recent industry prototypes also illustrate both the potential and the remaining challenges of future AR hardware. Meta's *Project Orion*, publicly demonstrated in 2024 as a research prototype, showcases a lightweight optical see-through AR system with an estimated field of view of approximately 70°—substantially wider than current waveguide-based consumer devices [116, 59]. Orion employs advanced silicon carbide waveguide optics and distributes com-

putation to an external wireless processing unit, while interaction is supported through hand tracking and a neural wristband that detects electromyography (EMG) signals from the user's forearm [116, 2]. Although Orion demonstrates that wide-field optical AR in a glasses-like form factor is technically achievable, it remains a research prototype rather than a consumer product. Its architecture—including specialised optics, external compute hardware, and experimental input systems—illustrates the complexity still involved in delivering wide-field AR in a socially acceptable wearable form factor [116, 59].

At the opposite end of the spectrum, commercially released smart glasses such as the 2025 Meta Ray-Ban Display glasses demonstrate a more incremental approach toward consumer AR. These glasses integrate a small monocular near-eye display with an approximate field of view of $\sim 20^{\circ}$, enabling glanceable information such as notifications, navigation cues, and AI-generated responses to appear within the user's field of view [115, 60]. However, this limited augmentation window functions more like a heads-up display than a spatial AR environment, supporting small information panels rather than large-scale environmental augmentation. Such devices highlight the current trade-off in consumer hardware between compact wearable form factors and the perceptual envelope required for richer AR experiences [115].

More broadly, research and survey literature indicates sustained effort toward (i) widening optical see-through FoV (often via waveguide/holographic or hybrid combiner architectures) and (ii) improving opacity control and occlusion capabilities [197, 201, 203]. Yet even optimistic trajectories suggest these capabilities may arrive with non-trivial constraints (e.g., eyebox limitations, optical efficiency, non-uniform resolution, power/thermal costs) that still matter for how ornamentations can be placed and perceived in ARTV.

It is possible that by 2030–2035, consumer ARTV platforms may plausibly offer wider augmentation envelopes and improved contrast/opacity control than today, but the co-occurrence of *wide FoV + outdoor-readable contrast + thin optical see-through form factors + robust occlusion* remains a repeatedly identified systems challenge [197]. This is precisely why it is useful to study ornamentation parameters now under controlled conditions (including VR-based approximations): as hardware constraints relax, the key design question for passenger ARTV shifts from whether a device can display augmentations at all to whether it can display *sufficient* augmentation to meaningfully influence the viewing experience. In other words, the limiting factor becomes less 'can the device show the augmentation?' and more 'how much augmentation should be shown, where should it appear within the field of view and how much of the field of view should it occupy, and what types of augmentation are effective in enhancing immersion while preserving environmental awareness?' These questions directly motivate the systematic investigation of ornamentation density, placement, and type undertaken in this thesis.

2.6.3 Implications for Passenger ARTV

The hardware picture set out above has a direct methodological consequence. Where an envisioned ARTV experience depends on display properties that no current optical see-through device provides—particularly wide field of view together with sufficient opacity for augmentations to read as solid—those properties cannot be studied on the hardware that would eventually deliver them. Investigating how ornamentation density, placement, and type shape the balance between immersion and awareness therefore requires either waiting for hardware that does not yet exist, or approximating it. Section 2.7 examines the second option, its precedents, and the questions it raises about the validity of the resulting findings.

2.7 Evaluating AR using VR

2.7.1 Two Uses of VR Simulation in AR Research

A substantial body of AR research is conducted not on AR hardware but within VR, with augmentations rendered as though overlaid on the real world. The practice has a long precedent, ranging from early work investigating AR object manipulation [143] to more recent studies simulating mixed-reality office environments [101], and it is common in passenger contexts specifically [17, 110]. It is worth distinguishing two quite different reasons for adopting it, because they raise different questions about validity.

The first is to approximate a display that does not yet exist. As section 2.6 sets out, an envisioned AR experience may depend on a combination of display properties—particularly wide field of view together with strong opacity or occlusion—that no current optical see-through consumer device provides simultaneously. Where this is the case, VR can be used to approximate the perceptual capabilities of a prospective AR display rather than to represent the surrounding environment. Because both the environment and the augmentations are rendered within the same pipeline, parameters such as effective field of view, augmentation density, opacity, and placement strategy can be manipulated experimentally in ways that shipping hardware does not permit.

The second is to approximate an environment that cannot readily be accessed or controlled. Evaluating AR experiences in authentic settings is logistically and technically demanding, particularly where those settings are dynamic, public, or difficult to instrument. Section 2.7.3 discusses the resulting practical and ethical constraints in detail. VR allows a representation of such an environment to be constructed and held constant across participants and conditions, so that variability arising from crowding, vehicle motion, or the behaviour of other people does not confound the manipulation of interest.

Pass-through AR occupies an intermediate position between the two. It renders the real environment through cameras and displays, affording a larger effective augmentation window

than optical see-through hardware and supporting visually solid ornamentations with correct compositing and occlusion in the video pipeline. The environment is genuine rather than reconstructed, but it reaches the user as a camera image, so perceptual quality depends on sensor dynamic range, latency, noise, and display characteristics [5]. It is also worth noting that optical see-through AR may face constraints that are not purely a matter of engineering maturity: achieving high-contrast, daylight-readable augmentations through thin waveguide or diffractive optics involves physical trade-offs between brightness, field of view, and form factor that may limit how closely optical see-through can approach the compositing quality of video pass-through or conventional displays [196]. Each of the three approaches therefore involves a different compromise, and the studies reported in this thesis span the range.

These two uses of simulation are conceptually separable but are frequently combined in practice, and they are combined in this thesis. Chapter 4 and chapter 5 represent a passenger carriage in VR while also manipulating field of view and augmentation opacity beyond what current optical see-through hardware supports; the same simulation therefore stands in for both an environment and a display. Chapter 6 separates them, deploying on pass-through hardware in a physical carriage so that the environment is real and only the display remains a compromise. The validity of each use, however, is threatened by different things.

2.7.2 Validity and Limitations of VR-Simulated AR

Where simulation stands in for a display, the principal threat to validity is perceptual: a simulated wide-field, high-opacity AR display is not optically equivalent to a real one, and differences in rendering, luminance, and the absence of true see-through may affect how augmentations are perceived. Where simulation stands in for an environment, the threat is instead social and attentional: a reconstructed carriage may reproduce the visual appearance of a transit setting without reproducing its demands. The evidence below bears on both, and the distinction matters when deciding what a given simulation-derived finding can support.

Lacoche et al. [85] examined the validity of the approach in the context of an AR smart home application by comparing a real AR system with a VR-based simulation of the same experience. Their results showed only minor differences between the two conditions in terms of task performance and perceived workload, suggesting that VR-based simulations may provide a reasonable approximation for certain types of AR interaction. However, the study also highlighted limitations; for example, passive haptic feedback from interacting with real objects—such as touching a physical sofa—could not be reproduced within the simulated environment.

Importantly, the environments examined in such validation studies are typically stable and predictable. Domestic and office settings lack the dynamic motion, unpredictable bystander behaviour, and shifting environmental conditions characteristic of passenger transit environments. As a result, it remains unclear whether VR-based simulations can accurately reproduce the perceptual and attentional demands associated with AR experiences deployed in real-world

transportation contexts.

While VR-based simulation can achieve high face validity—the extent to which a study appears realistic and representative to participants—this does not guarantee ecological validity, or that findings will generalise to real-world deployments. Dole and Ju [41] argue that with a high level of face validity a simulation can begin to approximate ecological validity. Empirical comparisons between real and simulated environments, however, suggest that behavioural outcomes may still differ in important ways. Kalantari et al. [75] compared way-finding tasks conducted in real-world environments with those performed in VR, and observed significant differences across all measures, including distance travelled, number of navigation errors, task completion time, spatial memory performance, and perceived cognitive workload. Although participants responded to the same environmental cues in both conditions, the extent of their responses differed significantly. Wayfinding is a particularly relevant example in transit contexts, where passengers must frequently navigate complex stations, interpret signage, and monitor their journey progress while engaging with other activities. Similarly, Pizzolante et al. [138] examined the product experience of a bicycle in both reality and VR, and found significantly higher levels of engagement and positive affect in VR than in the real world. Simulated environments may therefore diverge from their real counterparts not only in task performance but in affective response, and in either direction.

These concerns are sharpened when the experience being simulated depends on social presence, situated attention management, or environmental unpredictability rather than simply visual layout or task performance. The concern is not that VR simulation is inappropriate for AR evaluation in general, but that lightweight passenger simulations—those which prioritise place illusion and experimental control while populating the environment with scripted, predictable avatars—may reproduce the visual appearance of a transit environment without reproducing its social plausibility. Real bystanders behave unpredictably, introduce self-consciousness, and create social accountability that scripted passengers cannot replicate [170].

Technical differences between simulated and real AR systems may further influence validity. Differences in depth cues, tracking latency, peripheral vision, and occlusion handling between VR-simulated AR and actual pass-through AR may affect both perceptual and cognitive outcomes. Real-world AR deployments must additionally contend with imperfect models of the environment (incomplete or misaligned scene meshes), dynamic lighting conditions, and unpredictable bystander behaviour—factors typically absent from controlled VR simulations of transport environments. Pass-through AR introduces perceptual artefacts of its own, including camera latency, dynamic lighting variability, occlusion instability, and tracking inconsistencies, further widening the gap between simulated and situated experience [28].

Combined, these findings indicate that the adequacy of a simulation cannot be settled in general terms, but depends on the type of claim it is being used to support. Simulation may be appropriate for early-stage design exploration and for the controlled manipulation of visual

features such as augmentation density or placement. It is less clearly sufficient where the claim concerns ecological validity, social presence, situated attention management, or the perceptual fit between virtual content and a physical public environment. This thesis does not propose a general criterion for simulation adequacy, and treats adequacy as claim-relative rather than absolute; chapter 6 examines empirically which classes of finding transferred from simulation to a situated deployment and which did not.

2.7.3 Practical and Ethical Constraints of Real-World AR Evaluation

Beyond methodological concerns, evaluating AR experiences directly in public transit environments introduces practical and ethical challenges. Public transport spaces are shared environments involving other passengers who have not consented to participate in research studies. Conducting experiments in such settings may raise privacy concerns, particularly if sensors or cameras are used to capture environmental data or participant behaviour.

There are also safety considerations. Participants using HMDs in busy transport environments may experience reduced environmental awareness, increasing the risk of collisions, trips, or other accidents. Researchers must therefore carefully consider participant safety as well as the potential impact on other passengers sharing the environment.

Access and logistical constraints further complicate *in situ* studies. Conducting controlled experiments on operational transport systems often requires coordination with transport authorities, access permissions, and careful scheduling to avoid disrupting normal passenger services. Even when permission is granted, variability in crowd density, vehicle motion, and environmental conditions can make it difficult to maintain experimental consistency. These constraints are precisely why VR-based simulation is attractive as an intermediate step: it allows key variables to be controlled while avoiding many of the ethical and logistical difficulties of deploying experimental systems in live public settings.

If simulated AR over- or under-estimates immersion or awareness costs, however, conclusions about ARTV design may not generalise to real-world passenger use. This risk is particularly salient for experiences in which social dynamics and environmental unpredictability are central rather than incidental, as lightweight simulations may capture broad directions of effect while misrepresenting their magnitude or causes. Accordingly, this thesis treats VR-simulated AR not as a replacement for real-world evaluation, but as a controlled stepping stone toward understanding how immersive ARTV designs translate into authentic passenger environments—and the validity of that stepping stone is itself treated as an empirical question. This motivates a direct comparison between simulated AR and consumer AR deployment in a physically and socially situated, public-facing transit carriage:

RQ4: *How does the passenger immersion–awareness trade-off apply when moving from simulated AR-in-VR to AR experiences situated in the real world?*

2.8 Summary and Research Gap

Passengers increasingly use travel time for media consumption, creating an opportunity for XR technologies to transform the commuting experiences. However, immersive media technologies introduce a fundamental design challenge in passenger contexts: the tension between media immersion and environmental awareness. Fully occlusive VR maximises immersive properties but removes the visual access to the surroundings that passengers report needing for safety and social acceptability. AR offers a promising alternative by preserving awareness of the surrounding environment while enabling digital augmentations that extend media content into the passenger's surroundings.

This chapter reviewed literature across commuting behaviour, media immersion, XR passenger constraints, ARTV design frameworks, and the hardware and methodological realities of current AR systems. Together, these strands highlight a gap in understanding how AR-based environmental augmentations—ARTV ornamentations—can enhance media immersion while maintaining environmental awareness in passenger contexts. In particular, the perceptual parameters governing these experiences, including augmentation density, field of view, and ornamentation type, remain underexplored.

Additionally, while many AR studies evaluate designs using VR-based simulations to control experimental conditions, it remains unclear how accurately such simulations capture the immersion–awareness dynamics experienced in real-world AR deployments. Addressing these conceptual and methodological gaps, this thesis investigates how immersive media experiences can be designed and evaluated in passenger contexts, and chapter 7 develops the resulting account of the immersion–awareness relationship. This immersion-awareness tension motivates the empirical investigations presented in the following chapters.

Chapter 3

Current and Desired Commuter Activities: A Survey and Participatory Design Workshops on Design Opportunities for Passenger XR

3.1 Introduction

This chapter provides the empirical foundation for the thesis through two complementary studies: a survey of commuters, mainly in the UK and the USA, establishing how passengers currently use their travel time and where they would welcome improvement, and a series of participatory design workshops in which commuters generated design responses to the challenges the survey identified. This chapter therefore addresses RQ1 (*How could XR improve the current experience of immersive media on public transport?*) by empirically examining the gap between current commuter activities and activities that passengers would ideally like to undertake if technological constraints were removed, and by exploring with commuters how immersive technologies might close that gap. The findings identify behavioural patterns, experiential preferences, and design strategies that inform the direction taken in later chapters.

According to the 2023 National Travel Survey, the average one-way commute in England is now 27 minutes, amounting to almost an hour of travel for each day that a person commutes [48]. This is time that is often perceived as wasted [35]. In recent years there has been growing interest in the use of XR to enhance the passenger experience, such as using VR to create an immersive movie watching experience whilst allowing for intrusions from other passengers [188], or in a car creating an immersive shooting game in a virtual helicopter with its movements matched to the real world movements of the car [61]. However, the proposed activities often focus on either highly immersive experiences, or on simply transposing existing activities (e.g. watching movies) to a new context using spatial computing. This can be seen in the work of Medeiros et al.

[110], who bring a multi-monitor office setup to a train using XR. Little work has directly asked passengers what they might want to do during their commute if XR were available, and this absence means that research and industry may be overlooking ideas that passengers themselves could contribute.

Understanding commuter motivations, preferences, and constraints is essential for designing XR experiences that are both desirable and practical in transit environments. Without this grounding, passenger XR applications risk replicating domestic XR experiences without accounting for the unique constraints of public transport. These constraints—such as divided attention, limited personal space, and social exposure—are reflected in the activities commuters avoid, seek to improve, or describe as problematic.

The survey (section 3.2) examines the gap between the activities commuters currently undertake and those they would prefer to undertake if technological constraints were removed, extending the NRPS activity data analysed by Lyons et al. [103] to a wider range of transport modes. The workshops (section 3.3) then move from characterising that gap to exploring how immersive technologies might close it, using generative participatory design methods to elicit and develop design concepts with commuters rather than evaluating systems presented to them.

The two studies serve different purposes and neither alone determines the direction taken. The survey establishes which activities passengers value and where current provision falls short; the workshops surface the design strategies commuters themselves reach for, and reveal a consistent preference for selectively augmenting the travel environment while preserving awareness of the carriage and other passengers. Section 3.3.5 sets out how these findings, together with practical considerations of tractability and experimental control, led to media viewing being selected as the design context for the remainder of the thesis, in preference to the reading and scenery-related activities that the survey ranked more highly.

The concept of ARTV Ornamentation examined in subsequent chapters is developed from prior work in augmented television [183, 150]. What this chapter contributes is evidence that the design space it occupies (extending media content into the surrounding environment while retaining a view of it) corresponds to a genuine passenger need, together with a characterisation of the immersion and awareness constraints any such system must satisfy. This framing motivates the design investigations presented in chapters 4 and 5, which examine how AR-based environmental augmentation can enhance media immersion while preserving environmental awareness.

3.2 Desired Commuter Activity Survey

To address the absence of user-centred evidence on passenger XR needs, a survey was designed to move beyond descriptive activity patterns and identify the activities commuters would ideally undertake during travel, alongside the barriers that currently prevent them from doing so. The survey administered in this chapter was based upon the NRPS as analysed by Lyons et al. [103]

and Singleton's 2016 survey of travel-based multitasking among 656 commuters in the Portland, Oregon, region [161]. Some activities were modified and some taken verbatim (see table 3.1). Video conferencing was added as an activity due to its rise in popularity since 2014 and 2016 when the NRPS and Singleton's surveys were conducted. It rose to prominence during the COVID pandemic and remains a major productivity activity today. Online video conference and meeting provider Zoom reported 6 billion meeting minutes in 2016, which rose to 3.3 trillion by the end of 2021 [195], and online meetings remain common: in North America, 42% of meetings took place solely online [142]. Because of this, *Video Conferencing* was considered as a unique commuter activity and added to the survey.

Relaxation / Meditation was added as an activity because of the popularity of meditation apps such as *Calm* and *Headspace*. In 2019, 52 million users downloaded one of the top 10 meditation apps [135] and in 2020 the number of apps downloaded in the *Emotional Wellbeing* category rose 42% to 96 million [25].

Respondents in Singleton's survey [161] were presented with a list of activities and asked which they had done previously on a commute. For the survey here, that list was taken and participants instead had to categorise each activity as *Don't do and don't want to do*, *Currently Do*, *Currently do but want to do better*, or *Don't currently do but want to*. This was done to identify commuter desires and pain points. For each activity that was marked as either *Do but want to do better* or *Don't do but want to*, the participant was asked an open response question to understand their reasoning. The aim was to determine if and how XR could improve the commute experience in this regard. Respondents also selected the primary activity they undertook during their commute. This allows for comparison with the NRPS [103] and Singleton's survey [161]. The full survey can be found in section B.1. The raw data can be found at the Zenodo repository [33].

3.2.1 Recruitment and Scope

The survey was approved by the College of Science and Engineering Ethics Committee at the University of Glasgow (application 300220083). Respondents were recruited via university mailing lists, online forums such as reddit.com/takemysurvey, and posters put up at public transit stations such as university bus stops, subway stops, and local train stations. Posters were placed at public transport stops rather than car parks because the survey targeted passengers. Participants who provided an email address were entered into a raffle to win one of five £5 Amazon vouchers. The survey was active from February until May 2023.

A distinction was drawn between commuters who operate their own vehicle and those who travel as passengers. Respondents whose primary mode required active control of the vehicle, such as driving or cycling, were routed to a separate block of questions for vehicle operators, which is not analysed in this chapter. This reflects fundamental differences in attentional demand and activity structure between the two, as Singleton [161] documents in distinguishing

Table 3.1: The 22 commuting activities respondents categorised, the source of each item (NRPS [103], Singleton [161], modified from these, or added for this survey), and the short label used in this chapter.

Survey item	Short label	Source	Justification
Using social media websites or apps (Facebook, Twitter, LinkedIn, Tumblr, TikTok, Instagram)	Using Social Media	Singleton (TikTok added)	
Texting, emailing or other messaging	Messaging	Singleton	
Reading electronically	Reading Electronically	Singleton	
Playing game on electronic device	Playing Game on Electronic Device	Modified (Singleton; NRPS)	Distinguish electronic from physical games
Physical games (pen and paper, crosswords, board games, playing cards)	Physical Games	Modified (Singleton; NRPS)	Distinguish electronic from physical games
Listening to music or audio (podcast)	Listening to Music or Audio	NRPS	
Doing nothing	Doing Nothing	Singleton	
Viewing scenery	Viewing Scenery	Modified (Singleton; NRPS)	Split from a combined scenery and people-watching item
Watching people	Watching People	Modified (Singleton; NRPS)	Split from a combined scenery and people-watching item
Thinking or daydreaming	Thinking or Daydreaming	Singleton	
Reading print (newspaper, physical book etc.)	Reading Print	Singleton	
Eating or drinking	Eating or Drinking	NRPS	
Talking on the phone	Talking on Phone	Singleton	
Personal grooming (shaving, makeup etc.)	Personal Grooming	Singleton	
Sleeping / snoozing	Sleeping	NRPS	
Writing or editing paper documents	Writing or Editing Paper Documents	Singleton	
Writing or editing electronic documents	Writing or Editing Electronic Documents	Singleton	
Watching movie, TV or other video	Watching Video	Singleton	
Trying not to be anxious about the journey (e.g. delays or where to get off)	Trying Not to Be Anxious about the Journey	NRPS	
Talk to other passengers	Talk to Other Passengers	NRPS	
Video conferencing (Zoom, Skype, FaceTime)	Video Conferencing	Added	Video conferencing has distinct challenges
Relaxation / meditation	Relaxation / Meditation	Added	Associated apps becoming more popular

operating modes, where the mental and physical resources required by the travelling task leave room for little more than passive activities, from riding modes, where they do not. Operating modes require continuous engagement with navigation, hazard perception, and motor control, placing strict constraints on the types of activities that can be safely undertaken during travel. As a result, these contexts are not compatible with immersive media consumption in the form explored in this thesis. Research on non-driving-related tasks (NDRTs) in automated vehicles treats the activities drivers want once they no longer steer or monitor the vehicle, and which of those remain feasible, as a distinct design problem [136], with requirements that differ substantially from those addressed here. These contexts therefore fall outside the scope of this work. Respondents who mainly walked were not routed to the operator questions and were retained; most of them also used public transport as part of their commute. Passengers, by contrast with vehicle operators, do not have primary responsibility for navigation or vehicle control, allowing greater flexibility in how attention can be allocated during the journey. This makes passenger contexts particularly suitable for investigating immersive media consumption and the balance between engagement and environmental awareness, which chapter 7 develops as the PIAT.

3.2.2 Respondents

The survey received 108 responses. Respondents who did not commit to giving thoughtful answers, or who completed the survey in under 90 seconds, were removed, leaving 89, of whom 84 commuted at least one day per week. Respondents whose primary mode was driving (15) or cycling (2) were routed to the operator questions described above, and a further 9 did not reach the activity questions. The activity analyses therefore use the 58 passengers who answered them (57 for primary activity). Of these 58, 23 were aged 18–25, 21 were 26–34, 10 were 35–44, and 4 were 55–64; none were aged 45–54 or over 65. Thirty were men, 25 women, and 2 non-binary, and 1 did not state their gender. Their approximate locations, inferred by the survey platform from IP addresses, were the UK (32), the USA (12), East and South-East Asia (6), elsewhere in Europe (1), and unknown (7). Their primary modes were bus (25), walking (10), subway or metro (10), rail (9), and car passenger (3), with one not stated.

Transport Types

Figure 3.1 shows the transport modes used by all 84 commuters, including the drivers and cyclists. The most common transport type, both for part of the commute and as the primary mode, was the *Bus* (primary $n = 29$, uses $n = 52$). For primary mode, this was followed by *Car (driving)* ($n = 15$) and *Walking* ($n = 12$). For use as part of the commute, the second most common mode was *Walking* ($n = 44$) and *Subway / Metro* was third ($n = 34$). The high proportion of bus commuters is likely to reflect the placement of a survey flyer outside a highly trafficked university bus stop. The sample was therefore dominated by bus commuters, whereas the later studies

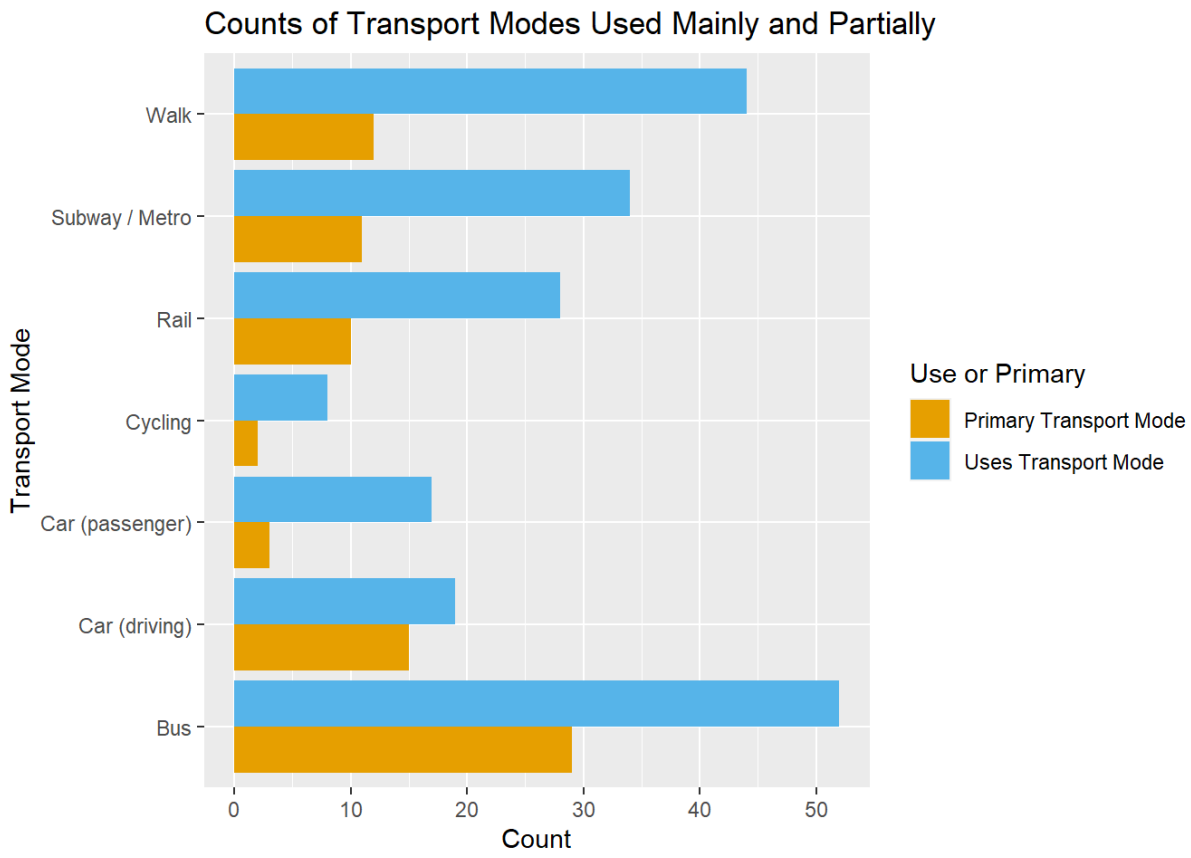


Figure 3.1: Transport modes used by the 84 commuters, including drivers and cyclists: modes used for any part of the commute (respondents could select several) and the primary mode (82 respondents).

use rail (chapters 4 and 5) and subway (chapter 6) settings. Seated passengers on these modes share the riding-mode conditions described above, but the survey cannot show whether activity preferences differ between them. Because of the small sample size, comparisons by transport mode are not made.

Commute Duration

Respondents were asked about their average commute duration in minutes. Figure 3.2 shows the distribution of commute duration for each primary transport mode among all commuters. The median commute duration in minutes by primary mode was 22.5 for *Walk* (12 respondents), 46.5 for *Subway / Metro* (11), 55 for *Rail* (10), 55 for *Cycling* (2), 50 for *Car Passenger* (3), 30 for *Car (driving)* (15), and 37.5 for *Bus* (29). Commute durations were given as free text; a few answers were ranges or implausible values (for example 480 or 2880 minutes), which were converted to a single plausible duration in minutes before analysis. Singleton [161] reported a mean travel time of 35.66 minutes ($SD = 21.27$) across all commuters. For the 58 passengers analysed here, the mean commute duration was 44.57 minutes ($SD = 31.95$, $n = 56$ with valid duration data), with a median of 38 minutes.

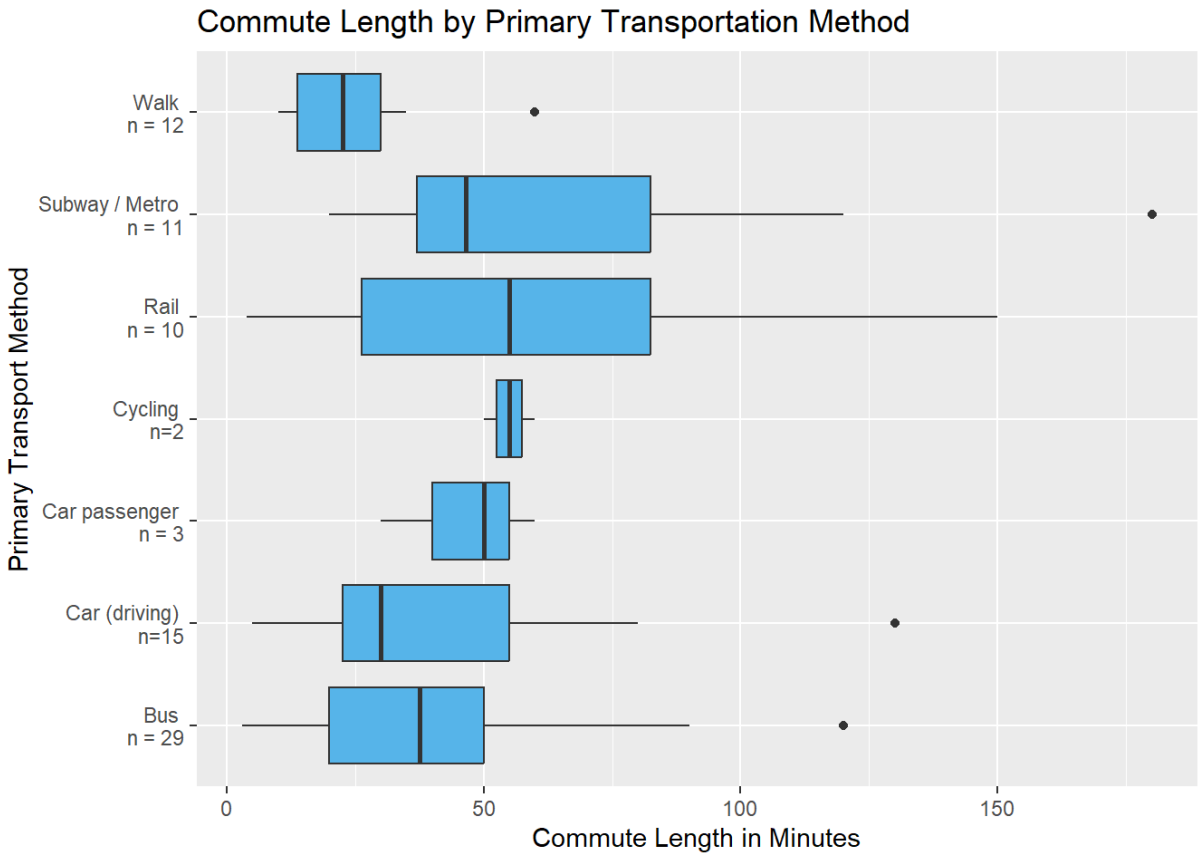


Figure 3.2: Commute duration by primary transport mode for all commuters with a valid duration ($n = 79$), including drivers and cyclists.

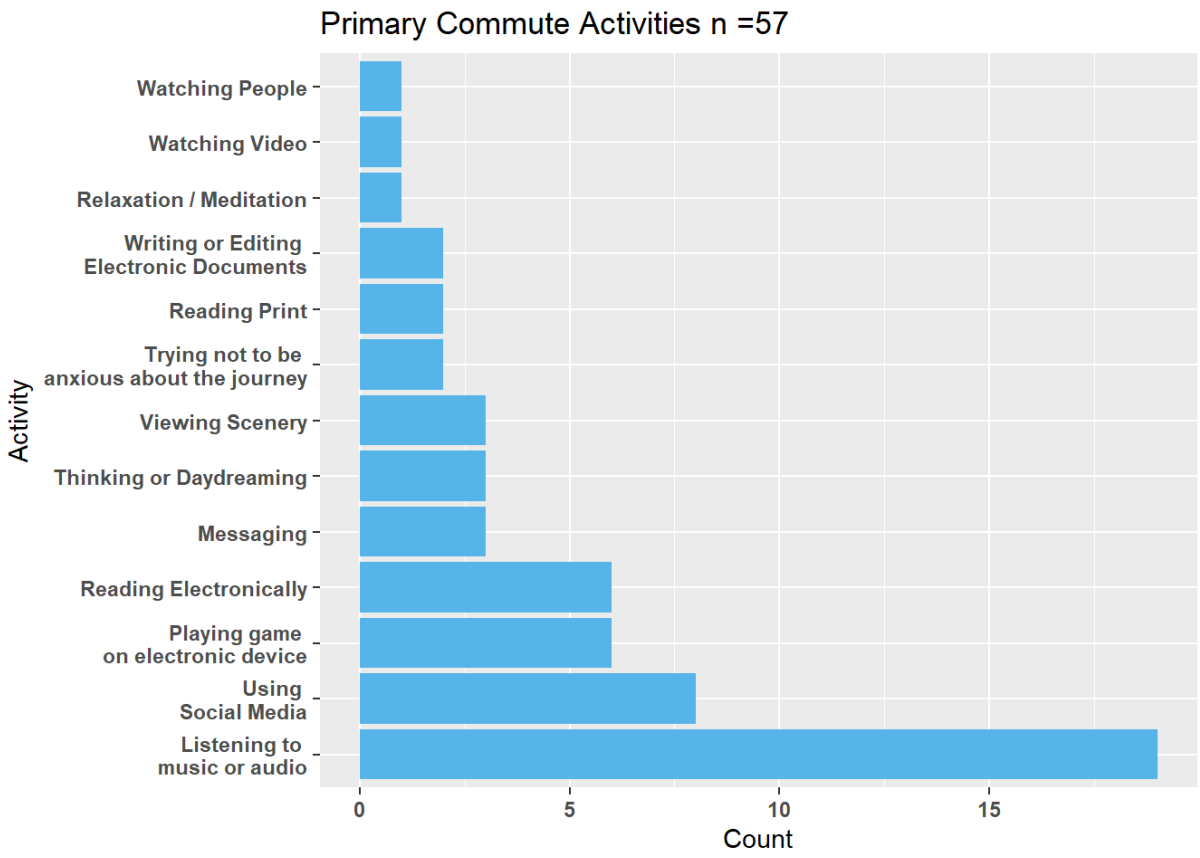


Figure 3.3: Primary commute activity of the passengers who answered this question ($n = 57$).

3.2.3 Existing Journey Experience

Journey Objectives

Respondents selected all applicable journey objectives from three options. Of the 58 respondents who answered this question, 43 selected *Pass the Time* (74.1%), 25 selected *Enjoy the Journey* (43.1%), and 17 selected *Be Productive* (29.3%).

Current Primary Activities During Commute

Respondents were asked to name their primary activity during their commute. As seen in fig. 3.3, the most common activity was *Listening to Music or Audio* ($n = 19$) followed by *Using Social Media* ($n = 8$), *Playing Game on Electronic Device* ($n = 6$), and *Reading Electronically* ($n = 6$).

In the 2014 wave of the NRPS [103], the activity on which commuters and all passengers spent most time was *Reading for Leisure* (37.2% and 32.6% respectively). In the present survey, primary commute activity was reported by 57 of the 58 passengers analysed. When *Reading Electronically* and *Reading Print* are combined into a broader *Reading* category, reading accounts for 8 responses (14.0%), making it joint second with *Using Social Media*. This lower proportion may reflect both the larger number of activity categories available in the present

survey and the much smaller sample size compared with the NRPS (57 valid primary-activity responses, compared with 19,091 NRPS respondents to the equivalent 2014 question).

Window Gazing / People Watching was the second most common main activity for all passengers in 2014 (18.6%) and the third most common for commuters (11.3%). With the present survey's categories of *Viewing Scenery* and *Watching People* combined, 4 respondents (7.0%) reported these as their primary activity.

Working / Studying was the third most common main activity for all passengers in 2014 (12.4%) and the second most common for commuters (12.7%). This was much less common in the present survey: only 2 of the 57 respondents who reported a primary commute activity selected *Writing or Editing Electronic Documents* (3.5%), and none selected *Writing or Editing Paper Documents* as their primary activity.

None of the three most common primary activities in the present survey was prominent in the NRPS data. The most common primary activity in the present survey, *Listening to Music or Audio*, accounted for 33.3% of valid primary-activity responses, compared with 9.4% of commuters in the 2014 NRPS (and a peak of 11.1% in 2010). Similarly, *Using Social Media*, which was the main activity of 1.8% of commuters in 2014, accounted for 14.0% in the present survey. Both differences are likely to reflect the place smartphones now occupy in how people connect with each other and consume content. *Playing Game on Electronic Device* was the primary activity for 10.5% of respondents in the present survey, compared with 2.0% of commuters in the 2014 NRPS for playing games (electronic or otherwise). This may reflect the younger profile of the present respondents or the rise and popularity of smartphone gaming [9].

3.2.4 Desired Journey Experience

Participants were asked to categorise each activity in the activities list as either *Currently Do*, *Do & Want To Do Better*, *Don't Do But Want To*, or *Don't Do & Don't Want To*. The activity-categorisation matrix was not forced-choice, so skipped rows were treated as non-selection and recoded as *Don't Do & Don't Want To*. Table 3.2 summarises these recoded responses using a fixed denominator of respondents who answered at least one item in the matrix ($n = 58$). To identify activities associated with unmet demand or scope for improvement, the *Improve* column combines responses for *Do & Want To Do Better* and *Don't Do But Want To*. These two categories represent distinct states: the former indicates a quality gap, where an activity is already undertaken but the experience falls short, while the latter indicates an access barrier, where the activity is not undertaken at all. They do not rank activities identically. *Using Social Media*, for instance, registers a substantial quality gap ($n = 15$) alongside no access barrier whatsoever ($n = 0$), while *Reading Print* ranks fifth on quality gap but first on access barrier. Both components are reported separately in table 3.2, and the distinction is returned to where it bears on interpretation. The highest combined improvement-interest values were for *Reading Electronically* ($n = 25$), *Reading Print* ($n = 22$), and *Viewing Scenery* ($n = 21$).

The activities most commonly categorised as *Currently Do* were *Messaging* ($n = 35$), *Using Social Media* ($n = 34$), and *Listening to Music or Audio* ($n = 33$). The activities most commonly categorised as *Don't Do & Don't Want To* were *Video Conferencing* ($n = 47$), *Writing or Editing Paper Documents* ($n = 43$), and both *Physical Games* and *Personal Grooming* ($n = 42$).

Table 3.2: Activity categorisation counts for respondents who answered at least one item in the activity-categorisation matrix ($n = 58$). Skipped rows were recoded as *Don't Do & Don't Want To*. *Improve* is calculated as *Do & Want To Do Better* plus *Don't Do But Want To*.

Activity	Do	Better	Want	Don't	Improve	Total
Reading Electronically	11	17	8	22	25	58
Reading Print	11	13	9	25	22	58
Viewing Scenery	27	17	4	10	21	58
Listening to Music or Audio	33	15	3	7	18	58
Writing or Editing Electronic Documents	7	12	4	35	16	58
Using Social Media	34	15	0	9	15	58
Thinking or Daydreaming	29	12	3	14	15	58
Relaxation / Meditation	10	12	2	34	14	58
Playing Game on Electronic Device	14	10	4	30	14	58
Watching People	27	11	2	18	13	58
Doing Nothing	15	9	4	30	13	58
Trying Not to Be Anxious about the Journey	14	9	4	31	13	58
Watching Video	11	9	4	34	13	58
Messaging	35	9	2	12	11	58
Eating or Drinking	11	9	2	36	11	58
Sleeping	9	9	2	38	11	58
Talking on Phone	11	8	1	38	9	58
Physical Games	7	7	2	42	9	58
Talk to Other Passengers	12	6	3	37	9	58
Personal Grooming	8	6	2	42	8	58
Writing or Editing Paper Documents	9	3	3	43	6	58
Video Conferencing	7	4	0	47	4	58

Note. Do = *Currently Do*; Better = *Do & Want To Do Better*; Want = *Don't Do But Want To*; Don't = *Don't Do & Don't Want To*; Improve = Better + Want.

Because the recoding of skipped rows shapes the resulting ranking, a sensitivity analysis was conducted to establish whether the conclusions depend on it. Recomputing improvement interest as a proportion of respondents who answered each item, rather than recoding blanks as non-selection, produced a strongly correlated ordering (Spearman's $\rho = 0.77$, $p < .001$), with

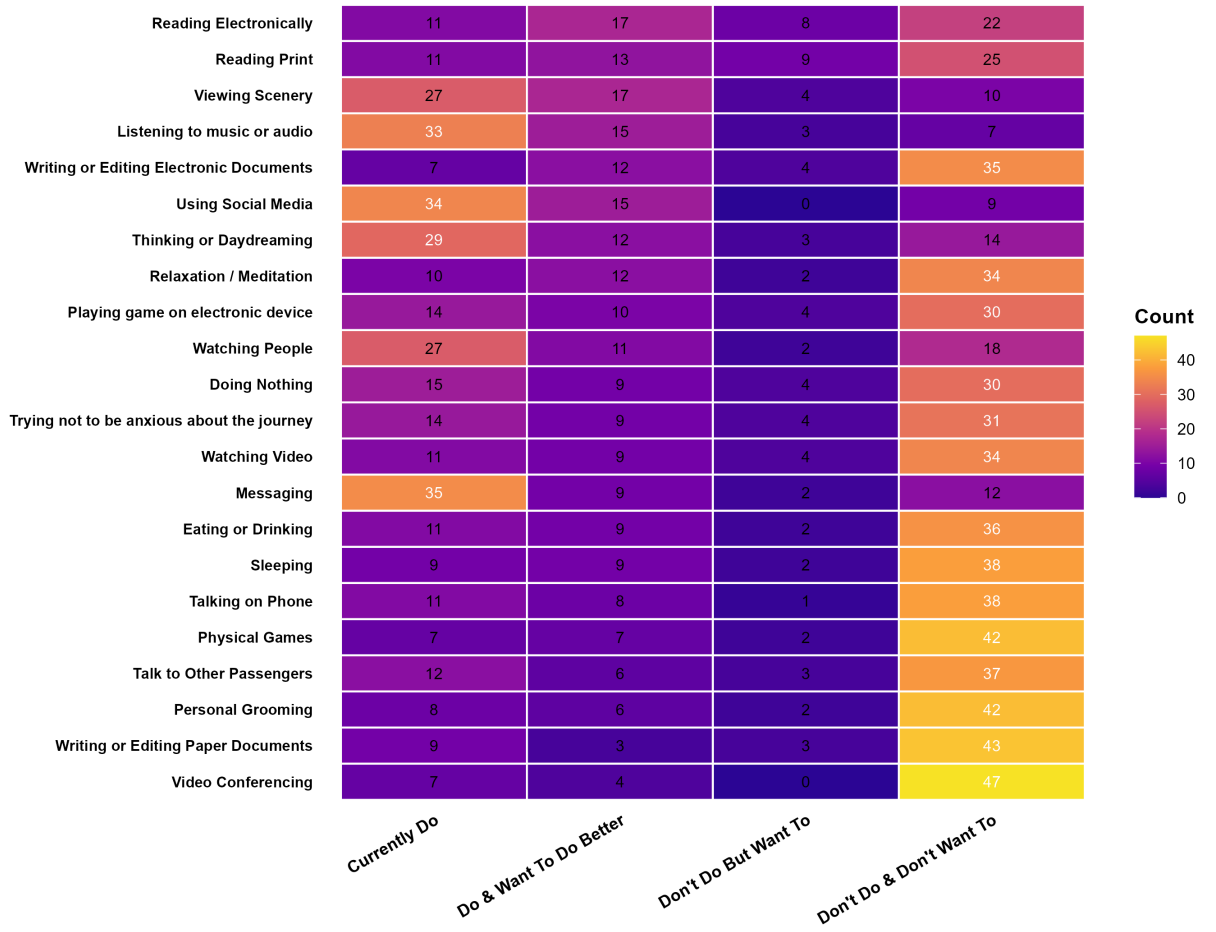


Figure 3.4: Recoded commute activity categorisation counts ($n = 58$). Skipped rows were recoded as *Don't Do & Don't Want To*. Activities are ordered by improvement interest, calculated as the sum of *Do & Want To Do Better* and *Don't Do But Want To*.

the three most requested activities unchanged in composition. Individual activities did shift: *Using Social Media* fell from joint sixth to eighteenth under the proportional treatment because of its comparatively high response rate ($n = 56$), while *Watching Video* ($n = 33$) rose from joint tenth to sixth. The overall pattern on which this chapter's conclusions rest is therefore not an artefact of the recoding, although the precise ordering of mid-ranked activities is sensitive to how skipped rows are treated.

Figure 3.4 visualises the same recoded counts as a heatmap, providing an overview of how response patterns differ across activities and categories.

In the 2014 wave of the NRPS, Lyons et al. [103] found that the activity on which the most rail passengers spent any time was *Window Gazing / People Watching* (43.7%), and that it was the second most common such activity for commuters (37.2%). This was split into two distinct activities in the present survey because it was hypothesised that XR would have different impacts on viewing scenery and watching people. Using the combined count of *Currently Do* and *Do & Want To Do Better* in table 3.2, *Viewing Scenery* was undertaken by 44 respondents (75.9%),

while *Watching People* was undertaken by 38 respondents (65.5%). These proportions are lower bounds, since skipped rows were counted as not done. They are higher than the corresponding NRPS figures, although the comparison should be interpreted cautiously because the present survey used different activity categories, a smaller sample, and a different response format. When passenger objectives are taken into account, with 43 respondents selecting *Pass the Time* (74.1% of those who answered that question), the prevalence of scenery-viewing and people-watching is plausible. These are activities that do not require active engagement or the use of any object.

Reading for Leisure was the most common activity for commuters in the 2014 NRPS, at 45.4% (down from 61.4% in 2004), and the second most common for all passengers, at 40.5%. Using the same combined measure in the present survey, *Reading Electronically* was undertaken by 28 respondents (48.3%), while *Reading Print* was undertaken by 24 respondents (41.4%). The present survey separated electronic and print reading to examine whether the two activities created different opportunities or constraints for XR. These figures suggest that reading remains a substantial commute activity, although comparisons with the NRPS should again be interpreted cautiously due to differences in categories and survey format.

The third most common activity for both all passengers and commuters in the 2014 NRPS was *Text Messages / Phone Calls – Personal*, at 27.6% and 30.5% respectively. For the same reasons as above, the present survey separated this into *Messaging* and *Talking on Phone*. In the present survey, *Messaging* was undertaken by 44 respondents (75.9%), while *Talking on Phone* was undertaken by 19 respondents (32.8%). *Messaging* was therefore one of the most common activities in the present survey, whereas talking on the phone was much less common.

The most common activity by this combined measure was *Using Social Media*. This was reported by 49 respondents (84.5%), compared with 15.5% of commuters who accessed social networking sites in the 2014 NRPS [103]. This difference is likely to reflect the growth of smartphone-mediated social media use, although the comparison should be treated cautiously because of the smaller sample and differing response format in the present survey.

The second most common activity by this combined measure was *Listening to Music or Audio*, reported by 48 respondents (82.8%). This compares with a peak of 26.7% of commuters in the 2010 NRPS, decreasing to 20.2% in 2014. This may reflect the increased availability and normalisation of mobile audio streaming services.

The transit-mode results in Table 3 of Singleton [161], which report the share of 175 bus and rail commuters who did each activity on their most recent commute, can be compared with the combined *Currently Do* and *Do & Want To Do Better* counts in table 3.2. Differences of note include 16.0% of Singleton's transit commuters for playing games, compared with 24 respondents in the present survey reporting *Playing Game on Electronic Device* as either currently done or done but needing improvement (41.4%); 32.6% for reading print, compared with 24 respondents in the present survey (41.4%); and 29.1% for listening to music, radio, or other audio, compared

with 48 respondents in the present survey (82.8%).

Open-text Reasons for Desired Activity Improvement

For activities that respondents categorised as either *Do & Want To Do Better* or *Don't Do But Want To*, respondents were asked to provide open-text reasons for their selection. These responses were not analysed as a formal qualitative dataset, but were reviewed to identify illustrative constraints that help explain why particular activities were difficult or desirable during commuting.

Several responses related to reading and screen-based activities identified practical constraints associated with the commute environment. For *Reading Electronically*, respondents referred to the need for a better screen and discomfort when reading on a small mobile screen. For *Writing or Editing Electronic Documents*, one respondent noted that they were usually limited to using their phone, which made writing difficult, while another cited the instability of a 'shaky bus'. For *Reading Print*, responses included carrying constraints and one mention of motion sickness. These comments suggest that screen size, device ergonomics, vehicle movement, and carrying additional materials can act as barriers to reading or working while travelling.

Open-text responses concerning *Viewing Scenery* also indicated that the activity was shaped by the physical conditions of the journey. Respondents referred to issues such as limited access to a suitable view, poor visibility, weather, distraction from other activities, and the repetitive or limited aesthetic quality of familiar routes. These responses provide context for the quantitative finding that 17 respondents reported viewing scenery but wanting to do it better, while a further 4 did not currently do it but wanted to.

These open-text responses should be interpreted cautiously because they were sparse, and many respondents did not provide answers to this section of the survey. Nevertheless, they do provide examples of the kinds of practical and environmental constraints that may explain why respondents wanted to improve otherwise common commute activities. These constraints are used in the following discussion to identify potential XR design opportunities, rather than as evidence of generalisable themes.

3.2.5 Discussion

In the years since the NRPS waves analysed by Lyons et al. [103] and Singleton's survey [161], there has been a marked increase in the use of social media and listening to music as part of the commute. This is possibly due to the way smartphones have affected people's lives and the tasks they do every day. The large change in activities between 2004 and 2023 brought on by the smartphone might be echoed over the coming decades by the introduction of XR as the platform matures and gains popularity. In a future with everyday wearable AR, AR glasses may be as common as a phone. This may afford many new opportunities to improve the experience of the

activities passengers choose to conduct during their commute.

Using Diminished Reality to Address Distraction

Reading Electronically and *Reading Print* were prominent activities in the survey, and the open-text responses included examples of distraction and environmental constraints as reasons why reading could be difficult when commuting. One possible XR approach to this issue is diminished reality (DR), a form of mediated reality originating in Mann's early work on wearable computing [105], in which elements of the physical environment are visually removed, hidden, or attenuated from the user's perception in real time [124]. In contrast to AR, which adds virtual content to the environment, DR reduces or suppresses selected parts of the perceived real world. Empirical work on DR as a cognitive aid has found that computationally diminishing visual distractions can lower subjective workload and, under some conditions, improve task performance and environment awareness compared to an unmodified distraction condition [109].

DR has already been demonstrated in vehicular settings, most visibly in Land Rover's Transparent Bonnet concept, which used grille-mounted cameras and a head-up display to render the bonnet and engine bay effectively invisible to the driver [88], and in subsequent research examining driver acceptance of see-through cockpits [98, 99]. These applications diminish parts of the vehicle itself in order to improve the driver's view of the road. In a commuting context, the same principle could instead be applied on the passenger's behalf, reducing the salience of distracting environmental features while preserving task-relevant awareness. For example, an XR device could identify a visually disruptive passenger and replace them with an approximation of the surrounding seat or carriage interior. This would make the passenger less visually salient while leaving the surrounding context intact. This is consistent with observed passenger behaviour in AR transit studies, where passengers naturally used virtual displays to visually shield themselves from other passengers [110], which suggests a pre-existing desire for selective visual attenuation of co-present others. DR could offer a more principled mechanism for achieving this than display placement alone, and might help passengers focus more effectively on reading while maintaining awareness of the transit environment.

Using Altered and Extended Reality to Improve or Replace the Surroundings

Viewing scenery was also a common commuter activity: 27 respondents categorised it as *Currently Do*, 17 as *Do & Want To Do Better*, and 4 as *Don't Do But Want To*. In total, 21 respondents expressed either unmet interest in viewing scenery or a desire to improve the activity. Open-text responses suggested several barriers to enjoying scenery during the commute, including distraction from other activities, limited access to a suitable view or window seat, poor visibility, weather, and the repetitiveness or limited aesthetic appeal of familiar routes. This suggests an opportunity for XR to enhance or alter the perceived external environment. For

example, XR could alter the view outside the window so that, instead of repeatedly seeing the same urban, industrial, or otherwise unappealing route, passengers could experience the vista of a train ride through the Alps, a rainforest, a futuristic cityscape, or even a journey through space. This could change regularly to avoid the user tiring of the same environment. Similarly, where passengers lack access to a window seat, AR could render a virtual window over a wall or in the aisle, thereby making every seat a potential window seat.

XR could also introduce virtual elements into the user's perception to enhance their experience and immerse them in their activity. For example, when watching a nature documentary, plants or environmental features from the on-screen jungle could be rendered within the train carriage, bringing elements of the video content into the passenger's surrounding environment. For someone reading *Murder on the Orient Express*, augmentations could make the commuter train resemble the one in the book, enhancing their immersion in it.

It is conceivable that this could even affect the experience of reading print. Meta's Depth API lets real objects in the passthrough view occlude virtual content, so a physical book can remain visible in front of a virtual scene [114]. Future wide-FoV AR glasses could render a virtual world around a physical book to enhance focus and immersion when reading.

Overcoming Display Constraints with Spatial Computing

Open-text responses suggested that some respondents experienced practical constraints when carrying out screen-based or reading activities during the commute. For *Reading Electronically*, respondents referred to the need for a 'good screen' and discomfort when reading on a 'short screen like mobile'. Similarly, for *Writing or Editing Electronic Documents*, one respondent noted that they were usually limited to using their phone, which made writing difficult. These responses suggest that the physical limitations of handheld devices may constrain some screen-based commute activities.

XR could address some of these constraints by reducing reliance on the physical size and placement of phone or laptop screens. With spatial computing, digital content is not confined to a single physical display; instead, virtual panels can be positioned around the user and scaled to suit the available space and task. This approach has been explored by Medeiros et al. [110], who examined the use of multi-display XR workspaces in train contexts. In a commuting context, similar techniques could allow passengers to view documents, social media feeds, videos, or e-books on larger virtual displays than would be possible with a handheld device or laptop.

The open responses also indicated that the physical conditions of travel can interfere with reading or working. For example, one respondent cited a 'shaky bus' as a barrier to *Writing or Editing Electronic Documents*, while others referred to carrying constraints when reading print materials. XR may help address some of these issues by presenting content on a stable virtual panel at a comfortable viewing distance and angle, rather than requiring passengers to hold a phone, balance a laptop, or carry additional printed material.

Motion discomfort was also mentioned by one respondent in relation to *Reading Print*, who referred to 'getting motion sickness'. While this was not a widespread theme in the open responses, it suggests a possible design consideration for spatial displays in transit. For example, placing text at eye level on a floating AR panel could reduce the need to look down at a book or laptop. However, further empirical work would be needed to determine whether such displays reduce or exacerbate motion discomfort in moving vehicles.

More generally, spatial displays could support a wider range of 2D activities during travel, including reading e-books, browsing social media, watching video, or editing documents. However, these possibilities should be understood as design opportunities suggested by the reported constraints, rather than as direct evidence that XR would necessarily resolve them.

Virtual Reality for Immersion and Concentration

Although listening to music or audio was already common, 18 respondents wanted to do it better or to start. Audio quality alone may not account for all of what respondents felt was missing; VR could extend the activity beyond playback entirely, for instance by placing the listener in a virtual concert venue or an environment matched to the content being played. Alternatively, visual distractions could be reduced or removed to support focused listening.

For productivity-related tasks, such as *Writing or Editing Electronic Documents*, VR could place the user in an environment they find conducive to focus, such as a quiet office or a natural setting. Similarly, passengers who want to relax or meditate could be placed in a virtual retreat or other calming environment.

Encouraging Better Social Connectivity

With the sensing capabilities of current and future XR devices, social media may not remain confined to photographs and video. Spatial video capture, introduced on the iPhone 15 Pro and viewable on the Apple Vision Pro, already allows moments to be recorded with stereoscopic depth and revisited as scenes rather than as flat images [6], and a plausible extension would place viewers within immersive dioramas of a friend's experience. Similarly, video calling need not remain a framed window. Apple's Spatial Personas render participants as three-dimensional avatars driven in real time by eye, face, and hand tracking, positioned within the viewer's own space with shared spatial context so that participants perceive objects at consistent positions relative to one another [8]. Extending such capabilities to the commute would allow passengers to maintain social connection during otherwise solitary travel, although doing so in a shared carriage raises questions of social acceptability that remain largely unexamined.

Summary and Transition to the Workshops

The survey shows that respondents mainly used travel time to pass the time, most often with audio, social media, and messaging, and that the largest unmet demand concerned reading and viewing scenery. It did not ask directly about awareness of the surroundings; the open-text responses point instead to distraction, limited views, and device constraints. How passengers would balance engagement with awareness when using XR was therefore explored in the workshops, and is developed later as the PIAT (chapter 7).

While the survey identifies a range of unmet needs and opportunities within current commuter experiences, it does not address how these might be effectively supported or reimagined through XR. In particular, it remains unclear what forms of interaction, content, or spatial augmentation would be appropriate for transit environments, and how such experiences might be meaningfully integrated into transit contexts.

To explore these questions, the following section presents a set of ideation-focused participatory design workshops, designed to examine how XR could be used to address the challenges identified in the survey and to generate insights into what forms these experiences might take in practice.

3.3 Participatory Design Workshops: Exploring XR Design Opportunities for Commuters

Two areas stood out in the survey: reading and working, where respondents described distraction and device constraints, and where reading had the highest improvement interest of any activity; and viewing scenery, where respondents described limited or unappealing views. While the survey identified these issues, it did not provide insight into how immersive technologies might address them. Both point to the same underlying constraint, in that commuters must allocate attention between an activity and their surroundings, and any XR intervention will alter how that attention is distributed. Chapter 7 develops this constraint as the PIAT.

To explore potential design directions for XR interventions within this space, a series of participatory design workshops were conducted. These sessions draw on established participatory and co-design approaches [156, 158], in which participants are actively involved in generating and developing design ideas. The aim of these workshops was not to evaluate specific systems, but to collaboratively generate and discuss possible approaches for using augmented and virtual reality to improve the commuting experience. In particular, the sessions explored how immersive technologies might support both focused activities and relaxation while maintaining an appropriate level of environmental awareness.

3.3.1 Participants

Three workshops were conducted, each consisting of four participants ($N = 12$), with a balanced gender distribution (2 male, 2 female per session). Participants were recruited through the university's participant mailing list. Ages ranged from 19 to 35 years. All participants reported commuting five days per week, with journey times ranging from 20 to 75 minutes (one way). The most common mode of transport was bus, followed by rail. Participants had varying levels of prior exposure to XR technologies, ranging from no previous experience to occasional use of VR headsets or AR applications. Participants were randomly assigned to workshops. Each session lasted approximately one hour. The study was approved by the College of Science and Engineering Ethics Committee at the University of Glasgow (application 300230196).

3.3.2 Procedure

Each workshop was a one-hour moderated group session combining individual ideation, group discussion, refinement, and voting, and was audio recorded for later transcription. At the beginning of the session, participants were shown a short set of videos demonstrating examples of AR [1] and VR [141] experiences to establish a shared understanding of the capabilities and differences between the two technologies. Participants were then verbally introduced to two commuting challenges identified in the earlier questionnaire study: difficulty maintaining focus while working or reading during travel, and the desire to relax or enjoy the surrounding travel environment. These challenges were presented as broad design prompts rather than fixed requirements. Participants were asked to generate ideas for how AR or VR technologies might address either or both of these challenges. The workshop activities were structured to support both individual idea generation and collaborative development, following generative participatory design practices [156]. The ideation, discussion, and refinement stages below ran once for each challenge:

Individual ideation (4 minutes) After the challenge prompt had been introduced, participants were first asked to generate ideas individually for how AR or VR technologies might address it. These ideas were written on individual post-it notes, with one idea per note.

Group discussion (10 minutes) Participants then presented their ideas to the group and discussed them collectively. During this stage, participants were encouraged to expand upon or combine ideas proposed by others.

Idea refinement (2 minutes) Following the discussion, participants were given an opportunity as a group to revise or extend their initial ideas.

Preference voting Finally, each participant voted for two ideas they considered most promising from those generated during the session.

Throughout the sessions, the moderator prompted participants to clarify or elaborate on their suggestions where necessary. Transcripts, notes, and artefacts generated during the sessions were later analysed to identify recurring themes across groups.

3.3.3 Analysis

The transcripts, discussions, and artefacts produced during the workshops were analysed by the author using an inductive thematic grouping approach. Ideas and observations from each session were reviewed and organised into clusters based on recurring concepts and concerns. This process is consistent with affinity diagramming techniques commonly used in human-centred design to synthesise qualitative data and identify patterns across participant contributions [21]. Given the exploratory nature of the workshops, the analysis focused on identifying broad themes and design opportunities rather than applying a formal coding framework. Themes were developed from all three workshops; quotations below are drawn from the first two and are labelled by workshop and participant (e.g. W1-3 is participant 3 in workshop 1).

3.3.4 Findings

Several recurring themes emerged across the workshops. These themes provide insight into how commuters envision immersive technologies being used to support both productivity and relaxation during travel. The ideas generated across the workshops reflected a range of approaches to balancing immersive engagement with environmental awareness.

Customisable Virtual Work Environments

Participants widely agreed that immersive environments could support focused activities such as reading or working. However, preferences for the type of environment varied substantially between individuals. Some participants suggested environments resembling traditional workspaces such as offices or libraries, while others preferred more relaxing environments such as beaches, forests, or mountain landscapes. One participant described this variability directly, proposing ‘a customisable environment that each individual can create exactly what they want’ (W2-3). This diversity of preferences suggests that any immersive commuting experience would benefit from a high degree of personalisation. In addition, participants suggested that the environment could be adapted to match the theme of the content being consumed. One participant proposed that ‘if you’re reading some sort of sci-fi book or something, you could have a customisable setting that is related to that sci-fi scenery,’ reasoning that this ‘would really focus you in, because the surroundings would be similar’ (W2-2). The same participant went further, suggesting that such augmentations could be packaged by title as well as by genre: ‘there could also

be some templates, like if you're reading Harry Potter, you could do this, and then there's this package' (W2-2). These observations highlight the importance of customisable environments in immersive travelling scenarios, and anticipate the distinction between genre-level and title-level augmentation examined in chapter 5.

Maintaining Awareness of Surroundings

Participants were nonetheless consistently concerned about losing awareness of their physical surroundings. Maintaining awareness was considered important for several reasons, including personal safety, awareness of other passengers, and knowing when to disembark. One participant summarised the resulting design problem as one of degree rather than kind: 'it's just balancing how much of reality you could still see, so you'd be able to decide how much basically you want to block out' (W1-3).

As a result, many of the proposed solutions incorporated mechanisms to preserve awareness of the real environment. For example, participants suggested representing nearby passengers as simplified silhouettes, avatars, or blurred figures. These approaches were intended to reduce distraction while still providing awareness of other people in the space, and participants distinguished between awareness of people and awareness of place: one described wanting to be 'aware of them, but you don't actually necessarily have to see the people,' adding that 'that way you also know where you are' (W1-2). This concern aligns with prior work on maintaining environmental awareness in immersive systems, such as the concept of reality anchors described by Bajorunaite et al. [17] and discussed in chapter 2.

Concerns about awareness were often expressed in terms of specific journey-management tasks. Participants in different workshops independently proposed location-based stop notifications, with one suggesting that 'the glasses could connect, they would have a GPS' so that 'even if you're fully immersed you could just get a notification saying your bus stop is coming up' (W1-3), and another proposing a link to mapping software so the system knows 'when you're close to your stop' (W2-3). Others wanted a mechanism to interrupt immersion, describing the need for 'something that brings you back to reality, whether it's like an emergency or just somebody bumping into you' (W2-3), and identified personal belongings as elements that should remain visible (W2-2). One participant revised their position over the course of the discussion, having initially favoured full immersion: 'my answer kind of changed from the first thing we talked about, to keep that sense of reality, because you still want to know you're on a train, after we've had all the discussions, I do feel like going all the way might be dangerous' (W1-3). Reduced awareness was also described as carrying an attentional cost of its own, with one participant observing that 'my lack of control would end up being the biggest distraction where I can't read' (W2-3).

Environmental Filtering and Focus

A frequently suggested concept was selectively filtering or modifying the surrounding environment in order to reduce distraction. Rather than completely replacing the physical environment with a virtual one, participants proposed techniques such as blurring or dimming elements of the environment that were not directly relevant to the user's current activity. One suggestion involved maintaining a clear region immediately surrounding the user while gradually blurring or fading elements further away, described as leaving 'space from you that can still be clear, but anything further would be like a blur' such that 'you would see what's happening around you but it wouldn't be distracting you' (W1-1). This approach was described as creating a form of 'focus bubble' that could help the user concentrate on their task. A participant in the second workshop arrived at a comparable proposal independently, drawing an analogy with 'how depth of field blur works in video games' to 'blur out the other passengers' while remaining able to 'see your physical book and around you' (W2-1). Two features of these proposals are notable. First, participants consistently described graded rather than binary reduction, suggesting that figures could be made partially transparent so that 'you were aware they're there but it doesn't block your scenery anymore' (W1-1). Second, several participants applied the same logic to sound as to vision, proposing to 'not switch off the sounds around you, but the same way as you're not completely switching off the visuals, just audio blurring' (W2-3). Complete removal of other passengers, by contrast, was rejected. Participants anticipated that removing people visually while their presence remained audible or physically detectable would be disconcerting, noting that one might 'still hear the voice somewhere' while they are 'still kicking your legs' (W2-3), and that being 'alone in a hole' would be unsettling (W2-2). Such ideas highlight a preference for partial environmental modification rather than complete immersion.

Modifying Exterior Scenery

For the relaxation use case, participants proposed modifying the scenery visible outside the vehicle while leaving the interior of the carriage unchanged. For example, the view through the train window could be replaced with a more visually appealing landscape while the interior of the carriage and other passengers remained visible. Participants described this as preserving the character of the journey: one noted a preference for 'changing the outside' while retaining the interior because 'every transport is different' (W1-1), and another reported not wanting to change the interior 'because it adds to the realness of it,' adding that changing the vehicle 'too much would just feel weird' (W1-2). A participant in the second workshop described the same arrangement as remaining 'sat in whatever mode of transportation you are, but replace what's happening outside of the windows' (W2-3).

Others proposed modifying the interior itself, but in ways that retained its identity as a vehicle and preserved awareness of other passengers. One suggested altering the apparent travel class so that 'you change the whole interior, so you're maybe like in first class,' anticipating

that this 'could create like a placebo effect' (W2-3). Another described changing the interior décor while explicitly retaining passenger visibility: 'you can still see the passengers, so you still know if these people are here, but they're sitting on some fancier place' (W2-2). Participants suggested that these approaches could provide a more relaxing visual experience while preserving awareness of the immediate surroundings.

3.3.5 Implications for XR Media Design

The participatory design discussions revealed a consistent tension between immersive experiences and the need to remain aware of the surrounding environment. While participants expressed interest in immersive environments that could support focus or relaxation, they also demonstrated a strong preference for retaining some connection to the physical world. Many of the proposed ideas involved selectively modifying or augmenting the existing environment rather than replacing it entirely. This suggests that XR systems designed for commuting contexts should prioritise approaches that balance immersion with environmental awareness, rather than treating these as mutually exclusive states.

Several design directions emerged from the discussions. Participants frequently proposed highly personalised immersive environments to support focused activities such as reading or working, including virtual offices, libraries, and natural environments. However, the diversity of these preferences highlighted a key challenge: the effectiveness of such environments appeared to be highly dependent on individual taste. As a result, designing and evaluating a single representative prototype for these use cases would require a level of personalisation beyond the scope of a controlled experimental study. Consequently, while these ideas are valuable in illustrating user desires, they were not pursued further within this thesis as a primary design direction.

Another strong theme was the desire to reduce distractions by filtering or obscuring elements of the surrounding environment, such as blurring other passengers or representing them as simplified figures. Although this direction directly addresses the problem of environmental distraction, it primarily constitutes a modification of visual fidelity rather than a richer transformation of the media experience. As such, it offered limited scope for sustained investigation beyond a single study, and was therefore not selected as the central focus of the thesis.

In contrast, a more promising direction emerged from suggestions that involved selectively augmenting elements of the environment while preserving the overall context of the journey. In particular, participants proposed modifying the scenery outside the vehicle while maintaining awareness of the carriage interior and other passengers. This approach offered a balance between immersion and awareness, and suggested a design space in which digital content could be integrated with the physical environment rather than replacing it entirely.

Given these insights, the research direction was refined toward media consumption as a tractable context for investigating the PIAT. Several survey- and workshop-derived opportunities were considered. Environmental filtering and scenery replacement were promising, but

each suggested a relatively narrow intervention: reducing visual distraction or altering the exterior view. These directions were therefore less suitable as the central design mechanism for a sustained programme of experimental work. Reading was also strongly supported by the survey, but studying reading in AR would first require establishing an appropriate baseline and addressing substantial display-legibility constraints, including text resolution, comfort, and comparison with physical or electronic reading media.

Video viewing was therefore selected not because the survey ranked it highly (it was the primary activity of only one respondent), but because it provided a controllable media-consumption case through which environmental augmentation could be systematically investigated. A video stimulus allows the core media content to remain constant while the surrounding environment is augmented to different degrees, making it suitable for examining how ARTV Ornamentation affects immersion, enjoyment, and environmental awareness. Audio-based media, including audiobooks, remained a plausible future direction, particularly given the popularity of listening to music or audio in the survey, but was not pursued within this thesis because the development and evaluation of the video-based studies required greater scope than initially anticipated.

This shift in focus led to the development of the concept of ARTV Ornamentation, in which video content is augmented through spatial extensions into the surrounding environment. By building on the idea of selectively augmenting the physical space, particularly elements such as the carriage interior or exterior view, this approach aims to enhance engagement with media while maintaining awareness of the real-world context. These design considerations directly informed the development of the first experimental study, which explores how such augmentations influence the experience of media consumption in a simulated commuting environment.

3.4 Conclusions

This chapter addresses RQ1 (*How could XR improve the current experience of immersive media on public transport?*) by characterising how passengers currently use their commuting time and where they would welcome improvement. In doing so it narrows a design space that remains relatively underexplored: while prior research has examined particular XR applications in transit settings, the breadth of passenger-centred activities and design opportunities has not been systematically characterised. The survey showed unmet demand for more engaging and better-supported use of travel time, and the workshop participants wanted more immersive experiences while remaining sensitive to constraints imposed by safety, social interaction, and journey management, consistently seeking immersion while retaining awareness of their surroundings. This is the tension the thesis investigates, and chapter 7 returns to it in light of the empirical findings.

Analysis across both studies identified three broad XR design opportunities: reducing or filtering environmental distractions, enhancing or replacing external scenery, and extending media content into the surrounding environment. The first two offer comparatively limited scope for

systematically investigating the immersion–awareness relationship, and the workshop data revealed a further difficulty with immersive work environments such as virtual offices or natural settings: preferences varied so substantially between individuals that no representative prototype could be defined for controlled evaluation.

Extending media content into the surrounding environment therefore offered the strongest fit with the thesis's research problem. It connects to the media-consumption activities that commuters already undertake and wish to improve, it permits systematic modification of the surrounding environment, and it inherently preserves visual access to the real world. This allows both immersion and environmental awareness to be measured. Workshop participants had themselves proposed selectively augmenting the environment while preserving awareness of the carriage interior, which suggested a design space in which digital content is integrated with the physical environment rather than replacing it. ARTV Ornamentation was therefore adopted as the primary design mechanism investigated in the chapters that follow.

What the present studies cannot establish is whether environmental augmentation measurably increases immersion without a corresponding cost to awareness, or how that balance is affected by the constraints of contemporary AR hardware. Chapter 4 addresses both, examining how device FoV and ornamentation intensity shape enjoyment, engagement, and environmental awareness.

Chapter 4

Using ARTV Ornamentations to Balance Media Immersion and Environmental Awareness for Passengers on Public Transport

4.1 Introduction

Chapter 3 found that workshop participants wanted more immersive experiences during travel, but only if they remained aware of their surroundings, and consistently proposed augmenting the travel environment rather than replacing it. These findings highlighted both the potential for XR technologies to support passenger entertainment and the need to balance engagement with environmental awareness.

In response to this tension, this thesis introduces the concept of *ARTV Ornamentation*: ARTV elements that function as themed, aesthetic, and environment-integrated augmentations. By spatially embedding media theming within the user's physical surroundings, ARTV Ornamentation is hypothesised to deepen immersion without fully occluding reality. This approach looks to combine the high-immersion affordances of VR with the awareness-preserving characteristics of AR.

This chapter investigates ARTV Ornamentation as a design mechanism for balancing immersion and awareness within passenger transportation. While prior ARTV research has primarily focused on domestic viewing, this work provides the first empirical investigation of ARTV Ornamentation in a simulated transport context. Specifically, the chapter examines how two factors, FoV and ornamentation intensity, shape the balance between immersion and environmental awareness. These factors are central to the use of ARTV Ornamentation on public transport because they determine how much digital content can be presented, where it appears in relation to the user's surroundings, and how strongly it competes with the physical environment for atten-

tion. FoV is treated as a hardware-imposed constraint on the visible relationship between digital content and the real world. A narrow FoV may limit the effectiveness of ARTV Ornamentation, as the primary media display can occupy a large proportion of the available digital view, leaving limited space for peripheral augmentations. Conversely, a wider FoV may allow ornamentations to be perceived more naturally in peripheral vision (seen 'out of the corner of the eye'), which is hypothesised to increase media immersion without requiring direct foveal attention. The other variable, ornamentation intensity, is investigated as the primary design-led variable, controlling the degree to which media content is embedded within and visually transforms the surrounding environment. The two factors are likely to interact: while a wider FoV allows for more peripheral immersion, it also increases the potential for digital content to interfere with environmental awareness. By manipulating both parameters, this chapter examines how the balance between immersion and awareness changes as ornamentation increases, and whether this depends on FoV.

The experiment described in this chapter addresses RQ2 (*How can ARTV with Ornamentation be used to balance media immersion and environmental awareness in passenger transportation?*) and RQ3 (*What is the impact of device FoV, ornamentation density, and ornamentation type on the media immersion–environmental awareness trade-off?*), contributing evidence on the effectiveness of ARTV Ornamentation as a passenger media solution. These results establish the foundational knowledge required for the design refinements and extensions explored later in chapter 5.

4.2 Study Overview

To evaluate ARTV Ornamentation, this study examines passengers' experience of ARTV video content within a simulated train environment. The commuter activity findings from chapter 3 identified media consumption as a popular and expandable passenger activity, motivating investigation into ARTV Ornamentation as a mechanism for enhancing the commute experience. This study investigates whether ornamentation can increase immersion and media engagement while maintaining some level of environmental awareness. In doing so, the study considers whether ARTV Ornamentation can move passenger media experiences toward greater immersion, and at what cost to awareness of the travel environment.

Due to limitations of contemporary AR hardware at the time of this research (restricted FoV, limited display luminance, and the difficulty of experimentally manipulating display characteristics on consumer AR devices), this study evaluates ARTV Ornamentation using AR simulated within VR. Simulating AR in VR enables systematic manipulation of design variables such as FoV and ornamentation density while maintaining experimental consistency and repeatability. This approach has been used successfully in prior passenger and XR studies [188, 110], and is reviewed in section 2.7.

The sub-research questions for this study are:

RQ2-1 **AR vs VR:** How does ARTV Ornamentation compare to immersive VR in terms of enjoyment, engagement, and immersion?

RQ2-2 **Degree of ARTV Ornamentation:** How does increasing the degree of ARTV ornamentation (DoA) affect media immersion, engagement, and environmental awareness, and how does this shape the trade-off between them?

RQ3-1 **FoV:** Does device FoV affect the media immersion of a user consuming ARTV content?

4.3 Experiment

4.3.1 Design

A within-subjects design compared eight viewing conditions: two baselines and six ARTV Ornamentation conditions. Participants completed all conditions seated in a lab while wearing a VR headset.

Media content. Participants watched clips from the BBC nature documentary *Planet Earth II* [19]. Its jungle settings provided a clear visual theme for the ornamentations, and nature documentaries have previously been used as augmentable content in ARTV research [150]. Participants watched a different clip in each condition. Clip assignment was rotated so that each condition was paired with a different clip before any clip was reused; no participant saw the same clip twice.

Journey framing. In all conditions except the VR Baseline, participants viewed a virtual train carriage, which they were asked to treat as the real travel environment. Rail was chosen because it is a common context for video consumption and passengers report relatively high comfort watching video there [161]. As in Medeiros et al. [110], the scenario was framed as a long-distance journey, which supports extended media engagement while still requiring awareness of nearby passengers [17]. In Bajorunaite et al.'s terms the journey was externally managed: participants were seated and did not need to monitor stops, make route decisions, or prepare to alight. The carriage contained seats, tables, overhead storage, windows, and neighbouring passenger avatars, whose models and idle animations were sourced from Adobe Mixamo [3]. These features provided the transit context and the surfaces around which ornamentations were placed by hand.

Baselines. The VR Baseline presented a fully virtual jungle containing a large 2D screen; the carriage was not visible. The screen measured 150 inches diagonally and was positioned 4.5 m

from the participant, following home-cinema recommendations [38]; it therefore subtended approximately 40.5° horizontally and 23.4° vertically. The VR Baseline was the condition with the strongest immersive properties while retaining a conventional 2D video format. The AR Baseline presented a 2D screen overlaid on the carriage with no ornamentation, reflecting the floating-screen approach of AR glasses such as the Xreal Air 2 Ultra [199]. This screen measured 0.92 m × 0.53 m and was positioned 1 m from the participant, subtending approximately 49.4° horizontally and 29.8° vertically; the same screen was used in all ARTV conditions. The two baselines provided reference points at either end of the conditions: full replacement of the environment and minimal augmentation of it.

Factors. Two factors were manipulated in the remaining six conditions: DoA (three levels) and FoV (two levels), giving the conditions shown in fig. 4.1 and table 4.1.

The three ornamentation levels did not vary along a single dimension. At **Low** ornamentation, 12 plant models [173] were placed on the participant's table and the overhead storage, providing a minimal thematic anchor that did not extensively obscure the carriage. At **Medium** ornamentation, the number of models increased to 55, mixing plants, rocks, and foliage debris throughout the carriage, while leaving gaps through which the carriage remained visible. **High** ornamentation used the same 55 models as Medium and additionally retextured the carriage surfaces to match the jungle theme. The step from Low to Medium therefore changed the density of ornamentation, whereas the step from Medium to High changed the treatment of the carriage itself; the implications are discussed in section 4.5.5. With most visible surfaces retextured, the High level moves towards augmented virtuality on Milgram's reality–virtuality continuum [122, 120].

To simulate overlaying a texture onto an object in AR, the jungle textures in the High conditions were rendered at 15% transparency so that the original carriage textures remained visible, and all added models were rendered at 5% transparency. These values were chosen through experimentation with a Microsoft HoloLens 2, keeping the ornamentations noticeably distinct from the carriage. They model a display with controlled augmentation opacity rather than the optical behaviour of current additive see-through combiners, in which dark image regions render as transparent and contrast falls as ambient light rises. Opacity and occlusion control are active engineering targets for OST displays [146, 203, 204, 67], and are already partly addressed in consumer devices through electrochromic dimming [27]. The conditions therefore approximate a class of display in which such control is available (section 4.5.5).

The two FoV levels span the augmentation envelopes of current and plausible near-term displays. The High FoV condition used the Quest 2 FoV (104° horizontal, 98° vertical), representing the wider envelope of passthrough AR; this was also the physical limit of the headset, so it bounded the visible area in every condition. The Low FoV condition restricted ornamentation rendering to 50° horizontal and 40° vertical, approximating a near-term OST glasses display.

Table 4.1: The FoV within which ornamentations were rendered, and the augmentation used, in each condition.

Condition	**FoV (H×V, °)**	**Augmentation**
AR Baseline	104×98	None
Low FoV – Low	50×40	12 models
Low FoV – Medium	50×40	55 models
Low FoV – High	50×40	55 models and retextured carriage
High FoV – Low	104×98	12 models
High FoV – Medium	104×98	55 models
High FoV – High	104×98	55 models and retextured carriage
VR Baseline	104×98	Fully virtual environment, no carriage

This value sits within the range of announced and demonstrated hardware: consumer glasses currently ship at approximately 20° [60, 115], HoloLens 2 provides 43°×29°, and wide-field prototypes such as Meta's Project Orion and Lumus's ZOE waveguide reach approximately 70° [116, 178]. It is also consistent with the increments between successive OST generations: HoloLens 1 to HoloLens 2 increased the FoV from 30°×17.5° to 43°×29°, and a comparable further step places a subsequent generation in this region.

In the Low FoV conditions, ornamentations were rendered only within this restricted region while the carriage remained visible across the full Quest 2 FoV (fig. 4.1, panels c, e, and g). This produced a visible boundary at the edge of the augmentation region that a physical OST display would not show, since the real environment is seen across the eye's natural field of view. Results for the Low FoV conditions therefore reflect both a smaller augmentation envelope and the presence of this boundary. In all conditions the video screen lay entirely within the assigned FoV, so it was clipped only if participants looked away from it.

Only one AR Baseline was included, using the High FoV setting. Because it contained no ornamentation, a Low FoV version would have looked essentially the same, and the High FoV version provided the more relevant comparison with current passthrough viewing.

4.3.2 Measures

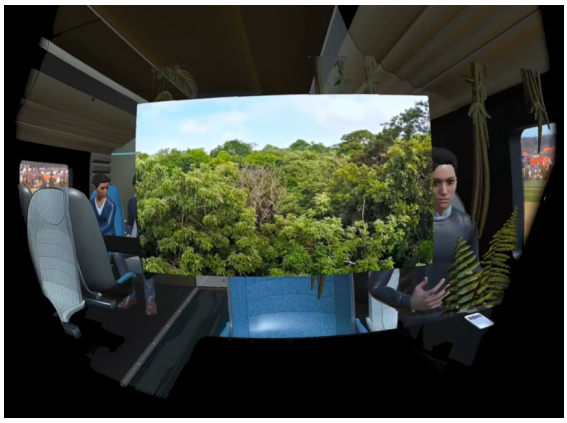
Demographic data were collected before the study. After each condition, participants answered a combination of pre-existing questionnaires and two additional questions developed for this study. The questionnaires were the Film-IEQ [149], administered with 23 of its 24 items (section 4.3.5), the UEQ-S [157], and the UES-SF [130], following a previous evaluation of ARTV [150]. The two additional questions asked participants how much they agreed, on a 5-point Likert scale from Strongly Agree to Strongly Disagree, with two statements: 'I was aware of other passengers whilst watching the video' and 'I was aware of the virtual modifications (trees and plants) whilst watching the video'. These served two distinct purposes. The first item provided



(a) AR Baseline



(b) VR Baseline



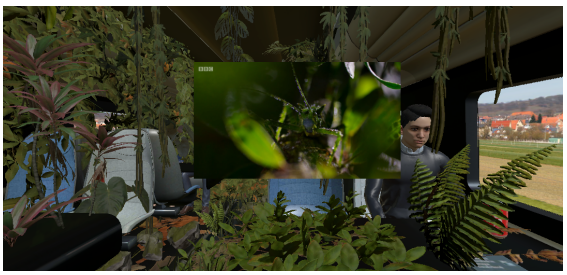
(c) Low FoV – Low



(d) High FoV – Low



(e) Low FoV – Medium



(f) High FoV – Medium



(g) Low FoV – High



(h) High FoV – High

Figure 4.1: The eight viewing conditions. (a) The AR Baseline: a floating screen in the carriage with no ornamentation. (b) The VR Baseline: the same video in a fully virtual jungle. (c–h) The remaining six conditions cross FoV (High, Low) with DoA levels (Low, Medium, High). Each row pairs the Low and High FoV versions of the same DoA level, with the Low FoV condition shown on the left and the High FoV condition on the right. In the Low FoV conditions, jungle ornamentation is rendered only within a central $50^\circ \times 40^\circ$ region, leaving a visible boundary with the unmodified carriage.

a measure of environmental awareness, the construct of interest in this study. The second functioned as a manipulation check, confirming that participants perceived the ornamentations at the intended level; it is not a measure of environmental awareness and is not interpreted as one.

After completing all experimental conditions, participants engaged in a semi-structured exit interview and questionnaire. The semi-structured interview asked participants to reflect on their preferred and least preferred conditions, the extent to which each condition affected their immersion in the video, their awareness of the train environment and other passengers, any sources of distraction or discomfort, and whether they would use similar ARTV experiences during real train travel. In the exit questionnaire, participants ranked the eight conditions from most to least preferred and indicated, for each condition, whether they would use it instead of a phone or laptop to watch video on a train.

The interviews were analysed with a single-coder, inductive thematic analysis at a semantic level, focusing on participants' explicit descriptions of immersion, awareness, distraction, comfort, and preference:

1. **Transcription and data preparation.** Audio recordings were transcribed with *WhisperX* [14], checked manually against the audio, and anonymised. The audio files were then deleted in accordance with the data management plan.

2. **First-cycle coding.** Following Saldaña [154], an inductive coding pass assigned open codes to transcript segments, allowing concepts to emerge from participants' descriptions rather than from predefined categories.

3. **Second-cycle pattern coding.** Codes were grouped into broader categories over several iterations; for example, comments about peripheral distractions and the visibility of exits were grouped under environmental awareness.

4. **Thematic synthesis.** Categories were synthesised into themes addressing the research questions, and checked for consistency against the transcripts.

4.3.3 Participants

Participants were recruited through the university's participant mailing list and flyers posted on campus. The 24 participants were aged 20 to 60 ($M = 31.79$, $SD = 10.87$); 12 were men and 12 were women. All 24 completed all eight conditions and neither withdrew nor asked to stop during the study.

Participants received a £10 Amazon voucher for the 75-minute study.

4.3.4 Procedure

Participants read an information sheet, consented to the study and to audio recording of the interview, and answered the demographic questions. They then sat on a chair wearing a Quest 2 headset connected to a computer running the study via Quest Link, with all conditions rendered at 72 Hz, the headset's default refresh rate. Condition order and clip assignment were counterbalanced across participants using a Latin square. The study was approved by the College of Science and Engineering Ethics Committee at the University of Glasgow (application 300220080).

4.3.5 Analysis

Scale and subscale scores were calculated following the published procedures for each questionnaire [149, 157, 130], with two exceptions for the Film-IEQ. It was scored on its 23 administered items, because item 13 (Captivation) was omitted from the questionnaire in error. Its three Real-World Dissociation items were reversed by negation rather than by subtraction from 8, so Real-World Dissociation and overall Film-IEQ scores are a constant 24 points lower than under the published key; this shifts every condition equally and does not affect any comparison. Internal consistency was estimated within each condition, so that each estimate used one response per participant (table 4.2).

Factorial analysis. The six ARTV Ornamentation conditions formed a 2 (FoV) $\times$ 3 (DoA) within-subjects design; the two baselines were excluded from this analysis because they are not levels of either factor. Overall Film-IEQ scores did not deviate from normality (Shapiro–Wilk, $W = 0.99$, $p = .13$) and were analysed with a two-way repeated-measures analysis of variance (ANOVA), applying the Greenhouse–Geisser correction where Mauchly's test indicated a violation of sphericity. *Post hoc* comparisons were paired and Bonferroni-corrected. All other measures violated normality and were analysed with the Aligned Rank Transform (ART) [193], which transforms the data before a repeated-measures ANOVA with participant as the error term. *Post hoc* comparisons used Aligned Rank Transform Contrast (ART-C) contrasts [43], Bonferroni-corrected for the three DoA levels. Where the FoV $\times$ DoA interaction was significant, simple effects (pairs of conditions sharing a level of one factor) were tested: for the Film-IEQ, DoA comparisons were Bonferroni-corrected within each FoV level and the single FoV comparison at each DoA level was left unadjusted; for the other measures, all simple effects were Holm-corrected.

Comparisons with the baselines. Each ARTV condition was compared with each baseline within a one-way repeated-measures model over all eight conditions (parametric for the Film-IEQ; ART with ART-C contrasts otherwise), using Dunnett-adjusted treatment-versus-control con-

Table 4.2: Internal consistency of each scale, estimated separately within each of the eight conditions ($n = 24$ per estimate). Values are the median across conditions, with the range in brackets.

Scale	Items	Cronbach's α	McDonald's ω_t
Film-IEQ	23	.86 [.80, .91]	.90 [.88, .92]
Captivation	11	.88 [.80, .92]	.91 [.88, .93]
Real-World Dissociation	3	.86 [.79, .92]	.86 [.79, .93]
Comprehension	4	−.05 [−.40, .28]	.82 [.75, .91]
Transportation	5	.80 [.72, .93]	.82 [.75, .93]
UEQ-S	8	.90 [.83, .93]	.91 [.86, .94]
Pragmatic	4	.87 [.83, .93]	.88 [.85, .94]
Hedonic	4	.89 [.85, .96]	.90 [.86, .96]
UES-SF	12	.92 [.88, .95]	.93 [.89, .96]
Focused Attention	3	.87 [.82, .92]	.88 [.83, .92]
Perceived Usability	3	.85 [.63, .89]	.88 [.74, .92]
Aesthetics	3	.94 [.80, .96]	.94 [.84, .96]
Reward	3	.95 [.75, .97]	.96 [.80, .97]

The negative α for Comprehension reflects the published scoring key: item 12 asks how *challenging* the content was, whereas the other three items ask how easy it was to understand and how well it was understood, but only the Real-World Dissociation items are reversed [149]. Items worded in opposite directions lower α, which is sensitive to direction; ω_t is not. Rigby et al. [149] likewise report their lowest reliability for this subscale ($\alpha = .66$). With item 12 reversed, the median α for Comprehension was .79 [.62, .90]; the published key was retained for all reported analyses.

trasts. The single awareness items were also fitted with cumulative link mixed models as a robustness check; these agreed with the results reported below.

Sensitivity and equivalence. No *a priori* power analysis was conducted. With 24 participants, a single paired comparison had 80% power to detect effects of $d_z \geq 0.60$ at $\alpha = .05$; for a family of seven comparisons against one control (Bonferroni approximation), the corresponding threshold is $d_z \geq 0.78$. Non-significant comparisons with the VR Baseline were therefore followed by two one-sided tests (TOST) equivalence tests [87], with bounds set in raw scale units as $d_z = 0.60$ multiplied by the standard deviation of the paired differences.

Preferences. Rankings were compared with a Friedman test (effect size Kendall's W), followed by Wilcoxon signed-rank tests with Holm correction. Prospective use (yes/no) was compared with Cochran's Q, followed by exact McNemar tests with Holm correction.

4.4 Results

Results are presented by factor, followed by comparisons with the baselines, preferences, and the interview themes. Table 4.3 summarises the factorial analysis and table 4.4 the comparisons

Table 4.3: Effects of FoV and DoA within the six ARTV Ornamentation conditions, with significant *post hoc* comparisons (p in brackets; adjusted unless marked ‡). The Film-IEQ was analysed with a repeated-measures ANOVA (Greenhouse–Geisser corrected); the other measures with the ART and ART-C contrasts.

Measure	Effect	F	df	p	Significant *post hoc*
Film-IEQ	FoV	13.18	1, 23	.001	High > Low (.001)
	DoA	7.49	1.42, 32.58	.005	High > Low (.013)
	Interaction	4.30	1.97, 45.20	.020	At High FoV: Medium > Low (.017), High > Low (.004); none at Low FoV; at Medium and High DoA: High FoV > Low FoV (.027‡, <.001‡)
UEQ-S	FoV	20.78	1, 115	<.001	High > Low (<.001)
	DoA	4.73	2, 115	.011	High > Low (.009)
	Interaction	2.81	2, 115	.064	–
UES-SF	FoV	20.88	1, 115	<.001	High > Low (<.001)
	DoA	8.81	2, 115	<.001	High > Low (<.001)
	Interaction	2.77	2, 115	.067	–
Awareness of other passengers	FoV	11.82	1, 115	<.001	Low > High (<.001)
	DoA	8.95	2, 115	<.001	Low > Medium (.001), Low > High (<.001)
	Interaction	4.32	2, 115	.016	At High FoV: Low > Medium (.003), Low > High (.002); at Medium and High DoA: Low FoV > High FoV (.007 each)
Awareness of augmentations (manipulation check)	FoV	20.42	1, 115	<.001	High > Low (<.001)
	DoA	7.67	2, 115	<.001	Medium > Low (.007), High > Low (.001)
	Interaction	0.33	2, 115	.72	–

‡Unadjusted: a single FoV comparison at each DoA level (table D.3).

with the baselines; full results for all subscales, with effect sizes, are in appendix D.

4.4.1 Effect of FoV

Media immersion. FoV had a significant effect on the overall Film-IEQ, $F(1,23) = 13.18$, $p = .001$, with higher scores in the High FoV conditions. The *Real-World Dissociation* ($F(1,115) = 9.24$, $p = .003$) and *Transportation* ($F(1,115) = 17.90$, $p < .001$) subscales showed the same pattern; *Captivation* and *Comprehension* did not differ.

Experience. High FoV conditions scored higher on the UEQ-S, $F(1,115) = 20.78$, $p < .001$, including its *Pragmatic* and *Hedonic* subscales.

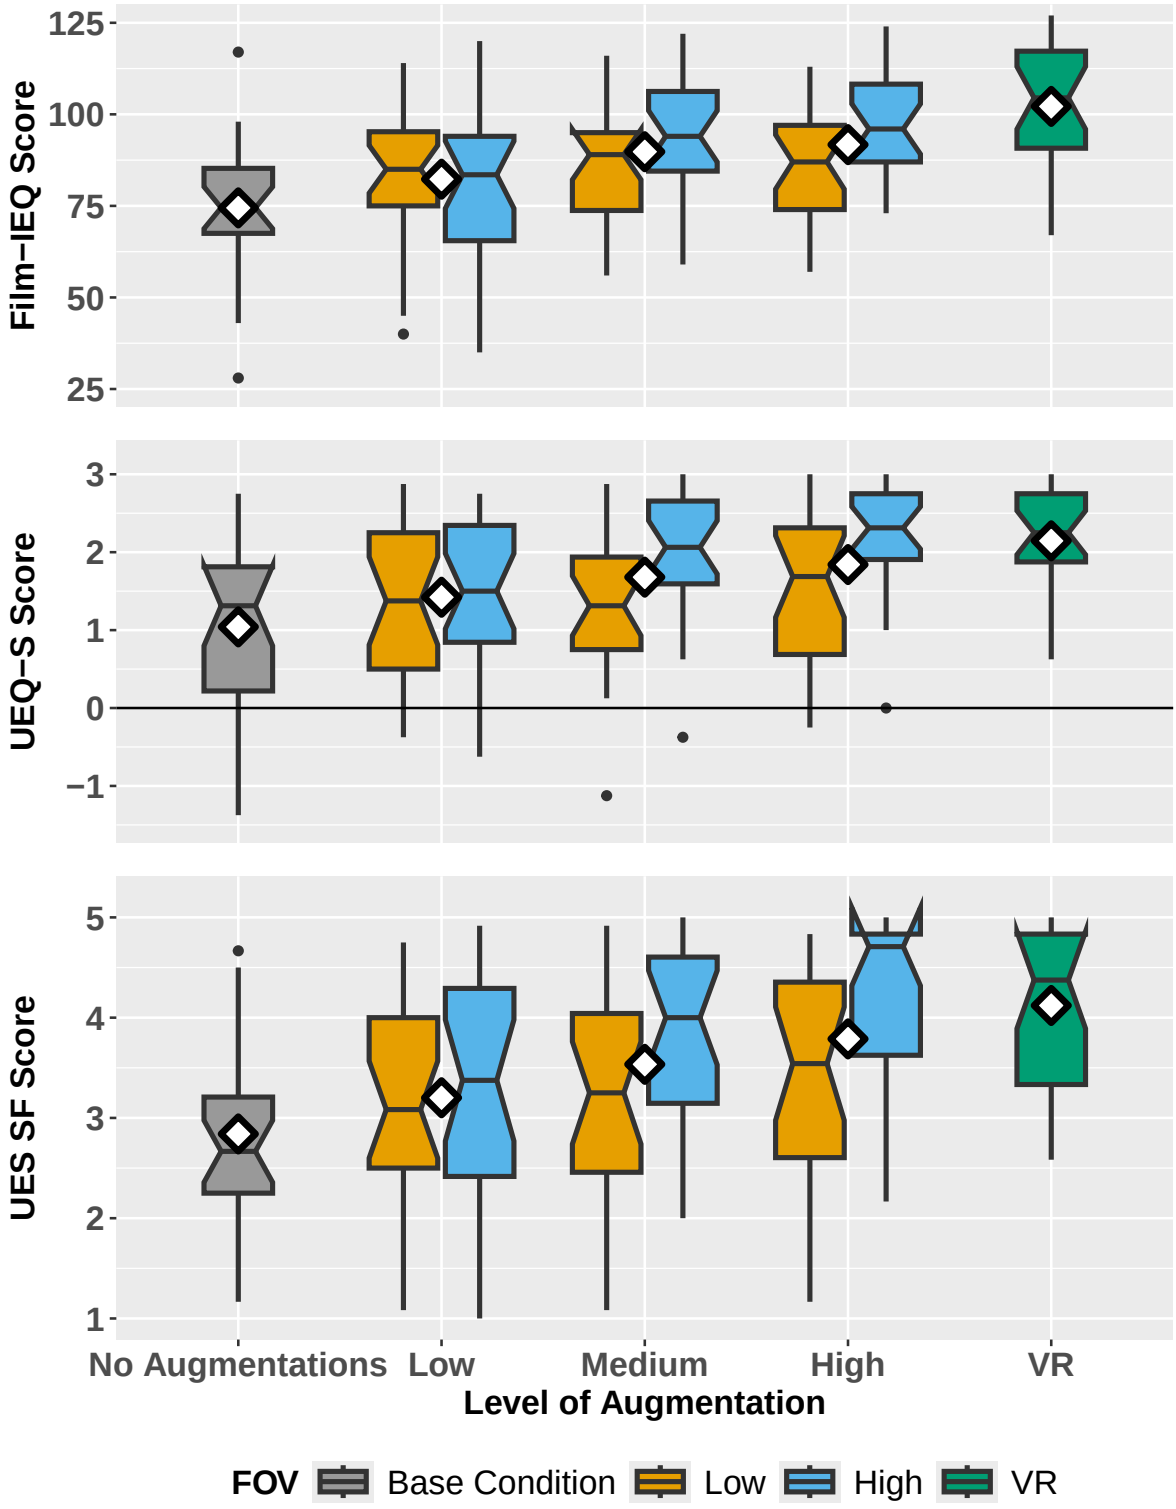


Figure 4.2: Notched box plots of overall Film-IEQ, UEQ-S, and UES-SF scores by DoA and FoV, with the AR and VR Baselines for reference. Notches indicate an approximate 95% confidence interval around the median.

Table 4.4: Conditions that differed significantly from each baseline (Dunnett-adjusted $p < .05$ within a repeated-measures model; parametric for the Film-IEQ, ART-C otherwise); arrows give the direction relative to the baseline. HF/LF: High/Low FoV; L/M/H: Low/Medium/High ornamentation.

Measure	vs AR Baseline	vs VR Baseline
Film-IEQ	↑ HF-H, HF-M, LF-H, LF-L	↓ HF-L, LF-L, LF-M, LF-H
UEQ-S	↑ HF-H, HF-M	↓ HF-L, LF-L, LF-M, LF-H
UES-SF	↑ HF-H, HF-M	↓ HF-L, LF-L, LF-M, LF-H
Awareness of other passengers	↓ HF-H, HF-M	↑ all six
Awareness of augmentations (manipulation check)	↑ all six	↓ HF-L, LF-L, LF-M, LF-H

Engagement. High FoV conditions scored higher on the UES-SF, $F(1,115) = 20.88$, $p < .001$, and on its *Focused Attention* subscale, $F(1,115) = 21.40$, $p < .001$. *Perceived Usability* showed the opposite direction, being higher in the Low FoV conditions, $F(1,115) = 13.15$, $p < .001$.

Awareness. Self-reported awareness of other passengers was lower in the High FoV conditions, $F(1,115) = 11.82$, $p < .001$. As intended, participants reported greater awareness of the ornamentations with the High FoV (manipulation check; $F(1,115) = 20.42$, $p < .001$).

4.4.2 Effect of Degree of ARTV Ornamentation

Media immersion. DoA had a significant effect on the overall Film-IEQ, $F(1.42,32.58) = 7.49$, $p = .005$; High ornamentation scored higher than Low ($p = .013$), while Medium did not differ significantly from either. *Real-World Dissociation* ($F(2,115) = 5.38$, $p = .006$; High > Low) and *Transportation* ($F(2,115) = 9.00$, $p < .001$; Medium and High > Low) showed the same pattern.

Experience. DoA affected the UEQ-S, $F(2,115) = 4.73$, $p = .011$ (High > Low), driven by the *Hedonic* subscale ($F(2,115) = 11.05$, $p < .001$; Medium and High > Low); the *Pragmatic* subscale did not differ.

Engagement. DoA affected the UES-SF, $F(2,115) = 8.81$, $p < .001$ (High > Low), and its *Focused Attention*, *Aesthetics*, and *Reward* subscales (all $p < .01$); *Perceived Usability* did not differ ($p = .14$). *Aesthetics* was the only measure on which High ornamentation scored significantly higher than Medium ($p = .008$).

Awareness. DoA affected self-reported awareness of other passengers, $F(2,115) = 8.95$, $p < .001$: awareness was higher at Low ornamentation than at Medium ($p = .001$) or High ($p < .001$), with no difference between Medium and High. Awareness of the ornamentations increased from Low to Medium ($p = .007$) and from Low to High ($p = .001$).

4.4.3 Interaction between FoV and Degree of ARTV Ornamentation

The interaction was significant for the Film-IEQ ($F(1.97,45.20) = 4.30$, $p = .020$) and for awareness of other passengers ($F(2,115) = 4.32$, $p = .016$), but not for the UEQ-S, the UES-SF, or the manipulation check. In both cases, ornamentation mattered only when the FoV was wide. With High FoV, Medium and High ornamentation produced higher Film-IEQ scores than Low ($p = .017$ and $p = .004$) and lower awareness of other passengers ($p = .003$ and $p = .002$). With Low FoV, ornamentation level made no significant difference to either measure. Correspondingly, FoV affected immersion and awareness at Medium and High ornamentation but not at Low. The interaction was also significant for three subscales, *Transportation*, *Hedonic* quality, and *Reward*, with the same pattern: at High FoV, Medium and High ornamentation scored above Low, with no differences at Low FoV (tables D.2 and D.3). For *Perceived Usability* the interaction was also significant ($p = .038$), but its only significant simple effect was lower usability for High FoV–High than Low FoV–High.

4.4.4 Comparisons with the AR and VR Baselines

The baselines serve as reference points rather than levels of the factorial design: the AR Baseline represents conventional floating-screen viewing and the VR Baseline full replacement of the environment (table 4.4).

Compared with the AR Baseline. High FoV–High and High FoV–Medium ornamentation produced higher Film-IEQ scores than the AR Baseline (mean differences of 22.8 and 19.8 points; both $p < .001$). Two Low FoV conditions were also higher: High ornamentation (11.7 points, $p = .013$) and Low ornamentation (8.5 points, $p = .030$); Low FoV–Medium did not reach significance ($p = .057$). Of these Low FoV results, only Low FoV–High was robust to the method of analysis: under the ART, Low FoV–Medium was significant ($p = .039$) and Low FoV–Low was not ($p = .084$). The two High FoV conditions with Medium or High ornamentation were the only ones to score higher than the AR Baseline on the UEQ-S and UES-SF, and the only ones to reduce awareness of other passengers relative to it (both $p < .001$). All six ARTV conditions increased awareness of the ornamentations relative to the AR Baseline, indicating that participants perceived the manipulation.

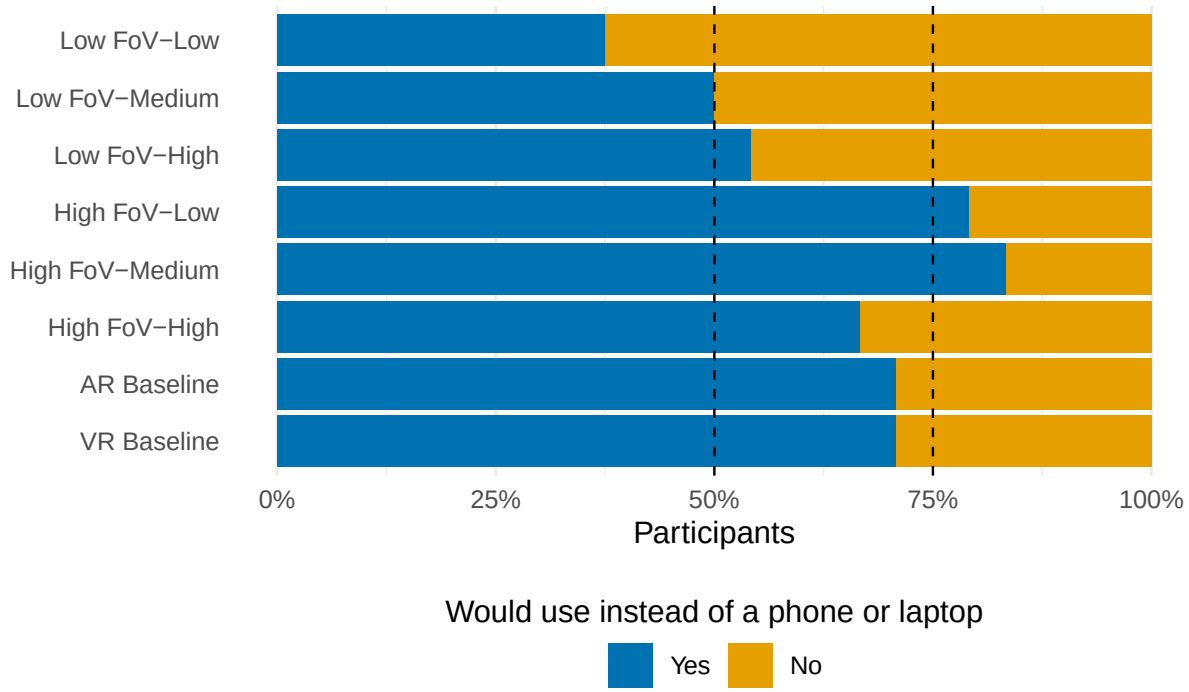


Figure 4.3: Percentage of participants ($N = 24$) who would use each condition instead of a phone or laptop to watch video on a train. Dashed lines mark 50% and 75%.

Compared with the VR Baseline. All six ARTV conditions produced significantly greater awareness of other passengers than the VR Baseline (all $p < .001$). On the Film-IEQ, UEQ-S, and UES-SF, the High FoV–Medium and High FoV–High conditions were the only ones not significantly different from the VR Baseline after correction. Both nonetheless scored lower on the Film-IEQ: by 7.9 points for Medium (uncorrected $p = .019$; Dunnett $p = .099$) and 4.9 points for High (uncorrected $p = .049$; Dunnett $p = .22$). Equivalence tests were inconclusive for both (Medium: 90% CI $[-13.24, -2.51]$, bounds ± 9.16, $p_{\text{TOST}} = .34$; High: 90% CI $[-8.89, -0.86]$, bounds ± 6.85, $p_{\text{TOST}} = .20$; table D.6). These conditions therefore cannot be said to match the VR Baseline; the data are consistent with a small-to-moderate shortfall ($g_z \approx -0.4$ to -0.5) that the study was not powered to detect after correction. At the subscale level, *Real-World Dissociation* remained lower than in the VR Baseline in every ARTV condition, as would be expected when the carriage remains visible. For awareness of the ornamentations, the High FoV–Medium and High FoV–High conditions did not differ from the fully virtual jungle, whereas the other four conditions were lower.

4.4.5 Condition Preferences

Preferences broadly followed the questionnaire results (fig. 4.4; table D.7). Rankings differed across conditions, $\chi^2(7) = 58.0$, $p < .001$, Kendall's $W = .35$. After Holm correction, High FoV–High, High FoV–Medium, and the VR Baseline were each ranked above the AR Baseline,

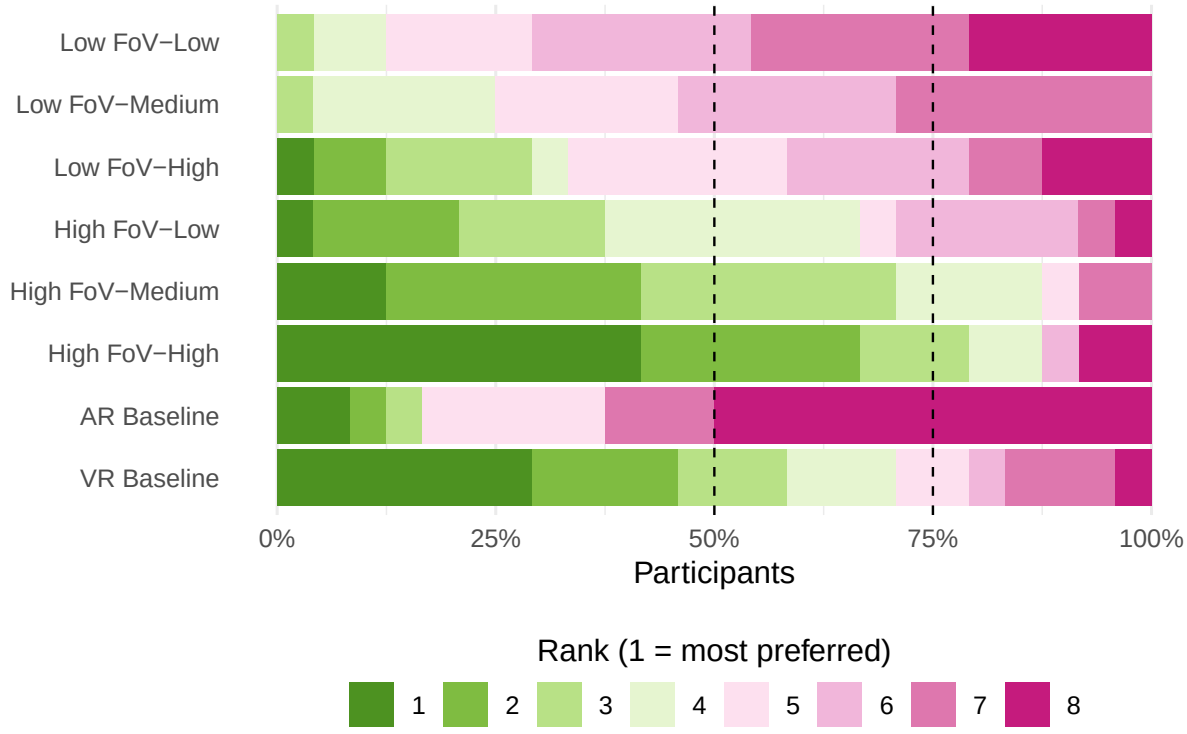


Figure 4.4: Distribution of participants' rankings of the eight conditions, from 1 (most preferred) to 8 (least preferred). Dashed lines mark 50% and 75% of participants.

and did not differ from one another. High FoV–High was also ranked above all three Low FoV conditions, and High FoV–Medium above Low FoV–Low and Low FoV–Medium (table D.8). Within each FoV level, mean ranks improved with ornamentation, but these differences were not significant. High FoV–High and High FoV–Medium each received 21 of 24 top-four placements, ahead of the VR Baseline (17) and High FoV–Low (16). High FoV–High was most often ranked first (10 participants), followed by the VR Baseline (7), but it was also ranked last by two participants, whereas High FoV–Medium was never ranked below seventh. The AR Baseline was ranked last by half of the participants (12).

Prospective use followed a different pattern (fig. 4.3). It differed across conditions, $Q(7) = 19.8$, $p = .006$, but no pairwise comparison survived Holm correction, so the following ordering is descriptive. High FoV–Medium (20 of 24, 83%) and High FoV–Low (19, 79%) were the conditions most often chosen as an alternative to a phone or laptop, followed by the AR and VR Baselines (17 each, 71%) and High FoV–High (16, 67%); the Low FoV conditions were chosen least often (9–13). Although most participants ranked the AR Baseline low, most would still use it on a train. This suggests that participants distinguished between the condition they preferred in the study and what they would use on a journey.

4.4.6 Qualitative Results

The themes below were developed inductively across the full interview dataset (section 4.3.2). Because each participant experienced every condition, responses often compared conditions directly; themes are therefore reported under the condition group or factor in which they were most clearly expressed.

Low FoV Conditions

Users found the limited FoV distracting ($n = 7$) Seven participants found it distracting that ornamentation was rendered only within the restricted rectangle (fig. 4.1c, e, and g). P11 said '*I felt kind of distracted by the square. I'm not going to lie.*'; P14 agreed: '*I like the idea of the low one, but I felt as if it was kind of distracting sometimes. When you moved, the trees would disappear*'.

No peripheral awareness of the ornamentation ($n = 7$) Seven participants noted that with the Low FoV much of the ornamentation was not visible while watching. P2 said '*When I was watching the TV, I didn't really see anything else. I wasn't really seeing the plants around it.*' P11 described the same problem: '*sometimes I wanted to see the plants, like, to get more, like, feeling I was in there, but it was like, I have to move my head and then I lose the sight of what I'm watching.*'

Too much contrast ($n = 4$) P1, P17, P19, and P20 described the borders between the video, the ornamentation, and the carriage as confusing and distracting. As P1 put it: '*I think the low field of view was just a bit distracting. The contrast between the train and the field of view [of the ornamentations]. And then there's three separation barriers, you know, for the movie and the field of view [within which the ornamentations are rendered] and the train*'.

Restricted FoV increased awareness of the train environment ($n = 6$) Six participants valued the restricted region because it left more of the carriage visible than the High FoV conditions. P2 explained: '*I have more of an understanding of who's around me and what's happening, which is super important to me if I'm on a train. But I still have all the benefits of being in VR, seeing my world augmented in such a way*'. P3 agreed: '*So I put the low FoV and the high modification first [in my ranking] because I felt like I could still see like if the door was opening or like a stop was coming up on the train or people were there ... but I could also still see all the cool flora around. So that was pretty cool.*'

High FoV Conditions

Increasing FoV increased immersion ($n = 4$) Four participants said the wider FoV improved their immersion. P2 explained: '*I felt a lot more immersed. Because even when I wasn't looking to the left, I could still see there were plants to the left.*'

Participants accepted reduced awareness when immersion gains were clear ($n = 4$) Four participants described the High FoV conditions as a worthwhile balance between stronger immersion and reduced, but in their view still adequate, awareness of the carriage. P3 explained their ranking of High FoV–High ornamentation: '*Then I put the high FoV, high modification, because it just looked really cool and it was probably worth losing a bit of awareness, because you could still see some of the foreign things*'. P12 said '*you feel in the scene, but you can still see something around you*', and P16 that the experience let them '*still keep an eye on like what's happening out in the world*' while feeling immersed.

High Ornamentation

Highest level of ornamentation distracting ($n = 5$) Five participants, including P2, P7, and P20, found the highest level distracting. P2 said '*It's a lot more distracting, and it took me a second to get into what was happening*'.

Reduced awareness of the environment ($n = 4$) Four participants said the ornamentation reduced their awareness of the carriage. As P19 put it, '*personally I wouldn't prefer this because it obscures the real world and I still want to be involved in the real world*'.

Awareness of the train beneficial ($n = 6$) Six participants valued remaining aware of the train while immersed. P23 commented:

> *High FoV high modification was top because it was the best balance of having a more immersive experience but also still staying aware of my surroundings. I think if I was on a train I might need to keep track of people coming and going and where the train is and so on. So I wouldn't necessarily want to be entirely immersed in the scene. So I think I went [for] the best in both worlds.*

Increased immersion ($n = 11$) Eleven participants described the High level as increasing their immersion. P12 said '*I could still see what's happening around me, but it felt more like you're in the scene. So I really liked that*', and P16 '*when you had a lot I felt like it was really kind of dragging you into it*'. P21 added: '*I'm not sure how, but it didn't distract me at all. I'd say it definitely enhanced it*'. P22 captured the intention of the condition: '*it is very aesthetically pleasing and it really gave more sense of what was in the video, what was being discussed, especially because it was more connected to the environment*'.

Texture change improved the experience ($n = 12$) Twelve participants singled out the retexturing of the carriage as a notable improvement. P3 said '*I think it made things a lot more worth it, it was very cool*', and P5 '*because the table was green and the chairs were too it genuinely was so cool*'. P12, P21, P23, and P24 commented on the seats: '*the seats, I think they were bluish, so that is very much a colour that you don't really see in the jungle ... So I think really changing that [texture] did make a difference*' and '*I didn't realise how much making the chairs blend in made a difference*'. P14, P16, and P21 described its effect on the experience: '*I felt it did add to the experience a lot*'; '*it looks like kind of like almost like a zombie apocalypse ... if this was a jungle and this train had derailed*'; and '*it's interesting how much it then added to it, just having the different sort of concordant colours there*'.

Occluding passengers could be beneficial ($n = 10$) Ten participants saw ornamentation partly hiding other passengers as a positive. P6 said '*If someone was next to me or around me and was moving vividly then it would be nice if there were some plants around so it would feel more immersed*', and P24 '*I don't want to see you. It's nice to have a leaf in front of [them]*'. P16 related this to their ADHD: '*I felt like ... especially with having ADHD, all of these, I felt like I saw everybody and I would watch them, whereas these, the plants kind of covered them up a bit. So that was a good thing*'.

Low Ornamentation

Low ornamentation afforded awareness ($n = 3$) P2, P4, and P23 valued the awareness of the train and other passengers that Low ornamentation left them. P23 explained: '*I'd maybe rather just have a slight modification to keep good awareness of everything, while just watching a movie with a little bit of ornamentation*'.

Unconvincing immersion ($n = 6$) Six participants felt little immersion with Low ornamentation. P14 said '*the low one was a bit too low, in my opinion ... what was the point in [it] being there?*', and P16 '*when it was just a little bit you're like what's going on? It's kind of like somebody that buys like one house plant and puts it in the living room*'.

Attitudes towards the VR Baseline

Lack of awareness concerning ($n = 6$) Six participants were concerned about their lack of awareness while watching in VR. P2 said '*I have no idea of what's happening around me, who's around me.*'

VR was most immersive ($n = 5$) Five participants described the VR scene as the most immersive. P10 summarised: '*The VR scene, that was more or less, you were wrapped in cotton*

wool . . . On the train one, you had your plants and that. But I would definitely say that the VR was more immersive'.

4.5 Discussion

The results support ARTV Ornamentation as a way of increasing media immersion while keeping the carriage in view, with two qualifications: the benefit depended on a wide FoV, and it came with reduced awareness of other passengers. With a wide FoV, Medium and High ornamentation raised immersion, user experience, and engagement above the floating-screen AR Baseline, and participants described being drawn into the content. These were also the conditions most often ranked highly and, in the case of High FoV–Medium, the most often chosen as an alternative to a phone or laptop, although differences in prospective use were not significant after correction.

4.5.1 FoV Moderated the Benefit of Ornamentation

The clearest result is the interaction between the two factors: ornamentation increased immersion, and reduced awareness, only when the FoV was wide. With the restricted FoV, adding ornamentation made no significant difference to either measure. Low FoV–High ornamentation still scored higher than the AR Baseline, so ornamentation was not without effect at the narrower FoV, but the gain was about half that of the equivalent High FoV condition and did not extend to the overall user experience or engagement scores. The interviews help explain this. Participants found the edge of the augmented region distracting, and noted that they had to turn away from the video to see the ornamentation at all.

This answers RQ3-1 (*Does device FoV affect the media immersion of a user consuming ARTV content?*): it does, chiefly by limiting how much ornamentation can contribute. Immersive ARTV of the kind studied here is therefore likely to be less effective on devices with narrow fields of view, such as the Xreal Air or HoloLens 2, and more effective on devices whose FoV approaches that of current VR headsets, such as passthrough headsets or future OST displays. Part of the Low FoV penalty, however, may reflect the visible rendering boundary, which physical OST devices do not show (section 4.5.5).

4.5.2 Comparison with VR

The High FoV–Medium and High FoV–High conditions were the only ARTV conditions that could not be distinguished from the VR Baseline on immersion, user experience, or engagement after correction for multiple comparisons. This should not be read as equivalence. Both conditions scored somewhat lower than VR on the Film-IEQ, the uncorrected differences were

significant, and equivalence tests were inconclusive. Real-world dissociation, in particular, remained lower in every ARTV condition, as would be expected when the carriage stays visible. Meanwhile, every ARTV condition preserved substantially more awareness of other passengers than VR, and six participants raised concerns about the lack of awareness in the VR Baseline.

This answers RQ2-1 (*How does ARTV Ornamentation compare to immersive VR in terms of enjoyment, engagement, and immersion?*): with a wide FoV and Medium or High ornamentation, ARTV came closer to VR than the other AR conditions did, and no difference could be detected after correction, but the study cannot establish that it matched VR. What it does show is that these conditions narrowed much of the gap between floating-screen AR and VR while retaining far greater awareness of the carriage.

4.5.3 Balancing Awareness against Media Immersion

Increasing ornamentation increased immersion, hedonic quality, and engagement, but reduced awareness of other passengers, and both effects were concentrated in the step from Low to Medium ornamentation at the wide FoV. The step from Medium to High, which retextured the carriage without adding models, produced no further significant change in immersion or awareness; only perceived aesthetics increased. High FoV–Low ornamentation, by contrast, did not differ significantly from the AR Baseline on the overall Film-IEQ, UEQ-S, or UES-SF scores, or on awareness of other passengers, which suggests that sparse ornamentation is insufficient to alter the experience meaningfully.

Preferences point the same way. Although four participants said the High level was worth the reduced awareness, and High FoV–High was most often ranked first, High FoV–Medium matched it in top-four rankings, was never ranked last, and was the condition most often chosen as an alternative to a phone or laptop (descriptively; no pairwise difference in prospective use was significant).

This answers RQ2-2 (*How does increasing the DoA affect media immersion, engagement, and environmental awareness, and how does this shape the trade-off between them?*): more ornamentation increased immersion and engagement at the cost of self-reported awareness of other passengers, with most of both changes occurring by the Medium level. High FoV–Medium therefore offers the best balance observed in this study: it achieved most of the immersion gain of the High level with comparable awareness, and it was the condition most often chosen for use on a train. Neither Medium nor High ornamentation, however, preserved awareness at the level of the AR Baseline.

4.5.4 Design Implications

Participants differed in how much ornamentation they were comfortable with: some prioritised immersion despite reduced visibility of the carriage, while others preferred minimal augmen-

tation. ARTV Ornamentation should therefore be adjustable. Apple's Vision Pro presents a control that sets how much of the real world is mixed with the virtual [5]; an analogous control for ARTV could adjust ornamentation intensity, letting passengers choose the level of environmental transformation that suits their comfort with reduced awareness while keeping the carriage in view.

Designs must also account for the display. Ornamentation that relies on extensive spatial augmentation is likely to be ineffective or distracting on narrow-FoV devices, and many current and near-term OST glasses remain limited in FoV, brightness, and occlusion. Wider-FoV VST devices such as the Apple Vision Pro and Quest 3 can support more spatially integrated ornamentation, but the present results suggest this will also bring greater reductions in awareness. ARTV Ornamentation cannot therefore be designed independently of the class of device.

More broadly, the study extends ARTV beyond overlays that present supplementary information alongside media. Augmentation that extends the theme of the content into the surrounding environment increased immersion, engagement, and hedonic quality relative to a floating screen. Such augmentation could therefore complement existing overlay approaches.

The same approach need not be limited to video. A passenger reading, on paper or on an AR panel, might benefit from augmentations themed to the book; a passenger simply passing time might benefit from calming augmentations unrelated to any content, such as a stream running down the aisle or a koi pond on the table.

4.5.5 Limitations

Lab setting and simulated AR

The study was conducted in a lab using AR simulated in VR, so results may differ in a real train. The passenger avatars had idle animations but did not leave their seats or move through the carriage, which does not capture the dynamics of a real journey; passengers moving around the carriage are likely to affect how ornamentation is experienced, as chapter 6 explores. The ornamentations were rendered with controlled opacity (section 4.3.1), which current additive OST displays cannot fully reproduce. The findings therefore motivate, rather than substitute for, evaluation on real devices in real vehicles.

Ornamentation levels combined density and retexturing

The three ornamentation levels did not vary along a single dimension: Low to Medium increased the number of models, while Medium to High kept the models and retextured the carriage. Effects of the Medium-to-High step, including the aesthetic benefit that twelve participants attributed to the texture change, cannot be attributed to density, and effects spanning Low to High combine both changes. Future work should vary density and retexturing independently.

Low FoV boundary

The Low FoV conditions restricted only the ornamentation, leaving a visible rectangular boundary that seven participants found distracting. A physical OST display would not produce this boundary, so part of the Low FoV penalty may be an artefact of the simulation.

Baseline screen sizes

The screen used in the AR Baseline and ARTV conditions subtended approximately $49.4^\circ \times 29.8^\circ$, larger than the VR Baseline screen (approximately $40.5^\circ \times 23.4^\circ$). The angular sizes were not matched. Because screen size is known to affect immersion (section 2.3.2), this difference would, if anything, have favoured the AR and ARTV conditions. The lower Film-IEQ scores of every ARTV condition relative to the VR Baseline therefore cannot be attributed to screen size, although the slightly larger AR screen may have narrowed the gap.

Clip effects

Each condition was paired with different clips across participants, but clip was not modelled in the analysis, so differences between clips form part of the condition effects.

Measurement

The Film-IEQ was administered without item 13, so its overall and Captivation scores are not directly comparable with published values, and its Comprehension subscale showed low internal consistency under the published key (table 4.2); no effects on Comprehension were found. Real-World Dissociation scoring differed from the published key by a constant (section 4.3.5), which does not affect comparisons but does affect comparison with published values. Environmental awareness was measured with a single self-report item; the study included no behavioural awareness task, and self-reported awareness may not correspond to what participants would actually notice. The interviews were analysed by a single coder, so the themes reflect one analyst's interpretation of the data. Finally, the modest sample limits sensitivity: after correction, differences smaller than about $d_z = 0.78$ could have gone undetected, which is why the comparisons with VR are reported as inconclusive rather than as equivalence.

Duration, content, and ornamentation type

Participants watched short clips, so the experience remained novel, and effects over a full programme were not measured. A single genre was used; chapter 5 examines genre. The ornamentations were passive environmental elements with limited ambient motion, chosen so that the conditions differed mainly in the amount and treatment of ornamentation rather than in behaviour or interactivity. More active or narrative-responsive augmentations, such as the swim-

ming turtle of Saeghe et al. [153], may affect engagement, distraction, and awareness differently, so the findings apply primarily to passive, environmentally themed ornamentation. The trade-off between immersion and awareness may generalise to other ornamentation types, but its specific form is likely to vary.

No comparison with other awareness techniques

ARTV Ornamentation was not compared with VR awareness techniques such as Reality Anchors [17] or the Apple Vision Pro's control over how much of the real world is visible [5]. Whether such techniques could make passthrough-style awareness unnecessary remains open.

Journey type

The study simulated a long, externally managed journey [17]: participants did not need to monitor stops, navigate, prepare to alight, or respond to time-critical events. Short, self-managed journeys place greater demands on awareness, and lower levels of ornamentation may be more appropriate there. Future work should examine how ornamentation should adapt across journey types.

4.6 Conclusions

This chapter presented the first evaluation of ARTV Ornamentation for passenger media consumption. Its conclusions for the sub-questions are:

- **RQ3-1 (FoV):** FoV affected immersion, user experience, engagement, and awareness, and it moderated the effect of ornamentation, which increased immersion only when the FoV was wide.
- **RQ2-1 (AR vs VR):** with a wide FoV and Medium or High ornamentation, ARTV could not be distinguished from immersive VR after correction, but equivalence could not be established and its immersion scores were somewhat lower. It preserved far greater awareness of other passengers than VR.
- **RQ2-2 (Degree of ornamentation):** increasing ornamentation raised immersion and engagement and lowered self-reported awareness of other passengers, with most of both changes occurring by the Medium level. Medium ornamentation with a wide FoV offered the best balance observed, though awareness remained below that of a floating screen.

For the thesis, these results support the premise that environment-integrated ornamentation can increase media immersion without fully occluding the passenger's surroundings, and they show that immersion and awareness moved in opposite directions as ornamentation increased.

The evidence is simulation-based: it does not show how ARTV Ornamentation would function during an actual journey, on a deployed AR display, or among passengers who move. Because a narrow FoV limited what ornamentation could achieve, subsequent chapters hold FoV constant at the wider setting.

Two questions follow. The first is whether ornamentation must be tailored to specific content, or whether genre-level designs are sufficient, particularly for content whose setting changes more than the nature clips used here; chapter 5 addresses this, along with transitions between ornamentation styles. The second is whether findings from simulation hold in a real passenger setting, which chapter 6 examines.

Chapter 5

Ornamentation Specificity and Adaptivity for Multi-Scene Narrative Media

5.1 Introduction

Chapter 4 showed that ARTV Ornamentation can increase media immersion, engagement, and user experience relative to a floating screen, and that these gains depended on a wide FoV and came with reduced awareness of other passengers. That evaluation used a single genre, nature documentary, whose on-screen setting stayed largely constant. It did not address how ornamentation should be designed for content whose setting changes from scene to scene. In particular, it remained unclear whether ornamentation must be tailored to specific shows, can be generalised across a genre, or should adapt as the on-screen scene changes.

For ARTV Ornamentation to be practical, these questions matter. If ornamentation must be authored for each programme, it is costly to produce; if genre-level designs suffice, reusable asset sets could serve many programmes. Narrative video also changes setting frequently, raising the question of whether ornamentation should remain static or change with the scene. This chapter uses the term *biome* for the distinct visual and environmental character of a scene's setting, for example a snowy exterior, a castle interior, or a desert landscape, rather than in its ecological sense.

Whereas chapter 4 varied the amount and extent of ornamentation, this chapter varies what the ornamentation represents and whether it changes with the content. It contributes to RQ3 (*What is the impact of device FoV, ornamentation density, and ornamentation type on the media immersion–environmental awareness trade-off?*) by examining ornamentation type, and to RQ2 (*How can ARTV with Ornamentation be used to balance media immersion and environmental awareness in passenger transportation?*) by examining whether passengers would choose to use ornamented ARTV. The study addresses the following sub-questions, numbered to follow those of chapter 4:

RQ3-2 **Ornamentation specificity:** Do show-specific ornamentations produce different immer-

sion, engagement, awareness, or preference from genre-based ornamentations?

RQ3-3 **Biome adaptation:** Does changing the ornamentation to match the biome of the current scene affect immersion, engagement, awareness, or preference, compared with static ornamentation?

RQ2-3 **Willingness to use:** Would passengers choose ornamented ARTV over a phone or laptop for watching video on a train, and does this differ from a floating screen?

5.2 Study: ARTV Objects and Biomes

Prior ARTV research, including chapter 4, has used nature documentaries, whose on-screen environment remains visually consistent, so ornamentation does not need to change to stay thematically coherent. Many narrative programmes instead move between interiors, landscapes, and locations. This study compared ornamentation strategies for such content using an extract from *Game of Thrones* [20], as a representative case of multi-scene narrative media, viewed in a simulated train carriage. The independent variable was the ornamentation strategy, with four levels (section 5.2.3); the dependent variables were immersion, user experience, engagement, awareness, and preference.

The four conditions do not form a factorial design. Instead, they were constructed to support two planned contrasts: genre-based against show-specific ornamentation, which isolates *specificity*; and show-specific against biome-specific ornamentation, which isolates *adaptivity*. The adaptive condition changed several things at once (the carriage textures, particle effects, and the timing of both), so the second contrast tests the combined adaptive treatment rather than any one of its components.

5.2.1 Design

A within-subjects design compared four conditions. As in chapter 4, AR was simulated in VR, an approach used in prior passenger XR research [17, 188, 110] and discussed in section 2.7. This allowed the ornamentation to be manipulated systematically while holding the environment constant, providing controlled evidence before ornamentation is examined in a real passenger setting in chapter 6. Participants were seated in a virtual train carriage, which they were asked to treat as their real travel environment, and the scenario was framed as the same long, externally managed journey as in chapter 4: participants did not need to monitor stops or prepare to alight. The other passengers were avatars from Adobe Mixamo [**?**] with idle animations, such as shifting their weight and moving their arms slightly; they did not stand up or move about the carriage.

As in chapter 4, retextured surfaces were rendered at 15% transparency, so that the original carriage textures remained slightly visible, and added objects at 5% transparency. These values

model a display with controlled augmentation opacity rather than the optical behaviour of current additive see-through displays (section 4.5.5). Objects were placed throughout the carriage, mostly in the participant's immediate surroundings and around the screen, following the layout principles of chapter 4: visibility around the central video, use of the surrounding space, and aesthetic integration. Ornamentation was rendered across the full headset FoV, corresponding to the wider setting of chapter 4.

5.2.2 Media Content

Participants watched the same continuous clip from *Game of Thrones*, Season 2 Episode 2 (8:19 to 17:32), in every condition. *Game of Thrones* is a widely viewed medieval fantasy drama with distinctive iconography, a genre not previously examined in ARTV research, and it moves frequently between visually distinct settings. The clip contained three biomes: a medieval interior (the Lannisters' castle at King's Landing), a snowy exterior (Jon Snow at Craster's Keep), and a sandy exterior (Daenerys Targaryen in the Red Waste). All scenes were dialogue-based with no fighting. A single continuous clip was used to give a coherent narrative and avoid breaks in presence [104] from switching between unrelated clips. Participants were required to have watched *Game of Thrones*, so that they could recognise the iconography in the show-specific condition. Condition order was counterbalanced using a Latin square.

5.2.3 Conditions

The four conditions are shown in fig. 5.1. In the results they are referred to as Baseline, Genre, Show, and Biome.

Baseline. Participants watched the clip on a floating screen in the unmodified carriage, which remained unchanged throughout. This replicates the AR Baseline of chapter 4, representing current floating-screen AR viewing, and served as the control condition. The same screen was used in all four conditions: $0.92\,\text{m} \times 0.53\,\text{m}$ at 1 m from the participant, subtending approximately $49.4^\circ \times 29.8^\circ$.

Genre-Based Augmentations. In the genre-based augmentation (GBA) condition, carriage surfaces were retextured to resemble a medieval interior (for example stone walls and wooden surfaces), and generic medieval objects were placed throughout the carriage, including books, coins, leather pouches and bags, tankards, and a dagger. The theme remained static throughout the clip.

Show-Specific Augmentations. The show-specific augmentation (SSA) condition used the same layout and static medieval retexturing as GBA, but the generic objects were replaced with



(a) Baseline: unmodified carriage, no ornamentation, and a floating screen.



(b) Genre-Based Augmentations: medieval retexturing and generic medieval objects, including coins, a leather bag, a dagger, and a book.

Figure 5.1: The participant's view in each condition. (a) Baseline; (b) Genre-Based Augmentations; (c) Show-Specific Augmentations; (d, e) Biome-Specific Augmentations during the snow and desert scenes.



(c) Show-Specific Augmentations: the same medieval retexturing as (b), with objects from the show, including a Lannister shield, a Stark wolf, and a Wildling shield and dagger.



(d) Biome-Specific Augmentations during a snow scene: snow, stone, and ice textures with the show-specific objects.

Figure 5.1: The participant's view in each condition (continued).



(e) Biome-Specific Augmentations during a desert scene: sand and sandstone textures with the show-specific objects.

Figure 5.1: The participant's view in each condition (continued).

objects from *Game of Thrones*, including house sigils (the Stark direwolf and the Lannister lion), shields, and weapons such as a Wildling shield and dagger. The theme remained static throughout the clip.

Biome-Specific Augmentations. The biome-specific augmentation (BSA) condition extended SSA by changing the carriage textures when the scene changed biome. During the King's Landing scene it was identical to SSA; at each change to the snowy or desert exterior, the carriage textures cross-faded to a matching set by modulating the transparency of layered texture sets. Particle effects were also added: falling snow during the scene in the North and blowing sand during the desert scene. These were introduced because both pilot participants noted a sense of dissonance when on-screen weather was not reflected in the surrounding environment. The particles could overlap the screen.

Rationale for the contrasts. The comparison between Genre and Show tests whether bespoke content is necessary. If generic ornamentation performs comparably to show-specific ornamentation, assets could be developed per genre rather than per programme; if show-specific iconography performs better, the additional production cost may be warranted. The comparison between Show and Biome tests whether static ornamentation is sufficient for content that moves between locations, or whether the carriage should change with the scene.

Objects and textures. Objects and textures were combined in every ornamented condition rather than varied independently. In piloting, placing a number of *Game of Thrones* objects

comparable to the medium plant density of chapter 4 produced visual clutter rather than greater immersion: unlike foliage, which blends together, distinct and semantically rich artefacts such as swords, shields, and crests each draw attention. An object-only condition dense enough to be comparable with the retextured conditions was therefore not viable for this content. In chapter 4, adding retexturing to the same objects did not significantly change immersion, although it increased aesthetic ratings and was singled out by twelve participants. The comparison here was therefore made between genre-based and show-specific objects within a retextured environment, as they would most likely be deployed.

5.2.4 Measures

After each condition, participants completed the Film-IEQ [149], the UEQ-S [157], and the UES-SF [130], as in chapter 4. Their subscales were:

- **Film-IEQ:** Captivation, Real-World Dissociation, Comprehension, and Transportation;
- **UEQ-S:** Pragmatic and Hedonic quality;
- **UES-SF:** Focused Attention, Perceived Usability, Aesthetics, and Reward.

As in chapter 4, the Film-IEQ was administered without item 13. UES-SF items were answered on a 7-point scale.

Participants also rated their agreement, on a 5-point scale from Strongly Agree to Strongly Disagree, with two statements: '*I was aware of other passengers whilst watching the video*' and '*I was aware of the augmentations*'. The first measured environmental awareness. The second was a manipulation check confirming that participants perceived the ornamentation, and is not interpreted as a measure of environmental awareness.

After all four conditions, participants ranked the conditions in order of preference and, for each condition, rated their agreement on a 5-point scale with the statement '*I would be interested in watching a video on a train using this condition instead of a phone or laptop*'. A short semi-structured interview followed, built around three questions:

1. Can you explain why you ranked the conditions in the order that you have?
2. Did you find the transition between scenes distracting?
3. What did you think of the sand and snow particle effects in the sand and snow scenes during the Biome-Specific Augmentations condition?

Follow-up questions explored participants' reasoning. Recordings were transcribed with WhisperX [14], and the transcripts were analysed with a single-coder inductive thematic analysis in QCAmap following Saldaña [154].

5.2.5 Participants

Twenty participants (14 men and 6 women), aged 19 to 60 ($M = 31$, $SD = 10.4$), were recruited through the university's participant mailing list. All had watched at least part of *Game of Thrones*: 13 had watched all of it, 2 most of it, and 2 several series; 2 had watched a few episodes and 1 a single episode. All 20 completed all four conditions; none withdrew or reported discomfort. Participants received a £10 Amazon voucher.

5.2.6 Procedure

Participants read an information sheet, consented to the study and to audio recording of the interview, and provided their age, gender, and familiarity with *Game of Thrones*. They sat on a static chair wearing a Meta Quest 3 headset connected to a computer running the Unity scene via Quest Link, with all conditions rendered at 72 Hz, the headset's default refresh rate. After each condition they completed the questionnaires and awareness items; after all four they completed the exit questionnaire and interview. The study took 60 minutes and was approved by the College of Science and Engineering Ethics Committee at the University of Glasgow (application 300230047).

5.2.7 Analysis

Scale and subscale scores were calculated following the published procedures for each questionnaire [149, 157, 130], with the Film-IEQ scored on its 23 administered items. Internal consistency was estimated within each condition (table 5.1). Scores were not normally distributed (Shapiro–Wilk), so each measure was analysed with the ART [193] followed by a repeated-measures ANOVA with participant as the error term. Pairwise comparisons used ART-C contrasts [43] with Tukey correction. Effect sizes for contrasts are reported as Hedges' g_z computed from the raw paired differences.

No *a priori* power analysis was conducted. With $N = 20$ and $\alpha = .05$, a single paired comparison had 80% power to detect $d_z \geq 0.66$. Where a comparison of awareness with the Baseline was not significant, TOST equivalence tests [87] were used, with bounds set at this smallest detectable effect ($d_z = 0.66$ multiplied by the standard deviation of the paired differences). A result within these bounds therefore excludes only differences the design could have detected, not all differences of practical interest. The awareness items were also fitted with cumulative link mixed models as a robustness check.

Rankings and interest ratings from the exit questionnaire were compared with Friedman tests (effect size Kendall's W), followed by Wilcoxon signed-rank tests with Holm correction.

Table 5.1: Internal consistency of each scale, estimated separately within each of the four conditions ($n = 20$ per estimate). Values are the median across conditions, with the range in brackets.

Scale	Items	Cronbach's α	McDonald's ω_t
Film-IEQ	23	.86 [.82, .89]	.89 [.86, .94]
Captivation	11	.84 [.76, .90]	.88 [.82, .93]
Real-World Dissociation	3	.86 [.74, .91]	.87 [.78, .91]
Comprehension	4	.23 [−.11, .51]	.57 [.45, .91]
Transportation	5	.81 [.74, .90]	.83 [.81, .90]
UEQ-S	8	.79 [.68, .85]	.85 [.68, .87]
Pragmatic	4	.94 [.88, .95]	.94 [.88, .96]
Hedonic	4	.73 [.69, .79]	.77 [.75, .82]
UES-SF	12	.95 [.90, .97]	.95 [.90, .97]
Focused Attention	3	.89 [.83, .93]	.89 [.85, .93]
Perceived Usability	3	.87 [.52, .95]	.89 [.65, .95]
Aesthetics	3	.91 [.86, .93]	.92 [.88, .93]
Reward	3	.93 [.91, .94]	.93 [.91, .94]

As in chapter 4, the low α for Comprehension reflects the published key, which does not reverse item 12 although it is worded in the opposite direction to the other three items; with item 12 reversed, the median α was .47 [.25, .88]. The published key was retained for all reported analyses.

5.3 Results

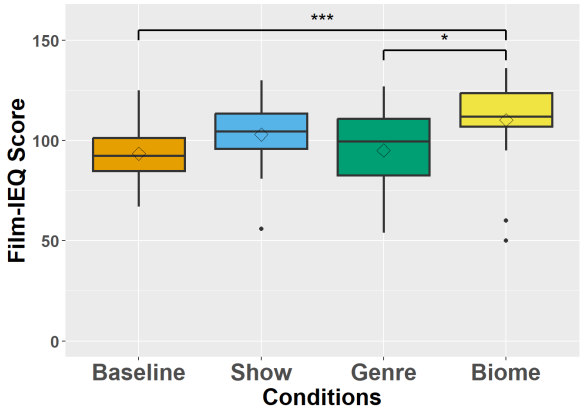
Table 5.2 summarises the omnibus tests and the significant contrasts, and fig. 5.2 shows the overall scores. Complete results, including descriptive statistics and all contrasts, are in appendix E. The raw data are available at the Zenodo repository [31].

5.3.1 Comparisons with the Baseline

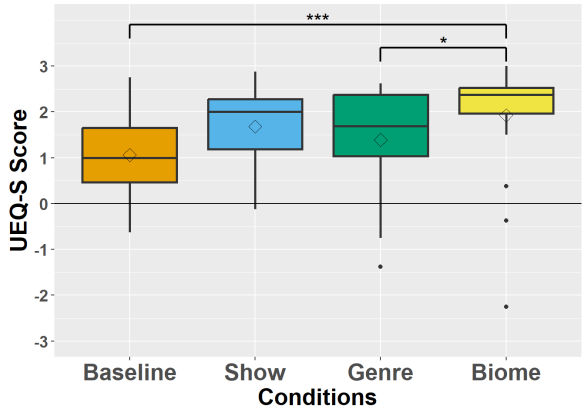
Biome scored significantly higher than the Baseline on the overall Film-IEQ, UEQ-S, and UES-SF, and on the Captivation, Transportation, Hedonic, Focused Attention, Aesthetics, and Reward subscales. Show scored higher than the Baseline on Transportation, the overall UES-SF, Aesthetics, and Reward. Genre scored higher than the Baseline only on Aesthetics; its difference on Reward was marginal ($p = .054$). All three ornamented conditions increased awareness of the augmentations, indicating that participants perceived the ornamentation.

5.3.2 Genre-Based and Show-Specific Augmentations

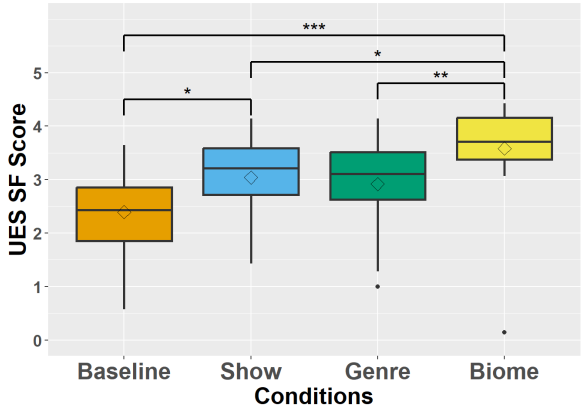
No measure differed significantly between Genre and Show (all $p > .18$). Using objects and iconography from the show did not produce a detectable improvement over generic medieval objects, although the absence of a significant difference does not establish that the two were equivalent.



(a) Film-IEQ



(b) UEQ-S



(c) UES-SF

Figure 5.2: Overall (a) Film-IEQ, (b) UEQ-S, and (c) UES-SF scores by condition ($N = 20$). Diamonds mark the mean, the central line the median, and the box the interquartile range. Asterisks mark significant ART-C contrasts: * $p < .05$, ** $p < .01$, *** $p < .001$.

Table 5.2: Omnibus ART ANOVAs ($df = 3, 57$) and ART-C contrasts (Tukey-corrected). Contrasts are listed where $p < .05$, and where $p < .06$ (marked †, marginal). g_z is Hedges' paired effect size from the raw scores; negative values indicate that the first-named condition scored lower. BL: Baseline; G: Genre; S: Show; B: Biome.

Measure		***F***	***p***	η_p^2	**Contrast**	***p***	g_z
Film-IEQ	Overall	6.28	<.001	.25	BL – B	<.001	−0.58
					G – B	.011	−0.73
	Captivation	3.84	.014	.17	BL – B	.030	−0.29
					G – B	.023	−0.50
	Real-World Dissociation	1.18	.32	.06			
	Comprehension	0.95	.42	.05			
	Transportation	11.89	<.001	.38	BL – S	.006	−0.76
					BL – B	<.001	−1.16
					G – B	.005	−0.89
UEQ-S	Overall	6.40	<.001	.25	BL – B	<.001	−0.49
					G – B†	.051	−0.65
	Pragmatic	1.82	.15	.09			
	Hedonic	11.65	<.001	.38	BL – B	<.001	−0.79
					S – B	.006	−0.61
					G – B	.002	−0.83
UES-SF	Overall	11.96	<.001	.39	BL – S	.038	−0.50
					BL – B	<.001	−0.74
					S – B	.012	−0.57
					G – B	.003	−1.00
	Focused Attention	12.80	<.001	.40	BL – B	<.001	−0.82
					S – B	.002	−1.02
					G – B	<.001	−1.11
	Perceived Usability	2.52	.067	.12			
	Aesthetics	12.39	<.001	.39	BL – S	.004	−0.78
					BL – G	.011	−0.58
					BL – B	<.001	−0.98
					G – B	.032	−0.76
	Reward	11.27	<.001	.37	BL – S	.031	−0.52
					BL – G†	.054	−0.39
					BL – B	<.001	−0.78
					S – B	.022	−0.68
					G – B	.012	−0.69
Awareness of other passengers		4.34	.008	.19	BL – B	.029	0.51
					G – B	.019	0.77
Awareness of augmentations (manipulation check)		14.09	<.001	.43	BL – S	.001	−0.86
					BL – G	<.001	−1.02
					BL – B	<.001	−1.24

5.3.3 Biome-Specific Augmentations

Biome scored significantly higher than Genre on the overall Film-IEQ, Captivation, and Transportation; Hedonic quality; and the overall UES-SF, Focused Attention, Aesthetics, and Reward. Its advantage on the overall UEQ-S was marginal ($p = .051$), and it produced lower awareness of other passengers. Biome also scored higher than Show on Hedonic quality, the overall UES-SF, Focused Attention, and Reward. Changing the carriage with the scene therefore improved several aspects of the experience beyond static ornamentation, although, as noted in section 5.2, this effect combines the texture changes with the particle effects.

5.3.4 Environmental Awareness

Self-reported awareness of other passengers was significantly lower in Biome than in the Baseline ($g_z = 0.51$) and than in Genre ($g_z = 0.77$). Neither Genre nor Show differed significantly from the Baseline.

Equivalence tests distinguished between these two non-significant comparisons. Genre was equivalent to the Baseline within bounds of ± 0.85 scale points ($t(19) = -2.61$, $p = .009$; raw difference 0.10, 90% CI $[-0.40, 0.60]$, $g_z = 0.07$): any difference in awareness between the two was smaller than this design could detect. The Show comparison was inconclusive, being neither significantly different from the Baseline nor equivalent to it ($p = .144$; $p_{\text{TOST}} = .084$). Cumulative link mixed models gave the same pattern. Biome, which differed significantly from the Baseline, moved the experience further towards immersion at the cost of reduced awareness of the surrounding passengers.

5.3.5 Condition Preferences

Rankings differed across conditions, $\chi^2(3) = 27.5$, $p < .001$, Kendall's $W = .46$ (fig. 5.4). Biome was ranked first by 15 participants and in the top two by 18 (mean rank 1.40), and was ranked above every other condition (Holm-corrected: Baseline $p = .003$, Genre $p < .001$, Show $p = .045$). Show (mean rank 2.35) was ranked above the Baseline (3.50; $p = .011$). Genre (2.75) did not differ significantly from the Baseline ($p = .121$) or from Show ($p = .274$). The one participant who ranked Biome last disliked the concept of ornamentation as a whole (section 5.3.6).

Interest in using each condition instead of a phone or laptop also differed, $\chi^2(3) = 19.4$, $p < .001$, $W = .32$ (fig. 5.3). Most participants agreed or strongly agreed for every condition: Biome 19 of 20 (95%, of whom 15 strongly agreed), Show 16 (80%), the Baseline 13 (65%, of whom 5 strongly agreed), and Genre 12 (60%). Interest in Biome was significantly higher than in Genre ($p = .005$) and Show ($p = .026$); its difference from the Baseline was marginal ($p = .053$). No other pair differed.

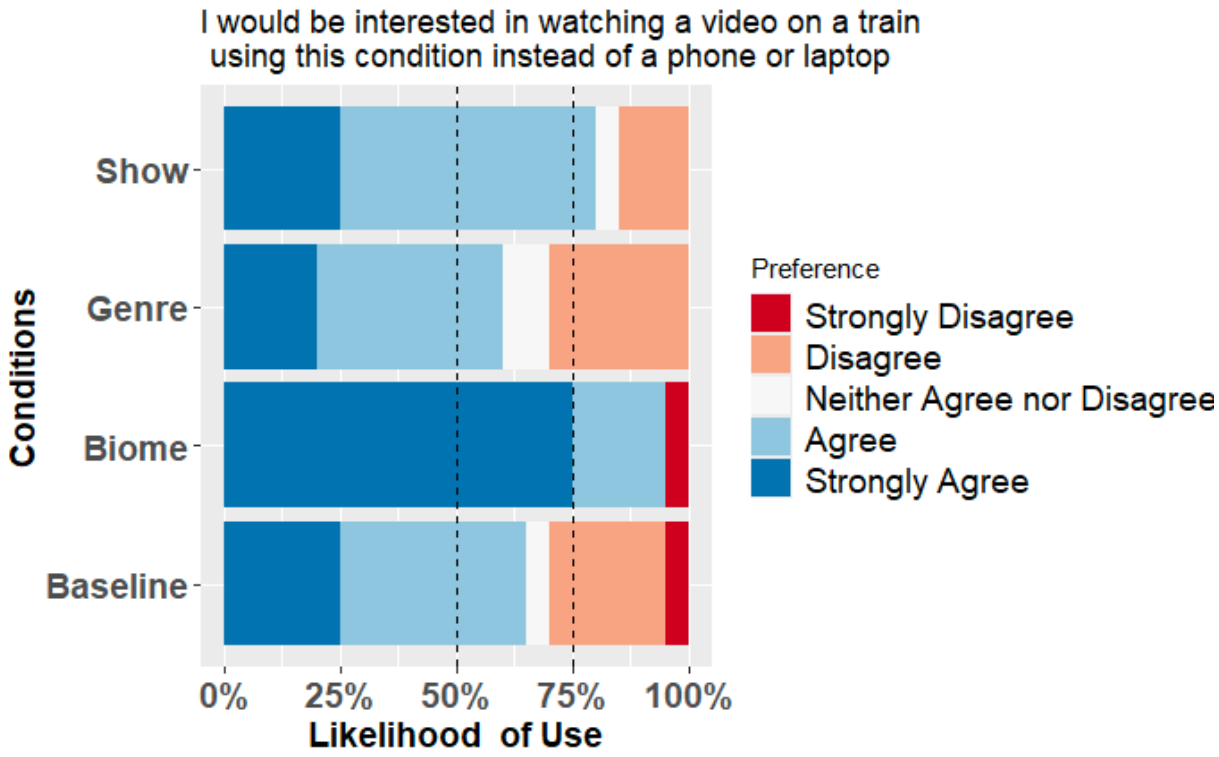


Figure 5.3: Agreement with the statement 'I would be interested in watching a video on a train using this condition instead of a phone or laptop' ($N = 20$). Dotted lines mark 50% and 75%.

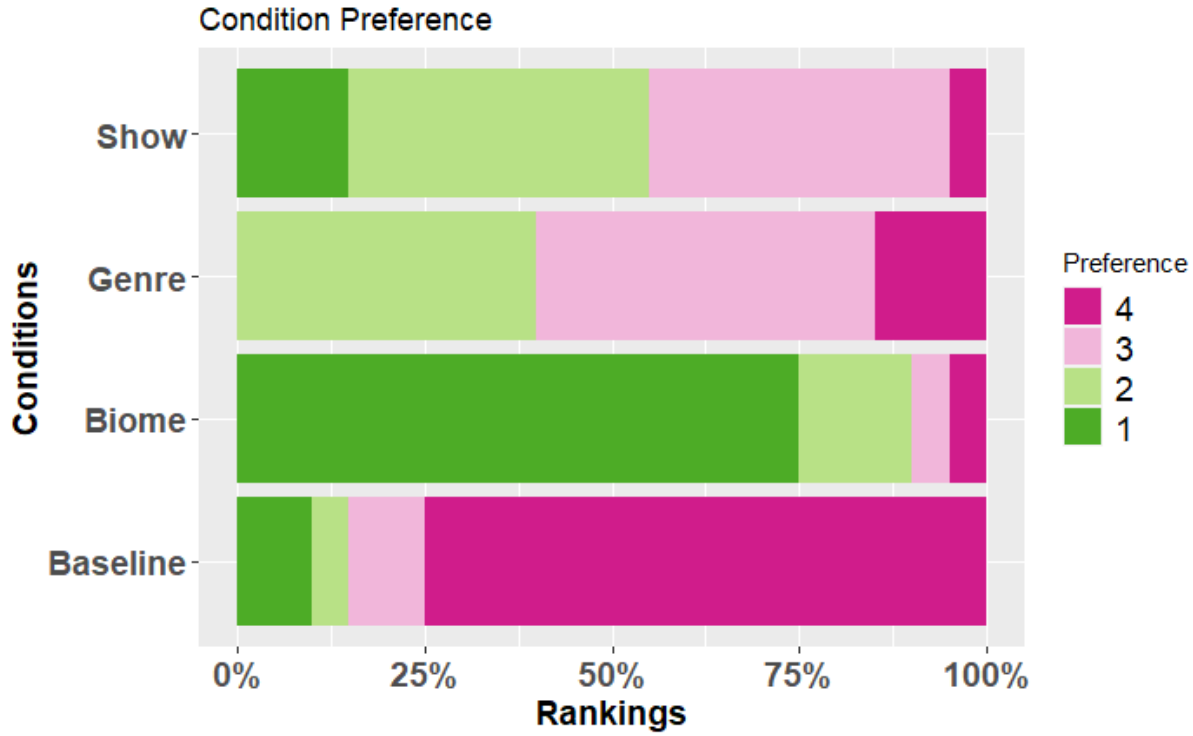


Figure 5.4: Participants' rankings of the four conditions, from 1 (most preferred) to 4 (least preferred) ($N = 20$).

5.3.6 Qualitative Results

The themes below were developed from the post-study interviews. Participant quotations have been lightly edited for clarity and brevity.

Retextured environments enhance immersion through environmental alignment ($n = 13$) Participants reported that retexturing the surroundings to match the on-screen content increased immersion, and described Biome as more engaging because it followed the scene. P16 explained: '*I liked that, when it was snowing, I kind of always imagined feeling cold and when it was in the desert, the sort of the heat rising. Yeah, almost got a smell. It really evoked the senses and I liked that.*' For these participants, immersion depended not only on the presence of ornamentation but on its coherence with the content.

Ornamentation competes for attention ($n = 11$) Participants also reported that ornamentation competed with the video for visual attention. Show-specific objects and environmental effects were frequently described as distracting. P2 noted that '*with objects laying around, every once in a while I was just tempted to see what was there, which then takes the focus away from the screen*', and P16 that '*[it's] maybe a bit more distracting if there are things that are particular because you're like, oh, why is that there? What's that? It distracts your brain.*' Richer ornamentation can therefore add to the experience while also drawing attention away from the content.

Show-specific ornamentation adds thematic engagement but risks incongruity ($n = 9$) Some participants found show-specific objects more engaging because they were recognisable. P12 explained: '*I like the Game of Thrones objects more than the medieval objects because that was something you can relate with the show and these are like iconic things that are there.*' Others found the same objects out of place when they did not fit the current scene. P20 said: '*in the first scene, where it was at King's Landing, there were Stark banners which were incongruous to the scene itself.*'

Generic and show-specific ornamentation involve a trade-off ($n = 13$) Preferences between generic and show-specific objects were divided. Some participants preferred generic objects because they suited every scene; as P4 noted, '*the generic medieval sort of fits all three scenes whereas the Game of Thrones things only really fit mostly the second scene.*' Others preferred show-specific objects for their relevance. P5 explained: '*The generic objects after a while just faded away. I didn't really care about them. But the Game of Thrones specific objects were much more interesting and it was much more in theme.*'

Environmental effects add realism but need careful placement ($n = 10$) Participants frequently described the snow and sand particles as making the experience feel more alive; P18 said '*those particle effects made it feel more real overall for me.*' Several, however, found particles that overlapped the screen distracting, and suggested keeping them in the background or away from the display.

Alternative perspective: aversion to environmental augmentation ($n = 1$) One participant (P19) stood apart from the rest. P19's scores for Biome were extreme outliers on the Film-IEQ, UEQ-S, and UES-SF, and their scores for Genre and Show were also outliers, though less pronounced. These scores were retained in all analyses. Excluding P19 reversed no contrast and strengthened several; the effect of condition on awareness of other passengers increased from $F(3,57) = 4.34$, $p = .008$, to $F(3,54) = 7.70$, $p < .001$. Without P19, the reduction in awareness under Show relative to the Baseline also became significant (Tukey-corrected $p = .038$), so the inconclusive Show result above depends on this participant's inclusion. Retaining P19 was the more conservative choice.

P19 found all ornamented conditions distracting and preferred conventional viewing: '*I just didn't like the moving things. It was distracting and fake and not for me*', and '*I just like the way television is, and I don't think [ornamentation] adds anything else.*' P19 had no prior experience with AR or VR headsets and described a general aversion to 3D cinema, particularly content that felt intrusive or moved towards them. This is a single participant's account and is not offered as evidence of a general disposition, but it indicates that some viewers may not want any augmentation. For most participants, the balance between immersion and awareness was a matter of how much ornamentation to use; for P19, any ornamentation was unwelcome. This suggests, tentatively, that a viewer's disposition towards augmented media may itself shape where that balance is struck, so that an acceptable configuration may vary by viewer as well as by context and device.

5.4 Discussion

The results show that ARTV Ornamentation can be applied beyond nature documentaries to narrative content with several settings. They also show that what the ornamentation represents, and whether it changes with the content, affect the balance between immersion and awareness.

5.4.1 RQ3-2: Should Ornamentation Be Show-Specific, or Is Genre Sufficient?

No measure differed significantly between Genre and Show. Both improved on the Baseline for aesthetic appeal, but Show's gains over the Baseline were broader, extending to Transportation,

overall engagement, and Reward, whereas Genre's were limited to Aesthetics. Show was also ranked above the Baseline, while Genre was not. The answer to RQ3-2 is therefore that show-specific objects produced no detectable advantage over generic ones in direct comparison, but the evidence that genre-level ornamentation alone improves on a floating screen is limited to its aesthetic appeal.

This result should be read in light of the implementation. The show-specific objects were held constant across scenes, so participants who knew the show saw iconography from one house during scenes centred on another, and several found this incongruous. In terms of presence, such mismatches are failures of coherence, the property of a virtual environment on which the plausibility illusion depends [162], and which Skarbez et al. [163] add as a dimension of their revised reality–virtuality taxonomy. Generic objects, being plausible in any of the scenes, did not raise this problem. Show-specific ornamentation may therefore need to change with the narrative as well as the setting, which would increase production effort (for example, creating and time-stamping assets). Programme makers, who know the storyline, may be best placed to author such ornamentation; this is a design implication rather than a finding of the study. Whether scene-matched show-specific objects would outperform generic ones remains to be tested.

If generic and show-specific ornamentation do produce similar experiences, reusable sets of models and textures could cover many programmes within a genre, without the cost of curating assets for each one. The present results are consistent with this possibility but do not establish it.

5.4.2 RQ3-3: Should Ornamentation Adapt to Changes in Scene Biome?

Biome scored higher than Genre on immersion (Film-IEQ) and engagement (UES-SF), and on Hedonic quality; its advantage on the overall UEQ-S was marginal. It also scored higher than Show on Hedonic quality, overall engagement, Focused Attention, and Reward. Participants ranked Biome above every other condition, and 13 described scene-matched retexturing as increasing their immersion. The answer to RQ3-3 is that changing the ornamentation with the scene improved immersion, engagement, and preference relative to static ornamentation. Because the adaptive condition combined texture changes, particle effects, and the timing of both, these gains are those of this combined treatment and cannot be attributed to biome matching, or to any one component, alone.

The gains were not without cost: Biome was the only condition to reduce self-reported awareness of other passengers significantly. Scene-matched augmentation can therefore move the experience further towards immersion, and further along the trade-off that chapter 7 develops as the PIAT.

Participants' views on the particle effects differed. Many appreciated them, but several found particles that overlapped the screen distracting. Whether particles should cover the screen may

therefore be best left to the viewer. P18 suggested that particles could also appear outside the train windows; tools such as Meta's Depth API [114], which allow virtual content to be occluded by real objects, could make this possible.

5.4.3 RQ2-3: Would Passengers Choose Ornamented ARTV?

Most participants would consider using any of the conditions instead of a phone or laptop, including 65% for the floating-screen Baseline. Biome attracted the strongest interest (95% agreement, 75% strong agreement) and was the preferred condition, but its interest ratings differed only marginally from the Baseline after correction. The answer to RQ2-3 is therefore that passengers were open to AR video on a train in general, and preferred biome-adaptive ornamentation among the options offered, but the evidence that ornamentation increases willingness to use AR beyond a floating screen is marginal. Interest ratings also reflect stated intentions after a short laboratory session, not behaviour on real journeys.

Participants' comments nonetheless show a desire to bring the setting of the content into their own surroundings, and ornamentation that keeps the carriage visible allows this to be done while retaining some view of the real environment. Awareness of the environment is a key factor in passengers' comfort with, and the perceived acceptability of, XR use [16, 188].

5.4.4 Applicability across Genres

Together, chapter 4 and this chapter suggest that both nature documentaries and narrative drama can benefit from ARTV Ornamentation. Other dramas, which are similarly narrative-driven and move between locations, may benefit in the same way. Genres with different characteristics may not. Sitcoms, for example, tend to use a small number of sets, which limits the scope for the scene-matched changes of the Biome condition; whether watching *Seinfeld* while surrounded by objects from Jerry's apartment would increase enjoyment and immersion is an open question. Some genres do not prioritise visual immersion at all. Modleski [123] argues that soap operas were constructed so that viewers could follow the plot while occupied with other tasks, relying more on sound than image, and streaming services have similarly been reported to favour content that a distracted, 'second screen' viewer can follow [72]. For such content, increased immersion may be neither achievable nor wanted. Nor is ornamentation appropriate for all content: for news reporting from an active war zone, for example, greater immersion could heighten distress. Future work should examine where across the range of genres ornamentation helps, and where it hinders.

5.4.5 Design Guidelines

Within the bounds described in section 5.2.7, awareness of other passengers under Genre was equivalent to the Baseline, which is closest to how passengers currently watch video in AR. Genre's gains, however, were limited to aesthetic appeal. The comparison for Show was inconclusive, and Biome produced both the largest gains and a significant reduction in awareness. For those designing ARTV Ornamentation, the findings suggest the following:

Guideline 1: Designers of ARTV content should consider changing ornamentation to match the on-screen setting, which in this study gave the most immersive and most preferred experience. Because this also reduced awareness of other passengers, scene-matched ornamentation should be offered with a way to reduce or disable it.

Guideline 2: Designers should weigh the development cost of show-specific ornamentation against its benefit, which was not detectable here when the objects were held constant across scenes. If show-specific objects are used, they should match the current scene to avoid incongruity.

Guideline 3: Designers may consider particle effects that match the setting (for example, blowing sand in a desert scene), which participants described as adding realism, but should let viewers choose whether particles overlap the screen or turn them off. The contribution of particles was not isolated from that of the texture changes in this study.

5.4.6 Limitations

Single programme and simulated AR

The evaluation used a single programme, *Game of Thrones*, whose medieval-fantasy setting and strong variation in biome suited the present comparisons but may not represent other narrative genres. As in chapter 4, AR was simulated in VR, with ornamentation at controlled opacity and passenger avatars that did not move about the carriage, which limits ecological validity.

Combined design changes

The conditions combined several changes rather than isolating them. Objects and retexturing varied together, and the Biome condition added particle effects and scene-synchronised change. The relative contribution of each component cannot be separated here.

Repeated clip and sample

Participants watched the same clip in all four conditions. Order was counterbalanced, but repeated exposure may have affected engagement independently of the ornamentation. The sample of 20 was predominantly male (14 of 20), and all participants had watched at least part of the

show, though three had seen only a few episodes and may not have recognised all of the iconography. The design could detect only relatively large effects ($d_z \geq 0.66$ for a single comparison), so smaller differences, including between Genre and Show, may have gone undetected.

Measurement

Environmental awareness was measured with a single self-report item, with no behavioural measure. The equivalence bounds for awareness were set at the smallest effect the design could detect, and so do not exclude smaller differences. The Film-IEQ was administered without item 13, and its Comprehension subscale showed low internal consistency under the published key (table 5.1); no effects on Comprehension were found. The inconclusive awareness result for Show depended on the inclusion of one participant (P19). Interest in future use was self-reported after a short laboratory session. The interviews were analysed by a single coder, so the themes reflect one analyst's interpretation.

These limitations, together with directions for future work, are considered further in chapter 7.

5.5 Conclusion

This chapter extended ARTV Ornamentation from a single-setting documentary to narrative content that moves between settings. Its conclusions for the sub-questions are:

- **RQ3-2 (Specificity):** show-specific objects produced no detectable advantage over generic genre-based objects when held constant across scenes. Genre-based ornamentation improved on a floating screen only in aesthetic appeal, and did not reduce awareness of other passengers.
- **RQ3-3 (Biome adaptation):** changing the carriage to match the on-screen setting increased immersion, engagement, and preference relative to static ornamentation, but was the only condition to reduce self-reported awareness of other passengers.
- **RQ2-3 (Willingness to use):** most participants would consider watching video in AR on a train, and biome-adaptive ornamentation was the preferred option, but its advantage over a floating screen in stated interest was marginal.

For RQ3, these results show that the balance between immersion and awareness depends not only on the amount and extent of ornamentation examined in chapter 4, but also on what the ornamentation represents and on whether it changes as the content changes. Ornamentation that follows the content more closely moved the experience further towards immersion and away from awareness. For RQ2, passengers valued this, but some, like P19, did not want augmentation at all, which supports allowing ornamentation to be adjusted or turned off. The results

also support exploring reusable libraries of genre- and setting-level assets as an alternative to programme-specific authoring, although they do not show that generic and show-specific ornamentation are equivalent.

The evidence remains simulation-based. Chapter 6 examines ornamentation in a real passenger setting, where passengers move and must watch media while also carrying out a real-world awareness task.

Chapter 6

ARTV Ornamentations In Situ with Realistic Awareness Task

6.1 Introduction

Chapters 4 and 5 showed that ARTV Ornamentation can raise immersion and engagement above a floating screen, and examined how it can be matched to content that moves between settings, but did so entirely in VR simulation, in which both the augmentation and the carriage around it were virtual. Simulation affords control, but it cannot show what happens when augmentation must be registered to real surfaces, shared with real people, and experienced alongside the noise and unpredictability of a public space. It is therefore not known whether the earlier findings hold once the carriage is real.

This chapter deploys ARTV Ornamentation using pass-through AR in a preserved subway carriage open to the public, and compares it with the simulated approach of the earlier chapters and with two fully immersive VR viewing formats. Participants also performed a stop-announcement task, adding an explicit demand on journey awareness that the earlier studies lacked. The study therefore examines whether the immersion–awareness relationship observed in simulation holds under more realistic attentional demands, and how pass-through AR compares with the simulated and fully immersive alternatives in user experience, engagement, awareness, task performance, and preference.

Conducting the study on an operating service was pursued but not possible. Access was requested from both ScotRail, which operates the national rail network in Scotland, and Strathclyde Partnership for Transport, which operates the Glasgow Subway, and both declined permission; SPT cited its focus on the Subway modernisation programme, including the introduction of automated platform screen doors. A preserved museum carriage was therefore used as the closest available alternative. It provides the physical dimensions, seating arrangement, and public presence of a transit environment, and so a more ecologically valid setting than the earlier simulations, but it is stationary and outside live operation. The consequences of this for interpreting

the findings are considered in section 6.4.3.

This chapter answers **RQ4** from the thesis: *How does the passenger immersion–awareness trade-off apply when moving from simulated AR-in-VR to AR experiences situated in the real world?* To answer this, the chapter considers two sub-research questions:

- **RQ4-1**: To what extent do findings from prior ARTV studies conducted within VR-simulated transit environments, transfer to ARTV experiences enacted within a real physical transit setting using pass-through AR?
- **RQ4-2**: How do user experiences of ARTV in a transit context compare to existing commercially available XR media-viewing approaches for in-transit entertainment?

6.2 Methodology

6.2.1 Study Design

A within-subjects study was conducted in which participants experienced six conditions spanning three modes of reality representation: pass-through AR in the physical carriage, simulated AR in VR, and fully immersive VR media viewing. The study was approved by the College of Science and Engineering Ethics Committee at the University of Glasgow (application 300230236).

The conditions supported two comparisons. To examine the validity of VR-based simulation for ARTV research (RQ4-1), the four AR and simulated-AR conditions formed an embedded 2×2 design crossing Medium (pass-through AR, simulated AR) with Ornamentation (baseline, ornamented), to determine which findings from simulation transfer to a physical carriage. To situate ARTV within existing immersive media practice (RQ4-2), these conditions were compared with two fully immersive VR formats common in commercial XR applications: a virtual cinema and a themed panoramic environment. These represent broad categories of immersive viewing and were not intended to match the production fidelity of specific commercial applications; they provide a reference for evaluating the relative strengths and limitations of ARTV in transit contexts.

In every condition, participants watched the same set of video clips while performing a stop-announcement awareness task. Outcomes comprised user experience and engagement, self-reported awareness of other passengers and of the augmentations, awareness-task performance, condition preference, usage intention, and semi-structured interview responses.

6.2.2 Conditions

The six experimental conditions were designed to systematically vary two primary factors: (1) the mode of reality representation, comprising pass-through AR, simulated AR-in-VR, and fully

immersive VR media viewing; and (2) the presence or absence of ARTV Ornamentation.

The six conditions were as follows (see figs. 6.1 and 6.2 for participant views):

AR Baseline. Participants were seated in the physical train carriage and watched clips on a floating virtual screen using pass-through AR. The surrounding environment remained unmodified, allowing pass-through visibility of the real-world carriage and any bystanders present. See fig. 6.3 for the unmodified train carriage.

ARTV Ornamented. Identical to AR Baseline, except that the physical carriage environment was augmented using ARTV Ornamentation. The carriage interior was visually transformed to match the biome of the current clip, and virtual objects specific to the show being watched and matched to the current clip (e.g. vegetation, props) were placed throughout the environment. Augmentations were positioned within the participant's immediate surroundings and around the periphery of the screen, prioritising visibility of the central video content while using the surrounding space for aesthetic augmentation. This arrangement and density were informed by the best-performing configurations identified in chapters 4 and 5, which kept the central display clear while integrating augmentation with the surrounding environment.

Simulated AR Baseline. Participants were seated in a nearby private room within the museum. VR was used to simulate an AR experience within a virtual reconstruction of the train carriage, following the approach used in chapters 4 and 5 and in prior transit XR research [17, 188, 110]. Video content was presented on a floating screen positioned in front of the participant, mirroring AR Baseline. The virtual carriage interior remained visually unchanged throughout the condition.

Simulated ARTV Ornamented. Identical to Simulated AR Baseline, but with the addition of ARTV Ornamentation corresponding to that used in ARTV Ornamented. Virtual textures and objects were applied throughout the simulated carriage to match the biome of the current clip, following the same layout principles used in the physically situated ARTV Ornamented condition.

VR Cinema. Participants were seated in the physical train carriage and watched clips within a fully immersive virtual reality cinema, without visual access to the real world. This condition emulated commercially available VR media-viewing experiences, such as watching videos through applications like Bigscreen VR [22].[1]

[1] `https://www.bigscreenvr.com/`

VR Videosphere. Participants were seated in the physical train carriage and watched the same video content on a floating panel while surrounded by a static 360° panoramic image matching the biome of the current clip. This condition approximates the contribution of environmental surroundings to immersion, although it also removes visual access to the carriage, so it differs from the AR conditions in occlusion as well as in surround. A direct comparison between 2D video and true 360° video would introduce substantial confounds related to content framing, camera control, and narrative structure; this controlled proxy therefore allowed us to evaluate the role of the environmental context whilst maintaining content consistency across conditions.

Exclusion of Smartphone Viewing

A direct smartphone viewing condition was not included because the study was designed to compare XR media-viewing configurations, rather than to test whether XR viewing is more immersive than handheld viewing. Existing work suggests that display size and field of view influence presence and immersion, so a handheld baseline would likely have introduced a predictable but less theoretically informative contrast [64, 148, 185]. Instead, perceived suitability relative to everyday device use was captured through the post-study usage-intent questionnaire, in which participants indicated whether they would use each condition in place of a phone or laptop. This provided a comparison grounded in adoption and preference rather than expected perceptual differences.

6.2.3 Apparatus and Environment

All conditions were conducted using a Meta Quest 3 headset rendering at 72 Hz and used Sony MDR-ZX310 on-ear headphones. Closed headphones were used in preference to the headset's built-in audio to minimise interference from ambient noise in the carriage and to ensure consistent audio presentation across participants.

For pass-through AR conditions, participants viewed the real-world carriage through the headset cameras. In the ARTV Ornamented condition, Meta's Mixed Reality Utility Kit (MRUK) [113][2] was used to support scene understanding and environmental augmentation, enabling virtual textures and objects to be spatially registered with real-world surfaces within the carriage; the kit used a room model captured before the sessions rather than continuous geometry reconstruction. This can be seen in fig. 6.1b.

The pass-through AR and fully immersive VR media-viewing conditions were conducted inside a decommissioned subway carriage housed within a public museum exhibit at the Riverside Museum, Glasgow (as seen in fig. 6.3b). The carriage is Glasgow Subway car 39T, built in 1898 and used until the Subway closed for modernisation in 1977, after which it was restored

[2] https://developers.meta.com/horizon/documentation/unity/unity-mr-utility-kit-overview/

to its original condition by Glasgow Museums staff [54]. It has a wooden interior with a bench seat along each side, the same seating arrangement used on current Glasgow Subway trains. The carriage's external dimensions are 2680 × 7630 × 2300 mm (height × length × width). Although not an operational transport environment, and not representative of modern rolling stock, the carriage preserved key spatial and social characteristics of a transit environment, including seating layout, spatial constraints, sightlines, and passenger proximity. The carriage remained open to the public throughout the study and members of the public entered, exited, and moved through the carriage during the study. This introduced natural bystander presence, movement, and ambient noise, contributing to the ecological realism of the experience in ways that are difficult to reproduce convincingly in laboratory simulations; in the pass-through conditions this activity was both visible and audible, whereas in VR Cinema and VR Videosphere it was audible only. The study was conducted with the agreement of the museum, and museum staff were aware of it throughout: posters outside the carriage informed visitors that a study was in progress, visitors were asked not to interact with the participant or the researcher, and the researcher remained beside the participant throughout and was responsible for their safety while members of the public were present.

Because the study was conducted in a public museum exhibit, the physical carriage remained accessible to visitors throughout data collection. To characterise the social environment participants experienced during the physical-carriage conditions, a bystander count was recorded for each trial. This was defined as the total number of members of the public who were either present within the carriage at the start of a trial or who entered it during the trial. Counts were recorded by the experimenter and are reported descriptively; because visitor flow was outside experimental control, bystander presence was treated as a naturally occurring feature of the ecologically situated setting, rather than a manipulated variable. These records were retained to allow their potential influence on awareness-related outcomes to be examined during analysis; descriptive counts per condition are reported in section 6.3.

The simulated AR-in-VR conditions were administered in a nearby private room within the museum. This controlled private-room setting allowed for consistent presentation of simulated conditions without the variability introduced by real-world public activity. The simulated conditions used a virtual train carriage designed to approximate the spatial arrangement and visual characteristics of the physical carriage (as seen in fig. 6.1c). Simulated passengers were included in the scene as they were in chapters 4 and 5; however, in an effort towards heightened realism, the passengers were more animated and on walk cycles. The simulated passengers were implemented as animated character models using assets generated with Adobe Mixamo [**?**],[3] following prior VR research in which Fuse- and Mixamo-based virtual humans have been used to populate simulated environments with pedestrians and bystanders [79, 80], and consistent with transit-focused XR work in which simulated passenger representations have been used to

[3] `https://www.mixamo.com/`

study presence and awareness in shared transit contexts [17, 110]. The simulated passengers did not speak, but would stand up and move about the carriage at randomised intervals before returning to their seats. When seated, they exhibited basic idle animations such as posture shifts and small arm movements, approximating low-level visual occupancy and ambient motion within the carriage, rather than the full behavioural unpredictability or social accountability of real passengers [83]. This dual-context design enabled comparison between simulated and physically situated AR experiences, while maintaining experimental control over condition presentation.

6.2.4 Participants

Twenty-one people took part in the study, recruited from university mailing lists. One pilot session was not analysed, and one participant was excluded because they did not appear to understand the stop-task instructions and did not report the target stop in any condition.[4] The remaining nineteen participants (11 men and 8 women) ranged in age from 19 to 55 ($M = 28.95$, $SD = 9.02$). Most reported limited prior experience with XR head-mounted displays: eleven had used them once or twice and one had never used them, while the remaining seven reported more regular use ranging from a few times per month to daily. Nine participants wore glasses. One participant had also taken part in the study reported in chapter 5. After completing the study, participants were given a £10 Amazon voucher as compensation.

6.2.5 Tasks

Primary Task: Video Viewing

Each condition required participants to watch three approximately two-minute video clips that faded into one another. Clips were selected from media previously used in chapters 4 and 5: Game of Thrones, the Star Wars original trilogy, and David Attenborough nature documentaries. To support environment-based theming, each clip was limited to a single biome. Each grouped set of clips included one clip from each content source and one biome type (e.g. a Star Wars snow scene, a Game of Thrones forest scene, and a nature documentary desert scene).

Secondary Task: Stop-Announcement Awareness

During each condition, participants also completed a secondary awareness task intended to approximate the attentional demands involved in self-managed passenger journeys, where media engagement must be balanced against awareness of journey progress. The task required participants to monitor station information and identify when a predefined target stop was reached.

[4] Participant identifiers are not consecutive, and the identifier of the excluded participant was reassigned to a later participant.



(a) AR Baseline



(b) ARTV Ornamented



(c) Simulated AR Baseline



(d) Simulated ARTV Ornamented

Figure 6.1: Overview of the AR and simulated AR experimental conditions. Top row: AR Baseline and ARTV Ornamented, showing a floating video screen in the real train carriage with and without environment-matched augmentations. Bottom row: Simulated AR Baseline and Simulated ARTV Ornamented, showing the virtual train carriage with and without ornamentation.

In the pass-through AR conditions, station information was delivered via a physical 10-inch tablet positioned on the window opposite the participant within the carriage. The tablet acted as a scrolling train sign, displaying next-station information while playing corresponding speech announcements. In the Simulated AR Baseline and Simulated ARTV Ornamented conditions, this signage was replicated within the virtual train environment as a scrolling information display, so participants had access to both visual and auditory station information. In the VR Cinema and



(a) VR Videosphere (b) VR Cinema

Figure 6.2: The two fully immersive VR media-viewing conditions. Participants were seated in the physical train carriage for both. Left: VR Videosphere, showing the video on a floating panel inside a panoramic image matching the setting of the clip; right: VR Cinema, showing the same clip on a large screen in a virtual cinema.

VR Videosphere conditions, participants did not have access to visual station signage, because the surrounding environment was replaced by the immersive VR media-viewing environment. Station information was therefore delivered through audio announcements only, reflecting the reduced access to external visual journey information that occurs when using occlusive VR in transit contexts. In every condition, the announcements were played from the tablet's speaker at full volume and were heard through the participant's headphones.

Participants were instructed to monitor this information throughout each condition and to notify the experimenter when the train arrived at the target station. Performance was recorded as a binary outcome (pass/fail): if the participant did not alert the experimenter within 10 seconds of the target station being announced, a failure was recorded. Station names were designed to sound plausible but did not correspond to real locations, and were changed between conditions to prevent memory from aiding performance. Each condition comprised 5 stops, announced at intervals of approximately between 35 and 75 seconds, and the target was the second, third, or fourth stop, selected at random for each condition.

This task was not intended as a fine-grained measure of attentional load or cognitive performance. Rather, it was designed to capture whether participants could maintain awareness of contextually relevant journey information while engaging with media content, reflecting the ecologically meaningful attentional demands associated with public transport use. It is therefore a controlled proxy for journey monitoring: participants were not travelling, and missing the target station carried none of the practical consequences of missing a stop on a real journey, a



(a) Interior view of the subway carriage used in the study.



(b) Exterior view of the subway carriage used in the study.

Figure 6.3: Interior and exterior views of the Glasgow Subway carriage exhibit at the Riverside Museum used as the physical study environment for the experiment reported in this chapter. The person in the photos has been obscured for anonymity. Images captured by the author.

constraint returned to in section 6.4.3.

6.2.6 Procedure

Participants were first welcomed, given an overview of the study and asked to provide informed consent. They then completed a short demographic questionnaire and were fitted with the headset and headphones. Participants were instructed on both the primary media-viewing task and the secondary awareness task before the conditions began.

Due to the fact that moving participants repeatedly between the physical carriage and the nearby private room used for the simulated conditions would have substantially increased setup time and participant fatigue, full counterbalancing across all six conditions was not feasible. Instead, conditions were grouped by location, with ten participants beginning with the physical-carriage block and nine beginning with the simulated block. Within the physical-carriage block, the four conditions followed a balanced Latin square (a Williams design [187]), such that across four orders each condition appeared once in each position and once after every other condition; the order of the two simulated conditions alternated between participants. Because nineteen of the twenty-four planned participants were analysed, these orders were not perfectly balanced, with the four carriage orders used six, five, four and four times. This block structure reduced movement and recalibration between environments, but meant that full randomisation across all six conditions was not feasible. For one participant, the questionnaire software recorded the same condition label for all six conditions; these labels were corrected from the session order sheet.

Participants completed all six conditions, each consisting of three short video clips presented sequentially. After each condition, participants completed quantitative questionnaires assessing experience, engagement, and awareness. Following all six conditions, participants completed an exit questionnaire and took part in a semi-structured interview.

6.2.7 Measures

After experiencing each condition, participants completed the short versions of the UEQ-S (8 items, semantic differential scales) [157] and the UES-SF (12 items, Likert scales) [130]. The UEQ-S captures pragmatic and hedonic quality, while the UES-SF captures engagement-related constructs including focused attention, perceived usability, aesthetic appeal, and reward. Together, these instruments provide complementary perspectives on overall experience and engagement, and have previously proven effective in distinguishing between AR and VR media conditions in chapters 4 and 5.

The Film-IEQ [149] was not included in this study, in order to keep a six-condition session to a manageable length; administering it would have added approximately twenty minutes to study duration. In chapter 5, its results were largely convergent with those of the UEQ-S and UES-SF.

In chapter 4, however, the interaction between field of view and degree of ornamentation was significant on the Film-IEQ but not on the UEQ-S or UES-SF overall scores, although it did appear on their Hedonic and Reward subscales (and on the Film-IEQ Transportation subscale). This chapter therefore reports user experience and engagement rather than media immersion as the Film-IEQ measures it, a limitation returned to in section 6.4.3.

Participants also rated their agreement with two awareness statements after each condition, on a 5-point Likert scale (1 = Strongly Disagree, 5 = Strongly Agree):

1. '*I was aware of other passengers whilst watching the video*'
2. '*I was aware of the augmentations*'

These items were included to support both research questions by capturing awareness of the surrounding environment and of the augmentations themselves, enabling comparison between simulated and situated AR experiences (RQ4-1) and across XR viewing modalities (RQ4-2). In the physical-carriage conditions, 'other passengers' referred to members of the public, who were visible in the pass-through conditions and audible only in the two VR conditions; in the simulated conditions it referred to the animated passengers. The second item was only meaningful in the three conditions that added visual content around the screen (the two ornamented conditions, and VR Videosphere, where it referred to the panoramic image), and is analysed for those conditions only.

After completing all conditions, participants filled in an exit questionnaire in which they ranked the six conditions in order of preference. They also indicated, for each condition, their agreement with the statement: '*I would be interested in watching a video on a train using this condition instead of a phone or laptop*'. These responses captured perceived suitability for real-world use.

Participants then took part in a short semi-structured interview concerning their experiences across conditions. The four core interview questions were:

1. Tell me about your experience in [CONDITION]
2. How did the VR conditions with the virtual train compare to the AR conditions in which you could see the real train?
3. Can you explain why you ranked the conditions in the order that you have?
4. Do you think your awareness of which stop to get off at was the same across all conditions?

Follow-up questions were asked as needed to clarify and expand on participants' responses.

6.2.8 Analysis

Perceived Usability items were reverse-scored before subscale and total scores were calculated. Missing questionnaire responses (21 item responses, 0.9% of the total, and four passenger-awareness ratings) were imputed using multilevel predictive mean matching (`mice` [181], method

`2l.pmm`, five imputations), and the first imputed dataset was analysed. Repeating every omnibus test on the remaining four imputations, and on the observed data without imputation, produced the same pattern of significant results.

For the UES-SF, UEQ-S, and each of their subscales, a Shapiro-Wilk test for normality was conducted to confirm the necessity of the ART [193]. For RQ4-1, a two-way ART ANOVA crossing Medium (pass-through, simulated) with Ornamentation (baseline, ornamented) was conducted on the four AR and simulated-AR conditions, with ART-C [43] contrasts, Holm-corrected, used to compare pairs of cells. For RQ4-2, a one-way ART ANOVA across all six conditions was followed by ART-C contrasts with Tukey correction. This approach is consistent with the analyses reported in chapters 4 and 5. Effect sizes are reported as partial η^2 for ANOVA effects and as Hedges' g_z [86] for contrasts, calculated from participants' raw paired differences. With nineteen participants, a single paired comparison had 80% power to detect effects of $g_z \geq 0.68$ at $\alpha = .05$, or $g_z \geq 0.99$ with a Bonferroni correction for the fifteen pairwise comparisons, so smaller differences may not have been detected. Internal consistency was acceptable: median Cronbach's α across conditions was .80 for the UEQ-S and .89 for the UES-SF, and between .80 and .89 for the subscales, with the lowest value in any single condition being .57 for Perceived Usability.

Several supplementary tests were also conducted. A further ART ANOVA including block order (simulated conditions first or second) as a between-participants factor found no order $\times$ condition interaction on any measure (all $p \geq .075$); participants who completed the simulated block first gave higher Focused Attention ratings across all conditions ($F(1,17) = 4.72$, $p = .044$), but because block order also coincided with when participants took part, this may reflect differences between the two groups rather than order itself. Awareness of the augmentations was compared across the three conditions in which the item was meaningful using a Friedman test with Holm-corrected Wilcoxon signed-rank tests, awareness-task outcomes were compared across conditions using Cochran's Q, and rankings and usage-intent ratings were compared using Friedman tests with Holm-corrected Wilcoxon signed-rank tests.

For the physical carriage conditions, bystander count was also examined as an additional factor to assess whether variation in public presence influenced the measured outcomes. This analysis was included because bystander presence varied naturally across trials as a consequence of conducting the study in a public museum exhibit. Cumulative link mixed models [30] of passenger awareness, with a random intercept for participant, were compared with and without a term for bystander count.

Interview recordings were transcribed offline using WhisperX [14].[5] The transcripts of all nineteen participants were analysed using a single-coder hybrid inductive–deductive thematic analysis approach [154], involving iterative cycles of first-cycle and second-cycle coding. In the first cycle, coding was conducted inductively to capture segments relevant to participants'

[5] `https://github.com/m-bain/whisperX`

experiences across conditions. In the second cycle, these codes were refined, merged where overlapping, and developed into higher-level themes. A deductive layer was then applied to organise themes in relation to the study's research questions, specifically around: (1) differences between pass-through AR and simulated AR, (2) the effects of ornamentation, (3) awareness and attention management, and (4) perceived suitability for public transport use. To support consistency, codes and theme definitions were revisited iteratively across transcripts, with analytic memos used to document coding decisions. Inter-rater reliability was not calculated, as the analysis was intended to provide an interpretive account of participants' experiences rather than a reliability-tested coding scheme. The resulting themes were used to contextualise the quantitative findings.

6.3 Results

The results are presented in terms of how different viewing conditions shape the balance between immersive engagement and environmental awareness under realistic attentional demands.

6.3.1 Quantitative Results

Scores for the overall scales and sub-scales were calculated according to their respective procedures [157, 130]. Median scores per condition are reported in table 6.1. Results of the two-way analysis for RQ4-1 are reported in table 6.2. Significant ART ANOVA results and pairwise contrasts for the six-condition analysis are reported in table 6.3; a full table of all contrasts including non-significant results is provided in table F.4 in appendix F. Raw experiment data can be found archived at the Zenodo repository [32].

Table 6.1: Median scores per condition for UEQ-S subscales (range -3 to $+3$), UES-SF subscales (range 1–7), and awareness items (range 1–5). Awareness-P = awareness of other passengers; Awareness-A = awareness of augmentations, analysed only for the three conditions containing added visual content.

Condition	UEQ-S			UES-SF					Awareness	
	Pragmatic	Hedonic	Total	Foc. Att.	Perc. Usab.	Aesthetics	Reward	Total	Awr-P	Awr-A
AR Baseline	1.75	1.00	1.25	4.33	7.00	4.33	5.00	5.17	5	–
ARTV Ornamented	0.75	1.75	1.25	4.33	6.33	4.00	5.00	4.92	4	5
Simulated AR Baseline	1.50	0.75	0.75	4.67	7.00	4.00	4.33	4.83	3	–
Simulated ARTV Ornamented	1.00	2.50	1.75	5.33	6.33	5.00	5.67	5.42	3	5
VR Cinema	1.00	2.00	1.50	6.33	6.67	6.00	5.67	6.17	2	–
VR Videosphere	0.75	2.00	1.13	5.33	6.33	5.33	5.33	5.42	2	5

RQ4-1: Medium and Ornamentation Comparisons

The following results examine how pass-through AR and simulated AR conditions compared with and without ornamentation, using the two-way analysis reported in table 6.2. No Medium $\times$ Ornamentation interaction was significant on any measure (all $p > .12$).

Table 6.2: Two-way ART ANOVA results for the four AR and simulated-AR conditions, crossing Medium (pass-through, simulated) with Ornamentation (baseline, ornamented). All effects have df $= (1, 54)$. Significant effects are shown in bold.

	Medium			Ornamentation			Medium × Ornamentation		
Measure	F	p	η_p^2	F	p	η_p^2	F	p	η_p^2
UEQ-S Total	0.11	.74	.00	3.17	.081	.06	1.76	.19	.03
UEQ-S Pragmatic	0.93	.34	.02	**4.96**	**.030**	**.08**	0.62	.43	.01
UEQ-S Hedonic	0.00	.95	.00	**20.60**	**< .001**	**.28**	0.72	.40	.01
UES-SF Total	1.81	.18	.03	2.49	.12	.04	2.15	.15	.04
Focused Attention	**6.37**	**.015**	**.11**	3.04	.087	.05	0.58	.45	.01
Perceived Usability	0.24	.63	.00	4.01	.050	.07	0.33	.57	.01
Aesthetic Appeal	1.63	.21	.03	3.66	.061	.06	2.11	.15	.04
Reward	0.42	.52	.01	**5.94**	**.018**	**.10**	2.44	.12	.04
Awareness of passengers	**11.29**	**.001**	**.17**	0.05	.82	.00	1.59	.21	.03

Hedonic Quality. The clearest ornamentation effect was observed for Hedonic quality ($F(1,54) = 20.60$, $p < .001$, $\eta_p^2 = .28$). Ornamented conditions scored higher than their baselines in both media: in simulation ($p = .002$, $g_z = 0.65$) and in pass-through AR ($p = .034$, $g_z = 0.62$), providing further evidence that the ornamentation benefit observed in earlier chapters persists both under simulation and in a physical carriage.

Reward and Pragmatic Quality. Ornamentation also increased Reward ($F(1,54) = 5.94$, $p = .018$, $\eta_p^2 = .10$), although the difference within a medium reached significance only in simulation ($p = .041$, $g_z = 0.46$; pass-through $p > .99$, $g_z = 0.22$). Ornamentation reduced pragmatic quality ($F(1,54) = 4.96$, $p = .030$, $\eta_p^2 = .08$), with neither within-medium difference significant (pass-through $p = .18$, $g_z = -0.53$; simulated $p = .88$, $g_z = -0.22$), and the effect on Perceived Usability was marginal and in the same direction ($p = .050$). Ornamentation therefore improved hedonic aspects of the experience at some cost to its perceived practicality.

Medium Effects. Focused Attention was higher in the simulated medium than in pass-through AR ($F(1,54) = 6.37$, $p = .015$, $\eta_p^2 = .11$), with no significant effect of ornamentation ($p = .087$) and no interaction. Awareness of other passengers ran in the opposite direction, being higher in pass-through AR ($F(1,54) = 11.29$, $p = .001$, $\eta_p^2 = .17$), again with no significant effect of ornamentation ($p = .82$) and no interaction; this difference was significant between the two baseline conditions ($p = .003$, $g_z = 0.85$) but not between the two ornamented conditions ($p = .38$, $g_z = 0.36$). Because the two media also differed in setting, as the simulated conditions were run in a private room rather than the public carriage, neither difference can be attributed to the medium alone.

Awareness of Augmentations. Ratings of awareness of the augmentations were high in all three conditions containing added visual content (median 5 in each), but differed between them

Table 6.3: ART ANOVA results and significant pairwise contrasts for the six-condition analysis. F statistics have df $= (5, 90)$ throughout. Contrasts are ART-C with Tukey correction; g_z is Hedges' g_z calculated from raw paired differences; negative g_z indicates the first-named condition scored lower. Non-significant measures (UEQ-S Overall, UEQ-S Pragmatic, UES-SF Perceived Usability) are omitted; see appendix F for the complete results.

		ART ANOVA			Significant ART-C Contrasts		
Questionnaire	**Subscale**	F	p	η_p^2	Comparison	p	g_z
Short User Experience Questionnaire	**Hedonic**	5.72	$< .001$	0.24	AR Baseline − Sim. ARTV Ornamented	.009	−0.58
					AR Baseline − VR Videosphere	.013	−0.78
					ARTV Ornamented − Sim. AR Baseline	.037	0.64
					Sim. AR Baseline − Sim. ARTV Ornamented	.003	−0.65
					Sim. AR Baseline − VR Videosphere	.004	−0.81
User Engagement Scale Short Form	**Overall**	5.28	$< .001$	0.23	AR Baseline − VR Cinema	.002	−1.30
					ARTV Ornamented − VR Cinema	.004	−0.69
					VR Cinema − Sim. AR Baseline	.003	1.04
	Focused Attention	8.34	$< .001$	0.32	AR Baseline − VR Cinema	$< .001$	−1.79
					AR Baseline − Sim. ARTV Ornamented	.022	−0.60
					AR Baseline − VR Videosphere	.019	−0.64
					ARTV Ornamented − VR Cinema	$< .001$	−1.26
					VR Cinema − Sim. AR Baseline	$< .001$	1.26
	Aesthetics	5.53	$< .001$	0.23	AR Baseline − VR Cinema	.001	−1.10
					ARTV Ornamented − VR Cinema	.007	−0.65
					VR Cinema − Sim. AR Baseline	.001	0.91
	Reward	3.11	.012	0.15	VR Cinema − Sim. AR Baseline	.028	0.85
I was aware of other passengers whilst watching the video		6.13	$< .001$	0.25	AR Baseline − VR Cinema	$< .001$	1.02
					AR Baseline − Sim. AR Baseline	.013	0.85
					AR Baseline − VR Videosphere	$< .001$	0.96
					ARTV Ornamented − VR Cinema	.034	0.70
					ARTV Ornamented − VR Videosphere	.034	0.63

($\chi^2(2) = 10.8$, $p = .005$, Kendall's $W = .28$). Participants reported greater awareness of the augmentations in Simulated ARTV Ornamented than of the panoramic image in VR Videosphere ($p = .024$); ARTV Ornamented did not differ significantly from VR Videosphere ($p = .074$) or from Simulated ARTV Ornamented ($p = .49$).

Bystander Counts Bystander counts were recorded for eighteen of the nineteen participants and varied across the four physical-carriage conditions. Mean counts per trial were: ARTV Ornamented ($M = 8.89$, $SD = 5.90$, range 2–29), AR Baseline ($M = 7.28$, $SD = 3.66$, range 1–14), VR Cinema ($M = 10.28$, $SD = 4.87$, range 4–22), and VR Videosphere ($M = 10.28$, $SD = 6.43$, range 3–23). Including bystander count in a cumulative link mixed model of passenger awareness improved model fit ($\chi^2(1) = 4.37$, $p = .037$): participants reported greater awareness of other passengers when more members of the public were present (odds ratio 1.96 per standard deviation of bystander count, $p = .046$), with the differences between conditions unchanged.

Awareness Task Performance. Eight of the 114 trials (7%) ended in a missed stop. Each condition produced at least one failure, with Simulated ARTV Ornamented and VR Videosphere each producing two, and failure rates did not differ significantly across conditions (Cochran's $Q(5) = 1.05$, $p = .96$). With so few failures, the task had little capacity to detect differences between conditions, and these results are reported descriptively. They indicate that most partic-

ipants noticed the target stop in every condition; they do not establish that the conditions placed equivalent demands on attention. Within the physical-carriage conditions, the five missed stops occurred in trials with fewer bystanders on average ($M = 6.4$) than trials in which the stop was noticed ($M = 9.4$), but there were too few failures to test this.

RQ4-2: ARTV Compared with VR Viewing Conditions

Pairwise contrasts examined how the AR-based conditions compared with the two immersive VR viewing conditions, contextualising ARTV within the existing forms of immersive media viewing.

Hedonic Quality. The UEQ-S Hedonic subscale was significant ($F(5,90) = 5.72$, $p < .001$, $\eta_p^2 = 0.24$). All conditions produced positive hedonic scores, reflecting that even unaugmented pass-through AR was experienced as an enjoyable viewing mode. In comparative terms, AR Baseline (median = 1.00) scored lower than Simulated ARTV Ornamented ($p = .009$, $g_z = -0.58$) and VR Videosphere ($p = .01$, $g_z = -0.78$), but did not differ significantly from VR Cinema ($p = .21$). The lowest-scoring condition was Simulated AR Baseline (median = 0.75), which scored significantly lower than Simulated ARTV Ornamented ($p = .003$, $g_z = -0.65$) and VR Videosphere ($p = .004$, $g_z = -0.81$), whilst ARTV Ornamented scored higher than Simulated AR Baseline ($p = .04$, $g_z = 0.64$). The highest hedonic scores were recorded for the ornamented and panoramic conditions, with Simulated ARTV Ornamented (2.50) and both VR conditions (2.00) leading. UEQ-S Pragmatic quality and UEQ-S Overall were not significant ($F(5,90) = 1.41$, $p = .23$; $F(5,90) = 0.83$, $p = .53$ respectively).

Overall Engagement. UES-SF Overall was significant ($F(5,90) = 5.28$, $p < .001$, $\eta_p^2 = 0.23$), with VR Cinema scoring higher than AR Baseline ($p = .002$, $g_z = -1.30$), ARTV Ornamented ($p = .004$, $g_z = -0.69$), and Simulated AR Baseline ($p = .003$, $g_z = 1.04$). No significant differences were found between the AR conditions themselves on this measure.

Focused Attention. UES-SF Focused Attention was the most strongly significant subscale ($F(5,90) = 8.34$, $p < .001$, $\eta_p^2 = 0.32$). The pattern here reflects the degree of environmental surround or occlusion rather than ornamentation: both real AR conditions scored equally (median = 4.33), while conditions offering greater environmental surround or occlusion had higher medians: Simulated AR Baseline (4.67), VR Videosphere (5.33), Simulated ARTV Ornamented (5.33), and VR Cinema (6.33). VR Cinema produced significantly higher Focused Attention than AR Baseline ($p < .001$, $g_z = -1.79$) and ARTV Ornamented ($p < .001$, $g_z = -1.26$), while Simulated AR Baseline was also significantly lower than VR Cinema ($p < .001$, $g_z = 1.26$). AR Baseline additionally scored lower than both Simulated ARTV Ornamented ($p = .02$, $g_z = -0.60$) and VR Videosphere ($p = .02$, $g_z = -0.64$). No significant difference

was found between the two AR conditions or between the two simulated conditions on this measure.

Aesthetic Appeal and Reward. Aesthetic Appeal was significant ($F(5,90) = 5.53$, $p < .001$, $\eta_p^2 = 0.23$), with VR Cinema rated higher than AR Baseline ($p = .001$, $g_z = -1.10$), ARTV Ornamented ($p = .007$, $g_z = -0.65$), and Simulated AR Baseline ($p = .001$, $g_z = 0.91$). Reward was significant ($F(5,90) = 3.11$, $p = .01$, $\eta_p^2 = 0.15$), with only VR Cinema scoring higher than Simulated AR Baseline ($p = .03$, $g_z = 0.85$) as the sole significant contrast. Perceived Usability was not significant ($F(5,90) = 0.86$, $p = .51$).

Awareness of Other Passengers. Awareness of other passengers was significant ($F(5,90) = 6.13$, $p < .001$, $\eta_p^2 = 0.25$). AR Baseline produced significantly higher awareness than VR Cinema ($p < .001$, $g_z = 1.02$), Simulated AR Baseline ($p = .01$, $g_z = 0.85$), and VR Videosphere ($p < .001$, $g_z = 0.96$), while ARTV Ornamented demonstrated higher awareness than VR Cinema ($p = .03$, $g_z = 0.70$) and VR Videosphere ($p = .03$, $g_z = 0.63$). Both pass-through conditions therefore preserved more reported awareness of other passengers than either fully immersive VR condition.

6.3.2 Preferences

Preference rankings and usage-intent responses are shown in figs. 6.4 and 6.5. Neither differed significantly between conditions: rankings (Friedman $\chi^2(5) = 6.34$, $p = .28$, Kendall's $W = .07$) and usage-intent ratings ($\chi^2(5) = 4.95$, $p = .42$, $W = .05$), with no pairwise comparison remaining significant after Holm correction. The orderings below are therefore descriptive. ARTV Ornamented and VR Cinema were each ranked in the top three by 12 of 19 participants, followed by Simulated ARTV Ornamented and VR Videosphere (10 each), AR Baseline (9), and Simulated AR Baseline (4); Simulated ARTV Ornamented was most often ranked first (7 participants), and Simulated AR Baseline had the lowest mean rank (4.37, compared with 3.11–3.63 for the other conditions).

Interest in using a condition in place of a phone or laptop was highest for the AR conditions: 15 of 19 participants (79%) agreed or strongly agreed for AR Baseline, and 13 (68%) for ARTV Ornamented. The lower intent for ARTV Ornamented was attributed in interviews to implementation limitations, including texture instability (Z-fighting: a visual artefact in which two surfaces at nearly the same depth cause flickering as the renderer alternates between them), augmentation clipping, and lower visual fidelity, rather than rejection of the ornamentation concept. Ten participants (53%) agreed or strongly agreed for each of VR Cinema and VR Videosphere, and nine (47%) for each of the simulated conditions, though several participants appeared to interpret the simulated conditions hypothetically, imagining the simulation layered onto a real journey rather than as a standalone condition. The usage-intent item did not specify a journey

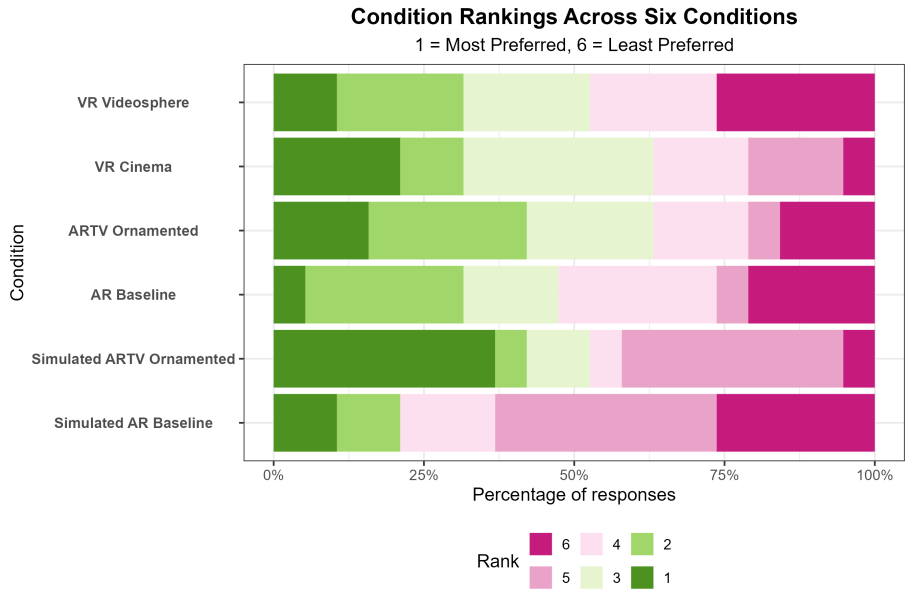


Figure 6.4: Participant ordering of the conditions from 1 (*Most Preferred*) to 6 (*Least Preferred*). Differences between conditions were not statistically significant.

type; in interviews, willingness to use the fully immersive conditions depended strongly on journey length, with ten participants saying they would consider them for a long journey on which they did not have to watch for their stop. Overall, twelve participants stated they would be interested in using some form of ARTV Ornamentation on public transport if current implementation issues were resolved.

6.3.3 Interviews

Interview data were collected for each of the conditions. Quotations have been edited for clarity, with words added for sense shown in square brackets. Longer quotations are presented as indented blocks. The following themes were identified through the hybrid inductive–deductive thematic analysis described in section 6.2.8, and are organised in relation to the two research questions. The raw interview transcripts can be found at the Zenodo repository [32].

RQ4-1: Comparing Pass-Through AR and Simulated AR Conditions

Simulation Transferred at the Level of Concept, but Not Situated Experience. When asked directly whether the simulated and pass-through AR conditions were interchangeable, six of the sixteen participants asked responded yes. Three additional participants felt they would have been interchangeable had the fidelity of the AR ornamentation matched the simulated version, and two felt they would have been were it not for the presence of real visitors. P21 stated: *'I think they are very similar. I mean, aside from the blurring of reality being as convincing, I'd say they were quite similar. [The simulated ornamented] condition was really, really good. The*

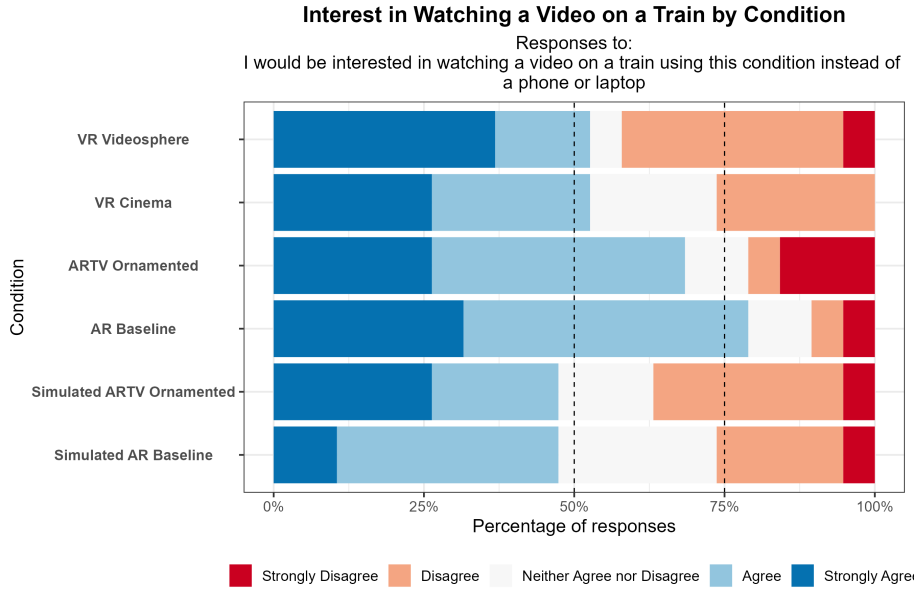


Figure 6.5: Participant responses to the statement *I would be interested in watching a video on a train using this condition instead of a phone or laptop*. Dashed black lines indicate 50% and 75% agreement thresholds. Differences between conditions were not statistically significant.

only issue is just the fidelity being better — it can make you feel more immersed'. P1 noted that '*the augmentations were pretty similar across them both*'. Five participants, however, felt the differences were too significant for the conditions to be considered equivalent.

Participants also identified important environmental and social differences. The physical carriage was described as noisier, more unpredictable, and more socially dynamic than the simulated version. P16 captured the perceptual contrast: '*There was definitely kind of the real world kind of slipping in and out a bit... less immersive, but also easier to keep an eye on kind of like where you were on the train line*'. The private room was not experienced as neutral either: P2 described feeling afraid there, '*maybe because it was empty*', and relief in the carriage at being able to '*see who's coming, who's going, who's like speaking around me*'.

Pass-Through AR Exposed Implementation Issues That Simulation Did Not Capture. Ten participants reported issues with ARTV Ornamented, including augmentation clipping into physical objects, flickering textures (Z-fighting), and lower visual fidelity than the simulated equivalent (P1, P4, P6, P8, P9, P10, P16, P19, P20, P21). Seven specifically described problems with the overlaid textures (P1, P4, P9, P10, P16, P19, P20), which were sometimes drawn in front of or through people in the carriage, including the researcher, so that faces and bodies appeared and disappeared. P4 explained:

> *I think it's a limitation of the technology. The textures kept moving here and there. Sometimes a person's face was occluded by the texture and sometimes it was not. So it was quite distracting for me and I was finding it hard to focus on the clip because*

the textures were moving around.

Mismatches between virtual content and the carriage could also create social ambiguity: P4 saw children in the carriage pointing and was briefly unsure whether they were pointing at a virtual object that only P4 could see. Four participants noted that Simulated ARTV Ornamented provided higher-fidelity augmentation (P6, P16, P19, P21), with P6 observing: '*I think the detail in the [simulated ornamented] condition was more defined*'. The simulated conditions had their own difficulties, however, which concerned the presentation of journey information rather than the augmentations: the virtual station sign did not appear to update for P5, lagged behind the clip change for P20, who missed the stop in that condition, and was difficult to see from a seated position for P6 and P19, with P3 and P5 describing how they relied on the announcements alone as a result. In the physical carriage, the floating screen could also partially obscure the tablet (P19).

Participants also raised costs arising from the headset itself rather than from the augmentations: P19 found it heavy, P16 reported eye strain and a headache by the end of the session, and P3 found it difficult to obtain a comfortable view through varifocal glasses. P16 and P6 both raised how conspicuous the headset would be to other passengers, with P16 noting that it is '*too expensive a piece of property*' to use on a train.

Hearing and Seeing Real Passengers Made the Physical Carriage Feel Fundamentally Different. Four participants described the virtual passengers as unconvincing or did not notice them at all (P1, P13, P14, P18), and two described them as less noticeable than real visitors because they neither spoke nor came close (P8, P15). P14 noted: '*I could also see people in virtual reality. It didn't feel real, but I could see them*'. Not all participants agreed: P16 found the simulated carriage '*very realistic*' and was occasionally startled when a virtual passenger appeared nearby. Real passengers were considerably more noticeable in the AR conditions, particularly because of ambient conversation and movement. P14 described how '*other real passengers made it very realistic — I could hear people's voices talking to each other*'. Eight participants mentioned the noise of real visitors as a key difference between the physical and simulated settings (P1, P9, P13, P14, P15, P17, P20, P21), and three found the virtual train unconvincing overall (P5, P6, P13), citing either the visual quality of the environment or the predictable behaviour of the avatars as failing to disrupt their experience in the way real passengers did. Visitors occasionally went further than being present: despite the request not to interact with participants, one entered the carriage and spoke to P21 during a pass-through condition.

RQ4-2: ARTV Compared with VR Viewing Conditions

Ornamentation Increased Immersion but Added Visual Complexity. Eleven participants clearly preferred ornamented conditions, often because environmental theming made the media feel more absorbing (P1, P3, P5, P7, P14, P15, P16, P18, P19, P20, P21). Several described

the experience of having the carriage themed to the content as a qualitatively different kind of engagement. P18 explained:

> *I particularly really enjoyed when the train was decorated around the movie that you were watching and then it changed based on each movie. I found that just really fun and interactive and you felt more immersed in the film.*

Others described a similar sense that ornamentation added meaningfully to the viewing experience. P6 noted that '*it kind of complemented what was on the screen — what I was watching — I thought it was very immersive*'. P1 described the ornamented condition as '*a lot cooler*' and feeling '*much more immersive*'. P15 similarly felt the augmented version '*really did feel exciting and innovative,*' while P20 described how, when a clip changed, they would look around for what had changed in the environment: '*I like that it adds to the clip as well*'. P7 noted that ornamentation was '*more immersive because it looks like you're in the set or something — because everything is more interesting than only looking at the usual trains*'.

However, eight participants preferred baseline conditions because they valued simplicity or awareness (P2, P4, P6, P8, P9, P10, P13, P17). P8 explained: '*I felt personally if there was a bit too much detail and some of the things going on in the immersive one, it was distracting from the actual video clips*'. P19 similarly stated: '*I do have a preference of being aware of my surroundings. So I really did like the Simulated AR Baseline*'. Two participants, P6 and P18, split their answers, preferring the ornamented condition in the simulated setting but the baseline for real travel. These responses support the quantitative pattern: ornamentation improved hedonic aspects of viewing but introduced some practical and attentional costs.

Fully Immersive VR Supported Focus but Was Less Suitable for Travel. Thirteen participants preferred VR Cinema over VR Videosphere and four preferred VR Videosphere, with the remaining two not comparing them directly. Those preferring VR Cinema cited immersion ($n = 10$), screen size ($n = 6$), and similarity to a real cinema ($n = 5$). VR Cinema was described as coherent and focused. P14 explained their preference: '*I think the cinema was pretty interesting for me because watching a movie in a cinema is a very great experience for me... I really liked the fact that you can be in a cinema and watch a movie*'. P6, who described themselves as a regular cinema-goer, felt VR Cinema '*did mimic that very well*'. P8 described it as producing '*total immersiveness*'. P17 said simply: '*It really felt like you were in a cinema*'.

By contrast, some participants found VR Videosphere less integrated with the content. P19 noted: '*You're putting some 360 image on a totally separate movie. It's not made for it*'. P4 found the scale of the surrounding imagery unconvincing, noting that trees appeared oversized and that sudden brightness changes in the snow biome were distracting. P18 offered a nuanced distinction:

> *I think I preferred the cinema because I was putting all of my attention into that one screen. Whereas with the video sphere, it was really cool, but I was taking my*

> *eyes off the screen and putting my attention into the certain things that were around me... If I was using the system just purely to watch a film, then I would definitely opt for the cinema. Whereas if I was just bored and thought, let's put on something random, then yeah, the video sphere would be cool.*

However, despite its strong engagement and focused attention scores, VR Cinema was described as less acceptable for everyday transit use, as its attentional capture came at the cost of environmental awareness. Only P13 suggested using it on an ordinary journey, and then only with the volume turned down, whereas ten participants said they would consider the fully immersive conditions for a longer journey on which they did not need to monitor their stop (P1, P2, P6, P9, P10, P13, P14, P16, P17, P21).

Participants Valued Flexible Awareness of Passengers, Belongings, and Journey Cues. Awareness was central to participant accounts of in-transit viewing. Fourteen participants reported that awareness varied across conditions, three reported that it did not, and 11 explicitly wanted to remain aware of other passengers (P1, P4, P6, P7, P8, P13, P15, P17, P18, P19, P21). P1 explained their reasoning: '*I wanted some level of awareness of passengers if someone is trying to speak to me or someone needs to sit there next to me, or someone needs to move my bag, or something — that's the kind of stuff I'd want to be notified about*'. P7 valued the AR conditions because '*I can still be aware of the surroundings while I'm enjoying the movie*', and P15 stated: '*I would rather have the awareness of what's going on, even if it comes with a bit of distraction*'. A minority of participants ($n = 5$) were content to disengage from their surroundings entirely (P3, P13, P14, P16, P20), with P14 explaining: '*I liked the fact that I could just forget about everything in the environment and enjoy a movie while I'm inside the train*'.

Ten participants specifically mentioned wanting to track journey progress and upcoming stops (P1, P3, P5, P8, P9, P13, P16, P17, P18, P19). P19 contrasted the AR and VR conditions in terms of their utility for this:

> *For the cinema and the 360 thing, it seems like the VR headset is made for me to just sit down there, just watch, and don't bother with anything else. The AR helps me be aware so that I'm not late to where I'm supposed to go. It makes me feel more on time. Cinema, it's just purely immersing — not for awareness.*

Participants Identified a Direct Immersion–Awareness Trade-Off. Twelve participants described a direct trade-off between immersion and awareness (P1, P4, P6, P8, P10, P13, P14, P16, P17, P18, P19, P21). P13 summarised the progression across conditions:

> *The level of awareness definitely changed. I was more aware of the environment when there wasn't any modifications. I was definitely aware of my environment and knew who was there. As you started adding graphical elements to it, it became a*

little bit harder to keep track, and then in the cinema experience it was impossible, basically.

P4 described the contrast between the simulated and pass-through conditions concisely: '*I was more focused on my clip in the VR condition and I was more focused on other people in the AR condition*'. P1 observed that '*the more virtual the environments, the less aware I was of the passengers*'. P16 reflected: '*Less immersive, but also easier to keep an eye on kind of like where you were on the train line in the journey... it's a toss-up between wanting to be immersed and feeling unsafe*'. P13 also noted a potential safety concern with fully immersive VR: '*I could see the cinema conditions potentially having some negative consequences... it's so immersive that you sort of forget what's going on around you. That's a potential hazard, I think*'.

In the occlusive VR conditions, five participants described relying more deliberately on the audio announcements (P8, P14, P18, P19, P21). P14 explained that '*the primary strategy was to hear and to listen to the announcement*', but that in VR Cinema and VR Videosphere '*it was a bit more difficult because I was more immersed and it was a bit difficult to stay focused [on stop information]*'. P18 recalled: '*I hit a point and I was like, I need to start listening out, because my stop's going to be soon*', and reflected that they '*would be a lot more likely to miss my stop*' in the more immersive conditions. P21 found that even with the audio available, without a visual reference '*it became quite hard to know what stop is what*'. Three participants described this monitoring as effortful: P8 had to '*really put more effort into hearing*', P14 found it '*a bit stressful*', and P19 stated '*I need to listen to it, which I don't want to do*'. Some participants instead kept track by counting announcements rather than monitoring station names, with P5 describing how they would '*listen out for the train station announcement and count them*'.

6.4 Discussion

6.4.1 RQ4-1: Comparing Simulated and Physically Situated AR

This research question sought to understand not only whether VR-simulated AR can reproduce prior ARTV findings from chapters 4 and 5, but also where such simulations begin to diverge when deployed in socially dynamic, real-world contexts. Rather than treating equivalence as a binary property, our findings suggest that simulation fidelity operates along three distinct dimensions that collectively shape how ARTV experiences are perceived and evaluated: *perceptual accountability*, *social plausibility*, and *attentional consequence*. Each dimension captures a different way in which the physical transit environment introduces demands that could be challenging and costly for VR-based simulation to fully reproduce or simulate. These findings suggest that the apparent balance between immersion and awareness is shaped not only by augmentation design itself, but also by the ecological fidelity of the evaluation context in which that balance is assessed.

Perceptual Accountability

Perceptual accountability refers to the requirement that virtual content remain visually and spatially consistent with the physical world, such that misalignment, flicker, occlusion errors, or registration instability become immediately noticeable when augmentations coexist with real environments. In simulation, augmentations can be authored together with the virtual environment, ensuring that textures, objects, and surfaces align by design. In pass-through AR, by contrast, augmentations must be registered against a physical world that is only partially sensed and reconstructed, a structurally different challenge that simulation does not reproduce.

This distinction was directly visible in participant responses. The most frequently reported issues with the ARTV Ornamented condition were imperfect texture mapping onto physical surfaces, Z-fighting (where textures flickered as they competed with real-world geometry at similar depth values), and augmentation elements clipping through or being occluded by physical objects. Retextured surfaces were also sometimes drawn in front of or through real people, including the researcher. These artefacts arose from the use of Meta's Mixed Reality Utility Kit, which was at the time in beta, and which relied on pre-scanned room models rather than real-time geometry reconstruction. Had a custom implementation using continuous SLAM-based tracking been employed, some of these issues may have been reduced. The sporadic but noticeable nature of these failures was itself significant: participants could accept a lower-fidelity augmented environment, but were disrupted by inconsistency, as each instance of flickering or misalignment drew attention away from the media and toward the system's limitations.

The quantitative results are mixed as to how far these artefacts weakened the ornamentation benefit. The hedonic benefit of ornamentation was of comparable size in both media ($g_z = 0.62$ in pass-through AR and 0.65 in simulation), but the Reward benefit was significant only in simulation, and ornamentation reduced pragmatic quality in both. Participants themselves distinguished between the concept and its implementation: twelve of nineteen reported they would be interested in using ARTV Ornamentation on public transport if the technical issues were resolved. That interest is, however, both conditional and hypothetical, and the artefacts affected only one condition. The study cannot separate their effect from the effect of ornamenting a physical space as such: augmenting real surroundings may compete with those surroundings in ways that simulation does not reveal, in which case even a stable implementation might yield a smaller benefit than simulation suggests.

Social Plausibility

Social plausibility refers not merely to the presence of other people in the environment, but to whether surrounding behaviour, responsiveness, and unpredictability align with users' expectations of shared public space. The simulated conditions included animated passenger avatars generated using Adobe Mixamo, which moved about the carriage at randomised intervals and exhibited basic idle animations when seated. While these agents provided visual occupancy,

they did not speak, react to the participant, or generate ambient noise. Several participants described noticing that the avatars moved on predictable cycles and never initiated social behaviour, although one participant found them realistic enough to be startled by them, indicating that scripted agents are not uniformly unconvincing. Prior work supports this: Freeman et al. [49] found that neutral avatars in a virtual London Underground carriage were interpreted as hostile by a sizeable minority of participants, so scripted passengers can provoke social responses even when their behaviour is fixed.

The significance of this gap was not merely aesthetic. Real passengers in the physical carriage introduced genuine attentional and social demands: they moved unpredictably, spoke to one another, entered and exited the space, generated ambient noise, and on one occasion approached a participant directly. Eight participants identified the noise of real visitors as a key difference between the settings, describing how hearing other passengers' voices and conversations made the physical carriage feel meaningfully different as a social environment. This aligns with the difference observed between the two baseline conditions on self-reported passenger awareness, and with the finding that awareness ratings increased with the number of visitors present. Both patterns are consistent with the functional social demands that real bystanders created and that scripted avatars could not replicate, although the two media also differed in setting, so the comparison reflects the simulated environment as a whole rather than the avatars alone.

These findings suggest that for transit contexts, where social accountability is an ongoing background demand, scripted agents can approximate the spatial character of a carriage, but may not reproduce the behavioural fidelity needed to evaluate situated awareness. Researchers using simulated passengers should consider richer behavioural models that include conversational activity, unpredictable movement, and responsive idle behaviour. Spatial audio, representing ambient conversation and movement noise, may also be important given how consistently participants identified sound as a key differentiator between the physical and simulated environments.

Attentional Consequence

Attentional consequence refers to whether environmental events meaningfully affect how users allocate attention between media content, journey information, other passengers, and the physical setting. It captures not just what is present in the environment, but whether what is present places genuine demands on the user's attention that must be actively managed.

Evidence for this dimension comes primarily from the interviews. In the simulated conditions, the environment was largely predictable: avatars moved on known cycles, no unexpected events occurred, and the private room setting removed ambient unpredictability. Participants could therefore direct attention more freely toward media content. In the physical carriage by contrast, real bystanders could approach, sit nearby, speak, or move through the space at any

moment, and several participants explicitly reported having to monitor the carriage alongside the media. Participants further expressed concern about missing stops in more immersive conditions, a concern that reflects genuine attentional pressure rather than hypothetical awareness loss. The two settings also differed in how reliably journey information was presented: the physical tablet was continuously visible in the pass-through conditions, whereas several participants reported that the virtual sign lagged, failed to update, or was awkward to see. The implication is that simulation may underrepresent not only awareness itself, but the cognitive effort of maintaining awareness alongside immersive media, although the present study cannot separate that effort from the other differences between the two settings.

Simulation Validity Depends on the Research Claim

These dimensions point to conditional rather than general equivalence. Simulation reproduced the hedonic benefit of ornamentation, with no interaction between medium and ornamentation on any measure, which supports its use for early comparison of design variants such as layout, density, or thematic style. It did not reproduce the situated experience: the physical carriage introduced real passengers, ambient noise, unpredictability, and the need for augmentation to hold together against real surfaces. Simulation may therefore overestimate how well augmentation integrates with a real environment, and underrepresent the awareness demands it carries.

This is the distinction between face validity and ecological validity set out in section 2.7.2. A simulated carriage can look enough like a train to place participants in the scenario without capturing the social, perceptual, and attentional contingencies that shape real use. Consistent with Dole and Ju's argument that face validity is a useful proxy only when a simulation represents the features relevant to the research claim [41], simulation of the kind used in chapters 4 and 5 appears well suited to comparing ARTV variants, but gives only a partial account of social awareness, situated attention, or the perceptual accountability of pass-through AR; for those claims, *in situ* validation is needed.

For researchers designing simulation-based studies in transit contexts, two practical implications follow. First, simulation is best suited to evaluations where social dynamics are not a primary focus, unless the simulation is sufficiently sophisticated to portray and validate more socially realistic situations. Where social context is central to the research question, considerable care must be given to the behavioural fidelity of simulated passengers: not merely their visual appearance, but their responsiveness, unpredictability, and acoustic presence. Second, when implementing pass-through AR augmentation, registration fidelity and perceptual stability matter substantially to user experience. Researchers should distinguish clearly between limitations of the augmentation concept and limitations of the specific implementation pipeline used to evaluate it, without assuming in advance which of the two accounts for a weaker result.

6.4.2 RQ4-2: Immersion–Awareness Trade-offs Across XR Viewing Modalities

This research question examined how AR-based media experiences compare to currently available XR viewing approaches in a transit context, where entertainment must coexist with ongoing demands for environmental monitoring and journey management. This comparison examines how different XR viewing modalities affect the balance between media engagement and environmental awareness, and what this means for their practical suitability for everyday transit use. Because passenger media consumption is already well supported by smartphones and laptops, which are lightweight, socially familiar, and require no additional equipment, the XR conditions evaluated here must also be interpreted against this pragmatic baseline of existing handheld media use, and not only against each other.

ARTV Ornamentation Supports Flexible Immersion

A central finding of this study is that the value of ARTV Ornamentation in transit contexts does not lie in rivalling the focused immersion of VR Cinema, but in enriching engagement with media content while preserving the user's ability to re-enter the physical environment when needed. These findings motivate the term *flexible immersion* for this design strategy: increased immersion that does not require disengagement from the surrounding environment. The term describes the pattern observed here rather than a construct this study set out to validate. Participants valued ornamentation not because it maximised attentional capture, but because it improved the viewing experience while leaving the physical world accessible. Augmentations were experienced as an extension of the carriage environment rather than a replacement of it, allowing users to shift attention between media content and surroundings without the experience breaking down.

This was reflected in both the quantitative and preference data. Ornamentation significantly increased hedonic quality in both media, and increased Reward overall, while ARTV Ornamented preserved higher reported passenger awareness than either fully immersive VR condition. Although ARTV Ornamented did not significantly outscore AR Baseline on focused attention or on any other questionnaire measure, and although preference rankings did not differ significantly between conditions, ARTV Ornamented was among the conditions most frequently ranked in the top three and was the second most acceptable alternative to a phone or laptop. Participants described the carriage-themed augmentation as immersive, fun, and qualitatively different from unaugmented viewing, while still feeling connected to their surroundings. This suggests that ARTV Ornamentation may be best understood as a design strategy for balancing engagement and awareness, rather than for maximising immersion alone. This distinction that is particularly important in transit contexts, where full attentional capture is unlikely to suit most passengers.

Focused Immersion Alone Is Insufficient for Transit Suitability

The comparison between VR Cinema and the AR-based conditions illustrates why focused engagement alone is an insufficient criterion for evaluating in-transit XR media. VR Cinema produced the highest median scores across overall engagement, focused attention, and aesthetic appeal, and was described by many participants as coherent, cinematic, and absorbing. However, despite its strong performance on these engagement-related measures, only one participant suggested they would use VR Cinema on an ordinary journey, while ten said they would consider the fully immersive conditions for longer-distance travel. The same property that made the condition effective for sustained, focused viewing, its occlusion of the surrounding environment, also made it poorly suited to transit contexts where attention must periodically return to passengers, belongings, and journey progress.

This contrast highlights a fundamental tension in mobile media design: systems optimised for sustained, uninterrupted focus are often poorly matched to environments that require intermittent awareness. Participants who expressed willingness to use VR Cinema did so with explicit caveats about journey length and attentional demands, noting they would only consider it for longer or less demanding journeys, where monitoring the environment was less critical. Journey length was not manipulated in this study, so the suggestion that AR-based conditions are better suited to short, variable, or socially dynamic journeys rests on participants' own accounts rather than on measured differences. The immersion and awareness cost of VR Cinema also becomes clearer when compared against the pragmatic baseline of phone and laptop use: a device that occludes the environment and requires conspicuous specialist equipment offers a high barrier to adoption unless its immersive benefits substantially outweigh the costs, and handheld viewing was not itself evaluated here.

Environmental Surround Alone Was Not Sufficient

VR Videosphere produced significantly higher focused attention and hedonic scores than AR Baseline, but was not preferred to the other conditions, and did not differ from ARTV Ornamented on any questionnaire measure other than awareness of other passengers. This pattern suggests that environmental surround alone is not sufficient to make a viewing experience compelling in transit contexts. The Videosphere condition surrounded participants with a static 360-degree panoramic image matched to the biome of the current clip, but this image was not semantically connected to the specific media content being watched. Participants who responded negatively to the condition found the disconnect between the surrounding imagery and the narrative content unconvincing, with some describing the static biome image as being placed around a completely separate film rather than integrated with it. Others found the scale of surrounding imagery unconvincing, or were distracted by sudden brightness changes between biomes.

ARTV Ornamentation combined show-specific objects with retexturing matched to each clip, and was described more positively in the interviews despite providing less total environ-

mental coverage. One explanation is that the benefit of ornamentation does not arise simply from increasing visual coverage or surrounding the viewer, but from the semantic relationship between the augmentations and the media being consumed. This would be broadly consistent with chapter 5, in which the adaptive Biome treatment, combining retexturing with content-matched objects, was rated above the other conditions, although that study could not separate the contributions of its components. The present study does not test this explanation directly, however: the two conditions also differed in medium and in occlusion, and the questionnaire measures did not distinguish them.

VR Videosphere was not intended as a representative benchmark for all forms of 360-degree or panoramic media. Rather, it served as a controlled comparison approximating the contribution of environmental surround alone, decoupled from content-specific augmentation. Its performance relative to the ornamented conditions should therefore be interpreted as evidence about environmental theming without narrative integration, rather than as a general verdict on immersive panoramic media.

Maintaining Awareness Took Effort

Both AR conditions produced significantly higher self-reported passenger awareness than VR Cinema and VR Videosphere, as expected given that AR preserved visual access to the physical carriage. However, awareness task failure rates did not differ significantly across conditions, including the fully occlusive VR conditions. With only eight failures across 114 trials, and with some participants keeping track by counting announcements, the task had little capacity to detect differences, and it therefore cannot establish whether reduced visual access affected journey management.

Interview data indicate that participants in VR conditions relied more deliberately on the audio stop announcements, shifting their monitoring strategy from a visual to an auditory channel. This shift was partly required by the design rather than freely chosen: VR Cinema and VR Videosphere presented station information through audio alone, whereas the remaining four conditions also presented it visually, though the virtual sign in the simulated conditions was not always reliable. Modality is therefore confounded with occlusion in these comparisons. With that caveat, participants' accounts are consistent with the view that situational awareness involves not only perceiving environmental information but actively managing how attention is allocated across available cues [177].

This compensation should not be interpreted as cost-free. Three participants described having to listen more deliberately, one described reconstructing missed information retrospectively, and several noted they would be more likely to miss their stop in the more immersive conditions. For transit XR design, this suggests that robust and salient auditory journey cues (stop announcements, alerts, and schedule information) should be treated as a core design requirement when visual immersion is high, rather than an optional supplement. Auditory compensation does not

eliminate the awareness cost of occlusive VR; it relocates that cost to a different cognitive channel, and one which participants found demanding even over a short simulated journey.

Hardware and Pragmatic Costs of Current Pass-Through AR

The use of pass-through AR introduced costs that shaped user experience independently of augmentation design. Participants raised the weight of the headset, eye strain, difficulty obtaining a comfortable view through varifocal glasses, and how conspicuous the device would be to other passengers. These hardware-level constraints are distinct from the augmentation-specific issues discussed under perceptual accountability: they affect the baseline AR viewing experience regardless of whether or not ornamentation is present. As noted above, ARTV Ornamentation must ultimately be judged not only against fully immersive VR alternatives but against the practical baseline of smartphones and laptops. AR-based viewing can offer benefits that handheld devices cannot (larger perceived display size, spatial integration with the carriage, and content-themed augmentation that enriches the surrounding environment), but these benefits must outweigh the pragmatic costs of current hardware, including restricted field of view, camera-mediated visual quality, device weight, and social visibility. The present findings suggest that while participants reported that they would rather use AR-based viewing than fully immersive VR for everyday transit use, current pass-through hardware introduces perceptual and pragmatic costs that limit its immediate viability as a straightforward replacement for handheld viewing. As field of view, camera fidelity, and registration stability improve in future devices, many of these costs may diminish, a possibility returned to in chapter 7. The findings from this study should therefore be interpreted in light of the rapid pace of hardware development, while recognising that the study cannot show whether a stable implementation would remove these costs altogether.

6.4.3 Limitations

While this chapter advances the ecological validity of ARTV Ornamentation evaluation through deployment in a physically and socially situated environment, several limitations of the study design should be acknowledged. First, the subway carriage was stationary and housed within a museum exhibit rather than operating as part of a live transport service. Vehicle motion, real journey duration, and the temporal pressures associated with boarding and alighting were therefore absent. Participants were aware they were not on a real journey, which may have reduced the ecological salience of the awareness task: while the stop-announcement paradigm provided a controlled proxy for journey monitoring, missing the target station carried no genuine consequence, and monitoring behaviour may have been more relaxed than it would be under real travel conditions. Nor were participants responsible for disembarking at an actual destination, managing luggage, or sustaining XR use among real co-passengers, and the cost of failure differs substantially in operational transit, where missing a station may mean delayed onward travel

or feeling vulnerable in an unfamiliar place. The task therefore approximated one component of passenger awareness under controlled conditions, rather than reproducing the motivational, temporal, and consequential structure of live public transport use; this remaining gap, which might be termed *journey consequence*, sits alongside the three fidelity dimensions discussed in section 6.4.1.

Second, the bystanders present during the physical-carriage conditions were museum visitors rather than co-passengers sharing a journey. While their presence introduced genuine ambient activity, movement, and noise, they did not exhibit the social behaviours characteristic of shared transit use, such as occupying adjacent seats, stowing luggage, or the mutual social accountability of fellow travellers. Visitors were also asked not to interact with participants, although at least one did so, and visitor numbers varied over the course of data collection, including across the end of the school summer holidays, with higher mean counts in the two fully immersive VR conditions than in the pass-through conditions. Their presence therefore contributed to the ecological character of the setting, without fully reproducing the social dynamics of a live passenger environment, and should be distinguished from the unconvincing nature of the simulated avatars discussed in section 6.4.1.

Third, the medium was confounded with both setting and order. The pass-through and fully immersive VR conditions took place in the public carriage while the simulated conditions took place in a private room, and conditions were blocked by location rather than fully randomised. Differences between pass-through and simulated AR may therefore reflect the room, the presence of the public, or the order of the blocks, in addition to the medium itself. No order by condition interaction was found, but block order coincided with when participants took part, and at least one participant described the empty private room as itself uncomfortable.

Fourth, the awareness task had limited sensitivity. Failures were rare, and some participants reported counting announcements rather than monitoring station names. Station information was audio-only in the two fully immersive VR conditions by design, and several participants reported that the virtual sign in the simulated conditions lagged, failed to update, or was difficult to see, so the presentation of journey information was not fully equivalent across conditions. One participant was excluded for failing every awareness task, which may bias the observed failure rate downwards. More broadly, real passengers may monitor stops intermittently rather than continuously, rely on platform announcements, check route information, respond to crowding or service disruption, or alter their willingness to use XR depending on journey familiarity, perceived safety, time pressure, and social exposure, none of which this paradigm could capture.

Fifth, with nineteen participants the study was powered to detect only relatively large within-participant differences, and the Film-IEQ was not administered, so media immersion was not measured directly. One participant had previously taken part in the study reported in chapter 5, and was therefore already familiar with ARTV Ornamentation.

Sixth, a limitation specific to the simulated conditions concerns participant understanding of

the experimental setup. Several participants appeared to interpret the simulated AR conditions hypothetically, imagining the simulation layered onto a real journey rather than evaluating it as a standalone viewing mode. This likely influenced preference and usage-intent responses for those conditions, where ratings may reflect enthusiasm for the concept of AR on a real train rather than assessment of the simulated condition itself. Clearer pre-condition briefing or a brief familiarisation trial could help establish the correct mental model in future studies using simulated AR as a proxy for in-situ deployment.

Finally, the interview data were coded by a single coder and inter-rater reliability was not assessed. These limitations, and their broader implications for the evaluation of passenger XR systems, are returned to in section 7.6.1.

6.5 Conclusions

First, the design-level benefit of ARTV Ornamentation survives the move from simulation to a physical carriage, but in a narrower form. Ornamentation raised hedonic quality by comparable amounts in both media, with no interaction between medium and ornamentation, yet its effect on reward held only in simulation, and it lowered pragmatic quality in both. Evidence from simulation can therefore be relied on for the direction of a design effect, but not for the full profile of benefits and costs that a situated deployment will show.

Second, the validity of simulation depends on the claim it is used to support. It is adequate for comparing ARTV design variants, but not for claims about social awareness, situated attention, or the stability of augmentation registered to a real environment. Perceptual accountability, social plausibility, and attentional consequence identify where the two settings diverged, and journey consequence identifies a further gap that neither setting reached. Because medium was confounded with setting, these dimensions locate the divergence rather than explain it, and they are offered as a checklist for future passenger XR evaluations rather than as established causes.

Third, for everyday transit use, the design target is not maximum immersion. VR Cinema produced the strongest engagement, but participants reserved it for long journeys on which they did not need to watch for their stop, whereas both pass-through conditions preserved substantially more reported awareness. The behavioural task could not show whether that difference matters for journey management, and participants' accounts suggest that the cost of occlusion shifts to effortful listening rather than disappearing. Passenger media designs should therefore aim for the level of engagement that can be sustained while the passenger keeps the awareness their journey requires, and should treat redundant, salient journey cues as a requirement whenever vision is occluded.

These conclusions are bounded by a beta-stage implementation whose registration artefacts affected only the ornamented pass-through condition, and by a stationary carriage with no real journey. Chapter 7 considers them alongside the findings of the earlier chapters.

Chapter 7

General Discussion

7.1 Introduction

This chapter synthesises the findings presented throughout this thesis and develops their broader implications for the design of immersive media in passenger contexts. While the preceding chapters examined specific aspects of immersive media consumption through individual empirical studies, this chapter integrates those findings to provide a unified account of how immersive XR experiences can be designed to support both the media engagement and the environmental awareness that the passenger context requires.

The work undertaken in this thesis was motivated by the growing potential of XR technologies to transform media consumption beyond traditional stationary environments. Public transport contexts present a particularly promising yet challenging setting for immersive media, as passengers often seek engaging entertainment experiences while simultaneously needing to remain aware of their physical surroundings [17]. Chapter 2 demonstrated that existing immersive media approaches frequently prioritise either immersion or awareness, creating a tension that limits the practical deployment of XR systems in mobile and socially shared environments. Through a series of studies exploring commuter behaviour, immersive media augmentation strategies, and scalable ornamentation techniques, this thesis investigated how XR experiences can be designed to balance these competing requirements.

The chapter is organised in five parts. Section 7.2 develops the PIAT as a design concept describing how immersive engagement and environmental awareness relate to one another in passenger XR, drawing together evidence from across the research programme and establishing what kind of contribution the concept represents. Section 7.3 then sets out a design framework for ARTV in passenger contexts, defining the dimensions through which ARTV systems can respond to the trade-off that the PIAT describes. Section 7.4 considers coherence, plausibility, and the authoring of augmentations, which cut across those dimensions. Section 7.5 translates these into practical guidance for the design of passenger XR media systems, and section 7.6 considers the limitations of the research and the directions for future work that follow from

them.

7.2 The Passenger Immersion–Awareness Trade-off

Findings reported across the research programme collectively highlight the presence of a fundamental tension within passenger XR experiences, described throughout this research as the PIAT. This trade-off represents the competing demands between immersive engagement with digital media and the ongoing requirement for passengers to remain aware of their physical environment for personal safety and security, management of journey progress, and maintenance of socially appropriate behaviour within shared transit spaces. While immersive XR technologies are often designed to maximise presence and engagement, typically achieved through replacing the user's view of reality with a wholly or predominantly virtual scene, passenger contexts impose additional cognitive, social, and safety requirements that necessitate continued environmental awareness. This section establishes how the trade-off was observed across the studies, how it should be characterised, and what kind of contribution it represents. The design decisions through which ARTV systems can respond to it are set out separately in section 7.3.

7.2.1 Origins of the PIAT

Prior studies of passenger media use show that commuters already divide attention between media content and awareness of their surroundings [17, 110, 133]. Traditional mobile devices, such as smartphones and tablets, allow passengers to consume entertainment while maintaining peripheral awareness of their environment. However, these devices impose constraints on immersion due to their limited display size [64], restricted spatial interaction (where media interaction is largely confined to a handheld screen rather than being integrated with the surrounding physical environment via HMD), and fragmented attention between immersion in the content and awareness of their surroundings and journey management. Within this thesis, some participants were willing to accept less immersion in exchange for awareness: in chapter 4, the intermediate ornamentation level was chosen most often for use on a train, and in chapter 6, only one participant said they would use VR Cinema on an ordinary journey, despite its higher engagement scores.

The workshops reported in chapter 3 indicate that passengers already understood the problem in these terms before any system had been presented to them, one participant describing the design challenge as 'just balancing how much of reality you could still see' (W1-3). Proposals raised during those sessions consistently sought to preserve some access to the surrounding environment rather than to remove it, and one participant revised an initial preference for full immersion over the course of the discussion on the grounds that losing all sense of reality would be unsafe.

XR technologies offer the potential to increase passenger media immersion substantially by providing large-scale spatial displays, environmental augmentation, and interactive media experiences. However, prior chapters have demonstrated that increasing immersion through XR also introduces the risk of diminishing environmental awareness. Fully immersive XR experiences can occlude real-world visual and auditory cues, potentially reducing users' ability to monitor surroundings, respond to social signals, or detect environmental changes. This establishes a core tension within passenger XR design, in which increasing immersion can simultaneously increase the risk of awareness loss.

7.2.2 The Relationship Observed Across the Research Programme

Figure 7.1 positions every condition from the three empirical studies according to the two constructs at issue: focused attention, measured using the UES-SF Focused Attention subscale and used here as the engagement measure common to all three studies, and environmental awareness, measured through self-reported awareness of other passengers. As the three studies recruited separate participant groups, they are presented as independent samples rather than combined into a single analysis. Focused attention is used in place of the Film-IEQ because chapter 6 did not administer that instrument; where both were collected, in chapters 4 and 5, the two track one another closely.

Across the three studies, the condition means show a broadly consistent inverse pattern, in which those producing the strongest focused attention also produced the weakest reported awareness of the surrounding carriage, while those preserving awareness did so at a cost to engagement. In chapter 4 this ordering is close to monotonic across all eight conditions, running from the unornamented AR baseline, through progressively more heavily ornamented conditions, to the fully occlusive VR baseline. Chapter 5 and chapter 6 reproduce the same ordering, despite differences in media content, augmentation strategy, display hardware, and whether the surrounding environment was simulated or physically present. As the relationship persists across these differences, it is unlikely to be attributable to the particular characteristics of any single study.

Two conditions depart from this ordering. In chapter 5, the Genre-Based condition produced both higher focused attention and marginally higher reported awareness than the unornamented baseline, and in chapter 6 the Simulated ARTV Ornamented condition did likewise relative to the Simulated AR Baseline. Neither difference was statistically reliable and both are small in magnitude; in the case of Genre-Based, equivalence testing established only that its awareness did not differ from the baseline by more than ± 0.85 scale points, which was the smallest effect that study could detect [87]. These departures indicate that the trade-off operates as a tendency across conditions rather than as a constraint governing every pairwise comparison.

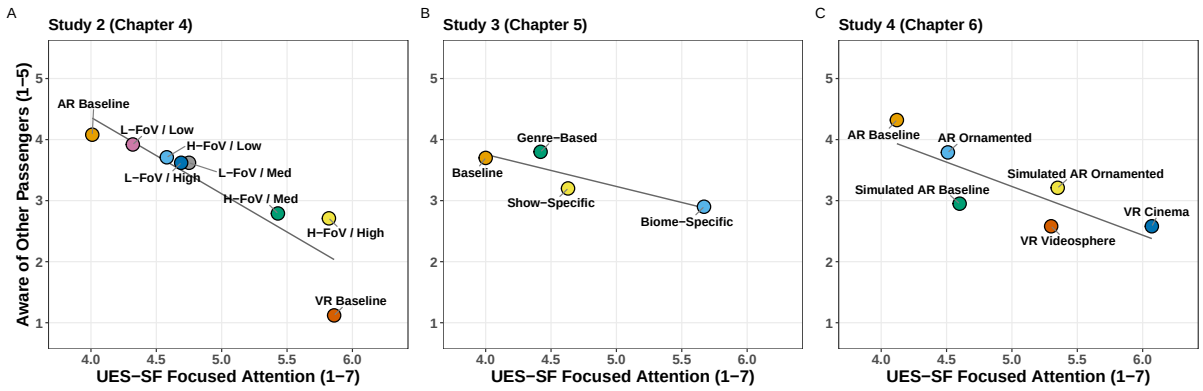


Figure 7.1: Condition means for each empirical study, positioned by focused attention (UES-SF) and self-reported awareness of other passengers. Chapter 4 condition labels abbreviate field of view and degree of ornamentation, such that 'H-FoV / Med' denotes high field of view with medium ornamentation. The three studies recruited separate participants and are not pooled; fitted lines are descriptive rather than inferential. Focused attention is used as the engagement measure throughout for comparability, as chapter 6 did not administer the Film-IEQ.

7.2.3 How Immersion and Awareness Vary Together

While fig. 7.1 shows engagement and awareness moving in opposite directions across most conditions, two further findings indicate that the relationship between them is neither strictly determined nor uniform in its consequences.

The first is that the two measures do not move in lockstep. VR Cinema and VR Videosphere in chapter 6 differed in focused attention while producing identical reported awareness, and the two conditions noted above moved in the same direction on both measures rather than in opposite directions. These effects are small, and neither was statistically reliable, but they indicate that knowing how immersive a configuration is does not determine how much awareness it leaves the passenger. Immersion and environmental awareness were therefore measured separately throughout this thesis, and are treated as two outcomes of the design decisions set out in section 7.3 rather than as two ends of a single scale.

The second is that a difference in reported awareness does not establish its behavioural consequences. Chapter 6 identified significant differences in self-reported awareness across conditions, but the stop-detection task produced only eight failures across 114 trials and could not detect a corresponding difference in performance. That null cannot be read as evidence that the reductions were inconsequential, and the task is better understood as too coarse to answer the question. What the design of the conditions does establish is that every condition, including those that occluded the carriage entirely, left intact the auditory channel through which station announcements were delivered; indeed, in the fully occlusive conditions that channel was the only route by which journey information was available. Which perceptual channels a design leaves available is therefore a separate question from how much awareness it removes overall, a distinction developed further in section 7.2.4.

In interviews, several participants described shifting their attention between the media and their surroundings during a condition, rather than holding a fixed balance throughout; in chapter 6, for example, one participant described deciding part-way through that they needed to start listening for their stop (section 6.3.3). ARTV Ornamentation offers one way of supporting such shifts by embedding media augmentation within the physical passenger environment. Rather than replacing the real-world environment, ornamentation integrates digital narrative elements into existing surroundings, so that users retain visual and spatial reference to the carriage while engaging with the content. This may reduce the need to remove the HMD in order to re-establish awareness, although this was not tested.

7.2.4 The Composition of Environmental Awareness

Treating environmental awareness as a single undifferentiated quantity obscures three distinctions that emerged across the research programme. These are the distinctions the studies and workshops surface, and they are not presented as an exhaustive decomposition of the construct.

The first concerns the objects of awareness. Chapter 6 measured social awareness and wayfinding separately, through the passenger-awareness item and the stop-detection task respectively, and participants themselves treated these as distinct concerns: several described wanting to know who was nearby in case they were spoken to or needed to move a bag, while others were concerned principally with not missing their stop, and a minority wanted neither. What a passenger is aware of is therefore not a single undifferentiated field. Workshop participants had anticipated this distinction, one describing a wish to be 'aware of them, but you don't actually necessarily have to see the people' while continuing to know 'where you are' (W1-2), and identifying personal belongings as a further category requiring separate treatment. Custodial awareness of belongings was raised by participants but was not measured in any of the studies reported here.

The second concerns modality. A design that occludes vision does not necessarily remove awareness in aggregate, and in chapter 6 the conditions that removed visual access to the carriage nonetheless left auditory access intact, with the safety-relevant information reaching participants through that route. This should be interpreted with care, as those conditions also presented station information by audio alone, so reliance on listening was in part required by the design rather than freely adopted. Participants nonetheless described that reliance as effortful, which indicates that the substitution was not cost-free. Workshop participants had anticipated this consideration, proposing that environmental audio be dampened rather than silenced and observing that physical sensation persists irrespective of what a display presents (W1-3, W2-3). The design question is consequently not only how much awareness a given configuration costs in aggregate, but which perceptual channels it leaves available and what it costs to use them.

The third concerns degree. Chapter 4 manipulated ornamentation at three levels, and at the wide field of view reported awareness declined across them rather than switching between pres-

ence and absence, a pattern visible in fig. 7.1, where the high-FoV conditions move stepwise through the space as ornamentation increases. At the restricted field of view the same manipulation produced no significant change, so gradation in the design does not guarantee gradation in the outcome. Participants in the same study ranked the intermediate level as highly as the most intense one and chose it most often for use on a train. Rankings for these two levels did not differ significantly, so this does not establish a preferred intermediate point, but it is consistent with passengers valuing a degree of awareness rather than treating it as all-or-nothing. Workshop participants had proposed graded mechanisms on the same reasoning, suggesting that other passengers might be rendered partially transparent such that 'you were aware they're there but it doesn't block your scenery anymore' (W1-1), and rejecting complete removal as disconcerting, since a person removed visually would remain audible and physically present (W2-3).

7.2.5 Contextual Variation in Awareness Demands

The region of this space that a design may acceptably occupy is not fixed, but depends upon the demands that a passenger's situation places upon them. Passenger environments are inherently dynamic, with changing social configurations, unpredictable movement, and varying levels of safety risk, and awareness demands are often higher during boarding, disembarking, or crowded travel conditions, whereas longer or more stable journey phases may allow deeper immersive engagement.

Workshop participants distinguished several such conditions without prompting. Journey duration was identified as decisive, with one participant observing that immersive experiences 'might work on sort of longer rides, or like a plane ride', whereas on 'a 20 minute commute then you would want sort of the blurred setting because you need to be aware of what's going on and you can't really sink into it' (W2-2). The same participant treated physical enclosure and custodial exposure as a separate consideration, noting that on a flight 'I'm not as worried about somebody stealing my stuff, for example — someone will not like walk off the flight' (W2-2). Crowding and posture were raised independently by a participant in the other workshop, who identified the case of 'sitting in a middle aisle and everyone's standing around you' (W1-1), while others reasoned about gating immersive use by setting altogether, suggesting that full virtuality might be permissible 'only in a safe environment' (W1-3).

These conditions, comprising journey duration, the density of disembarkation decisions, physical enclosure, crowding and posture, and social exposure, differ in kind rather than solely in degree. Standing alters which perceptual channels are available to the passenger, enclosure alters the purpose that custodial attention serves, and the frequency of disembarkation decisions alters how often attention must be reclaimed from the media experience. Each therefore calls for a different design response rather than simply more or less of the same one.

One of these conditions has within-study support. In chapter 6, the number of members of the public present in the carriage varied naturally across trials, and including that count in a model of

reported passenger awareness improved its fit, with participants reporting greater awareness of other passengers when more people were present. The surrounding social environment therefore appears to shape where within the space a given design is experienced as sitting, although social density was not manipulated and this remains an association rather than a demonstration.

The remaining conditions rest on participants' reasoning alone, as none was manipulated in the studies reported here, each of which was conducted with seated participants on journeys of comparable character, who were neither standing, nor crowded, nor responsible for luggage. The contextual component of the trade-off is accordingly presented to motivate investigation rather than as a finding of the present work. These conditions are also properties of the journey rather than of the passenger, and chapter 5 indicates that individual disposition varies independently of them, with one participant rejecting ornamentation in any form.

7.2.6 Relation to Established Models

The PIAT is positioned in relation to existing accounts of immersion, presence, and mixed reality rather than in competition with them.

Presence-based frameworks emphasise perceptual realism and the extent to which a user responds to a virtual environment as though located within it [192, 165]. These models have largely been developed in stationary or controlled settings and focus primarily on perceptual immersion within virtual environments, where withdrawal from the surrounding physical environment is treated as an indicator of success rather than as a cost. The passenger context inverts this assumption, in that the surrounding environment must remain perceptually available throughout.

As discussed in chapter 2, Milgram and Kishino [121] proposed both the Reality–Virtuality continuum and a three-dimensional taxonomy of mixed reality displays, of which the field largely adopted the former while setting aside the latter [163]. Skarbez et al. revise that taxonomy into three dimensions, comprising extent of world knowledge, immersion, and coherence, while observing that all three describe properties of systems and that the corresponding user-side constructs remain undeveloped. They note in particular that established accounts of presence contain no notion of extent of world knowledge, that little research has addressed the dimension since Milgram's original papers, and that no constructs have been proposed which combine it with immersion or coherence.

Extent of world knowledge is the dimension of most direct relevance to passenger ARTV, as it describes how much a system knows about, and is able to respond to, the real environment in which it operates. In a passenger context this determines the extent to which augmentation can be registered to the carriage rather than merely displayed within it. The Spatial Integration dimension developed in section 7.3 describes this same property. The conditions compared in chapter 6 differed in it, in that ARTV Ornamented registered content to the surfaces of the physical carriage whereas VR Videosphere presented a themed panorama that took no account

of the carriage's geometry; this difference followed from how the conditions were constructed rather than from a deliberate manipulation of registration, and participants did not control where content was placed in any study. That these conditions produced comparable engagement while differing substantially in retained awareness is consistent with world knowledge being consequential. However, as they also differed in whether object-level augmentation was present and in whether the surrounding environment was preserved or replaced, this comparison should be interpreted as indicative rather than as isolating the contribution of world knowledge.

The account developed in this thesis concerns experiences that unfold alongside ongoing real-world demands, incorporating the influence of mobility, social context, and environmental variability on how attention is distributed. It should be noted that mobility itself was not examined empirically, as no study involved a moving vehicle; what the studies address is the shared, socially exposed, and journey-structured character of passenger travel rather than motion as such. In doing so, the account treats passenger XR as a setting with design requirements of its own, in which designs must allow engagement and awareness to be balanced rather than maximising either in isolation. The contribution of the PIAT in relation to existing models is therefore to connect a system dimension that those taxonomies identify, but leave unattached, to a user-side outcome with consequences in the passenger context. It does not propose a further system dimension, and it does not extend presence theory, which addresses a different question.

7.2.7 The Status and Contribution of the PIAT

The PIAT is a design concept, and specifically a strong concept in the sense described by Höök and Löwgren [63]: intermediate-level knowledge which is more abstract than any particular design instance, but which does not aspire to the generality of a theory, and which is assessed according to its generativity, or its capacity to be taken up in the design of new instances. It is not a predictive model, as it specifies no functional form relating immersion to awareness and this thesis does not claim one, and neither is it a measurement instrument nor an account of presence.

Four contributions are claimed, each of which stands independently of the others.

The first is the characterisation of the relationship itself. That immersive engagement and environmental awareness stand in tension is documented in prior work [17, 188, 153], but is generally expressed as a single trade-off between comparable quantities. Treating them instead as two separately measurable outcomes arising from a common set of design decisions accommodates findings that a single quantity is unable to represent, and makes available the distinctions of object, modality, and degree set out in section 7.2.4.

The second is the connection between extent of world knowledge and retained environmental awareness. Skarbez et al. [163] identify extent of world knowledge as a dimension of mixed reality systems that has received little attention since Milgram and Kishino [121] proposed it, and note that no constructs have been advanced which combine it with immersion or coherence.

The PIAT addresses the user-side half of that gap, in that the degree to which a system registers augmentation to the physical environment, rather than displaying it independently of that environment, bears directly on what a passenger retains of their surroundings. The passenger context makes this connection consequential in a way that domestic settings do not, since what a user retains of a living room carries no comparable cost.

The third is positioning. The PIAT permits designs to be located with respect to the trade-off rather than merely described as subject to it, as fig. 7.1 demonstrates for eighteen conditions across three studies. Whereas prior work establishes that reintroducing awareness into VR incurs a cost to immersion [17], the account developed here distinguishes designs according to what they leave available to the passenger, which is a separate question.

The fourth is contextual variation. The acceptable region depends upon properties of the journey, and those properties differ in kind rather than in degree, so a framing that treated context as a single threshold would be unable to distinguish between them. This contribution rests on participant reasoning and, for social density, on a within-study association rather than on experimental evidence.

The dimensions identified in this section are those which the research programme surfaced. They are not offered as exhaustive, and the account is intended to be extended as further parameters, modalities, and contexts are examined. This openness reflects the kind of contribution the PIAT represents rather than a deficiency in it, as a strong concept is assessed according to whether it generates new instances rather than whether it enumerates all of them [63].

Conditions Under Which the Account Would Not Hold

The PIAT advances claims that are capable of failing. Were the manipulation of the immersive properties of a passenger media experience to produce no systematic difference in the balance between engagement and environmental awareness, the central relationship would be disconfirmed. Were awareness to prove inseparable by object or by modality, such that occluding one perceptual channel reduced awareness uniformly irrespective of which other channels remained available, the decomposition set out in section 7.2.4 would be incorrect; testing this requires a behavioural measure sensitive enough to detect the difference, which the stop-detection task used in chapter 6 was not. Were passengers' preferred position to prove invariant across contexts differing substantially in their awareness demands, the contextual component would be disconfirmed. This represents the claim most in need of testing. Testing it requires journeys in which these conditions vary of their own accord and can be recorded as they occur, rather than the uniform seated conditions under which the present studies were conducted, and it is therefore pursued most readily through the in-the-wild deployment considered in section 7.6.2.

7.2.8 Applying the PIAT to Unstudied Contexts

The value of a strong concept lies in its application to cases beyond those from which it was derived [63]. Three passenger contexts, none of which was studied in this thesis, illustrate the reasoning that the account supports. The conclusions drawn are illustrative rather than empirical.

A rush-hour underground service with standing room only imposes high demands on each of the contextual conditions identified in section 7.2.5. Disembarkation decisions arrive every few minutes, the passenger is standing and must maintain balance in a manner that commits proprioceptive attention, crowding restricts sightlines and requires frequent physical accommodation of others, custodial exposure is elevated by the continuous boarding and alighting of other passengers, and social density remains high throughout the journey. The PIAT suggests that acceptable designs occupy a narrow region under these conditions, in that augmentation should be sparse, should not occlude the sightlines through which a standing passenger monitors both signage and other people, and should avoid committing the auditory channel, which in a crowded and noisy carriage is already degraded. Fully occlusive designs are excluded altogether.

An off-peak suburban rail journey undertaken from a seat, with stops at intervals of ten to twenty minutes, presents a substantially different profile. Disembarkation decisions are infrequent and predictable, posture is stable, crowding is low, and custodial exposure is moderate. The acceptable region is correspondingly wider, and the PIAT suggests that moderate ornamentation is supportable, which is consistent with the preference pattern observed in chapter 4 under comparable conditions and with workshop participants' reasoning about longer and quieter journeys (W2-2).

A long-haul flight presents a third profile, in which disembarkation is enforced and singular, posture is fixed for extended periods, the cabin is enclosed such that custodial exposure is minimal, a consideration workshop participants raised unprompted (W2-2), and journey duration is sufficient that sustained absorption is both feasible and desirable. Almost the entire space is available under these conditions, and fully occlusive VR becomes defensible, which is consistent with the commercial positioning of such systems for in-flight entertainment [188].

These cases differ in kind rather than solely in the quantity of awareness each requires. A framing that treated context as a single threshold would be unable to distinguish between them, and the design responses each context calls for differ accordingly.

7.3 A Design Framework for ARTV in Passenger Contexts

Building on the conceptualisation of the PIAT, this section formalises how ARTV experiences can be designed for passenger contexts. While the PIAT describes the underlying experiential tension between immersive engagement and environmental awareness, it does not by itself provide a structured account of the design decisions through which ARTV systems can respond to that tension.

This thesis argues that ARTV cannot be fully understood through existing design frameworks when applied to passenger contexts. Prior ARTV frameworks [152] primarily characterise augmentation in terms of content, interaction, and presentation dimensions, but do not explicitly account for the attentional, social, and environmental constraints inherent to mobile, shared transport settings.

Across the empirical studies presented in this thesis, ARTV Ornamentation was shown to affect the balance between immersive engagement and environmental awareness, formalised here as the PIAT. These findings indicate that ARTV design decisions are not purely representational: they influence how users allocate attention between media content and their surroundings.

This section therefore reconceptualises ARTV as a structured design space for responding to the PIAT and proposes a unified framework for ARTV in passenger contexts. Building on the framework proposed by Saeghe et al. [152], this work introduces a set of design dimensions that characterise ARTV experiences in passenger contexts more broadly, rather than limiting the framework to specific implementations such as ornamentation. While ARTV Ornamentation is used throughout this thesis as the primary instantiation of these dimensions, the framework may also generalise to other forms of ARTV, including minimally augmented viewing, spatial display-based configurations, and hybrid approaches. These extensions capture design considerations specific to mobile and shared passenger environments, including *Content Placement*, *Bystander Integration*, and *Confined Space Interaction*, as well as the role of ARTV Ornamentation as a distinct form of spatial augmentation. In doing so, the section formalises ARTV as a design space structured around the immersion–awareness trade-off, defined through a set of core dimensions that describe how immersive media can be integrated with the physical environment, while preserving environmental awareness in passenger contexts.

7.3.1 ARTV as a Design Space for Responding to the PIAT

In passenger contexts, ARTV should not be understood solely as a technique for extending media content into the surrounding environment. Rather, it acts as a design layer through which immersive media is integrated with the passenger's surrounding physical environment. The significance of ARTV in this thesis therefore lies not only in what is augmented, but in how augmentation reshapes the user's attentional relationship with their surroundings.

The empirical studies presented in this thesis demonstrate that ARTV experiences can increase engagement, relative to baseline AR viewing, while preserving greater environmental awareness than fully immersive virtual reality. However, these outcomes are not uniform. The extent to which ARTV supports engagement or awareness depends on how augmentation is designed, including its intensity, spatial configuration, relationship to content, and responsiveness to context. On this basis, ARTV is more appropriately conceptualised as a design space defined by a set of interacting dimensions that influence how the PIAT is negotiated. This section defines five such dimensions: *Ornamentation Intensity*, *Spatial Integration*, *Content Coupling*, *Adap-*

tivity, and *Awareness Preservation*. These dimensions do not operate independently; rather, they jointly shape how immersive media is integrated with the physical environment and how attention is distributed between media content and real-world context.

Three further considerations extend the framework rather than standing alongside these five as dimensions in their own right. *Content Placement* and *Confined Space Interaction* elaborate Spatial Integration, addressing respectively where augmentation is positioned and what interaction the physical constraints of a carriage permit; *Bystander Integration* elaborates Awareness Preservation, addressing how co-present others are treated within the augmented scene. Each is defined alongside its parent dimension below.

The subsections that follow define these design dimensions and describe how they extend existing ARTV frameworks by incorporating constraints and requirements specific to mobile, shared passenger environments.

7.3.2 Dimension 1: Ornamentation Intensity

Ornamentation Intensity refers to the quantity, visual prominence, and experiential dominance of augmentation within the user's environment. This includes both the amount of AR content present and the degree to which that content competes with, complements, or recedes behind the primary media experience.

Findings from the controlled and ecological evaluations (chapters 4 and 6) demonstrate that ornamentation intensity is a primary driver of immersive engagement in passenger ARTV, although its effects are conditional. In chapter 4, increasing ornamentation raised reported immersion, engagement, and enjoyment only when the field of view was wide; at the restricted field of view the same manipulation produced no significant change. Where ornamentation did raise engagement it also lowered self-reported awareness of other passengers, indicating attentional competition between media content and the surrounding carriage. In chapter 6, ornamentation raised hedonic quality in both the physical and simulated settings without a detectable effect on self-reported awareness, but lowered pragmatic quality and, marginally, perceived usability. Taken together, these results establish ornamentation intensity as a key control dimension within the PIAT, whose costs fall on awareness, on perceived practicality, or on both, depending on the configuration and the display.

Within this dimension, ornamentation takes several forms that plausibly differ in visual salience and spatial footprint: discrete 3D objects, the re-texturing or transformation of environmental surfaces, and atmospheric or particle-based effects. Their relative contribution to perceived intensity is a design consideration rather than an empirical finding, as no study reported here isolated them. Chapter 4 found that adding retexturing to objects produced no further gain in immersion beyond a rise in aesthetic appeal, and chapter 5's strongest condition combined retexturing, particles, and timed transitions as a single bundled treatment.

Rather than representing a binary design choice, ornamentation intensity should be under-

stood as a graded parameter that can be adjusted to balance immersive engagement with environmental awareness. Where along that parameter passengers place themselves is itself informative: participants in chapter 4 ranked the intermediate level alongside the highest level and chose it most often as an alternative to a phone or laptop, yet at the wide field of view that same level reduced awareness of other passengers significantly relative to the unornamented AR baseline. The preferred configuration and the awareness-preserving configuration are therefore not the same, and effective ARTV design requires calibrating intensity against both rather than assuming that satisfying one satisfies the other.

7.3.3 Dimension 2: Spatial Integration

The second dimension concerns how ARTV Ornamentation is spatially integrated within the user's environment. While prior ARTV frameworks account for where augmentations may appear relative to the display, they do not explicitly address the implications of spatial placement in mobile, shared passenger contexts. This dimension is derived principally from prior literature rather than from manipulation within this thesis, and is included because the studies reported here bear on it indirectly.

In such environments, spatial integration becomes a critical design consideration due to the presence of other passengers, limited personal space, and the need to maintain awareness of the surrounding environment. Prior work highlights that the positioning of virtual content is influenced by both environmental and social factors. For example, participants in previous studies have expressed preferences for placing virtual content in ways that minimise distraction or avoid obstructing nearby passengers whose presence, movement, or potential interaction may require attention [110]. These findings suggest that spatial placement is not purely a technical concern, but is closely tied to social context and environmental awareness.

The studies reported here bear on this dimension without manipulating it, as each presented fixed configurations of ARTV Ornamentation and left placement outside participant control. The conditions did nonetheless differ in the degree to which augmentation was registered to the physical environment: in chapter 6, ARTV Ornamented mapped textures and objects onto the surfaces of the real carriage, whereas VR Videosphere surrounded the viewer with imagery that took no account of the carriage at all. The comparison is suggestive rather than conclusive: ARTV Ornamented and VR Videosphere did not differ on any questionnaire measure other than awareness of other passengers, which ARTV Ornamented preserved better, although the two conditions also differed in occlusion and in object-level augmentation (section 7.2.6). Registration also carried a cost specific to this dimension: because augmentation was mapped onto the physical carriage, registration errors could place it in front of people, as happened in chapter 6.

Spatial integration therefore provides one route through which ARTV designs can negotiate the PIAT: the placement and extent of augmentations within the user's field of view shape how attention is distributed between media content and the real environment. Because placement

was predefined in every study, this work says nothing about how users would position or adapt content themselves.

Based on these findings and prior literature, this thesis extends the ARTV framework to include *Content Placement* as an explicit design consideration. This captures decisions regarding where augmentations and displays are positioned within the user's environment, including their proximity, occlusion relationships, and alignment with physical space. Additionally, this thesis introduces *Confined Space Interaction* to account for the constraints imposed by limited physical space in passenger contexts. Prior research has shown that interaction techniques suitable for home environments may not translate effectively to constrained settings due to risks of collision or intrusion into the personal space of others [111]. As such, ARTV Ornamentation in passenger contexts should prioritise minimal or low-amplitude interaction, favouring passive or lightweight engagement over physically demanding input.

Together, these extensions reposition spatial integration as a central component of ARTV design in passenger contexts. Rather than treating placement as a secondary implementation detail, it becomes a key factor in managing the balance between immersion and awareness within the PIAT.

7.3.4 Dimension 3: Content Coupling

Content Coupling refers to the relationship between ARTV Ornamentation and the media content being consumed. Specifically, it captures the extent to which augmentation is aligned with the media at different levels of specificity, ranging from loosely related environmental decoration to tightly integrated, content-specific augmentation. This includes distinctions between broader genre-level augmentation, scene- or biome-level augmentation, and augmentation tied to specific media properties, objects, settings, or narrative events. These distinctions reflect different levels of abstraction in how ornamentation is designed.

The design exploration presented in chapter 5 demonstrates that ARTV Ornamentation can be constructed at different levels of abstraction, for example reflecting broader genre characteristics or elements specific to a given piece of media content. Content coupling is therefore a choice about what the augmentation refers to: the genre, the setting of the current scene, or objects specific to the programme. The findings of that study were more qualified than the distinction itself might suggest. Show-specific objects produced no detectable advantage over generic genre-level objects on any measure, which is not the same as establishing that the two are equivalent; genre-level ornamentation improved on the unornamented baseline only in aesthetic appeal; and the condition that changed the surrounding environment with the scene outperformed both, while being the only condition to reduce reported awareness of other passengers. Participants were nonetheless sensitive to how well augmentation aligned with the media, and incongruity between the two was raised spontaneously by nine participants.

When considered alongside the broader findings of this thesis, content coupling can be un-

derstood as influencing the coherence between the media content and the surrounding augmentation, a relationship developed further in section 7.4. Augmentation that is strongly coupled to scene, biome, or content-specific elements is more likely to be experienced as an extension of the media, whereas weakly coupled augmentation may be perceived as generic decoration or visual distraction. In combination with dimensions such as ornamentation intensity and spatial integration, the degree of coupling may therefore affect both immersive engagement and the distribution of attention between media and environment. In light of this influence on the PIAT, this thesis includes *Content Coupling* as a dimension within the ARTV design space, capturing how augmentation relates to media content at different levels of abstraction rather than treating augmentation as independent of the media it accompanies.

7.3.5 Dimension 4: Adaptivity

Adaptivity refers to the extent to which ARTV Ornamentation changes over time rather than remaining fixed throughout an experience. In passenger ARTV, this can occur in two related ways: media-driven adaptation, where ornamentation changes in response to the content being watched, and context-aware adaptation, where ornamentation changes in response to the passenger environment or the user's attentional demands. The studies presented in this thesis examined media-driven adaptation only, while the findings motivate context-aware adaptation as a direction for future ARTV systems.

In chapter 5, ornamentation was implemented using predefined transitions that reflected changes in media content, allowing augmentation to shift between distinct thematic states over time. These transitions were authored as part of the experimental conditions, rather than dynamically generated or user-driven, and formed part of a bundled treatment that also included environmental retexturing and particle effects. That condition produced the strongest engagement and preference of the four, but because its components were not varied independently, the contribution of temporal change specifically cannot be isolated.

The evaluation presented in chapter 6 motivates the extension from media-driven change to context-aware adaptation. Several participants there described shifting attention between the media and the carriage according to what the environment demanded at a given moment, and reported doing so with greater effort in the more occlusive conditions. If awareness demands vary within a journey in this way, then an ornamentation configuration that is appropriate at one point may not be appropriate at another. This suggests that future ARTV systems may need to adapt not only to the temporal structure of the media, but also to changing awareness demands during travel, such as crowding, movement in the surrounding environment, or phases of the journey requiring greater monitoring.

Adaptivity is therefore included as a dimension within the ARTV design space because temporal change can influence both the perceived intensity of augmentation and the distribution of attention between media and environment. By distinguishing between static, media-responsive,

and context-responsive ornamentation, this dimension highlights how change over time may shape immersive media experiences in passenger contexts and contribute to the ongoing negotiation of the PIAT.

7.3.6 Dimension 5: Awareness Preservation

Awareness Preservation refers to the extent to which ARTV experiences support the user's ability to remain aware of their physical and social environment while engaging with immersive media. This includes awareness of surrounding passengers, environmental changes, and journey-related information, all of which are critical in passenger contexts.

Across the research programme, the importance of maintaining environmental awareness is consistently reflected in both quantitative and qualitative findings. Chapter 4 showed that ornamentation at a wide field of view reduced self-reported awareness of other passengers relative to unornamented viewing, and chapter 5 found the same for its scene-adaptive condition. Chapter 6 showed that both pass-through conditions preserved substantially more reported awareness than either fully immersive alternative, while participants described monitoring the journey as more effortful once the carriage was occluded. These findings directly underpin the PIAT, establishing awareness preservation as a central requirement for ARTV design in passenger contexts.

While earlier design dimensions describe how augmentation is constructed and presented, awareness preservation captures how ARTV experiences remain compatible with real-world passenger demands. The results of this thesis indicate that users do not disengage entirely from their environment during immersive media consumption, but instead distribute attention between media content and surrounding context. This suggests that ARTV experiences must be designed to support ongoing attentional negotiation, rather than attempting to maximise immersion at the expense of awareness.

Prior work similarly shows that social and environmental factors shape how XR systems are used and perceived in shared settings. The presence of other passengers introduces considerations around visibility, occlusion, personal space, and social appropriateness, which influence how augmentation should be designed and presented [110, 16]. This thesis therefore extends the ARTV framework to include *Bystander Integration*: the treatment of co-present others within the augmented scene, spanning designs that leave bystanders wholly unmodified, those that attenuate or partially conceal them, those that replace them with representations, and those that actively signal their presence or proximity to the viewer. Bystander Integration was not manipulated in any study reported here, and is included on the basis of prior literature and of participants' own reasoning: workshop participants in chapter 3 proposed partial transparency in preference to removal, and participants in chapter 6 objected specifically to augmentation that occluded people in the carriage.

The findings from this thesis also indicate substantial variation in user preference regarding acceptable levels of augmentation, suggesting that no single configuration optimally balances

immersion and awareness for all users. One way of enacting awareness preservation within the ARTV design space may therefore be through personalisation or lightweight user control, allowing passengers to adjust augmentation density, visibility, or prominence to match their individual comfort and awareness requirements, whether directly or through automatic modulation of ornamentation level. While this thesis does not implement adaptive or user-controlled systems, these observations suggest that awareness preservation should be treated as a configurable property of passenger immersive media experiences rather than as a fixed threshold. This dimension therefore captures the requirement for ARTV experiences to support continued engagement with the physical and social environment. In doing so, it positions awareness not as an external constraint on immersion, but as a co-existing component of passenger immersive media experiences.

7.3.7 Implications of the Framework

The five dimensions presented in this section define a structured design space for ARTV in passenger contexts. Rather than representing independent features, these dimensions describe interdependent aspects of how immersive media can be integrated with the physical environment while maintaining environmental awareness. A key implication of this framework is that ARTV design should not be understood in terms of maximising individual properties such as immersion or visual fidelity. Instead, effective ARTV experiences depend on balancing multiple interacting factors, including augmentation intensity, spatial configuration, content alignment, temporal behaviour, and awareness requirements. As section 7.2.3 establishes, immersion and environmental awareness are separate outcomes rather than two ends of a single scale, and these factors shape how each is affected by a given configuration.

The framework is not restricted to ornamentation-based ARTV experiences. At one extreme, fully immersive VR media experiences such as virtual cinemas or videospheres can be understood as configurations with high immersive enclosure, low awareness preservation, and minimal integration with the physical passenger environment. At the other, lightweight AR configurations such as floating displays with minimal environmental augmentation preserve awareness but provide weaker spatial extension of the media. ARTV Ornamentation, as explored in this thesis, occupies an intermediate region in which augmentation is integrated with the physical environment to support both immersive engagement and environmental awareness.

The awareness measures collected across the research programme support this positioning. In chapter 6, awareness of the augmentation layer was high in every condition that supplied one, and ARTV Ornamented was the only such condition in which participants also reported continued awareness of other passengers: VR Videosphere achieved comparable awareness of its surrounding imagery with substantially lower awareness of the carriage, while the unornamented AR baseline preserved awareness of the carriage but offered no augmentation layer to attend to. The distinction between integrating augmentation with the passenger environment and replacing that environment is therefore visible in what passengers reported attending to, and not only in

the design intent behind each condition.

The framework also extends prior ARTV models by incorporating the constraints of mobile and shared passenger environments. By introducing dimensions such as spatial integration, adaptivity, and awareness preservation, it accounts for factors that are critical in passenger contexts but are not fully addressed in existing ARTV frameworks. This makes the framework useful not only for analysing the studies presented in this thesis, but also for guiding future research and design. By articulating key dimensions of passenger ARTV experiences, it supports systematic exploration of how different configurations affect user experience, and how XR systems might be adapted to varying contexts, content types, and user preferences. This positions ARTV not simply as a form of augmented media, but as a context-sensitive design problem in which immersive experiences must be compatible with real-world behavioural and environmental demands.

7.3.8 ARTV Ornamentation as the Primary Framework Instantiation

Within this framework, ARTV Ornamentation is the implementation developed and evaluated in this thesis. Rather than constructing a fully authored virtual scene or presenting media as a detached overlay, it embeds media-related visual elements into the passenger's surroundings: atmospheric effects, vegetation, surface transformations, or scene-matched objects anchored to carriage walls, windows, seats, and other surfaces. This allows the passenger environment to become part of the media experience while remaining recognisable as itself, so users experience the media as spatially extended without fully losing environmental reference.

The three empirical studies instantiate several of these dimensions directly: Ornamentation Intensity in chapter 4, and Content Coupling and Adaptivity in chapter 5, through the scene- and biome-level transitions described above, though as section 7.4 discusses, a static biome can become incongruent as the narrative moves on. Density is also the parameter to which the journey-context reasoning of section 7.2.5 and the design implications of section 7.5 apply. In this sense, ornamentation shifts the design challenge away from constructing complete virtual worlds and toward developing reusable, flexible augmentation systems that can be scaled across content, contexts, and passenger needs; scaling it beyond hand-authored assets toward automated or generative pipelines is considered in section 7.6.2.

7.4 Coherence, Plausibility, and the Authoring of Augmentations

The five dimensions describe what a passenger ARTV system presents and how much of the environment it leaves available. They do not by themselves address whether what is presented holds together, either as a depiction of a world or as something that could plausibly be occurring

in the carriage the passenger occupies. This section considers that question, together with the practical matter of who authors augmentations and how they are produced at scale, and the social conditions under which they would be used repeatedly.

7.4.1 Coherence and Plausibility in Passenger ARTV

Skarbez et al. [163] identify coherence as one of the three dimensions of their revised mixed reality taxonomy, describing the extent to which the elements of an environment behave in a manner consistent with one another and with the user's expectations. In their earlier survey, Skarbez et al. [162] distinguish this from immersion, arguing that a system may be highly immersive yet incoherent, and that coherence is what supports plausibility illusion (the sense that events occurring within an environment are really happening), as distinct from place illusion. This distinction is directly relevant to passenger ARTV, where augmentation is not presented within an authored virtual world but within a real carriage that the passenger can see, hear, and feel throughout.

Two findings from this research programme bear on coherence. The first concerns coherence between augmentation and media. In chapter 5, nine participants raised incongruity between the ornamentation and the content on screen, most directly where the ornamentation displayed the iconography of one house while the scene depicted another, a mismatch that P20 identified explicitly. The augmentation in that condition was no less immersive in construction than in others; what it lacked was consistency with the narrative it accompanied. The second concerns coherence between augmentation and the physical environment. In chapter 6, retextured surfaces were at times drawn in front of or through people in the carriage, and the augmentation shifted or flickered against the real geometry. Participants tolerated a lower-fidelity augmented environment but were disrupted by inconsistency, which drew attention toward the system rather than toward either the media or the carriage.

These are distinct failures with a common structure. In each case the augmentation was rendered as specified, and in each case it failed to behave as the surrounding context implied it should. Coherence in passenger ARTV is therefore doubly constrained, in that augmentation must remain consistent both with the media it extends and with the physical and social environment into which it is placed, and neither constraint can be satisfied by increasing visual fidelity alone. This also indicates a limit to what the Content Coupling dimension can deliver on its own: coupling determines what an augmentation refers to, whereas coherence determines whether what it refers to is sustained as the content and the carriage change around it.

7.4.2 Authoring Augmentations at Scale

The augmentations evaluated in this thesis were authored by the researcher for specific clips, which is not a viable production model. Chapter 5 addresses this by examining ornamentation

authored at different levels of abstraction, and its findings bear directly on who could produce such content and how. Show-specific objects produced no detectable advantage over generic genre-level objects, which suggests that reusable libraries of genre or setting assets may be sufficient for a substantial part of the design space, without requiring bespoke authoring per programme. This should be read with care, as the absence of a detectable difference is not equivalence, and the genre-level condition improved on unornamented viewing only in aesthetic appeal.

Where augmentation must track the content more closely, as the scene-adaptive condition did, authoring becomes a question of who holds the necessary knowledge of the media. Programme makers are positioned to supply scene and setting metadata as part of production, in the way that subtitle and audio description tracks are already produced, and this would allow ornamentation to be triggered by the content itself rather than inferred from it. This is a design implication rather than a finding of this thesis, and it raises questions of editorial control that the studies reported here do not address: an augmentation layer authored by a broadcaster carries its intent into the passenger's physical space, where a viewer may reasonably expect to control what appears around them.

The scalability question also interacts with coherence. Reusable genre libraries reduce authoring cost but increase the risk of the mismatch identified in chapter 5, since a generic medieval asset cannot distinguish between the factions or settings that a viewer familiar with the programme will recognise. Automated or generative approaches to this problem are considered in section 7.6.2.

7.4.3 Social Effects and Adoption Over Time

Passenger ARTV is used in the presence of others, and the studies reported here touch on this in two ways. Participants in chapter 6 were conscious of how the headset would appear to other passengers, raising both its visibility and its value as a target for theft, and one visitor to the carriage approached a participant mid-condition without realising they could be seen. Participants in chapter 5 varied considerably in their disposition towards augmentation as such, with one rejecting it in any form, which indicates that acceptability is not solely a property of the design.

Neither the social reception of ARTV by co-present passengers nor its acceptability over repeated use was measured in this thesis, and both are consequential for whether such systems would be adopted. The relevant directions for future work are set out in section 7.6.2, which considers bystander perception and longitudinal use together, since the norms governing conspicuous device use in shared transit spaces are established socially rather than individually.

Table 7.1: Relationship between the framework dimensions (section 7.3) and the design implications below, with the primary source of evidence for each.

Design implication	Dimensions	Principal evidence
Graduated immersion rather than binary modes	Ornamentation Intensity; Awareness Preservation	Chapter 4 (three ornamentation levels; intermediate level ranked alongside the highest and most often chosen for use)
Dynamic and context-aware augmentation	Adaptivity; Awareness Preservation	Chapter 5 (media-driven transitions); chapter 6 (interviews)
Preserve environmental reference through spatial integration	Spatial Integration; Awareness Preservation	Chapter 6 (pass-through vs occlusive conditions)
Design XR media as socially situated experiences	Awareness Preservation (Bystander Integration)	Prior literature; chapter 3 workshops; chapter 6 interviews
Prioritise perceptual stability and environmental consistency	Spatial Integration; Content Coupling	Chapter 6 (registration artefacts); chapter 5 (incongruity)
Support flexible attention management	Awareness Preservation; Adaptivity	Chapter 6 (interviews; reported awareness)

7.5 Design Implications for Passenger XR Systems

The findings of this thesis provide a set of design implications for immersive XR media systems intended for passenger contexts. Public transport environments present unique interaction constraints characterised by mobility, shared social space, and the requirement to remain aware of the surrounding environment. The empirical and conceptual contributions of this thesis indicate that immersive passenger XR systems must be designed to support dynamic attentional negotiation, rather than maximising immersion as a singular objective. The following design implications synthesise insights derived from the research programme and provide guidance for the development of XR media experiences in mobility-driven environments. Table 7.1 relates each implication to the framework dimensions from which it follows and to the evidence supporting it.

7.5.1 Design for Graduated Immersion

The Ornamentation Intensity dimension (section 7.3) established immersion as a graded property rather than a binary state: intermediate levels of ornamentation were ranked as highly as the most intense level and chosen most often for use on a train, even though the most intense level did not preserve awareness at the level of unaugmented viewing. Passenger XR systems should therefore avoid rigid transitions between fully immersive and fully reality-focused modes.

Designers should prioritise XR media systems that support graduated levels of immersion through adjustable augmentation density, visual complexity, and spatial extent. This may include layered ornamentation strategies in which components are adjusted separately, for example reducing atmospheric or particle effects while retaining 3D objects, removing foreground objects while preserving low-salience retexturing, or limiting augmentation to peripheral surfaces when greater awareness is required; and adaptive systems that reduce augmentation as the passenger approaches their stop, as crowding increases, or when nearby movement demands more attention.

7.5.2 Support Dynamic and Context-Aware Augmentation

Graduated immersion implies a further requirement: that the level of augmentation should be adjustable during a journey rather than fixed at its start. Chapter 6 found that reported passenger awareness rose with the number of people present in the carriage, and participants described shifting attention between media and surroundings during a condition. Journey phase, crowding, and social density were not manipulated in this thesis, so this requirement is motivated by these findings rather than established by them.

ARTV systems may therefore benefit from mechanisms that adjust augmentation density, visibility, or spatial prominence in response to changing attentional and environmental demands, for example reducing ornamentation during boarding or when a carriage becomes crowded, and restoring it during quieter phases of travel. Such context-aware strategies have the potential to maintain engagement while reducing attentional competition, though their effectiveness remains an open question for future research.

7.5.3 Preserve Environmental Reference Through Spatial Integration

The Spatial Integration dimension (section 7.3) showed that augmentation registered to the physical environment, rather than replacing it, allows users to retain visibility of real-world elements while engaging with media, contributing to the balance between immersive engagement and environmental awareness captured by the PIAT.

Passenger XR systems should therefore prioritise spatially anchored augmentation that aligns digital content with real-world surfaces, objects, and spatial structures, supporting peripheral awareness of the surroundings during media use. Designs that treat the physical environment as part of the narrative experience may suit public transport contexts better than approaches based on environmental replacement.

7.5.4 Design XR Media as Socially Situated Experiences

Passenger XR use occurs within shared public environments where social visibility, bystander presence, and shared-space etiquette shape user behaviour and technology acceptance [110, 16]. In the ecological evaluation, reported awareness of other passengers rose with the number of people present, and several participants were conscious of how the headset itself would appear to others (chapter 6). In passenger contexts, awareness necessarily includes other passengers, so XR media experiences should be designed as socially situated rather than purely individual.

The goal should not be to maximise awareness of other passengers at all times. For some users and journeys, reducing the salience of nearby passengers may be part of what makes immersive media appealing; for others, maintaining awareness of movement, proximity, gaze, or potential interaction matters for comfort and safety. Augmentation strategies should therefore preserve access to socially relevant cues without unnecessarily foregrounding bystanders as distractions, as the Bystander Integration extension (section 7.3) describes: avoiding occlusion of nearby passengers, maintaining visibility of peripheral movement, or letting users tune how much passenger activity remains visible during viewing.

7.5.5 Prioritise Perceptual Stability and Environmental Consistency

Chapter 6 found that registration drift, occlusion inconsistencies, and visual artefacts, including augmentation drawn across people in the carriage, were disruptive to participants. That study could not separate the effect of these artefacts from the effect of ornamenting a physical space as such, so their contribution should not be overstated; what it does establish is that participants noticed and were disrupted by them, and that this bears on coherence in the sense set out in section 7.4.

Designers of passenger XR systems should therefore prioritise robust environmental tracking, reliable occlusion handling, and consistent spatial alignment of digital content. In passenger contexts, where environmental unpredictability and user motion are common, such artefacts may be more noticeable and more disruptive than in stationary use, and perceptual stability of this kind may contribute to user trust and long-term usability.

7.5.6 Support Flexible Attention Management

Chapter 6 found that participants did not disengage fully from their surroundings during immersive viewing, but shifted attention between the media and ongoing environmental monitoring, a pattern described there as *flexible immersion*: increased immersion that does not require disengagement from the surroundings.

Passenger XR media systems should therefore be designed to support layered attention management rather than sustained attentional capture. This may include maintaining access to environmental cues, both auditory and peripheral visual, enabling rapid transitions between immer-

sive and awareness-oriented states, and avoiding augmentation that fully dominates the user's perceptual field. Where visual access is removed, auditory journey information becomes the only remaining channel, and participants described monitoring it as effortful, so redundancy across channels should be treated as a requirement rather than a convenience. XR experiences that support flexible attentional switching are more likely to match how participants described managing their attention; whether such designs also make travel safer was not measured in this thesis.

7.6 Limitations and Future Research Directions

While this thesis provides empirical, conceptual, and design contributions to immersive passenger XR interaction, several limitations should be acknowledged. These limitations provide important context for interpreting the findings and highlight directions for future research on immersive media experiences in mobility-driven environments.

7.6.1 Limitations of the Research

Methodological Limitations

Although this research incorporated both controlled experimentation and ecologically representative evaluation environments, the studies were conducted within simulated or semi-controlled transit settings rather than fully operational public transport systems. The decommissioned subway carriage used in chapter 6 enabled investigation of social presence, spatial layout, and environmental unpredictability beyond laboratory simulation, but it did not fully capture the behavioural and environmental complexity of live transport infrastructure. In particular, the research did not examine the influence of vehicle motion, variable journey duration, real scheduling pressures, or the temporal rhythms associated with boarding and disembarking. These factors may significantly influence passenger attention, comfort, and XR system usability. Conducting research within operational transport environments presents ethical, logistical, and safety challenges, but remains an important step for validating immersive passenger XR systems under fully realistic conditions.

It is useful to distinguish between two related but distinct concepts: *ecological realism* (the extent to which perceptual, social, and environmental conditions resemble those of real-world use) and *operational realism*, which concerns deployment within fully functioning public transport infrastructure. The evaluation presented in chapter 6 achieves greater ecological realism than prior simulation-based studies by introducing real bystanders, ambient noise, and physical spatial constraints, but does not achieve operational realism. Evaluation aboard an operational commuter service would be necessary to capture the full complexity of social dynamics, scheduling pressures, and passenger flow associated with live transit use. At the time this research was

conducted, requests to conduct in-situ evaluation on active trains were declined by the local transport operators, reflecting broader constraints faced by researchers investigating emerging XR technologies in public infrastructure contexts. This constraint reinforces the importance of the hybrid evaluation approach taken in this thesis, which sought to advance ecological validity incrementally while acknowledging the gap that remains between a preserved carriage and fully operational deployment.

Chapter 6 also identified three respects in which simulation diverged from situated deployment, which qualify the evidence obtained from the simulated studies that preceded it. *Perceptual accountability* concerns the requirement for virtual content to remain consistent with a physical world that is only partially sensed; *social plausibility* concerns whether the behaviour of other people matches expectations of shared public space; and *attentional consequence* concerns whether environmental events place genuine demands on attention. A fourth, *journey consequence*, remained beyond the reach of that study as well, since no participant faced any cost for failing to manage their journey. Chapters 4 and 5 were conducted wholly in simulation and are therefore subject to all four, which is the principal reason their findings are treated in this thesis as evidence about relative design performance rather than about real-world suitability.

A further methodological limitation concerns the confounding of medium with setting in chapter 6. The pass-through and fully immersive conditions were experienced in the public carriage while the simulated conditions were experienced in a private room, and conditions were blocked by location rather than fully randomised. Differences attributed to the medium may therefore also reflect the room, the presence of the public, or the order in which the blocks were completed. This is a general hazard for research that compares simulated and situated AR, since the two are rarely separable in practice, and it is noted here because the comparison is a contribution of this thesis rather than an incidental feature of one study.

A separate limitation concerns the design of the ARTV Ornamentation stimuli themselves. The design and placement of ornamentation elements were determined by the researchers and not systematically optimised. Some augmentations may have been too visually prominent, drawing attention away from the media content or surrounding environment, while others may have been insufficiently aligned with the narrative context, potentially introducing distraction or incongruity. This reflects a broader challenge in ARTV design, where visual prominence and spatial placement must be carefully balanced to support engagement without disrupting attention. In some cases, ornamentation elements were perceived as thematically incongruent with the underlying media, as when Stark iconography appeared during Lannister scenes in chapter 5, highlighting the importance of content coupling in maintaining coherence. Future work should explore systematic approaches to ornamentation design, including user-driven placement and perceptual optimisation strategies.

The media stimuli were also constrained. In chapter 5, it was not possible to find four clips of suitable length covering the same three settings without objectionable content, so the same

clip was shown in every condition, with counterbalancing used to mitigate the effects of repeat viewing. The order in which the settings changed was nonetheless identical across conditions, and familiarity with that sequence may have influenced the results. Applying ornamentation to general content, for example by using scene information carried in a programme's metadata to determine when the augmentation should change, would remove this constraint.

A further limitation concerns the stop-announcement awareness task used in chapter 6. Although the task provided a functional proxy for journey monitoring and enabled controlled comparison across XR conditions, it does not capture the full range of attentional behaviours exhibited during real-world travel, and it proved too coarse to detect differences between conditions, producing only eight failures across 114 trials. Passenger awareness during live travel is typically more intermittent, self-directed, and shaped by changing journey goals. In practice, passengers may shift attention in response to announcements, route familiarity, crowding, proximity to their destination, personal belongings, or unexpected social activity. The task therefore approximates one important component of passenger awareness, but may not fully represent the spontaneous, intermittent, and multi-layered awareness strategies used during operational journeys.

A further methodological concern specific to the simulated conditions evaluated in chapter 6 relates to participant understanding of the experimental setup. In some cases, participants appeared to misunderstand the nature of the simulated AR conditions, believing they were watching a virtual screen while physically seated on a real train, rather than experiencing a simulation of what AR-based media viewing might look like in a transit context. This misunderstanding likely affected how some participants interpreted and rated the simulated conditions, particularly for usage-intent and preference measures, where responses may reflect hypothetical enthusiasm for the concept of AR on a real journey rather than evaluation of the simulated condition as a stand-alone experience. Clearer pre-condition briefing or the use of interactive familiarisation trials prior to data collection may help establish the correct mental model in future studies that ask participants to engage with simulated AR as a proxy for real-world deployment.

Finally, the comparison this thesis draws is between ARTV Ornamentation and fully occlusive VR, and it does not test the obvious alternative: occlusive VR supplemented with awareness cues, such as reality anchors or selective pass-through of people and signage. Such designs might achieve a comparable balance by a different route, and the absence of that comparison bounds the claims made for spatial integration here.

Statistical Limitations

Environmental awareness was measured throughout by a single self-report item concerning awareness of other passengers, administered after each condition. A single item provides no internal consistency estimate, cannot distinguish between the objects of awareness identified in section 7.2.4, and is subject to the interpretive variation that any post-hoc self-report carries.

Chapter 6 added a behavioural measure, but as noted above it lacked the sensitivity to corroborate or contradict the self-reports.

The three empirical studies used 24, 20, and 19 participants respectively, and a single paired comparison at 80% power could detect effects of $d_z \geq 0.60$, 0.66, and $g_z \geq 0.68$ in each. Smaller differences may therefore have gone undetected, and non-significant comparisons should not be read as evidence of no difference, except where equivalence testing was applied and its bounds are stated. Correction for multiple comparisons was applied throughout, which further reduces sensitivity, and several of the comparisons on which the thesis relies fall close to conventional significance thresholds in either direction.

Internal consistency was acceptable for the instruments used, with two exceptions that bear on interpretation. The Film-IEQ was administered without item 13 in chapters 4 and 5, so its overall and Captivation scores are not directly comparable with published values, and its Comprehension subscale showed low internal consistency in both studies, which is why Comprehension is not relied upon in the synthesis presented here. Chapter 6 did not administer the Film-IEQ at all, so media immersion as that instrument defines it was not measured in the situated study, and the synthesis in section 7.2.2 uses Focused Attention as the common measure instead.

Technical and Hardware Limitations

The findings of this thesis are shaped by the capabilities and limitations of contemporary XR hardware, particularly video pass-through AR devices such as the Meta Quest 3 used in this research. Participants reported perceptual constraints associated with restricted field of view, camera-mediated visual artefacts, and occasional instability in spatial registration, along with the weight of the device and, in one case, eye strain. These limitations influenced user experience and may have affected perceptions of immersion and environmental integration.

The use of emerging development tools, including beta-stage scene understanding and environmental mapping frameworks, also introduced technical artefacts such as texture instability and occlusion inconsistencies. While these challenges reflect the current state of consumer XR technology, rapid hardware and software development may alter the feasibility and user acceptance of immersive passenger XR systems. As such, the generalisability of some findings may evolve as XR display technologies mature.

Scope and Contextual Limitations

The research examined entertainment-focused media consumption on journeys framed to participants as long enough for sustained viewing, consistent with the scope set out in chapter 1. Other passenger contexts, including short and frequently interrupted journeys, work-oriented XR use, and accessibility-focused applications, were not explored, and the findings should not be assumed to generalise to these settings without further investigation. No study involved a

moving vehicle, and none administered a simulator sickness measure, so the motion sickness considerations raised in chapter 2 remain untested here; this is a substantial gap, since motion is among the conditions most likely to affect both comfort and the usability of head-mounted displays in transit.

Participant samples were drawn from university and urban commuter populations across all four studies, which limits demographic representativeness. Gender composition ranged from near-parity in chapters 4 and 6 (12 men and 12 women; 11 men and 8 women respectively) to a male-skewed sample in chapter 5 (14 men, 6 women), and all samples skewed younger, reflecting the university recruitment context. More fundamentally, the commuter framing adopted throughout this thesis, centred on regular, point-to-point transit journeys, may itself reflect a gendered conception of passenger travel. As Criado Perez [36] documents, women's travel more commonly involves trip-chaining across multiple destinations for caregiving and other responsibilities, a pattern that is systematically under-represented in transport research and not captured by UK transit data, which remains largely non-gender-disaggregated. Future work should examine whether ARTV design priorities and awareness demands differ across these more complex journey types.

Furthermore, while this thesis examines awareness in relation to safety, social interaction, and journey management, it does not address broader regulatory, ethical, or infrastructural considerations associated with large-scale XR deployment in public transport systems.

7.6.2 Future Research Directions

Building on the empirical and conceptual contributions of this thesis, future research must move beyond controlled evaluation toward adaptive, scalable, and socially situated XR media systems. The following directions outline how ARTV ornamentation can evolve in response to real-world deployment, advances in adaptive interfaces, and emerging content generation paradigms.

In-the-Wild Deployment of Passenger XR Systems

While this thesis provides a structured empirical investigation of ARTV ornamentation through controlled and physically situated studies, the majority of findings are derived from simulated or semi-controlled environments. This reflects a broader methodological pattern within XR research, where practical, ethical, and technical constraints often necessitate the use of laboratory-based or simulated evaluations. As discussed in chapter 6, such approaches enable systematic manipulation of variables and reproducibility, but inevitably abstract away key characteristics of real-world passenger environments, including vehicle motion, variable journey duration, unpredictable passenger behaviour, and the temporal rhythms associated with boarding and disembarking.

Future research must therefore prioritise the deployment of immersive XR media systems

within fully operational transport infrastructures. In-the-wild studies would enable examination of how ARTV systems perform under dynamic and non-deterministic conditions, where attentional demands are shaped not only by media content but by real-time environmental events. Prior work in mobile and situated XR has demonstrated that augmentation density, placement, and anchoring strategies can significantly affect users' ability to maintain awareness of their surroundings during movement [77, 145]. Extending such findings to passenger contexts is critical, as transport environments introduce additional complexities including constrained personal space, social proximity to unfamiliar others, and the need to monitor journey progress.

In-the-wild deployment is not simply a matter of ecological validation, but raises new research questions that cannot be adequately addressed in controlled settings. For example, how do passengers dynamically regulate attention across extended journeys when immersive media is continuously available? Laboratory studies necessarily use short exposure windows, which capture initial responses but cannot model the attentional rhythms that emerge across a full journey, including fatigue, habitual checking behaviour, and the gradual recalibration of comfort with the technology. To what extent do users adapt their behaviour over time, developing strategies for balancing immersion and awareness? This is intrinsically a longitudinal question: participants in a single lab session are always encountering the technology for the first time, meaning the adaptive strategies that develop through repeated use, such as selectively engaging ornamentation, timing media consumption around stops, or adjusting density based on journey phase, are structurally invisible to controlled evaluation. How do transient environmental factors, such as crowding, service disruptions, or unexpected events, interact with XR media consumption? The unpredictability of such events is central to their attentional significance; staged or simulated disruptions carry none of the genuine urgency or social stakes that shape real passenger responses. This last question bears directly on the contextual component of the PIAT set out in section 7.2.5. The conditions identified there, including crowding, posture, custodial exposure, and the frequency of disembarkation decisions, were not manipulated in any study reported here, each of which was conducted with seated participants under uniform and uncrowded conditions. They are also conditions that resist manipulation, since staging them would remove the unpredictability that gives them their attentional force. Establishing whether they shape the configuration passengers find acceptable, and whether they do so in kind rather than merely in degree, therefore depends on observing them as they arise across real journeys. Addressing these questions therefore requires deployment within live transport environments where duration, variability, and real consequences are not approximated but present.

However, conducting research within live transport environments presents significant challenges. Safety considerations are paramount, particularly when evaluating technologies that may influence situational awareness in shared public spaces. Ethical concerns arise around bystander consent, privacy, and the recording of behaviour in uncontrolled environments. Logistical constraints, including access to transport infrastructure, variability in service conditions, and the

difficulty of standardising experimental protocols, further complicate deployment; the refusals encountered during this research illustrate that such access cannot be assumed. These challenges suggest that future work should explore hybrid evaluation methodologies that combine controlled experimentation with opportunistic in-the-wild data collection, as well as collaborations with transport providers to enable safe and scalable studies.

More broadly, advancing in-the-wild evaluation will require a shift in how passenger XR systems are conceptualised, not as isolated experimental artefacts, but as components embedded within socio-technical infrastructures. This includes integration with real-time transport information systems, consideration of device onboarding and usability in transient settings, and alignment with operational constraints such as peak-time crowding and service variability. By moving beyond controlled environments, future research can more fully characterise the PIAT as it manifests under real-world conditions, providing a more robust foundation for the design of XR media systems intended for everyday use.

Adaptive and Intelligent Ornamentation Systems

A central implication of this thesis is that ARTV ornamentation should not be treated as a fixed, statically authored layer, but as a dynamic system that responds to changing passenger context. The findings presented throughout this work indicate that the balance between media immersion and environmental awareness is not fixed, but varies with environmental conditions, user behaviour, and media characteristics. This suggests that future ARTV systems must move beyond static configurations of ornamentation, toward adaptive mechanisms that actively regulate this balance over time.

This direction aligns with a broader shift within XR research toward adaptive and context-aware interfaces. Prior work has explored optimisation-based approaches to mixed reality interface design, where elements such as placement, visibility, and level of detail are dynamically adjusted based on task requirements and environmental constraints [100, 29, 46]. More recent systems extend this paradigm by incorporating semantic understanding of the environment and, increasingly, social context, enabling interfaces to adapt in ways that preserve visibility, avoid occlusion of important real-world elements, and remain appropriate within shared spaces [97]. In parallel, emerging work on attention-aware XR systems has begun to explore how user attentional state can be inferred and used to modulate the presentation of virtual content in real time [77, 145].

However, much prior work has been motivated by productivity, navigation, or information management tasks, where adaptation is typically framed as a means of improving usability or efficiency. In contrast, the challenge addressed in this thesis is fundamentally different: rather than optimising task performance, passenger XR media systems must balance immersive engagement with the need to remain aware of a dynamic and socially shared environment. This introduces a distinct design objective for adaptive XR systems, in which adaptation must be

sensitive not only to environmental and task context, but also to the user's ongoing allocation of attention between media and reality.

Future research should therefore investigate adaptive ornamentation as a mechanism for real-time attention management within ARTV systems. Central open questions concern how contextual change should be detected and acted upon: what signals, whether environmental, behavioural, or journey-derived, are reliable proxies for shifting awareness demand, and at what thresholds should they trigger adaptation? How abruptly or gradually should ornamentation change in response, and how visible should that change be to the user? Similarly, it remains unclear how semantically selective adaptation should be, that is, whether systems should preserve visibility of specific real-world elements such as signage or nearby passengers, or whether coarser adjustments to overall density are sufficient. These questions require empirical investigation that goes beyond demonstrating that adaptation is desirable, toward understanding the perceptual and attentional mechanics that would make a deployed adaptive system effective.

Alongside environmental context, user state could itself drive adaptation. Behavioural indicators, such as head movement patterns or interaction frequency, as well as physiological signals associated with cognitive load or attentional engagement, could be used to infer when users are over- or under-engaged with the media content. This would enable closed-loop systems that modulate ornamentation in response to inferred attentional demand, shifting ARTV from a passive media enhancement to an active attentional support system. Such approaches build on emerging work in attention-aware XR but extend it into the domain of immersive media consumption, where the goal is not simply to minimise distraction but to maintain an appropriate balance between competing attentional demands.

A key open question is how control over adaptation should be distributed between system and user. Fully automated adaptation reduces the cognitive overhead of managing ornamentation, but risks automation complacency: if passengers assume the system is handling awareness demands, they may disengage from environmental monitoring in ways that undermine the safety benefits ARTV is designed to preserve [44]. Conversely, user-driven control maintains transparency and agency but places additional attentional demands on users in contexts where attention is already divided. Future work should therefore explore mixed-initiative approaches that combine system-driven adaptation with lightweight user control, and investigate how transparency and predictability influence trust calibration in adaptive XR media systems.

More broadly, adaptive ornamentation positions ARTV within an emerging class of intelligent XR systems that continuously mediate between virtual content and real-world context. In this framing, ornamentation is not simply decorative, but functional: it becomes a mechanism through which XR systems actively manage the relationship between immersion and awareness. Advancing this direction will require integrating insights from adaptive interface design [100, 29], attention-aware computing [77, 145], and socially situated XR [97], ultimately enabling passenger media experiences that are responsive, contextually appropriate, and capable

of operating robustly within dynamic real-world environments.

Automation and Content Authoring for Scalable ARTV Deployment

A key limitation of the ARTV paradigm, as explored in this thesis, is the reliance on manually authored ornamentation. While controlled studies permit careful design and evaluation of augmentation strategies, such approaches do not scale to the diversity of media content, environments, and user contexts encountered in real-world deployment: each new title or genre would require bespoke asset creation, spatial configuration, and scene-level authoring, imposing a per-title production overhead that is commercially impractical at scale. As discussed in section 7.4, chapter 5 indicates that reusable genre-level libraries may carry part of this burden, but also that the resulting augmentation risks incongruity with specific scenes. As ARTV systems move beyond experimental prototypes, the development of automated and intelligent content authoring pipelines could become essential for enabling widespread adoption.

This challenge aligns with a broader shift in XR research toward scalable and adaptive content generation. Prior work has explored techniques for authoring augmented reality experiences that can generalise across environments through semantic understanding of space and content. Systems such as ScalAR enable designers to demonstrate AR interactions in one environment and automatically adapt them to new settings using learned semantic relationships [58]. Similarly, research in scene-aware storytelling has shown how virtual content can be generated and positioned based on the structural and semantic properties of real-world environments, enabling narrative experiences that adapt dynamically to their surroundings [93, 96]. These approaches suggest a path toward ARTV systems in which ornamentation is not manually created and placed, but generated and adapted in response to both media content and physical context.

More recent developments in artificial intelligence further accelerate this trajectory. The integration of machine learning and generative models into XR pipelines has begun to enable systems that can infer appropriate augmentations from high-level descriptions of content, scene structure, or user intent. In this context, ARTV ornamentation could evolve from static, pre-authored assets into dynamically generated media layers that respond to scene changes, narrative progression, and user context in real time. For example, semantic analysis of video content could be used to identify scene transitions, object categories, or emotional tone, which in turn could drive the generation of contextually appropriate environmental augmentations. Such approaches would extend the content-adaptive strategies explored in this thesis, moving from predefined genre or scene-based mappings toward fully generative, data-driven systems. They would also bear directly on the coherence problem identified in section 7.4, since a system able to identify which faction, setting, or narrative moment is on screen is in a position to avoid the mismatches that participants noticed.

Automation in ARTV is not solely a technical challenge, but also a design problem. Decisions must be made regarding the degree of control retained by content creators versus that

delegated to automated systems. Fully automated generation risks producing augmentations that are inconsistent with artistic intent or that introduce visual clutter, while fully manual authoring limits scalability and adaptability. Future work should therefore investigate hybrid pipelines in which designers specify high-level constraints, styles, or rules, and automated systems generate ornamentations within these boundaries. This mirrors emerging approaches in XR authoring tools, where semantic abstractions and rule-based systems are used to balance creative control with scalability [58, 96].

A further consideration concerns personalisation. Automated ARTV systems could tailor ornamentation not only to content and environment, but also to individual user preferences, attentional patterns, or accessibility needs. For example, users may prefer different levels of augmentation density, stylistic themes, or degrees of environmental integration, and chapter 5 found that at least one participant preferred no ornamentation at all. Incorporating personalisation into automated pipelines raises additional challenges around user modelling, preference learning, and system transparency, but also offers the potential for highly customised immersive media experiences.

More broadly, advances in automated content generation position ARTV within an emerging paradigm of generative XR, where experiences are continuously constructed rather than statically authored. In this framing, ornamentation becomes a dynamic and contextually responsive layer that is co-produced by content, environment, and system intelligence. Future research should therefore explore how automated and generative approaches can be integrated with adaptive ornamentation strategies, ultimately enabling ARTV systems that are scalable, context-aware, and capable of supporting diverse media experiences across real-world passenger environments.

Advancements in XR Display Technologies

The findings of this thesis are grounded in the capabilities and limitations of contemporary XR hardware, particularly VST augmented reality systems. As discussed in chapter 2, current devices are constrained by factors such as limited field of view, camera-mediated visual artefacts, latency, and imperfect depth representation, all of which influence the perceptual integration of virtual content with the physical environment. These constraints directly impact the effectiveness of ARTV ornamentation, shaping both immersion and the user's ability to maintain awareness of their surroundings.

Rapid advances in XR display technologies are likely to significantly influence the viability of immersive passenger media systems. Improvements in field of view, display resolution, camera fidelity, depth sensing, and real-time spatial mapping have the potential to reduce many of the perceptual limitations observed in this work. The registration failures described in chapter 6 are instructive in this respect: they arose from a scene-understanding pipeline that relied on a room model captured in advance rather than on continuous reconstruction. Furthermore, a con-

tinuous reconstruction approach could be combined with light detection and ranging (LiDAR) sensing, which would reduce the depth errors that produce misalignment between virtual and real surfaces. Additionally, real-time semantic segmentation distinguishes structural surfaces from transient objects such as people, which is precisely the distinction that failed when augmentation was drawn across passengers. As these capabilities mature, the perceptual accountability failures identified in chapter 6 are likely to diminish, allowing hardware constraints to be distinguished more clearly from the limits of the ornamentation concept itself.

At the same time, advances in OST AR present an alternative trajectory for immersive media experiences. Unlike VST systems, OST displays preserve direct visual access to the physical world, potentially mitigating some of the perceptual artefacts associated with camera passthrough, such as reduced field of view and latency. However, OST technologies introduce their own challenges, including limited brightness, reduced contrast in varying lighting conditions, and persistent ambiguity in occlusion and depth perception. These trade-offs suggest that different display paradigms may support different forms of ARTV experiences, with VST systems offering greater control over visual presentation and OST systems providing stronger perceptual continuity with the real world.

Lightweight display glasses are the form factor most likely to be adopted for passenger use, since the weight and social visibility of a headset were both raised as barriers by participants in chapter 6. They cannot, however, deliver ARTV Ornamentation as it is defined in this thesis. Their fields of view fall at or below the restricted condition of chapter 4, at which ornamentation produced no significant gain in immersion, and their additive displays can add light to a scene but cannot remove it, so surfaces can be overlaid but not retextured, darkened, or occluded. The two constraints compound: what such a device can display is a small, translucent region of added content, which is closer to the floating-screen baseline than to any ornamented condition evaluated here. This is a substantive limit rather than a temporary one, and it identifies a question this thesis cannot answer: whether a form of ornamentation that relies only on additive light within a narrow field of view can produce any of the engagement benefit observed here, or whether the benefit depends on the environmental transformation that only occlusion-capable displays permit. Until wide field-of-view devices capable of occlusion become socially acceptable in public, passenger ARTV faces a choice between a device that can deliver the experience and a device that passengers would be willing to wear.

While hardware improvements may alleviate many current limitations, they do not eliminate the fundamental design challenge addressed in this thesis. The PIAT arises not only from technical constraints, but from the intrinsic demands of mobile, shared environments in which users must divide attention between media content and their surroundings. Even with near-perfect display fidelity and environmental integration, passengers will still need to monitor journey progress, respond to social cues, and maintain awareness of their environment. As such, the need for carefully designed and potentially adaptive ornamentation strategies is likely to persist

even as hardware evolves.

Future research should therefore aim to disentangle limitations arising from current hardware constraints from those that are fundamental to passenger XR contexts. Replicating and extending the studies presented in this thesis using next-generation devices will help clarify which observed effects, such as the reduced effectiveness of ornamentation under limited field of view, are likely to diminish with improved hardware, and which reflect deeper characteristics of attention, perception, and behaviour in mobile environments. Comparative evaluations across VST and OST systems may be particularly valuable in this regard, enabling a more nuanced understanding of how different display paradigms shape the balance between immersion and awareness.

Advancements in XR display technologies should be viewed not as a replacement for interaction design considerations, but as an evolving substrate upon which such considerations must operate. The design principles and frameworks developed in this thesis, including ARTV ornamentation and the PIAT, provide a foundation for understanding how immersive media systems can be structured across both current and future hardware.

Longitudinal and Social Acceptability Studies

While this thesis provides insight into short-term user responses to ARTV Ornamentation, the long-term use of immersive XR media in public transport contexts remains largely unexplored. Passenger XR systems are not used in isolation, but within socially shared and repeatedly experienced environments, where patterns of behaviour, expectations, and norms may evolve over time. Understanding how users adapt to and integrate XR media into their daily routines therefore represents a critical direction for future research.

Longitudinal studies are particularly important for examining how the PIAT is negotiated across repeated journeys and extended periods of use. Initial novelty effects may influence user preferences, immersion, and willingness to use XR systems, potentially leading to different outcomes compared to sustained, habitual use. Over time, users may develop strategies for managing attention, such as selectively engaging with ornamentation, adjusting system settings, reducing use during demanding journey phases, or timing media consumption around stops and transfers. Future research should therefore examine ARTV use over longer durations and repeated journeys to understand how engagement, cognitive load, attentional strategy, and awareness management change once ornamentation becomes familiar rather than novel.

In addition to individual adaptation, the social context of public transport introduces further complexity. XR media consumption in shared environments raises questions around bystander perception, social comfort, and interaction etiquette. Prior work has shown that the visibility and legibility of mixed reality interfaces can influence how bystanders interpret and respond to XR users, affecting both user comfort and social acceptability [172]. Similarly, research on privacy-aware adaptation in AR highlights the need for systems to account for the presence

of others, dynamically adjusting content to respect social and contextual boundaries [144]. In passenger contexts, where users are often in close proximity to unfamiliar individuals, these considerations become particularly salient, and the incident in chapter 6 in which a member of the public addressed a participant who could see them through pass-through, although this was not apparent from outside the headset, illustrates the ambiguity that co-present others face. Future work should therefore examine not only the passenger's long-term adaptation to ARTV, but also how co-present passengers perceive, tolerate, or contest sustained XR media use in shared transit spaces.

The relationship between XR use and social norms over time represents an additional challenge. The acceptability of technologies such as smartphones and headphones in public transport has evolved significantly, with behaviours that were once unusual becoming commonplace. XR devices may follow a similar trajectory, but this process is likely to depend on factors such as device form factor, visibility of interaction, and perceived disruption to others. Longitudinal research can help identify how norms around XR media use emerge, stabilise, or shift, and how system design can support socially appropriate behaviour during this process.

Future work should also examine how XR systems can actively support social awareness and co-presence, rather than treating other passengers solely as sources of distraction or constraint. This includes exploring mechanisms for integrating bystanders into the XR experience in subtle and non-intrusive ways, as described by the Bystander Integration extension in section 7.3, as well as ensuring that critical real-world elements, such as nearby passengers or personal belongings, remain perceptually accessible. Such considerations align with emerging work on socially situated XR, which emphasises that virtual content must coexist with, rather than replace, the social dynamics of physical environments [97].

More broadly, advancing the long-term adoption of passenger XR systems requires understanding not only how individuals interact with the technology, but how these interactions are embedded within broader social and cultural contexts. This includes considerations of trust, predictability, and perceived appropriateness, as well as the development of shared expectations around when and how immersive media should be used in public spaces. By examining XR use over time and within real social settings, future research can provide a more comprehensive understanding of how immersive passenger media systems can be designed to be not only effective, but also socially sustainable.

Chapter 8

Conclusions

8.1 Answering the Research Questions

This thesis investigated how immersive XR media experiences can be designed to support passenger engagement while preserving environmental awareness. This section answers each research question from the evidence reported in chapters 3 to 6, cross-referencing rather than restating the studies themselves.

8.1.1 Research Question 1

> How could XR improve the current experience of immersive media on public transport?

The survey and workshops reported in chapter 3 show that passengers already treat travel time as media time, using phones, tablets, and laptops, and that they do so while continuing to monitor journey progress, other passengers, and their belongings. Those devices constrain engagement through their display size and their separation from the surrounding space, while fully occlusive VR conflicts with the monitoring that participants described as necessary: workshop participants consistently proposed designs that preserved some access to the surroundings rather than removing it, and one revised an initial preference for full immersion on the grounds that losing all sense of reality would be unsafe. The limitations of handheld devices are drawn from prior work (chapter 2) rather than measured here.

The opportunity for XR therefore lies in augmenting the passenger environment rather than replacing it. Extending media content beyond the display and into the surrounding space may offer engagement that a handheld screen cannot, while leaving the environment perceptually available; handheld viewing was not itself compared in this thesis. This is the tension that chapter 7 develops as the PIAT, and ARTV Ornamentation is the design response this thesis evaluates against it.

8.1.2 Research Question 2

> How can ARTV with Ornamentation be used to balance media immersion and environmental awareness in passenger transportation?

ARTV Ornamentation balances the two by integrating media content into the passenger environment rather than substituting for it, which leaves the carriage visible while extending the media beyond the screen. Chapter 4 found that, at a wide field of view, Medium and High ornamentation raised immersion, user experience, and engagement above the floating-screen AR baseline, and that these were the only ARTV conditions that could not be distinguished from fully immersive VR after correction for multiple comparisons. That is not equivalence: both scored lower than VR on the Film-IEQ, the uncorrected differences were significant, and equivalence testing was inconclusive. Chapter 6 found the ornamentation benefit again in a physical carriage, where it raised hedonic quality in both pass-through and simulated AR, and where both pass-through conditions preserved substantially more reported awareness of other passengers than either fully immersive alternative.

The balance is not costless. In chapter 4, the ornamentation that raised engagement also reduced reported awareness of other passengers relative to unornamented viewing; in chapter 5, the scene-adaptive condition was the most engaging and the most preferred, and the only one to reduce awareness; and in chapter 6, ornamentation lowered perceived pragmatic quality. What ARTV Ornamentation offers is therefore not the removal of the trade-off but a position within it that fully immersive alternatives cannot occupy: engagement above that of a floating screen, with the surrounding environment still available. Participants in chapter 6 described moving their attention between the media and the carriage, which is the behaviour this design position is intended to support.

8.1.3 Research Question 3

> What is the impact of device FoV, ornamentation density, and ornamentation type on the media immersion–environmental awareness trade-off?

Field of view moderated the effect of ornamentation. Chapter 4 found an interaction between field of view and degree of ornamentation: at the wide field of view, higher levels of ornamentation raised immersion and lowered reported awareness, whereas at the restricted field of view the differences between ornamentation levels were not significant on either measure. Ornamentation was not without effect at the restricted field of view, as the High ornamentation condition still exceeded the floating-screen baseline, but the gain was smaller. Part of that shortfall may be an artefact of the simulation, since the restricted conditions showed a visible boundary that a physical display would not (section 4.5.5). The result therefore bounds, rather than rules out, what narrow field-of-view displays can deliver.

Ornamentation level behaved as a graded parameter rather than a switch, but the two steps differed in kind. The step from Low to Medium increased the number of objects, whereas the step from Medium to High held object count constant and added retexturing, which produced no further gain in immersion beyond a rise in aesthetic appeal. Participants ranked the intermediate level as highly as the most intense one and chose it most often for use on a train, yet at the wide field of view that level reduced awareness of other passengers significantly below the unornamented baseline. The configuration passengers prefer and the configuration that preserves awareness are therefore not the same.

Between object-level approaches, the specificity of the objects made no detectable difference. Chapter 5 found no detectable difference between show-specific and generic genre-level objects on any measure, which is not the same as establishing that the two are equivalent, and genre-level ornamentation improved on the unornamented baseline only in aesthetic appeal. The scene-adaptive Biome condition outperformed both on immersion, engagement, and preference, while also reducing reported awareness. That condition, however, changed several things at once: it retextured the carriage, added particle effects, and changed with the scene over time. Its advantage therefore shows that this combination is more effective than static object placement, but not which of its components is responsible, or that authoring at the level of the scene is the cause.

8.1.4 Research Question 4

> How does the passenger immersion–awareness trade-off apply when moving from simulated AR-in-VR to AR experiences situated in the real world?

The design-level findings transferred; the situated experience did not. Chapter 6 found that ornamentation raised hedonic quality by comparable amounts in pass-through AR and in simulation, with no interaction between medium and ornamentation on any measure, so the effect established in the simulated studies was reproduced in a physical carriage. Beyond that, the two settings diverged in three respects: pass-through deployment exposed registration and rendering failures that simulation could not, in the form of texture instability and augmentation drawn across real people; the scripted passengers of the simulation did not create the social demands that real visitors did, as reflected in higher reported awareness in the carriage and in its association with the number of visitors present; and the physical setting placed demands on attention that the predictable private room did not. Section 6.4.1 developed these as perceptual accountability, social plausibility, and attentional consequence.

The behavioural evidence for the trade-off remains weaker than the self-report evidence. The stop-detection task used in chapter 6 produced eight failures across 114 trials and could not distinguish between conditions, so it cannot establish whether the reported reductions in awareness carried behavioural consequences. What participants did describe was that monitor-

ing the journey became more effortful once the carriage was occluded, and that they reverted to listening for announcements when they could no longer see the signage. The ecological validity achieved is accordingly intermediate: the study added a physical carriage, members of the public, pass-through hardware, and an explicit monitoring demand, but not a moving vehicle, fellow passengers on a shared journey, or any consequence for missing a stop.

8.1.5 Summary of Research Question Outcomes

These answers describe a design position rather than a solution. Passengers already divide attention between media and environment, and value doing so; ARTV Ornamentation raises engagement above conventional AR viewing while leaving the environment available, at a cost to reported awareness and to perceived practicality; that cost depends on field of view and on how much and what kind of augmentation is applied, and, on participants' own accounts, on the conditions of the journey; and the hedonic benefit survives the move from simulation to a physically situated setting, although the situated conditions that shape it do not reduce to what simulation reproduces.

8.2 Contributions of the Thesis

The contributions below follow the order set out in chapter 1.

8.2.1 C1: An Empirical Characterisation of Opportunities for Immersive XR Media in Public Transport Contexts

Chapter 3 provides a characterisation of how passengers currently consume media in transit and what they would require of an immersive alternative. Its contribution is not that passengers watch video while travelling, which is established, but that they articulate the immersion–awareness problem themselves, in terms of what they need to keep track of and why, before any system is put in front of them. The workshops identify conditions under which passengers said they would accept greater immersion, including journey duration, crowding, and enclosure; these form the contextual component of the PIAT (section 7.2.5), although none was varied in the studies that followed.

8.2.2 C2: The Introduction and Systematic Evaluation of ARTV Ornamentation as a Design Mechanism for Balancing Immersion and Awareness

This thesis introduces ARTV Ornamentation, in which media-themed augmentation is placed into the passenger's surroundings so that immersion arises through environmental integration

rather than environmental substitution, and evaluates it across three studies. Chapter 4 establishes the effect and its dependence on field of view; chapter 5 examines the level of abstraction at which ornamentation can be authored, showing that a scene-adaptive treatment, combining retexturing, particle effects, and timed transitions, outperformed both genre-level and show-specific object placement, although its components were not separated, and that reusable genre-level objects alone improved on unornamented viewing only in aesthetic appeal; and chapter 6 evaluates the paradigm on pass-through hardware in a physical carriage. The paradigm is offered as a design approach with an identified position in the trade-off and identified costs, rather than as a validated system.

8.2.3 C3: An Analysis of How Key Design and Technical Parameters Shape the Immersion–Awareness Trade-off

The thesis quantifies how field of view, ornamentation density, and ornamentation type shape reported immersion, engagement, and awareness, as summarised under RQ3. Two results are of most consequence for design. The first is the interaction between field of view and ornamentation, which means that the paradigm's viability is bounded by display hardware as well as by content design. The second is the separation between the preferred configuration and the awareness-preserving configuration, which suggests that user preference should not be assumed to indicate situational appropriateness. Both outcomes concern self-reported awareness; chapters 4 and 5 included no behavioural measure, and the measure introduced in chapter 6 proved too coarse to corroborate the self-reports.

8.2.4 C4: A Methodological Case Study of VR-Simulated AR as a Tool for Evaluating AR Media Experiences

Chapter 6 compares VR-simulated AR with pass-through AR in a physically situated setting, using the same content, the same measures, and the same participants. Simulation reproduced the direction and approximate size of the ornamentation effect on hedonic quality, but did not reproduce the perceptual accountability of augmentation registered to a real environment, the social demands created by real bystanders, or the attentional consequence of an unpredictable setting. The contribution is a case of conditional rather than general equivalence, together with three dimensions along which the divergence can be anticipated in other passenger XR evaluations. Because medium was confounded with setting in that comparison, the case identifies where simulation and situated deployment diverged rather than isolating why.

8.2.5 C5: The Development of the Passenger Immersion–Awareness Trade-off as a Design Concept for Passenger XR

The PIAT, developed in chapter 7, treats immersive engagement and environmental awareness as two outcomes of a common set of design decisions rather than as two ends of a single scale. It allows designs to be positioned with respect to the trade-off rather than merely described as subject to it, distinguishes awareness by its object, its modality, and its degree, and connects the extent to which a system registers content to the physical world, a dimension that existing mixed reality taxonomies identify but leave unattached to any user-side construct, to what a passenger retains of their surroundings. It is a design concept in the sense of intermediate-level knowledge, assessed by whether it generates further design instances rather than by predictive power, and its contextual component rests on participant reasoning and one within-study association rather than on experimental evidence.

8.3 Conclusion

This thesis set out to establish how immersive media might work for passengers who cannot afford to stop attending to where they are. Its answer is that immersion in this setting is better treated as a relationship to be managed than as a quantity to be maximised. ARTV Ornamentation demonstrates one way of managing it, by extending media into the carriage rather than replacing the carriage. In the controlled studies, it produced greater immersion and engagement than a floating screen, at a cost to reported awareness and perceived practicality. Those gains were largest on a wide field-of-view display, and depended on how much augmentation was applied and of what kind. The PIAT generalises the problem beyond this particular paradigm, and the comparison between simulated and situated evaluation indicates how much of such work can be done before a real carriage becomes necessary.

What remains open is the part of the problem that could not be reached from a stationary carriage. No study here involved a moving vehicle, a journey with consequences, or passengers who had somewhere to be, and the conditions that participants themselves identified as decisive, including crowding, posture, and the frequency of decisions about when to get off, were never varied. As XR displays move toward forms that passengers might plausibly wear in public, the question is less whether immersive media can be made compatible with travel than which parts of that compatibility hold when the journey is real.

Appendix A

Raw Datasets

1. The raw dataset for Chapter 3 is available at DOI: https://doi.org/10.5281/zenodo.20573749
2. The raw dataset for Chapter 4 is available at DOI: https://doi.org/10.5281/zenodo.20574145
3. The raw dataset for Chapter 5 is available at DOI: https://doi.org/10.5281/zenodo.20574411
4. The raw dataset for Chapter 6 is available at DOI: https://doi.org/10.5281/zenodo.20583166

Appendix B

Survey Instruments

This appendix contains the survey questions used across the studies.

B.1 Study 1: Commuter Behaviour Survey

This appendix presents the survey instrument used to investigate commuter activities and preferences during passenger journeys. The survey was administered using Qualtrics. The structure below reflects the question logic and response formats used in the study.

B.1.1 Activity Engagement During a Journey

Participants were asked to consider a typical journey (approximately 30 minutes) as a passenger in a car, bus, or train.

Prompt:

> For this question think about a 30 minute long journey as a passenger in a car, bus, or train, and the activities you undertake. For each of the following activities move the slider to the corresponding category reflecting if you currently don't do this activity and don't want to do it, do this activity, do it but think it could be done better, or don't do it but would like to.

Response Scale

- 0 = Don't do and don't want to
- 1 = Currently do
- 3 = Currently do but want to do better
- 4 = Don't currently do but want to

Activities

Using social media websites or apps (Facebook, Twitter, LinkedIn, Tumblr, TikTok, Instagram)
Texting, emailing or other messaging
Reading electronically
Playing games on an electronic device
Listening to music or audio (podcasts)
Doing nothing
Viewing scenery
Watching people
Thinking or daydreaming
Reading print (newspaper, physical book, etc.)
Eating or drinking
Talking on the phone
Personal grooming (e.g., shaving, makeup)
Sleeping / snoozing
Writing or editing paper documents
Writing or editing electronic documents
Watching movies, TV, or other video
Trying not to be anxious about the journey (e.g., delays or where to get off)
Video conferencing (e.g., Zoom, Skype, FaceTime)
Relaxation / meditation
Physical games (e.g., pen and paper, crosswords, board games, playing cards)
Talking to other passengers
Other (please specify)

B.1.2 Follow-Up Questions (Conditional)

Based on participant responses, additional open-ended questions were presented.

If Participant Selected: "Currently do but want to do better"

Why do you not do this activity as often as you would like / as well as you would like?

If Participant Selected: "Don't currently do but want to"

What factors prevent you from doing this activity?

B.1.3 Current Commute Version

A second version of the question asked participants to reflect on their *current commute* rather than a hypothetical journey.

> For this question think about your current commute and the activities you undertake.
> For each of the following activities move the slider to the corresponding category...

The same activity list and response scale were used as described above.

B.1.4 Notes on Survey Structure

- The activity list was presented using a slider-based Likert-style interface in Qualtrics.
- Follow-up questions were conditionally displayed using Qualtrics loop and branch logic.
- Multiple versions of the question were used (e.g., hypothetical journey vs. current commute), but the core instrument remained consistent across conditions.

B.2 Study 2: ARTV Ornamentation Evaluation

This section presents the survey instruments used in Chapter 4 to evaluate ARTV Ornamentations for immersion and awareness. All surveys were administered using Qualtrics. Participants completed a demographic questionnaire and consent form at the start of the session, a per-condition battery after each of the eight experimental conditions, and an end-of-study questionnaire after experiencing all conditions.

B.2.1 Participant Information Sheet and Consent Form

The following information sheet and consent form were administered digitally via Qualtrics at the start of the session. Participants were required to indicate consent before proceeding.

Purpose

The purpose is to understand the impact virtual content has on immersion and awareness, and user preferences for how to increase immersion.

What is Expected

In this study you will be watching video clips in a VR headset and asked to answer some questions regarding them. There will also be a short interview at the end. No previous experience with VR is required to participate.

Benefits and Data Handling

You will be provided with a £10 Amazon voucher. The results will be held by the researcher and the University of Glasgow, and may be published in a journal or presented at a conference. Results will always be presented such that individual participants cannot be identified. Upon publication, anonymous data will be deposited in a suitable repository.

Withdrawal and Ethics

Participation is voluntary; you may withdraw at any time without explanation. This project has been reviewed by the College Ethics Committee. The research is funded by ERC Horizon 2020 project ViAjeRo #835197.

Consent Items

Participants were required to confirm all of the following statements and then select *I consent* or *I do not consent*:

- I confirm that I have read and understood the participant information sheet for the above study and have had the opportunity to ask questions.
- I understand that my participation is voluntary and that I am free to withdraw at any time, without giving any reason.
- I understand that the data collected may be used in publications or presentations.
- I am at least 18 years old.

B.2.2 Demographic Questionnaire

Prior to beginning the experimental conditions, participants completed the following demographic questions.

Prompt:

What is your age?

Free-text numeric entry.

Prompt:

What is your gender?

Single-select: Male; Female; Non-binary / third gender; Prefer not to say.

B.2.3 Per-Condition Questionnaire Battery

The following battery was administered after each of the eight experimental conditions. All items were presented as horizontal sliders in Qualtrics. The battery comprised four instruments: the Film Immersive Experience Questionnaire (Film-IEQ), custom Passenger Awareness items, the User Engagement Scale Short Form (UES-SF), and the User Experience Questionnaire Short (UEQ-S).

Film Immersive Experience Questionnaire (Film-IEQ)

The Film-IEQ [149] was used to assess the participant's immersive experience of the media content after each condition. The canonical scale comprises 24 items across four subscales: Captivation (items 1–4, 13, 17–19, 21–24), Real-World Dissociation (items 5–7, negatively scored), Comprehension (items 12, 14–16), and Transportation (items 8–11, 20).

Prompt:

> For the purposes of this study, "real world" refers to the train interior with no modifications. For the following questions place the slider such that it corresponds to your answer. 1 being Not at all / Very Little and 7 being A lot / Very.

Response scale: 1 (Not at all / Very Little) – 7 (A lot / Very), continuous horizontal slider. The 24 items are:

1. To what extent did the movie, TV show, or clip hold your attention?
2. To what extent did you feel you were focused on the movie, TV show, or clip?
3. How much effort did you put into watching the movie, TV show, or clip?
4. Did you feel that you were trying your best to follow the events of the movie, TV show, or clip?
5. To what extent did you feel consciously aware of being in the real world whilst watching?
6. To what extent were you aware of yourself in your surroundings?
7. To what extent did you notice events taking place around you?
8. To what extent could you picture yourself in the scene of the events shown in the movie, TV show, or clip?
9. To what extent did you feel as though you were separated from your real-world environment?
10. To what extent did you feel that the movie, TV show, or clip was something you were experiencing, rather than something you were just watching?

11. To what extent was your sense of being in the environment shown in the movie, TV show, or clip stronger than your sense of being in the real world?
12. To what extent did you find the concepts and themes of the movie, TV show, or clip challenging?
13. To what extent did you feel motivated to keep on watching?
14. To what extent did you find the concepts and themes of the movie, TV show, or clip easy to understand?
15. To what extent did you feel like you were making progress towards understanding what was happening during the movie, TV show, or clip, and what you thought might happen at the end?
16. How well do you think you understood what happened?
17. To what extent were you interested in seeing how the events shown in the movie, TV show, or clip would progress?
18. How much did you want the events in the movie, TV show, or clip to unfold successfully for the main characters involved?
19. Were you in suspense about how the events would unfold?
20. At any point did you find yourself become so involved that you wanted to speak to the movie, TV show, or clip directly?
21. To what extent did you enjoy the cinematography, graphics, and/or imagery?
22. How much would you say you enjoyed watching the movie, TV show, or clip?
23. When it was over, were you disappointed that you had to stop watching?
24. Would you like to watch more of this, or similar content, in the future?

Custom Passenger Awareness Items

Two bespoke items assessed awareness of co-present passengers and of the applied augmentations.

Prompt:

> Rate how much you agree with these statements from Strongly Disagree (1) to Strongly Agree (5).

Response scale: 1 (Strongly Disagree) – 5 (Strongly Agree), continuous horizontal slider. Items:

1. I was aware of other passengers whilst watching the video.
2. I was aware of the virtual modifications (trees and plants).

User Engagement Scale – Short Form (UES-SF)

The UES-SF [130] was administered after each condition to assess engagement with the viewing experience.

Prompt:

> Rate the following statements from 1 (Strongly Disagree) to 7 (Strongly Agree).

Response scale: 1 (Strongly Disagree) – 7 (Strongly Agree), continuous horizontal slider. Items are listed below by subscale.

Focused Attention (FA):

1. I lost myself in this experience.
2. The time I spent watching the clip just slipped away.
3. I was absorbed in this experience.

Perceived Usability (PU):

4. I felt frustrated whilst watching the clip.
5. I found it confusing watching the clip.
6. Watching the clip was taxing.

Aesthetics (AE):

7. Watching the clip in this manner was attractive.
8. Watching the clip in this manner was aesthetically appealing.
9. Watching the clip in this manner appealed to my senses.

Reward Factor (RW):

10. Watching the clip in this manner was worthwhile.
11. My experience was rewarding.
12. I felt interested in this experience.

User Experience Questionnaire – Short (UEQ-S)

The UEQ-S [157] was administered after each condition using a 7-point semantic differential format.

Prompt:

> This question consists of pairs of contrasting attributes that may apply to the product. The circles between the attributes represent gradations between the opposites. You can express your agreement with the attributes by ticking the circle that most closely reflects your impression. Please decide spontaneously; there is no right or wrong answer.

Response scale: 7-point scale between each attribute pair. Items by subscale:

Pragmatic Quality:

Negative pole	Scale	Positive pole
Obstructive	1–2–3–4–5–6–7	Supportive
Complicated	1–2–3–4–5–6–7	Easy
Inefficient	1–2–3–4–5–6–7	Efficient
Confusing	1–2–3–4–5–6–7	Clear

Hedonic Quality:

Negative pole	Scale	Positive pole
Boring	1–2–3–4–5–6–7	Exciting
Not interesting	1–2–3–4–5–6–7	Interesting
Conventional	1–2–3–4–5–6–7	Inventive
Usual	1–2–3–4–5–6–7	Leading edge

B.2.4 End-of-Study Questionnaire

After completing all eight conditions, participants completed the following questions.

Condition Preference Ranking

Prompt:

> Rank the conditions you have experienced in order of preference. 1 being your favourite.

Presented as a drag-and-drop ranking interface. The eight conditions were:

- Low FOV, Low Modification

- Low FOV, Medium Modification
- Low FOV, High Modification
- High FOV, Low Modification
- High FOV, Medium Modification
- High FOV, High Modification
- Baseline (AR, no ornamentation)
- VR Scene

Interest in Real-World Use

Prompt:

> Would you be interested in using this condition over a phone / laptop when on the train?

Presented as a matrix with the eight conditions above as rows. Response options: Yes; No.

B.3 Study 3: ARTV Ornamentations for a New Genre

This section presents the survey instruments used in Chapter 5 to investigate ARTV Ornamentation specificity and dynamic scene adaptation. The core per-condition questionnaire battery was identical to that described in Section B.2.3, with minor modifications to two items to reflect the changed stimulus content. These differences are noted explicitly below.

B.3.1 Participant Information Sheet and Consent Form

The following information sheet and consent form were administered digitally via Qualtrics. The study received ethics approval from the College of Science and Engineering Ethics Committee, University of Glasgow (ref: 300230047).

Purpose

The purpose of this study is to determine if there is a difference in the level of immersion a user experiences when watching Augmented Reality Television when the augmentations are specific to what is being shown on screen, or when the augmentations are more generally based on the theme of the content.

What is Expected

Firstly, you will be asked to fill in a questionnaire with basic demographic information. Then you will be fitted with the head-mounted display (HMD) such that it sits comfortably on your head, and you can see the virtual view correctly. This will occur at the School of Computing Science (Sir Alwyn Williams Building). As part of the study you will watch clips from television shows under four different conditions: with no augmentations; with augmentations chosen to match the general theme of the content; with augmentations tailored to what is currently on screen; and again with no augmentations. After each condition you will answer a series of questions. After experiencing all conditions you will answer a further set of questions and take part in a short interview. The interview will be voice recorded throughout.

If you experience any of the following symptoms during this study, remove the head-mounted display immediately and inform the researcher: (1) altered vision; (2) lightheadedness; (3) dizziness; (4) involuntary movements such as eye or muscle twitching; (5) confusion; (6) nausea; (7) loss of awareness; (8) convulsions; (9) cramps; (10) disorientation; (11) motion sickness; (12) discomfort or pain in head or eyes; or (13) any other symptoms you would class as atypical or unusual.

Benefits and Data Handling

You will be provided with a £10 Amazon voucher. Results will be held by the researcher and the University of Glasgow, and may be published in a journal or presented at a conference. Results will always be presented such that individual participants cannot be identified. Upon publication, anonymous data will be deposited in a suitable repository. The research is funded by ERC Horizon 2020 project ViAjeRo #835197.

Consent Items

Participants confirmed all of the following statements and selected *I consent* or *I do not consent*. Those who did not consent were unable to proceed.

- I confirm that I have read and understood the participant information sheet for the above study and have had the opportunity to ask questions.
- I understand that my participation is voluntary and that I am free to withdraw at any time, without giving any reason.
- I understand that the data collected may be used in publications or presentations and will be anonymised and stored within a research repository.
- I understand that all data collected from me will be treated confidentially and anonymised, will be seen in its raw form only by the experimenters, and if published will not be identifiable as coming from me.

- I am at least 18 years old.
- I understand that I should not log in to, or use, any personal web-based services on supplied devices during the course of the experiment, as all usage is being recorded and this could compromise my privacy.
- I agree to take part in the above study.

B.3.2 Demographic Questionnaire

The Study 3 demographic questionnaire extended that used in Study 2 with two additional items assessing prior participation and media familiarity.

Prompt:

What is your age?

Free-text numeric entry.

Prompt:

What is your gender?

Single-select: Male; Female; Non-binary / third gender; Prefer not to say.

Prompt:

Did you take part in the previous study?

Single-select: Yes; No. Participants who indicated prior participation were excluded.

Prompt:

Rate how familiar you are with the following TV shows from “Never Heard of It” to “Watched It All”.

Presented as a matrix. Rows: Game of Thrones; House of the Dragon.

Response Scale

- 1 = Never heard of it
- 2 = Never seen it
- 3 = Watched an episode
- 4 = Watched a few episodes
- 5 = Watched a series
- 6 = Watched several series
- 7 = Watched most of it
- 8 = Watched it all

B.3.3 Per-Condition Questionnaire Battery

The per-condition battery was identical to that described in Section B.2.3, comprising the Film-IEQ, the UES-SF, and the UEQ-S with the same response scales and item wording, with the following modification:

Modified Custom Passenger Awareness Item

In Study 3, awareness item 2 was updated to refer to the specific augmentations used:

Prompt:

> Rate how much you agree with these statements from Strongly Disagree (1) to Strongly Agree (5).

1. I was aware of other passengers whilst watching the video.
2. I was aware of the virtual modifications (texture changed, Game of Thrones items).

All other items were identical to those listed in Sections B.2.3–B.2.3.

Notes on Survey Structure

- The per-condition battery was administered four times, once per condition.
- The UES-SF and UEQ-S were identical across Studies 2 and 3.
- Only the second Passenger Awareness item differed between studies to reflect the changed ornamentation content.

B.3.4 End-of-Study Questionnaire

After completing all four conditions, participants completed the following questions.

Condition Preference Ranking

Prompt:

> Rank the conditions you have experienced in order of preference. 1 being your favourite.

Presented as a drag-and-drop ranking interface. The four conditions were:

- Baseline (no ornamentation)
- Interior Change With Scene
- Constant Generic Medieval Objects
- Constant Game of Thrones Objects

Interest in Real-World Use

Prompt:

> I would be interested in watching a video on a train using this condition instead of watching using a phone or laptop.

Presented as a matrix with the four conditions above as rows.

Response Scale

- Strongly Agree
- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

B.4 Study 4: ARTV In Situ with Realistic Awareness Task

This section presents the survey instruments used in Chapter 6 to evaluate ARTV Ornamentation in a physically situated, ecologically valid passenger context. The core per-condition battery was again closely based on that described in Section B.2.3, with one extension to the Film-IEQ and a modified awareness item. Modifications are noted explicitly below.

B.4.1 Participant Information Sheet and Consent Form

The following information sheet and consent form were administered digitally via Qualtrics prior to the session. The study received ethics approval from the College of Science and Engineering Ethics Committee, University of Glasgow (ref: 300230236).

Purpose

The purpose of this study is to understand the opinions of users when watching Augmented Reality Television in different transport settings.

What is Expected

You will meet the experimenter outside the Riverside Museum in Glasgow. You will watch a series of video clips under six conditions. For four conditions you will be seated in a vintage subway car exhibit, with members of the public coming and going as normal, to mimic a public transportation environment. For the remaining two conditions you will be seated in a quiet, private room. You will answer a short questionnaire after each condition and again after experiencing all conditions. At the end of the study you will take part in a short interview. The session will last approximately one hour and will include a continuous voice recording.

If you experience any of the following symptoms during this study, remove the head-mounted display immediately and inform the researcher: (1) altered vision; (2) lightheadedness; (3) dizziness; (4) involuntary movements such as eye or muscle twitching; (5) confusion; (6) nausea; (7) loss of awareness; (8) convulsions; (9) cramps; (10) disorientation; (11) motion sickness; (12) discomfort or pain in head or eyes; or (13) any other symptoms you would class as atypical or unusual.

Benefits and Data Handling

You will be provided with a £10 Amazon voucher. Results will be held by the researcher and the University of Glasgow, and may be published in a journal or presented at a conference. Results will always be presented such that individual participants cannot be identified. Upon publication, anonymous data will be deposited in a suitable repository. The research is funded by ERC Horizon 2020 project ViAjeRo #835197.

Consent Items

Participants confirmed all of the following statements and selected *I consent* or *I do not consent*. Those who did not consent were unable to proceed.

- I confirm that I have read and understood the participant information sheet for the above study and have had the opportunity to ask questions.
- I understand that my participation is voluntary and that I am free to withdraw at any time, without giving any reason.
- I understand that the data collected may be used in publications or presentations and will be anonymised and stored within a research repository.
- I understand that all data collected from me will be treated confidentially and anonymised, will be seen in its raw form only by the experimenters, and if published will not be identifiable as coming from me.
- I am at least 18 years old.

- I understand that I should not log in to, or use, any personal web-based services on supplied devices during the course of the experiment, as all usage is being recorded and this could compromise my privacy.
- I agree to take part in the above study.

B.4.2 Demographic Questionnaire

The Study 4 demographic questionnaire extended that of Study 3 with items assessing visual correction use and prior XR device experience, reflecting the use of pass-through AR headsets in this study.

Prompt:

What is your age?

Free-text numeric entry.

Prompt:

What is your gender?

Single-select: Male; Female; Non-binary / third gender; Prefer not to say.

Prompt:

Do you wear glasses or contacts?

Response Options

- I wear glasses
- I wear contact lenses
- I wear both
- I do not wear glasses and I do not wear contact lenses

Prompt:

Please indicate your experience with Augmented Reality and/or Virtual Reality Head-Mounted Devices (e.g., Microsoft HoloLens, Meta Quest).

Response Options

- Never used them before
- Used them once or twice
- I use them a few times a year
- I use them a few times a month
- I use them once per week
- I use them a few times per week
- I use them daily

Prompt:

> Did you take part in a previous Game of Thrones study using the Quest 3 VR headset?

Single-select: Yes; No. One participant indicated prior participation in the study reported in Appendix 5. They were retained in the analysed sample, and their prior exposure to ARTV Ornamentation is reported in section 6.2.4 and acknowledged among the limitations in section 6.4.3.

Prompt:

> Rate how familiar you are with the following TV shows from "Never Heard of It" to "Watched It All".

Matrix question using the same 8-point familiarity scale as in Study 3 (Section B.3.2). Rows: Game of Thrones; House of the Dragon.

B.4.3 Per-Condition Questionnaire Battery

The per-condition battery was closely based on that described in Section B.2.3. Two modifications were made to reflect the ecologically valid setting and the use of physically situated pass-through AR.

Film Immersive Experience Questionnaire (Film-IEQ)

All 24 canonical Film-IEQ items were included, see section B.2.3 for the full item list.

Modified Custom Passenger Awareness Item

Awareness item 2 was updated to reflect the pass-through AR context:

1. I was aware of other passengers whilst watching the video.
2. I was aware of the virtual modifications (texture changed).

Additional Standalone Awareness Item

A single additional item was presented separately after the battery to capture passenger awareness independently from the other awareness items:

Prompt:

> Rate how much you agree with these statements from Strongly Disagree (1) to Strongly Agree (5).

1. I was aware of other passengers whilst watching the video.

Response scale: 1 (Strongly Disagree) – 5 (Strongly Agree), continuous horizontal slider.

All remaining instruments (UES-SF, UEQ-S) were identical to those described in section B.2.3.

Notes on Survey Structure

- The per-condition battery was administered six times, once per condition.
- The UES-SF, and UEQ-S were identical to those used in Studies 2 and 3.
- The second Passenger Awareness item was modified to reflect the AR-specific augmentation context.

B.4.4 End-of-Study Questionnaire

After completing all six conditions, participants completed the following questions.

Condition Preference Ranking

Prompt:

> Rank the conditions you have experienced in order of preference. 1 being your favourite.

Presented as a drag-and-drop ranking interface. The six conditions were:

- AR Baseline
- ARTV Ornamented
- Simulated AR Baseline
- Simulated ARTV Ornamented
- VR Cinema
- VR Videosphere

Interest in Real-World Use

Prompt:

> I would be interested in watching a video on a train using this condition instead of watching using a phone or laptop.

Presented as a matrix with the six conditions above as rows.

Response Scale

- Strongly Agree
- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

Appendix C

Interview Guides

This appendix presents the semi-structured interview guides used in the experimental studies reported in Chapters 4, 5, and 6. Each guide was designed to elicit qualitative reflections on participant experience after the experimental conditions had been completed. Interviews were voice-recorded and subsequently transcribed for thematic analysis.

C.1 Study 1: ARTV Ornamentation Evaluation (Chapter 4)

This section presents the semi-structured interview guide used in Chapter 4. The interview was conducted immediately after participants completed the end-of-study questionnaire, with the interviewer referencing the participant's submitted condition ranking on screen as an anchor for discussion. Questions were semi-structured: the four core questions below were asked of all participants, with follow-up probes used as relevant based on individual responses. The interview was voice-recorded and subsequently transcribed for thematic analysis.

C.1.1 Opening

The interview opened by asking the participant to walk through their submitted condition ranking:

> Can you talk me through the reasoning behind your rankings? Starting with why you ranked [favourite condition] first.

Follow-up probes used as relevant:

- Why did you rank [condition X] above [condition Y]?
- Why did you rank the [low / high] FOV conditions where you did?
- Why did you rank the VR Scene where you did?

- Would you use any of those conditions instead of a phone or laptop on the train? Which ones, and why?
- If you could modify the system in any way, what would you change?

C.1.2 Plant Density: Low vs. Medium

Could you tell the difference between the low amount of plants and the medium amount of plants?

Follow-up probes used as relevant:

- Was the difference obvious, or quite subtle?
- Did you prefer one over the other? Why?

C.1.3 Augmentation Level: Medium vs. High

The difference between the medium and high conditions was that we changed the textures of the seat backs and the table surfaces — so the seats went from blue cloth to a stone or bark texture. Did you notice that change?

Follow-up probes used as relevant:

- Did that improve the experience, or was the difference negligible?
- Did you prefer the medium or the high? Why?

C.1.4 Passenger Awareness

Did you find that the augmentations affected your awareness of other passengers in the carriage? Were there any points where the plants were covering or obfuscating the other passengers?

Follow-up probes used as relevant:

- Was that a good or bad thing for you?
- If you were using this on a real train, would that be a problem?
- How would you want the system to handle passenger visibility?

C.1.5 Closing

Do you have any other comments, feedback, or questions?

Notes on Interview Structure

- The interview was conducted with reference to the participant's submitted condition preference ranking, displayed on screen.
- The four core questions (plant density, texture change, passenger awareness, and closing) were asked of all participants. The opening ranking discussion and follow-up probes were adapted to each participant's responses.
- The real-world use probe ("Would you use any of these conditions over a phone or laptop?") was asked when not already addressed in the ranking discussion.
- Interviews were voice-recorded and transcribed for thematic analysis, as described in Section 4.3.2.

C.2 Study 2: ARTV Ornamentations for a New Genre (Chapter 5)

Note: The interview was conducted after participants submitted the end-of-study condition-preference ranking. The four conditions were: Baseline (no ornamentation); Genre-Based Augmentations (generic medieval objects and textures, constant throughout); Show-Specific Augmentations (Game of Thrones themed objects, constant throughout); and Biome-Specific Augmentations (objects and textures that changed with the on-screen scene). The interviewer used the participant's submitted ranking as an anchor for the discussion.

C.2.1 Opening: Condition Preference Ranking

Prompt:

> Can you tell me about why you ranked the conditions in the order that you ranked them?

Follow-up probes used as relevant:

- Why was [condition] your favourite / least favourite?
- Why did you rank [condition X] above [condition Y]?
- Would you use any of those conditions on a train instead of watching on a phone or laptop? Which ones, and why?

C.2.2 Biome Transitions: Scene Changes

Prompt:

> When the scene changed — going from the medieval setting to the snow and then to the desert — did you find that transition distracting at all?

Follow-up probes used as relevant:

- Did it improve the experience overall, or detract from it?
- Did it take you out of the content, or add to it?

C.2.3 Particle Effects: Snow and Sand

Prompt:

> When the particle effects — the snow or the sand — came in front of the screen, would you have preferred that those didn't occlude the screen? So the screen would always be fully visible with no particles in front of it?

Follow-up probes used as relevant:

- Was the snow more distracting than the sand, or about the same?
- Would you prefer the particles to disappear when they reach the screen, or to fade/become more transparent?
- Or were you okay with the particles passing in front of the screen as they were?

C.2.4 Object Specificity: Show-Specific vs. Genre-Based

Prompt:

> Did you notice a difference between when the objects were Game of Thrones themed — the banners, the wolf head, the Iron Throne — and when they were just generic medieval objects like books and coins? Do you think having show-specific objects improved the experience compared to just having generic medieval ones?

Follow-up probes used as relevant:

- Was that a meaningful difference, or quite small?
- If you were watching a different fantasy show with the same generic medieval objects, do you think that would still add to the experience?

C.2.5 Object Adaptation: Static vs. Scene-Responsive

Prompt:

> In the Show-Specific Augmentations condition, the Game of Thrones objects stayed the same throughout even as the scene moved between different locations — King's Landing, the Stark territories, the desert. Did you notice that the objects weren't changing with the scene?

Follow-up probes used as relevant:

- Would you have preferred that the objects changed with the location — for example, more Lannister-themed items at King's Landing and Stark-themed items in the North?
- Do you think that would be more immersive, or more distracting?
- How might that transition be done without breaking immersion?

C.2.6 Real-World Use and Generalisation

Prompt:

> Would you be interested in using this kind of augmentation technology when watching content — either on a train or at home?

Follow-up probes used as relevant:

- If it were lightweight glasses rather than a headset, would that change your answer?
- Do you think genre-based augmentations would work well across different shows in the same genre — for example, watching different fantasy shows in the same medieval environment?

C.2.7 Repeated Clip

Prompt:

> Having watched the same clip four times across the conditions, did you find that got boring? Do you think that affected how you answered the questionnaires?

C.2.8 Closing

> Do you have any other comments, feedback, or questions?

Notes on Interview Structure

- The interview was anchored to the participant's submitted condition preference ranking, displayed on screen.
- The opening ranking discussion, biome transitions, particle effects, and object specificity questions were asked of all or nearly all participants. The object adaptation, real-world use, and repeated clip probes were raised as relevant based on individual responses.
- The awareness of other passengers was not explicitly probed in this study, as the simulated passengers were animated models rather than real co-present bystanders.
- Interviews were voice-recorded and transcribed for thematic analysis, as described in Section 5.2.5.

C.3 Study 3: ARTV In Situ with Realistic Awareness Task (Chapter 6)

The following interview guide was used in Chapter 6. The interview was conducted after participants completed the end-of-study questionnaire and condition preference ranking. Conditions are referred to below using the thesis names; the corresponding internal labels used during data collection were: Lab Alpha (Simulated ARTV Ornamented), Lab Beta (Simulated AR Baseline), AR Alpha (ARTV Ornamented), AR Beta (AR Baseline), Cinema (VR Cinema), and Videosphere (VR Videosphere).

C.3.1 Simulated AR Conditions (Lab)

Prompt:

> Can you tell me about your experience in the VR train carriage conditions? Those were the two conditions in the quiet room.

Follow-up probes used as relevant:

- What did you like or dislike about those conditions?
- Did the augmentations in Simulated ARTV Ornamented make a difference compared to the Simulated AR Baseline?

C.3.2 Physically Situated AR Conditions (Subway Carriage)

Prompt:

> Tell me about your experience in the augmented reality conditions in the physical train carriage — the ones where you could still see me and the tablet.

Follow-up probes used as relevant:

- Were there any issues with the visuals — did anything look wrong or distracting?
- Did the augmentations affect how you felt about being in a public space?

C.3.3 VR Cinema Condition

Prompt:

> Tell me about your experience in the cinema condition — where you were fully in VR but seated in the physical train.

Follow-up probes used as relevant:

- Did you feel safe or comfortable not being able to see your surroundings?
- Would you be happy using that condition on a real train with other passengers coming and going?

C.3.4 VR Videosphere Condition

Prompt:

> Tell me about your experience in the videosphere condition — where you had the floating screen with a 360 degree image surrounding you.

Follow-up probes used as relevant:

- How did it compare to the cinema condition?
- Did you have a preference between the two?

C.3.5 Comparison: Simulated AR vs. Physically Situated AR

Prompt:

> How did the VR train carriage conditions compare to the augmented reality conditions in the physical train?

- What were some differences between them?
- What were some similarities?
- Would you say they were interchangeable?

C.3.6 Comparison: Train Conditions vs. Cinema and Videosphere

Prompt:

> How did the train conditions — both the simulated and the physical — compare to the VR Cinema and VR Videosphere conditions?

Follow-up probes used as relevant:

- Which did you find more immersive?
- Which felt more comfortable for use in public?

C.3.7 Stop Awareness Task

Prompt:

> Do you think your awareness of which stop to get off at was the same across all six conditions? Why, or why not?

Follow-up probes used as relevant:

- Did it change depending on which condition you were in?
- Were some conditions easier or harder for managing the task?

C.3.8 Real-World Use

Prompt:

> Which condition would you choose to use in public transport, and why?

Follow-up probes used as relevant:

- Would your answer change for a short journey versus a long one?
- Are there any conditions you would not use, and why?

C.3.9 Condition Preference Ranking

Prompt:

> In the exit questionnaire you were asked to rank the six conditions in order of preference. Can you explain how you ranked them and why?

Follow-up probes used as relevant:

- Why was [favourite condition] your top choice?
- Why was [least favourite condition] at the bottom?

C.3.10 Closing

> Any final comments, queries, or questions?

Notes on Interview Structure

- The interview was conducted after participants completed the end-of-study questionnaire and condition ranking.
- All nine topic areas above were covered with each participant; probes were deployed as relevant based on individual responses.
- The stop-awareness probe was sometimes integrated into the condition comparison discussion rather than asked as a standalone question.
- Interviews were voice-recorded and transcribed for thematic analysis, as described in Section 6.2.8.

Appendix D

Chapter 4 - ARTV Ornamentations for Immersion and Awareness - Complete Statistical Results

Table D.1: Descriptive statistics (Median, IQR, Mean, *SD*) for all measures by condition ($n = 24$ per condition).

Measure	Condition	Median	IQR	Mean	*SD*
Film-IEQ Overall	AR Baseline	74.50	17.80	74.50	19.20
	HF-H	96.00	21.20	97.30	14.60
	HF-L	83.50	28.50	81.60	20.40
	HF-M	94.00	21.80	94.30	16.00
	LF-H	87.00	23.00	86.20	15.20
	LF-L	85.00	20.20	83.00	17.80
	LF-M	89.00	21.20	85.30	16.00
	VR Baseline	104.0	26.50	102.0	17.20
Film-IEQ Captivation	AR Baseline	55.00	14.20	56.40	12.20
	HF-H	63.50	10.50	63.80	8.62
	HF-L	58.50	20.80	57.40	12.80
	HF-M	63.50	13.50	62.70	9.93
	LF-H	60.50	14.50	60.20	10.00
	LF-L	61.50	12.50	59.60	10.60
	LF-M	60.50	15.20	60.70	9.57
	VR Baseline	65.00	13.80	62.90	10.00
Film-IEQ Real World Diss.	AR Baseline	−14.50	7.25	−14.80	4.68
	HF-H	−8.00	6.50	−9.83	5.51
	HF-L	−13.00	5.25	−12.70	4.34
	HF-M	−10.50	4.50	−10.30	3.95
	LF-H	−11.50	5.00	−11.90	4.42
	LF-L	−12.00	7.25	−13.00	4.66
	LF-M	−13.00	6.50	−12.70	4.82
	VR Baseline	−5.50	6.00	−6.46	3.40
Film-IEQ Comprehension	AR Baseline	21.00	3.25	20.20	2.19
	HF-H	21.00	2.00	20.40	2.06
	HF-L	21.00	3.00	20.20	2.57
	HF-M	22.00	2.25	20.80	2.11

Table D.1 continued

Measure	Condition	Median	IQR	Mean	*SD*
	LF-H	20.00	3.00	20.30	1.90
	LF-L	21.00	3.00	20.40	2.19
	LF-M	20.00	3.25	20.00	2.20
	VR Baseline	22.00	2.00	21.00	2.31
Film-IEQ Transportation	AR Baseline	12.00	10.50	12.70	7.02
	HF-H	22.00	6.25	23.00	5.84
	HF-L	17.00	12.50	16.80	7.09
	HF-M	20.00	7.50	21.20	6.53
	LF-H	16.00	8.00	17.60	6.05
	LF-L	15.00	10.50	16.00	7.11
	LF-M	19.00	8.00	17.30	5.97
	VR Baseline	25.00	7.50	24.70	6.02
UEQ-S Overall	AR Baseline	1.31	1.59	1.04	1.16
	HF-H	2.31	0.84	2.21	0.76
	HF-L	1.50	1.50	1.47	1.05
	HF-M	2.06	1.06	2.00	0.84
	LF-H	1.69	1.62	1.47	1.01
	LF-L	1.38	1.75	1.38	0.95
	LF-M	1.31	1.19	1.36	1.00
	VR Baseline	2.25	0.88	2.15	0.72
UEQ-S Pragmatic	AR Baseline	1.75	1.81	1.49	1.12
	HF-H	2.00	1.25	1.99	1.06
	HF-L	1.75	2.12	1.64	1.10
	HF-M	2.00	1.31	1.89	1.03
	LF-H	1.50	2.06	1.27	1.29
	LF-L	1.62	1.56	1.38	1.02
	LF-M	1.25	1.38	1.27	1.07
	VR Baseline	2.50	1.19	2.06	0.94
UEQ-S Hedonic	AR Baseline	0.75	2.12	0.59	1.59

Table D.1 continued

Measure	Condition	Median	IQR	Mean	*SD*
	HF-H	2.75	1.00	2.44	0.66
	HF-L	1.25	2.00	1.30	1.22
	HF-M	2.00	1.31	2.11	0.78
	LF-H	1.62	1.31	1.67	0.87
	LF-L	1.25	2.06	1.38	1.17
	LF-M	1.62	1.62	1.45	1.11
	VR Baseline	2.38	1.25	2.23	0.79
UES-SF Overall	AR Baseline	2.67	0.96	2.84	1.02
	HF-H	4.71	1.21	4.21	0.86
	HF-L	3.38	1.88	3.27	1.09
	HF-M	4.00	1.46	3.87	0.88
	LF-H	3.54	1.75	3.36	1.06
	LF-L	3.08	1.50	3.13	1.05
	LF-M	3.25	1.58	3.20	1.10
	VR Baseline	4.38	1.50	4.12	0.83
UES-SF Focused Attention	AR Baseline	3.83	2.00	4.01	1.34
	HF-H	6.17	1.83	5.82	1.15
	HF-L	4.67	1.83	4.58	1.47
	HF-M	5.67	1.67	5.43	1.24
	LF-H	5.00	2.17	4.69	1.34
	LF-L	4.33	2.17	4.32	1.38
	LF-M	5.00	1.42	4.75	1.30
	VR Baseline	6.17	1.42	5.86	1.07
UES-SF Perceived Usability	AR Baseline	1.00	0.75	1.47	0.75
	HF-H	1.00	0.42	1.40	0.71
	HF-L	1.00	0.75	1.47	0.75
	HF-M	1.00	1.08	1.57	1.01
	LF-H	1.33	1.00	1.76	0.99
	LF-L	1.00	1.00	1.62	1.07
	LF-M	1.17	1.42	1.78	1.04

Table D.1 continued

Measure	Condition	Median	IQR	Mean	*SD*
	VR Baseline	1.00	0.75	1.50	0.82
UES-SF Aesthetics	AR Baseline	4.17	2.42	4.14	1.63
	HF-H	6.83	1.83	6.10	1.22
	HF-L	4.67	2.33	4.85	1.51
	HF-M	6.00	1.75	5.67	1.32
	LF-H	5.50	2.00	5.18	1.54
	LF-L	5.00	2.00	4.65	1.60
	LF-M	4.83	2.25	4.58	1.59
	VR Baseline	6.17	1.50	5.92	1.15
UES-SF Reward	AR Baseline	5.00	2.42	4.68	1.60
	HF-H	6.83	1.33	6.33	0.85
	HF-L	5.50	2.33	5.12	1.48
	HF-M	6.00	1.75	5.94	0.97
	LF-H	5.67	1.67	5.35	1.38
	LF-L	5.67	2.42	5.18	1.62
	LF-M	5.17	2.25	5.25	1.48
	VR Baseline	6.67	1.08	6.21	1.03
Aware of Other Passengers	AR Baseline	5.00	1.25	4.08	1.25
	HF-H	2.50	2.25	2.71	1.40
	HF-L	4.00	2.00	3.71	1.27
	HF-M	2.50	2.00	2.79	1.22
	LF-H	4.00	2.00	3.62	1.21
	LF-L	4.00	2.00	3.92	1.06
	LF-M	4.00	1.25	3.62	1.24
	VR Baseline	1.00	0.00	1.12	0.34
Aware of Augmentations	AR Baseline	1.00	1.00	1.75	1.26
	HF-H	5.00	1.00	4.50	0.88
	HF-L	4.00	2.00	3.58	1.32
	HF-M	4.00	1.00	4.21	1.02

Table D.1 continued

Measure	Condition	Median	IQR	Mean	*SD*
	LF-H	4.00	2.25	3.62	1.28
	LF-L	3.00	2.00	2.96	1.20
	LF-M	4.00	1.25	3.62	1.21
	VR Baseline	5.00	1.00	4.46	0.83

Note. Median and IQR are reported as primary descriptives; Mean and *SD* are included for reference. HF/LF: High/Low FoV; L/M/H: Low/Medium/High ornamentation. Overall Film-IEQ scores did not deviate from normality ($W = .989$, $p = .133$); all other measures violated normality and were analysed using the ART. The Film-IEQ was administered without item 13 (Captivation). Real-World Dissociation items were reverse-scored by negation rather than by subtraction from 8, so Real-World Dissociation and overall Film-IEQ scores are 24 points lower than under the published key; differences between conditions are unaffected.

Table D.2: Omnibus tests for the 2 (FoV) $\times$ 3 (DoA) factorial design over the six ARTV Ornamentation conditions. The overall Film-IEQ was analysed with a repeated-measures ANOVA (Greenhouse–Geisser corrected); all other measures with the ART. $\eta_p^2 = F \cdot df_1/(F \cdot df_1 + df_2)$.

Measure	**Effect**	df_1	df_2	F	p	η_p^2
Film-IEQ	FoV	1	23	13.18	.001**	.364
	DoA	1.42	32.58	7.49	.005**	.246
	Interaction	1.97	45.20	4.30	.020*	.158
Captivation	FoV	1	115	1.33	.251	.011
	DoA	2	115	2.57	.081	.043
	Interaction	2	115	2.15	.121	.036
Real-World Dissociation	FoV	1	115	9.24	.003**	.074
	DoA	2	115	5.38	.006**	.086
	Interaction	2	115	2.18	.118	.037
Comprehension	FoV	1	115	1.55	.215	.013
	DoA	2	115	0.04	.962	.001
	Interaction	2	115	0.54	.582	.009
Transportation	FoV	1	115	17.90	<.001***	.135
	DoA	2	115	9.00	<.001***	.135
	Interaction	2	115	3.21	.044*	.053
UEQ-S	FoV	1	115	20.78	<.001***	.153

Table D.2 continued

Measure	**Effect**	df_1	df_2	F	p	η_p^2
	DoA	2	115	4.73	.011*	.076
	Interaction	2	115	2.81	.064	.047
Pragmatic	FoV	1	115	21.38	<.001***	.157
	DoA	2	115	0.65	.524	.011
	Interaction	2	115	1.23	.296	.021
Hedonic	FoV	1	115	15.81	<.001***	.121
	DoA	2	115	11.05	<.001***	.161
	Interaction	2	115	4.05	.020*	.066
UES-SF	FoV	1	115	20.88	<.001***	.154
	DoA	2	115	8.81	<.001***	.133
	Interaction	2	115	2.77	.067	.046
Focused Attention	FoV	1	115	21.40	<.001***	.157
	DoA	2	115	9.20	<.001***	.138
	Interaction	2	115	2.49	.087	.042
Perceived Usability	FoV	1	115	13.15	<.001***	.103
	DoA	2	115	1.97	.144	.033
	Interaction	2	115	3.35	.038*	.055
Aesthetics	FoV	1	115	15.17	<.001***	.117
	DoA	2	115	10.97	<.001***	.160
	Interaction	2	115	2.21	.114	.037
Reward	FoV	1	115	6.47	.012*	.053
	DoA	2	115	4.86	.009**	.078
	Interaction	2	115	3.15	.047*	.052
Awareness of other passengers	FoV	1	115	11.82	<.001***	.093
	DoA	2	115	8.95	<.001***	.135
	Interaction	2	115	4.32	.016*	.070
Awareness of augmentations[†]	FoV	1	115	20.42	<.001***	.151
	DoA	2	115	7.67	<.001***	.118
	Interaction	2	115	0.33	.723	.006

[†]Manipulation check. * $p < .05$; ** $p < .01$; *** $p < .001$.

Table D.3: *Post hoc* comparisons for the factorial design. DoA comparisons are Bonferroni-corrected; simple effects (pairs of conditions sharing a level of one factor, tested where the interaction was significant) are Holm-corrected. For the overall Film-IEQ, paired comparisons of estimated marginal means are reported, with simple effects of DoA Bonferroni-corrected within each FoV level. Other measures use ART-C contrasts. Negative t indicates that the first-named level scored lower. HF/LF: High/Low FoV; L/M/H: Low/Medium/High ornamentation.

Measure	**Family**	**Contrast**	t	df	p_{adj}
Film-IEQ	DoA	Low – Medium	−2.55	23	.053
	DoA	Low – High	−3.17	23	.013*
	DoA	Medium – High	−1.26	23	.665
	FoV	Low – High	−3.63	23	.001**
	Simple effect	LF-L – LF-M	−0.68	23	>.999
	Simple effect	LF-L – LF-H	−1.11	23	.837
	Simple effect	LF-M – LF-H	−0.27	23	>.999
	Simple effect	HF-L – HF-M	−3.06	23	.017*
	Simple effect	HF-L – HF-H	−3.63	23	.004**
	Simple effect	HF-M – HF-H	−1.33	23	.590
	Simple effect	LF-L – HF-L‡	0.51	23	.618
	Simple effect	LF-M – HF-M‡	−2.36	23	.027*
	Simple effect	LF-H – HF-H‡	−4.08	23	<.001***
Captivation	DoA	Low – Medium	−1.86	115	.196
	DoA	Low – High	−2.05	115	.127
	DoA	Medium – High	−0.19	115	>.999
	FoV	Low – High	−1.15	115	.251
Real-World Dissociation	DoA	Low – Medium	−1.83	115	.211
	DoA	Low – High	−3.27	115	.004**
	DoA	Medium – High	−1.45	115	.453
	FoV	Low – High	−3.04	115	.003**
Comprehension	DoA	Low – Medium	−0.27	115	>.999
	DoA	Low – High	−0.19	115	>.999
	DoA	Medium – High	0.08	115	>.999
	FoV	Low – High	−1.25	115	.215
Transportation	DoA	Low – Medium	−3.08	115	.008**
	DoA	Low – High	−4.06	115	<.001***
	DoA	Medium – High	−0.98	115	.988

Table D.3 continued

Measure	**Family**	**Contrast**	t	df	p_{adj}
	FoV	Low – High	−4.23	115	<.001***
	Simple effect	HF-H – HF-L	4.96	115	<.001***
	Simple effect	HF-H – HF-M	1.66	115	.497
	Simple effect	HF-H – LF-H	4.37	115	<.001***
	Simple effect	HF-L – HF-M	−3.29	115	.009**
	Simple effect	HF-L – LF-L	0.62	115	>.999
	Simple effect	HF-M – LF-M	2.94	115	.024*
	Simple effect	LF-H – LF-L	1.20	115	.923
	Simple effect	LF-H – LF-M	0.23	115	>.999
	Simple effect	LF-L – LF-M	−0.98	115	.993
UEQ-S	DoA	Low – Medium	−1.88	115	.187
	DoA	Low – High	−3.05	115	.009**
	DoA	Medium – High	−1.17	115	.737
	FoV	Low – High	−4.56	115	<.001***
Pragmatic	DoA	Low – Medium	−0.46	115	>.999
	DoA	Low – High	−1.13	115	.778
	DoA	Medium – High	−0.67	115	>.999
	FoV	Low – High	−4.62	115	<.001***
Hedonic	DoA	Low – Medium	−2.81	115	.018*
	DoA	Low – High	−4.67	115	<.001***
	DoA	Medium – High	−1.86	115	.196
	FoV	Low – High	−3.98	115	<.001***
	Simple effect	HF-H – HF-L	5.84	115	<.001***
	Simple effect	HF-H – HF-M	1.97	115	.256
	Simple effect	HF-H – LF-H	4.50	115	<.001***
	Simple effect	HF-L – HF-M	−3.87	115	.001**
	Simple effect	HF-L – LF-L	−0.20	115	>.999
	Simple effect	HF-M – LF-M	3.42	115	.005**
	Simple effect	LF-H – LF-L	1.13	115	>.999
	Simple effect	LF-H – LF-M	0.89	115	>.999
	Simple effect	LF-L – LF-M	−0.25	115	>.999
UES-SF	DoA	Low – Medium	−2.14	115	.104

Table D.3 continued

Measure	Family	Contrast	t	df	p_{adj}
	DoA	Low – High	−4.20	115	<.001***
	DoA	Medium – High	−2.06	115	.125
	FoV	Low – High	−4.57	115	<.001***
Focused Attention	DoA	Low – Medium	−3.27	115	.004**
	DoA	Low – High	−4.04	115	<.001***
	DoA	Medium – High	−0.78	115	>.999
	FoV	Low – High	−4.63	115	<.001***
Perceived Usability	DoA	Low – Medium	−0.87	115	>.999
	DoA	Low – High	−1.98	115	.150
	DoA	Medium – High	−1.11	115	.803
	FoV	Low – High	3.63	115	<.001***
	Simple effect	HF-H – HF-L	−0.99	115	>.999
	Simple effect	HF-H – HF-M	−0.73	115	>.999
	Simple effect	HF-H – LF-H	−3.33	115	.011*
	Simple effect	HF-L – HF-M	0.27	115	>.999
	Simple effect	HF-L – LF-L	−0.08	115	>.999
	Simple effect	HF-M – LF-M	−1.62	115	.757
	Simple effect	LF-H – LF-L	2.26	115	.207
	Simple effect	LF-H – LF-M	0.98	115	>.999
	Simple effect	LF-L – LF-M	−1.28	115	>.999
Aesthetics	DoA	Low – Medium	−1.53	115	.383
	DoA	Low – High	−4.60	115	<.001***
	DoA	Medium – High	−3.07	115	.008**
	FoV	Low – High	−3.90	115	<.001***
Reward	DoA	Low – Medium	−1.88	115	.187
	DoA	Low – High	−3.10	115	.007**
	DoA	Medium – High	−1.21	115	.682
	FoV	Low – High	−2.54	115	.012*
	Simple effect	HF-H – HF-L	4.73	115	<.001***
	Simple effect	HF-H – HF-M	1.89	115	.309
	Simple effect	HF-H – LF-H	4.03	115	<.001***
	Simple effect	HF-L – HF-M	−2.85	115	.037*

Table D.3 continued

Measure	Family	Contrast	t	df	p_{adj}
	Simple effect	HF-L – LF-L	−0.40	115	>.999
	Simple effect	HF-M – LF-M	2.44	115	.097
	Simple effect	LF-H – LF-L	0.31	115	>.999
	Simple effect	LF-H – LF-M	0.30	115	>.999
	Simple effect	LF-L – LF-M	−0.01	115	>.999
Awareness of other passengers	DoA	Low – Medium	3.59	115	.001**
	DoA	Low – High	3.74	115	<.001***
	DoA	Medium – High	0.15	115	>.999
	FoV	Low – High	3.44	115	<.001***
	Simple effect	HF-H – HF-L	−3.82	115	.002**
	Simple effect	HF-H – HF-M	−0.11	115	>.999
	Simple effect	HF-H – LF-H	−3.37	115	.007**
	Simple effect	HF-L – HF-M	3.70	115	.003**
	Simple effect	HF-L – LF-L	−0.70	115	>.999
	Simple effect	HF-M – LF-M	−3.31	115	.007**
	Simple effect	LF-H – LF-L	−1.14	115	>.999
	Simple effect	LF-H – LF-M	−0.05	115	>.999
	Simple effect	LF-L – LF-M	1.09	115	>.999
Awareness of augmentations[†]	DoA	Low – Medium	−3.10	115	.007**
	DoA	Low – High	−3.62	115	.001**
	DoA	Medium – High	−0.52	115	>.999
	FoV	Low – High	−4.52	115	<.001***

[†]Manipulation check. [‡]Unadjusted (single comparison per DoA level). * $p < .05$; ** $p < .01$; *** $p < .001$.

Table D.4: Comparisons of each condition with the AR Baseline within a repeated-measures model over all eight conditions, Dunnett-adjusted for seven comparisons. For the overall Film-IEQ (parametric), the estimate is the mean difference in scale points; for all other measures (ART-C), it is the difference in aligned ranks. Positive estimates indicate that the listed condition scored higher than the AR Baseline. HF/LF: High/Low FoV; L/M/H: Low/Medium/High ornamentation.

Measure	Condition	Estimate	t	df	p
Film-IEQ	HF-L	7.12	2.40	23	.125
	HF-M	19.83	5.00	23	<.001***

Table D.4 continued

Measure	Condition	Estimate	t	df	p
	HF-H	22.83	4.98	23	<.001***
	LF-L	8.54	3.08	23	.030*
	LF-M	10.83	2.78	23	.057
	LF-H	11.71	3.44	23	.013*
	VR Baseline	27.71	5.51	23	<.001***
Captivation	HF-L	5.3	0.52	161	.971
	HF-M	29.7	2.91	161	.025*
	HF-H	35.0	3.43	161	.005**
	LF-L	14.3	1.40	161	.562
	LF-M	17.6	1.72	161	.364
	LF-H	15.7	1.54	161	.476
	VR Baseline	30.6	3.00	161	.019*
Real-World Dissociation	HF-L	19.4	1.79	161	.327
	HF-M	47.1	4.34	161	<.001***
	HF-H	55.8	5.15	161	<.001***
	LF-L	18.9	1.74	161	.353
	LF-M	20.8	1.92	161	.261
	LF-H	29.4	2.71	161	.044*
	VR Baseline	90.3	8.32	161	<.001***
Comprehension	HF-L	2.9	0.27	161	.997
	HF-M	17.2	1.60	161	.439
	HF-H	6.5	0.61	161	.954
	LF-L	4.1	0.39	161	.989
	LF-M	−2.6	−0.25	161	.997
	LF-H	1.8	0.17	161	.999
	VR Baseline	22.0	2.04	161	.208
Transportation	HF-L	30.4	3.11	161	.014*
	HF-M	61.5	6.29	161	<.001***
	HF-H	75.9	7.76	161	<.001***
	LF-L	24.4	2.50	161	.075
	LF-M	33.7	3.44	161	.005**
	LF-H	36.0	3.68	161	.002**

Table D.4 continued

Measure	Condition	Estimate	t	df	p
	VR Baseline	89.8	9.18	161	<.001***
UEQ-S	HF-L	20.1	1.95	161	.247
	HF-M	48.8	4.74	161	<.001***
	HF-H	62.6	6.08	161	<.001***
	LF-L	12.1	1.18	161	.704
	LF-M	12.9	1.26	161	.656
	LF-H	18.6	1.80	161	.320
	VR Baseline	57.4	5.57	161	<.001***
Pragmatic	HF-L	6.6	0.61	161	.952
	HF-M	20.6	1.91	161	.267
	HF-H	27.3	2.53	161	.070
	LF-L	−7.1	−0.66	161	.940
	LF-M	−11.6	−1.08	161	.764
	LF-H	−7.6	−0.71	161	.926
	VR Baseline	29.6	2.74	161	.040*
Hedonic	HF-L	21.4	2.10	161	.185
	HF-M	58.5	5.74	161	<.001***
	HF-H	77.8	7.63	161	<.001***
	LF-L	23.8	2.34	161	.110
	LF-M	26.1	2.56	161	.064
	LF-H	34.1	3.35	161	.007**
	VR Baseline	66.1	6.48	161	<.001***
UES-SF	HF-L	21.4	2.20	161	.148
	HF-M	51.7	5.33	161	<.001***
	HF-H	74.1	7.65	161	<.001***
	LF-L	13.8	1.42	161	.550
	LF-M	18.3	1.89	161	.275
	LF-H	25.4	2.62	161	.055
	VR Baseline	66.3	6.85	161	<.001***
Focused Attention	HF-L	22.9	2.30	161	.119
	HF-M	56.1	5.66	161	<.001***

Table D.4 continued

Measure	Condition	Estimate	t	df	p
	HF-H	71.3	7.19	161	<.001***
	LF-L	11.6	1.17	161	.710
	LF-M	28.1	2.83	161	.031*
	LF-H	26.0	2.62	161	.055
	VR Baseline	73.6	7.42	161	<.001***
Perceived Usability	HF-L	2.0	0.23	161	.998
	HF-M	−0.3	−0.04	161	>.999
	HF-H	−6.6	−0.78	161	.901
	LF-L	2.5	0.30	161	.995
	LF-M	13.2	1.56	161	.462
	LF-H	21.6	2.56	161	.065
	VR Baseline	0.4	0.05	161	>.999
Aesthetics	HF-L	21.0	2.05	161	.202
	HF-M	51.3	5.02	161	<.001***
	HF-H	70.2	6.87	161	<.001***
	LF-L	16.7	1.64	161	.413
	LF-M	13.3	1.30	161	.628
	LF-H	35.2	3.44	161	.005**
	VR Baseline	60.9	5.96	161	<.001***
Reward	HF-L	14.5	1.41	161	.559
	HF-M	43.8	4.24	161	<.001***
	HF-H	63.3	6.13	161	<.001***
	LF-L	18.5	1.79	161	.328
	LF-M	19.3	1.87	161	.285
	LF-H	21.9	2.12	161	.178
	VR Baseline	58.5	5.66	161	<.001***
Awareness of other passengers	HF-L	−16.5	−1.68	161	.387
	HF-M	−50.9	−5.19	161	<.001***
	HF-H	−53.1	−5.42	161	<.001***
	LF-L	−9.4	−0.96	161	.824
	LF-M	−20.1	−2.05	161	.202
	LF-H	−19.9	−2.03	161	.213

Table D.4 continued

Measure	Condition	Estimate	t	df	p
	VR Baseline	−107.9	−11.01	161	<.001***
Awareness of augmentations†	HF-L	60.5	5.27	161	<.001***
	HF-M	84.2	7.35	161	<.001***
	HF-H	99.2	8.65	161	<.001***
	LF-L	36.1	3.14	161	.012*
	LF-M	59.9	5.22	161	<.001***
	LF-H	61.9	5.40	161	<.001***
	VR Baseline	96.0	8.37	161	<.001***

Under the ART, the overall Film-IEQ comparisons for LF-L and LF-M were $p = .084$ and $p = .039$ respectively. †Manipulation check. * $p < .05$; ** $p < .01$; *** $p < .001$.

Table D.5: Comparisons of each condition with the VR Baseline within a repeated-measures model over all eight conditions, Dunnett-adjusted for seven comparisons. For the overall Film-IEQ (parametric), the estimate is the mean difference in scale points; for all other measures (ART-C), it is the difference in aligned ranks. Positive estimates indicate that the listed condition scored higher than the VR Baseline. HF/LF: High/Low FoV; L/M/H: Low/Medium/High ornamentation.

Measure	Condition	Estimate	t	df	p
Film-IEQ	HF-L	−20.58	−4.72	23	<.001***
	HF-M	−7.88	−2.52	23	.099
	HF-H	−4.88	−2.08	23	.221
	LF-L	−19.17	−4.57	23	<.001***
	LF-M	−16.88	−4.20	23	.002**
	LF-H	−16.00	−4.64	23	<.001***
	AR Baseline	−27.71	−5.51	23	<.001***
Captivation	HF-L	−25.3	−2.48	161	.079
	HF-M	−0.9	−0.09	161	>.999
	HF-H	4.4	0.43	161	.984
	LF-L	−16.3	−1.59	161	.440
	LF-M	−13.0	−1.27	161	.645
	LF-H	−14.9	−1.46	161	.525
	AR Baseline	−30.6	−3.00	161	.019*
Real-World Dissociation	HF-L	−70.9	−6.54	161	<.001***

Table D.5 continued

Measure	**Condition**	**Estimate**	t	df	p
	HF-M	−43.2	−3.98	161	<.001***
	HF-H	−34.5	−3.18	161	.011*
	LF-L	−71.4	−6.58	161	<.001***
	LF-M	−69.5	−6.41	161	<.001***
	LF-H	−60.9	−5.62	161	<.001***
	AR Baseline	−90.3	−8.32	161	<.001***
Comprehension	HF-L	−19.1	−1.77	161	.335
	HF-M	−4.8	−0.44	161	.982
	HF-H	−15.4	−1.43	161	.542
	LF-L	−17.8	−1.65	161	.404
	LF-M	−24.6	−2.29	161	.124
	LF-H	−20.1	−1.87	161	.284
	AR Baseline	−22.0	−2.04	161	.208
Transportation	HF-L	−59.4	−6.07	161	<.001***
	HF-M	−28.2	−2.89	161	.026*
	HF-H	−13.9	−1.42	161	.553
	LF-L	−65.4	−6.68	161	<.001***
	LF-M	−56.1	−5.74	161	<.001***
	LF-H	−53.8	−5.50	161	<.001***
	AR Baseline	−89.8	−9.18	161	<.001***
UEQ-S	HF-L	−37.3	−3.63	161	.003**
	HF-M	−8.6	−0.83	161	.882
	HF-H	5.2	0.50	161	.974
	LF-L	−45.2	−4.39	161	<.001***
	LF-M	−44.5	−4.32	161	<.001***
	LF-H	−38.8	−3.77	161	.002**
	AR Baseline	−57.4	−5.57	161	<.001***
Pragmatic	HF-L	−22.9	−2.12	161	.175
	HF-M	−9.0	−0.83	161	.883
	HF-H	−2.2	−0.21	161	.998
	LF-L	−36.7	−3.40	161	.005**
	LF-M	−41.2	−3.81	161	.001**

Table D.5 continued

Measure	**Condition**	**Estimate**	t	df	p
	LF-H	−37.2	−3.45	161	.005**
	AR Baseline	−29.6	−2.74	161	.040*
Hedonic	HF-L	−44.7	−4.38	161	<.001***
	HF-M	−7.6	−0.74	161	.915
	HF-H	11.7	1.15	161	.723
	LF-L	−42.2	−4.14	161	<.001***
	LF-M	−40.0	−3.92	161	<.001***
	LF-H	−32.0	−3.14	161	.013*
	AR Baseline	−66.1	−6.48	161	<.001***
UES-SF	HF-L	−45.0	−4.64	161	<.001***
	HF-M	−14.7	−1.51	161	.491
	HF-H	7.8	0.81	161	.892
	LF-L	−52.6	−5.42	161	<.001***
	LF-M	−48.0	−4.96	161	<.001***
	LF-H	−40.9	−4.22	161	<.001***
	AR Baseline	−66.3	−6.85	161	<.001***
Focused Attention	HF-L	−50.8	−5.12	161	<.001***
	HF-M	−17.5	−1.76	161	.341
	HF-H	−2.3	−0.24	161	.998
	LF-L	−62.0	−6.25	161	<.001***
	LF-M	−45.6	−4.59	161	<.001***
	LF-H	−47.6	−4.80	161	<.001***
	AR Baseline	−73.6	−7.42	161	<.001***
Perceived Usability	HF-L	1.5	0.18	161	.999
	HF-M	−0.8	−0.09	161	>.999
	HF-H	−7.1	−0.83	161	.882
	LF-L	2.1	0.24	161	.997
	LF-M	12.8	1.51	161	.495
	LF-H	21.2	2.50	161	.074
	AR Baseline	−0.4	−0.05	161	>.999
Aesthetics	HF-L	−39.9	−3.91	161	<.001***

Table D.5 continued

Measure	**Condition**	**Estimate**	t	df	p
	HF-M	−9.6	−0.94	161	.836
	HF-H	9.3	0.91	161	.848
	LF-L	−44.2	−4.32	161	<.001***
	LF-M	−47.6	−4.66	161	<.001***
	LF-H	−25.7	−2.52	161	.071
	AR Baseline	−60.9	−5.96	161	<.001***
Reward	HF-L	−44.0	−4.26	161	<.001***
	HF-M	−14.7	−1.42	161	.550
	HF-H	4.8	0.46	161	.980
	LF-L	−40.1	−3.88	161	.001**
	LF-M	−39.2	−3.79	161	.001**
	LF-H	−36.6	−3.55	161	.003**
	AR Baseline	−58.5	−5.66	161	<.001***
Awareness of other passengers	HF-L	91.4	9.33	161	<.001***
	HF-M	57.0	5.82	161	<.001***
	HF-H	54.8	5.59	161	<.001***
	LF-L	98.5	10.05	161	<.001***
	LF-M	87.8	8.96	161	<.001***
	LF-H	88.1	8.99	161	<.001***
	AR Baseline	107.9	11.01	161	<.001***
Awareness of augmentations[†]	HF-L	−35.5	−3.09	161	.015*
	HF-M	−11.7	−1.02	161	.793
	HF-H	3.3	0.29	161	.996
	LF-L	−59.9	−5.22	161	<.001***
	LF-M	−36.1	−3.15	161	.012*
	LF-H	−34.1	−2.97	161	.021*
	AR Baseline	−96.0	−8.37	161	<.001***

[†]Manipulation check. * $p < .05$; ** $p < .01$; *** $p < .001$.

Table D.6: Equivalence tests (TOST) for the High FoV–Medium and High FoV–High conditions against the VR Baseline. Bounds are $d_z = 0.60$ multiplied by the standard deviation of the paired differences, in raw scale units. Mean differences are condition minus VR Baseline, with 90% confidence intervals; p_t is the paired t-test.

Measure	**Condition**	**Diff.**	**90% CI**	**Bounds**	g_z	p_t	p_{TOST}
Film-IEQ	HF-M	−7.87	[−13.24, −2.51]	±9.16	−0.50	.019	.343
	HF-H	−4.88	[−8.89, −0.86]	±6.85	−0.41	.049	.204
Awareness of	HF-M	1.67	[1.21, 2.12]	–	1.23	<.001	.999
other passengers	HF-H	1.58	[1.07, 2.10]	–	1.04	<.001	.986

Neither Film-IEQ comparison established equivalence. The awareness comparisons are included for completeness: both conditions produced clearly greater awareness than the VR Baseline, so equivalence was not expected.

Table D.7: Exit questionnaire responses ($N = 24$): rankings of the eight conditions (1 = most preferred; columns give the median and mean rank and the number of participants ranking each condition first, in their top four, and last) and whether participants would use each condition instead of a phone or laptop to watch video on a train. HF/LF: High/Low FoV; L/M/H: Low/Medium/High ornamentation.

Condition	**Median**	**Mean**	**1st**	**Top 4**	**Last**	**Would use, n (%)**
HF-H	2	2.54	10	21	2	16 (67)
HF-M	3	3.04	3	21	0	20 (83)
VR Baseline	3	3.38	7	17	1	17 (71)
HF-L	4	4.12	1	16	1	19 (79)
LF-H	5	4.96	1	8	3	13 (54)
LF-M	6	5.54	0	6	0	12 (50)
LF-L	6	6.21	0	3	5	9 (38)
AR Baseline	7.5	6.21	2	4	12	17 (71)

Rankings: Friedman $\chi^2(7) = 58.0$, $p < .001$, Kendall's $W = .35$. Prospective use: Cochran's $Q(7) = 19.8$, $p = .006$.

Table D.8: Pairwise comparisons of the exit questionnaire, Holm-corrected across the 28 pairs. Rankings: Wilcoxon signed-rank tests; a positive median difference means Condition A was ranked lower (less preferred) than Condition B. Prospective use: exact McNemar tests on the discordant pairs (participants who would use only A, or only B).

		Rankings			**Prospective use**		
A	**B**	**Med. diff.**	V	p_{Holm}	**A only**	**B only**	p_{Holm}
LF-L	LF-M	1.0	223.0	.289	1	4	>.999
LF-L	LF-H	2.0	226.5	.289	2	6	>.999
LF-L	HF-L	3.0	264.5	.022*	1	11	.178
LF-L	HF-M	4.0	283.0	.004**	2	13	.199
LF-L	HF-H	5.0	278.5	.006**	3	10	>.999
LF-L	AR Baseline	−1.0	152.0	>.999	1	9	.559
LF-L	VR Baseline	3.0	265.0	.022*	1	9	.559
LF-M	LF-H	1.0	214.0	.519	2	3	>.999
LF-M	HF-L	2.0	234.0	.208	3	10	>.999
LF-M	HF-M	3.0	273.0	.011*	1	9	.559
LF-M	HF-H	4.0	279.5	.006**	3	7	>.999
LF-M	AR Baseline	−2.0	106.0	>.999	4	9	>.999
LF-M	VR Baseline	2.5	246.5	.092	5	10	>.999
LF-H	HF-L	1.5	187.0	>.999	5	11	>.999
LF-H	HF-M	2.5	244.0	.102	2	9	>.999
LF-H	HF-H	3.0	272.0	.011*	2	5	>.999
LF-H	AR Baseline	−2.0	94.0	.883	5	9	>.999
LF-H	VR Baseline	2.5	220.5	.447	5	9	>.999
HF-L	HF-M	1.0	247.5	.069	2	3	>.999
HF-L	HF-H	2.0	245.0	.094	6	3	>.999
HF-L	AR Baseline	−2.0	45.5	.050	4	2	>.999
HF-L	VR Baseline	1.0	194.0	>.999	4	2	>.999
HF-M	HF-H	1.0	198.0	>.999	5	1	>.999
HF-M	AR Baseline	−4.0	38.0	.027*	6	3	>.999
HF-M	VR Baseline	0.0	130.0	>.999	5	2	>.999
HF-H	AR Baseline	−5.5	41.5	.036*	5	6	>.999
HF-H	VR Baseline	−1.0	106.5	>.999	5	6	>.999
AR Baseline	VR Baseline	4.0	272.5	.011*	3	3	>.999

* $p < .05$; ** $p < .01$; *** $p < .001$.

Appendix E

Chapter 5 – ARTV Ornamentations for a New Genre – Complete Statistical Results

This appendix reports the complete Chapter 5 statistical output: descriptive statistics for each condition, omnibus ART ANOVAs, and all Tukey-corrected ART-C pairwise contrasts for all measures. The main results table in Chapter 5 reports only the significant subset. Conditions are abbreviated as: **BL** = Baseline; **SSA** = Show-Specific Augmentations (Constant Game of Thrones Objects); **GBA** = Genre-Based Augmentations (Constant Generic Medieval); **BSA** = Biome-Specific Augmentations (Interior Change With Scene).

Table E.1: Descriptive statistics (Median, IQR, Mean, SD) for all Chapter 5 measures by condition ($n = 20$ per condition). Film-IEQ subscale scores are summed item totals; UEQ-S subscales scored -3 to $+3$; UES-SF items were answered on a 7-point scale and scored 1–7, with Perceived Usability items reversed; awareness items scored 1–5.

	BL				**SSA**				**GBA**				**BSA**			
Measure	Med	IQR	Mean	SD	Med	IQR	Mean	SD	Med	IQR	Mean	SD	Med	IQR	Mean	SD
Film-IEQ Overall	92.5	16.5	93.4	15.2	104	17.5	103.0	16.7	99.5	28.2	95.1	20.1	112	16.8	110.0	22.1
Film-IEQ Captivation	51.0	9.75	51.0	9.42	56.0	11.8	54.2	8.97	55.5	17.8	49.0	12.6	58.5	9.00	56.2	13.6
Film-IEQ Real World Diss.	10.5	8.25	10.7	4.86	13.0	6.25	12.2	4.59	11.5	7.50	11.6	4.69	13.0	4.50	12.7	4.39
Film-IEQ Comprehension	19.0	3.75	18.8	2.85	20.0	5.25	19.0	3.39	19.5	7.25	18.5	4.49	21.5	4.00	20.2	3.13
Film-IEQ Transportation	12.0	4.50	12.8	4.76	19.0	9.75	17.8	6.87	16.0	7.50	16.0	5.59	23.0	6.75	21.0	6.83
UEQ-S Overall	1.00	1.19	1.06	0.873	2.00	1.09	1.68	0.928	1.69	1.34	1.39	1.20	2.38	0.56	1.92	1.29
UEQ-S Pragmatic	1.25	1.31	1.29	1.08	2.00	1.81	1.71	1.12	1.75	1.62	1.29	1.38	2.12	1.19	1.51	1.71
UEQ-S Hedonic	1.00	1.69	0.84	1.35	1.88	1.44	1.65	0.944	1.75	1.31	1.49	1.17	2.62	0.75	2.34	1.02
UES-SF Overall	4.54	1.00	4.52	0.913	5.54	0.88	5.25	1.01	5.29	0.89	5.10	1.12	6.17	0.79	5.92	1.30
UES-SF Focused Attention	4.00	1.17	4.00	1.26	5.00	1.17	4.63	1.23	4.83	1.50	4.42	1.33	6.00	1.00	5.67	1.25
UES-SF Perceived Usability	6.67	1.08	6.33	0.684	6.67	1.67	6.17	1.09	6.50	1.17	6.08	1.15	7.00	0.33	6.30	1.56
UES-SF Aesthetics	3.17	1.33	3.45	1.32	5.00	1.42	4.92	1.16	4.83	2.00	4.77	1.26	6.00	1.75	5.70	1.48
UES-SF Reward	4.50	1.75	4.28	1.31	5.33	1.08	5.28	1.11	5.33	1.50	5.13	1.38	6.17	1.33	6.00	1.32
Passenger Awareness	4.00	2.00	3.70	1.34	3.00	2.25	3.20	1.32	4.00	0.25	3.80	1.01	2.00	2.00	2.90	1.33
Augmentation Awareness	1.50	3.00	2.50	1.70	4.50	1.00	4.10	1.21	5.00	1.00	4.20	1.11	5.00	0.00	4.75	0.550

Table E.2: Omnibus ART ANOVA results for all Chapter 5 measures. F statistics have df $= (3, 57)$ throughout; η_p^2 is partial eta squared.

Measure	F	p	η_p^2	Sig.
Film-IEQ Overall	6.28	.0009	0.25	***
Film-IEQ Captivation	3.84	.0143	0.17	*
Film-IEQ Real World Diss.	1.18	.3249	0.06	
Film-IEQ Comprehension	0.95	.4234	0.05	
Film-IEQ Transportation	11.89	$< .001$	0.38	***
UEQ-S Overall	6.40	.0008	0.25	***
UEQ-S Pragmatic	1.82	.1533	0.09	
UEQ-S Hedonic	11.65	$< .001$	0.38	***
UES-SF Overall	11.96	$< .001$	0.39	***
UES-SF Focused Attention	12.80	$< .001$	0.40	***
UES-SF Perceived Usability	2.52	.0668	0.12	
UES-SF Aesthetics	12.39	$< .001$	0.39	***
UES-SF Reward	11.27	$< .001$	0.37	***
Passenger Awareness	4.34	.0080	0.19	**
Augmentation Awareness	14.09	$< .001$	0.43	***

Significance codes: $< .001$ ***, $< .01$ **, $< .05$ *

Table E.3: Complete ART-C pairwise contrasts for all Chapter 5 measures. Contrasts are Tukey-corrected; df = 57 throughout. Estimates, standard errors, and t are on the aligned-rank scale; g_z is Hedges' paired effect size computed from the raw scores. Negative values indicate the first-named condition scored lower. Significant contrasts are marked in the final column. Conditions: BL = Baseline; SSA = Show-Specific Augmentations; GBA = Genre-Based Augmentations; BSA = Biome-Specific Augmentations.

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
Film-IEQ Overall	BL − SSA	−13.62	6.07	57	−2.246	.1232	−0.54	
	BL − GBA	−4.83	6.07	57	−0.795	.8562	−0.07	
	BL − BSA	−24.45	6.07	57	−4.031	.0009	−0.58	***
	SSA − GBA	8.80	6.07	57	1.451	.4736	0.33	
	SSA − BSA	−10.82	6.07	57	−1.785	.2912	−0.27	
	GBA − BSA	−19.62	6.07	57	−3.235	.0106	−0.73	*
Film-IEQ Captivation	BL − SSA	−8.38	6.23	57	−1.344	.5393	−0.29	
	BL − GBA	0.58	6.23	57	0.092	.9997	0.16	
	BL − BSA	−17.80	6.23	57	−2.856	.0296	−0.29	*
	SSA − GBA	8.95	6.23	57	1.436	.4824	0.32	
	SSA − BSA	−9.43	6.23	57	−1.512	.4369	−0.13	
	GBA − BSA	−18.38	6.23	57	−2.949	.0233	−0.50	*
Film-IEQ Real World Diss.	BL − SSA	−7.62	5.72	57	−1.333	.5459	−0.27	
	BL − GBA	−4.55	5.72	57	−0.796	.8561	−0.16	
	BL − BSA	−10.22	5.72	57	−1.788	.2896	−0.35	
	SSA − GBA	3.08	5.72	57	0.538	.9494	0.12	
	SSA − BSA	−2.60	5.72	57	−0.455	.9685	−0.10	
	GBA − BSA	−5.67	5.72	57	−0.992	.7543	−0.25	
Film-IEQ Comprehension	BL − SSA	−2.55	6.17	57	−0.413	.9760	−0.04	
	BL − GBA	−0.95	6.17	57	−0.154	.9987	0.07	
	BL − BSA	−9.40	6.17	57	−1.523	.4305	−0.31	
	SSA − GBA	1.60	6.17	57	0.259	.9938	0.12	
	SSA − BSA	−6.85	6.17	57	−1.110	.6848	−0.29	

Continued on next page

Table E.3: Complete ART-C pairwise contrasts for all Chapter 5 measures (continued).

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
	GBA − BSA	−8.45	6.17	57	−1.369	.5234	−0.35	
Film-IEQ Transportation	BL − SSA	−16.90	4.88	57	−3.465	.0055	−0.76	**
	BL − GBA	−11.65	4.88	57	−2.388	.0909	−0.45	
	BL − BSA	−28.65	4.88	57	−5.873	< .001	−1.16	***
	SSA − GBA	5.25	4.88	57	1.076	.7053	0.33	
	SSA − BSA	−11.75	4.88	57	−2.409	.0868	−0.50	
	GBA − BSA	−17.00	4.88	57	−3.485	.0051	−0.89	**
UEQ-S Overall	BL − SSA	−14.20	5.63	57	−2.522	.0671	−0.51	
	BL − GBA	−9.32	5.63	57	−1.656	.3560	−0.21	
	BL − BSA	−24.18	5.63	57	−4.294	.0004	−0.49	***
	SSA − GBA	4.88	5.63	57	0.866	.8223	0.26	
	SSA − BSA	−9.97	5.63	57	−1.772	.2973	−0.20	
	GBA − BSA	−14.85	5.63	57	−2.638	.0511	−0.65	
UEQ-S Pragmatic	BL − SSA	−10.15	5.69	57	−1.783	.2921	−0.31	
	BL − GBA	−2.40	5.69	57	−0.422	.9746	0.00	
	BL − BSA	−10.75	5.69	57	−1.888	.2446	−0.11	
	SSA − GBA	7.75	5.69	57	1.361	.5285	0.32	
	SSA − BSA	−0.60	5.69	57	−0.105	.9996	0.12	
	GBA − BSA	−8.35	5.69	57	−1.467	.4641	−0.21	
UEQ-S Hedonic	BL − SSA	−12.53	5.31	57	−2.357	.0974	−0.54	
	BL − GBA	−10.60	5.31	57	−1.995	.2021	−0.33	
	BL − BSA	−30.88	5.31	57	−5.810	< .001	−0.79	***
	SSA − GBA	1.92	5.31	57	0.362	.9836	0.15	
	SSA − BSA	−18.35	5.31	57	−3.453	.0056	−0.61	**
	GBA − BSA	−20.27	5.31	57	−3.815	.0019	−0.83	**
UES-SF Overall	BL − SSA	−15.57	5.65	57	−2.756	.0381	−0.50	*
	BL − GBA	−13.05	5.65	57	−2.310	.1078	−0.33	
	BL − BSA	−33.58	5.65	57	−5.942	< .001	−0.74	***

Continued on next page

Table E.3: Complete ART-C pairwise contrasts for all Chapter 5 measures (continued).

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
	SSA − GBA	2.52	5.65	57	0.447	.9700	0.13	
	SSA − BSA	−18.00	5.65	57	−3.186	.0122	−0.57	*
	GBA − BSA	−20.52	5.65	57	−3.632	.0033	−1.00	**
UES-SF Focused Attention	BL − SSA	−11.50	5.67	57	−2.030	.1893	−0.35	
	BL − GBA	−8.07	5.67	57	−1.425	.4892	−0.20	
	BL − BSA	−33.52	5.67	57	−5.917	$< .001$	−0.82	***
	SSA − GBA	3.42	5.67	57	0.604	.9302	0.15	
	SSA − BSA	−22.02	5.67	57	−3.887	.0015	−1.02	**
	GBA − BSA	−25.45	5.67	57	−4.492	.0002	−1.11	***
UES-SF Perceived Usability	BL − SSA	−1.50	4.72	57	−0.318	.9888	0.15	
	BL − GBA	1.18	4.72	57	0.249	.9945	0.20	
	BL − BSA	−10.47	4.72	57	−2.220	.1301	0.02	
	SSA − GBA	2.67	4.72	57	0.567	.9415	0.08	
	SSA − BSA	−8.97	4.72	57	−1.902	.2386	−0.08	
	GBA − BSA	−11.65	4.72	57	−2.469	.0758	−0.24	
UES-SF Aesthetics	BL − SSA	−20.90	5.82	57	−3.598	.0037	−0.78	**
	BL − GBA	−18.90	5.82	57	−3.237	.0106	−0.58	*
	BL − BSA	−35.30	5.82	57	−6.062	$< .001$	−0.98	***
	SSA − GBA	2.10	5.82	57	0.361	.9838	0.10	
	SSA − BSA	−14.30	5.82	57	−2.464	.0766	−0.54	
	GBA − BSA	−16.40	5.82	57	−2.825	.0321	−0.76	*
UES-SF Reward	BL − SSA	−16.20	5.71	57	−2.835	.0313	−0.52	*
	BL − GBA	−14.93	5.71	57	−2.612	.0544	−0.39	
	BL − BSA	−33.17	5.71	57	−5.805	$< .001$	−0.78	***
	SSA − GBA	1.27	5.71	57	0.223	.9960	0.10	
	SSA − BSA	−16.98	5.71	57	−2.970	.0220	−0.68	*
	GBA − BSA	−18.25	5.71	57	−3.194	.0119	−0.69	*
Passenger Awareness	BL − SSA	9.35	4.98	57	1.878	.2491	0.33	

Continued on next page

Table E.3: Complete ART-C pairwise contrasts for all Chapter 5 measures (continued).

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
	BL − GBA	−0.75	4.98	57	−0.151	.9988	−0.07	
	BL − BSA	14.30	4.98	57	2.872	.0285	0.51	*
	SSA − GBA	−10.10	4.98	57	−2.028	.1898	−0.58	
	SSA − BSA	4.95	4.98	57	0.994	.7534	0.23	
	GBA − BSA	15.05	4.98	57	3.022	.0191	0.77	*
Augmentation Awareness	BL − SSA	−19.68	4.94	57	−3.984	.0011	−0.86	**
	BL − GBA	−21.55	4.94	57	−4.363	.0003	−1.02	***
	BL − BSA	−31.27	4.94	57	−6.332	$< .001$	−1.24	***
	SSA − GBA	−1.88	4.94	57	−0.380	.9812	−0.07	
	SSA − BSA	−11.60	4.94	57	−2.349	.0991	−0.49	
	GBA − BSA	−9.72	4.94	57	−1.969	.2118	−0.53	

Appendix F

Chapter 6 - ARTV in situ with Realistic Awareness Task – Complete Statistical Results

This appendix reports the complete Chapter 6 statistical output: medians for each condition, omnibus ART ANOVAs, all Tukey-corrected ART-C pairwise contrasts for the six-condition analysis, and the supplementary tests reported in the chapter. The main-results table in Chapter 6 reports only the significant subset, and the two-way (Medium $\times$ Ornamentation) analysis used for RQ4-1 is reported in full in the chapter itself.

Table F.1: Median scores per condition for Chapter 6 measures. UEQ-S subscales are scored from -3 to $+3$, UES-SF subscales from 1–7, and awareness items from 1–5. Awareness of the augmentations was only meaningful in the three conditions containing added visual content.

Condition	n	UEQ-S			UES-SF					Awareness	
		Prag.	Hed.	Overall	Foc. Att.	Perc. Usab.	Aesth.	Reward	Overall	Pass.	Aug.
AR Baseline	19	1.75	1.00	1.25	4.33	7.00	4.33	5.00	5.17	5	–
ARTV Ornamented	19	0.75	1.75	1.25	4.33	6.33	4.00	5.00	4.92	4	5
Sim. AR Baseline	19	1.50	0.75	0.75	4.67	7.00	4.00	4.33	4.83	3	–
Sim. ARTV Ornamented	19	1.00	2.50	1.75	5.33	6.33	5.00	5.67	5.42	3	5
VR Cinema	19	1.00	2.00	1.50	6.33	6.67	6.00	5.67	6.17	2	–
VR Videosphere	19	0.75	2.00	1.13	5.33	6.33	5.33	5.33	5.42	2	5

Table F.2: Omnibus ART ANOVA results for the six-condition analysis of Chapter 6. F statistics have df $= (5, 90)$ throughout; η_p^2 is partial eta squared.

Measure	F	p	η_p^2	Sig.
UEQ-S Overall	0.83	.5297	0.04	
UEQ-S Pragmatic	1.41	.2294	0.07	
UEQ-S Hedonic	5.72	$<$ .001	0.24	***
UES-SF Overall	5.28	$<$ .001	0.23	***
UES-SF Focused Attention	8.34	$<$ .001	0.32	***
UES-SF Perceived Usability	0.86	.5128	0.05	
UES-SF Aesthetics	5.53	$<$ .001	0.23	***
UES-SF Reward	3.11	.0122	0.15	*
Awareness: other passengers	6.13	$<$ .001	0.25	***

Table F.3: Supplementary tests reported in Chapter 6. Awareness of the augmentations was compared across the three conditions containing added visual content; the stop task was compared across all six; the bystander model covers the four physical-carriage conditions; and the exit-survey measures cover all six conditions for the 19 analysed participants.

Measure	Test	Statistic
Awareness of augmentations	Friedman	$\chi^2(2) = 10.8, W = .2$
Sim. ARTV Orn. vs VR Videosphere	Wilcoxon (Holm)	—
ARTV Orn. vs VR Videosphere	Wilcoxon (Holm)	—
ARTV Orn. vs Sim. ARTV Orn.	Wilcoxon (Holm)	—
Stop-task failures	Cochran's Q	$Q(5) = 1.05$
Bystander count (passenger awareness)	Cumulative link mixed model, likelihood ratio	$\chi^2(1) = 4.37$
Preference rankings	Friedman	$\chi^2(5) = 6.34, W = .0$
Usage intent	Friedman	$\chi^2(5) = 4.95, W = .0$

Table F.4: Complete ART-C pairwise contrasts for the six-condition analysis of Chapter 6. Contrasts are Tukey-corrected. Estimates, standard errors and t are on the aligned-rank scale; g_z is Hedges' paired effect size computed from the raw scores. Negative values indicate that the first-named condition scored lower. Significant contrasts are marked in the final column.

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
UEQ-S Overall	AR Baseline − ARTV Ornamented	-2.37	8.89	90	-0.27	.9998	-0.10	
	AR Baseline − VR Cinema	-3.08	8.89	90	-0.35	.9993	-0.16	
	AR Baseline − Sim. AR Baseline	9.34	8.89	90	1.05	.8990	0.34	
	AR Baseline − Sim. ARTV Ornamented	-7.45	8.89	90	-0.84	.9596	-0.13	
	AR Baseline − VR Videosphere	-4.03	8.89	90	-0.45	.9975	-0.17	
	ARTV Ornamented − VR Cinema	-0.71	8.89	90	-0.08	1.0000	-0.02	
	ARTV Ornamented − Sim. AR Baseline	11.71	8.89	90	1.32	.7745	0.26	
	ARTV Ornamented − Sim. ARTV Ornamented	-5.08	8.89	90	-0.57	.9926	-0.06	
	ARTV Ornamented − VR Videosphere	-1.66	8.89	90	-0.19	1.0000	-0.06	
	VR Cinema − Sim. AR Baseline	12.42	8.89	90	1.40	.7283	0.47	
	VR Cinema − Sim. ARTV Ornamented	-4.37	8.89	90	-0.49	.9964	-0.02	
	VR Cinema − VR Videosphere	-0.95	8.89	90	-0.11	1.0000	-0.05	
	Sim. AR Baseline − Sim. ARTV Ornamented	-16.79	8.89	90	-1.89	.4155	-0.32	
	Sim. AR Baseline − VR Videosphere	-13.37	8.89	90	-1.50	.6625	-0.39	
	Sim. ARTV Ornamented − VR Videosphere	3.42	8.89	90	0.38	.9989	0.00	
UEQ-S Pragmatic	AR Baseline − ARTV Ornamented	18.84	9.50	90	1.98	.3595	0.53	
	AR Baseline − VR Cinema	14.39	9.50	90	1.52	.6551	0.31	
	AR Baseline − Sim. AR Baseline	11.13	9.50	90	1.17	.8488	0.33	
	AR Baseline − Sim. ARTV Ornamented	20.21	9.50	90	2.13	.2824	0.56	
	AR Baseline − VR Videosphere	21.16	9.50	90	2.23	.2354	0.48	
	ARTV Ornamented − VR Cinema	-4.45	9.50	90	-0.47	.9971	-0.11	
	ARTV Ornamented − Sim. AR Baseline	-7.71	9.50	90	-0.81	.9646	-0.17	
	ARTV Ornamented − Sim. ARTV Ornamented	1.37	9.50	90	0.14	1.0000	0.01	
	ARTV Ornamented − VR Videosphere	2.32	9.50	90	0.24	.9999	0.01	
	VR Cinema − Sim. AR Baseline	-3.26	9.50	90	-0.34	.9993	-0.06	

Continued on next page

Table F.4: Complete ART-C pairwise contrasts for all Chapter 6 measures (continued).

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
	VR Cinema − Sim. ARTV Ornamented	5.82	9.50	90	0.61	.9899	0.13	
	VR Cinema − VR Videosphere	6.76	9.50	90	0.71	.9800	0.23	
	Sim. AR Baseline − Sim. ARTV Ornamented	9.08	9.50	90	0.96	.9304	0.22	
	Sim. AR Baseline − VR Videosphere	10.03	9.50	90	1.06	.8973	0.21	
	Sim. ARTV Ornamented − VR Videosphere	0.95	9.50	90	0.10	1.0000	0.01	
UEQ-S Hedonic	AR Baseline − ARTV Ornamented	-20.89	7.82	90	-2.67	.0911	-0.62	
	AR Baseline − VR Cinema	-17.97	7.82	90	-2.30	.2056	-0.83	
	AR Baseline − Sim. AR Baseline	2.76	7.82	90	0.35	.9993	0.08	
	AR Baseline − Sim. ARTV Ornamented	-27.50	7.82	90	-3.52	.0087	-0.58	**
	AR Baseline − VR Videosphere	-26.55	7.82	90	-3.40	.0127	-0.78	*
	ARTV Ornamented − VR Cinema	2.92	7.82	90	0.37	.9990	0.13	
	ARTV Ornamented − Sim. AR Baseline	23.66	7.82	90	3.03	.0370	0.64	*
	ARTV Ornamented − Sim. ARTV Ornamented	-6.61	7.82	90	-0.84	.9582	-0.13	
	ARTV Ornamented − VR Videosphere	-5.66	7.82	90	-0.72	.9786	-0.18	
	VR Cinema − Sim. AR Baseline	20.74	7.82	90	2.65	.0956	0.79	
	VR Cinema − Sim. ARTV Ornamented	-9.53	7.82	90	-1.22	.8269	-0.18	
	VR Cinema − VR Videosphere	-8.58	7.82	90	-1.10	.8814	-0.40	
	Sim. AR Baseline − Sim. ARTV Ornamented	-30.26	7.82	90	-3.87	.0028	-0.65	**
	Sim. AR Baseline − VR Videosphere	-29.32	7.82	90	-3.75	.0041	-0.81	**
	Sim. ARTV Ornamented − VR Videosphere	0.95	7.82	90	0.12	1.0000	-0.01	
UES-SF Overall	AR Baseline − ARTV Ornamented	-1.66	8.49	90	-0.20	1.0000	-0.13	
	AR Baseline − VR Cinema	-33.68	8.49	90	-3.97	.0020	-1.29	**
	AR Baseline − Sim. AR Baseline	-0.84	8.49	90	-0.10	1.0000	-0.05	
	AR Baseline − Sim. ARTV Ornamented	-20.79	8.49	90	-2.45	.1505	-0.37	
	AR Baseline − VR Videosphere	-17.24	8.49	90	-2.03	.3331	-0.52	
	ARTV Ornamented − VR Cinema	-32.03	8.49	90	-3.77	.0038	-0.69	**
	ARTV Ornamented − Sim. AR Baseline	0.82	8.49	90	0.10	1.0000	0.08	
	ARTV Ornamented − Sim. ARTV Ornamented	-19.13	8.49	90	-2.25	.2238	-0.30	
	ARTV Ornamented − VR Videosphere	-15.58	8.49	90	-1.84	.4483	-0.35	

Continued on next page

Table F.4: Complete ART-C pairwise contrasts for all Chapter 6 measures (continued).

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
	VR Cinema − Sim. AR Baseline	32.84	8.49	90	3.87	.0027	1.04	**
	VR Cinema − Sim. ARTV Ornamented	12.89	8.49	90	1.52	.6526	0.25	
	VR Cinema − VR Videosphere	16.45	8.49	90	1.94	.3861	0.56	
	Sim. AR Baseline − Sim. ARTV Ornamented	-19.95	8.49	90	-2.35	.1851	-0.36	
	Sim. AR Baseline − VR Videosphere	-16.39	8.49	90	-1.93	.3897	-0.43	
	Sim. ARTV Ornamented − VR Videosphere	3.55	8.49	90	0.42	.9983	0.04	
UES-SF Focused Attention	AR Baseline − ARTV Ornamented	-5.32	8.27	90	-0.64	.9874	-0.27	
	AR Baseline − VR Cinema	-44.37	8.27	90	-5.36	$<$ .0001	-1.79	***
	AR Baseline − Sim. AR Baseline	-8.18	8.27	90	-0.99	.9202	-0.32	
	AR Baseline − Sim. ARTV Ornamented	-26.50	8.27	90	-3.20	.0224	-0.60	*
	AR Baseline − VR Videosphere	-26.95	8.27	90	-3.26	.0191	-0.64	*
	ARTV Ornamented − VR Cinema	-39.05	8.27	90	-4.72	$<$ .001	-1.25	***
	ARTV Ornamented − Sim. AR Baseline	-2.87	8.27	90	-0.35	.9993	-0.06	
	ARTV Ornamented − Sim. ARTV Ornamented	-21.18	8.27	90	-2.56	.1179	-0.46	
	ARTV Ornamented − VR Videosphere	-21.63	8.27	90	-2.62	.1041	-0.52	
	VR Cinema − Sim. AR Baseline	36.18	8.27	90	4.37	$<$ .001	1.26	***
	VR Cinema − Sim. ARTV Ornamented	17.87	8.27	90	2.16	.2667	0.39	
	VR Cinema − VR Videosphere	17.42	8.27	90	2.11	.2935	0.56	
	Sim. AR Baseline − Sim. ARTV Ornamented	-18.32	8.27	90	-2.21	.2415	-0.37	
	Sim. AR Baseline − VR Videosphere	-18.76	8.27	90	-2.27	.2179	-0.34	
	Sim. ARTV Ornamented − VR Videosphere	-0.45	8.27	90	-0.05	1.0000	0.03	
UES-SF Perceived Usability	AR Baseline − ARTV Ornamented	11.11	8.14	90	1.36	.7481	0.39	
	AR Baseline − VR Cinema	8.87	8.14	90	1.09	.8844	0.15	
	AR Baseline − Sim. AR Baseline	4.03	8.14	90	0.49	.9963	0.17	
	AR Baseline − Sim. ARTV Ornamented	15.16	8.14	90	1.86	.4322	0.42	
	AR Baseline − VR Videosphere	8.84	8.14	90	1.09	.8857	0.19	
	ARTV Ornamented − VR Cinema	-2.24	8.14	90	-0.27	.9998	-0.21	
	ARTV Ornamented − Sim. AR Baseline	-7.08	8.14	90	-0.87	.9528	-0.14	
	ARTV Ornamented − Sim. ARTV Ornamented	4.05	8.14	90	0.50	.9961	-0.01	

Continued on next page

Table F.4: Complete ART-C pairwise contrasts for all Chapter 6 measures (continued).

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
	ARTV Ornamented − VR Videosphere	-2.26	8.14	90	-0.28	.9998	-0.24	
	VR Cinema − Sim. AR Baseline	-4.84	8.14	90	-0.59	.9911	0.04	
	VR Cinema − Sim. ARTV Ornamented	6.29	8.14	90	0.77	.9715	0.17	
	VR Cinema − VR Videosphere	-0.03	8.14	90	-0.00	1.0000	0.02	
	Sim. AR Baseline − Sim. ARTV Ornamented	11.13	8.14	90	1.37	.7462	0.14	
	Sim. AR Baseline − VR Videosphere	4.82	8.14	90	0.59	.9914	-0.02	
	Sim. ARTV Ornamented − VR Videosphere	-6.32	8.14	90	-0.78	.9709	-0.19	
UES-SF Aesthetics	AR Baseline − ARTV Ornamented	-4.39	8.36	90	-0.53	.9950	-0.19	
	AR Baseline − VR Cinema	-34.53	8.36	90	-4.13	.0011	-1.10	**
	AR Baseline − Sim. AR Baseline	-0.58	8.36	90	-0.07	1.0000	-0.00	
	AR Baseline − Sim. ARTV Ornamented	-21.39	8.36	90	-2.56	.1186	-0.38	
	AR Baseline − VR Videosphere	-19.00	8.36	90	-2.27	.2165	-0.56	
	ARTV Ornamented − VR Cinema	-30.13	8.36	90	-3.60	.0066	-0.65	**
	ARTV Ornamented − Sim. AR Baseline	3.82	8.36	90	0.46	.9974	0.16	
	ARTV Ornamented − Sim. ARTV Ornamented	-17.00	8.36	90	-2.03	.3324	-0.33	
	ARTV Ornamented − VR Videosphere	-14.61	8.36	90	-1.75	.5053	-0.32	
	VR Cinema − Sim. AR Baseline	33.95	8.36	90	4.06	.0014	0.91	**
	VR Cinema − Sim. ARTV Ornamented	13.13	8.36	90	1.57	.6201	0.26	
	VR Cinema − VR Videosphere	15.53	8.36	90	1.86	.4355	0.57	
	Sim. AR Baseline − Sim. ARTV Ornamented	-20.82	8.36	90	-2.49	.1384	-0.40	
	Sim. AR Baseline − VR Videosphere	-18.42	8.36	90	-2.20	.2468	-0.47	
	Sim. ARTV Ornamented − VR Videosphere	2.39	8.36	90	0.29	.9997	0.05	
UES-SF Reward	AR Baseline − ARTV Ornamented	-5.61	8.17	90	-0.69	.9831	-0.22	
	AR Baseline − VR Cinema	-22.34	8.17	90	-2.73	.0783	-0.94	
	AR Baseline − Sim. AR Baseline	3.16	8.17	90	0.39	.9988	0.07	
	AR Baseline − Sim. ARTV Ornamented	-18.61	8.17	90	-2.28	.2141	-0.37	
	AR Baseline − VR Videosphere	-12.03	8.17	90	-1.47	.6827	-0.37	
	ARTV Ornamented − VR Cinema	-16.74	8.17	90	-2.05	.3236	-0.33	
	ARTV Ornamented − Sim. AR Baseline	8.76	8.17	90	1.07	.8909	0.27	

Continued on next page

Table F.4: Complete ART-C pairwise contrasts for all Chapter 6 measures (continued).

Measure	Comparison	Est.	SE	df	t	p	g_z	Sig.
	ARTV Ornamented − Sim. ARTV Ornamented	-13.00	8.17	90	-1.59	.6063	-0.23	
	ARTV Ornamented − VR Videosphere	-6.42	8.17	90	-0.79	.9692	-0.04	
	VR Cinema − Sim. AR Baseline	25.50	8.17	90	3.12	.0283	0.85	*
	VR Cinema − Sim. ARTV Ornamented	3.74	8.17	90	0.46	.9974	0.06	
	VR Cinema − VR Videosphere	10.32	8.17	90	1.26	.8042	0.43	
	Sim. AR Baseline − Sim. ARTV Ornamented	-21.76	8.17	90	-2.66	.0929	-0.46	
	Sim. AR Baseline − VR Videosphere	-15.18	8.17	90	-1.86	.4341	-0.38	
	Sim. ARTV Ornamented − VR Videosphere	6.58	8.17	90	0.81	.9659	0.16	
Awareness: other passengers	AR Baseline − ARTV Ornamented	11.37	8.49	90	1.34	.7627	0.56	
	AR Baseline − VR Cinema	37.29	8.49	90	4.39	$< .001$	1.02	***
	AR Baseline − Sim. AR Baseline	28.74	8.49	90	3.38	.0131	0.85	*
	AR Baseline − Sim. ARTV Ornamented	23.16	8.49	90	2.73	.0798	0.71	
	AR Baseline − VR Videosphere	37.29	8.49	90	4.39	$< .001$	0.96	***
	ARTV Ornamented − VR Cinema	25.92	8.49	90	3.05	.0343	0.70	*
	ARTV Ornamented − Sim. AR Baseline	17.37	8.49	90	2.05	.3256	0.54	
	ARTV Ornamented − Sim. ARTV Ornamented	11.79	8.49	90	1.39	.7339	0.36	
	ARTV Ornamented − VR Videosphere	25.92	8.49	90	3.05	.0343	0.63	*
	VR Cinema − Sim. AR Baseline	-8.55	8.49	90	-1.01	.9144	-0.17	
	VR Cinema − Sim. ARTV Ornamented	-14.13	8.49	90	-1.66	.5589	-0.27	
	VR Cinema − VR Videosphere	-0.00	8.49	90	-0.00	1.0000	0.00	
	Sim. AR Baseline − Sim. ARTV Ornamented	-5.58	8.49	90	-0.66	.9861	-0.14	
	Sim. AR Baseline − VR Videosphere	8.55	8.49	90	1.01	.9144	0.17	
	Sim. ARTV Ornamented − VR Videosphere	14.13	8.49	90	1.66	.5589	0.31	

Bibliography

[1] Microsoft hololens official announcement trailer hd youtube, 2024. URL https://www.youtube.com/watch?v=qyzzdMkU7Fw.

[2] J. et al. Abrams. Non-invasive neural interfaces for human-computer interaction. *Nature Communications*, 12:1–11, 2021.

[3] Adobe. Mixamo. https://www.mixamo.com/, 2023.

[4] Medhat Alaker, Greg R. Wynn, and Tan Arulampalam. Virtual reality training in laparoscopic surgery: A systematic review & meta-analysis. *International Journal of Surgery*, 29:85–94, 2016. doi: 10.1016/j.ijsu.2016.03.034. Listed as free article in PubMed.

[5] Apple. Apple vision pro. https://www.apple.com/apple-vision-pro/, 2024.

[6] Apple Inc. Apple introduces spatial video capture on iPhone 15 Pro. https://www.apple.com/uk/newsroom/2023/12/apple-introduces-spatial-video-capture-on-iphone-15-pro/, December 2023. Accessed: 2026-09-12.

[7] Apple Inc. Apple previews new entertainment experiences launching with Apple Vision Pro. Apple Newsroom, January 2024. URL https://www.apple.com/newsroom/2024/01/apple-previews-new-entertainment-experiences-launching-with-apple Accessed: 2026-09-15.

[8] Apple Inc. Use a spatial Persona in FaceTime on Apple Vision Pro. https://support.apple.com/en-gb/guide/apple-vision-pro/tana1ea03f18/visionos, 2024. Apple Vision Pro User Guide. Accessed: 2026-09-12.

[9] M. Armstrong. Smartphones are taking over the global games market, 2018. URL https://www.statista.com/chart/12955/smartphones-are-taking-over-the-global-games-market/.

[10] Sarah Atkinson and Helen W. Kennedy. From conflict to revolution: The secret aesthetic, narrative spatialisation and audience experience in immersive cinema design. *Participations: Journal of Audience & Reception Studies*, 13(1):252–279, May 2016. Published May 2016.

[11] Sarah Atkinson and Helen W. Kennedy, editors. *Live Cinema: Cultures, Economies, Aesthetics*. Bloomsbury Academic, New York, 2017. ISBN 9781501324833.

[12] Sarah Atkinson and Helen W. Kennedy. *Secret Cinema and the immersive experience economy*. Manchester University Press, 2022. ISBN 9781526140173. URL http://www.jstor.org/stable/jj.21996671.

[13] Audi, Feb 2024. URL https://www.audi.com/en/innovation/future-technology/virtual-reality/holoride-virtual-reality-meets-the-real-world.html.

[14] Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. Whisperx: Time-accurate speech transcription of long-form audio. *INTERSPEECH 2023*, 2023. doi: 10.21437/interspeech.2023-78.

[15] Laura Bajorunaite. *Reality anchors: investigating the use of reality cues for socially acceptable immersive technologies in transit*. Phd thesis, University of Glasgow, 2025. URL https://theses.gla.ac.uk/85221/. Available at University of Glasgow Enlighten Theses repository.

[16] Laura Bajorunaite, Stephen Brewster, and Julie R. Williamson. Virtual reality in transit: how acceptable is vr use on public transport? In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pages 432–433, 2021. doi: 10.1109/VRW52623.2021.00098.

[17] Laura Bajorunaite, Stephen Brewster, and Julie R. Williamson. Reality anchors: Bringing cues from reality to increase acceptance of immersive technologies in transit. volume 7, New York, NY, USA, September 2023. Association for Computing Machinery. doi: 10.1145/3604266. URL https://doi.org/10.1145/3604266.

[18] Laura Bajorunaite, Julie R. Williamson, and Stephen Brewster. Vr headsets in-the-wild: Qualitative insights on safety, awareness, and social challenges from real train journeys. In *Proceedings of the 23rd International Conference on Mobile and Ubiquitous Multimedia*, MUM '24, page 82–94, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400712838. doi: 10.1145/3701571.3701576. URL https://doi.org/10.1145/3701571.3701576.

[19] BBC Natural History Unit. Planet earth ii. Television series, BBC One, 2016. Narrated by David Attenborough; series producer Tom Hugh-Jones.

[20] David Benioff and D.B. Weiss. Game of Thrones. Television series, 2011–2019. Created by David Benioff and D.B. Weiss.

[21] Hugh Beyer and Karen Holtzblatt. *Contextual Design: Defining Customer-Centered Systems*. Morgan Kaufmann, San Francisco, CA, 1998. ISBN 9781558604117.

[22] Bigscreen. Bigscreen vr cinema software. `https://www.bigscreenvr.com/software`, 2026.

[23] Mark Billinghurst, Adrian Clark, and Gun Lee. A survey of augmented reality. *Found. Trends Hum.-Comput. Interact.*, 8(2–3):73–272, March 2015. ISSN 1551-3955. doi: 10.1561/1100000049. URL `https://doi.org/10.1561/1100000049`.

[24] Doug A. Bowman and Ryan P. McMahan. Virtual reality: How much immersion is enough? *Computer*, 40(7):36–43, 2007. doi: 10.1109/MC.2007.257.

[25] Laura Ceci. Global wellness app downloads by category 2020, Sep 2021. URL `https://www.statista.com/statistics/1262405/global-downloads-health-and-medical-apps-by-category/`.

[26] Chiao-Ju Chang, Yu Lun Hsu, Wei Tian Mireille Tan, Yu-Cheng Chang, Pin Chun Lu, Yu Chen, Yi-Han Wang, and Mike Y. Chen. Exploring augmented reality interface designs for virtual meetings in real-world walking contexts. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, DIS '24, page 391–408, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400705830. doi: 10.1145/3643834.3661538. URL `https://doi.org/10.1145/3643834.3661538`.

[27] Jingwei Chen, Shengliang Zhang, Sheng Cao, Mengjie Zhu, Wu Zhang, Abdulhakem Y. Elezzabi, Bin Wang, Linhua Liu, William W. Yu, and Haizeng Li. Electrochromic devices for dimming and display in augmented reality smart glasses. *Device*, 2(12):100470, 2024. ISSN 2666-9986. doi: https://doi.org/10.1016/j.device.2024.100470. URL `https://www.sciencedirect.com/science/article/pii/S2666998624003454`.

[28] Y. F. Cheng, A. Carden, H. Cho, C. G. Fidalgo, J. Wieland, and D. Lindlbauer. Augmented reality productivity in-the-wild: A diary study of usage patterns and experiences of working with ar laptops in real-world settings. *IEEE Transactions on Visualization and Computer Graphics*, 31(10):9195–9212, 2025. doi: 10.1109/TVCG.2025.3592962.

[29] Yifei Cheng, Yukang Yan, Xin Yi, Yuanchun Shi, and David Lindlbauer. Semanticadapt: Optimization-based adaptation of mixed reality layouts leveraging virtual-physical semantic connections. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST '21, page 282–297, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450386357. doi: 10.1145/3472749.3474750. URL https://doi.org/10.1145/3472749.3474750.

[30] Rune Haubo Bojesen Christensen. *ordinal: Regression Models for Ordinal Data*. URL https://CRAN.R-project.org/package=ordinal.

[31] Iain Christie, Stephen Brewster, and Mark McGill. Genre-based, show-specific, and biome-specific artv ornamentations: Raw dataset from a vr-simulated passenger train study, 2025. URL https://doi.org/10.5281/zenodo.20574411.

[32] Iain Christie, Stephen brewster, and Mark McGill. Artv ornamentations in situ with realistic awareness task: Raw dataset from a physical subway carriage study, 2026. URL https://doi.org/10.5281/zenodo.20583166.

[33] Iain Christie, Stephen Brewster, and Mark McGill. Commuter survey data - thesis chapter 3, 2026. URL https://doi.org/10.5281/zenodo.20573749.

[34] Devin Coldewey. StarVR's one headset flaunts eye-tracking and a double-wide field of view. https://techcrunch.com/2018/08/14/starvrs-one-headset-flaunts-eye-tracking-and-a-double-wide-field- August 2018. Accessed: 2026-09-08.

[35] Yannick Cornet, Giuseppe Lugano, Christina Georgouli, and Dimitris Milakis. Worthwhile travel time: a conceptual framework of the perceived value of enjoyment, productivity and fitness while travelling. *Transport Reviews*, 42(5):580–603, 2022. doi: 10.1080/01441647.2021.1983067. URL https://doi.org/10.1080/01441647.2021.1983067.

[36] Caroline Criado Perez. *Invisible Women: Exposing Data Bias in a World Designed for Men*. Chatto & Windus, London, 2019. ISBN 9781784741723.

[37] James J. Cummings and Jeremy N. Bailenson. How immersive is enough? a meta-analysis of the effect of immersive technology on user presence. *Media Psychology*, 19 (2):272–309, 2016. doi: 10.1080/15213269.2015.1015740.

[38] Scott Davies. Calculating the ideal screen size from your seating position, Jul 2022. URL https://www.customht.com.au/blogs/ht-hifi/calculating-the-ideal-screen-size-from-your-seating-position.

[39] Cyriel Diels and Jelte E. Bos. Self-driving carsickness. *Applied Ergonomics*, 53:374–382, March 2016. ISSN 0003-6870. doi: 10.1016/j.apergo.2015.09.009. URL https://www.sciencedirect.com/science/article/pii/S0003687015300818.

[40] Martin Dodd, Iain Knight, Alix Edwards, Mervyn Edwards, Peter Ball, Josh Appleby, and Jack Radcliffe. Bus safety standard: Bus braking data analysis. Technical Report PPR1013, TRL Limited, 2022. URL https://www.trl.co.uk/uploads/trl/documents/PPR1013---Bus-Safety-Standard---Bus-Braking-Data-Analysis_1.pdf. Prepared for Transport for London (TfL); Report date: 03/02/2022.

[41] Lorin Dole and Wendy Ju. Face and ecological validity in simulations: Lessons from search-and-rescue hri. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, page 1–8, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450359702. doi: 10.1145/3290605.3300681. URL https://doi.org/10.1145/3290605.3300681.

[42] Pouya Eghbali, Kaisa Väänänen, and Tero Jokela. Social acceptability of virtual reality in public spaces: experiential factors and design recommendations. In *Proceedings of the 18th International Conference on Mobile and Ubiquitous Multimedia*, MUM '19, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450376242. doi: 10.1145/3365610.3365647. URL https://doi.org/10.1145/3365610.3365647.

[43] Lisa A. Elkin, Matthew Kay, James J. Higgins, and Jacob O. Wobbrock. An aligned rank transform procedure for multifactor contrast tests. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST '21, page 754–768, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450386357. doi: 10.1145/3472749.3474784. URL https://doi.org/10.1145/3472749.3474784.

[44] Mica R. Endsley. Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1):32–64, 1995. doi: 10.1518/001872095779049543. URL https://doi.org/10.1518/001872095779049543.

[45] Dick Ettema, Tommy Gärling, Lars E. Olsson, and Margareta Friman. Out-of-home activities, daily travel, and subjective well-being. *Transportation Research Part A: Policy and Practice*, 44(9):723–732, 2010. ISSN 0965-8564. doi: https://doi.org/10.1016/j.tra.2010.07.005. URL https://www.sciencedirect.com/science/article/pii/S0965856410000960.

[46] João Marcelo Evangelista Belo, Mathias N. Lystbæk, Anna Maria Feit, Ken Pfeuffer, Peter Kán, Antti Oulasvirta, and Kaj Grønbæk. Auit – the adaptive user interfaces toolkit for designing xr applications. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, UIST '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393201. doi: 10.1145/3526113.3545651. URL https://doi.org/10.1145/3526113.3545651.

[47] Gary W Evans and Richard E Wener. Rail commuting duration and passenger stress. *Health psychology : official journal of the Division of Health Psychology, American Psychological Association*, 25(3):408–412, 2006. doi: doi:10.1037/0278-6133.25.3.408.

[48] Office for National Statistics (ONS). Office for national statistics: Travel survey 2022, Aug 2023. URL https://www.gov.uk/government/statistics/national-travel-survey-2022.

[49] Daniel Freeman, Katherine Pugh, Angus Antley, Mel Slater, Paul Bebbington, Matthew Gittins, Graham Dunn, Elizabeth Kuipers, David Fowler, and Philippa Garety. Virtual reality study of paranoid thinking in the general population. *British Journal of Psychiatry*, 192(4):258–263, 2008. doi: 10.1192/bjp.bp.107.044677.

[50] Damian Frej, Paweł Grabski, and Emilia M. Szumska. Experimental study on longitudinal acceleration of urban buses and coaches in different road maneuvers. *Sensors*, 23(6): 3125, 2023. doi: 10.3390/s23063125.

[51] Gensler Research Institute and Immersive Experience Institute. Evolving immersive: The 2025 immersive entertainment & culture industry report. Technical report, Gensler, 2025. URL https://www.gensler.com/gri/immersive-industry-report-2025. Accessed: 2026-01-19.

[52] Yiannis Georgiou and Eleni A. Kyza. The development and validation of the ari questionnaire. *Int. J. Hum.-Comput. Stud.*, 98(C):24–37, February 2017. ISSN 1071-5819. doi: 10.1016/j.ijhcs.2016.09.014. URL https://doi.org/10.1016/j.ijhcs.2016.09.014.

[53] Florin-Timotei Ghiurău, Mehmet Aydın Baytaş, and Casper Wickman. Arcar: On-road driving in mixed reality by volvo cars. In *Adjunct Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, UIST '20 Adjunct, page 62–64, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450375153. doi: 10.1145/3379350.3416186. URL https://doi.org/10.1145/3379350.3416186.

[54] Glasgow Museums. Subway car 39t. Glasgow Museums Collections Navigator, object MTEMP.6097. Maker: Hurst, Nelson & Company, Motherwell, 1898. On display at the

Riverside Museum, Glasgow. URL https://collections.glasgowmuseums.com/mwebcgi/mweb?request=record;id=27588;type=101. Accessed [11 09 2026].

[55] Karl M. Guttag. HoloLens 2 display evaluation (part 6: Microsoft's FUD on photographs). KGOnTech, September 2020. URL https://kguttag.com/2020/09/06/hololens-2-display-evaluation-part-6-microsofts-fud-on-photograph Blog post. Accessed 10 September 2026.

[56] Erik Hansson, Kristoffer Mattisson, Jonas Björk, Per-Olof Östergren, and Kristina Jakobsson. Relationship between commuting and health outcomes in a cross-sectional population survey in southern sweden. *BMC Public Health*, 11:834, 2011. doi: 10.1186/1471-2458-11-834. URL https://doi.org/10.1186/1471-2458-11-834. Open Access; PMCID: PMC3226527.

[57] Devindra Hardawar. Xreal's One Pro intrigues me in ways Meta's smart glasses don't. Engadget, September 2025. URL https://www.engadget.com/ar-vr/xreals-one-pro-intrigues-me-in-ways-metas-smart-glasses-dont-120 html. Accessed: 2026-05-12.

[58] Fengming He, Xiyun Hu, Tianyi Wang, Ananya Ipsita, and Karthik Ramani. ScalAR: Authoring semantically adaptive augmented reality experiences in virtual reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22. ACM, 2022. doi: 10.1145/3491102.3517665.

[59] David Heaney. Meta orion ar glasses explained, 2024. URL https://www.uploadvr.com/meta-orion-ar-glasses/. Accessed: 2026-03-10.

[60] David Heaney. Meta ray-ban display glasses hands-on, 2025. URL https://www.uploadvr.com/meta-ray-ban-display-glasses/. Accessed: 2026-03-10.

[61] Philipp Hock, Sebastian Benedikter, Jan Gugenheimer, and Enrico Rukzio. Carvr: Enabling in-car virtual reality entertainment. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, page 4034–4044, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450346559. doi: 10.1145/3025453.3025665. URL https://doi.org/10.1145/3025453.3025665.

[62] Holoride, 2023. URL https://www.holoride.com/.

[63] Kristina Höök and Jonas Löwgren. Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Trans. Comput.-Hum. Interact.*, 19(3), October 2012.

ISSN 1073-0516. doi: 10.1145/2362364.2362371. URL https://doi.org/10.1145/2362364.2362371.

[64] Jinghui Hou, Yujung Nam, Wei Peng, and Kwan Min Lee. Effects of screen size, viewing angle, and players' immersion tendencies on game experience. *Computers in Human Behavior*, 28(2):617–623, 2012. ISSN 0747-5632. doi: https://doi.org/10.1016/j.chb.2011.11.007. URL https://www.sciencedirect.com/science/article/pii/S0747563211002512.

[65] IMAX Corporation. The IMAX experience, n.d. URL https://www.imax.com/en/gb/the-imax-experience. Accessed: 2026-01-19.

[66] International Data Corporation. XR market expands 44.4% in 2025 as smart glasses take center stage, March 2026. URL https://www.idc.com/promo/arvr/. Accessed: 2026-05-12.

[67] Yuta Itoh, Tobias Langlotz, Daisuke Iwai, Kiyoshi Kiyokawa, and Toshiyuki Amano. Light attenuation display: Subtractive see-through near-eye display via spatial color filtering. *IEEE Transactions on Visualization and Computer Graphics*, 25(5):1951–1960, 2019. doi: 10.1109/TVCG.2019.2899229.

[68] Juliet Jain and Glenn Lyons. The gift of travel time. *Journal of Transport Geography*, 16(2):81–89, 2008. ISSN 0966-6923. doi: https://doi.org/10.1016/j.jtrangeo.2007.05.001. URL https://www.sciencedirect.com/science/article/pii/S0966692307000531.

[69] Charlene Jennett, Anna L. Cox, Paul Cairns, Samira Dhoparee, Andrew Epps, Tim Tijs, and Alison Walton. Measuring and defining the experience of immersion in games. *International Journal of Human-Computer Studies*, 66(9):641–661, 2008. ISSN 1071-5819. doi: https://doi.org/10.1016/j.ijhcs.2008.04.004. URL https://www.sciencedirect.com/science/article/pii/S1071581908000499.

[70] Brett Jones, Rajinder Sodhi, Michael Murdock, Ravish Mehra, Hrvoje Benko, Andrew Wilson, Eyal Ofek, Blair MacIntyre, Nikunj Raghuvanshi, and Lior Shapira. Roomalive: Magical experiences enabled by scalable, adaptive projector-camera units. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, UIST '14, page 637–644, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450330695. doi: 10.1145/2642918.2647383. URL https://doi.org/10.1145/2642918.2647383.

[71] Brett R. Jones, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. Illumiroom: peripheral projected illusions for interactive experiences. In *Proceedings of the SIGCHI Conference*

on Human Factors in Computing Systems, CHI '13, page 869–878, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450318990. doi: 10.1145/2470654.2466112. URL https://doi.org/10.1145/2470654.2466112.

[72] Ralph Jones. 'not second screen enough': is netflix deliberately dumbing down tv so people can watch while scrolling? *The Guardian*. URL https://www.theguardian.com/tv-and-radio/2025/jan/17/not-second-screen-enough-is-netflix-deliberately-dumbing-down-tv-

[73] D. Kahneman and A. B. Krueger. Developments in the measurement of subjective well-being. *Journal of Economic Perspectives*, 20 (1):3–24, 2006. doi: DOI:10.1257/089533006776526030.

[74] D. Kahneman, A. B. Krueger, D. A. Schkade, N. Schwarz, and A. A. Stone. A survey method for characterizing daily life experience: The day reconstruction method. *Science (New York, N.Y.)*, 306,5702:1776*1780, 2004. doi: doi:10.1126/science.1103572.

[75] Saleh Kalantari, Bill Tong Xu, Armin Mostafavi, Anne Seoyoung Lee, and Qi Yang. Comparing spatial navigation in a virtual environment vs. an identical real environment across the adult lifespan. *Computers in Human Behavior*, 2024. doi: 10.1016/j.chb.2024.108104.

[76] Konstantina Kilteni, Raphaela Groten, and Mel Slater. The sense of embodiment in virtual reality. *Presence: Teleoper. Virtual Environ.*, 21(4):373–387, December 2012. ISSN 1054-7460. doi: 10.1162/PRES_a_00124. URL https://doi.org/10.1162/PRES_a_00124.

[77] You-Jin Kim, Radha Kumaran, Jingjing Luo, Tom Bullock, Barry Giesbrecht, and Tobias Höllerer. On the go with ar: Attention to virtual and physical targets while varying augmentation density. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400713941. doi: 10.1145/3706598.3714289. URL https://doi.org/10.1145/3706598.3714289.

[78] Marion Koelle. What makes wearable technologies socially acceptable? *XRDS*, 29 (2):30–35, January 2023. ISSN 1528-4972. doi: 10.1145/3571299. URL https://doi.org/10.1145/3571299.

[79] Alexandros Koilias, Michael Nelson, Sahana Gubbi, Christos Mousas, and Christos-Nikolaos Anagnostopoulos. Evaluating human movement coordination during immersive walking in a virtual crowd. *Behavioral Sciences*, 10(9):130, 2020. doi: 10.3390/bs10090130.

[80] Claudia Krogmeier, Christos Mousas, and David Whittinghill. Human–virtual character interaction: Toward understanding the influence of haptic feedback. *Computer Animation and Virtual Worlds*, 30(3-4):e1883, 2019. doi: 10.1002/cav.1883.

[81] Jan-Carlos Kucharek. Look at what you can't have now: The intense reality of the Las Vegas Sphere, February 2024. URL https://www.ribaj.com/buildings/las-vegas-sphere/. RIBA Journal. Accessed: 2026-01-19.

[82] Ouren X. Kuiper, Jelte E. Bos, and Cyriel Diels. Looking forward: In-vehicle auxiliary display positioning affects carsickness. *Applied Ergonomics*, 68:169–175, 2018. ISSN 0003-6870. doi: https://doi.org/10.1016/j.apergo.2017.11.002. URL https://www.sciencedirect.com/science/article/pii/S000368701730251X.

[83] Christos Kyrlitsias and Despina Michael-Grigoriou. Social interaction with agents and avatars in immersive virtual environments: A survey. *Frontiers in Virtual Reality*, 2: 786665, 2022. doi: 10.3389/frvir.2021.786665.

[84] Annemarie Künn-Nelen. Does commuting affect health? *Health Economics*, 25(8): 984–1004, 2016. doi: 10.1002/hec.3199. URL https://doi.org/10.1002/hec.3199.

[85] Jérémy Lacoche, Éric Villain, and Anthony Foulonneau. Evaluating usability and user experience of ar applications in vr simulation. *Frontiers in Virtual Reality*, 3: 881318, 2022. doi: 10.3389/frvir.2022.881318. URL https://doi.org/10.3389/frvir.2022.881318.

[86] Daniël Lakens. Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 4:863, 2013. doi: 10.3389/fpsyg.2013.00863.

[87] Daniël Lakens. Equivalence tests: A practical primer for *t* tests, correlations, and meta-analyses. *Social Psychological and Personality Science*, 8(4):355–362, 2017. doi: 10.1177/1948550617697177.

[88] Land Rover. Land rover debuts invisible car technology. https://media.jlr.com/corporate/news/2014/04/land-rover-debuts-invisible-car-technology, April 2014. Transparent Bonnet concept, Discovery Vision Concept. Accessed: 2026-09-12.

[89] Ed Lantz. A survey of large-scale immersive displays. In *Proceedings of the 2007 Workshop on Emerging Displays Technologies: Images and Beyond: The Future of Displays and Interaction*, EDT '07, New York, NY, USA, 2007. ACM. doi: 10.1145/1278240.1278241.

[90] Bettina Laugwitz, Theo Held, and Martin Schrepp. Construction and evaluation of a user experience questionnaire. In *Proceedings of the 4th Symposium of the Workgroup Human-Computer Interaction and Usability Engineering of the Austrian Computer Society on HCI and Usability for Education and Work*, USAB '08, page 63–76, Berlin, Heidelberg, 2008. Springer-Verlag. ISBN 9783540893493. doi: 10.1007/978-3-540-89350-9_6. URL https://doi.org/10.1007/978-3-540-89350-9_6.

[91] Dong-Wook Lee, Je-Yeon Yun, Nami Lee, and Yun-Chul Hong. Association between commuting time and depressive symptoms in 5th korean working conditions survey. *Journal of Transport & Health*, 34:101731, 2024. ISSN 2214-1405. doi: https://doi.org/10.1016/j.jth.2023.101731. URL https://www.sciencedirect.com/science/article/pii/S2214140523001676.

[92] Hsinju Lee, Fang-Hsin Hsu, Wei-Ko Li, Jie Tsai, Ying-Yu Chen, and Yung-Ju Chang. Get distracted or missed the stop? investigating public transit passengers' travel-based multitasking behaviors, motives, and challenges. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394215. doi: 10.1145/3544548.3581391. URL https://doi.org/10.1145/3544548.3581391.

[93] Changyang Li, Wanwan Li, Haikun Huang, and Lap-Fai Yu. Interactive augmented reality storytelling guided by scene semantics. *ACM Transactions on Graphics*, 41(4), 2022. doi: 10.1145/3528223.3530061.

[94] Jingyi Li, Ceenu George, Andrea Ngao, Kai Hollander, Stefan Mayer, and Andreas Butz. An exploration of users' thoughts on rear-seat productivity in virtual reality. In *12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, AutomotiveUI '20, page 92–95, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450380669. doi: 10.1145/3409251.3411732. URL https://doi.org/10.1145/3409251.3411732.

[95] Jingyi Li, Ceenu George, Andrea Ngao, Kai Holländer, Stefan Mayer, and Andreas Butz. Rear-Seat Productivity in Virtual Reality: Investigating VR Interaction in the Confined Space of a Car. *Multimodal Technologies and Interaction*, 5(4):15, April 2021. ISSN 2414-4088. doi: 10.3390/mti5040015. URL https://www.mdpi.com/2414-4088/5/4/15. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.

[96] Wanwan Li, Changyang Li, Minyoung Kim, Haikun Huang, and Lap-Fai Yu. Location-aware adaptation of augmented reality narratives. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. ACM, 2023. doi: 10.1145/3544548.3580978.

[97] Zhipeng Li, Christoph Gebhardt, Yves Inglin, Nicolas Steck, Paul Streli, and Christian Holz. SituationAdapt: Contextual UI optimization in mixed reality with situation awareness via LLM reasoning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, UIST '24. ACM, 2024. doi: 10.1145/3654777.3676470.

[98] Patrick Lindemann and Gerhard Rigoll. A diminished reality simulation for driver-car interaction with transparent cockpits. In *2017 IEEE Virtual Reality (VR)*, pages 305–306, 2017. doi: 10.1109/VR.2017.7892298.

[99] Patrick Lindemann, Dominik Eisl, and Gerhard Rigoll. Acceptance and user experience of driving with a see-through cockpit in a narrow-space overtaking scenario. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 1040–1041, 2019. doi: 10.1109/VR.2019.8798069.

[100] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. Context-aware online adaptation of mixed reality interfaces. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, UIST '19, page 147–160, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450368162. doi: 10.1145/3332165.3347945. URL https://doi.org/10.1145/3332165.3347945.

[101] Tianyu Liu, Weiping He, and Mark Billinghurst. Using virtual reality as a simulation tool for augmented reality virtual windows: Effects on cognitive workload and task performance. *arXiv preprint*, 2024. URL https://arxiv.org/abs/2409.16037.

[102] Matthew Lombard, Robert D. Reich, Maria Elizabeth Grabe, Cheryl Campanella Bracken, and Theresa Bolmarcich Ditton. Presence and television: The role of screen size. *Human Communication Research*, 26(1):75–98, 2000. doi: 10.1111/j.1468-2958.2000.tb00750.x.

[103] Glenn Lyons, Juliet Jain, and Iain Weir. Changing times – A decade of empirical insight into the experience of rail passengers in Great Britain. *Journal of Transport Geography*, 57:94–104, December 2016. ISSN 09666923. doi: 10.1016/j.jtrangeo.2016.10.003. URL https://linkinghub.elsevier.com/retrieve/pii/S0966692316305798.

[104] Christian Mai, Niklas Thiem, and Heinrich Hussmann. Drawingpresence: A method for assessing temporal fluctuations of presence status in a vr experience. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 3(4), sep 2020. doi: 10.1145/3369827. URL https://doi.org/10.1145/3369827.

[105] Steve Mann. Mediated reality. Technical Report TR 260, MIT Media Lab, Perceptual Computing Section, Cambridge, MA, 1994.

[106] John Mateer. Directing for cinematic virtual reality : how traditional film director's craft applies to immersive environments and notions of presence. *Journal of Media Practice*, 18, 05 2017. doi: 10.1080/14682753.2017.1305838.

[107] Mark McGill, Julie R. Williamson, Angelica Ng, Frank Pollick, and Stephen Brewster. Challenges in passenger use of mixed reality headsets in cars and other transportation. *Virtual Reality*, 24(4):583–603, 2020. doi: 10.1007/s10055-019-00420-x. URL https://doi.org/10.1007/s10055-019-00420-x.

[108] Mark McGill, Graham Wilson, Daniel Medeiros, and Stephen Anthony Brewster. Passengxr: A low cost platform for any-car, multi-user, motion-based passenger xr experiences. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, UIST '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393201. doi: 10.1145/3526113.3545657. URL https://doi.org/10.1145/3526113.3545657.

[109] Anne Collins McLaughlin, Maribeth Gandy Coleman, Vicky Byrne, Rachel Benton, Frank Lodge, and Trevor Patten. Cognitive aid design using diminished reality to support selective attention by reducing distraction. *Human Factors*, 67(3):937–961, 2025. doi: 10.1177/00187208251325169.

[110] Daniel Medeiros, Mark McGill, Alexander Ng, Robert McDermid, Nadia Pantidi, Julie Williamson, and Stephen Brewster. From Shielding to Avoidance: Passenger Augmented Reality and the Layout of Virtual Displays for Productivity in Shared Transit. *IEEE Transactions on Visualization and Computer Graphics*, 28(11):3640–3650, November 2022. ISSN 1941-0506. doi: 10.1109/TVCG.2022.3203002. Conference Name: IEEE Transactions on Visualization and Computer Graphics.

[111] Daniel Medeiros, Graham Wilson, Mauricio Sousa, Nadia Pantidi, Mark McGill, Diego Drago, and Stephen Brewster. Exploring user placement for vr remote collaboration in a constrained passenger space. In *Proceedings of the 30th ACM Symposium on Virtual Reality Software and Technology*, VRST '24, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400705359. doi: 10.1145/3641825.3687722. URL https://doi.org/10.1145/3641825.3687722.

[112] Meta, May 2023. URL https://about.fb.com/news/2023/05/meta-bmw-ar-vr-experiences/.

[113] Meta. Mixed reality utility kit overview. https://developer.oculus.com/documentation/unity/unity-mr-utility-kit-overview/, 2024.

[114] Meta. Developer center – unity depth api. https://developer.oculus.com/documentation/unity/unity-depthapi/, 2024. Accessed: 2026-03-10.

[115] Meta. Ray-ban meta display glasses announcement, 2025. URL https://about.fb.com/news/2025/09/meta-ray-ban-display-ai-glasses/. Accessed: 2026-03-10.

[116] Meta Reality Labs. Introducing orion: Our most advanced augmented reality glasses prototype, 2024. URL https://about.fb.com/news/2024/09/introducing-orion-our-most-advanced-ar-glasses-prototype/. Accessed: 2026-03-10.

[117] Myriam Metzulat, Barbara Metz, Andreas Landau, Alexandra Neukum, and Wilfried Kunde. Does the visual input matter? influence of non-driving related tasks on car sickness in an open road setting. *Transportation Research Part F: Traffic Psychology and Behaviour*, 104:234–248, 2024. ISSN 1369-8478. doi: https://doi.org/10.1016/j.trf.2024.06.002. URL https://www.sciencedirect.com/science/article/pii/S136984782400127X.

[118] Microsoft. HoloLens 2 hardware. https://learn.microsoft.com/en-us/hololens/hololens2-hardware, 2023. Accessed: 2026-03-10.

[119] Microsoft Technology Licensing LLC. Head-mounted display with electrochromic dimming module for augmented and virtual reality perception, 2016. URL https://patents.google.com/patent/WO2016028828A1/en. International publication of patent application PCT/US2015/045779 describing an AR/VR headset with electrochromic ambient light dimming module controlling transmissive optical density.

[120] Paul Milgram and Herman Colquhoun. A taxonomy of real and virtual world display integration. *Mixed Reality*, page 5–30, 1999. doi: 10.1007/978-3-642-87512-0_1.

[121] Paul Milgram and Fumio Kishino. A taxonomy of mixed reality visual displays. *IEICE Transactions on Information and Systems*, E77-D(12):1321–1329, 1994.

[122] Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. Augmented reality: a class of displays on the reality-virtuality continuum. *SPIE Proceedings*, 1995. doi: 10.1117/12.197321.

[123] Tania Modleski. *Regarding Television: Critical Approaches–An Anthology*, volume 2. University Publications of America, 1983. ISBN 9780313270093.

[124] Shohei Mori, Sei Ikeda, and Hideo Saito. A survey of diminished reality: Techniques for visually concealing, eliminating, and seeing through real objects. *IPSJ Transactions on Computer Vision and Applications*, 9:17, 2017. doi: 10.1186/s41074-017-0028-1.

[125] Myura Nagendran, Kurinchi Selvan Gurusamy, Rajesh Aggarwal, Marilena Loizidou, and Brian R. Davidson. Virtual reality training for surgical trainees in laparoscopic surgery. *Cochrane Database of Systematic Reviews*, (8):CD006575, 2013. doi: 10.1002/14651858.CD006575.pub3.

[126] National Research Council. *Virtual Reality: Scientific and Technological Challenges*. National Academies Press, Washington, DC, 1995. doi: 10.17226/4761. Open access via NAP.edu.

[127] Teron Nguyen, Nen NguyenDinh, Bernhard Lechner, and Yiik Diew Wong. Insight into the lateral ride discomfort thresholds of young-adult bus passengers at multiple postures: Case of singapore. *Case Studies on Transport Policy*, 7(3):617–627, 2019. doi: 10.1016/j.cstp.2019.07.002.

[128] Niels Nilsson, Rolf Nordahl, and Stefania Serafin. Immersion revisited: A review of existing definitions of immersion and their relation to different theories of presence. *Human Technology*, 12(2):108–134, Nov. 2016. doi: 10.17011/ht/urn.201611174652. URL `https://ht.csr-pub.eu/index.php/ht/article/view/176`.

[129] Heather L. O'Brien and Elaine G. Toms. The development and evaluation of a survey to measure user engagement. *Journal of the American Society for Information Science and Technology*, 61(1):50–69, 2010. doi: https://doi.org/10.1002/asi.21229. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.21229`.

[130] Heather L. O'Brien, Paul Cairns, and Mark Hall. A practical approach to measuring user engagement with the refined user engagement scale (ues) and new ues short form. *International Journal of Human-Computer Studies*, 112:28–39, 04 2018. doi: 10.1016/j.ijhcs.2018.01.004.

[131] Ofcom. Smartphone ownership penetration in the united kingdom (uk) in 2012-2023, by age, 2023. URL `https://www.statista.com/statistics/271851/smartphone-owners-in-the-united-kingdom-uk-by-age/`.

[132] Joseph O'Hagan, Mohamed Khamis, Mark McGill, and Julie R. Williamson. Exploring attitudes towards increasing user awareness of reality from within virtual reality. In *Proceedings of the 2022 ACM International Conference on Interactive Media Experiences*, IMX '22, page 151–160, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392129. doi: 10.1145/3505284.3529971. URL `https://doi.org/10.1145/3505284.3529971`.

[133] Joseph O'Hagan, Julie R. Williamson, Florian Mathis, Mohamed Khamis, and Mark McGill. Re-evaluating vr user awareness needs during bystander interactions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New

York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394215. doi: 10.1145/3544548.3581018. URL https://doi.org/10.1145/3544548.3581018.

[134] Tom Parkinson and Ian Fisher. Tcrp report 13: Rail transit capacity. Transit Cooperative Research Program (TCRP) Report 13, Transportation Research Board, National Research Council, Washington, DC, 1996. URL https://onlinepubs.trb.org/onlinepubs/tcrp/tcrp_rpt_13-a.pdf.

[135] Sarah Perez. Top 10 meditation apps pulled in 195 million usd in 2019, Feb 2020. URL https://techcrunch.com/2020/01/30/top-10-meditation-apps-pulled-in-195m-in-2019-up-52-from-2018/.

[136] Bastian Pfleging, Maurice Rang, and Nora Broy. Investigating user needs for non-driving-related activities during automated driving. In *Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia*, MUM '16, pages 91–99, New York, NY, USA, 2016. ACM. doi: 10.1145/3012709.3012735.

[137] Pimax. Pimax crystal super — technical specifications. https://eu.pimax.com/pages/pimax-crystal-super, 2025. Accessed: 2026-09-08.

[138] Marta Pizzolante, Sabrina Bartolotta, Eleonora Diletta Sarcinella, Alice Chirico, and Andrea Gaggioli. Virtual vs. real: exploring perceptual, cognitive and affective dimensions in design product experiences. *BMC Psychology*, 12(1):10, 2024. doi: 10.1186/s40359-023-01497-5. URL https://doi.org/10.1186/s40359-023-01497-5.

[139] Katharina Margareta Theresa Pöhlmann, Marc Stephan Kurt Auf Der Heyde, Gang Li, Frans Verstraten, Stephen Anthony Brewster, and Mark McGill. Can visual motion presented in a vr headset reduce motion sickness for vehicle passengers? In *Adjunct Proceedings of the 14th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, AutomotiveUI '22, page 114–118, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450394284. doi: 10.1145/3544999.3552488. URL https://doi.org/10.1145/3544999.3552488.

[140] Irina Popovici and Radu-Daniel Vatavu. Understanding users' preferences for augmented reality television. *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2019. doi: 10.1109/ismar.2019.00024.

[141] Meta Quest. Oculus quest 2: First steps, 2024. URL https://www.youtube.com/watch?v=60yP8f5E-B4.

[142] Bobby Rae, Dec 2023. URL https://doodle.com/en/state-of-meetings-report-2023/.

[143] Eric Ragan, Curtis Wilkes, and Doug A. Bowman. Simulation of augmented reality systems in purely virtual environments. In *Proceedings of the 2009 IEEE Virtual Reality Conference (VR)*, pages 287–288. IEEE, 2009. doi: 10.1109/VR.2009.4811049.

[144] Shwetha Rajaram, Macarena Peralta, Janet G. Johnson, and Michael Nebeling. Exploring the design space of privacy-driven adaptation techniques for future augmented reality interfaces. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, pages 1226:1–1226:19. ACM, 2025. doi: 10.1145/3706598.3713320.

[145] Julian Rasch, Matthias Wilhalm, Florian Müller, and Francesco Chiossi. Ar you on track? investigating effects of augmented reality anchoring on dual-task performance while walking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400713941. doi: 10.1145/3706598.3714258. URL https://doi.org/10.1145/3706598.3714258.

[146] Kishore Rathinavel, Gordon Wetzstein, and Henry Fuchs. Varifocal occlusion-capable optical see-through augmented reality display based on focus-tunable optics. *IEEE Transactions on Visualization and Computer Graphics*, 25(11):3125–3134, 2019. doi: 10.1109/TVCG.2019.2933120.

[147] James T. Reason and J. J. Brand. *Motion Sickness*. Academic Press, 1975. ISBN 9780125840507.

[148] Jacob M. Rigby, Duncan P. Brumby, Anna L. Cox, and Sandy J. J. Gould. Watching movies on Netflix: Investigating the effect of screen size on viewer immersion. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*, MobileHCI '16 Adjunct, pages 714–721. ACM, 2016. doi: 10.1145/2957265.2961843.

[149] Jacob M. Rigby, Duncan P Brumby, Sandy J. J. Gould, and Anna L Cox. Development of a questionnaire to measure immersion in video media: The film ieq. In *Proceedings of the 2019 ACM International Conference on Interactive Experiences for TV and Online Video*, TVX '19, page 35–46, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450360173. doi: 10.1145/3317697.3323361. URL https://doi.org/10.1145/3317697.3323361.

[150] Pejman Saeghe, Sarah Clinch, Bruce Weir, Maxine Glancy, Vinoba Vinayagamoorthy, Ollie Pattinson, Stephen Robert Pettifer, and Robert Stevens. Augmenting television with

augmented reality. *Proceedings of the 2019 ACM International Conference on Interactive Experiences for TV and Online Video*, 2019. doi: 10.1145/3317697.3325129.

[151] Pejman Saeghe, Gavin Abercrombie, Bruce Weir, Sarah Clinch, Stephen Pettifer, and Robert Stevens. Augmented reality and television: Dimensions and themes. In *Proceedings of the 2020 ACM International Conference on Interactive Media Experiences*, IMX '20, pages 13–23, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379762. doi: 10.1145/3391614.3393649. URL https://doi.org/10.1145/3391614.3393649.

[152] Pejman Saeghe, Mark McGill, Bruce Weir, Sarah Clinch, and Robert Stevens. Evaluating and updating a design space for augmented reality television. In *Proceedings of the 2022 ACM International Conference on Interactive Media Experiences*, IMX '22, pages 79–94, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392129. doi: 10.1145/3505284.3529965. URL https://doi.org/10.1145/3505284.3529965.

[153] Pejman Saeghe, Bruce Weir, Mark McGill, Sarah Clinch, and Robert Stevens. Augmenting a nature documentary with a lifelike hologram in virtual reality. *ACM International Conference on Interactive Media Experiences*, 2022. doi: 10.1145/3505284.3532974.

[154] Johnny Saldaña. *The coding manual for qualitative researchers*. Langara College, 2022.

[155] Maria V. Sanchez-Vives and Mel Slater. From presence to consciousness through virtual reality. *Nature Reviews Neuroscience*, 6(4):332–339, 2005. doi: 10.1038/nrn1651. Review article.

[156] Elizabeth B.-N. Sanders. From user-centered to participatory design approaches. In Jorge Frascara, editor, *Design and the Social Sciences*. Taylor & Francis, 2002.

[157] Martin Schrepp, Andreas Hinderks, and Jörg Thomaschewski. Design and evaluation of a short version of the user experience questionnaire (ueq-s). *International Journal of Interactive Multimedia and Artificial Intelligence*, 4(6):103, 01 2017. doi: 10.9781/ijimai.2017.09.001.

[158] Douglas Schuler and Aki Namioka, editors. *Participatory Design: Principles and Practices*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1993. ISBN 9780805809510.

[159] Neal E. Seymour, Anthony G. Gallagher, Sanziana A. Roman, Michael K. O'Brien, Vipin K. Bansal, Dana K. Andersen, and Richard M. Satava. Virtual reality training improves operating room performance: results of a randomized, double-blinded study. *Annals of Surgery*, 236(4):458–463, 2002. doi: 10.1097/00000658-200210000-00008.

[160] Patrick A. Singleton. Discussing the "positive utilities" of autonomous vehicles: will travellers really use their time productively? *Transport Reviews*, 39(1):50–65, 2019. doi: 10.1080/01441647.2018.1470584. URL https://doi.org/10.1080/01441647.2018.1470584.

[161] Patrick A. Singleton. Multimodal travel-based multitasking during the commute: Who does what? *International Journal of Sustainable Transportation*, 14(2):150–162, 2019. doi: 10.1080/15568318.2018.1536237.

[162] Richard Skarbez, Frederick P. Brooks, Jr., and Mary C. Whitton. A survey of presence and related concepts. *ACM Comput. Surv.*, 50(6), November 2017. ISSN 0360-0300. doi: 10.1145/3134301. URL https://doi.org/10.1145/3134301.

[163] Richard Skarbez, Missie Smith, and Mary C. Whitton. Revisiting Milgram and Kishino's Reality-Virtuality Continuum. *Frontiers in Virtual Reality*, 2:647997, March 2021. ISSN 2673-4192. doi: 10.3389/frvir.2021.647997. URL https://www.frontiersin.org/articles/10.3389/frvir.2021.647997/full.

[164] Mel Slater. A note on presence terminology. *Presence Connect*, 3, 01 2003.

[165] Mel Slater. Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364:3549–3557, 12 2009. doi: 10.1098/rstb.2009.0138.

[166] Mel Slater and Sylvia Wilbur. A framework for immersive virtual environments (five): Speculations on the role of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 6(6):603–616, 1997. doi: 10.1162/pres.1997.6.6.603.

[167] Sphere Entertainment Co. The science behind Sphere, n.d. URL https://www.thesphere.com/science. Accessed: 2026-01-19.

[168] Sphere Studios. The wizard of oz at sphere, 2025. URL https://www.thesphere.com/shows/wizard-of-oz-experience.

[169] Sphere Studios and Google Cloud. Ai in the wizard of oz at sphere, 2025. URL https://www.thesphere.com/shows/wizard-of-oz-experience/ai.

[170] Helen Stefanidi, Jan-Hendrik Sünderkamp, Markus Tatzgern, Alina Itzlinger, and Alexander Meschtscherjakov. You're making things ar-kward: Exploring augmented reality in-the-wild. *Proc. ACM Hum.-Comput. Interact.*, 9(5), September 2025. doi: 10.1145/3743740. URL https://doi.org/10.1145/3743740.

[171] Hans Strasburger. Peripheral vision and pattern recognition: A review. *Journal of Vision*, 24(4):15, 2024. doi: 10.1167/jov.24.4.15. Erratum published for this article (see Corrections to: *Peripheral vision and pattern recognition*).

[172] Paul Streli et al. Shadows of reality: Enhancing bystander awareness of mixed reality interfaces. In *Proceedings of the ACM Symposium on Spatial User Interaction*, SUI '25. ACM, 2025. doi: 10.1145/3694907.3765920.

[173] SeedMesh Studio. URL https://assetstore.unity.com/packages/3d/environments/jungle-tropical-vegetation-178966.

[174] Joel Teixeira, Sebastien Miellet, and Stephen Palmisano. Effects of vection type and postural instability on cybersickness. *Virtual Reality*, 28(2):82, 2024. doi: 10.1007/s10055-024-00969-2. URL https://doi.org/10.1007/s10055-024-00969-2.

[175] The Walt Disney Company. Disney+ on Apple Vision Pro ushers in a new era of storytelling innovation and immersive entertainment. Press release, January 2024. URL https://press.disneyplus.com/news/disney-plus-on-apple-vision-pro-ushers-in-a-new-era-of-storytelli Accessed: 2026-09-15.

[176] Transport Scotland. Scotrail franchise – franchise agreement (redacted), 2016. URL https://www.transport.gov.scot/media/2302/abellio-scotrail-redacted-franchise-agreement-web-version-revised pdf. Contains rolling stock requirements including a maximum jerk rate of 0.5 m/s3 for normal traction and braking.

[177] Camille Truong-Allié, Martin Herbeth, and Alexis Paljic. A study of the influence of ar on the perception, comprehension and projection levels of situation awareness. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pages 541–551, 2023. doi: 10.1109/VR55154.2023.00069.

[178] UBI Research. The evolution of optical technologies for ar glasses at ces 2026, January 2026. URL https://en.ubiresearchnet.com/ces-2026-ar-glasses-optics-technology-waveguide-dimming/. Trade press. LUMUS ZOE geometric waveguide, >70° FOV.

[179] UK Department for Transport. Transport statistics great britain 2024: Domestic travel. https://www.gov.uk/government/statistics/transport-statistics-great-britain-2025/transport-statistics-great-britain-2024-domestic-travel, 2025. Accessed: 2026-03-09.

[180] UploadVR. StarVR HMD with 210-degree FOV is $3,200 for developers. https://www.uploadvr.com/starvr-one-3200-developers/, November 2018. Accessed: 2026-09-08.

[181] Stef van Buuren and Karin Groothuis-Oudshoorn. mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software*, 45(3):1–67, 2011. doi: 10.18637/jss.v045.i03.

[182] Radu-Daniel Vatavu. There's a world outside your tv: Exploring interactions beyond the physical tv screen. In *Proceedings of the 11th European Conference on Interactive TV and Video*, EuroITV '13, page 143–152, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450319515. doi: 10.1145/2465958.2465972. URL https://doi.org/10.1145/2465958.2465972.

[183] Radu-Daniel Vatavu, Pejman Saeghe, Teresa Chambel, Vinoba Vinayagamoorthy, and Marian F Ursu. Conceptualizing augmented reality television for the living room. In *Proceedings of the 2020 ACM International Conference on Interactive Media Experiences*, IMX '20, page 1–12, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379762. doi: 10.1145/3391614.3393660. URL https://doi.org/10.1145/3391614.3393660.

[184] L.A Warwick-Evans, N Symons, T Fitch, and L Burrows. Evaluating sensory conflict and postural instability. theories of motion sickness. *Brain Research Bulletin*, 47 (5):465–469, 1998. ISSN 0361-9230. doi: https://doi.org/10.1016/S0361-9230(98)00090-2. URL https://www.sciencedirect.com/science/article/pii/S0361923098000902.

[185] Sunu Wibirama and Hanung A. Nugroho. Towards understanding addiction factors of mobile devices: An eye tracking study on effect of screen size. In *Proceedings of the 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, EMBC '17, pages 2454–2457. IEEE, 2017. doi: 10.1109/EMBC.2017.8037353.

[186] Luke Wilkinson. Cupra exponential experience: Wearing vr goggles at the wheel, Sep 2023. URL https://www.carmagazine.co.uk/car-news/tech/cupra-vr-driving-experience/.

[187] E. J. Williams. Experimental designs balanced for the estimation of residual effects of treatments. *Australian Journal of Scientific Research, Series A*, 2(2):149–168, 1949. doi: 10.1071/CH9490149.

[188] Julie R. Williamson, Mark McGill, and Khari Outram. Planevr: Social acceptability of virtual reality for aeroplane passengers. In *Proceedings of the 2019 CHI Conference on*

Human Factors in Computing Systems, CHI '19, page 1–14, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450359702. doi: 10.1145/3290605.3300310. URL https://doi.org/10.1145/3290605.3300310.

[189] Andrew Wilson, Hrvoje Benko, Shahram Izadi, and Otmar Hilliges. Steerable augmented reality with the beamatron. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*, UIST '12, page 413–422, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450315807. doi: 10.1145/2380116.2380169. URL https://doi.org/10.1145/2380116.2380169.

[190] Christopher Wilson, Diane Gyi, Andrew Morris, Robert Bateman, and Hiroyuki Tanaka. Non-driving related tasks and journey types for future autonomous vehicle owners. *Transportation Research Part F: Traffic Psychology and Behaviour*, 85:150–160, 2022. doi: 10.1016/j.trf.2022.01.004.

[191] Graham Wilson, Mark McGill, Daniel Medeiros, and Stephen A. Brewster. A lack of restraint: Comparing virtual reality interaction techniques for constrained transport seating. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2390–2400, 2023. doi: 10.1109/TVCG.2023.3247084. URL https://doi.org/10.1109/TVCG.2023.3247084.

[192] Bob G. Witmer and Michael J. Singer. Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments*, 7(3):225–240, Jun 1998. doi: 10.1162/105474698565686.

[193] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011. doi: 10.1145/1978942.1978963.

[194] Julia Woodward and Jaime Ruiz. Analytic review of using augmented reality for situational awareness. *IEEE Transactions on Visualization and Computer Graphics*, 29(4): 2166–2183, April 2023. ISSN 1077-2626. doi: 10.1109/TVCG.2022.3141585. URL https://doi.org/10.1109/TVCG.2022.3141585.

[195] Matthew Woodward, Dec 2023. URL https://www.searchlogistics.com/learn/statistics/zoom-user-statistics/.

[196] Xinxing Xia, Frank Yunqing Guan, Yiyu Cai, and Nadia Magnenat Thalmann. Challenges and advancements for ar optical see-through near-eye displays: A review. *Frontiers in Virtual Reality*, Volume 3 - 2022, 2022. ISSN 2673-4192. doi: 10.3389/frvir.2022.838237. URL https://www.frontiersin.org/journals/virtual-reality/articles/10.3389/frvir.2022.838237.

[197] Jianghao Xiong, En-Lin Hsiang, Ziqian He, Tao Zhan, and Shin-Tson Wu. Augmented reality and virtual reality displays: emerging technologies and future perspectives. *Light: Science & Applications*, 10(1):216, Oct 2021. ISSN 2047-7538. doi: 10.1038/s41377-021-00658-8. URL https://doi.org/10.1038/s41377-021-00658-8.

[198] XReal. Xreal air 2. https://www.xreal.com/air2, 2024.

[199] XREAL. XREAL Air 2 Ultra, 2024. URL https://next.xreal.com/us/air2ultra/. Product page. Accessed: 2025-10-04.

[200] XREAL. XREAL One Pro, 2025. URL https://www.xreal.com/one-pro. Product page. Accessed: 2026-05-12.

[201] Qian Yang, Yi Ding, and Shin-Tson Wu. Full-color, wide field-of-view single-layer waveguide for augmented reality displays. *Journal of the Society for Information Display*, 32(5):247–254, 2024. doi: 10.1002/jsid.1288.

[202] Kun Yin, En-Lin Hsiang, Junyu Zou, Yannanqi Li, Zhiyong Yang, Qian Yang, Po-Cheng Lai, Chih-Lung Lin, and Shin-Tson Wu. Advanced liquid crystal devices for augmented reality and virtual reality displays: principles and applications. *Light: Science & Applications*, 11:161, 2022. doi: 10.1038/s41377-022-00851-3.

[203] Yan Zhang, Xiaodan Hu, Kiyoshi Kiyokawa, and Xubo Yang. Add-on occlusion: Turning off-the-shelf optical see-through head-mounted displays occlusion-capable. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2700–2709, May 2023. ISSN 1077-2626. doi: 10.1109/TVCG.2023.3247064. URL https://doi.org/10.1109/TVCG.2023.3247064.

[204] Yan Zhang, Rundong Chu, Qingtai Dong, Xiaodan Hu, Keyao You, Zijie Zhou, Zixuan Guo, Hangyu Zhou, Kiyoshi Kiyokawa, and Xubo Yang. Mask balancing: Perception-driven dynamic visibility enhancement for occlusion-capable optical see-through head-mounted displays. *IEEE Transactions on Visualization and Computer Graphics*, 32(5): 3465–3474, 2026. doi: 10.1109/TVCG.2026.3679903.