



Locher, Maik (2025) *Micro-optics: From wormholes to medical applications*. PhD thesis.

<https://theses.gla.ac.uk/85459/>

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 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

Micro-optics: From wormholes to medical applications

Maik Locher

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

School of Physics and Astronomy
College of Science and Engineering
University of Glasgow



University
of Glasgow

May 2025

Abstract

Lying within the overarching boundaries of micro-optics, this thesis begins by exploring two different “pixelated” approaches to medical spectacles. The first was aimed as a treatment alternative to surgical corrections for torsional diplopia (double vision due to a relative eyeball rotation around the back-to-front axis). To do this, the spectacles were designed to approximately rotate the view seen through them. In the very basic design consisting of simple wedges (each of which we call a pixel), the component does not actually perform imaging leading to a reduced visual acuity of about 6/420. This was worse than desired, which is why the latter half of the chapter was designated to implementing methods that may improve the visual acuity, consisting of a pupil restriction and a new approach we call derivative control. Using this, the visual acuity was improved to a value of about 6/35, but at the cost of added complexity and bulkiness. The second spectacle pair was designed to help with uncorrected refractive errors by permitting patients to adjust the focusing power of the spectacles through simply rotating two cylindrical lens spirals relative to one another. The combination yields an equivalent Fresnel lens, which is the reason we call these adaptive Fresnel lenses. These were simulated in the ray optical limit, yielding an expected feasible variable range of ± 2 diopters. Several configurations were explored where the distances and spiral types were varied, with ray trace simulations confirming the expected view through the adaptive lens.

The other half of the text is dedicated to an extension of the previously published ideal lens cloak. The cloaking principles were explored, yielding conditions for which a light ray within the cloak must remain within the cloak and hence travel in a closed ray trajectory. These were used in the theoretical creation of a novel “extreme omnidirectional ideal lens cloak” which hides an object within from all viewing directions. Furthermore, attractor-like properties were found within the cloak, trapping some light rays entering from the outside. These properties were also found in a more realizable cloak we call the “shifty cloak”, which suggests significant similarities between the two cloaks. The shifty cloak was used to construct a Janus device and an optical wormhole, the latter of which was adversely affected by the attractor properties.

Contents

Abstract	i
Acknowledgements	xxxiv
Declaration	xxxv
Publications	xxxvi
Introduction	xxxvii
1 View rotation spectacles	1
1.1 Introduction	1
1.2 Simple view rotator construction	4
1.3 Simple view rotator results	7
1.3.1 Off-axis view	9
1.3.2 Further variations	11
1.3.3 Apparent magnification and rotation	14
1.4 The problem: view rotation vs image rotation	19
1.5 Improving the view rotator	20
1.5.1 Pupil restriction	21
1.5.2 Derivative control	24
1.5.3 Spaceplates	31
1.6 Discussion and conclusion	34
2 Adaptive spectacles	36
2.1 Introduction	36
2.2 Basic construction	38
2.3 Alvarez-Lohmann winding focusing	48
2.4 More winding focusing	49
2.4.1 Winding focusing by focal length	50
2.4.2 Winding focusing through separation	51
2.4.3 Combined winding focusing	53

2.5	Other spiral types	54
2.5.1	Archimedean spiral	55
2.5.2	Hyperbolic spiral	58
2.6	Physical restrictions and feasibility	61
2.7	Discussion and conclusions	69
3	Closed-loop light-ray trajectories in lens cloaks	73
3.1	Introduction	73
3.2	Review of the ideal-lens cloak	74
3.3	Closed-loop conditions	77
3.3.1	3D generalization	85
3.4	Omnidirectional cloaks	96
3.5	Attractors	100
3.6	Discussion and conclusions	103
4	Shifty cloak	106
4.1	Introduction	106
4.2	Shifty cloak	107
4.3	Optical wormhole	109
4.4	Janus device	116
4.5	Discussion and conclusion	118
	Final conclusions	i
A	Ideal lens wormhole	1
B	Ray tracing simulations and software	4
B.1	Dr Tim	4
B.2	Web browser simulations	5
B.3	Unreal engine 5	5
C	Wave optics of the adaptive Fresnel lens	9

List of Tables

3.1 The shorthand notation used throughout this chapter, specifically, for Fig. 3.10
and the rules involving the closed-loop light-ray trajectories. 84

List of Figures

- 1.1 A simple demonstration of the expected view for someone suffering from torsional diplopia. A scene containing a lattice as a view object and the university quadrangle as a backdrop was chosen for demonstration purposes. Here (a) represents the view seen by the healthy unrotated eye, while (c) is an eye that has been rotated such that the view seen appears to be rotated by 10° in the clockwise direction. To estimate the impact this may have on the binocular vision (vision due to both eyes), the views seen in (a) and (c) were overlayed with equal intensity to form (b). It can be seen that there is severe diplopia. 2
- 1.2 A cross-sectional sketch of the view rotator. The individual pixels (here shown as blue quadrilaterals) and the eye position are not to scale for illustration purposes. Green light rays originating at the eye at a distance of d_{eye} from the ocular (closer to the eye) surface of the view rotator, are shown to pass through the ocular surface without any change in light-ray direction. At the objective surface, this ray splits into a branch ray shown as a green dotted ray, which indicates the unrotated ray direction, explained further in Fig. 1.3. The solid green lines represent rays that have been refracted along a direction of the rotated view by the respective pixels. It can be seen that the central pixel does not in fact redirect the ray through it as expected. The individual pixels are separated according to the black dashed lines in order to minimize non-standard refracted rays. These would be light rays that enter one pixel and pass through a boundary to a neighboring pixel. A pixel thickness is described between the ocular and objective (further from the eye) pixel centers as t . Each pixel also comes with an index described by (n_x, n_y) which may be used to determine the objective surface normals, however here only the central ones are shown ranging from $(0, -1)$ to $(0, 1)$. The span, s , of the pixels is constant across the whole view rotator and dictates the pixel size. 5

- 1.3 The construction used to determine the ray direction change and corresponding objective-surface pixel normals. A light ray, shown as a dotted green line, is traced through a pixel center without any deviations in the same way seen in Fig. 1.2. The ray is then intersected with a plane at a distance, d_{rot} also known as the rotation plane distance, from the objective surface. This plane is where the view should appear to be rotated. The intersection position is then rotated about the center by θ . Therefore the normal at the objective surface should be set to redirect an undeviated ray with direction vector $\hat{\mathbf{d}}_1$ (shown as the dotted green line) onto the desired rotated direction, $\hat{\mathbf{d}}_2$, shown as a solid green line. 6
- 1.4 The view seen through the simple view rotator at a 130° field of view. In (a), the binocular image can be seen with a dotted border representing the edge of the spectacles. This was constructed by equally overlaying the view seen in (b) and (c). Here (b) is the control view expected of a healthy eye without any component or rotations added. On the other hand, (c) represents the view seen with a 10° torsion in the clockwise direction, when looking through the simple view rotator designed to rotate the view by 10° anticlockwise. Another dotted border was added around the spectacle perimeter in (c) to indicate where the correction ends. The slight blurring of the view through the view rotator in (c) is due to the extended aperture size of the eye, which was 4 mm in diameter. As a result, a seemingly upright image is seen when looking through the rotator. The sides remain rotated as the view rotator was chosen to not fill the entire field of view. 8
- 1.5 Off-axis images as seen through the view rotator ((b) and (c)) and with the “healthy” unrotated eye on the bottom of (a), with the corresponding eye (which directly relates to the camera) position above. The camera was placed such that the aperture coincided with the pupil and the back-to-front axis was the camera’s optical axis. For simplicity, we will refer to this camera setup as the eye. The top row indicates the eyeball rotation axis using an arrow. In (a), the eye was rotated about the y direction, which represents the up direction of a healthy (non-cyclodeviated) eye. This can be compared to a cyclodeviated (unhealthy) eye rotated about the same up direction axis in (b). The effective up direction for a cyclodeviated eye will no longer align with the up direction of a healthy eye. Instead, the new direction is tilted by the cyclodeviation angle and can be seen as y' in the top row. Since it is inherently unclear (which may be an interesting future research question) which axis the unhealthy eye will rotate about, both were shown with (c) showing the alternative to (b), where the rotation axis was chosen to be y' . It can be seen that (c) more closely resembles the desired view shown in (a). 10

- 1.6 The binocular view when looking at an angle of 17° off axis, about the y direction in (a) and about the y' direction in (b) and (c). In (a), the binocular image corresponds to the combination of the bottom images in Fig. 1.5(a) and (b), taking half the pixel value from each. Similarly, (b) is the combination of Fig. 1.5(a) and (c), showing a significant improvement to the binocular view. Lastly, a comparison to the expected view seen without the component, also rotated along the y' direction is given in (c). 11
- 1.7 A demonstration of the apparent rotation when the view rotator is seen from a distance further (a) and closer (b) than it was designed for. In (a), the distance to the eye was 2 cm while in (b) it was 1 cm. In both cases, the view rotator was designed for an eye distance (d_{eye} in Fig. 1.2) of 1.5 cm, roughly denoting the normal pupil-to-spectacle distance. The most notable change is an apparent over (a) and under (b) rotation. The corresponding binocular images, produced by overlaying the views of (a) and (b) with an upright view, are shown in (c) and (d) respectively. 12
- 1.8 The view seen (by an eye with a 10° cyclodeviation) of a horizontal line through the view rotator. The view rotator used was designed for an eye distance of 1.5 cm, while the eye was placed at 1 cm and 2.5 cm in (a) and (b) respectively. This resulted in the horizontal black line no longer appearing horizontal but curved instead. This curvature can be categorized into two sections, the inner section represented by the green shaded area spanning a FoV of 47° , and the outer red shaded region containing the FoV in the range of 47° - 130° . Within the green region, it was estimated that the line is roughly straight to determine the apparent over (in (b)) or under (in (a)) rotation. In both of these this is shown as ϕ , where $\phi \approx 3^\circ$ in (a) and $\phi \approx 5^\circ$ in (b). In total, this results in an adaptive range of $\sim \pm 4^\circ$ 13
- 1.9 Raytrace simulations of the view rotator at a 5° (a) and 15° (b) rotation angle. In both cases, the simulations were set up to, first, induce a respective 5° and 15° cyclodeviation and then a correction using the appropriate view rotator. It can be seen that a lower rotation angle and cyclodeviation in (a) results in a sharper view compared to a much more blurred view in (b). 14

- 1.10 In (a), a modification to the setup seen in Fig. 1.3 where a sample plane has been added in grey at a distance of d_s away from the view rotator. Note here that it is shown before the rotation plane but can in practice be placed anywhere in front of the view rotator. The light ray intersections at the sample plane for the rotated (in blue) and unrotated (in green) rays are shown as dots of the respective colors. Within the sample plane in (b) the two intersections can be seen to lie at $\hat{\mathbf{v}}$ and at $\hat{\mathbf{u}}$ for the unrotated and rotated ray respectively. These vectors can then be used to find the apparent rotation, θ_a , by considering the angle between them. Similarly, the ratio of the magnitudes between $\hat{\mathbf{v}}$ and $\hat{\mathbf{u}}$ may be used to find the apparent magnification. 15
- 1.11 Graphs showing the additional rotation and apparent magnification at a probing plane an infinite distance away in (a, b) and (c) respectively, corresponding to the equations found in Eqns (1.5) and (1.6). In order to understand the behavior at this plane, the rotation plane distance, d_{rot} , could not be set at infinity (the perhaps most intuitive distance to use) as this would return the exact desired magnifications and rotations, not representing the issues encountered along the view. For all cases, the eye distance, d_{eye} was set to 1.5 cm while d_{rot} was set to 1 m. In (a), an approximate straight line is shown to follow the curve, showing the additional rotation. The term additional rotation is used as it refers to the apparent rotation minus the expected rotation angle. Therefore, the graphs in (a) and (b) are based on the equation shown in Eqn (1.5) but only after subtracting the desired view rotation angle. The curved solid line represents the exact solution in (a) also for zero thickness while in (b) it was plotted for a thickness of 1 mm. The plot in (b) also shows where the straight line approximation breaks down when plotting it for larger view rotation angles. Lastly, (c) demonstrates this same effect for the apparent magnification, plotting the exact curve described in Eqn (1.6) as a dashed line, it can be seen that the thickness does matter, but for the range over which the view rotator operates it remains a good approximation. For more information and derivations of the plots and equations see the Mathematica document called `ApparentRotationViewRotator` in [1]. 17

- 1.12 A mathematical model for the apparent rotation and magnification of the view rotator in (a) and (b) respectively, for a constant eye-to-spectacle distance (d_{eye}). The “distance from objective surface” represents the distance of the sample plane to the view rotator, specifically the objective surface of the view rotator, d_r . The amount by which the view appears rotated or magnified is shown over a 0.3 m range of d_r . It can be seen that as expected at a distance of 0 m the rotation is 0° and the magnification is 1. Beyond this point, the curves shown agree with the simulation results shown as points along the curve. These were obtained by simulating a resolution test target at the corresponding depth and manually rotating and aligning them with an upright image at the same position. The error bars are an estimation based on the observed range where the two images aligned most closely. Since the curves in both cases are not simply a constant value, the comparison to the original *view* rotator in [2] can be drawn again where this would be the case. 18
- 1.13 (a) demonstrates the central surface profile of the view rotator. It can be seen that it consists of individual refractive wedges, where the objective surface is in view. The view through this is then shown in (b) using a pinhole camera, a reduced FoV(30°), larger pixel size (1 mm), and ignoring any diffraction considerations. By simulating it in this manner, it can be seen that the individual pixels do not rotate the image, but instead shift it to the appropriate rotated position. 19
- 1.14 A Snellen chart when viewed through the view rotator by an eye that has a 10° cyclodeviation simulated in (a). The same Snellen chart when viewed by a healthy eye can be seen in (b). By comparing the two charts, an estimate of the visual acuity can be made that shows the view rotator reduces the visual acuity to around 6/420. 21
- 1.15 A Snellen chart when seen through a view rotator in (a) when using a pupil restriction, which simulates a reduced pupil diameter of 1 mm, the letters ‘L’ and ‘P’ in the line corresponding to visual acuity 6/53 can now (just) be deciphered, which suggests that the view rotator now reduces the visual acuity to 6/53. There are some details visible in the 6/35 line but they are not sufficiently clear to claim a visual acuity of 6/35. In (b) the same Snellen chart is viewed with a healthy eye (without view rotator). 22
- 1.16 (a) shows a binocular image comprising of Fig. 1.15 (a) and (b) in equal parts (where each images pixel values are halved and then added together). It can be seen that the view appears clear, with a slight blur behind each letter. This is due to the corrected eye. In the same way, (b) shows the binocular image but now without any view rotator in place demonstrating the achieved improvement. . . 23

- 1.17 Improvement due to derivative control. The view seen through an idealized view rotator, (a), here shown as a phase hologram where each point redirects a light ray to the appropriate rotated view position, akin to an idealized version of the view rotator design presented in this chapter, when viewed using an extended aperture. This then represents an unphysical view rotator that is planar with infinitely small pixel size that would, in reality be subject to diffractive effects, which are not considered here. The camera focusing distance was set to 10, which coincides with a plane containing the front of the lattice structure. Unlike the simulations shown before this, the view rotator is no longer set up close to the camera (or eye), resulting in a larger camera distance and aperture size. Within this sub-subsection, all simulations are performed with a different simulator, namely the `DerivativeControlSurfaceExplorer`, also available in [3]. In (b) the same scene is viewed again but now with a derivative control surface added to the view rotation surface. The Jacobian of the derivative control surface was chosen to follow the aperture correction, given in Eqn (1.10), with θ_a being set to 20° , which can be shown to correct for an extended aperture by bringing the front of the lattice structure into focus. This was not possible in (a) regardless of the camera focusing distance. The simulation was set up for a 20° rotation, an eye distance of 3, and a rotation plane distance of 10. The lattice was placed such that the front of the lattice coincided with the rotation plane at a distance of 10 from the rotator. 26
- 1.18 The view seen through an idealized pixellated phase hologram of a view rotator is shown in (a) for a 20° rotation. Unlike the simulations shown before this, the view rotator is no longer set up to represent an eye. Here this manifests as a pinhole aperture and an increased camera distance. In (b) a derivative control surface was added to the view rotator with a Jacobian shown in Eqn (1.10). This is the Jacobian for an aperture correction, which as expected does not correct the pixelation perfectly, but is an improvement compared to (a). Lastly, the pixel correction Jacobian Eqn (1.11) is employed in (c) where it can now be seen to match neighboring pixels together creating a seamless view. In (c), the Jacobian rotation angle was set to 25.89° , while the magnification factor was set to 1.0183. We note that ignoring diffractive effects and dealing with pinhole cameras and planar holograms is clearly not physically feasible, but it instead serves as a proof of concept of the required Jacobian. 28

- 1.19 A demonstration of the two correction types for an extended aperture viewing a pixellated view rotator. This combines the scenarios outlined in Figs 1.17 and 1.18, where in (a) a pupil correction is shown. In (b) a pixel correction is used instead. In either case, the view seen is improved compared to Fig. 1.18(a) but still blurred. Ideally, pupil correction and pixel correction would be applied simultaneously, but as they are subtly different this is not possible, resulting in a view that is still slightly blurred. 29
- 1.20 A simple way to induce an image rotation using two cylindrical lens telescopes. Here the end surfaces of each block of refractive material act as telescopic cylindrical-lens pairs, with the cylindrical lens nodal line shown by a red solid line. By configuring them as shown in (a), with the second telescope rotated relative to the first one along the optical axis shown as a red arrow through the centers, the image gets rotated. This is done through a series of image flips, each due to a single cylindrical telescope. Following (b) through to (d) this is shown. First, the image of a beetle is flipped about the vertical axis shown as a solid red line in (b). The flipped image can be seen in (c) along with the next “flipping axis”, rotated by θ with respect to the previous flipping axis, again shown as a red arrow. This axis lies along the cylindrical lens cylinder axis. When the image is then again flipped about this axis, the resulting image can be seen in (d) where it can also be seen that it effectively has been rotated by an angle of 2θ above the horizontal (and initial image) orientation. 30
- 1.21 A sample of the nine central most pixels of a potential derivative controlled view rotator, shown in (a). Each pixel consists of two cylindrical lens telescopes set up to rotate the view seen through them. It should be noted that this was made using `ViewRotatorRefractiveDerivativeControl`, available in [1], with the intention of creating an accurate 3D model of the rotator. In (b) a simulation of the resulting view through such a component can be seen. It is diffuse and blurred, but the top line may barely be interpreted as a “Z” which would give it a visual acuity of roughly 6/89. While the second line is also slightly shown, it is not sufficient to accredit it with the respective visual acuity. 31
- 1.22 The view seen through the derivative controlled view rotator with a restricted pupil size. The derivative control chosen was a pixel correction with a unit magnification. It can be seen that the letters not only appear rotated but also in the appropriate location. This culminates in a visual acuity of $\sim 6/35$, shown in the third row of the Snellen chart. To achieve the pixel correction a Jacobian with a rotation angle of 10.15° was chosen. In practice, this was simulated using two cylindrical lens telescopes rotated by 5.075° relative to one another as described in Fig. 1.20. 32

- 1.23 The same off-axis test as carried out in Fig. 1.5 was done but now with the derivative control. In (a) the desired off-axis view can be seen. In (b) the view seen through the derivative controlled view rotator is shown using a restricted pupil size. This is the same view rotator used to create the view seen in Fig. 1.20. It can be seen that the view is littered with double images and severe blurring. This is mainly due to crosstalk between telescopes. 32
- 1.24 The view seen through the view rotator when a heavily idealized spaceplate is added. At first, a Snellen chart is viewed without any component or torsion in (a). This can be viewed as a control to understand the achieved visual acuities. In (b), the view rotator and spaceplate are added but with the spaceplate adding no additional distance. From (c) through to (e), the spaceplate was used to add an additional 1 cm, 2 cm, and 3 cm respectively. It can be seen that the visual acuity is improved with each bit of space added, but the FoV is also decreased as a larger fraction exits the spaceplate through a side, shown as a black region in the simulations. It should be noted that the spaceplate used in the simulation was heavily idealized and infinitely thin. The spaceplate used throughout, consisted of a single planar surface, which effectively shifted the light rays incident on it to the appropriate position along the same plane, coinciding with the expected transverse position after a ray had traveled the specified extra distance while retaining the initial ray direction. 33
- 2.1 A simple lens is shown in (a) with parallel red light rays incident on the left-hand side coming to a focus on the right-hand side after transmission. This lens was then turned into a Fresnel lens in (b) by removing some of the central thickness of the lens, leading to steps in the surface profile. Each part of the Fresnel lens separated by a step will be referred to as a “Fresnel facet”. The same light rays were incident and can be seen to focus down to a single point once again. A further simplification may be made where the surface profile is approximated not by curved surfaces but by straight lines, turning the Fresnel facets into approximate “Fresnel wedges”. This is shown in (c) with the same light rays incident once again. However, since the Fresnel wedges are no longer adding any focusing power to the system, the light rays are redirected into a “focal region” instead of a focal point. If the wedges are sufficiently small, this region can be made small to approximate a Fresnel lens more accurately. 37

- 2.2 A set of complementary cylindrical lenses with a cross-section, perpendicular to the optical power direction, is shown on the left-hand side. A 3D sketch is provided on the right-hand side showing the equivalent lens by stretching it along the direction of the nodal line, a line on the principal point along the direction of zero optical power. In (a), the complementary parts are shown to be aligned, with one half having a focal length of f while the other has a focal length of $-f$. Therefore, as expected, the system is telescopic, and the light rays (shown as red lines) retain their direction upon transmission. When the lenses are now offset along the optical power direction in (b), the system is still telescopic (indicated by the rays remaining parallel after transmission), but the outgoing direction is changed. The direction change is directly proportional to the phase gradient of the light rays, which in turn relates to the relative length of refractive material a light ray passes through. Adding up the refractive material along the optical axis direction yields a simple refractive wedge shown in (c). The light rays, upon transmission through this wedge, are then redirected nearly identically to the cylindrical lens setup. 40
- 2.3 A complementary set of cylindrical lenses arranged into an array (which is typically called a lenticular array) is shown in (a). As expected, light rays incident on this (shown in red) retain their initial direction upon transmission in all cases. However, in Fig. 2.2, it was demonstrated that a simple set of complementary cylindrical lenses may be turned into a wedge. Therefore, only considering the 2D cross-section of an approximate Fresnel lens shown in Fig. 2.1(c), it may be possible to rearrange the lenticular array such that the addition of an offset complementary array creates an equivalent (in 2D) wedge array as the approximate Fresnel lens. This is shown in (b) with the solid red rays being focused onto a single point. Of course, since this is an approximation, another ray incident on the same telescopic cylindrical lens pair (here shown as a green light ray) will not focus on the same point but into the focal region mentioned in Fig. 2.1(c). Lastly, since this is now a telescopic system, this will be subject to stray light rays, which enter one telescope and leave the system through a neighboring one, shown by a dashed red ray in (b). 41

- 2.4 A 3D generalization to extend the complimentary lenticular arrays from Fig. 2.3 into a full 3D realization. First, a normal cylindrical lens is considered in (a), with the nodal line being shown as a black line along the surface of the cylindrical lens. This nodal line is then wound into the concentric circles shown in (b). When adding the equivalent cylindrical lenses back, a 3D realization of concentric cylindrical lenses can be made. A sketch of this is shown in (c), which can be called a “concentric lenticular array”. When simply adding the complementary concentric lenticular array to this, the same as in Fig. 2.3(a) would be achieved, namely a component that is equivalent to just a flat refractive slab. The more interesting case, however, is shown in (d) where a complementary concentric lenticular array is added but with the concentric nodal lines shifted outwards in proportion to the radial distance from the center. When looking at the cross-section in (d), it can be seen that the same pattern as in the lenticular arrays in which one array is stretched (Fig. 2.3(b)) appears, and therefore, it can be expected to act as an approximate Fresnel lens, but now this holds for all cross-sections through the center. In the cross-section, the local misalignment, ΔR , was shown in two places. The misalignment, ΔR , will be positive if the cylindrical lens with the negative focal length (top in (d)), will be shifted towards a greater radius relative to the positive focal length cylindrical lens (bottom in (d)). This will then corresponds to an overall positive focal length Fresnel lens. 43
- 2.5 A similar consideration as in Fig. 2.4 was designed, where the same cylindrical lens and nodal line are shown in (a). Unlike the aforementioned figure, the nodal line is no longer wound into concentric circles but instead into a logarithmic spiral, shown in (b). When adding the cylindrical lens back to the nodal line in (c), a cylindrical lens wound into a spiral was made. Now, unlike the concentric circle, the exact complementary component was added back. When nothing further is done, it will act like a solid slab of flat refractive material; however, by simply rotating the two components relative to one another, the same radial offset as shown previously may be achieved in (d). The offset is shown as ΔR in the cross-section, and is proportional to the radial distance and the relative rotation angle. Therefore, ΔR may be adjusted simply by rotating the components. 45

.....

- 2.6 A plot of two logarithmic spirals rotated relative to one another about an axis passing through the center into the page. The relative rotation is angle is $\Delta\phi$, whereby the second, dashed, spiral is rotated in the clockwise direction. As a result of this, the radial distance of the two spirals along any direction, ϕ , will no longer be the same. This is shown by R_1 , and R_2 for a particular ϕ no longer being of the same length. Therefore, the radial offset between the two spirals at this point may then also be given by $R_2 - R_1$ 46
- 2.7 A ray-traced simulation of an adaptive Fresnel lens using the open-source ray-tracer called Dr. Tim [4]. Here, (a) shows the scene (a “Tim head” on a checker-board floor) being viewed without any component. When adding the complementary set of cylindrical lenses wound into a spiral shown in Fig. 2.5(d), but not rotated relative to one another, (b) was made. As expected, no significant changes (other than a slight darkening of the view due through the components due to the lens components being made slightly absorptive in order to be visible) are noticeable. However, upon a relative rotation, $\Delta\phi$, of 30° in (c), the image seen appears significantly larger. In fact, this should be the same image seen as when looking through an ideal lens of focal length 12.7 (calculated from Eqn (2.12). To check this, exactly such an ideal lens was added in (d), confirming the result. The parameters chosen for this simulation were $b = 0.015$ and $f = 0.1$. Some imperfections are clearly visible in (c), where the images of neighboring windings do not perfectly match up. The spiral boundaries are also very clearly visible, as this is where the stray light rays shown in Fig. 2.3(b) will occur. We note that the focal lengths do not have units, as the raytraces operate in dimensionless units, where a floor tile is a standard unit of length 1. Throughout this chapter, this will be the standard unless stated otherwise. All the ray trace simulations in this chapter were performed using a runnable Java archive file called `SpiralLensVisualiser`, available at [3], unless otherwise stated. 47
- 2.8 The exact same setup as shown in Fig. 2.7(C), but now when using Alvarez-Lohmann winding focusing. It can be seen that the stray light rays (or ghost image) at the spiral boundaries remain, but now the images between neighboring windings match up almost perfectly. 50

- 2.9 In a similar way to Fig. 2.8, here another winding focusing method is demonstrated. This was done by varying the cylindrical focal length, f , of the second negative focal length cylindrical lens spiral slightly. It can be seen that the image produced is equivalent to that shown in the aforementioned figure. However, this method requires a change in f , the focal length of the cylindrical lens, depending on the relative rotation angle, θ . Since the focal length of the cylindrical lenses is supposed to be locally static within the spiral, specifically, it may vary with ϕ , the unbound azimuthal angle, but not with θ , this type of winding focusing is not feasible. 52
- 2.10 The last winding focusing method, which adds a small amount of space between the two complementary cylindrical lens spirals, such that they are no longer perfectly telescopic. As a result, some focusing power is added, leading to the separation winding focusing method shown here. It can be seen that the image produced matches that of the other winding focusing methods, with the same shortcomings due to a ghost image. 53
- 2.11 A demonstration of a combination of separation winding focusing, and a change in focal length of the second cylindrical lens spiral. In (a), a simple complementary cylindrical lens spiral of focal lengths $f_1 = 0.1$ and $f_2 = -0.1$ was taken and separated by a distance 0.0048, leading to a significantly deteriorated image as it was no longer telescopic. When adjusting the focal length of f_1 to $f_1 = 0.104$, and separating the two components by 0.004, the system becomes telescopic once more, leading to an improved image shown in (b), however, throughout the neighboring winding images do not align perfectly as previously seen in the case where no winding focusing was used. To account for this, separation winding focusing was added back, adjusting the relative separation to 0.0048 for the same focal lengths. As a result, (c) was observed. It can be seen that this is still not perfect, especially further from the center due to the additional magnification added by the telescopic pair, and the assumption made when calculating the required additional distance in Eqn (2.23), but it is significantly better than both (a) and (b). 55
- 2.12 A plot of two Archimedean spirals rotated relative to one another by $\Delta\phi$. It can be seen that this specific spiral type has a constant winding width, with the line spacing being equidistant. Furthermore, when rotating the spirals relative to one another, the radial offset, ΔR , can be seen to be the same everywhere. This means it should be independent of radial distance from the center, unlike what was observed in Fig. 2.6. 56

- 2.13 A ray-traced simulation of an adaptive Fresnel lens using cylindrical lenses wound into an Archimedean spiral. It can be seen that the overall image does show the expected Tim head, but there are significant issues with the image quality as the elemental images do not align well to form an integral image. . . . 57
- 2.14 The exact same setup as in Fig. 2.13, but now when adding the Alvarez-Lohmann winding focusing to the cylindrical lenses. As desired, this leads to a significantly improved integral image where the elemental images of each winding now align properly. As was the case previously, this should mimic a Fresnel lens of focal length 12.7, which is shown in Fig. 2.7(d), which it does nearly perfectly except for a small, unrecognizable, spiral shaped slither of a ghost image and a very distinct point in the very center, where the spiral ends in a 90° pitch. 59
- 2.15 A plot of a pair of hyperbolic spirals rotated relative to one another (shown as a solid and as a dotted line respectively). The hyperbolic spiral is wound the opposite way to the previous spirals, where it approaches the center from infinity when increasing the winding number. This is an inverse relationship with ϕ which leads to the windings becoming (infinitely) small towards the center and their separation shrinking accordingly. This also occurs in the logarithmic spiral but is shown at a decreased rate for spiral parameters that may be attractive for spectacles. Due to this, the center of the spiral appears to be black, representative of a very tightly wound spiral. When looking at the separation between the two spirals given by ΔR , it can be seen that this varies greatly with the distance from the center and in conjunction with ϕ 60
- 2.16 A ray-traced look through an adaptive Fresnel lens where the cylindrical lenses are wound into a hyperbolic spiral. In this case, the same equivalent Fresnel focal length lens as in the previous ray trace figures is desired. It can be seen to not only reproduce the desired focal length as well as the other focused spirals, but it does so without the use of any additional winding focusing. Therefore, this specific spiral type can be seen to be entirely self-focusing. 62

- 2.17 A brief investigation into the ghost image found in all adaptive Fresnel lenses, irrespective of spiral type. In (a), an adaptive Fresnel lens using a logarithmic spiral winding was viewed with the two components rotated by 180° with respect to one another. This then led to the case where the ghost image and the desired image take up an equal fraction of the overall view. It can be seen that two Tim heads become visible, one being larger and inverted, while the other is small and upright. These should correspond to the images of a lens of a Fresnel focal length found by Eqn (2.12) and Eqn (2.44) respectively. To check this, two ideal lenses were set up with their focal lengths set according to this. In (b), it can be seen that the image seen matches that of the inverted Tim head, while (c) shows that of the smaller upright Tim head. 63
- 2.18 An extension to allow the adaptive Fresnel focal length to work best around non-zero focusing power. In (a), a cross-section of two complementary lenticular arrays is shown, which has a focusing power of zero, demonstrated by red light rays passing through it undeviated. By adding two lens surfaces on the flat regions of the lens, the same complementary setup now has a non-zero focusing power indicated by the red rays meeting at a focal point after transmission. . . 64
- 2.19 A demonstration of the added focusing power. In (a), an adaptive Fresnel lens is shown without any relative rotation angle, where the component was previously doing nothing, but now with the addition of two lens surfaces as shown in Fig. 2.18(b). Therefore, the view seen is significantly changed relative to the setup before the additional focusing power shown in Fig. 2.7(b). Furthermore, by now rotating the components relative to one another, the focusing power can be adaptively increased further, demonstrated in (b), where the Tim head appears closer (and hence larger) than in (a). 65

- 2.20 A phase profile plot of one part of a spiral lens. In all cases, the phase profile divided by the wave number, k , is plotted as a function of the spiral radius, R . In (a), a logarithmic spiral is used with a cylindrical lens focal length and a winding parameter of 0.1. The same cylindrical lens focal length and winding parameter are used in (b), but now for a hyperbolic spiral. Lastly, in (c), the same cylindrical lens focal length is used, but with a winding parameter of 0.05 and an Archimedean spiral. In (a), the individual windings, shown as blue bumps, meet exactly halfway with the same phase, indicated by two neighboring cylindrical lens centers being separated by D , while they meet at a distance of $D/2$ from each other. The same halfway meeting point is used in the other spiral cases, (b) and (c), but in these cases the phases of neighboring windings no longer have the same value at this cutoff point. As a result, both (b) and (c) contain numerous discontinuities (one per winding cutoff), where one such discontinuity is indicated by a red arrow. These may be resolved by constructing the spirals with a different winding cutoff point. 66
- 2.21 A ray-traced simulation indicating how an extended aperture size may change the result. To do this, a different ray tracer was used, available at [5]. This ray tracer makes use of the GPU, allowing real-time interactive simulations. A picture of a Tim head was viewed in (a). For a pinhole aperture in (b), after adding the adaptive Fresnel lens, the Tim head appears larger, but the ghost image (mostly showing up as a brown spiral, due to the color of the table and the background in all directions) can also be seen. When now extending the aperture size and adjusting the focusing distance (which is typically done automatically by the eye) to that of the Tim head image location, the ghost image becomes significantly blurred out to the point where it is barely noticeable, shown in (c). 68
- 2.22 A demonstration as to why a certain bit of separation is not desired but necessary for a refractive realization of an adaptive Fresnel lens. In (a), two complementary cylindrical lens spirals are shown with a separation of ~ 0 . It can be seen that they do not overlap in any part indicated by the color mixing (or lack thereof). When rotating the same components relative to one another in (b). It can be seen that significant regions of the cylindrical lenses overlap, shown as the darker regions in the 3D schematic and highlighted by a yellow area in the cross-sectional cut. Therefore, the only way in which this can be refractively feasible is for the two components to be separated by some larger distance. . . . 70

- 2.23 The relationship between the integral image quality of a complementary setup when adding some separation. In all cases, the focal length was set to $f = 0.1$ with the complementary component having a focal length of $f = -0.1$. In (a), a separation of 1 % of the cylindrical lens focal length was used. It can be seen that this actually looks as good, if not better, than no separation. This is because it approximates the required distance for the separation winding focusing. However, in (b) the separation is already significantly larger than the winding focusing separation at 2 % of the cylindrical lens focal length, and the image quality is significantly worse with the elemental images, especially further from the center, no longer aligning at all. When increasing the separation to just 3 % of the cylindrical lens focal length in (c), the image becomes even worse, as expected, but not desired. This shows that either the parameters should be chosen to allow for a small separation, or the combined winding focusing method outlined in section 2.4.3 has to be used. 71
- 3.1 A schematic showing a 2D construction of an ideal-lens cloak in (a). Each straight blue line represents an ideal lens. The heights of the inner vertices, shown as h_1 , h_2 , and h , are the vertical distances from the principal point on the bottom lens, p_1 , and the principal points of the inner lenses, p_2 , p_3 , and p_4 . A schematic of a sample light ray trajectory is also shown to pass through the cloak as a red line. As expected of a TO cloak, the ray exits the cloak at the intersection between the outermost lens and the incident light ray, while maintaining the incident light ray direction. This is highlighted with a red dotted line. Within the cloak, such a light ray can travel along a different direction, explored further at a later point. A basic 3D construction of the ideal-lens cloak is seen in (b). Here the red rods represent the net of the tetrahedron. Each polygon bound by 3 of these rods is filled with a lens (although omitted here for illustration purposes). The same principal points as in the 2D case are shown, but now the base vertices of the bottom lens are also shown as v_{1-3} 76
- 3.2 The same 2D ideal-lens cloak as shown in Fig. 3.1(a), but now each lens has been named L_{nm} where n and m are the cells bound by the lenses. These cells are labelled with a cell number, with cell 0 being the outside space, and cells 1-3 as shown in the figure. Within each cell type, there may be multiple cells, which are differentiated through the use of a letter, shown for example as 2a and 2b. A special name is given to the bottom lens, namely the base lens. 77

- 3.3 An ideal-lens cloak with an additional line drawn along the base lens. This black line extends to infinity, indicated by I. The real image of I is shown within cells 2 of the cloak as I'_{2a} and I'_{2b} . The exact same images would form for a point at negative infinity. To find the image positions, two rays were taken from I, and passed through lens L_{03b} , shown as a solid red line, to find the first image, which forms specially outside of, but still in relation to, cell 3b and is labeled as I_{3b} . Here the red solid line turns into a red dashed line indicating the ray will no longer intersect another lens as it is already within the correct cell. To now obtain the image within cell 2b, two rays are taken from I_{3b} and passed through to lens L_{2b3b} , which connects cells 3b and 2b. This is again shown by a dashed red line turning in to a solid red line upon the intersection. Finally, these rays now meet at the labeled image position of I'_{2b} , which when passing through the cloak, is also shown to be imaged onto I'_{2a} 78
- 3.4 In (a), the image of the base lens within cells 3 is shown(dashed line). The whole base lens is imaged into this cell, with half of it lying in cell 3a and the other half in 3b. The area bound by the base lens and its images in cells 3 is shaded in, shown in (b). 79
- 3.5 A ray-traced diagram of an ideal-lens cloak. The dark red line within cells 1 and 2 represents a ray traveling in a closed-loop trajectory. The image of this ray in cells 0 and 3 is shown as a dashed red line. Since the image is parallel to the base lens, the intersection with a line along the base lens occurs at infinity. Therefore it can be seen that the ray within cells 2 must pass through I'_{2a} and I'_{2b} , the images of I. The raytracing was done using an open-source software called Dr. Tim [4], with a runnable Java archive file called RayPlay2D available in [3]. Within this chapter, all 2D ray traced scenes were also performed using this software unless otherwise stated. 79
- 3.6 A modification to Fig. 3.3, with the addition of R, a line that runs along the base lens to and from infinity but excludes the base lens. The images of R are shown within cells 2 as a black dashed line called R'. The images of the base lens are now also shown in both cells 2 and 3. It can be seen that the combination of R and the base lens, imaged into cells 2, split the type 2 cells. Similarly, the image of the base lens in cells 3 splits the type 3 cells. 80
- 3.7 A ray-traced simulation of an ideal-lens cloak, showing multiple rays in both red and green. The red lines represent rays that follow a closed-loop trajectory while the green rays do not. The rays were started in the center of cell 1. Upon closer inspection, it can be seen that all the red rays pass through R', shown as a black dashed line, while the rays that escape in green pass through the base lens, and its images which are shown as blue dashed lines. 81

- 3.8 An ideal-lens cloak with a ray traced through it in (a). The cells that a ray passes may only be used once due to the loop imaging condition unless the ray travels along the same trajectory again. Any cell which the ray passes through can be thought of as used up, these are highlighted in grey. In this drawing, the ray intersects the fewest number of cells possible for any ray entering the base lens for this particular ideal-lens cloak. This stems from the base lens image positions, with an artistic representation shown in (b) as green and blue dashed lines. It can be seen that the base lens is split up into two halves, each being imaged into cells 2 and cells 3 once in its entirety. A ray passing the base lens at P must then pass all the images of P namely P'_1 , P'_{2b} , and P'_{3a} along its trajectory through the cloak, or have a real image that can pass these. The latter becomes impossible for a ray passing the base lens as it would then inhibit the real ray from passing through the cloak, as it would be required to cross the image of the base lens at a different point which in turn would correspond to a different point on the base lens. Traveling around the base lens is also not permitted due to the R' line, shown in Figs 3.6 and 3.7, causing it to travel in a closed-loop trajectory. 82
- 3.9 The same ideal-lens cloak is shown in both (a) and (b). In (a), a ray is traced through lens L_{3a} , and leaves the cloak through L_{3b} . It can be seen that it never crosses into the area bound by the base lens and its images in cells 3, here shown as the same bounded area previously described in Fig. 3.4. Instead, a real image of this ray forms within the cloak. The image can be seen to travel in a closed-loop trajectory presented as a red dashed line. The same holds in the inverse case, with a real ray starting within the cloak in (b), and traveling along the same closed-loop trajectory as shown in (a). The image of this ray is then shown to be the exact same as the real ray in (a), demonstrating that we can interchange the real ray with the image freely. 83
- 3.10 A schematic of all the most important lines, lenses, and cells in an ideal-lens cloak, used to derive a set of rules. The images of the base lens, now called D, and the line along the base lens but excluding it called R are shown as the primed version within the cloak as blue and black dashed lines respectively. Combining D and R yields B, a line spanning the entirety of the space along the base lens. The cells are labeled as previously described in Fig. 3.2. This cloak schematic can be referred to in conjunction with the definition of each used letter in Table 3.1 and the rules described in the main text. 84

- 3.11 A drawing of an ideal-lens cloak containing the image of every position infinitely far away from the cloak in all directions as a green dashed line. In cell 1 this line lies on the focal plane (shown as a black dotted line) as the base lens will image anything within the focal plane to infinity. In cells 2a and 2b, the line crosses the I'_{2a} and I'_{2b} points respectively which represent the points at infinity along the base lens as previously described in Fig. 3.6. 86
- 3.12 The image of an R plane section within one type 2 cell of the ideal-lens cloak. A wireframe model of the cloak is shown with the corresponding R plane parts imaged inside of it. Note that we only shown a part of it for clarity, but the other unshaded regions are imaged into the cloak as well. The light red part of the R plane is imaged onto the light red part within the cloak while the dark red part is imaged onto the dark red part inside. In (a) the cloak is viewed from an angle where the location of the images can be seen to lie within a cell of type 2, while (b) is a top-down version of the same scene. The top-down version may be used to demonstrate that the section shown is exactly $2/3$ of the whole R plane being imaged into just one type 2 cell. 86
- 3.13 In (a), the base lens is split up into 3 sections, each imaged into the corresponding cells 2, with the colors indicating the respective positions. The side view shows that the base lens is indeed imaged into the type 2 cells while the top-down view in (b) may be used to show the symmetry of the images within the cloak. Each image appears across the central symmetry axis from the corresponding position on the base lens. Note that in (b), the colors in cells 2 are representative of the base lens colors, and due to color mixing and scene illumination, the corresponding red part appears dimmer. 87
- 3.14 (a) top-down version of the colored-in base lens. Unlike the segmentation shown in Fig. 3.13, this is now split up to be fully in a single-letter cell when viewed from the top-down i.e. the top-down view of the A cells intersects only the red portion of the base lens. The image of just one of these sections is shown in (b), where it can be seen to be fully imaged into a single type 3 cell. The parts not shown were omitted to permit a view into the inside of the wireframe cloak model. 87

- 3.15 A visual representation of the special significance associated with the base lens. In (a), a setup of the cloak containing a ball is shown. The ball is placed within cell 1 of the cloak. A 2D schematic is shown in (b), which also contains an estimate to the expected ray directions, with only the rays through the base lens, in red, intersecting the sphere. An eyeball was placed at the specific viewing position for which the rays, shown as the red area, were drawn. Moving the camera to this position and viewing the same setup shown in (a) yields (c). It can be seen that part of the sphere is visible above the base lens and a part below when using the wireframe model of the cloak. Replacing the wireframe model with a cloak made of lenses then results in (d). Here, Part of the sphere can be seen, corresponding to the part that would have been visible through the base lens only. The distortion occurs due to the imaging properties of the cloak. . . . 88
- 3.16 A topological reduction of a 3D, $n = 8$ base vertex, cloak into a 2D schematic. Each blue straight line represents a lens labeled by the corresponding lens name. The central blue circle also represents a lens, but this is the special case of the base lens, D. Within the center and the outside of the cloak there are two cells labeled as 0. These represent the space outside of the cloak, with the central enclosed region representing the space below the base lens of the cloak, specifically the space in line with the base lens. Therefore any ray entering through the base lens must first travel in the enclosed 0 cell space. All the other cells are labeled in the same manner as for the simple 2D cloak, with letters differentiating between the same type of cell. This breakdown will be used to investigate the allowed ray trajectories and resulting rules in a simpler way. 90
- 3.17 An example of the simplest ray-trajectory when drawn in the topological reduction of an $n = 8$ base vertex cloak. This is the simplest as it crosses the base lens and intersects the least number of cells, where a solid red line denotes the ray trajectory. The longest green-dotted line represents the image of the base lens that it crossed within cell 3, the same as D'_3 shown previously. As the ray crosses this, it must then also cross the corresponding image in cell 2, shown as a shorter green dotted line. To reach this, the ray may travel only within either the type 2 cells or, as shown here, within the type 3 cells. The final image it will then have to cross is the corresponding region on the base lens, given by a curved green dotted line before leaving the cloak into cell 0 as shown. To permit the light ray to exit the corresponding region of the base lens, we permit it to curve within cell 1. 91

- 3.18 The same reduced cloak as in Figs. 3.16 and 3.17 but now showing the images of all of the base lens, D, and the R plane. The images are given by a black dashed line and a green-dotted line for the R plane and base lens respectively. It can be seen that the combination of the two images completely split the cloak into 3 regions, one below the combination of the R'_2 and D'_2 lines, a region bound between this and the D'_3 lines, and lastly the region above the D'_3 lines. These images dictate how a ray may travel through the cloak. 92
- 3.19 The possible closed-loops within an $n = 8$ ideal-lens cloak, shown in the topological reduced version. In (a), the case using the least number of cells using cell 1 as part of it. If cell 1 cannot be used, the closed-loop must form entirely within cells 2 and 3. This is shown in (b) where it can be seen to use all of the greyed-out cells in the process. In both cases, the images of R'_2 in the relevant cells are shown as black dashed lines, indicating what cell the ray must travel to. 93
- 3.20 An example of an attempted closed-loop trajectory which is hindered by the used cell location rather than the available cell number. In both cases, there is a light ray passing through the cloak shown as a light red line, the same ray which has been shown in Fig. 3.17. Additionally, a dark red ray can be seen indicating a possible image of the light red ray. In (a), one of the special cases is shown whereby R'_2 is crossed in 4 points indicated by a dark red ray crossing the black and gray dotted line. These dark red lines would then be a connected pair of closed loops which are imaged into one another. One of the loops could successfully form in the half of the cloak containing the free cells. The second loop cannot occur as there are no neighboring free cells indicated by the red cross where the dark red ray has to meet the used cells. A similar scene can be seen in (b), but instead, an attempt was made to produce a closed-loop as shown in Fig. 3.19(b). To complete the loop, the ray would once again be required to travel into a used cell, indicated by a red cross. These regions are unavoidable regardless of trajectory choice simply through the constraint that the pair of R'_2 images must be crossed. The corresponding pair of R'_2 lines is indicated by the shade of grey used in the diagram. 94
- 3.21 A schematic of a possible ray trace trajectory of a lens star within cells 2 of the ideal-lens cloak. In this case the trajectory can be seen to cross the R'_2 lines in two places across the centre, upholding the imaging conditions outlined. The cells the ray travelled through are greyed out to indicate that these have been used. It can be seen that a lens star trajectory will split the cloak fully, meaning any ray entering through the base lens into cell 1 cannot travel to the outside again. 96

- 3.22 A ray-traced lens star trajectory within an $n = 3$ base lens ideal-lens cloak. In (a), the lens star is identified with the grey triangle forming a possible plane in which the lens-star trajectory can occur. It can be seen that the lenses, shaded in cyan, intersect this plane as required for the formation of a lens star. To confirm the existence of closed-loop trajectories following lens stars, a light ray was placed onto the lens star plane within the cloak in (b). This was then ray traced through the scene, shown as a red line. The ray travels along a closed-loop intersecting each lens once before returning to the starting point as expected. Subsequent round trips follow the same trajectory. In (b), the lenses were not shown for clarity but all lenses were included in the ray trace of the scene. 97
- 3.23 A construction of the original omnidirectional lens cloak proposed in [6]. In (a), the basic construction and concept of an omnidirectional ideal-lens cloak is shown. An inner cloak is placed within an outer cloak such that it is imaged outside of the outer cloak. The image can be seen as a small dashed cyan triangle below the outer cloak. Within the inner cloak, a point, P , can be seen. As expected, the image of this point due to the outer cloak lies within the image of the inner cloak and is labeled as P'_{oi} . The combination of the cloaks image this point further onto the point P'_{oi} . A red line connecting these two points can be seen to miss the outer cloak. A viewing position and direction along this line has a line of sight through both images of P , but never through the physical cloak. Therefore, it is fully cloaked from these directions. Since any other view direction does not contain a line of sight where the images of P align, P must be hidden for *all* view directions. This is confirmed using a ray trace simulation in (b). It can also be seen that this coincides with the condition of all rays having to cross the R'_2 lines, here given as black dashed lines. 98
- 3.24 The simplest version of an omnidirectional cloak found. The basic ideal-lens cloak, shown as a construction of cyan lenses, was found to act as an omnidirectional cloak for a specific line (or surface in 3D). This line is the R'_2 line shown here as a black dashed line in one of the type 2 cells. When a light source is placed onto this line and traced, all rays cross the R'_2 line and as a result follow a closed-loop trajectory. Therefore anything along this line will automatically be cloaked for all directions. 99

- 3.25 The setup for an *extreme* ideal-lens cloak. In (a), an inner cloak can be seen to be placed such that the base lens of the inner cloak overlaps the R'_2 line of the outer cloak here shown as a black dashed line. The resulting structure should therefore, according to the rules found, act as an omnidirectional cloak for any object placed within the cloaking area of the inner cloak. To confirm this, a light source was placed within the inner cloak in (b). The rays were traced and it was found that none of them escaped to the outside as hoped. 99
- 3.26 The same extreme ideal-lens cloak shown in Fig. 3.25, but now generalized to 3D. In (a), a wireframe model of the setup can be seen where the base lens of the inner red cloak was aligned to lie on top of the R'_2 surface of the outer blue cloak. Within the inner cloak, a sphere is shown. The goal is to hide the sphere from a camera position where it should be visible when using just one cloak. One such position is shown in (b) where the same scene is viewed. When the wireframe model is now replaced with ideal lenses, and the scene is ray traced again, (c) is formed. In (c), it can be seen that the sphere is completely invisible and the overall cloak acts as a TO device with the space behind the cloak appearing undistorted as required. To prove that it was indeed the additional cloak that allowed for the omnidirectional invisibility, a ray-traced simulation of the same scene but removing the inner cloak is shown in (d). Here the sphere is once again visible as expected when only using the outer cloak. 100
- 3.27 A graphical representation of the attractor-like behaviour shown through Eqn (3.3) in (a). It can be seen that with each round trip, shown as a step in the graph, the angle of refraction moves closer towards a final value of 90° which is also the angle obtained when taking the limit of $n \rightarrow \infty$. At 90° , the ray moves parallel to the lens or perpendicular to the optical axis. This can be seen in (b) where a red ray is incident on a trapping lens in cell 1. Upon each round trip, the ray direction has a larger component parallel to the lens, agreeing with the graphical results and the equation found. 102
- 3.28 An ideal-lens cloak with multiple ideal lenses placed within it. In (a), a simple schematic can be seen with each cyan line representing a lens. A single ray is then passed into the cloak in (b) and made to intersect all the lenses at least once. After subsequent round trips, it can be seen that the ray is fully attracted to just one lens (the left one) and no longer intersects the other lenses. 102

- 3.29 A setup of an ideal-lens cloak containing two additional lenses within, shown in (a). These extra lenses were now arranged such that they form a telescope, with the black dotted line representing the focal plane. The telescope was constructed to magnify the beam when passing left to right, making any incident ray parallel to the optical axis exit the second lens closer to the optical axis. When passing in a light ray at a specific direction, it can be made to run parallel to the telescope's optical axis upon reaching the telescope, as shown in (b). Due to telescopic properties, this means it must also leave the second lens parallel to the optical axis, and the process will begin anew upon the second round trip, with each iteration moving the ray closer to the optical axis. As can be seen in (b), this occurs until the ray appears to be traveling perfectly on the optical axis, always intersecting both lenses. 103
- 4.1 A schematic of the shifty cloak. The 4 grey-shaded regions represent homogeneous material slabs, each used to redirect light rays in such a manner that the interior region (white square) is imaged to the dashed region and thus "shifted" (in this case to the right). A green ball was placed within the inside region of the cloak, which for some viewing positions appears to be invisible. The image of the ball is also shown with a dashed outline. One such viewing position is shown below the cloak where the red-shaded regions represent the view seen. The darker red part of the view passes through the cloak but *around* the ball while the light red region never intersects the cloak and can therefore not see the image of the ball. To clarify how a ray transverses the cloak, a ray, shown as a dark red line was drawn traveling within within the dark red region before leaving the cloak such that it preserved the initial direction and projected position (shown by a red dashed line intersecting the side of the cloak). 108
- 4.2 A 3D model and raytraced simulation of the shifty cloak shown in Fig. 4.1 using the `PolyhedralShiftyCloakVisualiser`, available on [3]. All further simulations in this section were also carried out using this simulation file. In (a), a 3D render was used to demonstrate a green sphere placed within a wireframe model of a shifty cloak. The image of this ball, due to the shifty cloak, can be seen as a green-striped sphere to the right of the cloak. When disabling the wire frame and enabling the optical slabs the full shifty cloak effect comes into play. As demonstrated for the 2D case in Fig. 4.1, the sphere is now imaged to the striped sphere location, but since there is no direct line of sight from the camera that intersects both the image of the sphere and the cloak, the sphere is completely hidden. The only visible feature is a darker outline of the cloak due to simulated absorptions for illustration purposes. 109

- 4.3 An omnidirectional cloak demonstration, (a), and simulation, (b), using two nested shifty cloaks. An inner shifty cloak was nested within an outer shifty cloak as shown in (a) outlined using solid lines. A sphere was placed within the inner cloak as the view object. When taking the first image position of the inner cloak and sphere due to the outer cloak alone, the dashed image is created to the right of the system. A vertical shift due to the inner cloak is then considered, imaging the sphere and interior space of the inner cloak to the dotted second image positions above the first image. Since there is no line of sight connecting the first and second image positions with the physical cloak, it should be fully invisible from all viewing positions. In (b), a top-down version of a ray-traced diagram of a similar scene is shown. Three light rays were shown originating from the second image position, a striped sphere, incident on the physical cloak. These indicate possible lines of sight intersecting the nested cloak structure. As made apparent, the rays do not intersect the first image and hence show that there exists no line of sight between the two images and the nested cloak. As expected, when tracing these rays through the cloaks, it can be seen that none of them pass through to the inner space of the cloaks. 110
- 4.4 Two distant regions of a 2D space shown to be connected through a simplified wormhole model. The folded space, in this case, can be seen as a blue-tinted transparent surface embedded in 3D space. The wormhole is labeled and can be seen to connect two distant regions of the folded space. When considering two positions on the bottom half of the space (positions 1 and 2), viewing two distant positions on the top half (positions 1' and 2' respectively), it can be seen that the distance travelled by the red line connecting the number 1 points will be much shorter than the distance of the number 2 points. Optically, this means that point 1' will appear to be significantly closer to 1 than point 2' will be for 2, reducing the path length of the light taken through the wormhole. A way to replicate a similar effect purely ray optically would be to image 1' closer to 1 whereby the path length is not physically shortened, but it *appears* closer in proportion to the physical path length reduction. On the other hand, for some view positions, such as 2, this should not occur, and the wormhole should be invisible from the view. This ignores any possible distortions caused by a real wormhole. 111

- 4.5 In (a), a demonstration of a 4 ideal lens setup that images an object closer to an observer, therefore making it appear larger. The distance by which the the object, a simple model of a Tim head, is imaged closer is given by Δz which is analytically expressed in Eqn (4.1) for an observer shown as an eye. The relative lens separations are shown as the sum, and scaled sum of their focal lengths ultimately creating two telescopes separated by some intermediate distance that is also dependent on their focal length. A raytraced simulation confirming this is shown in (b) and (c) created using the `IdealLensWormholeVisualiser` available at [3]. Here, (b) shows the view through a lens setup with all focal lengths set to 1 unit. This should be equivalent to a z shift of 8 units closer to the original camera position. Note that the simulation software uses arbitrary units where one floor tile represents a length of 1. When moving the camera 8 units closer in (c), and removing the lenses, it can be seen that the Tim head and background are of the same size as in (b). The outer ring in (b) represents regions of space where a ray does not intersect all 4 lenses. 112
- 4.6 The setup for an optical wormhole. A nested shifty cloak construction is shown outlined by solid lines. Within the outer cloak, there is an inner cloak and two lenses, with the inner cloak containing an additional 2 lenses. Overall, each pair of lenses is separated by the sum of their focal lengths. The first images, due to the outer cloak only, are shown using dashed outlines, consisting of the inner cloak and all 4 lenses. The second images will be due to the combination of the outer and inner cloaks and are given by dotted outlines. The exact shift amount and hence image locations were constructed such that the separation between the first lens images and the second lens images was the required distance for the 4 lens construction shown in Fig. 4.5. This was done to allow for a slightly larger shift than simply the size of the inner cloak cavity. The light grey shaded region represents the view positions where a line of sight intersecting all image positions may occur further explained in Fig. 4.7. 113

- 4.7 An illustration and demonstration of the restricted viewing directions of a shifty cloak. Similar to the ideal-lens cloak, a window into the shifty cloak may be considered to lie along the surface of the cloak over which the shift occurs. This is shown as a blue-shaded region on the right-hand side of the cloak in (a). The image positions of the interior are also shown in (a), giving the image of the slab intersections and interior as dashed lines. Using the intersections between the ray and the image positions, the trajectory through the cloak was constructed indicated as a red line, while the red crosses indicate the corresponding intersections. Two viewing positions are given as 1 and 2, with 1 being placed within a grey-shaded region representing all lines of sight that can intersect both the window and the view object (a red sphere). Unlike 1, 2 is placed outside the shaded region and therefore cannot intersect the window and the image of the red sphere. A ray-traced simulation confirming this is shown in (b) for view position 1 and (c) for position 2. In (b) the red sphere can therefore be seen while in (c) it cannot. Once again, for this, the `PolyhedralShiftyCloakVisualiser` was used. 114
- 4.8 Simulating the shifty cloak wormhole setup outlined in Fig. 4.6 using the `ShiftyCloakApplicationVisualiser` simulator available at [3] yields all images shown here. In (a), the cloaks and lenses were removed to see a small Tim head in the distance. This represents view 2 in Fig. 4.4 and acts as the control. The same scene is then viewed again but now with a shifty cloak wormhole setup, first the on-axis view in (b), then from a position from where the interior of the cloak is expected to be invisible in (c), and finally slightly off-axis in (d). The Tim head appears to be closer in (b) as desired while the lenses are fully invisible in (c) as desired. However, in (d) a dark region appears where the lenses should be. The extent of this region may be better understood when viewing the video in [7], which shows a simulated fly-around view of the shifty cloak wormhole. 115
- 4.9 A simulation of an ideal-lens wormhole sharing some of the properties seen in the shifty cloak wormhole. Here a similar setup to Fig. 4.6 was used, with some differences in lens focal length and position. The exact details were previously covered in [8] and have been summarized in appendix A. The properties seen in (a) show a Tim head imaged closer to the observer in the same on-axis view used previously. To compare the size increase, a small portion of the original Tim head can be seen in (b) while also showing the same, but enlarged Tim head towards the center. Since this is now an off-axis view, a similar black region forms due to attractor-like properties. This comparison was made in the `IdealLensWormholeVisualiser`. 116

- 4.10 The Janus device setup in (a), shows two shifty cloak wire frame models each containing a sphere. One of the spheres is red and solid while the other is green striped with each of them being imaged into the same space shown as a red and green striped sphere. Three camera positions are shown with each viewing the image from a different direction. For the camera placed at 1, some sample rays were drawn to indicate the expected view and show that the camera will see the red sphere only. The same could be done for all other camera positions which can be checked using (c-e). When viewing the setup from camera position 1, shown in (b), the drawn trajectories can be confirmed by showing that the red-green striped sphere only lies in a line of sight with the cloak containing the red sphere. (c-d) correspond to camera positions 1-3 respectively, where (c) confirms the expected outcome based on the rays shown in (a), where only the red sphere is visible. As the camera is moved around the cloak either one sphere, the other sphere or both can be seen at any one time. When both can be seen as shown in (d), half of each sphere is shown. lastly, (e) as expected shows only the green striped sphere. A video showing a round trip around the cloak is given in [7] to better understand these properties. All the simulations in this section were carried out using the simulations software named ShiftyCloakApplicationVisualiser. 117
- A.1 A schematic of the ideal lens cloak wormhole. The solid lines represent the physical space objects, where the outer cloak is cyan colored and labelled as O, the inner cloak is blue and labelled as I, and the 4 lenses within are all green, labelled as L_{1-4} . The inner cloak and the lenses are then imaged to the outside as a first image shown as dashed lines, with the same labels but primed i.e. I' and L'_{1-4} . The first images (L'_{1-4}) of the 4 lenses are then imaged once again by I' , to their final image position given by L''_{1-4} . These are shown by dotted green lines, and can be seen to be separated by exactly the 4-lens relationship outlined in Fig. 4.5 as desired. The area of permitted viewing positions for which a view direction exists where the wormhole should be visible is given by a grey background. 2

- B.1 An UE5 simulation of the refractive view rotator. In (a) a scene with a green bar tilted at 10° can be seen. The same scene is then viewed but after placing a view rotator in front of the camera in (b). It can be seen that the green bar is approximately rotated back to a horizontal orientation. The slight offset from being perfectly horizontal is due to the camera distance not being placed exactly where the view rotator was designed for it to lie. This could be fixed by creating another 3D model of the view rotator but for a further eye distance. It is for the same reason that the edges appear to be bending, with a similar effect clearly noticeable in Fig. 1.8, we conclude that the view rotator is accurately (to any noticeable degree) simulated in UE5. 6
- B.2 three screenshots of an UE5 scene containing two lenslet arrays, forming a telescopic array, designed to change magnification when shifted by one period to the side. In (a), a view without the components can be seen of just a 3D model of a Tim head. (b) adds back the component but in the basic configuration where it should return a unit magnified view. Some gray stripes are visible which are baffles required for this construction and are there by design. Besides this, the view seen mostly aligns with the expected view in (a), barring some small double images in each telescope center. However, when shifting the component by one period in (c), the view seen is significantly deteriorated, past any expected levels when comparing it with a Dr Tim simulation. This might indicate some issues with using the UE5 in the way we have done here for accurate ray tracing. 7

Acknowledgements

None of the works presented here would have been possible without the support of others. Specifically, I would like to give my biggest (and first) thanks to my supervisor and mentor throughout my PhD period, Dr Johannes Courtial. Without his support, guidance, expertise, and overall just great character, I would never have applied for a PhD position and much less managed to complete it! His consistently kind and patient attitude has made this project a pleasure to work on and will serve as a role model for my future career, regardless of the direction it may take.

I would also like to thank other academics, namely Prof. Mgr. Tom Tyc and Dr Jakub Belin, who have helped and provided me with great advice for my research and contributed significantly to various projects and publications. Also, a big thanks to all the project and visiting students for working with me on various projects, evolving my own projects as it went on.

On a social level, I have to thank all the members of the office who have made it an awesome place to work over the past years. Specifically, I have to thank Lewis for golfing (and FF) with me, Andrew and Paul for the nice games of football, Karlin for the lovely camping trips, Conall for his book recommendations and knowledge, and lastly Giedre for helping me paint the real Dr Tim in Fig. 2.21. But most of all, I would like to thank you guys and everyone else in the office (who I will leave out for brevity) for being great friends to me and sticking with me throughout all my yapping! The same goes for all my non-office friends for everything they do with me.

The largest pillar of support through all my studies and life as a whole has been my family. Without them pushing me and guiding me in the right direction, none of this would have ever been possible. I love you all ∞ !

Finally, this work was supported by the Engineering and Physical Sciences Research Council (grant number EP/T517896/1)

Declaration

With the exception of the first section in chapters 2, 3, and 4 (specifically the introduction sections and 2.2, 3.2, and 4.2) or where outlined in the text, all the work in this thesis was carried out by the author. In all chapters, the writing was done by the author. The simulation software used was written by the author, building on the existing ray tracer Dr. Tim [4], unless otherwise stated.

Maik Locher

Publications

The following papers have been written over the PhD period and published or are in preparation for publication:

- Gordon Wells, YiChen Wu, Maik Locher, and Johannes Courtial, "Geometric construction of relativistic and non-relativistic distortion," J. Opt. Soc. Am. A 41, 1329-1336 (2024)
- Johannes Courtial, Jakub Běĺín, Matúš Soboňa, Maik Locher, and Tomáš Tyc, "Shifty invisibility cloaks," Opt. Express 32, 11-25 (2024)
- Ivan Armstrong, Maik Locher, and Johannes Courtial, "Adaptive Fresnel lens: basic theory," J. Opt. Soc. Am. A 42, 211-220 (2025)
- Maik Locher, Elise Blackmore, Zhao-Kai Wang, and Johannes Courtial, "Wave optics of spiral adaptive Fresnel lenses", In preparation (final draft completed)
- Maik Locher, Ivan Armstrong, Zhao-Kai Wang, Di Wu, and Johannes Courtial, "Adaptive Fresnel lens: Advanced topics", In preparation
- Maik Locher and Johannes Courtial, "View rotation spectacles", In preparation

Introduction

Micro optics deals with optical components that are, individually, relatively small (compared to historic optical devices) but may also be used in conjunction to form much larger arrays. Consisting of many components found in classical optics, but much smaller in size, such as lenses and mirrors, the field only began its rise to popularity in the 1980s, with the term first being used in [9]. Improvements in manufacturing capabilities, specifically in the semiconductor industry, permitted the production of smaller optical components. A working precision in the micrometer range then allowed for the construction of micro-lenses and mirrors. Nonetheless, at the size range considered, these miniature classical components act much in the same way as their larger counterparts, barring diffractive effects. The field quickly expanded, incorporating and molding several new components and concepts, such as the commercially available digital micromirror devices (DMD) [10], arrays of controllable micromirrors used in display technologies. The catalogue further grew with the inclusion of optical components that were invented before the term “micro optics” ever gained popularity, such as diffraction gratings (first studied in the 18th century) [11], and holography [12]. To us, a further notable addition are lenslet arrays. These consist of –in the micro optics regime– microscopic lenses placed next to one another to form a larger array.

Nowadays, these have been used in various fields, from thin virtual reality glasses [13] to LiDAR systems in autonomous vehicles [14], slimming down components or creating configurations that enable novel approaches to existing technologies. Another application was found as a Gabor superlens [15], a special lens made up of two lenslet arrays. In the Gabor superlens, each lenslet pair may be thought of as a pixel, forming an elemental image such that the array as a whole creates a clear integral image.

This “pixellated” approach to optics was and may be used in several novel ways, to create thinner components for previously impossible view transformations, such as a rotation and/or magnification, of the view. Utilizing small refractive wedges (instead of lenslets) as our pixels, we demonstrate, using ray trace simulations, an approximate view rotation device sufficiently thin to work for spectacles in chapter 1. These were theorized and designed to aid in the treatment of torsional diplopia, a condition that causes double vision due to a rotational eyeball misalignment. The view quality for these was evaluated using Snellen charts as an initial evaluation of whether or not these will be able to help patients suffering from torsional diplopia.

Staying with pixelated optics, in chapter 2, we explore a new form of an adaptive lens, whereby the focal length can be tuned through the relative rotation of two components. The pixellation here consists of individual windings of a cylindrical lens spiral, which, in conjunction with its counterpart, another cylindrical-lens spiral, forms an equivalent Fresnel lens. Unlike some of the current adaptive lenses in existence, these were specifically designed with the form factor and range of common refractive errors in mind, which affect billions of people globally.

The use of lenslet arrays has previously been considered as a possible way to create an ideal lens cloak (a special case of the omnidirectional cloak in [16]). However, this has not yet been realized due to a number of issues. Nonetheless, we retain an interest in studying these cloaks theoretically to further understand the cloaking properties in chapter 3, which were then used to describe some of the phenomena seen in another invisibility cloak we call the “shifty cloak” in chapter 4. Unlike the ideal lens cloak, this was designed entirely with (mostly) optically feasible materials in mind. These cloaks were then used to simulate two potentially exotic applications of the shifty cloak outside of just being an invisibly cloak, namely a Janus device and an optical wormhole.

Chapter 1

View rotation spectacles

1.1 Introduction

Upon every reflection or refraction, a light ray changes direction. In a purely geometric ray optical sense, it may seem like any possible transformation is permitted, but under closer wave optical inspection it becomes apparent that any transformation which does not maintain a continuous wavefront upon propagation should not be permitted by a single (or in some cases multiple) reflections or refractions. Mathematically, this can be understood as transforming a zero curl incoming light ray field into a non-zero curl field which cannot occur [17].

An exception is made with pixellated optics, whereby each pixel consists of a small piece of the overall wave front. And while the wavefronts are locally (within each pixel) required to be continuous, neighboring pixels do not need to align the wavefronts, causing the overall field to contain discontinuities [18]. Hence, without considering a pixelated approach, there exist certain optical transformations that are optically “forbidden” with a thin, planar, light-ray-direction-changing component. One such forbidden transformation is a simple image rotation. No single light-ray-direction-changing (but not translating) surface can optically rotate an image without introducing wavefront discontinuities. As a result, there are limited treatment options for certain medical conditions such as torsional diplopia.

Torsional diplopia is a condition caused by a cyclodeviation, whereby one eyeball is rotated with respect to the other along the back-to-front axis of the eye. Since the human brain is exceptional at fusing images, it is surprising that for any rotation $> 4^\circ$ [19] diplopia (double vision) may start to appear, and with a relative rotation of 6° [20] or more fusion of the view is not possible for most patients. To understand the impact this may have, a simulation showing the binocular view of someone suffering from torsional diplopia is demonstrated in Fig. 1.1. Besides this, studies have found that the quality of life (QoL) for patients with this condition is severely diminished, with reports stating that the difficulties manifest in both the near and far vision [21]. As a result, everyday tasks such as driving and even simply walking become difficult [22]. When questioned, patients also expressed a desire to decrease their life expectancy if it meant

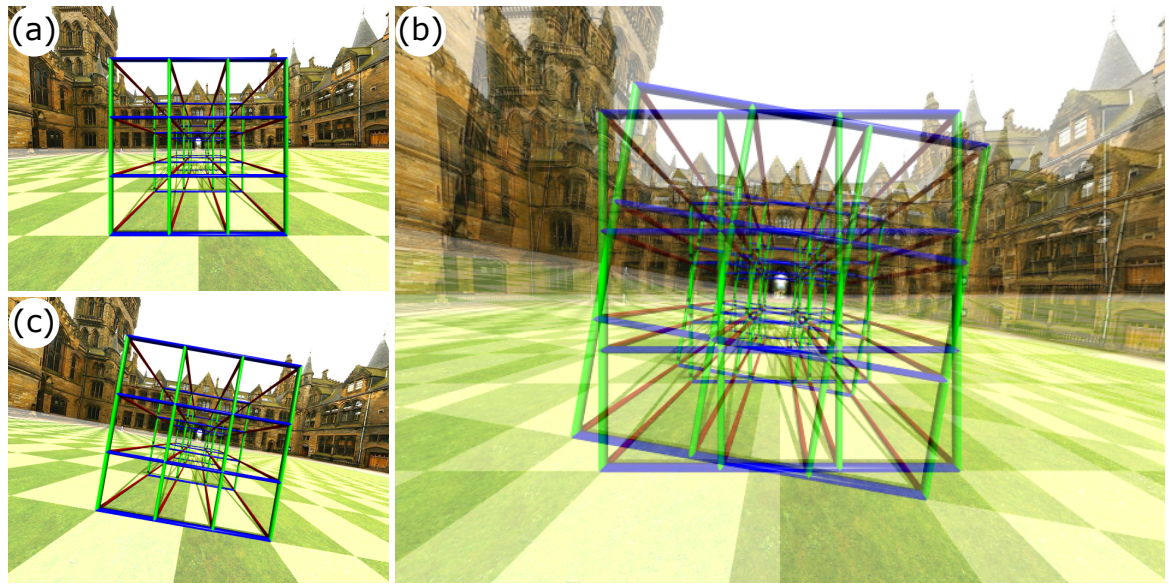


Figure 1.1: A simple demonstration of the expected view for someone suffering from torsional diplopia. A scene containing a lattice as a view object and the university quadrangle as a back-drop was chosen for demonstration purposes. Here (a) represents the view seen by the healthy unrotated eye, while (c) is an eye that has been rotated such that the view seen appears to be rotated by 10° in the clockwise direction. To estimate the impact this may have on the binocular vision (vision due to both eyes), the views seen in (a) and (c) were overlaid with equal intensity to form (b). It can be seen that there is severe diplopia.

increasing their QoL through a relief of their condition.

Torsional diplopia is possibly more widespread than one might think at first. A thorough search did not yield any estimate of the number of people suffering from torsional diplopia. But a very rough estimation to the number of people suffering from torsional diplopia can be made by combining and extrapolating related statistical data (from different cohorts) – not a great way of arriving at a good number, but the best we could come up with. An estimated 3% of reported diplopia cases are due to torsional diplopia [23], from a retrospective review with a sample group of 266 patients aged 15 years or older. From the year 2003 to 2012, an estimated 850,000 annual hospital visits [24] were reported within the United States(US) relating to diplopia, which relates to roughly 0.28% of the US population. Applying this to the whole world, an estimated 22 million diplopia cases should be reported globally per year. Now we note that this number may vary greatly since there is a lack of exact statistical data, and may already be globally inaccurate due to differing reporting rates across the world and access to health care. It should also be noted that the 850,000 estimation was based on diplopia being the primary reason for a hospital visit, not considering cases where diplopia may appear as a side effect of a more serious condition. The study was also not explicitly clear on the statistical impact of recurring hospital visits. Keeping this in mind, an estimated 660,000 global cases are due to torsional diplopia yearly. In [23], it was also found that roughly half of the torsional diplopia cases experienced spontaneous improvements in the condition within half a year, although this greatly depends

on the origin of the cyclodeviation. Again, applying this, we very roughly estimate that there are 330,000 global cases of torsional diplopia every year that do not improve without medical intervention.

The current treatment option of choice (where available) is a surgical correction called the Harada-Ito surgery, which was established in 1964 [25, 26]. Besides this, the remaining treatment options are typically less appealing, with the most common being simple visual occlusion [27], for example, by use of an eye patch. This means that in some instances where the surgical procedure cannot be done or where it simply is not available due to expense or health-care reasons, no modern alternative can be used to restore the vision and occlusion is the only option. This is the unfortunate reality in many of the developing regions of the world. Indeed, even in the modern world, the success rate of the Harada-Ito procedure is not always guaranteed. Currently, the success rate ranges between 43% to up to 93% [28, 29]. This large variation is strongly linked to the choice of patient, which can limit complications and the definition of a successful treatment. In all the remaining cases, the outlook is similar to that of the developing world; ocular occlusion. In the case of a successful surgery, the result is often an over-correction which the eye will naturally settle back into during the healing period. This over-correction can last up to 12 weeks [30], during which the patient may still experience intractable torsional diplopia. All of this raises the question of whether there is any alternative in which the view seen can be optically rotated, similar to spectacles being an alternative to laser eye surgery.

There exist a number of ways for an image to be rotated using optical elements. These can range from purely refractive approaches [31] to prismatic approaches, encompassing refractions and reflections in Dove prisms [32] or other prism types [33]. One issue with the prismatic approach is the bulkiness. When considering a single, or for the purposes of rotation in Dove prisms two, prisms, the size required to fill the clear aperture of a pair of spectacles results in an optical component too large to wear comfortably. Since the prism requires a minimum depth-to-aperture-diameter to perform the necessary transformations, this size cannot simply be reduced to fit.

One way to reduce the size is to use components consisting of optical-component arrays, which will have a smaller clear aperture to fill and hence permit a thinner form factor. One such attempt was to use telescopic arrays [31], called ray rotating sheets, consisting of two lenticular arrays. By a process we call “ray flipping”, whereby two reflections occur with respect to different planes, inducing a rotation, which is further explored and explained in section 1.5.2. Ultimately, the arrays are rotated with respect to one another, and the view through them appears rotated and magnified, as also seen in the Moiré magnifier [34] when constructed using two sets of lenslet arrays. Since the view becomes slightly magnified, binocular image fusion may be hindered. In both cases, these will also suffer from cross-talk, a commonly seen issue with lenslet arrays, leading to non-standard refraction that results in unwanted additional images [35]. Replacing the lenticular arrays with prism arrays [32] introduces other issues, such as

an increase in manufacturing complexity to ensure any required total internal reflections occur without leaking into neighboring prisms. Generally, the effective field of view of all the methods mentioned is also limited, making them potentially less feasible for the use in spectacles.

Ultimately, this all comes back to the issue that image rotation is an optically forbidden transformation by a single, light-ray-direction-changing surface. Here we will investigate a different, pixelated approach to create an approximate *view* rotator in section 1.2. The original view rotator was presented in [2], where the view seen through the view rotator component appears to rotate anything seen through it by the same amount, regardless of position past the component, but only when viewed from a specific distance from the component. In our case, this is not true, and as later explored in section 1.4, the apparent rotation varies with distance from the component. To avoid confusion, wherever talking about a view rotator, the approximate case will be referred to unless specifically stated otherwise.

The visual quality of the view rotation spectacles was determined through ray tracing simulations throughout section 1.3, using Dr. Tim [4], our open source raytracing software. This component deteriorates the visual quality, and later we show the effect on visual acuity using a Snellen chart [36]. The remainder of the chapter will be dedicated to improving the effect on visual acuity in section 1.5 using several methods, to finally simulate the component with the best visual acuity found. A discussion and conclusion will summarise it all in section 1.6 and give insights to the future outlooks and objectives to realize a view rotator that can end up aiding in the treatment of torsional diplopia.

1.2 Simple view rotator construction

As outlined in section 1.1, an image rotation cannot be carried out using a single refractive surface. Instead, a view rotation was proposed whereby each pixel redirects an incoming light ray bundle to the appropriate rotated position. Importantly, the individual pixels do not produce a rotated image but instead an upright but shifted one.

In our simulations, we consider a ray casting approach whereby the “camera” sends out light rays into the scene. Therefore, a single-point *emitter* was considered to construct the pixels. This source was placed at a distance, d_{eye} , from an ocular (eye-sided) surface as shown in Fig. 1.2. Furthermore, the ocular surface was chosen to be perpendicular to a light ray entering the pixel center from the point source. Snell’s law ensures that such rays, which we call 0^{th} order rays, pass through the ocular surface without any light ray direction changes. Therefore, the objective surface (object-sided) fully imposed the ray direction change needed to create the apparent view rotation.

A vectorial form of Snell’s law [37], given by

$$\hat{\mathbf{N}} \times \hat{\mathbf{d}}_2 - n\hat{\mathbf{N}} \times \hat{\mathbf{d}}_1 = 0 \Rightarrow \hat{\mathbf{N}} \propto \hat{\mathbf{d}}_2 - n\hat{\mathbf{d}}_1, \quad (1.1)$$




Figure 1.2: A cross-sectional sketch of the view rotator. The individual pixels (here shown as blue quadrilaterals) and the eye position are not to scale for illustration purposes. Green light rays originating at the eye at a distance of d_{eye} from the ocular (closer to the eye) surface of the view rotator, are shown to pass through the ocular surface without any change in light-ray direction. At the objective surface, this ray splits into a branch ray shown as a green dotted ray, which indicates the unrotated ray direction, explained further in Fig. 1.3. The solid green lines represent rays that have been refracted along a direction of the rotated view by the respective pixels. It can be seen that the central pixel does not in fact redirect the ray through it as expected. The individual pixels are separated according to the black dashed lines in order to minimize non-standard refracted rays. These would be light rays that enter one pixel and pass through a boundary to a neighboring pixel. A pixel thickness is described between the ocular and objective (further from the eye) pixel centers as t . Each pixel also comes with an index described by (n_x, n_y) which may be used to determine the objective surface normals, however here only the central ones are shown ranging from $(0, -1)$ to $(0, 1)$. The span, s , of the pixels is constant across the whole view rotator and dictates the pixel size.




Figure 1.3: The construction used to determine the ray direction change and corresponding objective-surface pixel normals. A light ray, shown as a dotted green line, is traced through a pixel center without any deviations in the same way seen in Fig. 1.2. The ray is then intersected with a plane at a distance, d_{rot} also known as the rotation plane distance, from the objective surface. This plane is where the view should appear to be rotated. The intersection position is then rotated about the center by θ . Therefore the normal at the objective surface should be set to redirect an undeviated ray with direction vector $\hat{\mathbf{d}}_1$ (shown as the dotted green line) onto the desired rotated direction, $\hat{\mathbf{d}}_2$, shown as a solid green line.

where $\hat{\mathbf{d}}_1$ and $\hat{\mathbf{d}}_2$ are the unit direction vectors of the incoming and outgoing ray respectively, n is the refractive index ratio and $\hat{\mathbf{N}}$ is the unit normal of the refracting surface, was used to find the objective normal. The required light-ray direction change was dictated by the 0^{th} order ray direction, and the refracted light-ray direction required to make a planar object behind the view rotator appear rotated when seen from the eye position as demonstrated in Fig. 1.3.

Furthermore, using the parameters shown in Figs. 1.2 and 1.3, the parameters in Eqn (1.1) can be rewritten as

$$\mathbf{d}_1 = (sn_x, sn_y, d_{\text{eye}}), \quad (1.2)$$

where the coordinate system is set up with the plane of the view rotator being perpendicular to the z direction (shown in Figs 1.2 and 1.3) and

$$\mathbf{d}_2 = \left(\frac{s}{d_{\text{eye}}}((t + d_{\text{eye}} + d_{\text{rot}})(n_x \cos \theta - n_y \sin \theta) - n_x(t + d_{\text{eye}})), \right. \\ \left. \frac{s}{d_{\text{eye}}}((t + d_{\text{eye}} + d_{\text{rot}})(n_y \cos \theta + n_x \sin \theta) - n_y(t + d_{\text{eye}})), d_{\text{rot}} \right), \quad (1.3)$$

where θ is the desired rotation angle, n_x and n_y are the pixel number, s the pixel size, and t, d_{eye} ,

and d_{rot} are the thickness, eye to ocular surface distance and objective surface to rotation plane distance respectively all shown in Figs. 1.2 and 1.3. Eqn (1.1) then becomes

$$\mathbf{N} = \frac{\mathbf{d}_2}{|\mathbf{d}_2|} - n_2 \frac{\mathbf{d}_1}{|\mathbf{d}_1|}, \quad (1.4)$$

where n_2 is the refractive index. This can be evaluated, but the result is messy and long. It should be noted that some geometric assumptions were made to simplify the result, such as a view rotator plane that is parallel to a plane containing the ocular pixel centres, also shown in Fig. 1.2. These are outlined in, a Mathematica [38] notebook, along with the derivation and the full form of Eqn (1.4), which is available in the supplemental material as

`NormalCalculationsViewRotator [1]`.

1.3 Simple view rotator results

Applying the process outlined throughout Eqns (1.2-1.4), a whole array of “pixel wedges” was constructed and simulated to induce a counterclockwise rotation of 10° . Throughout the chapter, the pixel size was chosen to be 0.3 mm, as it provided a view quality that appeared to yield the approximately largest visual acuity, determined by visual inspection, when considered with diffractive effects (explored in more details in section 1.4). The stack of images used to determine this value can be found in [39], with the respective pixel size given in the file names. Along with this, the distance between the simulation camera (which we will call the eye from here on) and the ocular surface, d_{eye} was chosen to be 1.5 cm, which roughly corresponds to the expected position of regular spectacles. The thickness, t , was set to be 1 mm to keep the spectacles small and therefore lightweight. It should be noted that this can be increased if required for manufacturability and structural integrity without any significant loss in performance. Lastly, the rotation plane (a plane where the rotation angle should be exactly 10°) distance, d_{rot} , was set to be 1 m in front of the view rotator. These represent the default parameters used in the simulation, any deviations from these will be noted in the text.

A model of the view rotator, using refractive components, was then constructed and simulated, where each light ray was redirected by a refractive wedge created based on the surface normals found when considering the relevant central “construction” ray. For the initial scene, the eye was rotated clockwise by 10° , resulting in the view rotator effectively correcting for this and making the view appear upright. The resulting view can be seen in Fig. 1.4(c), while a binocular view of this can be seen in (a). The eye pupil (or in simulation terms the camera aperture) was also chosen to represent that of a human eye and set to a 2 mm radius. Furthermore, to mimic the human vision, the field of view was set to 130° which is roughly that of a human eye.

Considering a side-by-side comparison of the expected binocular image in Fig. 1.1(b) and Fig. 1.4(a), it can clearly be seen that the view rotator successfully corrects for the rotation of

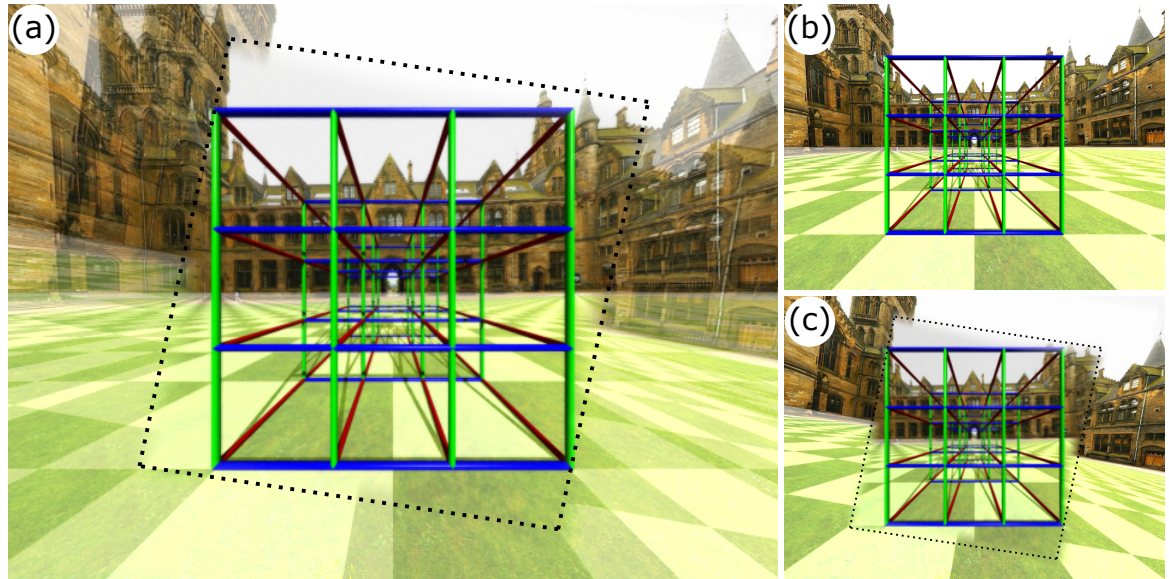


Figure 1.4: The view seen through the simple view rotator at a 130° field of view. In (a), the binocular image can be seen with a dotted border representing the edge of the spectacles. This was constructed by equally overlaying the view seen in (b) and (c). Here (b) is the control view expected of a healthy eye without any component or rotations added. On the other hand, (c) represents the view seen with a 10° torsion in the clockwise direction, when looking through the simple view rotator designed to rotate the view by 10° anticlockwise. Another dotted border was added around the spectacle perimeter in (c) to indicate where the correction ends. The slight blurring of the view through the view rotator in (c) is due to the extended aperture size of the eye, which was 4 mm in diameter. As a result, a seemingly upright image is seen when looking through the rotator. The sides remain rotated as the view rotator was chosen to not fill the entire field of view.

the eye. As this is for a large field of view (130°) but viewed as an image that takes up a smaller angular size, the blurring appears significantly reduced. This is further explored in section 1.4. Some further testing was done to determine how the simple view rotator behaved under non-standard, but no less relevant, circumstances.

1.3.1 Off-axis view

Initially, the off-axis view was considered. This was done to mimic the behaviour of natural eye movements. Generally, the pupil can automatically move up to 5° when focusing on a relatively still scene directly in front of the eye [40], to scan and obtain high-quality information about the scene (by imaging it into the fovea). Within this range, this occurs predominantly reflexively rather than any purposeful motion called saccades. For moving scenes, the amplitude of the saccades can be even larger ($\sim 15^\circ$). On the other hand, the limit is typically around 30° [41] before the head tilts to more comfortably view a certain direction, in order to avoid straining the eye.

To simulate this, the camera was placed onto an eyeball-sized sphere to accurately represent the position along with the viewing direction, with the pupils in Fig. 1.5 representing the camera aperture. The scene shown was rendered with a 17° angle between the viewing direction and the normal to the plane of the view rotator, resulting in the lower portion of Fig. 1.5. Note that 17° was chosen as it gives an indication of the expected off-axis view quality for the lower (5°) and upper (30°) limit of eyeball movements. For the large field of view case explored here, the results suggest that a small offset of just 5° should result in a nearly perfect image, while a 30° offset would result in a significantly worse image. Both of which may be tested using our open source ray tracer, specifically the runnable Java archive called `ViewRotationExplorerWithUnits` available in [3].

Interestingly, it appears as though the view seen through the view rotator was shifted downwards when comparing Fig. 1.5(a) and (b). This can be seen to a greater effect when overlaying both figures forming a binocular image (Fig. 1.6). In this case, the shift is manifested as double images, appearing vertically offset. Upon further investigation, it became clear that the offset was, at least in part, due to the way the eyeball movement was considered. A healthy eyeball will have an “up” axis along what will be referred to as the y axis. When the eyeball is initially rotated about the y axis, only the corresponding azimuthal direction is considered as shown in the top part of Fig. 1.5(a). However, this does not hold true when accounting for the torsion of the affected eye, which as a result will have a different “up” axis to the healthy eye. In fact, the up axis will be rotated by the torsion magnitude to have a new up axis, y' . This makes it inherently unclear if the affected eyeball will twist around the same up axis as the healthy eye, y , or if it will twist around the new up direction, y' . In terms of the other eye coordinates, the torsional effects rotates the eyeball along the back-to-front axis, leading to an inclination of the azimuthal plane (spanning the x and z axis in Fig. 1.5(a), top) by the torsion angle as made apparent when

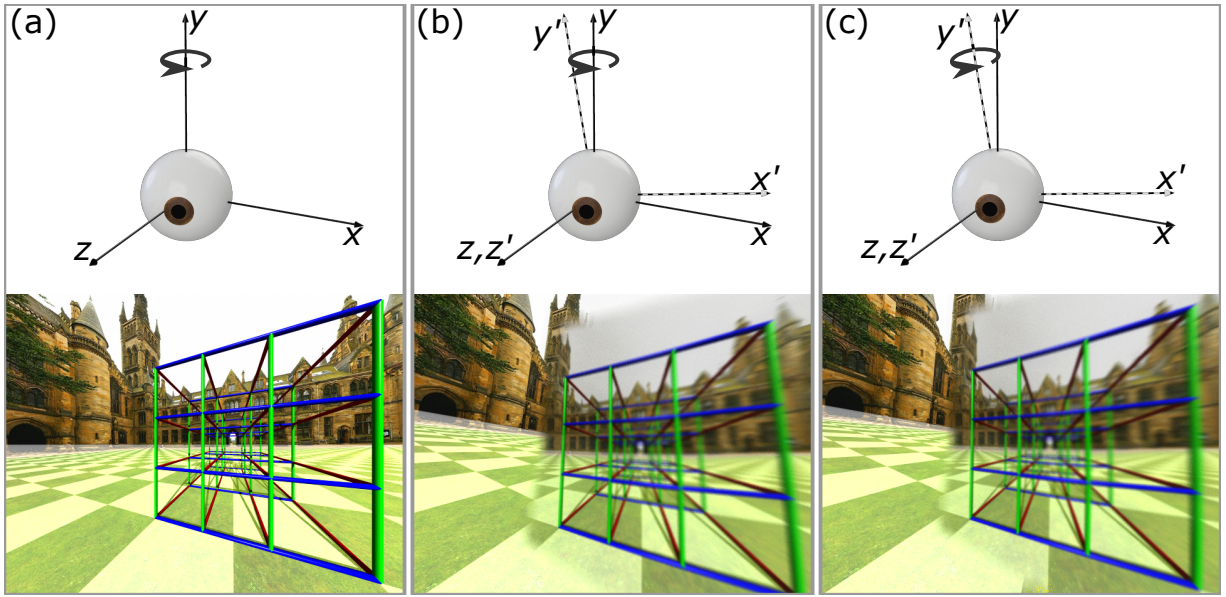


Figure 1.5: Off-axis images as seen through the view rotator ((b) and (c)) and with the “healthy” unrotated eye on the bottom of (a), with the corresponding eye (which directly relates to the camera) position above. The camera was placed such that the aperture coincided with the pupil and the back-to-front axis was the camera’s optical axis. For simplicity, we will refer to this camera setup as the eye. The top row indicates the eyeball rotation axis using an arrow. In (a), the eye was rotated about the y direction, which represents the up direction of a healthy (non-cyclodeviated) eye. This can be compared to a cyclodeviated (unhealthy) eye rotated about the same up direction axis in (b). The effective up direction for a cyclodeviated eye will no longer align with the up direction of a healthy eye. Instead, the new direction is tilted by the cyclodeviation angle and can be seen as y' in the top row. Since it is inherently unclear (which may be an interesting future research question) which axis the unhealthy eye will rotate about, both were shown with (c) showing the alternative to (b), where the rotation axis was chosen to be y' . It can be seen that (c) more closely resembles the desired view shown in (a).

comparing the new azimuthal plane spanning x' and z' in Fig. 1.5(b).

To adjust for this extra consideration, the eyeball around the y' direction, a shift of the pupil along the inclined azimuthal plane in Figs 1.5(c) and 1.6(b) was achieved. The resulting view much closer resembles the expected view (Fig. 1.5(a)). The remaining effects visible here are now entirely due to the interplay between the view rotator and the shift in the camera position (as opposed to a change in view direction) resulting in a lesser double image. One way to oppose the effects of the shift in pupil position may be to construct the view rotator around the point of rotation of the eye, which would remain fixed, instead of the centre of the pupil. This should then yield a view that remains unaffected through eyeball-movement-related shifts but retains the off-centre view properties explored here.

It can be inferred that for larger off-axis viewing angles the view rotator will not fully accurately replicate the expected view. While even in the best case scenario (Fig. 1.6(b)), the binocular view seen does not fully match. However, an improvement to the binocular view




Figure 1.6: The binocular view when looking at an angle of 17° off axis, about the y direction in (a) and about the y' direction in (b) and (c). In (a), the binocular image corresponds to the combination of the bottom images in Fig. 1.5(a) and (b), taking half the pixel value from each. Similarly, (b) is the combination of Fig. 1.5(a) and (c), showing a significant improvement to the binocular view. Lastly, a comparison to the expected view seen without the component, also rotated along the y' direction is given in (c).

without any component, shown in Fig. 1.6(c), can still be expected. While the off-axis view seen through the rotator is improved this may still be insufficient for the brain to fuse the images and cannot be tested without real clinical trials. For smaller angles, the view seen through the component becomes more akin to the desired unrotated view making image fusion more likely, but again not guaranteed.

1.3.2 Further variations

In addition to the eye movement, we also consider the change in view when the proposed spectacles are not used at the designed distance, that is if the eye distance does not match the d_{eye} , in Fig. 1.2, used when designing the spectacles. Since the spectacles were designed for a pinhole aperture at a set distance away, the expected result is a simple decrease in image quality. But interestingly, that is not the only observed property. When considering the same view rotator used in the simulation this far, and moving the viewing position further back than intended, the apparent rotation in the rotation plane is more than expected. In contrast, the apparent rotation is smaller when the viewing position is closer. This can be seen best in the center of the two views shown in Fig. 1.7 (a) and (b).

While this may show a decrease in binocular image quality shown in Fig. 1.7 (c) and (d), it may also add some flexibility to the design. In particular, when the angle required to correct the cyclodeviation is not exactly what the component is designed to do, a user could move the spectacles closer or further from the eye to achieve the desired rotation. To quantify this, images similar to the ones shown in Fig. 1.7 were taken for realistic eye distances (here chosen to be between 1 cm and 2.5 cm). The apparent rotation was estimated by measuring the angle above a horizontal line in Inkscape. It was found that the apparent rotation can vary by roughly $\pm 4^\circ$. It should be noted that the rotation achieved is not fully constant across the whole FoV, with

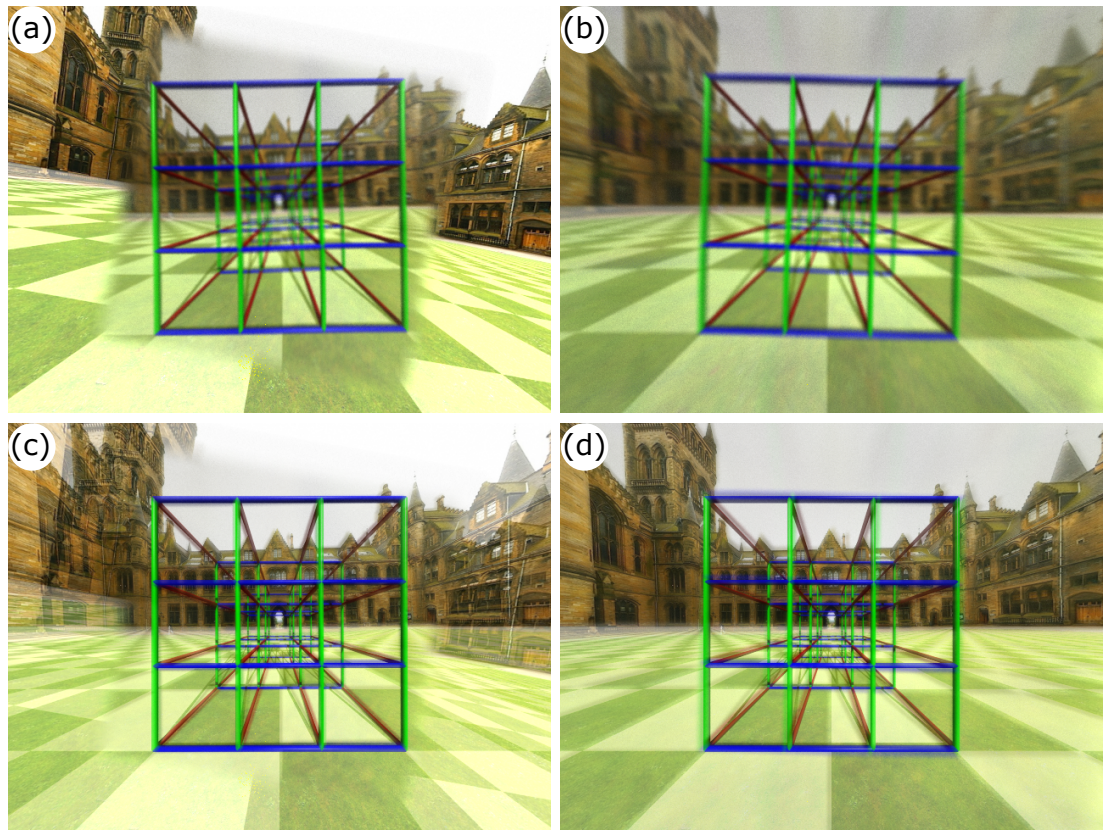


Figure 1.7: A demonstration of the apparent rotation when the view rotator is seen from a distance further (a) and closer (b) than it was designed for. In (a), the distance to the eye was 2 cm while in (b) it was 1 cm. In both cases, the view rotator was designed for an eye distance (d_{eye} in Fig. 1.2) of 1.5 cm, roughly denoting the normal pupil-to-spectacle distance. The most notable change is an apparent over (a) and under (b) rotation. The corresponding binocular images, produced by overlaying the views of (a) and (b) with an upright view, are shown in (c) and (d) respectively.

a smaller rotation further from the image center. In the case of a 2.5 cm pupil-to-spectacle distance, the extra rotation in the central $\approx 47^\circ$ FoV of the view is $\sim 5^\circ$ based on the curve seen in Fig. 1.8(b), where a black horizontal line was viewed with the same setups as shown in Fig. 1.7, using a blank background. The shape is indicative of the nature across the whole FoV of this additional rotation angle. Similarly, for the 1 cm pupil-to-spectacle distance this value was $\sim 3^\circ$, shown in Fig. 1.8(a). It should also be pointed out that this is precisely one of the ways in which we can describe that the view rotator is non-imaging. The rotation amount for a fully imaging view rotator should not change based on distance from the rotator, while also maintaining the appropriate parallax when shifting the view position, both of which the view rotator does not uphold.

Interestingly, there also appears to be a vertical shift in the position of the (roughly) horizontal line segment for any FoV angle $\gtrsim 47^\circ$. This was not significantly noticeable in the Fig. 1.7 simulations but appears to become more apparent with larger view rotation angles.

Nonetheless, this means that only a few standardized rotation amounts may need to be con-

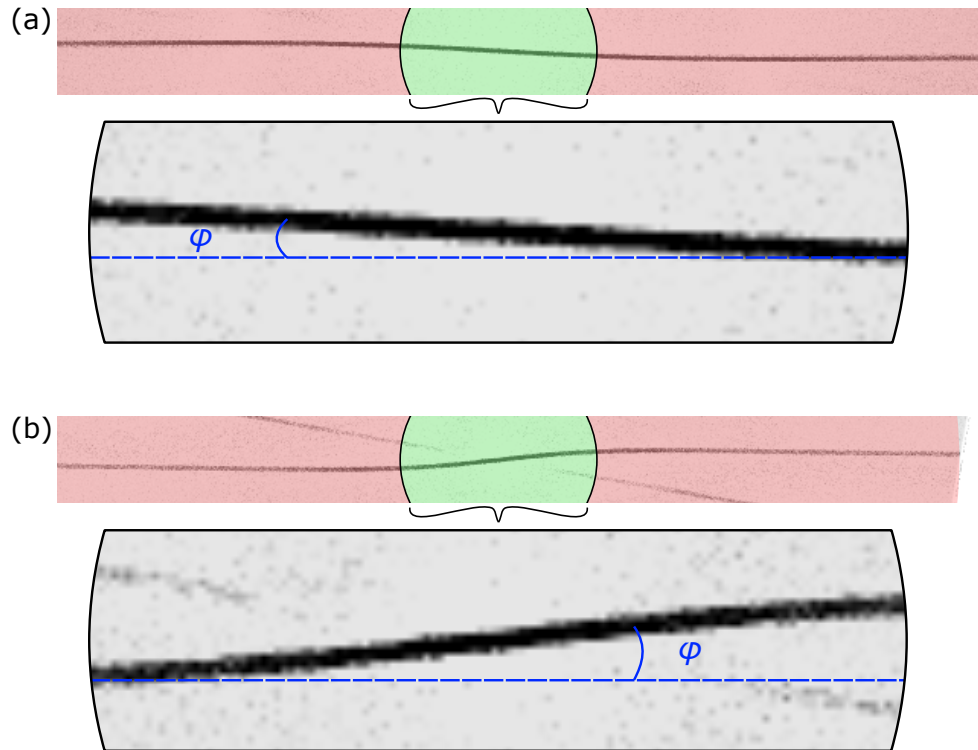


Figure 1.8: The view seen (by an eye with a 10° cyclodeviation) of a horizontal line through the view rotator. The view rotator used was designed for an eye distance of 1.5 cm, while the eye was placed at 1 cm and 2.5 cm in (a) and (b) respectively. This resulted in the horizontal black line no longer appearing horizontal but curved instead. This curvature can be categorized into two sections, the inner section represented by the green shaded area spanning a FoV of 47° , and the outer red shaded region containing the FoV in the range of 47° - 130° . Within the green region, it was estimated that the line is roughly straight to determine the apparent over (in (b)) or under (in (a)) rotation. In both of these this is shown as ϕ , where $\phi \approx 3^\circ$ in (a) and $\phi \approx 5^\circ$ in (b). In total, this results in an adaptive range of $\sim \pm 4^\circ$.

sidered when designing and producing a physical component which can, in turn, reduce the development and production cost, at the least for the central-most 47° .

However, before any realization can be considered more scenarios need to be simulated. One of these is the change in rotation amount. To this point, the simulations shown have all been for a 10° rotation angle and correction. Since the ultimate goal is to aid in the treatment of most torsional diplopia cases, some lesser and more severe cases need to be considered. To represent these, a cyclodeviation, and subsequent correction, of 5° and 15° respectively were chosen. The resulting views, shown in Fig. 1.9, clearly indicate that a larger correction angle results in a more blurred view through the component, while a smaller correction results in a clearer view. As a result, the angles that can be corrected appear to be limited, with larger corrections leading to a more blurred view and hence less useful information for the binocular image. A more detailed analysis of this and a discussion on general blurring is shown in section 1.4, where the pixelation of the component is studied in more detail.




Figure 1.9: Raytrace simulations of the view rotator at a 5° (a) and 15° (b) rotation angle. In both cases, the simulations were set up to, first, induce a respective 5° and 15° cyclodeviation and then a correction using the appropriate view rotator. It can be seen that a lower rotation angle and cyclodeviation in (a) results in a sharper view compared to a much more blurred view in (b).

1.3.3 Apparent magnification and rotation

As the component is designed to rotate the view at a defined view distance, indicated by the dashed plane in Fig. 1.3, any view distance outside of this plane will not necessarily be rotated or magnified by an appropriate amount. To fully understand this a mathematical analysis of the apparent rotation and magnification was carried out using Mathematica. To understand the process behind this, Fig. 1.3 was modified into Fig. 1.10 (a), where the same construction is used but a new sample plane is placed to sample the light ray intersections at any distance past the spectacles. The intersection between both the rotated and unrotated rays is then found within this sample plane, shown in Fig. 1.10 (b). From these points, it is possible to extract the magnification by considering the scaling between the radial distances ($|\hat{\mathbf{u}}|$ and $|\hat{\mathbf{v}}|$) and the center. For the apparent rotation, the angle of elevation of both $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ is found. The difference between these is then the apparent rotation angle, within the sample plane, indicated in Fig. 1.10(b) as θ_a . Now, the apparent rotation and magnification for any depth can be found by varying the sample plane distance without changing the spectacle setup. This results in two analytic expressions as well as an interactive graph, both of which can be found in the Mathematica document called `ApparentRotationViewRotator` in [1].

Using this sample plane, it was confirmed that the apparent rotation and magnification both reached the desired value at the rotation plane distance as expected. Furthermore, both the values remained close to the desired amount while not exceeding them by any significant amount when going to an infinity distance, provided reasonable values were chosen.

The exact relationship of the apparent rotation is complicated but can be approximated by

$$\theta_{a\infty} \approx \frac{d_{\text{rot}} + d_{\text{eye}}}{d_{\text{rot}} - t} \theta, \quad (1.5)$$




Figure 1.10: In (a), a modification to the setup seen in Fig. 1.3 where a sample plane has been added in grey at a distance of d_s away from the view rotator. Note here that it is shown before the rotation plane but can in practice be placed anywhere in front of the view rotator. The light ray intersections at the sample plane for the rotated (in blue) and unrotated (in green) rays are shown as dots of the respective colors. Within the sample plane in (b) the two intersections can be seen to lie at $\hat{\mathbf{v}}$ and at $\hat{\mathbf{u}}$ for the unrotated and rotated ray respectively. These vectors can then be used to find the apparent rotation, θ_a , by considering the angle between them. Similarly, the ratio of the magnitudes between $\hat{\mathbf{v}}$ and $\hat{\mathbf{u}}$ may be used to find the apparent magnification.

where the gradient (the terms besides θ) is an expression obtained by taking the derivative with respect to θ of the full apparent rotation at infinity expression and evaluating it at $\theta = 0$. Here $\theta_{a\infty}$ is the (approximate) apparent rotation at infinity, d_{rot} and d_{eye} are the rotation plane distance and pupil-to-spectacle distance described in Fig. 1.2 and Fig. 1.3 respectively, t is the thickness of the view rotator, and θ is the desired view rotation angle. The graph in Fig. 1.11(a) shows this equation working nicely for small rotation angles and (b) shows where it breaks down for larger rotation angles and a non-zero thickness. All of these graphs can also be interactively varied in the `ApparentRotationViewRotator` document in [1]. The apparent magnification at infinity, $M_{a\infty}$, yields a simpler expression,

$$M_{a\infty} \approx \frac{1}{\sqrt{1 + \frac{2d_{\text{eye}}}{d_{\text{rot}}} \left(1 + \frac{d_{\text{eye}}}{d_{\text{rot}}} - \cos(\theta) \left(1 + \frac{d_{\text{eye}}}{d_{\text{rot}}}\right)\right)}}, \quad (1.6)$$

which assumes a view rotator of zero thickness. Without further analysis, it is not very telling, but a good approximation as shown in Fig. 1.11(c) where it is compared to the non-zero thickness case. Since it shows that the magnification appears periodic, a worst case for the magnification

can be found for a view rotation angle, θ , of 180° , which, when replaced into Eqn (1.6), yields

$$M_{a\infty} \approx \frac{1}{1 + \frac{2d_{eye}}{d_{rot}}}, \quad (1.7)$$

where d_{eye} and d_{rot} are the same eye distance and rotation plane distance described in Eqn (1.5). This is a lower bound estimate as the expected maximum range of view rotation angles, θ , is between 4° and 30° . For this range Fig. 1.11(c) clearly shows that in reality, the magnification will be better (closer to 1) than this lower bound estimate. In all mathematical derivations, an assumption was made that a unit magnification was desired when the sample plane lies at the rotation plane ($d_s = d_{rot}$). We also note that this result is counterintuitive as the expected worst case, or largest amount of, curl occurs for a rotation transformation of 90° . Interestingly, the apparent magnification term, however, does not reflect this.

From Eqn (1.5) and (1.7), it can be seen that the ratio between the pupil-to-spectacle distance and the rotation plane distance needs to be small to achieve the best results. Luckily, that regime precisely aligns with the typical use case for spectacles, namely a much shorter pupil-to-spectacle distance compared to what a user would typically look at.

Unlike before, when considering viewing positions closer to the eye (where the view distance, d_s , is small), the previous plots in Fig. 1.11 (where d_s was taken to be at an infinite distance away) no longer hold. Therefore, to quantify the impact d_s has on the apparent rotation and magnification in a plane not an infinite distance away, a plot probing the apparent rotation and magnification at different view distances was calculated. It should be noted here that this view distance will, in most cases, not be the same as the rotation plane distance, d_{rot} , used to create the rotator. In fact, when the view distance is equal to the rotation plane distance, we get the expected result, namely, the magnification and rotation set when constructing the component. The resulting plots for the apparent rotation and magnification can be seen in Fig. 1.12 (a) and (b) respectively. Furthermore, these were confirmed in numerical experiments by simulating a resolution test target at some view distances away from the component. These simulations were then compared to an upright image with unit magnification taken at the same position but without the view rotator in place. By manually rotating and scaling the simulations an estimated apparent rotation angle and magnification was found. These were also placed on the plots in Fig. 1.12 and can be seen to line up well with the curves.

Having considered both of these, it is clear that for the use in spectacles, any view distances larger than or equal to the rotation plane distance will result in an adequate apparent magnification and rotation. Now in an ideal world, this should be the end of the chapter — but it is not.

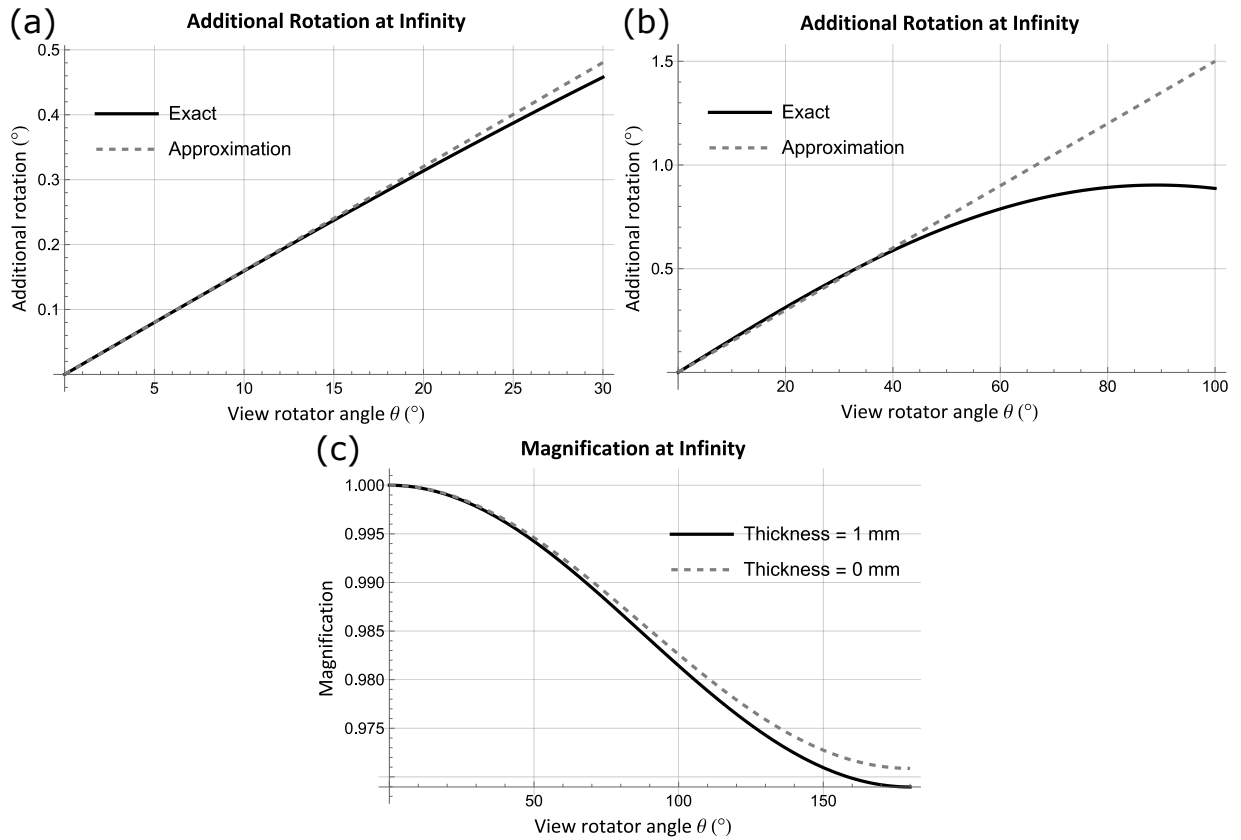


Figure 1.11: Graphs showing the additional rotation and apparent magnification at a probing plane an infinite distance away in (a, b) and (c) respectively, corresponding to the equations found in Eqns (1.5) and (1.6). In order to understand the behavior at this plane, the rotation plane distance, d_{rot} , could not be set at infinity (the perhaps most intuitive distance to use) as this would return the exact desired magnifications and rotations, not representing the issues encountered along the view. For all cases, the eye distance, d_{eye} was set to 1.5 cm while d_{rot} was set to 1 m. In (a), an approximate straight line is shown to follow the curve, showing the additional rotation. The term additional rotation is used as it refers to the apparent rotation minus the expected rotation angle. Therefore, the graphs in (a) and (b) are based on the equation shown in Eqn (1.5) but only after subtracting the desired view rotation angle. The curved solid line represents the exact solution in (a) also for zero thickness while in (b) it was plotted for a thickness of 1 mm. The plot in (b) also shows where the straight line approximation breaks down when plotting it for larger view rotation angles. Lastly, (c) demonstrates this same effect for the apparent magnification, plotting the exact curve described in Eqn (1.6) as a dashed line, it can be seen that the thickness does matter, but for the range over which the view rotator operates it remains a good approximation. For more information and derivations of the plots and equations see the Mathematica document called `ApparentRotationViewRotator` in [1].

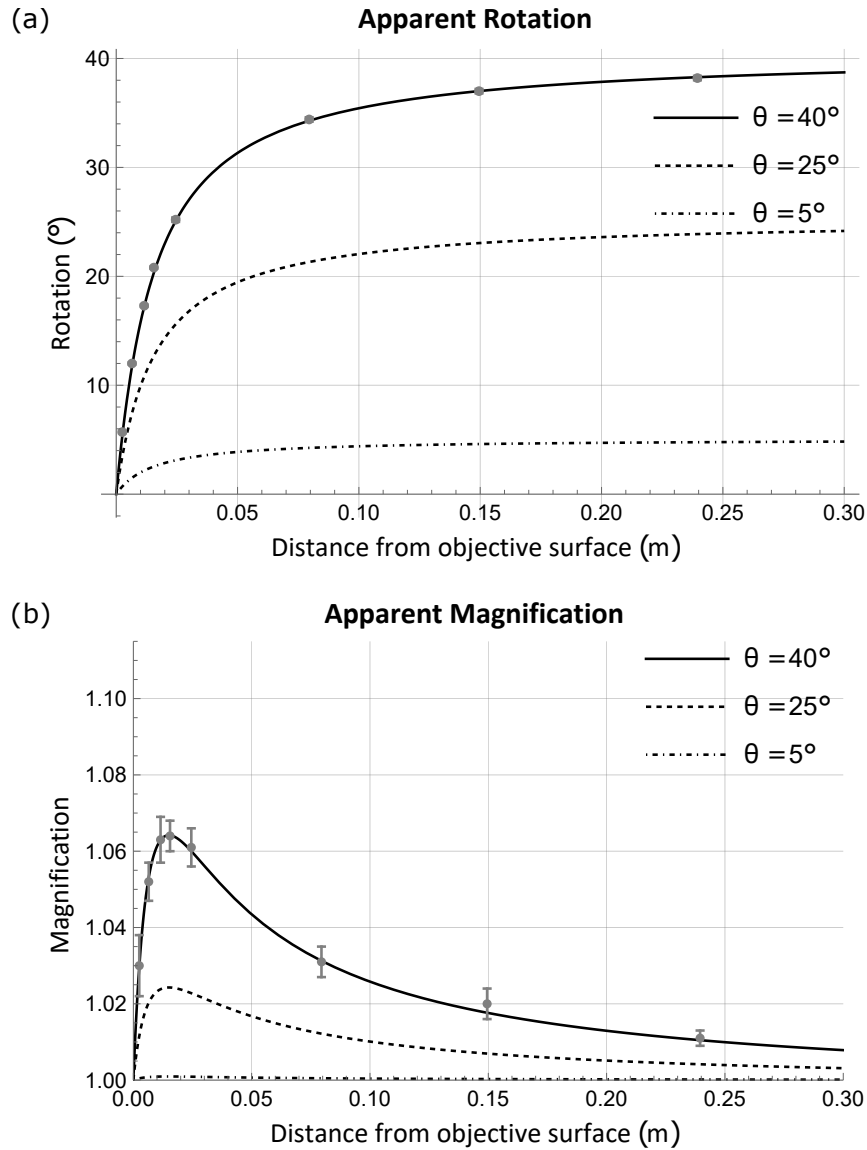


Figure 1.12: A mathematical model for the apparent rotation and magnification of the view rotator in (a) and (b) respectively, for a constant eye-to-spectacle distance (d_{eye}). The “distance from objective surface” represents the distance of the sample plane to the view rotator, specifically the objective surface of the view rotator, d_r . The amount by which the view appears rotated or magnified is shown over a 0.3 m range of d_r . It can be seen that as expected at a distance of 0 m the rotation is 0° and the magnification is 1. Beyond this point, the curves shown agree with the simulation results shown as points along the curve. These were obtained by simulating a resolution test target at the corresponding depth and manually rotating and aligning them with an upright image at the same position. The error bars are an estimation based on the observed range where the two images aligned most closely. Since the curves in both cases are not simply a constant value, the comparison to the original *view* rotator in [2] can be drawn again where this would be the case.

1.4 The problem: view rotation vs image rotation

The main issue, and also the main reason why we call it an approximate *view* rotator instead of an image rotator, becomes apparent when viewing the same component up close using a pinhole camera in Fig. 1.13. It can be seen that the individual pixels do not rotate anything but rather cause the image to be shifted to the corresponding rotated image position. In fact, even for infinitely small pixels and a construction designed with an eye position on the ocular surface, this would create a perfect *view* rotator (as described in [2]), but never an image rotator. Furthermore, due to diffractive considerations and an eye position that should not coincide with the ocular surface, our design can at best act as an approximate view rotator.

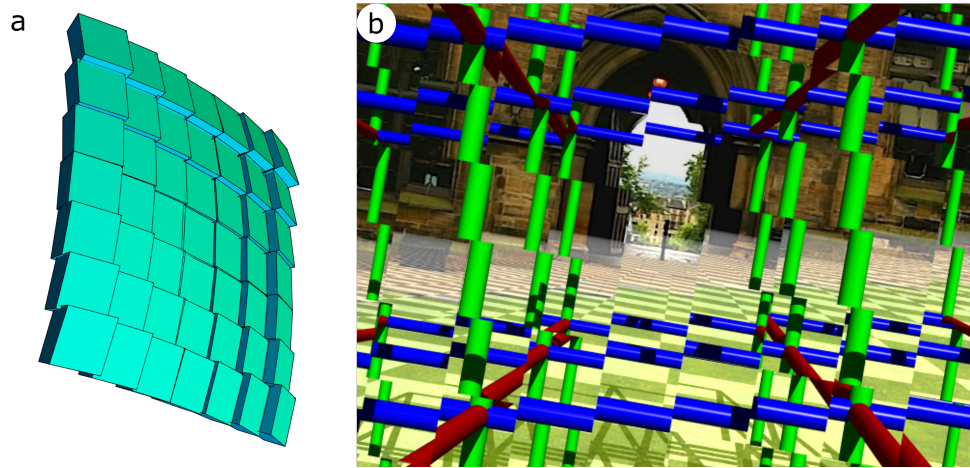


Figure 1.13: (a) demonstrates the central surface profile of the view rotator. It can be seen that it consists of individual refractive wedges, where the objective surface is in view. The view through this is then shown in (b) using a pinhole camera, a reduced FoV(30°), larger pixel size (1 mm), and ignoring any diffraction considerations. By simulating it in this manner, it can be seen that the individual pixels do not rotate the image, but instead shift it to the appropriate rotated position.

Besides diffractive issues, the fact that in reality the pupil acts as an extended aperture results in significant blurring. This is again due to the view rotator being inherently non-imaging (in fact, only certain points on the view rotator surface can be said “image” by being imaged back onto themselves), therefore, the view rotation from one point may appear to rotate differently than a nearby position. Provided that the view rotator was set up to rotate the view for a pinhole position on the pupil (approximately) correctly, it does not inherently mean that any nearby point will also be rotated correctly. In fact, it turns out that such a neighboring point will perform an approximate rotation but with a small shift in a direction that is inconsistent with the expected parallax and hence with the desired rotated position. Therefore, any light ray that does not enter (or leave from a ray casting approach) the pupil at the exact point it was designed to operate at, will inevitably be a (more or less) a “stray” light ray, with rays entering further from the pinhole resulting in worse effects to the overall view. Fundamentally, this all relates to the fact that an

image rotation is forbidden by a single refractive surface [18, 42]. Furthermore, for such stray light rays, there may also be several light ray direction changes when they pass through the view rotator, as they are no longer bound to travel through just a single pixel, as a refractive realisation hinders the pixels from being infinitely thin.

This then raises a question about the view quality of the view rotator this far. It seems to significantly improve the binocular view over the whole FoV of an eye by counteracting the view torsion, but that is only half of the story. As mentioned previously, the figures simulated over a large (130°) FoV do not take up an equally large FoV when viewed as a figure. If a smaller FoV is considered, more accurately the central Fov of an eye, the results can be seen to change significantly.

To investigate and optimize this, the simulated camera FoV was reduced to more closely resemble the human central FoV ($\sim 50^\circ$ with the most detail visible over $\sim 2^\circ$ in the center). Within this region, a Snellen chart was viewed to quantitatively determine the expected visual acuity when looking through the view rotator. The results can be seen in Fig. 1.14 and show that the expected visual acuity is roughly 6/420. The acuity was optimized by tuning the pixel size, ranging from 0.1 mm to 0.6 mm in steps of 0.05 mm, considering the trade-off between the extended aperture, diffractive effects, and the rotation quality of the view. This was done by simulating a number of images with a change in pixel size using the “movie” functionality in the simulation software called `ViewRotationExplorerWithUnits` [3]. The simulated images were then compared (all of which are available in [39], specifically the `2mm_aperture_radius` folder, which coincides with the pupil diameter of the eye), and the values associated with the image that had the apparent best visual acuity were chosen through visual inspection. However, the apparent best view is subjective here as a larger pixel size will yield a clearer view at the cost of apparent rotation. Nonetheless, we note that the method used may introduce a bias, which may potentially be improved in future by quantifying the image overlap between the ideal case Snellen chart and the Snellen chart obtained when viewed through the view rotator.

To put the visual acuity into perspective, the limit for being considered legally blind in the United Kingdom is set to 3/60 (or to align with our simulations 6/120). Therefore, the view rotator shows promise when considering the entire FoV but for any more detailed viewing, typically entirely performed by the central FoV, it shows that it can rotate the view, enabling a user to fuse the views seen by both eyes, but at the cost of very significantly blurring the view in the cyclodeviated eye. Whether or not this is preferable to alternative treatments is a question that needs to be answered in clinical testing.

1.5 Improving the view rotator

Since we found that the view rotator significantly blurs the view when considering any finer detail viewing, different ideas to improve it were considered. The main ideas are outlined and




Figure 1.14: A Snellen chart when viewed through the view rotator by an eye that has a 10° cyclodeviation simulated in (a). The same Snellen chart when viewed by a healthy eye can be seen in (b). By comparing the two charts, an estimate of the visual acuity can be made that shows the view rotator reduces the visual acuity to around 6/420.

demonstrated in this section with section 1.5.1 relating to restricting the pupil size, section 1.5.2 introduces a new concept called derivative control surfaces, and lastly section 1.5.3 briefly discusses other methods (such as spaceplates [43]) that effectively increase the distance between the eye and the view rotator. Within each of these, the improvements are shown by comparing Snellen charts.

1.5.1 Pupil restriction

The natural pupil diameter for a person is between 2 mm and 8 mm [44]. This was emulated in all the previous simulations where the camera aperture was set to a 4 mm diameter. However, from the design process illustrated in Fig. 1.2, a very simple way to improve the view rotator performance involving the pupil size can be inferred. As the rotator is designed for light rays passing through a point, simply restricting the pupil to a pinhole at that point would, in theory, produce a sharp, shifted view (where the individual pixels are unrotated). However, this will lead to additional diffraction issues. Therefore, instead of a pinhole aperture, a pupil restriction to a finite size may be more appropriate. This can be done physically through the use of a contact lens with a transmission function that only allows light within a small disc to pass or through pharmaceutical means that contract the eye's pupil. Since diffractive effects start to become significant before a pinhole size may be reached, these still had to be considered in the simulations. Therefore, pupil diffraction was added to the simulation software in addition to the pixel diffraction.

Several simulations were run to determine the ideal combination of pixel and aperture size, comparing the diffractive blurring due to both of these, until the best visual quality was achieved. It was found that the results with the largest visual acuity was visible when using a pupil restricted to ~ 1 mm in diameter along with the same 0.3 mm pixel span. This was based on a




Figure 1.15: A Snellen chart when seen through a view rotator in (a) when using a pupil restriction, which simulates a reduced pupil diameter of 1 mm, the letters 'L' and 'P' in the line corresponding to visual acuity 6/53 can now (just) be deciphered, which suggests that the view rotator now reduces the visual acuity to 6/53. There are some details visible in the 6/35 line but they are not sufficiently clear to claim a visual acuity of 6/35. In (b) the same Snellen chart is viewed with a healthy eye (without view rotator).

series of simulated images, varying the pupil size and pixel size until the largest visual acuity was observed by visual inspection of another image series taken with both, varying pixel size, and aperture size available in [39]. We note that this approach will lead to an approximate value, and that quantitative image comparison methods may yield a more exact result. However, for the purposes of this chapter, the returning benefits appear to be relatively low in comparison to the visual acuities that are obtained. As such, the improvements can be seen in Fig. 1.15, where a Snellen chart is once again viewed. It can be seen that the new visual acuity achieved is significantly improved to about 6/52 with some visible details in the 6/35 line. By only considering the visual acuity, it can therefore be said to no longer be legally blinding. To simulate the diffractive effects within all the simulations used in this chapter, the pixel was simulated as a square slit. A probability distribution based on single slit diffraction was then used to estimate the likelihood of a ray passing to one of the diffractive orders. By choosing a weighted random number, based on the probability distribution, a ray may then change direction accordingly. As expected, a smaller slit width yields a distribution with a lower central peak and hence larger diffractive effects. The calculations for this were performed considering a single wavelength of 632.8 nm.

Another prominent effect visible in the simulations is the orientation of the letters. Unlike in Fig. 1.14, the letters no longer appear to be fully upright and corrected, for the same pixel size of 0.3 mm. This is related to the explanation as to why this device is an approximate view rotator and not a perfect view rotator in section 1.4 and can also be seen in Fig. 1.13. Effectively, by reducing the pupil size the view becomes less blurred, and as a result, each individual pixel can

be seen more clearly. Since a single pixel does not rotate the view seen through it but instead shifts it, the pixels will retain a rotated view. In the binocular view of the same Snellen chart in Fig. 1.16 this can then be seen as a haze, but importantly without deteriorating the binocular view quality significantly when compared to an equivalent binocular image without the view rotator.

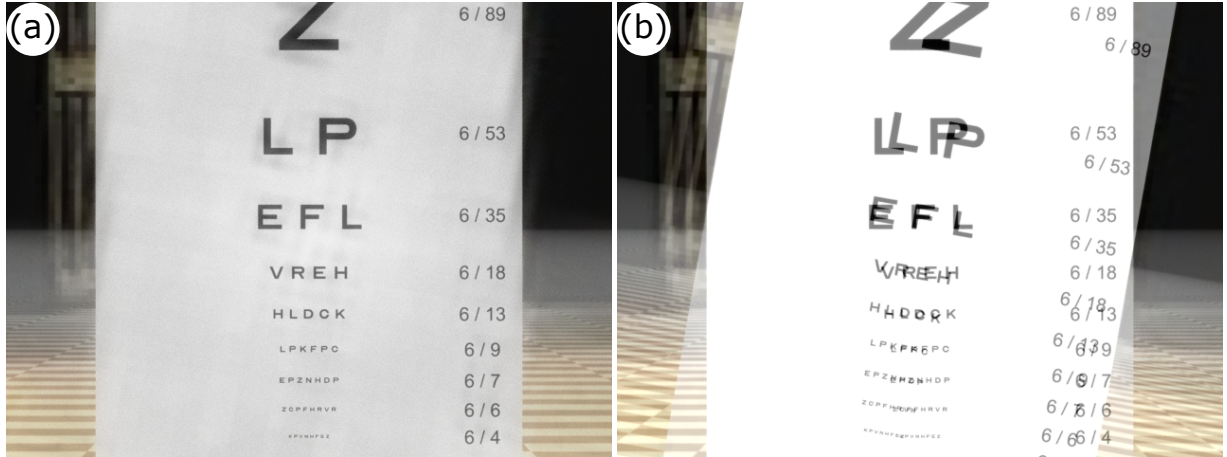


Figure 1.16: (a) shows a binocular image comprising of Fig. 1.15 (a) and (b) in equal parts (where each images pixel values are halved and then added together). It can be seen that the view appears clear, with a slight blur behind each letter. This is due to the corrected eye. In the same way, (b) shows the binocular image but now without any view rotator in place demonstrating the achieved improvement.

By restricting the pupil size, an additional dimming effect can be expected as less light is permitted to enter the eye. In proportion to the fractional area of the restricted pupil to the normal, unrestricted, pupil. Mathematically this can be expressed as

$$f_p = \frac{\text{restricted pupil area}}{\text{pupil area}} \approx \frac{\pi r_{rp}^2}{\pi r_p^2} = \left(\frac{r_{rp}}{r_p} \right)^2, \quad (1.8)$$

where f_p is the fractional area of the restricted pupil, r_{rp} is the restricted pupil radius, and r_p is the unrestricted pupil radius. The fraction then came to be 0.0625, when considering $r_p = 2$ mm and a $r_{rp} = 0.5$ mm. This would imply that the resulting view will be roughly 94% dimmer than the view without a pupil restriction. However, it is also true that the eye and brain are well equipped to deal with brighter or dimmer lighting conditions. Therefore it is not fully clear how much of an impact this will have on the overall view quality.

It is clear that this will lose some information from any light rays that would otherwise have entered into the pupil. However, in this particular case, rays further from the pinhole are typically refracted wrongly and therefore do not contribute significantly to the corrected view. This could mean that filtering out such rays is in fact an improvement to the binocular image quality as a whole.

1.5.2 Derivative control

In contrast to section 1.5.1, another way in which the visual acuity may be improved is by changing the view rotator itself to correct for an extended viewing position. For this, a new idea was devised in conjunction with a Masters project carried out by Ewan Eaglesham. We named this idea “derivative control surfaces”.

What is a derivative control surface?

To understand how these surfaces may aid in improving the view rotator we first describe how these work in a more general sense. A derivative control surface, as the name suggests, is a surface added to some component to further control certain properties of the component. More specifically, we want to retain the telescopic nature of an original transformation and only consider directional derivatives as opposed to positional derivatives. Therefore, we wish to control what we call first order light rays, where rays incident on different positions also have a slightly different direction that we can manipulate. One way to design some free form optical components is by considering an incoming light-ray bundle and deciding what the outgoing light ray bundle should look like. This very basic idea lays the ground work for some known optical design procedures such as the Simultaneous Multiple Surface (SMS) Optical Design method [45]. In fact, the SMS method allows for up to two light-ray bundles to be redirected by two optical surfaces (one front and one back surface of a component). We call the light rays which are used in the construction and are therefore redirected by the component perfectly, zeroth order rays.

Any light ray which is closely related to these zeroth order rays, that is in both direction and position, we call first order rays. By designing a surface with these in mind, it may be possible to extend the working range of a general free form optical component. When considering a single zeroth order light ray incident at some point on the surface with a direction D_{i0} , a small change in the originating position or direction may be added. Considering a first order ray incident at the same position with direction D_i , we can relate the outgoing first order ray direction, D_o , as

$$D_o = D_{o0} + J(D_i - D_{i0}), \quad (1.9)$$

where D_{o0} is the outgoing direction of the respective zeroth order light ray, and J is the Jacobian of the surface. Since the change in the direction between the zeroth and first order ray has to be very small, the Jacobian can be expressed as $JdD_i = dD_o$. In principle, this could also be extended to control higher-order derivatives (i.e. using the Hessian), but first, we attempt to make it work for the Jacobian. In many ways, this may be thought of as some type of refractive Taylor series, whereby the zeroth order base term corresponds to a zeroth order light ray direction that propagates correctly, but any higher order terms do not necessarily follow. Since the higher order terms are derivatives of the first term, we use the term derivative control surface. Also following on from the Taylor series idea, the Jacobian of a derivative control surface is really

what should be controlled. In the end, if the Jacobian of an additional surface can be controlled, a derivative control surface is made. What Jacobians are permitted is not well defined. Simple examples, such as a Jacobian of a telescope, clearly work, but more complex examples may not. Adding this derivative control surface to a free form component then allows for finer control of first order rays.

At this point, the required Jacobian needs to be found. This will depend on the desired transformation of the overall component and can therefore range from being very simple to impossibly complex. Currently, we do not have a method to determine the Jacobian required for a general transformation besides trial and error simulations along with estimated guessing. Extending this is one of the potential future research directions beyond the scope of this chapter.

Derivative control in context of view rotation

Molding the idea of derivative control to the view rotator, it was found, by Ewan Eaglesham, that there are two similar, but importantly different first order rays which we wish to control. As a result, there are also two Jacobians that need to be considered.

The first is related to an extended aperture size, which is why we call it aperture correction. It should be noted that in this subsection, we are no longer restricting ourselves to a setup that mimics an eye. To reflect this, *aperture* and *camera* will be used instead of *pupil* and *eye*. In this scenario, the first order rays correspond to rays which are emitted (in our ray-casting model) from any part on the aperture and incident on the same point of the view rotator. In order to give a sharp rotated view, the Jacobian needs to be a simple rotation matrix, here called J_a . The rotation angle of the Jacobian should be equivalent to the view rotator angle, and can therefore be expressed as

$$J_a = \begin{bmatrix} \cos \theta_a & -\sin \theta_a & 0 \\ \sin \theta_a & \cos \theta_a & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (1.10)$$

where θ_a is the rotation angle of the overall view rotator. This can, and has, been proven mathematically by Ewan Eaglesham in [46]. Independently, this can be confirmed through ray tracing simulations.

To do this, a simplified and heavily idealized view rotator was constructed with the property to redirect any incident ray to the rotated image position (akin to the view rotator which has been used throughout) while being infinitely thin and having a point-like pixel size, ignoring diffractive effects. On top of this, another infinitely thin surface was placed. This surface was the derivative control surface for which the Jacobian can be controlled. More detailed simulation parameters can be seen in the caption of Fig. 1.17.

The simulations show (Fig. 1.17) that when there is no derivative control, the image remains blurred regardless of the camera-focusing distance. When the derivative control surface is then




Figure 1.17: Improvement due to derivative control. The view seen through an idealized view rotator, (a), here shown as a phase hologram where each point redirects a light ray to the appropriate rotated view position, akin to an idealized version of the view rotator design presented in this chapter, when viewed using an extended aperture. This then represents an unphysical view rotator that is planar with infinitely small pixel size that would, in reality be subject to diffractive effects, which are not considered here. The camera focusing distance was set to 10, which coincides with a plane containing the front of the lattice structure. Unlike the simulations shown before this, the view rotator is no longer set up close to the camera (or eye), resulting in a larger camera distance and aperture size. Within this sub-subsection, all simulations are performed with a different simulator, namely the `DerivativeControlSurfaceExplorer`, also available in [3]. In (b) the same scene is viewed again but now with a derivative control surface added to the view rotation surface. The Jacobian of the derivative control surface was chosen to follow the aperture correction, given in Eqn (1.10), with θ_a being set to 20° , which can be shown to correct for an extended aperture by bringing the front of the lattice structure into focus. This was not possible in (a) regardless of the camera focusing distance. The simulation was set up for a 20° rotation, an eye distance of 3, and a rotation plane distance of 10. The lattice was placed such that the front of the lattice coincided with the rotation plane at a distance of 10 from the rotator.

activated to mimic the Jacobian shown in Eqn (1.10), the view seen becomes clearly focused onto the front of the lattice, which corresponds to the chosen focusing distance of the camera.

The second derivative control refers to rays which originate from some pinhole aperture (from a raycasting point of view), and are incident on the same pixel, but not at the same position on the pixel. In this consideration, the pixels are no longer considered point-like, as was the case in the aperture correction, but rather of a well-defined size. The aim of derivative control in this case is to make the image/view seen through the single pixel be rotated by the same angle as the overall image/view. Therefore it was also called pixel correction. By considering first order rays of this type, a new surface Jacobian, J_p , was found to perform the derivative control. This Jacobian was once again a rotation matrix but with an added magnification. Mathematically this can be expressed as

$$J_p = \begin{bmatrix} M \cos \theta_p & -M \sin \theta_p & 0 \\ M \sin \theta_p & M \cos \theta_p & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (1.11)$$

where θ_p is the rotation angle, and M is a magnification factor. Now unlike the case shown in Eqn (1.10), the rotation angle here is no longer simply the rotation angle of the view rotator. Instead, the angle follows

$$v = \tan^{-1} \left(\frac{\sin \theta_p}{z_2/z_1 + \cos \theta_p} \right), \quad (1.12)$$

an equation found in [31] as equation (4), where v is the apparent rotation at a distance z_1 past the rotator along the z -axis, z_2 is the distance between the camera and the rotator and θ_p is the rotator rotation angle. Now, θ_p may be obtained such that the apparent rotation angle, v , is that of the view rotator at the rotation plane a distance z_1 away.

From the same paper [31] equation (5), the magnification, M , can be obtained as

$$|M| = \left| \frac{1 + z_2/z_1}{1 + e^{i\theta_p} z_2/z_1} \right|, \quad (1.13)$$

where the symbols take the same meaning as in Eqn (1.12). The combination of these acts as an equivalent of taking a ray rotation sheet that has been modified to maintain a unit magnification. This can again be shown using ray tracing simulations.

In this case, the simulations were created using a pinhole aperture, without any diffractive effects, viewing an idealized view rotator of finite pixel size. Within each pixel, a derivative control surface was added to induce the Jacobian shown in Eqn (1.11) with the rotation angle and magnification derived from Eqn (1.12) and Eqn (1.13) respectively. The results shown in Fig. 1.18 demonstrate this for no derivative control in (a), derivative control but with the Jacobian shown in Eqn (1.10) in (b), and lastly the appropriate Jacobian Eqn (1.11) in (c). It can be seen that in the latter-most case, the view rotation appears to work perfectly.

However, the simple fact that the two Jacobians used are not the same means that there cannot be a single Jacobian that will correct for an extended aperture size and finite pixel sizes. In fact, it is even worse than this when considering that in the case of an extended aperture size and finite pixels, there will be rays that originate from anywhere on the aperture and are incident on different parts of a single pixel in which case these are rays which are not fully controlled by either pixel corrections or aperture corrections. This can be seen in Fig. 1.19 where another idealized view rotator is simulated using finite pixel sizes and viewed with an extended aperture. When applying either of the derivative controls outlined the view remains blurred. However, there does seem to be a significant improvement in the view quality regardless, and therefore, it may still prove to be useful when attempting to increase the visual acuity of the view seen through the view rotator.




Figure 1.18: The view seen through an idealized pixellated phase hologram of a view rotator is shown in (a) for a 20° rotation. Unlike the simulations shown before this, the view rotator is no longer set up to represent an eye. Here this manifests as a pinhole aperture and an increased camera distance. In (b) a derivative control surface was added to the view rotator with a Jacobian shown in Eqn (1.10). This is the Jacobian for an aperture correction, which as expected does not correct the pixelation perfectly, but is an improvement compared to (a). Lastly, the pixel correction Jacobian Eqn (1.11) is employed in (c) where it can now be seen to match neighboring pixels together creating a seamless view. In (c), the Jacobian rotation angle was set to 25.89° , while the magnification factor was set to 1.0183. We note that ignoring diffractive effects and dealing with pinhole cameras and planar holograms is clearly not physically feasible, but it instead serves as a proof of concept of the required Jacobian.

Implementing derivative control to the view rotator

Implementing a suitable derivative control in the actual view rotator requires some further considerations. Initially, the type of correction that seemed most appropriate for the physical view rotator had to be chosen. Since the final goal of the view rotator is its use in spectacles, considerations towards the binocular image have to be taken. Therefore, any transformation other than a simple rotation cannot be used to avoid further deterioration of the binocular image. As a result, the pixel correction, whose Jacobian is given in Eqn (1.11), cannot be used as it requires a non-unit magnification at the rotation plane. This stems from the expression of the ray-rotation sheet apparent rotation not matching that of the approximate view rotator apparent rotation outlined in section 1.3.3. An aperture correction (or in the case of an eye, a pupil correction) appears to




Figure 1.19: A demonstration of the two correction types for an extended aperture viewing a pixellated view rotator. This combines the scenarios outlined in Figs 1.17 and 1.18, where in (a) a pupil correction is shown. In (b) a pixel correction is used instead. In either case, the view seen is improved compared to Fig. 1.18(a) but still blurred. Ideally, pupil correction and pixel correction would be applied simultaneously, but as they are subtly different this is not possible, resulting in a view that is still slightly blurred.

be the most appropriate (and only other) choice. The required Jacobian is then a simple rotation matrix about the view axis with a rotation angle equivalent to the view rotator rotation angle.

A refractive realization was considered to feasibly suggest such an improvement. Besides prismatic approaches [33], we consider a simple configuration of refractive cylindrical-lens telescopes to induce the desired rotation. A pair of cylindrical-lens telescopes rotates the view seen through them by 2θ when one is rotated relative to the other by θ . This occurs as each telescope flips the view seen through it across the cylinder axis, as shown in Fig. 1.20, and two of these then rotate the overall view seen.

Placing two cylindrical-lens telescopes within each pixel, roughly shown in Fig. 1.21(a), then rotates the image locally. Since the pixel size is small, the clear aperture each telescope has to fill is also small, allowing for small cylindrical-lens focal lengths and therefore a relatively thin (~ 1 cm) improved view rotator, which is of course desirable for spectacles. A refractive version of this was programmed and subsequently simulated, shown in Fig. 1.21(b).

Unexpectedly, it shows that the achieved visual acuity is much lower than expected at $\sim 6/89$. Furthermore, the view is not only blurred but very diffused. This may be due to the use of a telescopic construction, which leads to a reduced field of view, which means an extended aperture will have various stray light rays. Furthermore, this also restricts the permitted eye movement further. The form factor of a refractive realization of this requires an overall thicker (and bulkier) construction than the simple view rotator, with an estimated thickness in the region of 1 – 2 cm.

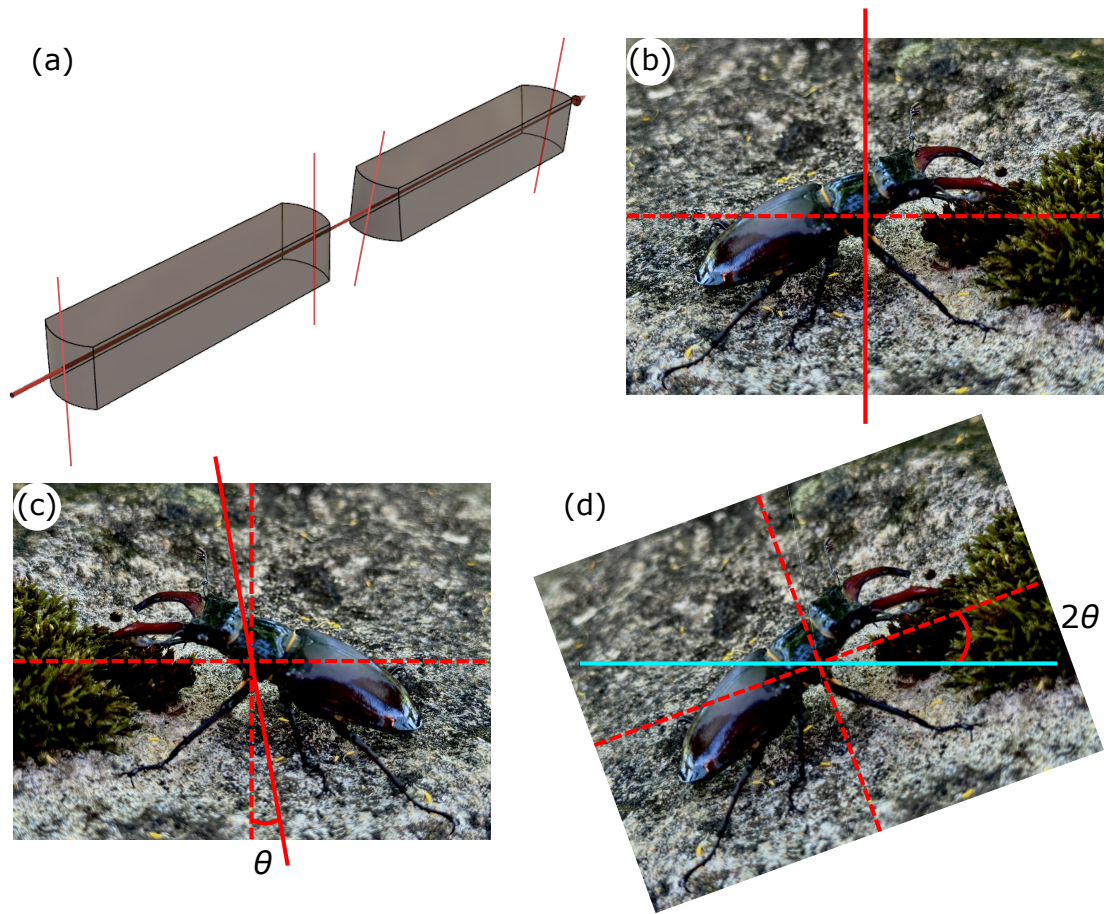


Figure 1.20: A simple way to induce an image rotation using two cylindrical lens telescopes. Here the end surfaces of each block of refractive material act as telescopic cylindrical-lens pairs, with the cylindrical lens nodal line shown by a red solid line. By configuring them as shown in (a), with the second telescope rotated relative to the first one along the optical axis shown as a red arrow through the centers, the image gets rotated. This is done through a series of image flips, each due to a single cylindrical telescope. Following (b) through to (d) this is shown. First, the image of a beetle is flipped about the vertical axis shown as a solid red line in (b). The flipped image can be seen in (c) along with the next “flipping axis”, rotated by θ with respect to the previous flipping axis, again shown as a red arrow. This axis lies along the cylindrical lens cylinder axis. When the image is then again flipped about this axis, the resulting image can be seen in (d) where it can also be seen that it effectively has been rotated by an angle of 2θ above the horizontal (and initial image) orientation.

Discussion

The final result for the full implementation is clearly not as good as expected based on the idealized preliminary simulations. The main reason for this is cross-talk between telescopes, also known as non-standard refraction. In the idealized case, the Jacobian was introduced through a single surface within the same plane as the rotation occurs. Therefore, any ray is redirected accordingly. However, this is not the case in the refractive case. For an extended pupil, most rays will not remain within a single telescope, instead, they may enter the pixel at one position




Figure 1.21: A sample of the nine central most pixels of a potential derivative controlled view rotator, shown in (a). Each pixel consists of two cylindrical lens telescopes set up to rotate the view seen through them. It should be noted that this was made using `ViewRotatorRefractiveDerivativeControl`, available in [1], with the intention of creating an accurate 3D model of the rotator. In (b) a simulation of the resulting view through such a component can be seen. It is diffuse and blurred, but the top line may barely be interpreted as a “Z” which would give it a visual acuity of roughly $6/89$. While the second line is also slightly shown, it is not sufficient to accredit it with the respective visual acuity.

and exit a completely different pixel altogether. This causes any number of issues ultimately manifesting as a diffuse and blurred image seen.

A way in which this could be improved is to use the previously shown method (section 1.5.1), where a restricted pupil will encourage fewer stray rays that jump between pixels. This seems a bit counterproductive as the main idea was to correct for an extended pupil size. However, if instead we adjust the rotation angle slightly to match the value for the pixel rotation and ignore the magnification term, a significantly improved view can be achieved. This can be seen in Fig. 1.22 which essentially shows the combined effects of a (unit magnification) pixel correction and a pupil restriction as used in section 1.5.1.

Besides the shortcomings mentioned in section 1.5.1 some other issues arise when investigating the off-axis view, as was previously done in Fig. 1.5. However, here the effects of looking away from the rotation axis become worse when adding the additional telescopes. The main culprit is once again non-standard refraction, which dominates the view seen.

A simulation is shown in Fig. 1.23, comparing the off-axis views. Ultimately this shows that every improvement thus far has a number of drawbacks associated with them. A further discussion on this and the general feasibility of the implementations is carried out in section 1.6.

1.5.3 Spaceplates

The last improvement relates to the use of spaceplates [43]. In the simplest terms, these are relatively thin optical components that act like a much thicker slice of space, and therefore, they




Figure 1.22: The view seen through the derivative controlled view rotator with a restricted pupil size. The derivative control chosen was a pixel correction with a unit magnification. It can be seen that the letters not only appear rotated but also in the appropriate location. This culminates in a visual acuity of $\sim 6/35$, shown in the third row of the Snellen chart. To achieve the pixel correction a Jacobian with a rotation angle of 10.15° was chosen. In practice, this was simulated using two cylindrical lens telescopes rotated by 5.075° relative to one another as described in Fig. 1.20.




Figure 1.23: The same off-axis test as carried out in Fig. 1.5 was done but now with the derivative control. In (a) the desired off-axis view can be seen. In (b) the view seen through the derivative controlled view rotator is shown using a restricted pupil size. This is the same view rotator used to create the view seen in Fig. 1.20. It can be seen that the view is littered with double images and severe blurring. This is mainly due to crosstalk between telescopes.

can increase the path length and ray offsets of a light ray beyond the expected amount for the size of the component. In our case, this would essentially mean the eye can be optically further away from the view rotator without changing the physical distance. As a result, the angular size of the pupil when seen from the view rotator will be smaller and more closely resemble a pinhole aperture, increasing the visual acuity.

Overall, the increase in visual acuity strongly correlates with the additional optical space the spaceplate adds. In Fig. 1.24 this can be seen, with the mentioned increase in visual acuity.

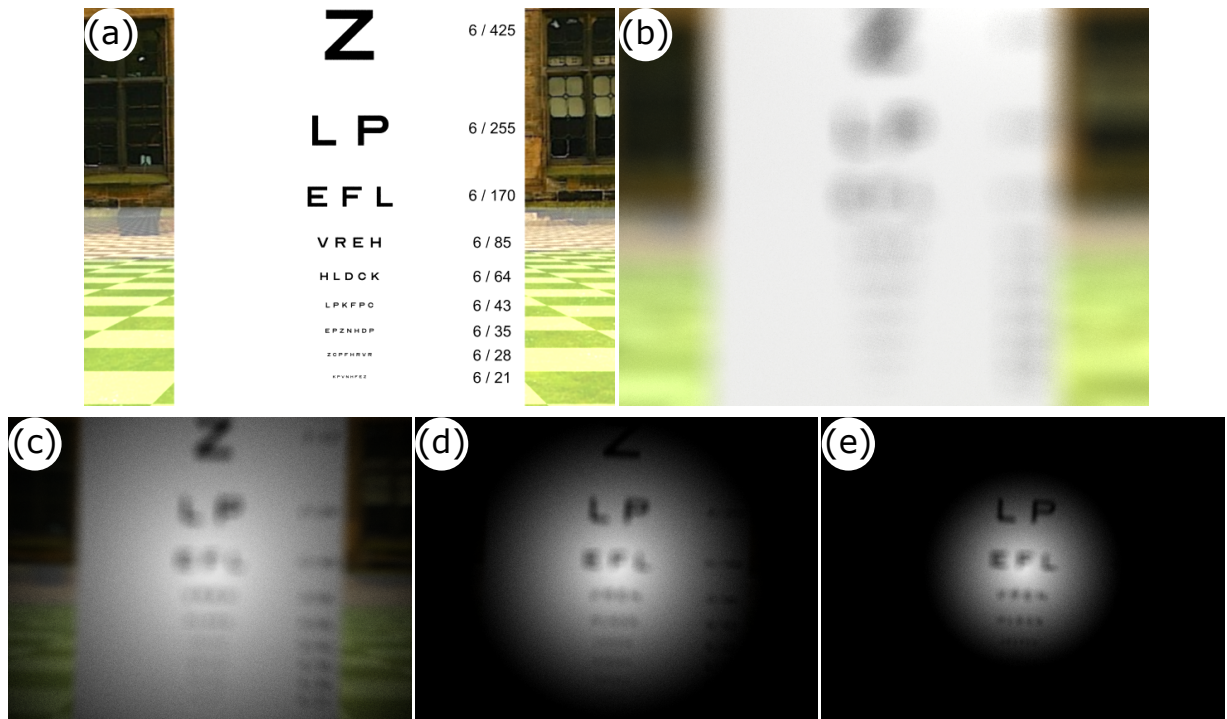


Figure 1.24: The view seen through the view rotator when a heavily idealized spaceplate is added. At first, a Snellen chart is viewed without any component or torsion in (a). This can be viewed as a control to understand the achieved visual acuities. In (b), the view rotator and spaceplate are added but with the spaceplate adding no additional distance. From (c) through to (e), the spaceplate was used to add an additional 1 cm, 2 cm, and 3 cm respectively. It can be seen that the visual acuity is improved with each bit of space added, but the FoV is also decreased as a larger fraction exits the spaceplate through a side, shown as a black region in the simulations. It should be noted that the spaceplate used in the simulation was heavily idealized and infinitely thin. The spaceplate used throughout, consisted of a single planar surface, which effectively shifted the light rays incident on it to the appropriate position along the same plane, coinciding with the expected transverse position after a ray had traveled the specified extra distance while retaining the initial ray direction.

There is a black region that forms around the edges of the spaceplate, which grows with increasing spaceplate strength. These regions correspond to rays that exit the sides of the spaceplate, assuming a perfectly cylindrical spaceplate setup. Of course, in reality, these may be free to travel but for simulation purposes, it was decided to absorb them and color the region black. By looking at the absorbed region it becomes clear that any increased spaceplate strength will as a drawback cause a smaller FoV. Therefore, an optimal strength would need to be found which will strongly depend on the preferences of a user.

Another important note to add is that the spaceplates simulated are idealizations of the currently available ones. In reality, spaceplates are currently limited to one or few wavelengths making them less attractive for the use in spectacles. It has been shown that the maximum additional space a spaceplate can add is heavily reliant on the complexity, often interlinked with the

number of layers, and the bandwidth [47]. To achieve the largest shift shown in Fig. 1.24, either a large compression ratio would be required or a thick spaceplate would be needed [43]. Both of these are not currently desirable for cheap wearable optics where a broad bandwidth is required.

1.6 Discussion and conclusion

Initially, it was shown that a simple view rotator can be constructed using an array of wedges, each of which we called a pixel, to redirect the view through them to the appropriate rotated image position. This rotator appeared to work very well over a large FoV, equivalent to that of the human eye when viewed in a simulation that takes up a much smaller field of view, effectively mitigating some of the (large) imperfections. When considering the more detailed central FoV, it became apparent that the view rotator results in a blurred view. The extent of this blur was tested using a Snellen chart and it was found that in the simplest case, it degrades the visual acuity well into the legally blind regime of 6/120.

At this point, it is unclear if this blurring will result in a binocular view that is worse than the alternative of ocular occlusion, by for example placing an eye patch over the afflicted eye. All of this can only be accurately tested when performing patient studies. To aid in this direction, the simulation software was rewritten using three.js [48], a JavaScript language that allows interfacing with the graphics processing unit (GPU). This was primarily done by Dr Courtial. By using a GPU-based approach, the view through the view rotator can be simulated in real-time, and even more convincingly using a virtual reality (VR) headset. Testing the view seen when using the VR headset, even the simple view rotator appears to improve the binocular view quality significantly. However, without any patient studies, these preliminary results must be taken with a healthy degree of caution. It was also not inherently clear how the depth perception was affected. Therefore, it may still simply be the case that the simple view rotator blurs out the view to an extent where it is as effective as an eyepatch. Answering this is a large part of our future research interests since a simpler component holds several advantages from manufacturability to cost and rigidity.

Since there is no certainty regarding the outcome of such tests, the alternative approach of attempting to improve the visual acuity of the view rotator was also taken. Here it was found that a simple improvement that can be taken is a physical pupil restriction, thus maintaining most of the advantages of keeping a simple view rotator while significantly increasing the view quality to a visual acuity of $\sim 6/53$.

Another improvement was our new derivative control approach. This involved introducing a surface in front of the view rotator, with full control over the Jacobian of the surface, and hence full control of so-called first-order rays. This proved very promising, however in order to realize the required Jacobians, namely predominantly two rotation matrices, a single refractive surface is not sufficient. Instead, each pixel was constructed of two cylindrical lens telescopes,

rotated relative to one another along the central view axis. As a result, a new problem came to light —cross-talk between neighboring telescopes— which manifested in the overall view as another source of blurring. To minimize this, a combination of the previously discussed pupil restriction and the derivative control was tested. This aimed to reduce the number of stray light rays, the underlying cause for cross-talk between telescopes. As a result, the best visual acuity was achieved at $\sim 6/35$.

Similar to the previous improvements this result compromises the feasibility of the view rotator in another way by resulting in significant off-axis blurring, again introduced due to cross-talk, but due to a different kind of stray ray. This time, the stray light rays originate due to the movement of the pupil as the eyeball rotates. As a result, even the center for which the whole view rotator was designed is no longer in the expected position.

This leads to more questions that need to be answered through trials; is it more comfortable to have a better view but a requirement to turn the head when looking at various parts of a scene? Can this work despite the natural eye movement when focusing on a scene? Similarly, it could be a case of adding the derivative control only to the central pixels, to allow for a detailed view directly in front of the user and the view without any derivative control in the peripheral vision. As a result, when viewing off-axis scenes, the view may be slightly worse but without any interference from the telescopes.

The last improvement outlined was the use of a spaceplate. To simulate this, a heavily idealized spaceplate was used as described in section 1.5.3. Despite this idealization, the best reasonable acuity that could be achieved was $\sim 6/170$, which despite being a significant improvement to the simple view rotator is significantly worse than the two methods outlined previously. Furthermore, the limited FoV due to rays that enter the spaceplate and leave it through a side, in the simulations of Fig. 1.24(c) to (e) shown as a black region, makes this method the least likely to be successfully implemented.

Overall, it was shown how a simple view rotator could be constructed and how it can function well over a large FoV, however, the more detailed central view cannot resolve any finer detail. This was visible when viewing a Snellen chart and concluding that the visual acuity was within the legally blind range for the United Kingdom. Without any improvements, this may still prove to be useful. Should this not be the case, several methods to improve the view rotator were outlined, lowering the visual acuity to at best $\sim 6/32$. It is not inherently clear if this value may be significantly improved further. One possibility to determine this is by calculating a best-case visual acuity, which may be possible by considering the curl of the vector field upon passing through a segmented rotating plane. In the future, we hope to study this along with conducting clinical trials to test the simple and improved view rotator as a treatment alternative or aid for torsional diplopia.

Chapter 2

Adaptive spectacles

2.1 Introduction

Globally, the most common visual impairments are refractive errors, specifically myopia and hyperopia (near and farsightedness), which make up the majority of the cases in the working population [49, 50]. The global economy loses an estimated \$269 billion annually due to uncorrected refractive errors. Moreover, this number predominantly affects the low to middle-income countries [51] where the lost economic output may be even more crucial. The idea to correct refractive errors is, of course, nothing novel, with the earliest spectacles being proposed back in the 13th century [52], where simple lenses were placed on a frame and worn. Over centuries of development, these spectacles were improved to the commonly known glasses we have today. In this time, specifically, the basic lens saw significant changes in material properties and manufacturing, leading to cheaper lenses and further solutions to correct refractive errors, such as contact lenses. Other routes innovated the overall structure of a lens with the invention of a Fresnel lens, first proposed by Count Buffon in 1748 and then independently re-invented by, and named after, Augustin-Jean Fresnel in 1822 [53].

The Fresnel lens was initially proposed for use in lighthouses to reduce the size and bulkiness of the relatively short focal length lenses used [54], but quickly gained popularity in other areas where size constraints were required. The light ray direction change of a thin refractive component can be directly given by the surface thickness gradient; in the case of a sufficiently thin lens, this means light rays get redirected purely according to the surface profile and independent of the space in between the surfaces. Removing this space entirely would then, in principle, create an infinitely thin lens, but this is not physical. Instead, several steps (which we will call “Fresnel facets”) where the thickness jumps are introduced to permit the appropriate surface profile, ultimately leading to a thinner and physically realizable lens. The exact thickness may be determined by the number of steps used, with a larger step count leading to a thinner Fresnel lens. An example of these Fresnel facets can be found in Fig. 2.1(b).

The advancements in lens technology and accessibility would therefore suggest that uncor-

rected refractive errors should be largely eradicated from current society, but that is not the case. The main reason for this is the cost associated with obtaining spectacles in the developing world. This can be split up primarily into two parts: the material and production part, and the human resources part [51]. Due to a lack of relevant healthcare workers in many developing regions, patients lack access to determine the required spectacle prescription strength. Training and maintenance of personnel and healthcare facilities to accommodate this was estimated to cost up to \$28 billion [55]. This boundary to accessible healthcare is one that should by no means be neglected, but it may also be beneficial to seek further development in lens design to offer alternative approaches to the issue. This was the intention of the proposed *adaptive* spectacles in [56]. As the name suggests, these are spectacles whose focal length may be adjusted by the user to suit their vision without the specialist knowledge required to determine their spectacle prescription strength, therefore offloading the burden of the lack in infrastructure onto the user.




Figure 2.1: A simple lens is shown in (a) with parallel red light rays incident on the left-hand side coming to a focus on the right-hand side after transmission. This lens was then turned into a Fresnel lens in (b) by removing some of the central thickness of the lens, leading to steps in the surface profile. Each part of the Fresnel lens separated by a step will be referred to as a “Fresnel facet”. The same light rays were incident and can be seen to focus down to a single point once again. A further simplification may be made where the surface profile is approximated not by curved surfaces but by straight lines, turning the Fresnel facets into approximate “Fresnel wedges”. This is shown in (c) with the same light rays incident once again. However, since the Fresnel wedges are no longer adding any focusing power to the system, the light rays are redirected into a “focal region” instead of a focal point. If the wedges are sufficiently small, this region can be made small to approximate a Fresnel lens more accurately.

There have been several different underlying techniques to achieve adaptive spectacles. One way to realize an adaptive lens is by using liquid lenses, which utilize liquid pressure to form two elastic transparent membranes into a lens by filling the space between them with a liquid [57]. Liquid-crystal lenses where the focal length may be controlled electrically also exist [58], along with classical refractive approaches such as a commercialized version of the Alvarez-Lohmann lens [59–62]. The Alvarez-Lohmann lens design works by having two complementary third-order-polynomial surfaces that can be shifted relative to one another to create the equivalent

phase profile of a classical (second-order-polynomial) lens. Diffractive optical elements have previously been shown to create a lens of adjustable focal length in [63]. These consist of two diffractive components, namely phase plates, that, when rotated with respect to one another, change their combined focal power. As it turns out, this approach is in many ways very similar to what we will show here. First, in the basic diffractive adjustable lens, a discontinuity exists within each phase plate, which upon rotation results in a section of the resulting lens having a different focusing power. As will become apparent, this section also exists in our approach, but wound into a spiral shape, where the fraction of the view taken up by the region is also directly correlated with the relative rotation angle. We also note that upon wave optical analysis (which is outside the scope of this thesis, but subject of a paper in preparation attached in appendix C), a limit exists where our lens approach yields the exact same diffractive lens as presented in [63]. However, generally our approach aims to create this component purely refractively as we are using a ray optical approach. This has the advantage that it should in theory work for a larger range of wavelengths, barring any chromatic aberration effects.

Adhering to the idea of using relatively shifted refractive components, the goal was to add a further candidate to the existing pool of adaptive lenses, specifically aimed to meet the requirements of aiding in the treatment of uncorrected refractive errors in the developing world. In much the same way as the Alvarez-Lohmann lens, shifting two second-order-polynomial surfaces with respect to one another creates the equivalent phase profile of a simple wedge. Using this under the approximation that the Fresnel facets are essentially just a series of approximate wedges, an adaptive version of a Fresnel lens can be made. The constraints imposed for this novel adaptive lens were that the final design should be relatively thin and lightweight, easy to use, conform to classic spectacle dimensions, work over a correction range of -6 D to 6 D, and be sufficiently simple to presumably allow cheap bulk production. The resulting adaptive lens is described in [64] and consists of two cylindrical lenses wound into logarithmic spirals, where a summary of the basic construction is given in section 2.2. Improving the performance using “winding focusing” was then the goal in sections 2.3 and 2.4. Extending the component further to include an Archimedean and hyperbolic spiral was the subject of sections 2.5.1 and 2.5.2, respectively, while a feasibility check to determine if our initial constraints could be met was carried out in section 2.6.

2.2 Basic construction

The underlying principle of the adaptive Fresnel spectacles was primarily developed in [64]. When considering classical refractive materials, the thickness profile of the material directly correlates to the imposed phase for any passing light ray. This is the reason why two complementary plano-lenses –lenses that have one flat side with no focusing power (the plano part of the lens) and the sides that face each other have the same focusing power magnitude but an

opposite signs (making one side convex and the other concave)—when placed such that their principal points overlap, act precisely as a slab of material of equal thickness. Light rays incident on such a component exit it without any light ray direction changes, ultimately meaning the component does not do very much. When shifting one of the lenses relative to the other such that the principal points no longer align but remain on the same principal plane, meaning the setup remains telescopic, the sum of the thickness profile changes, which then directly dictates how such a construction impacts the phase of any light rays incident upon it. As it turns out, the resulting thickness profile is that of a simple wedge with a constant phase gradient. Therefore, incident light rays will be redirected as if passing through an equivalent refractive wedge. The same applies to plano-cylindrical lenses, where one side is flat and the other acts as a cylindrical lens. Importantly, a cylinder only has a single meridian direction that runs along the surface of the cylinder where it is completely flat. Namely an axis direction that is perpendicular to a circular cross-section of the cylinder. Since there is no curvature, the focusing power along such an axis is precisely zero, when placing a line such that it lies on the principal point of a plano-cylindrical lenses cross-section and along a direction of zero focusing power, a line which we will refer to as the nodal line is formed. Along this line, all incoming light rays parallel to the optical axis will not be deviated by the plano-cylindrical lens. The axes perpendicular to the nodal line along the surface of the plano-cylindrical lens are where the surface curvature will be largest. Therefore, the resulting wedge can easily be understood by considering a 2D construction that can simply be extended along the nodal line. Such a construction is shown in Fig. 2.2, including a bundle of redirected light rays. For brevity, all further plano-cylindrical lenses considered will simply be called cylindrical lenses.

The well-known Fresnel lens may be approximated as an array of wedges. It should be noted that the exact surface will contain a slight curve to focus the individual segments (we consider a segment to be the lens part separated by neighbouring Fresnel facets) which we do not yet consider. The resulting lens will redirect individual ray bundles for each segment into an approximate focal region shown in Fig. 2.1. As a result of this approximation, there will not be a single focal point. Ignoring diffractive effects, the segment size may be reduced indefinitely, causing the focal region to grow ever closer to a single point. Taking this as the starting goal for an adaptive Fresnel lens, the goal was to create a component that mimics the behavior of the approximated Fresnel lens but allows for a variation in the wedge gradient and hence the focal length of the resulting Fresnel lens.

As explored previously, one way to produce an optical wedge is by using a set of complementary lenses. Therefore, by replacing each wedge with the corresponding set of complementary cylindrical lenses vertically shifted such that their optical axis is misaligned, the approximate Fresnel lens may be recreated. If the shift is increased or decreased, the resulting phase profile will mimic another Fresnel lens of greater or smaller focal power, respectively. This is summarized in Fig. 2.3, where this can be extended to a 3D schematic by simply stretching it along the

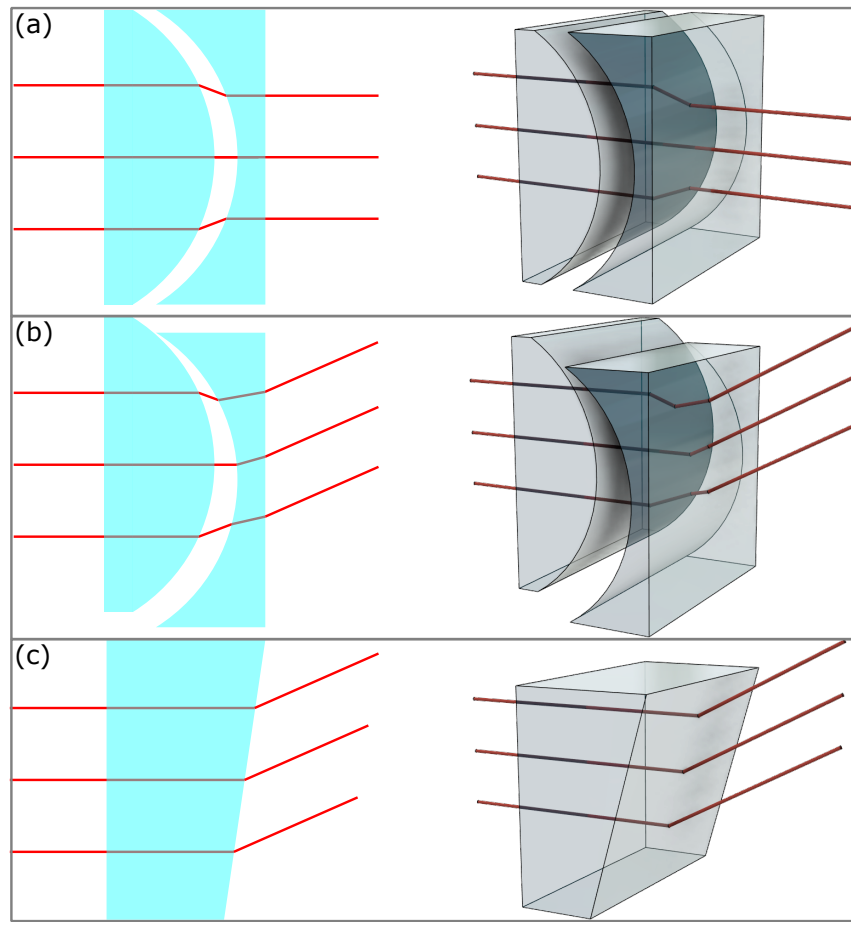


Figure 2.2: A set of complementary cylindrical lenses with a cross-section, perpendicular to the optical power direction, is shown on the left-hand side. A 3D sketch is provided on the right-hand side showing the equivalent lens by stretching it along the direction of the nodal line, a line on the principal point along the direction of zero optical power. In (a), the complementary parts are shown to be aligned, with one half having a focal length of f while the other has a focal length of $-f$. Therefore, as expected, the system is telescopic, and the light rays (shown as red lines) retain their direction upon transmission. When the lenses are now offset along the optical power direction in (b), the system is still telescopic (indicated by the rays remaining parallel after transmission), but the outgoing direction is changed. The direction change is directly proportional to the phase gradient of the light rays, which in turn relates to the relative length of refractive material a light ray passes through. Adding up the refractive material along the optical axis direction yields a simple refractive wedge shown in (c). The light rays, upon transmission through this wedge, are then redirected nearly identically to the cylindrical lens setup.

cylindrical axis of zero optical power in the same way shown in Fig. 2.2.

To understand the shift amount required and how closely this mimics a lens, the phase gradient of a lens can be compared to the phase gradient imposed by the cylindrical lens pairs. Starting from the expression of the phase change imposed on a light ray incident on a lens at a distance r from the center (and principal point) of the lens, given by

$$\Delta\Phi = \Delta\Phi_0 - \frac{kr^2}{2F}, \quad (2.1)$$




Figure 2.3: A complementary set of cylindrical lenses arranged into an array (which is typically called a lenticular array) is shown in (a). As expected, light rays incident on this (shown in red) retain their initial direction upon transmission in all cases. However, in Fig. 2.2, it was demonstrated that a simple set of complementary cylindrical lenses may be turned into a wedge. Therefore, only considering the 2D cross-section of an approximate Fresnel lens shown in Fig. 2.1(c), it may be possible to rearrange the lenticular array such that the addition of an offset complementary array creates an equivalent (in 2D) wedge array as the approximate Fresnel lens. This is shown in (b) with the solid red rays being focused onto a single point. Of course, since this is an approximation, another ray incident on the same telescopic cylindrical lens pair (here shown as a green light ray) will not focus on the same point but into the focal region mentioned in Fig. 2.1(c). Lastly, since this is now a telescopic system, this will be subject to stray light rays, which enter one telescope and leave the system through a neighboring one, shown by a dashed red ray in (b).

where $\Delta\Phi$ represents the phase offset, $\Delta\Phi_0$ is the uniform phase offset, which is just the phase at the center, F is the focal length of the lens and k is the wave number of the light. Taking the derivative of this then yields the corresponding phase gradient, Δk_r at some radius r given by

$$\Delta k_r = \frac{d}{dr} \Delta\Phi = -\frac{kr}{F}. \quad (2.2)$$

Since Eqn (2.2) applies to a parabolic lens, the expression holds true for all radial directions, forming *rings* of equal phase. Using only a simply stretched 3D version of the cylindrical lens array of the 2D version, shown in Fig. 2.3, will not allow this radial independence (in fact, the stretched version will mimic a cylindrical Fresnel lens). Therefore, for now, the focus was

to determine the required vertical shift for considering only a phase gradient along the optical power direction of the cylindrical lenses –the direction perpendicular to the nodal line, which will be referred to as the u direction for brevity. Rewriting Eqn (2.2) to reflect this for one of the cylindrical lenses yields

$$\Delta k_u = -\frac{ku}{f}, \quad (2.3)$$

where f is now the cylindrical lens focal length. Unlike F , which is the overall focal length of the desired Fresnel lens, f can be set for each cylindrical lens independently.

The expression for the phase gradient of two complementary cylindrical lenses can then, along the u direction, be given by the sum of their parts, namely the phase gradient due to the first cylindrical lens and then the second. In the case where the two complementary cylindrical lens optical axes align, the expression becomes

$$\Delta k_u = -\frac{ku}{f} + \frac{ku}{f} = 0, \quad (2.4)$$

due to their equal but opposite focal lengths. This result confirms the expected result stated previously, namely that two complementary cylindrical lenses aligned along their optical axis and sharing a common principal point do not induce any light ray direction change as their combined phase gradient is 0. The interesting case occurs when these are now shifted with respect to one another. Considering an equal but opposite shift along the u direction of $\pm\Delta u/2$ changes Eqn (2.4) into

$$\Delta k_u = -\frac{k(u + \Delta u/2)}{f} + \frac{k(u - \Delta u/2)}{f} = -\frac{k\Delta u}{f}, \quad (2.5)$$

where Δu is the relative shift between corresponding cylindrical lenses. Here, it is assumed that the shift is such that the second, negative focal length, term is the part that is shifted *further* away by Δu such that the aligned phase gradient parts are at $u + \Delta u/2$ for the first lens and $u - \Delta u/2$ for the second one. Equating Eqns (2.2) and (2.5) then results in

$$-\frac{k\Delta u}{f} = -\frac{kr}{F}, \quad (2.6)$$

provided that r lies along the u direction. In essence, this is a 2D equation of the cross-section of a lens as a cylindrical lens telescope. To generalize this further, the stretched cylindrical lens shown in Fig. 2.2 has to be wound into a ring such that u will always lie along the radial direction of the overall lens. Taking multiples of these at different radii and combining them into a single component then yields Fig. 2.4, a lens made of concentric cylindrical lens rings.

For these rings, we call R the radius of the nodal-line circles such that ΔR will be the radial difference between corresponding cylindrical lenses. In mathematical terms, this is represented




Figure 2.4: A 3D generalization to extend the complimentary lenticular arrays from Fig. 2.3 into a full 3D realization. First, a normal cylindrical lens is considered in (a), with the nodal line being shown as a black line along the surface of the cylindrical lens. This nodal line is then wound into the concentric circles shown in (b). When adding the equivalent cylindrical lenses back, a 3D realization of concentric cylindrical lenses can be made. A sketch of this is shown in (c), which can be called a “concentric lenticular array”. When simply adding the complementary concentric lenticular array to this, the same as in Fig. 2.3(a) would be achieved, namely a component that is equivalent to just a flat refractive slab. The more interesting case, however, is shown in (d) where a complementary concentric lenticular array is added but with the concentric nodal lines shifted outwards in proportion to the radial distance from the center. When looking at the cross-section in (d), it can be seen that the same pattern as in the lenticular arrays in which one array is stretched (Fig. 2.3(b)) appears, and therefore, it can be expected to act as an approximate Fresnel lens, but now this holds for all cross-sections through the center. In the cross-section, the local misalignment, ΔR , was shown in two places. The misalignment, ΔR , will be positive if the cylindrical lens with the negative focal length (top in (d)), will be shifted towards a greater radius relative to the positive focal length cylindrical lens (bottom in (d)). This will then corresponds to an overall positive focal length Fresnel lens.

by replacing the Δu term in Eqn (2.5) with ΔR , resulting in Eqn (2.6) being rewritten as

$$-\frac{k\Delta R}{f} = -\frac{kr}{F}. \quad (2.7)$$

Note that the substituting Δu by ΔR in Eqn (2.7) was possible as we consider the shift around the lens optical axis and therefore, substituting u by R , the radius at which the winding nodal line lies, into Eqn (2.5), along with a substitution of Δu by ΔR , which yields the left hand side of Eqn (2.7). For consistency, this also means that the radial shift is positive when the second negative focal length lens is radially further away from the center than the first positive focal length lens. Finally, this is rearranged into

$$\frac{\Delta R}{R} = \frac{f}{F}, \quad (2.8)$$

where r was assumed to be approximately equal to R , the radius at the winding nodal line. This is an important equation we call the *radial-offset equation*.

As demonstrated in Fig. 2.4 an approximate Fresnel lens equivalent was designed using two cylindrical lens arrays wound into rings. In the ideal case, this will then perform equivalently to the approximated Fresnel lens, and at worst, be a slightly worse lens with no added benefits. To turn this into a novel and useful idea, the concentric cylindrical lens array may be replaced by a single cylindrical lens. This lens is then wound into a spiral shape instead of a ring, and the same is done to the complementary cylindrical lens as shown in Fig. 2.5. This modification allows for a variation in ΔR according to the relative rotation between the two cylindrical lens spirals and as an extension of this, the Fresnel lens focal length, F , dictated by Eqn (2.8). This will inevitably introduce some aberration due to the pitch of the spiral. Since it does not form concentric rings, the azimuthal direction will no longer align with the u direction of the cylindrical lens. One way to mitigate this is by choosing spiral parameters that lead to a tighter wound spiral. Similarly, the spiral type will also influence the extent of the pitch and, hence, aberrations.

Initially, a logarithmic spiral shape was chosen, which follows

$$R_{log} = a \exp(b\phi) \quad (2.9)$$

in polar coordinates, where R_{log} is the radius from the center (with the center aligned to the optical axis), a is the spiral size or scale factor, while ϕ is the unbound azimuthal angle *not* restricted to a 2π range. The last parameter, b , is related to the winding density of the spiral, where smaller values correspond to a more tightly wound spiral. This is also why we will refer to b as the winding parameter from here on.

By relating the azimuthal angle, or more specifically, the difference in azimuthal angle $\Delta\phi$ between two relatively rotated spirals, the radial offset, ΔR , may be found. Mathematically, this can be expressed as

$$\begin{aligned} \Delta R &= R_2 - R_1 \\ &= R \left(\phi + \frac{\Delta\phi}{2} \right) - R \left(\phi - \frac{\Delta\phi}{2} \right), \end{aligned} \quad (2.10)$$

where R_1 and R_2 are the radial distances of the first, positive focal length, and second, negative




Figure 2.5: A similar consideration as in Fig. 2.4 was designed, where the same cylindrical lens and nodal line are shown in (a). Unlike the aforementioned figure, the nodal line is no longer wound into concentric circles but instead into a logarithmic spiral, shown in (b). When adding the cylindrical lens back to the nodal line in (c), a cylindrical lens wound into a spiral was made. Now, unlike the concentric circle, the exact complementary component was added back. When nothing further is done, it will act like a solid slab of flat refractive material; however, by simply rotating the two components relative to one another, the same radial offset as shown previously may be achieved in (d). The offset is shown as ΔR in the cross-section, and is proportional to the radial distance and the relative rotation angle. Therefore, ΔR may be adjusted simply by rotating the components.

focal length, cylindrical lens spirals respectively, highlighted in Fig. 2.6. Expressing these in terms of the spiral rotation angle, R_2 contained an extra rotation of $\Delta\phi/2$, while the extra rotation of R_1 was set to $-\Delta\phi/2$, such that their relative rotation was $\Delta\phi$. Replacing the expression for a logarithmic spiral for R , and simplifying it down, then yields

$$\begin{aligned}\Delta R_{log} &= R_{log}(\phi) 2 \sinh\left(\frac{b\Delta\phi}{2}\right) \\ &\approx R_{log}(\phi) b\Delta\phi,\end{aligned}\tag{2.11}$$

for $\Delta\phi$ in radians. The last line of Eqn (2.11) is an approximation and stems from the Taylor series expansion of the $2 \sinh(b\Delta\phi/2)$ term around $b\Delta\phi = 0$ and truncating it after the second order under the assumption that $b\phi \ll 1$. An additional graphical representation of this radial offset can be seen in Fig. 2.6 for a logarithmic spiral, indicating that it is indeed directly proportional

to the radial distance.




Figure 2.6: A plot of two logarithmic spirals rotated relative to one another about an axis passing through the center into the page. The relative rotation is angle is $\Delta\phi$, whereby the second, dashed, spiral is rotated in the clockwise direction. As a result of this, the radial distance of the two spirals along any direction, ϕ , will no longer be the same. This is shown by R_1 , and R_2 for a particular ϕ no longer being of the same length. Therefore, the radial offset between the two spirals at this point may then also be given by $R_2 - R_1$.

Finally, the expression for the logarithmic radial offset, ΔR_{log} , from Eqn (2.11) can be substituted into the radial-offset equation, Eqn (2.8), to relate the relative rotation to the resulting Fresnel lens focal length, F , as

$$F = \frac{f}{b\Delta\phi}. \quad (2.12)$$

This outlines a convenient realization of the chosen spiral, namely that in order to achieve a Fresnel lens focal length, F , that does not vary with radius (just like the focal length of a lens does not vary with radius) but only with the relative rotation angle $\Delta\phi$, a constant cylindrical lens focal length, f , must be chosen. Therefore, the optical power of the *adaptive* Fresnel lens should only and directly depend on the relative rotation angle between the two cylindrical lens spirals. To confirm this, ray trace simulations were performed in our open-source ray tracer called Dr. Tim [4]. These can be seen in Fig. 2.7, and show some of the expected shortcomings of the basic adaptive Fresnel lens. We note that for simulation purposes, the cylindrical lens spirals throughout this chapter were approximated by phase holograms of cylindrical lenses.

Two distinct shortcomings can be seen. First, due to the relative rotation, crosstalk between neighboring windings can occur, which manifests as a second image (a de-magnified image of the Tim head, visible in the centre of the nose of the main image) in the simulations. This is an issue for the original purpose of normal spectacles for UREs, as the additional double image is undesirable. There may be some other applications where this is instead a feature rather than a shortcoming, which are outlined in section 2.6. The second shortcoming stems from the images




Figure 2.7: A ray-traced simulation of an adaptive Fresnel lens using the open-source raytracer called Dr. Tim [4]. Here, (a) shows the scene (a “Tim head” on a checkerboard floor) being viewed without any component. When adding the complementary set of cylindrical lenses wound into a spiral shown in Fig. 2.5(d), but not rotated relative to one another, (b) was made. As expected, no significant changes (other than a slight darkening of the view due through the components due to the lens components being made slightly absorptive in order to be visible) are noticeable. However, upon a relative rotation, $\Delta\phi$, of 30° in (c), the image seen appears significantly larger. In fact, this should be the same image seen as when looking through an ideal lens of focal length 12.7 (calculated from Eqn (2.12)). To check this, exactly such an ideal lens was added in (d), confirming the result. The parameters chosen for this simulation were $b = 0.015$ and $f = 0.1$. Some imperfections are clearly visible in (c), where the images of neighboring windings do not perfectly match up. The spiral boundaries are also very clearly visible, as this is where the stray light rays shown in Fig. 2.3(b) will occur. We note that the focal lengths do not have units, as the raytraces operate in dimensionless units, where a floor tile is a standard unit of length 1. Throughout this chapter, this will be the standard unless stated otherwise. All the ray trace simulations in this chapter were performed using a runnable Java archive file called `SpiralLensVisualiser`, available at [3], unless otherwise stated.

within neighbouring windings not fully lining up, which creates a worse composite image. This is due to each winding acting like a radial wedge instead of a small part of a lens, and it does not conveniently hold an alternative use, which is why “winding focusing” was required.

2.3 Alvarez-Lohmann winding focusing

In order to properly focus light at a single focal plane, the approximations made to the Fresnel lens in Fig. 2.1(c) had to be undone. In a proper Fresnel lens, each of the facets redirects incident light rays toward the focal region but also has some focusing power equal to the inverse distance to the focal plane. In the adaptive Fresnel lens design, each winding by default corresponds to a zero focusing power Fresnel facet. To address this, winding focusing was introduced, which, as the name suggests, added back the desired focusing power to each winding. Effectively this then means that winding focusing changes the simplified (flat) Fresnel lens facets in Fig. 2.1(c) into the proper (curved) Fresnel lens facets shown in Fig. 2.1(b).

The focusing power added was also required to remain adaptive in order to maintain the general adaptive properties of the component. One way in which an adjustable focusing power may be achieved is through the use of an Alvarez-Lohmann lens [59, 60]. When two complementary third-order-polynomial refractive surfaces are shifted relative to one another, it is possible to introduce the equivalent focusing power. When restricting the Alvarez-Lohmann lens to only contain one axis with optical power, a type of Alvarez-Lohmann cylindrical lens may be created. This, in a similar way, will introduce focusing power but only along the optical power (or u) direction of the lens.

The basic non-focusing version of the adaptive Fresnel lens also operates by shifting the complementary cylindrical lenses with respect to one another. The phase change imparted by a cylindrical Alvarez-Lohmann lens is

$$\Delta\Phi_{AL}(u) = \frac{1}{6}Aku^3 \quad (2.13)$$

for each cylindrical lens, where u is the optical power direction, k is the wave number, and A is some constant we call the Alvarez-Lohmann coefficient. The combination of two such lenses shifted relative to one another is the sum of their imparted phase changes. Adding these in the same manner as in Eqn (2.5), (where the first component has a positive relative rotation and a positive A , while the second has a negative relative rotation and negative A) then gives

$$\Delta\Phi_{AL} = \Delta\Phi_{AL}\left(u + \frac{\Delta u}{2}\right) - \Delta\Phi_{AL}\left(u - \frac{\Delta u}{2}\right), \quad (2.14)$$

which, after a few simplification steps outlined in the supplemental information of [64], can be reduced to

$$\Delta\Phi_{AL} = \Delta\Phi_0 + \frac{ku^2}{2}A\Delta u, \quad (2.15)$$

where $\Delta\Phi_0$ is a spatially uniform phase shift. Comparing this with the phase change expression

of a lens of focal length f_{AL} given in Eqn (2.1), then gives

$$f_{AL} = -\frac{1}{A\Delta u}. \quad (2.16)$$

In the same way as shown in Eqn (2.7), this can be written for an Alvarez-Lohmann lens wound into the same spiral as the cylindrical lens, such that the shift Δu becomes an approximately radial shift term ΔR , where ΔR is the shift between corresponding windings, giving

$$f_{AL} = -\frac{1}{A\Delta R}. \quad (2.17)$$

In order to focus the individual windings, the focusing power of each winding pair must be that of the overall Fresnel lens focal length, F , obtained through the radial offset equation Eqn (2.8). When adding the Alvarez-Lohmann lens to each winding, the focusing power of each winding will be entirely due to the Alvarez-Lohmann lens (as the cylindrical lens spirals themselves do not produce any focusing power by design). Therefore, we can equate only the Alvarez-Lohmann focal length Eqn (2.17) to the Fresnel lens focal length to determine the Alvarez-Lohmann coefficient, A , required for the appropriate winding focusing, yielding

$$\frac{fR}{\Delta R} = -\frac{1}{A\Delta R}. \quad (2.18)$$

This conveniently contains two ΔR terms that cancel out one another to finally give an Alvarez-Lohmann coefficient, A , of

$$A = -\frac{1}{fR}, \quad (2.19)$$

which does not depend on any ΔR terms as desired. Therefore, regardless of rotation angle, the appropriate focusing power will be added if a cylindrical Alvarez-Lohmann lens is used with the coefficient changing according to Eqn (2.19).

In the case of a logarithmic spiral, f is constant and therefore the Alvarez-Lohmann coefficient simply changes with radial distance from the center. A ray optical simulation confirms this, showing a significantly improved image in Fig. 2.8.

2.4 More winding focusing

A property of telescopes, is that parallel incoming rays remain parallel upon transmission, resulting in zero focusing power. This can be shown by considering the expression for the combined focusing power of the system as

$$\frac{1}{F_{12}} = \frac{1}{f_1} + \frac{1}{f_2} - \frac{s}{f_1 f_2}, \quad (2.20)$$




Figure 2.8: The exact same setup as shown in Fig. 2.7(C), but now when using Alvarez-Lohmann winding focusing. It can be seen that the stray light rays (or ghost image) at the spiral boundaries remain, but now the images between neighboring windings match up almost perfectly.

where f_1 and f_2 are the focal lengths of the first and second lens, respectively, and s is the separation along the optical axis. For the expected telescopic case, f_2 will be the same as f_1 but with the opposite sign. Additionally, in the basic construction, their separation is such that they are telescopic, which is achieved by having a separation of zero. Therefore, it can be seen that Eqn (2.20) will be zero, as expected.

Winding focusing requires that the focal power of the individual windings is equal to the corresponding Fresnel lens focal length, F . Besides the previously discussed Alvarez-Lohmann winding focusing in section 2.3, the same winding focusing effect could be achieved by no longer restricting the corresponding windings, and hence, cylindrical lenses to be perfectly telescopic. Therefore, either their separation or focal lengths may be changed.

2.4.1 Winding focusing by focal length

In order to adjust the focal length of one cylindrical lens to no longer be of the same magnitude as the other ($f_1 \neq -f_2$), Eqn (2.5-2.8) need to be reevaluated for differing focal lengths. Starting at Eqn (2.5), which in turn becomes

$$\Delta k_u = -\frac{k((r-R) + \Delta R/2)}{f_1} + \frac{k((r-R) - \Delta R/2)}{f_2} = -\frac{k\Delta u}{f}, \quad (2.21)$$

where u and Δu were substituted for $(r-R)$ and ΔR respectively. Equating this to the phase gradient expression of a Fresnel lens, Eqn (2.2), leads to

$$-\frac{R}{F} = \frac{\Delta R(f_1 - f_2) - 2(r-R)(f_1 + f_2)}{2f_1 f_2}, \quad (2.22)$$

the new (and unfortunately significantly more complicated) radial offset equation.

Solving for the resulting Fresnel focal length of a logarithmic spiral following the same steps as outlined in section 2.2 ultimately leads to

$$F = -\frac{2f_1f_2}{b\Delta\phi(f_1 - f_2)}. \quad (2.23)$$

Note that in one of the steps outlined in `AdaptiveFresnelLensWF` [1] requires a simplification by assuming that $r \rightarrow R$. We permit this because this equality should hold for any radius, including the nodal line (or center) of each winding, where $r = R$. From Eqn (2.23), it becomes apparent that the Fresnel focal length varies directly with respect to $\Delta\phi$ as desired.

The focal power of the resulting Fresnel lens is the inverse of Eqn (2.23) and can be equated to the expression for the focal power due to two lenses, Eqn (2.20), to achieve winding focusing as explained in section 2.3. Solving the expression for f_2 gives

$$f_2 = \frac{f_1(2 + b\Delta\phi) - 2s}{b\Delta\phi - 2}. \quad (2.24)$$

A relative rotation of zero yields a Fresnel lens of infinite focal length, which is equivalent to a lens that performs no resulting light ray direction changes. Therefore, the expected solution for no relative rotation should be a winding-focused system that also has zero focusing power. Setting $\Delta\phi \rightarrow 0$, Eqn (2.24) then simply yields $f_1 - s$, which leads to a telescopic system. As a property of a telescope, parallel rays remain parallel, and therefore, the telescope system contains zero focusing power as desired. To check the expression for other rotation angles, it can be ray optically confirmed through ray trace simulations shown in Fig. 2.9 where the same setup as in Fig. 2.7(c) was taken and focused.

The most obvious shortcoming of using the cylindrical lens focal length changes to achieve winding focusing is the dependence on $\Delta\phi$, which means that a change in relative rotation would require a different cylindrical lens focal length. Since the Fresnel lens we are hoping to create should be *adaptive*, this is not a viable option to achieve winding focusing. Instead, this may be useful when attempting to make the adaptive Fresnel lens work best around a certain focal length or when attempting to allow for some extra separation between the components, discussed in more detail in section 2.6.

2.4.2 Winding focusing through separation

The second way in which the winding focusing may be achieved is through varying the separation between the cylindrical lens spirals. Similar to the method outlined in section 2.4.1, this was done to arrange the cylindrical lens spirals such that they were no longer perfectly telescopic. Assuming that the cylindrical lenses are perfectly complementary ($f_1 = -f_2$), the expression representing the combined focal power due to the combination of two cylindrical lenses, Eqn




Figure 2.9: In a similar way to Fig. 2.8, here another winding focusing method is demonstrated. This was done by varying the cylindrical focal length, f , of the second negative focal length cylindrical lens spiral slightly. It can be seen that the image produced is equivalent to that shown in the aforementioned figure. However, this method requires a change in f , the focal length of the cylindrical lens, depending on the relative rotation angle, θ . Since the focal length of the cylindrical lenses is supposed to be locally static within the spiral, specifically, it may vary with ϕ , the unbound azimuthal angle, but not with θ , this type of winding focusing is not feasible.

(2.20), can be rewritten as

$$F_{12} = \frac{f^2}{s}, \quad (2.25)$$

where f_1 was replaced by just f , the chosen focal length of one of the cylindrical lenses, and F_{12} is the combined focal length. The elementary images formed by the combination of cylindrical lenses can then be seamlessly fused together to create an integral image only if the combined cylindrical lens focal length found is the same as the Fresnel focal length of the spiral system. Expressing this mathematically leads to

$$F = \frac{f}{b\Delta\phi} = \frac{f^2}{s} = F_{12}, \quad (2.26)$$

an equality that allows for the separation, s , to be found as a function of the logarithmic spiral and focal length parameters. From all the spiral types investigated, outlined in section 2.5 in more detail, only the logarithmic spiral simplifies down nicely to

$$s = fb\Delta\phi, \quad (2.27)$$

a function that is only dependent on the relative rotation, $\Delta\phi$, and two constants for a logarithmic spiral lens, f and b . Therefore, to achieve winding focusing, the separation between the two cylindrical lens spirals needs to be adjusted linearly according to the relative rotation angle in radians. The rate at which the components separate depends on the spiral parameters and the

focal length of the cylindrical lens chosen. A basic check reveals that for a relative rotation of zero, the separation must also be zero. This is, of course, the case where the complementary lenses align perfectly to form a surface with no phase gradient and, hence, no resulting light ray direction change. Therefore, the winding focusing should also make the cylindrical lens pair focus at infinity, which is only achieved when the two cylindrical lenses are telescopic. For the complementary case considered, this occurs at zero separation, aligning with the mathematical expression found.




Figure 2.10: The last winding focusing method, which adds a small amount of space between the two complementary cylindrical lens spirals, such that they are no longer perfectly telescopic. As a result, some focusing power is added, leading to the separation winding focusing method shown here. It can be seen that the image produced matches that of the other winding focusing methods, with the same shortcomings due to a ghost image.

To test further rotation angles, ray-traced simulations were used to create Fig. 2.10, which shows the winding focusing effect for a 30° rotation angle that can be compared to the unfocused case along with an equivalent normal lens case shown in Fig. 2.7(c) and (d). A significant improvement can be seen, which becomes even more apparent for larger rotation angles, where the unfocused case shows several elemental images that do not align with the neighboring ones, leading to a degraded integral image.

2.4.3 Combined winding focusing

There are several factors to consider when determining the preferred winding focusing method. A full description of some of the feasibility issues with the adaptive Fresnel lens is given in section 2.6. To aid with some of the outlined issues, it is therefore advantageous to consider not only one but multiple ways of achieving winding focusing. Therefore, combining the ways this is performed in sections 2.4.1 and 2.4.2 can yield an enhanced winding focusing method that is based on separation focusing, which holds the advantage that it is fully adaptive (no cylindrical

lens focal lengths can change in a *static* spiral lens), but with the added complexity of permitting non complimentary cylindrical lenses i.e, $|f_2| \neq |f_1|$.

In this case, the Fresnel lens focal length Eqn (2.23), obtained using the new radial offset equation in Eqn (2.22), must be considered. In the same manner as outlined in section 2.4.2, the inverse of this new Fresnel lens focal can be equated to the focusing power of a combination of two cylindrical lenses given by Eqn (2.20). Equating the two yields

$$\frac{1}{F} = \frac{b\Delta\phi(f_1 - f_2)}{2f_1f_2} = \frac{1}{f_1} + \frac{1}{f_2} - \frac{s}{f_1f_2} = \frac{1}{F_{12}}. \quad (2.28)$$

Solving the system for s and simplifying it down, then yields an expression for s given by

$$s = f_1 + f_2 + \frac{b\Delta\phi}{2}(f_1 - f_2), \quad (2.29)$$

which only depends on the relative rotation, $\Delta\phi$, and constant cylindrical lens and spiral parameters (f_1 , f_2 , and b respectively) as desired. By changing f_1 and f_2 it is therefore possible to change the separation at no relative rotation and then further change the separation according to the relative rotation and rate of change. This lifts the initial restrictions of requiring the spiral principal points to lie within the same plane along the u direction, improving feasibility.

To test the result obtained in Eqn (2.29), ray trace simulations were used to create the images seen in Fig. 2.11. These prove not only that the combined winding focusing works but also highlight the significant improvement to the integral image formed, especially when compared to the equivalent separation of a setup without any winding focusing or telescopic adjustments.

2.5 Other spiral types

Up until this point, only the logarithmic spiral was considered. There are infinitely many ways in which a spiral may be wound. The logarithmic lens was chosen due to the constant cylindrical lens focal length it provided for the adaptive Fresnel lens. An alternative consideration is a spiral type with windings that are always equidistant from the neighboring winding, leading to a spiral whose windings will be the same thickness regardless of radial distance from the center. Such a spiral is called the Archimedean spiral. This spiral type will be investigated in section 2.5.1. Another interesting spiral was the hyperbolic spiral type. Due to the expression for the winding and cylindrical lens focal length variation of the hyperbolic spiral, it was considered to be inherently winding focused, leading to an Alvarez-Lohmann coefficient, A , of 0. This, along with the relevant ray trace simulation, is demonstrated in section 2.5.2.




Figure 2.11: A demonstration of a combination of separation winding focusing, and a change in focal length of the second cylindrical lens spiral. In (a), a simple complementary cylindrical lens spiral of focal lengths $f_1 = 0.1$ and $f_2 = -0.1$ was taken and separated by a distance 0.0048, leading to a significantly deteriorated image as it was no longer telescopic. When adjusting the focal length of f_1 to $f_1 = 0.104$, and separating the two components by 0.004, the system becomes telescopic once more, leading to an improved image shown in (b), however, throughout the neighboring winding images do not align perfectly as previously seen in the case where no winding focusing was used. To account for this, separation winding focusing was added back, adjusting the relative separation to 0.0048 for the same focal lengths. As a result, (c) was observed. It can be seen that this is still not perfect, especially further from the center due to the additional magnification added by the telescopic pair, and the assumption made when calculating the required additional distance in Eqn (2.23), but it is significantly better than both (a) and (b).

2.5.1 Archimedean spiral

The Archimedean spiral follows an expression for the radial distance, R_{Arch} , from the center of the spiral as

$$R_{Arch}(\phi) = b\phi, \quad (2.30)$$

where b and ϕ are the winding parameter and azimuthal angle, the same as in the logarithmic lens. When considering two identical spirals rotated relative to one another, the separation

between corresponding windings, ΔR_{arch} , can then be given by

$$\begin{aligned}\Delta R_{Arch} &= R_{Arch} \left(\phi + \frac{\Delta\phi}{2} \right) - R_{Arch} \left(\phi - \frac{\Delta\phi}{2} \right) \\ &= b\Delta\phi,\end{aligned}\tag{2.31}$$

where the same steps as in section 2.2, specifically Eqn (2.10), were followed. Plotting two spirals of the form given in Eqn (2.30) rotated relative to one another then yields the expected result shown in Fig. 2.12, where each of the spiral windings are all of the same width, and the separation between corresponding winding pairs is independent of radius and therefore constant for each relative rotation angle. This is a change compared to the logarithmic spiral, where the separation was proportional to the radius.




Figure 2.12: A plot of two Archimedean spirals rotated relative to one another by $\Delta\phi$. It can be seen that this specific spiral type has a constant winding width, with the line spacing being equidistant. Furthermore, when rotating the spirals relative to one another, the radial offset, ΔR , can be seen to be the same everywhere. This means it should be independent of radial distance from the center, unlike what was observed in Fig. 2.6.

This change in form leads to an updated Fresnel focal length of the system obtained by replacing ΔR in the radial offset equation and solving for F , the new Fresnel focal length, now called F_{Arch} for clarity. The expression found can be given as

$$F_{Arch} = \frac{fR}{b\Delta\phi},\tag{2.32}$$

and now depends on the radius from the center, R . Since we want the overall Fresnel lens focal length to be the same irrespective of the position on the adaptive lens, f must become a function

of R . To turn Eqn (2.32) into an equation independent of R , the focal length, $f(R)$, now as a function of R will be given by

$$f(R) = \frac{f_{R=1}}{R}, \quad (2.33)$$

where $f_{R=1}$ is a set focal length defined on the cylindrical lens spiral at radius 1. Therefore, the focal length across the cylindrical lens spiral must change, with the very center having zero focusing power. This is different from the logarithmic case, which was specifically chosen as the focal length of the cylindrical lens spiral may be constant.

Replacing the f term in Eqn (2.32) then yields the final Fresnel lens focal length as

$$F_{Arch} = \frac{f_{R=1}}{b\Delta\phi}. \quad (2.34)$$

The resulting view through such an adaptive lens can be seen in Fig. 2.13, which also demonstrates the equal winding thickness.




Figure 2.13: A ray-traced simulation of an adaptive Fresnel lens using cylindrical lenses wound into an Archimedean spiral. It can be seen that the overall image does show the expected Timex head, but there are significant issues with the image quality as the elemental images do not align well to form an integral image.

To improve this, winding focusing must be used again. Unfortunately, the winding focusing required changes with every spiral type. When attempting to perform winding focusing through separation or focal length adjustments, the resulting separation or focal length change in both cases depends on ϕ and hence on R . This may be possible, but is not desired, as it would require that the cylindrical lenses themselves be wound such that their windings gain (or lose) elevation further from the center. In the worst case, this can spiral out of control, where the component required becomes very thick, making it undesirable for use in spectacles. In general, this will increase the manufacturing complexity and is difficult to simulate by the currently used phase holograms.

As an alternative, Alvarez-Lohmann winding focusing may be used. For any relative rotation, $\Delta\phi$, that is not zero, the overlapping parts of the two spirals will not be perfectly complementary as the focal length of the individual cylindrical lenses changes with respect to ϕ and hence the radius. This results in some focusing power simply through rotation, which must be accounted for by replacing f_1 and f_2 in Eqn (2.20) with the expressions for the focal lengths of the Archimedean cylindrical lens spiral, Eqn (2.33). Taking the same relative rotation as shown in Eqn (2.31) and Fig. 2.12 for the radial distance, and assuming that the system is complementary at $\Delta\phi = 0$ then yields

$$\frac{1}{F_{12}} = \frac{R - \Delta R/2}{f_{R=1}} - \frac{R + \Delta R/2}{f_{R=1}} = -\frac{\Delta R}{f_{R=1}}. \quad (2.35)$$

This added focal power must then, together with the Alvarez-Lohmann focusing power (the inverse of Eqn (2.17)), equal the focusing power of the Archimedean Fresnel focal power (the inverse of Eqn (2.34)). Setting up this equality and solving for the Alvarez-Lohmann coefficient, A , then gives

$$\begin{aligned} \frac{1}{f_{AL}} + \frac{1}{F_{12}} &= \frac{1}{F_{Arch}} \\ \Rightarrow -A\Delta R_{Arch} - \frac{\Delta R_{Arch}}{f_{R=1}} &= \frac{\Delta R_{Arch}}{f_{R=1}} \\ \Rightarrow A &= -\frac{2}{f_{R=1}}, \end{aligned} \quad (2.36)$$

which only depends on $f_{R=1}$.

To see the effectiveness of this, another ray-trace simulation was performed in Fig. 2.14. This again serves to verify the mathematical construction and shows that a significantly improved image can be seen using the appropriate winding focusing. However, one unfortunate consequence of the chosen spiral type is the negative sign of the $1/F_{12}$ term in Eqn (2.36). With some subtle foreshadowing, if this were to somehow be positive, the Alvarez-Lohmann coefficient would simply be 0 and hence produce an adaptive Fresnel lens that will automatically perform winding focusing.

2.5.2 Hyperbolic spiral

Adapting the approach taken for the Archimedean spiral, where a constant winding width was desired, to a spiral that may be automatically winding focused lead to the Hyperbolic spiral. This was initially found by working backward from a zero Alvarez-Lohmann coefficient when equating the focal powers as shown in Eqn (2.36), but for consistency, this subsection will retain the structure of the previous spiral types. We are therefore, starting from the expression of such




Figure 2.14: The exact same setup as in Fig. 2.13, but now when adding the Alvarez-Lohmann winding focusing to the cylindrical lenses. As desired, this leads to a significantly improved integral image where the elemental images of each winding now align properly. As was the case previously, this should mimic a Fresnel lens of focal length 12.7, which is shown in Fig. 2.7(d), which it does nearly perfectly except for a small, unrecognizable, spiral shaped slither of a ghost image and a very distinct point in the very center, where the spiral ends in a 90° pitch.

a Hyperbolic spiral, where the radial distance of each winding is given by

$$R_{hyp} = -\frac{1}{b\phi}, \quad (2.37)$$

where the parameters are the same as in the other spiral types. To find the relevant expression for ΔR_{hyp} , we assume that the radial change is small, such that ΔR_{hyp} may be rewritten as dR_{hyp} . This can then be expressed by

$$\begin{aligned} dR_{hyp} &= \frac{dR_{hyp}}{d\phi} d\phi = \frac{1}{b\phi^2} d\phi = bR_{hyp}^2 d\phi \\ \Rightarrow \Delta R_{hyp} &\approx bR_{hyp}^2 \Delta\phi, \end{aligned} \quad (2.38)$$

where undoing the dR and $d\phi$ substitution reveals an approximate relationship between the radial offset of the hyperbolic spiral, ΔR_{hyp} , and the relative rotation angle between the two cylindrical lens spirals, $\Delta\phi$. A plot of such a spiral outlining the windings and the radial separation can be seen in Fig. 2.15.

It should be noted that, unlike the other two spiral types discussed, the hyperbolic spiral approaches the center from an infinite distance away with increasing ϕ (where it is at infinity for $\phi \rightarrow 0$ and at the centre for $\phi \rightarrow \infty$), producing windings of smaller width, eventually reaching a width of zero towards the center. Additionally, a negative sign was placed in front of the radial expression such that the final adaptive Fresnel lens focal length was of the same form as the other spirals, made apparent in Eqn (2.41).




Figure 2.15: A plot of a pair of hyperbolic spirals rotated relative to one another (shown as a solid and as a dotted line respectively). The hyperbolic spiral is wound the opposite way to the previous spirals, where it approaches the center from infinity when increasing the winding number. This is an inverse relationship with ϕ which leads to the windings becoming (infinitely) small towards the center and their separation shrinking accordingly. This also occurs in the logarithmic spiral but is shown at a decreased rate for spiral parameters that may be attractive for spectacles. Due to this, the center of the spiral appears to be black, representative of a very tightly wound spiral. When looking at the separation between the two spirals given by ΔR , it can be seen that this varies greatly with the distance from the center and in conjunction with ϕ .

Implementing the hyperbolic radial offset found into the radial offset equation, Eqn (2.8), then yields a Fresnel focal length of

$$F_{hyp} = \frac{f}{bR\Delta\phi}. \quad (2.39)$$

Similar to the Fresnel focal length of the Archimedean spiral, this depends on R , which is not desired. This can be offloaded into the f term by defining the cylindrical lens focal length as a function of R , $f(R)$ given by

$$f(R) = f_{R=1}R = -\frac{f_{R=1}}{b\phi}, \quad (2.40)$$

where $f_{R=1}$ is a constant focal length at radius one. The expression is also written in terms of b and ϕ , which will be useful when finding the combined focusing power. Inserting the expression

for f into Eqn (2.39) then gives

$$F_{hyp} = \frac{f_{R=1}}{b\Delta\phi} \quad (2.41)$$

as the final expression for the Fresnel focal length of a hyperbolic spiral. It can be seen that this is not only independent of R as desired but also has the same form as for the other spiral types.

In much the same way as in the Archimedean spiral type, this cannot easily be winding focused by separation due to a ϕ and hence R dependency when $f_2 \neq -f_1$. For equal focal lengths, and hence a separation of zero between the two cylindrical lens spirals, it is possible to test the focusing requirements by calculating the required Alvarez-Lohmann coefficient. To do this, the focusing power of each winding pair has to be found for any relative rotation angle. Using Eqn (2.20) this then gives

$$\frac{1}{F_{12}} = \frac{b(\phi + \Delta\phi/2)}{f_{R=1}} - \frac{b(\phi - \Delta\phi/2)}{f_{R=1}} = \frac{b\Delta\phi}{f_{R=1}} = \frac{1}{F_{hyp}} \quad (2.42)$$

where the focal lengths are written with the winding angle, ϕ and relative rotation angle $\Delta\phi$ instead of R and ΔR . This was done purely for simplification reasons, with the negative focal length cylindrical lens spiral having a negative relative rotation angle. This may seem to be contrary to the previous spiral types, but this is consistent with the negative focal length cylindrical lens spiral windings being radial further from the center for a hyperbolic spiral.

Combining the focusing power due to an Alvarez-Lohmann profile along with the focusing power found due to the combination found in Eqn (2.42) and equating it to the desired focal power of the overall Fresnel lens created by the cylindrical lens spiral pair gives

$$\begin{aligned} \frac{1}{f_{AL}} + \frac{1}{F_{12}} &= \frac{1}{F_{hyp}} \\ \Rightarrow -A\Delta R_{Arch} + \frac{b\Delta\phi}{f_{R=1}} &= \frac{b\Delta\phi}{f_{R=1}} \\ \Rightarrow A &= 0. \end{aligned} \quad (2.43)$$

This very clearly shows the desired result, namely, an Alvarez-Lohmann coefficient of zero. Therefore, the spiral should, in the case of two complementary cylindrical lens spirals at zero separation, automatically perform winding focusing. To confirm this, a ray trace simulation was performed in Fig. 2.16. This validates the theoretical findings by demonstrating an adaptive Fresnel lens that is focused without adding any separation or adjusting the phase profile.

2.6 Physical restrictions and feasibility

Regardless of spiral type, the best view *quality* (not necessarily the desired view) is always achieved at a relative rotation between the components of zero. Unfortunately, this also directly dictates that the preferred operating range is for large focal lengths since all expressions of the




Figure 2.16: A ray-traced look through an adaptive Fresnel lens where the cylindrical lenses are wound into a hyperbolic spiral. In this case, the same equivalent Fresnel focal length lens as in the previous ray trace figures is desired. It can be seen to not only reproduce the desired focal length as well as the other focused spirals, but it does so without the use of any additional winding focusing. Therefore, this specific spiral type can be seen to be entirely self-focusing.

Fresnel focal length found, regardless of spiral, have an inverse dependence on $\Delta\phi$. In fact, the best case is when the Fresnel focal length is infinite, meaning a component with zero focusing power, which means it acts essentially like a slab of refractive material. This is, of course, not ideal since the initial goal was for the adaptive Fresnel lens to work in spectacles.

The degradation comes in the form of a ghost image, where spirals of different windings overlap. Light rays that then pass through one winding in the first cylindrical lens spiral and a neighboring winding in the second can be said to undergo “winding hopping”. This can be clearly seen in Fig. 2.17, where a logarithmic spiral was simulated with a relative rotation of 180° , which creates a view composed of the ghost image and expected image taking up the same overall area. What this figure also shows is that the ghost image is, in fact, just the image due to a different lens, which has a Fresnel focal length given by

$$F_{alt} = \frac{f}{b(\Delta\phi \pm 2\pi)}, \quad (2.44)$$

where the “ \pm ” is chosen such that $\Delta\phi \pm 2\pi$ remains in the range of -2π to 2π (or in a compact term $-2\pi \leq \Delta\phi \pm 2\pi \leq 2\pi$). Therefore, this ghost image may, in some instances, not be considered an issue but a feature where a lens with dual focal lengths is desired. Similar lenses have previously been used for microscopy [65] to extend the depth of field.

Regardless of the potential applications of such a system, in the case of simple spectacles, this is not desirable. One option to remove the overlapping regions is by constructing a type of adaptive baffle using polarization outlined in [64]. This then simply makes the overlapping region appear black, effectively dimming the overall view for sufficiently small windings and




Figure 2.17: A brief investigation into the ghost image found in all adaptive Fresnel lenses, irrespective of spiral type. In (a), an adaptive Fresnel lens using a logarithmic spiral winding was viewed with the two components rotated by 180° with respect to one another. This then led to the case where the ghost image and the desired image take up an equal fraction of the overall view. It can be seen that two Tim heads become visible, one being larger and inverted, while the other is small and upright. These should correspond to the images of a lens of a Fresnel focal length found by Eqn (2.12) and Eqn (2.44) respectively. To check this, two ideal lenses were set up with their focal lengths set according to this. In (b), it can be seen that the image seen matches that of the inverted Tim head, while (c) shows that of the smaller upright Tim head.

relative rotations. However, this approach does not inherently address the fundamental issue, namely that at zero relative rotation angles, the component does “nothing” but works best.

When considering a physical construction of the adaptive spectacles as shown in Fig. 2.18(a), the ocular (closer to the eye) and objective (furthest from the eye) surfaces are both flat. By adjusting these instead, it may be possible to create a lens that no longer operates best at zero focusing power. Adding some focusing power to the flat surfaces effectively just adds some focusing power to the whole system, thereby creating an adaptive Fresnel lens that, when not rotated, is simply a common lens as highlighted in Fig. 2.18(b). However, once the parts are rotated relative to one another, the overall focusing power, F_{all} , is the sum of the adaptive Fresnel lens, F_{AF} , and the added lens focal length, F_0 , mathematically expressed as

$$\frac{1}{F_{all}} = \frac{1}{F_{AF}} + \frac{1}{F_0}, \quad (2.45)$$

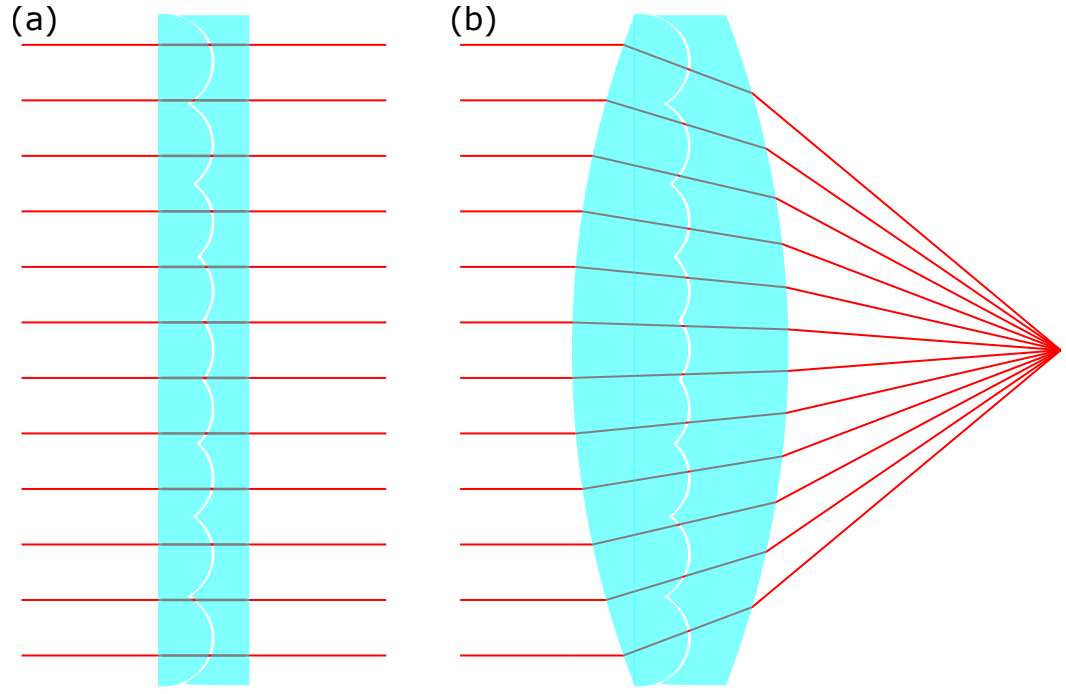


Figure 2.18: An extension to allow the adaptive Fresnel focal length to work best around non-zero focusing power. In (a), a cross-section of two complementary lenticular arrays is shown, which has a focusing power of zero, demonstrated by red light rays passing through it undeviated. By adding two lens surfaces on the flat regions of the lens, the same complementary setup now has a non-zero focusing power indicated by the red rays meeting at a focal point after transmission.

or, in terms of the corresponding focusing powers,

$$P_{all} = P_{AF} + P_0. \quad (2.46)$$

This is reflected in the ray-trace simulations shown in Fig. 2.19, demonstrating the added range and functionality around non-zero focusing power. In fact, Eqn (2.45) makes the overall result even better, as there were some previous restrictions to winding focusing by separation, which only permitted positive focal length winding focusing, as a negative adaptive Fresnel lens would have required a negative separation, which is unphysical. Therefore, using this method allows for the creation of a lens that works best around non-zero focusing power and permits negative focusing powers even in the case of separation winding focusing, significantly extending the operating range. From Eqn (2.45), it can then be seen that the exact range of the combined adaptive Fresnel lens is between just $1/F_0$ and $1/F_{all}$, which depends solely on the range achieved by the adaptive part of the lens. Theoretically, the range is restricted by the relative rotation angle, and the cylindrical lens spiral parameters, which currently can take any non-zero value. However, in reality, both the rotation angle and lens parameters are further restricted by various considerations.

The main issue with a refractive realization is the clear aperture that each winding has.




Figure 2.19: A demonstration of the added focusing power. In (a), an adaptive Fresnel lens is shown without any relative rotation angle, where the component was previously doing nothing, but now with the addition of two lens surfaces as shown in Fig. 2.18(b). Therefore, the view seen is significantly changed relative to the setup before the additional focusing power shown in Fig. 2.7(b). Furthermore, by now rotating the components relative to one another, the focusing power can be adaptively increased further, demonstrated in (b), where the Tim head appears closer (and hence larger) than in (a).

The exact width of each winding can be approximated by taking the difference in the radial distance between two neighboring nodal lines. (Details depend on how the boundary between neighboring windings is chosen.) In terms of angles, that means they are at a radial distance of $R(\phi)$ and $R(\phi + 2\pi)$, respectively. For the logarithmic spiral, this can be approximated by

$$w \approx 2\pi bR, \quad (2.47)$$

where w is the winding width at radius R . The exact expressions vary depending on the spiral type and the chosen boundary between neighboring windings. Additionally, this expression was a simplification made to the original condition outlined in the supplemental material of [64]. It was assumed that neighboring windings meet exactly halfway, which does not create issues for the logarithmic spiral lens as the focal length is constant; hence, at the halfway points, neighboring windings should have the same phase, demonstrated in in Fig. 2.20(a). This is no longer the case for the Archimedean or hyperbolic spiral, where the phase will not be the same, creating a phase jump that physically would correspond to a step in the refractive material, shown in Fig. 2.20(b) and (c) for the hyperbolic and Archimedean spirals respectively. To avoid such steps, a condition that matches the phase at the intersection can be used to define the winding boundaries.

The approximation taken for this is sufficient to give an insight into the rough feasible parameters of the adaptive Fresnel lenses when considered for use in spectacles. Specifically, the range of b and f parameters that may be used to create a cylindrical lens spiral that fills all windings for the entirety of the lens when considering the size of typical spectacles. We assume that a normal pair of spectacle lenses is about 0.05 m in diameter. Therefore, at the minimum, the

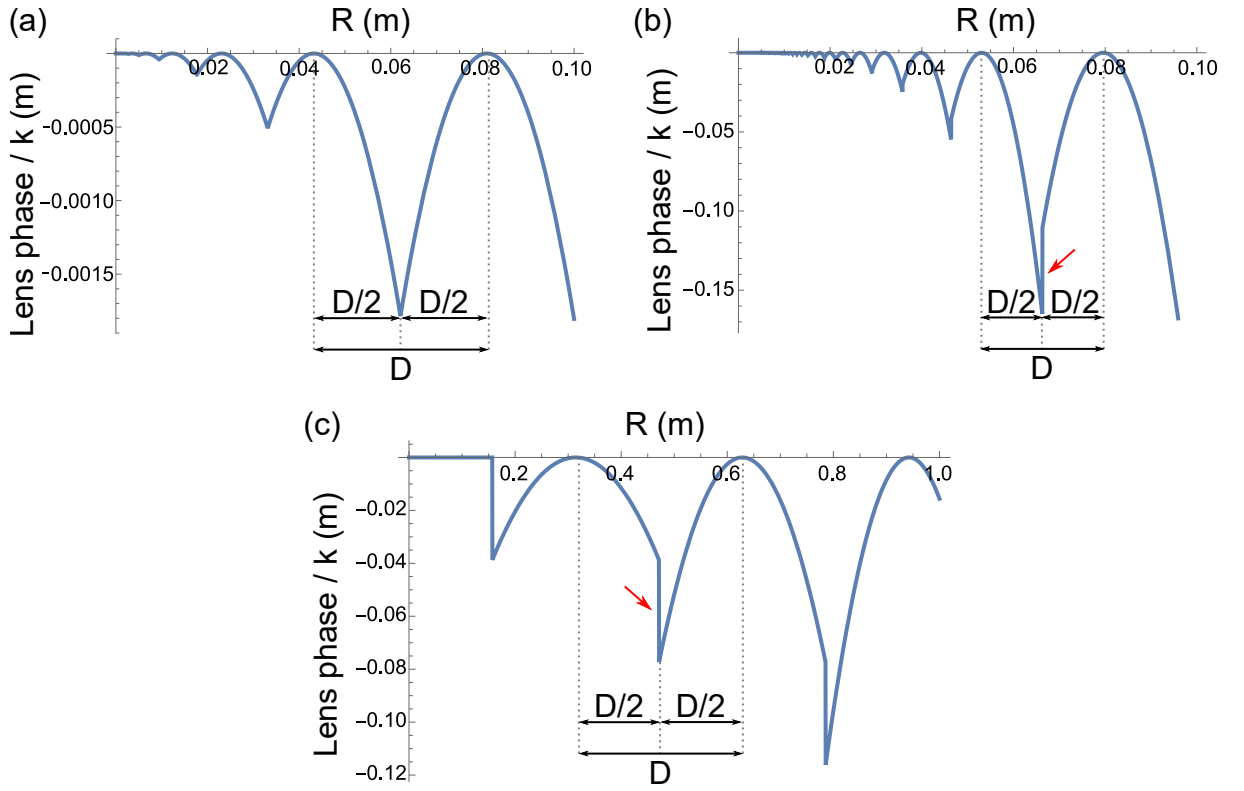


Figure 2.20: A phase profile plot of one part of a spiral lens. In all cases, the phase profile divided by the wave number, k , is plotted as a function of the spiral radius, R . In (a), a logarithmic spiral is used with a cylindrical lens focal length and a winding parameter of 0.1. The same cylindrical lens focal length and winding parameter are used in (b), but now for a hyperbolic spiral. Lastly, in (c), the same cylindrical lens focal length is used, but with a winding parameter of 0.05 and an Archimedean spiral. In (a), the individual windings, shown as blue bumps, meet exactly halfway with the same phase, indicated by two neighboring cylindrical lens centers being separated by D , while they meet at a distance of $D/2$ from each other. The same halfway meeting point is used in the other spiral cases, (b) and (c), but in these cases the phases of neighboring windings no longer have the same value at this cutoff point. As a result, both (b) and (c) contain numerous discontinuities (one per winding cutoff), where one such discontinuity is indicated by a red arrow. These may be resolved by constructing the spirals with a different winding cutoff point.

cylindrical lens used must fill the largest winding sitting a radius of about $R_{max} = 0.025$ m from the center. In fact, simply taking the limit where a refractive component can fill the aperture is not sufficient as lenses made of commonly used materials become thick and behave worse at increased surface curvature. Both of these are reflected in a lens's f-number (N), which is the ratio between the focal length and the width of the lens. Therefore, a larger f-number results in a thinner lens, with an infinite f-number simply being a flat piece of glass. The exact limit where the component works best cannot easily be determined, and therefore, we make an estimated guess based on the f-number of the human eye, setting our limit to be $N \geq 2$. Combining this limit with the maximum winding width in Eqn (2.47), then yields a minimum cylindrical lens focal length of

$$f_{min} \approx 2\pi b N_{min} R_{max} = 0.1\pi b, \quad (2.48)$$

where the minimum N was considered at a maximum radius R_{max} expected in spectacle lenses. This expression can be used to then find the minimum limit of the adaptive Fresnel lens focal length (or the maximum focusing power) such a cylindrical lens focal length permits by inserting f_{min} into Eqn (2.12), leading to

$$F_{AF} \geq \frac{0.1\pi}{\Delta\phi}, \quad (2.49)$$

To find an approximate range of allowed rotation angles, the fraction of light undergoing winding hopping had to be taken into consideration. From the expression of the winding width, Eqn (2.47), the fraction that will undergo winding hopping is given by $\Delta R/w$, which can be approximated to

$$\frac{\Delta R}{w} \approx \frac{\Delta\phi bR}{2\pi bR} = \frac{\Delta\phi}{2\pi}. \quad (2.50)$$

Therefore, if an adaptive Fresnel lens with at least 90% correctly refracted light rays is desired, the maximum allowed relative rotation is $|\Delta\phi_{max}| = 36^\circ$. Taking this into consideration for the focusing power range, P_{AL} , extracted from Eqn (2.49) then gives

$$|P_{AF}| = \frac{1}{|F_{AF}|} \leq 2, \quad (2.51)$$

Taking a similar consideration for the f-number of the additional lens profile, which is of course non-adaptive, can lead to a range of F_0 values restricted by $|F_0| \geq 0.1$ m. Combining all of this into Eqn (2.45) finally yields the expected range of optical powers permitted by feasible physical constraints and can be expressed by

$$\begin{aligned} -P_0 - P_{AF} &\leq P_{all} \leq P_0 + P_{AF} \\ \Rightarrow -P_0 - 2 &\leq P_{all} \leq P_0 + 2. \end{aligned} \quad (2.52)$$

This range can, when adding the appropriate additional lens profile, vary from -12 to 12 diopters, however the major contributor to this large range is the added focusing power, P_0 which is *non-adaptive*. The adaptive range will always be restricted to ± 2 diopters for the chosen constraints. As an example, an adaptive lens could be constructed to have an added focusing power, P_0 of -5 diopters in which case overall adaptive operating range would be -7 to -3 diopters. A further possible constraint to the adaptive range is that to perform winding focusing through separation, and negative Fresnel focal lengths would require the separation to be negative. This would mean the components have to pass through one another which is not physical. As a result, the adaptive focusing power range, P_{AL} , is further restricted to

$$0 \leq P_{AF} \leq 2, \quad (2.53)$$

when considering separation winding focusing.

These calculations were made under the assumption that the width follows Eqn (2.47), which

only holds for the logarithmic spiral. However, through some nice simplifications, the same final expression can also be obtained for the Archimedean spiral if it is also considered that the windings meet exactly halfway between one another. The only different result can be found for the hyperbolic spiral, where the range is dependent on b . Under the assumption of a relatively small value of b ($b \ll 1$), it yields approximately the same range as the other two spiral types, with a maximum focusing power of 2. Note that this was also derived under the assumption that neighboring windings meet halfway between them. This result follows the same calculations performed for the logarithmic spiral, but replacing the radius by the hyperbolic radius and the expression for the focal length by Eqn (2.40).




Figure 2.21: A ray-traced simulation indicating how an extended aperture size may change the result. To do this, a different ray tracer was used, available at [5]. This ray tracer makes use of the GPU, allowing real-time interactive simulations. A picture of a Tim head was viewed in (a). For a pinhole aperture in (b), after adding the adaptive Fresnel lens, the Tim head appears larger, but the ghost image (mostly showing up as a brown spiral, due to the color of the table and the background in all directions) can also be seen. When now extending the aperture size and adjusting the focusing distance (which is typically done automatically by the eye) to that of the Tim head image location, the ghost image becomes significantly blurred out to the point where it is barely noticeable, shown in (c).

Thus far, all simulations shown have been using a pinhole camera. To make sure the view quality of the adaptive Fresnel lens is not dependent on this, a new simulation was constructed to run on the GPU (graphics processing unit), therefore allowing for real-time ray-tracing. This can be found at [5], and has the advantages of being fully real-time and easily permitting any desired

aperture size. Furthermore, it is also virtual reality compatible, which may, in the future, aid in testing the spectacles before any physical prototyping is done. A comparison between a pinhole figure and an extended aperture is shown in Fig. 2.21. Not only does this show that the image quality does not deteriorate and that the adaptive Fresnel lens works for an extended aperture size, but the ghost image appears to fade away due to the camera being focused on the main image, which results in the ghost image being out of focus. Since in a human eye this focusing is done automatically, the simulation can be taken to (within some boundaries) accurately represent the expected view through the adaptive Fresnel lens, although it is not clear how likely the eye might be to focus on the ghost image instead. The simulation was still carried out with infinitely thin phase holograms, representative but not equivalent to real refractive components.

2.7 Discussion and conclusions

The initial goal of our approach on an adaptive Fresnel lens was to create spectacle lenses that may be adjusted freely by the user. This was mainly aimed at an application in the developing world, where access to an optician for to determine the refractive error correction required is not readily available. It should also allow a user to adjust their spectacles with, typically, worsening eye sight without a need to obtain new suitable spectacles once again. Therefore, the resulting lens should ideally be thin (which is also useful since we are typically working with thin components in our simulations), cheap to manufacture, and easy to use while fitting a form factor that permits it to be attractive in appearance (which happens to be a more important consideration that one might think). Existing solutions fall short on some of these. Building on the first iteration of our adaptive Fresnel lens described in section 2.2, this chapter introduces different winding focusing methods and two different spiral types that can be used to construct an adaptive Fresnel lens following the same principles as the original.

The three different focusing methods described were winding focusing through changing the cylindrical lens surface and adding an Alvarez-Lohmann lens surface to it, winding focusing through separating the two cylindrical lens spirals, and lastly, winding focusing through varying the cylindrical-lens focal length. Alvarez-Lohmann focusing is practical as it works in all spiral cases, although by design, the hyperbolic spiral does not require any focusing for a zero separation.

Focusing through separation is the most practical application for a number of reasons (outlined later in more detail), but it is not easily done for both the hyperbolic and Archimedean spirals since the separation required will contain a radius term, which means the separation required for different windings will not be the same. This is not inherently impossible to solve, but it will complicate the construction of these and may, in some instances, lead to bulky components, an undesirable trait in spectacle lenses.

Winding focusing through adjusting the focal length is practical since it could work for all

spiral types (although here it was only shown for the logarithmic type), but inherently does not solve the issue, since the initial goal was to create a lens with an adjustable focal length. Therefore, using a lens whose focal length must be adjusted with rotation angles defeats the purpose of the adaptive lens idea altogether. However, this consideration can be used in combination with the separation winding focusing to create the combined focusing method in section 2.4.3.

The combined focusing method primarily works best for the logarithmic lens again, as the separation term will not depend on radial distance from the center. However, it bears the advantage that focusing is possible while having two cylindrical lens spirals that are no longer complementary. This provides more flexibility with component separations, which becomes most important when considering the physical form factors of a refractive realization.




Figure 2.22: A demonstration as to why a certain bit of separation is not desired but necessary for a refractive realization of an adaptive Fresnel lens. In (a), two complementary cylindrical lens spirals are shown with a separation of ~ 0 . It can be seen that they do not overlap in any part indicated by the color mixing (or lack thereof). When rotating the same components relative to one another in (b). It can be seen that significant regions of the cylindrical lenses overlap, shown as the darker regions in the 3D schematic and highlighted by a yellow area in the cross-sectional cut. Therefore, the only way in which this can be refractively feasible is for the two components to be separated by some larger distance.

In a refractive version of the cylindrical lens spirals, a zero separation complementary setup cannot be rotated without increasing the separation to some degree. This is due to components overlapping, which has been shown in Fig. 2.22. The precise extra space required will depend on lens and spiral parameters, along with the f-number of the cylindrical lenses used, with a smaller f-number leading to larger separations required. To understand the significance this has on the view quality, a number of separations were shown in Fig. 2.23. These were carried out without

winding focusing for reference. It can be seen that even small separations can cause significant deterioration in the image quality. This is exactly why winding focusing by separation may be preferred as it permits some slight separation that improves the image quality, seen in Fig. 2.23(a), where the separation was close (but not exactly) the winding focusing distance. This is even more so the case when taking the combined focusing method into account, for which any (reasonable) separation is possible provided the focal length has been adjusted accordingly. Since this has been demonstrated in section 2.4.3, the concern about the physical allowance required to realize this refractively is minimized. A more in-depth analysis on the minimum required separation is given in [64], which ultimately leads to an adaptive focusing power range of about $|P_{AF}| \leq 1.26$ this is significantly less than the limit found in section 2.6 simply due to the separation minimization.




Figure 2.23: The relationship between the integral image quality of a complementary setup when adding some separation. In all cases, the focal length was set to $f = 0.1$ with the complementary component having a focal length of $f = -0.1$. In (a), a separation of 1 % of the cylindrical lens focal length was used. It can be seen that this actually looks as good, if not better, than no separation. This is because it approximates the required distance for the separation winding focusing. However, in (b) the separation is already significantly larger than the winding focusing separation at 2 % of the cylindrical lens focal length, and the image quality is significantly worse with the elemental images, especially further from the center, no longer aligning at all. When increasing the separation to just 3 % of the cylindrical lens focal length in (c), the image becomes even worse, as expected, but not desired. This shows that either the parameters should be chosen to allow for a small separation, or the combined winding focusing method outlined in section 2.4.3 has to be used.

Generally, common refractive errors can range from -6 D to about 6 D, with anything outside of this range counting as high myopia or hyperopia. The units used here are diopters, which are equivalent to the focusing power in base units. Based on section 2.6, no single reasonable adaptive Fresnel lens can span the whole range. Based on the range found of 2 D, 6 different configurations are required to span from -6 D to 6 D, where the P_0 term in Eqn (2.45) has to be adjusted. If we consider the same case but without the combined winding focusing, 9 different configurations would be required, with each spanning 1.26 D of the overall range. While it is unfortunate that using feasible parameters yields an adaptive Fresnel lens that cannot solely span the whole range, it may still be a significant improvement over currently available alternatives.

Finally, all of the analysis is primarily based on ray optical ideas and simulations. However, we can see that the very center of the adaptive Fresnel lens consists of narrow windings for both the hyperbolic and logarithmic spiral types. This then inherently raises concerns about diffractive issues along with potential evanescence effects, manifesting by making the combined component act similar to a mirror where light rays are redirected backwards. While this is not touched on in this chapter, it is the subject of a paper in preparation. (A completed draft of the current version is attached as appendix C.) It also demonstrates the resulting point spread functions and highlights further improvements to the adaptive Fresnel lens found when performing the wave optical analysis.

One of the future research directions is to extend the wave optical improvements to our ray trace software to qualitatively see how this impacts the image. Another direction which may be interesting to explore is a possibility of adding a read-out mechanism that can easily transmit a users preferred focusing power, allowing a simplified statistical data collection on how the users eyesight changes over time. While this could be interesting especially for medical studies, there are numerous unexplored challenges this poses from an engineering stand point. The combined winding focusing may also be analyzed further for the Archimedean and hyperbolic spiral to determine a range for which the resulting component complies with the desired form factor. This combined focusing may also benefit from including some flexibility in the cylindrical lens surface, which could be changed to have an Alvarez-Lohmann lens surface as well. Ultimately, this all aids in determining the best possible and most feasible refractive realization of an adaptive Fresnel lens. If all of these yield positive results, the final goal is to prototype such a component to demonstrate the viability of correcting refractive errors in the developing world.

Chapter 3

Closed-loop light-ray trajectories in lens cloaks

3.1 Introduction

Invisibility cloaks have long been a fictitious concept predominantly featured in literature and pop culture. Simply speaking, a “cloak” can be any box that hides an object inside by concealing it from all directions, where an observer can only see the box. This is, of course, not what is typically meant by cloaking in fiction or the field of optics. Instead, a proper cloak should hide an object but also remain invisible itself, preserving light paths upon transmission and appearing as if there was no cloak nor object for an outside observer. In recent times, material advances and the development of transformation optics (TO) have led to several different proper invisibility cloaking mechanisms [66]. The original cloak was proposed by Sir John Pendry [67], who, along with Prof. Ulf Leonhardt, pioneered the field of TO [68]. Generally, TO devices are a consequence of the form invariance of Maxwell’s equations under coordinate transformations. This permits the mapping between physical space and virtual space as a coordinate transform, which is where the “transformation” in TO comes from. The transformation linking the physical space –including the light ray trajectories mapped as required into this space– and the virtual space then yields the equivalent local material properties to construct a TO device. Since the light rays must be mapped accordingly, they can be precisely controlled through TO. This molding of light is what allows TO to become the cornerstone of several TO cloaks. In the case of the Pendry cloak, this transformation distorted the view seen through the device such that a finite space within the device can be seen as a single point when viewed from the outside. A practical realization of this came shortly after using metamaterials [69]. Other cloaking ideas constructed using TO devices exist such as the carpet cloak [70], which places an object below a “carpet” that can make the bump caused by the object look flat.

Narrowing this field down to specifically ray-optical TO devices, we can describe a new breed of cloaks. These are based mainly on the ray optical properties of TO and can therefore

allow wave optically forbidden light ray fields. This type of cloak will then, at least purely ray optically, completely hide an object for some or even all view directions. A famous example of such a cloaking device is the paraxial cloak [71], also known as the Rochester cloak. This is a pure lens construction with four lenses placed to form two sets of telescopes in series. When illuminated with rays parallel to the optical axis, the beam has a waist in the focal plane of both telescopes. When an object is placed into the space created surrounding this wait, which is not intersected by any paraxial light rays and viewed along the paraxial (hence the name) direction, the object appears to be invisible.

This cloak served as the inspiration for developing another lens cloak aimed at extending the viewing directions beyond the simple paraxial case to all possible angles around the cloak. A simplification was made by considering ideal-thin lenses instead of the physical, imperfect, lenses used in demonstrating the paraxial cloak. The resulting cloak was called, at first, an omnidirectional lens [16] and later appropriated to an ideal-lens cloak or a type of “abyss” cloak [6]. By imaging objects placed inside the ideal-lens cloak into an abyss below, it is possible to completely hide the object within for some viewing directions. In general, the omnidirectional lens can be changed into various known cloaks by adjusting the lens parameters, such as a lens-based Pendry cloak, or an ideal lens shifty cloak (explained further in chapter 4). Here, we are interested mainly in the case where it is set up to act as an abyss cloak.

Using two ideal-lens cloaks, an omnidirectional cloak can be created. Other applications, such as an image rotator or wormhole, have been considered, but due to dark areas appearing in certain viewing directions within the cloak, these were not fully successful. We note that the wormhole will be explored further in chapter 4. Why and how these occur was one of the questions posed, which we hope to explore in this chapter. A further question we wish to understand is how a light ray behaves within the cloak, specifically, what would an observer see when placed within the cloak, looking along a cloaked direction?

To answer these, the cloaking mechanisms of the ideal-lens cloak had to be understood and categorized. This was, at first, done for the 2D case in section 3.3, and then with some effort extended to the 3D cloak in section 3.3.1. Using the rules found, an additional novel omnidirectional cloak was theorized and demonstrated in section 3.4, while a clue to one of the questions posed was investigated in section 3.5. To facilitate these sections, a breakdown of the basic concept surrounding the ideal-lens cloak is provided in section 3.2.

3.2 Review of the ideal-lens cloak

To investigate closed-loop trajectories in these ideal-lens cloaks the basic operating principles of an ideal-lens cloak will be described here. It should be noted that this work was pioneered in [16] and [6]. As the name suggests, the ideal-lens cloak (otherwise known as an omnidirectional lens) is a construction based on idealized thin lenses. Unlike typical lenses, these are assumed to be

infinitely thin. Any light ray incident upon an ideal lens will therefore be redirected according to known lens theory but within a single plane. It also does not consider any aberrations or size restrictions based on focal length. More on this can be found in [72], where a further generalization is made to construct a generalized ideal thin lens or “glens” in short. Using ideal lenses, several rules were established that dictate the allowed combinations of lenses for which the system can be considered a transformation optics device. The requirement of unique transformation between the virtual and physical space then leads to the loop-imaging condition and the edge-imaging condition [16]. In short, the loop-imaging condition describes that along any loop within a lens TO device, the lenses encountered must image every position back onto itself, while the edge-imaging condition requires all lenses that meet at an edge to jointly yield the identity mapping. In fact, the former condition applies not only to lens TO devices but to any TO component (including metamaterial TO devices). These are further extended upon in [73] for the most general case of a combination of glenses, which includes combinations of ideal thin lenses as a special case, but for brevity, they will not be discussed further here. Instead, a simplified version of the allowed lens combinations may be derived from the conditions as:

1. Any combination of only 1 or 2 ideal lenses meeting cannot satisfy the conditions unless they possess zero focal power.
2. Any combination of 3 ideal lenses must have their principal points at the intersection where all 3 lenses meet.
3. Any combination of 4 ideal lenses must have their principal points such that they lie along a single straight line.

There exist rules for a larger number of ideal lens combinations in [73], but these are not required for the scope of this chapter. Combining these rules, the simplest possible shape that can be constructed is a triangle in 2D or a tetrahedron in 3D containing two extra internal vertices, as shown in Fig. 3.1. The principal points are also shown, following the required constraints.

Since we assume that the cloak operates as an abyss cloak where the interior is imaged to the outside, the base lens focal length, f_D , has to be set within the

$$0 < f_D \leq \frac{hh_1}{h - h_1} \quad (3.1)$$

range, where h is the height of the ideal-lens cloak and h_1 is the height of the lower inner vertex, both described and shown in 3.1. For all further considerations within the chapter, this was considered to hold true. The remaining focal lengths of the lenses were chosen to follow the loop-imaging and edge-imaging conditions, and therefore jointly create a full lens TO device. One of the properties of this is the maintained ray direction upon leaving the cloak, resulting in any light ray incident on the cloak leaving it again in the same direction and from a position that lies along the original ray trajectory. Such a sample ray is shown in Fig. 3.1(a). This follows

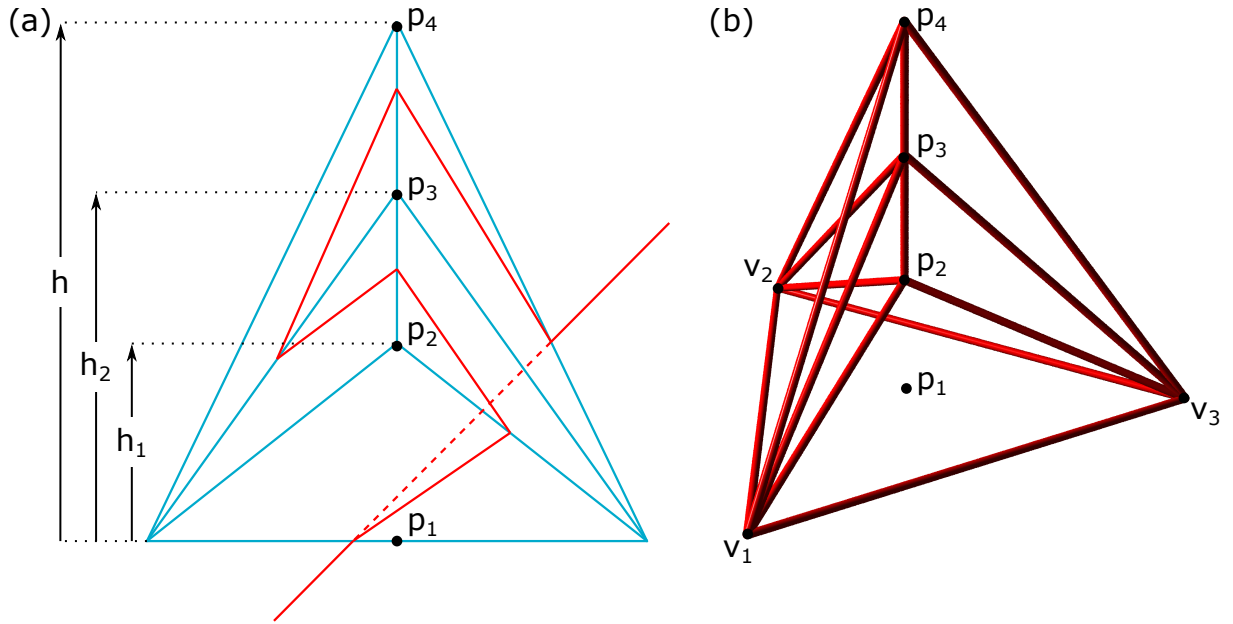


Figure 3.1: A schematic showing a 2D construction of an ideal-lens cloak in (a). Each straight blue line represents an ideal lens. The heights of the inner vertices, shown as h_1 , h_2 , and h , are the vertical distances from the principal point on the bottom lens, p_1 , and the principal points of the inner lenses, p_2 , p_3 , and p_4 . A schematic of a sample light ray trajectory is also shown to pass through the cloak as a red line. As expected of a TO cloak, the ray exits the cloak at the intersection between the outermost lens and the incident light ray, while maintaining the incident light ray direction. This is highlighted with a red dotted line. Within the cloak, such a light ray can travel along a different direction, explored further at a later point. A basic 3D construction of the ideal-lens cloak is seen in (b). Here the red rods represent the net of the tetrahedron. Each polygon bound by 3 of these rods is filled with a lens (although omitted here for illustration purposes). The same principal points as in the 2D case are shown, but now the base vertices of the bottom lens are also shown as v_{1-3} .

on from the fact that the bottom lens, containing p_1 in Fig. 3.1(a), has the same, but opposite in sign, focal length and position as the combination of all the other lenses together. Considering the imaging properties of each lens and the combination of lenses, the image of any object or ray within the cloak can be found at any part. In fact, the use of simplified idealized lenses makes this easier to calculate (and generally possible due to issues outside the scope of this chapter when using physical lenses). Imaging properties, in general, can be considered the underlying method to construct the foundational theorems and, hence, the rules for creating an ideal-lens cloak. Considering these further, an investigation into the underlying cloaking mechanisms was carried out.

3.3 Closed-loop conditions

Building on the previously published work summarized in sections 3.1 and 3.2, we first introduce new nomenclature for the ideal lens cloak. To begin, each of the empty spaces within the lens cloak, here called cells, were labeled as shown in Fig. 3.2. It should be noted that this was done for the 2D consideration only but will be extended later on for a 3D cloak. The boundary between neighboring cells is formed by ideal lenses, which are named according to the cells they separate. The names of the lenses are provided in the same figure. For any point outside of the cloak, there will be a corresponding image within each cell (and the “transformation” from the outside-space position to the image position is the “T” in TO), although some will not be real images. Jointly, these images indicate how a light ray will propagate through the ideal-lens cloak. The image position can be determined through each of the ideal lenses that split the space into individual cells by simply applying the imaging formulae and properties for a single ideal lens.




Figure 3.2: The same 2D ideal-lens cloak as shown in Fig. 3.1(a), but now each lens has been named L_{nm} where n and m are the cells bound by the lenses. These cells are labelled with a cell number, with cell 0 being the outside space, and cells 1-3 as shown in the figure. Within each cell type, there may be multiple cells, which are differentiated through the use of a letter, shown for example as 2a and 2b. A special name is given to the bottom lens, namely the base lens.

The concept of imaging within the ideal-lens cloak was used to first find a special point and then a line within the 2D model. We consider a line along the base lens of the ideal-lens cloak. This line is shown in Fig. 3.3 and can be considered to extend to infinity. A point along this line lying at infinity is then imaged into the ideal-lens cloak. Special interest was taken in cells 2a and 2b, where the image was real simultaneously in each cell i.e. Eqn (3.1) was fulfilled. These points are also shown in Fig. 3.3. Physically, this indicates that any light ray that passes the point at infinity must have an image of a ray that also passes the image of the point within

the ideal-lens cloak cell 2a, but also the real image in 2b.




Figure 3.3: An ideal-lens cloak with an additional line drawn along the base lens. This black line extends to infinity, indicated by I . The real image of I is shown within cells 2 of the cloak as I'_{2a} and I'_{2b} . The exact same images would form for a point at negative infinity. To find the image positions, two rays were taken from I , and passed through lens L_{03b} , shown as a solid red line, to find the first image, which forms specially outside of, but still in relation to, cell 3b and is labeled as I_{3b} . Here the red solid line turns into a red dashed line indicating the ray will no longer intersect another lens as it is already within the correct cell. To now obtain the image within cell 2b, two rays are taken from I_{3b} and passed through to lens L_{2b3b} , which connects cells 3b and 2b. This is again shown by a dashed red line turning in to a solid red line upon the intersection. Finally, these rays now meet at the labeled image position of I'_{2b} , which when passing through the cloak, is also shown to be imaged onto I'_{2a} .

Since the point lies along the base lens line, it can be shown that the light ray passing through it cannot physically cross that line again anywhere else. Specifically, as the ray crosses the plane of the base lens at infinity, it cannot pass through the base lens itself. This becomes very useful since we can show that the image of this base lens in cells 3a and 3b is also real and, importantly, divides the space within the cells into two parts, one part above the image and one below, shown in Fig. 3.4(a). Therefore, any ray, image or not, that crosses this image must also cross the base lens.

Since we have found that our initial sample ray that crosses a point at infinity has an image within the cloak that cannot cross the base lens, we can say that it can also not cross the image of the base lens in cells 3a and 3b. Looking at the enclosed area in Fig. 3.4(b) it then becomes clear that this image ray has to stay within the cloak in all cases. Reversing this, a ray starting within the cloak and traveling along the trajectory of the aforementioned image will stay within the cloak and produce an image in cell 0 along the same direction as the previous initial ray.




Figure 3.4: In (a), the image of the base lens within cells 3 is shown(dashed line). The whole base lens is imaged into this cell, with half of it lying in cell 3a and the other half in 3b. The area bound by the base lens and its images in cells 3 is shaded in, shown in (b).




Figure 3.5: A ray-traced diagram of an ideal-lens cloak. The dark red line within cells 1 and 2 represents a ray traveling in a closed-loop trajectory. The image of this ray in cells 0 and 3 is shown as a dashed red line. Since the image is parallel to the base lens, the intersection with a line along the base lens occurs at infinity. Therefore it can be seen that the ray within cells 2 must pass through I'_{2a} and I'_{2b} , the images of I . The raytracing was done using an open-source software called Dr. Tim [4], with a runnable Java archive file called RayPlay2D available in [3]. Within this chapter, all 2D ray traced scenes were also performed using this software unless otherwise stated.

Focusing on this trapped ray, we need to understand how this light ray can remain within the cloak. A ray optical simulation confirms what we can theoretically explain in Fig. 3.5. It reflects

the construction of the ideal-lens cloak, where the same ray must follow the same trajectory within each cell, a consequence of the edge and loop imaging conditions described in section 3.2. Therefore we can state that the ray must follow a complete closed-loop light-ray trajectory.




Figure 3.6: A modification to Fig. 3.3, with the addition of R , a line that runs along the base lens to and from infinity but excludes the base lens. The images of R are shown within cells 2 as a black dashed line called R' . The images of the base lens are now also shown in both cells 2 and 3. It can be seen that the combination of R and the base lens, imaged into cells 2, split the type 2 cells. Similarly, the image of the base lens in cells 3 splits the type 3 cells.

This argument can be further generalized to not only the point at infinity but to any point on the line along the base lens, provided the point does not lie on the base lens itself, and the image is real in cells 2. For simplicity, a line along the base lens but excluding the base lens will be called the R line (which stands for rest of the line of the base lens but not the base lens itself), shown as the black arrowed lines in Fig. 3.6. Taking the real images of the R line and the base lens within the cloak yields the dashed lines also shown in Fig. 3.6. Using the same arguments made for just the point, it can then be stated that any light ray passing the image of the R line (R' line) in cells 2 should be trapped and follow a closed-loop trajectory. This is confirmed again using a ray optical simulation in Fig. 3.7 where it can be seen that any ray incident on the R' line will remain trapped, while any of the rays missing this line leave the cloak *through* the base lens.

Another interpretation of this may result in the opposite situation —namely which light rays may never be trapped. Following on from the basic principle that light rays must be reversible, any light ray that enters the cloak through the base lens, must be reversible and also be able to leave through the base lens. Therefore, such a ray cannot form a closed loop light ray trajectory, as this would directly contradict the reversibility condition, and has to leave the cloak once again. Since this is a TO construction, the ray is limited further to only being allowed to cross the base lens once as the device has to restore the initial light ray direction upon leaving the




Figure 3.7: A ray-traced simulation of an ideal-lens cloak, showing multiple rays in both red and green. The red lines represent rays that follow a closed-loop trajectory while the green rays do not. The rays were started in the center of cell 1. Upon closer inspection, it can be seen that all the red rays pass through R' , shown as a black dashed line, while the rays that escape in green pass through the base lens, and its images which are shown as blue dashed lines.

cloak. Furthermore, any ray which should enter into cell 1 from the outside has to then pass through the base lens and its image in cells 3, which then dictates how such a ray has to pass through the construction, and as a consequence, it can be shown that even the image of this ray in other cells cannot form a closed loop as there are insufficient “free” cells. This can be seen in Fig. 3.8(a), where a sample trajectory is shown passing through the cloak, intersecting the minimum number of cells. The intersected cells are shaded in to indicate the used cells. Since there is only one unused cell it is insufficient to create a closed-loop for the image of this ray. In fact, the minimum number of cells required for a closed-loop to form is 3 when it includes cell 1 and 4 for all other cells. This comes as a consequence of light rays only changing direction upon intersecting a lens and because each light ray can only travel within a cell once unless it is along the same direction (and position). Generally, the minimum number of cells a light ray passes after entering (or leaving) through the base lens can be found by considering the images of the base lens more closely.

The base lens image in cells 2 is split such that each half is in only either cell 2a or cell 2b and will always appear to be across the central symmetry axis. In cells 3 each half of the base lens is imaged exclusively to cell 3a or 3b. Unlike the images in cells 2, here the base lens is not across the central symmetry axis. All of this is shown in Fig. 3.8(b). Utilizing this, a sample

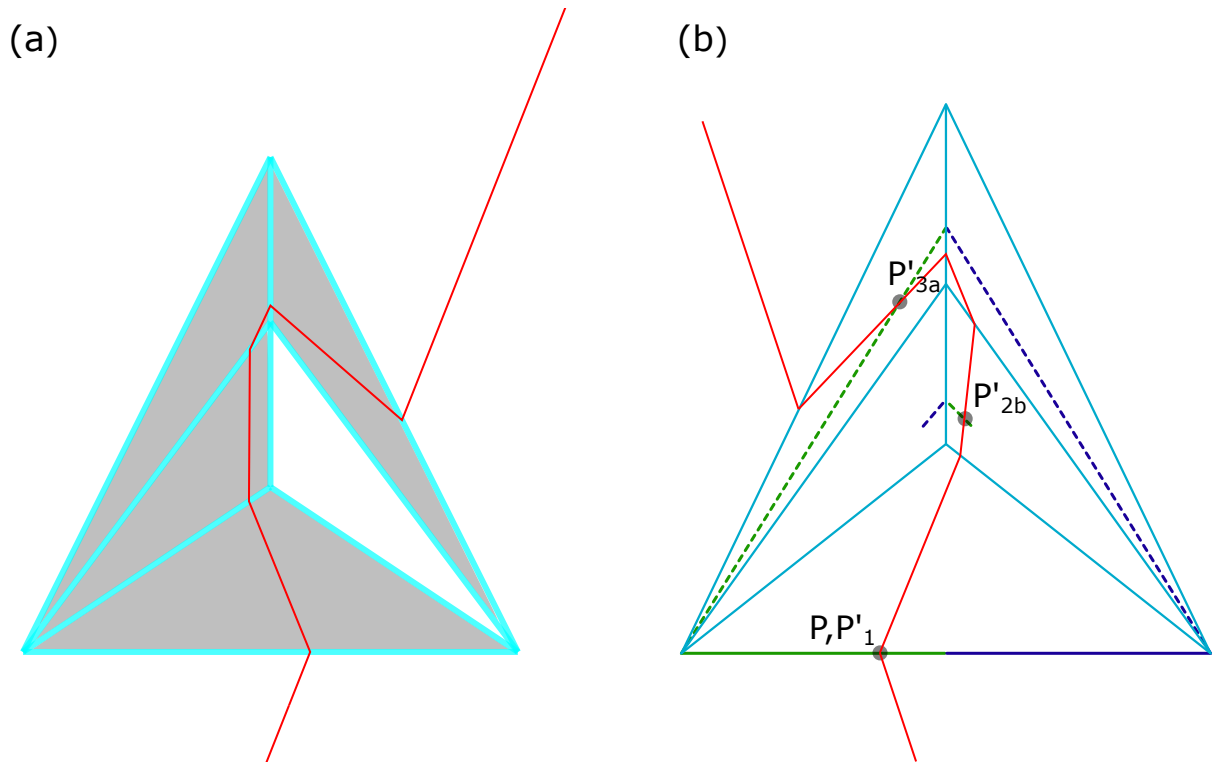


Figure 3.8: An ideal-lens cloak with a ray traced through it in (a). The cells that a ray passes may only be used once due to the loop imaging condition unless the ray travels along the same trajectory again. Any cell which the ray passes through can be thought of as used up, these are highlighted in grey. In this drawing, the ray intersects the fewest number of cells possible for any ray entering the base lens for this particular ideal-lens cloak. This stems from the base lens image positions, with an artistic representation shown in (b) as green and blue dashed lines. It can be seen that the base lens is split up into two halves, each being imaged into cells 2 and cells 3 once in its entirety. A ray passing the base lens at P must then pass all the images of P namely P'_1 , P'_{2b} , and P'_{3a} along its trajectory through the cloak, or have a real image that can pass these. The latter becomes impossible for a ray passing the base lens as it would then inhibit the real ray from passing through the cloak, as it would be required to cross the image of the base lens at a different point which in turn would correspond to a different point on the base lens. Traveling around the base lens is also not permitted due to the R' line, shown in Figs 3.6 and 3.7, causing it to travel in a closed-loop trajectory.

ray can be drawn by considering the image position of an arbitrary point on the base lens. A ray passing this point will first have to enter into cell 1. Therefore, no image of the ray can exist in cell 1, and we can consider it “used”. This ray also has to pass all the real images of the same point leading it to cross the symmetry axis, to the image in cells 2. The simplest case, in which the fewest cells are used, can be seen in Fig. 3.8(b), using either cell 2a or 2b. The ray then has to cross the symmetry axis once again to return to the half of the base lens where it was initially incident upon. Since we assume that it cannot pass through a vertex exactly it has to pass through one extra cell to reach a lens which permits it to finally enter the cell 3 containing its image. It should be noted that the vertex assumption is based on the idealization of the cloak

as a whole, consisting of infinitely thin idealized lenses, as a result, the vertex is a singular point. Therefore another two cells are used up, leading to a total of 5 used cells in the simplest case, which only leaves a single unused cell. Since a single cell is insufficient to produce closed ray trajectories no ray or its real image can exist within this cell, ultimately confirming what the simulations have shown in Fig. 3.8(a) for any ray.




Figure 3.9: The same ideal-lens cloak is shown in both (a) and (b). In (a), a ray is traced through lens L_{3a} , and leaves the cloak through L_{3b} . It can be seen that it never crosses into the area bound by the base lens and its images in cells 3, here shown as the same bounded area previously described in Fig. 3.4. Instead, a real image of this ray forms within the cloak. The image can be seen to travel in a closed-loop trajectory presented as a red dashed line. The same holds in the inverse case, with a real ray starting within the cloak in (b), and traveling along the same closed-loop trajectory as shown in (a). The image of this ray is then shown to be the exact same as the real ray in (a), demonstrating that we can interchange the real ray with the image freely.

Looking at a ray entering the cloak from the other side, any light ray that passes into the cloak through either lenses L_{03a} or L_{03b} and leaves through any lens other than the base lens cannot enter into cells 1 and 2 (and some parts of cells 3 since the image of the base lens splits this cell). This is again the shaded region in Fig. 3.4(b). In cell 0 such a ray must then pass the R line at some point, if it is parallel to the R line this would be at infinity. The ray must therefore always have an image within the shaded region that travels in a closed-loop trajectory. The same can be said and inverted for a ray that starts within the same shaded region. Namely, such a ray has to travel in a closed-loop trajectory and have an image that cannot be close-looped and travel in cell 0, crossing the R line. Taking the shaded region from Fig. 3.4 and overlaying it onto a ray trace simulation shows both of these in Fig. 3.9 where both the ray and its image are shown.

With these arguments, we can then create a number of rules and statements about the ideal leans cloak that dictate if any light ray (and its image) will travel in a closed-loop within the




Figure 3.10: A schematic of all the most important lines, lenses, and cells in an ideal-lens cloak, used to derive a set of rules. The images of the base lens, now called D , and the line along the base lens but excluding it called R are shown as the primed version within the cloak as blue and black dashed lines respectively. Combining D and R yields B , a line spanning the entirety of the space along the base lens. The cells are labeled as previously described in Fig. 3.2. This cloak schematic can be referred to in conjunction with the definition of each used letter in Table 3.1 and the rules described in the main text.

Shorthand	Description
D	The bottom lens on an ideal-lens cloak, sometimes also called the base lens.
D'	The real image of D within the ideal-lens cloak
R	A line that runs along and parallel to D , but excludes D .
R'	The images of R within the ideal-lens cloak.
B	The combined line of R and D , yielding a continuous line along and parallel to D .
I	A point that lies on B (and R) at an infinite distance away.

Table 3.1: The shorthand notation used throughout this chapter, specifically, for Fig. 3.10 and the rules involving the closed-loop light-ray trajectories.

cloak. Furthermore, the cloaking mechanism of this ideal-lens cloak can be explained using the same rules. A diagram showing the most important lines and images is shown in Fig. 3.10 to which all the rules can be related. A notable change is the use of “ D ” as the base lens for brevity. To summarize, of each symbol used can be seen in Table 3.1. The rules found are then as follows.

1. All rays have to pass through the B line. If they are parallel to the B line, the intersection occurs at I , an infinite distance away.
2. Any ray (image or not) that passes through the R' line must always be close-looped.

3. Any ray that is close-looped within the cloak passes through the R' line in cells 2. Such a ray will always have an image that crosses the corresponding point on the R line in cell 0 and is not close-looped. The opposite is also true where any ray that crosses the R line in cell 0 will have an image within the cloak that is close-looped and crosses the R' line in cells 2.
4. Any ray that crosses D cannot be close-looped and will always cross the image of D , D' , in cells 3.
5. If a ray crosses D it cannot have a real image that forms a closed loop, which consequently also means it cannot have a real image anywhere (as there are insufficient free cells in the 2D case).
6. Any object placed within cell 0 can only be visible when along a line of sight crossing D and the object. The cloaked area can be expanded to a region bound by D'_3 and D , which will only be visible by a ray also crossing D .

These arguments are based on the assumption that the images of the base lens and the R lines shown in Fig. 3.10 as R' and D' are real, which is automatically fulfilled when operating within the abyss cloak range highlighted in Eqn (3.1). It is also assumed that the *whole* ray in cell 0 will be imaged into the cloak. This follows from an argument involving the points at an infinite distance away being imaged into the cloak. The images of any point at infinity lie along the focal plane of the base lens within cell 1. In cells 2, they lie along a line connecting I' and the intersection between the aforementioned focal plane and the L_{12} lenses. This is shown in Fig. 3.11 and can be seen to form a triangle contained within cells 2 and 1. Since every ray in cell 0 travels to $\pm\infty$ it must also cross the images of infinity within the cloak. Moreover, since the ray is continuously imaged from $\pm\infty$ the entire ray must always be imaged into the cloak.

3.3.1 3D generalization

All of the rules regarding the closed-loop conditions can be shown to generalize into 3D, provided the imaging remains the same i.e. Eqn (3.1) holds. To start, rule 1 can simply be extended to change the B line to the B plane, spanning all of the space along the base lens of a 3D cloak. This must then be checked further to ensure the whole B plane is also imaged into the cloak.

A 3D raytracing model was used in Dr. Tim [4] where the points at infinity along with the base vertices (v_{1-3} in Fig. 3.1) of the cloak were imaged into each of the cells. The half-way points between the base vertices were also imaged in the same way to construct the complete image of the R plane (which follows the same law as in the 2D case, consisting of all of B but excluding the base lens) within cells 2. When rendering this for just one of the cell 2 cells, it can be seen, in Fig. 3.12, that $2/3$ of the R plane are imaged into just one of the type 2 cells. Therefore, considering all 3 type 2 cells, each point on the R plane must be imaged into cells 2




Figure 3.11: A drawing of an ideal-lens cloak containing the image of every position infinitely far away from the cloak in all directions as a green dashed line. In cell 1 this line lies on the focal plane (shown as a black dotted line) as the base lens will image anything within the focal plane to infinity. In cells 2a and 2b, the line crosses the I'_{2a} and I'_{2b} points respectively which represent the points at infinity along the base lens as previously described in Fig. 3.6.




Figure 3.12: The image of an R plane section within one type 2 cell of the ideal-lens cloak. A wireframe model of the cloak is shown with the corresponding R plane parts imaged inside of it. Note that we only show a part of it for clarity, but the other unshaded regions are imaged into the cloak as well. The light red part of the R plane is imaged onto the light red part within the cloak while the dark red part is imaged onto the dark red part inside. In (a) the cloak is viewed from an angle where the location of the images can be seen to lie within a cell of type 2, while (b) is a top-down version of the same scene. The top-down version may be used to demonstrate that the section shown is exactly $2/3$ of the whole R plane being imaged into just one type 2 cell.

twice, much the same as in the 2D consideration where it is now no longer the R' line but the R' surface. In the same manner, this can also be used to determine the light ray path within the cloak.

Imaging now the other part of the B plane, namely the base lens, which we call the D triangle, yields Fig. 3.13 for the images in cells 2 and Fig. 3.14 for cells 3. The equivalent 2D behavior of the images in cells 3 can be seen where each part of the base lens splits the corresponding type 3 cell into two parts. Therefore it also contains the same space within the cloak where a ray can




Figure 3.13: In (a), the base lens is split up into 3 sections, each imaged into the corresponding cells 2, with the colors indicating the respective positions. The side view shows that the base lens is indeed imaged into the type 2 cells while the top-down view in (b) may be used to show the symmetry of the images within the cloak. Each image appears across the central symmetry axis from the corresponding position on the base lens. Note that in (b), the colors in cells 2 are representative of the base lens colors, and due to color mixing and scene illumination, the corresponding red part appears dimmer.

only enter when it passes through the base lens using the same arguments made in the 2D case.




Figure 3.14: (a) top-down version of the colored-in base lens. Unlike the segmentation shown in Fig. 3.13, this is now split up to be fully in a single-letter cell when viewed from the top-down i.e. the top-down view of the A cells intersects only the red portion of the base lens. The image of just one of these sections is shown in (b), where it can be seen to be fully imaged into a single type 3 cell. The parts not shown were omitted to permit a view into the inside of the wireframe cloak model.

Ultimately, this can then be used to extend rules 2 and 3 to the 3D world. Similarly, since all of the D plane is imaged into cells 3 (and cells 2), rules number 4 and 6 can also be confirmed. An example of this “window” into the cloak, described by rule 6, is shown in Fig. 3.15.




Figure 3.15: A visual representation of the special significance associated with the base lens. In (a), a setup of the cloak containing a ball is shown. The ball is placed within cell 1 of the cloak. A 2D schematic is shown in (b), which also contains an estimate to the expected ray directions, with only the rays through the base lens, in red, intersecting the sphere. An eyeball was placed at the specific viewing position for which the rays, shown as the red area, were drawn. Moving the camera to this position and viewing the same setup shown in (a) yields (c). It can be seen that part of the sphere is visible above the base lens and a part below when using the wireframe model of the cloak. Replacing the wireframe model with a cloak made of lenses then results in (d). Here, Part of the sphere can be seen, corresponding to the part that would have been visible through the base lens only. The distortion occurs due to the imaging properties of the cloak.

The last rule that needs to be extended is rule number 5. Unfortunately, this is not as straightforward as the others. In the simplest case for the tetrahedral ideal-lens cloak, containing 3 base vertices as shown in Fig. 3.1(b), a light ray must use at the minimum 4 cells just like in the 2D case and due to the same imaging requirements of the ray having to travel across the central symmetry line (the line connecting the principal points). This originates from the image of the base lens within cells 2 and 3 which can be seen in Figs 3.13 and 3.14. Unlike the 2D case, the central symmetry line is no longer a simple lens. Therefore a light ray can no longer pass through it but must go *around* the symmetry line. The consequences of this are for the simple case not very interesting and ultimately lead to, just like before, the light ray having to use one additional cell which finally results in at least 4 cells being required for any ray passing through the base lens.

From the wireframe diagram in Fig. 3.1(b), a total of 7 cells excluding cell 0 can be counted for the 3D cloak. Therefore, if 4 are used it only leaves 3 free cells for a light ray to perform a closed-loop trajectory. The minimum number required for a closed-loop remains the same as in the 2D scenario where 3 cells are sufficient to form a closed-loop but only if one of these is of cell type 1 (excluding the special case discussed in the lens star section of 3.3.1). For any other closed-loop trajectories a minimum of 4 cells are required. Since the ray passes the base lens it must pass into cell 1 which forbids it from being used again for a closed-loop, and therefore resulting in a lack of free cells to host a closed-loop trajectory.

This works very nicely for the simple tetrahedral ideal-lens cloak but turns more complex when considering that the ideal-lens cloak can be extended to have any number of base vertices. These are significantly harder to analyze conceptually, and worse, no longer necessarily obey the simple argument used of a cloak having insufficient cells to host a closed-loop trajectory.

n base vertex ideal-lens cloaks

To extend the rules to any number of base vertices, a “topological” breakdown of a cloak was carried out. This can be done generally for any number of base vertices and significantly simplifies conceptualizing such cloaks. It also helps in demonstrating how a ray passes through a 3D cloak in 2D, which becomes messy when using a simple 3D to 2D projection. Since this effectively reduced planes to lines, we note that the light ray behavior will no longer need to appear physical.

For purposes outlined in a later part of this section, a cloak built using 8 base vertices was chosen. The 2D topological reduction can be seen in Fig. 3.16, where each lens and cell is labeled according to the 2D cloak model. Each lens and cell retains the relative position within the cloak, but the dimensions are reduced, whereby a (volume) cell is transformed into a 2D area (which we will also call cells for simplicity), and a lens is changed into a line in 2D. The lines, just like the lenses in 3D, split neighboring cells. The space outside the cloak, known as cell 0, is represented as a circular area in the center and the area outside.

The relevant images can also be shown in this reduction and just like before these dictate how a ray has to pass through the cloak and hence the reduced cloak. Due to the topological reduction used all rays will be a schematic ray trajectory and do not follow any strict ray tracing criteria. We note that cell 1 will behave slightly differently from the other cells. Namely, we allow a ray to curve “freely” in cell 1, with the only constraint being that it cannot remain just between cell 0 and cell 1. This curving is a consequence of the initial incident ray direction onto the base lens. Regardless of position, a ray incident on the base lens, D, in Fig. 3.16 may enter any other type 2 cell simply due to the freedom of incident direction. Since cell 0 has to be represented in the center, this means it may enter into cell 1 at the top (nearest to cell 2a) part and then continue into cell 2e. In the reduced form this would appear as though it travels back through the base lens unless ray curving is allowed within this cell. Considering this, a simple




Figure 3.16: A topological reduction of a 3D, $n = 8$ base vertex, cloak into a 2D schematic. Each blue straight line represents a lens labeled by the corresponding lens name. The central blue circle also represents a lens, but this is the special case of the base lens, D. Within the center and the outside of the cloak there are two cells labeled as 0. These represent the space outside of the cloak, with the central enclosed region representing the space below the base lens of the cloak, specifically the space in line with the base lens. Therefore any ray entering through the base lens must first travel in the enclosed 0 cell space. All the other cells are labeled in the same manner as for the simple 2D cloak, with letters differentiating between the same type of cell. This breakdown will be used to investigate the allowed ray trajectories and resulting rules in a simpler way.

ray trajectory was considered, where a ray passed through the base lens into the cloak and then out into cell 0 while intersecting the minimum number of cells, shown in Fig. 3.17.

It may appear as though there is a more direct path, however, the imaging properties do not allow the direct path passing from cell 1 to cell 0 simply through cells 2d and 3d. This can be seen clearly when considering the images of the base lens section outlined in Fig. 3.17 as a green-dotted line through which the ray must pass. The only alternative to this would be that the real image of the ray passes the images of the base lens instead. Since a ray entering through the cloak must leave the cloak along the same direction and position for it to be a TO device, it has to pass through to cells 3 and into cell 0 again in some way. When considering all of the relevant images in cells 2 and 3, it can be seen that the space is in fact completely split by the complete image of the base lens and R plane. This is shown in Fig. 3.18 as the dotted and dashed lines. It cannot pass through the R' line as any such ray must be closed-looped. The only other way is to pass through an image of the base lens, D'_2 . However since each point on the base lens is only imaged into cells 2 once, it must be the real ray that passes this point and not an image




Figure 3.17: An example of the simplest ray-trajectory when drawn in the topological reduction of an $n = 8$ base vertex cloak. This is the simplest as it crosses the base lens and intersects the least number of cells, where a solid red line denotes the ray trajectory. The longest green-dotted line represents the image of the base lens that it crossed within cell 3, the same as D'_3 shown previously. As the ray crosses this, it must then also cross the corresponding image in cell 2, shown as a shorter green dotted line. To reach this, the ray may travel only within either the type 2 cells or, as shown here, within the type 3 cells. The final image it will then have to cross is the corresponding region on the base lens, given by a curved green dotted line before leaving the cloak into cell 0 as shown. To permit the light ray to exit the corresponding region of the base lens, we permit it to curve within cell 1.

of the real ray. Otherwise, the real ray would not have any other point in cells 2 through which it can cross toward cells 3. If it did cross D'_2 at another point the corresponding point on the base lens would also need to be crossed leading to an additional base lens intersection which is impossible for the real ray or its image since both cell 1 and 0 are already used by the real ray.

We consider the minimum number of cells a ray passing through the base lens must cross. In the case of the 8 base vertex ideal-lens cloak, this is 7. Generalizing this, it can be found that the cloak must always use at least half of the cells as the images in cell 2 and cell 3 will always be directly across the central symmetry axis. In the case of an odd number of base vertices, this still holds but is no longer demonstrated as nicely since the image in cells 2 will be split across two cells with each imaging part of the base lens section just as shown in the 3 base vertex case in Figs 3.13 and 3.14. In addition to using half of the cloak's type 2 or type 3 cells, a ray must also enter into the cloak through the base lens and hence use cell 1 and at the minimum use one additional cell (either of type 2 or 3, as it must pass to the outermost layer of the cloak before it can leave through a type 3 cell). Ultimately this yields the minimum number of cells being used




Figure 3.18: The same reduced cloak as in Figs. 3.16 and 3.17 but now showing the images of all of the base lens, D , and the R plane. The images are given by a black dashed line and a green-dotted line for the R plane and base lens respectively. It can be seen that the combination of the two images completely split the cloak into 3 regions, one below the combination of the R'_2 and D'_2 lines, a region bound between this and the D'_3 lines, and lastly the region above the D'_3 lines. These images dictate how a ray may travel through the cloak.

by such a ray as $n_{\text{even}}/2 + 3$, where n_{even} is the even number of base vertices of the cloak. For odd numbers of base vertices, n_{odd} , this is given by $(n_{\text{odd}} + 1)/2 + 2$ which can be shown using a similar diagram.

Investigating the minimum number of cells used by a closed-loop yields a similar result. It is assumed that the ray must cross the image of the B plane (defined as the combination of the images of R and the base lens in Fig. 3.18). This leads on from the fact that the ray in cell 0 must travel above and below the B plane at some point. If it is parallel to the B plane this occurs at infinity. Therefore it is impossible to form a closed loop entirely in the space bound by the image of the base lens in cells 3 and the B' plane in cells 2 as any point further out (radially) from the B' image will be imaged to below the cloak and so will any point radially closer than the base lens image in cell 3. If a ray crosses the image of the base lens in cell 3 it must escape to the outside and cannot be close-looped, therefore it *must* cross the B' in cells 2 at some position.

The position at which it crosses can, for the moment, be anywhere. However, since the ray must return to the same position again to form a closed-loop it must cross B' twice. This can only happen if the part of B' crossed is the image of R in cell 2. If it were to cross at any of the base lens images it would be required to cross the same point again in the opposite direction as each point on the base lens is only imaged into cells 2 once. This is of course not allowed from the

basic construction of the cloak relating to the edge and loop imaging conditions. Considering the ray crosses a point on the R' line in cells 2, it must also cross the second image of the same point. This second image once again lies across the central symmetry line and therefore dictates the minimum number of cells used. If the use of cell 1 is permitted, this means at least $n_{\text{even}}/2 + 2$ or $(n_{\text{odd}} - 1)/2 + 2$ cells must be used. Since cell 1 is typically used already, the minimum number of cells used is increased and can be given by $n_{\text{even}} + 2$ or $n_{\text{odd}} + 1$. This can be achieved by traveling through cells 2 and 3 demonstrated in Fig. 3.19(b).

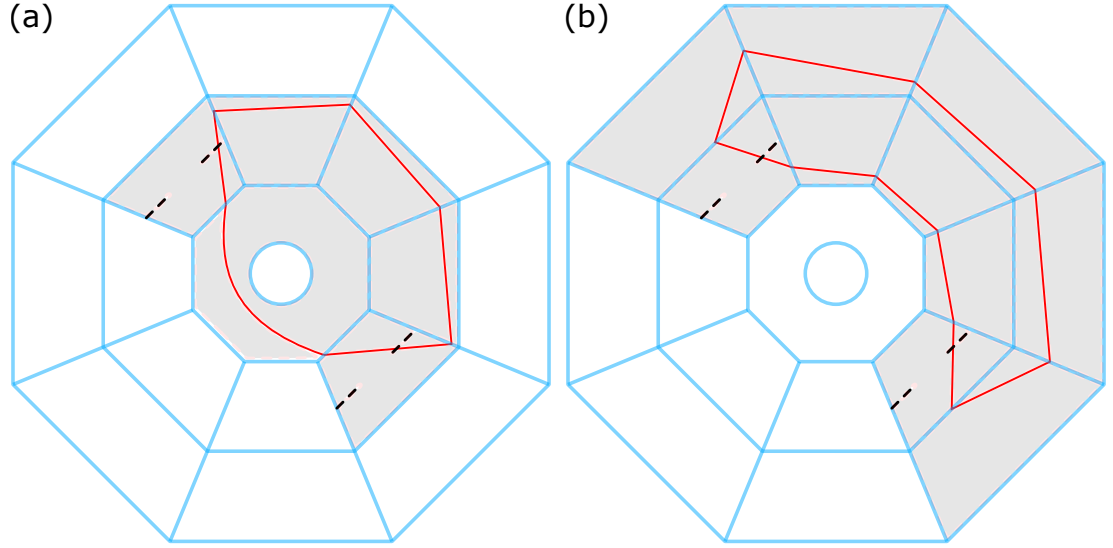


Figure 3.19: The possible closed-loops within an $n = 8$ ideal-lens cloak, shown in the topological reduced version. In (a), the case using the least number of cells using cell 1 as part of it. If cell 1 cannot be used, the closed-loop must form entirely within cells 2 and 3. This is shown in (b) where it can be seen to use all of the greyed-out cells in the process. In both cases, the images of R'_2 in the relevant cells are shown as black dashed lines, indicating what cell the ray must travel to.

We note two additional cases, a special occurrence of “lens stars” where a closed-loop travels entirely in cells 2 and hence uses n cells, discussed further in the lens star section of 3.3.1, and the possibilities that multiple closed-loops form to intersect the required images. When considering two separated closed-loop trajectories, the minimum number of cells used will always be 8, provided cell 1 cannot be used, with each loop using 4 cells and intersecting R' twice, such that the 4 total intersections across both loops are imaged onto one another.

The cloak has a total of $2n + 1$ cells, where n is simply the number of base vertices regardless of parity, and therefore the minimum number of base vertices for which the simple rule of insufficient free cells can apply is only valid for $n_{\text{even}} < 8$ and $n_{\text{odd}} < 5$. This agrees well with Fig. 3.19, where an 8-base vertex cloak is shown. It can be seen that the minimum number of cells used by a ray passing the base lens is 7, leaving 10 unused cells. As the minimum number of cells necessary for a closed-loop is 10 (or in the special double loop case 8), rule 5 no longer holds and permits the formation of closed-loop trajectories.

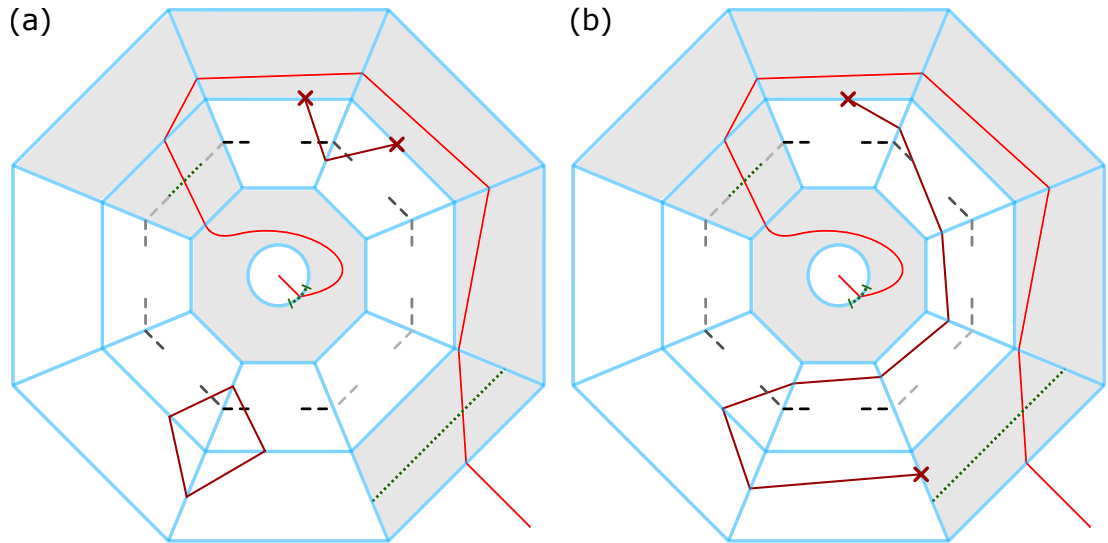


Figure 3.20: An example of an attempted closed-loop trajectory which is hindered by the used cell location rather than the available cell number. In both cases, there is a light ray passing through the cloak shown as a light red line, the same ray which has been shown in Fig. 3.17. Additionally, a dark red ray can be seen indicating a possible image of the light red ray. In (a), one of the special cases is shown whereby R'_2 is crossed in 4 points indicated by a dark red ray crossing the black and gray dotted line. These dark red lines would then be a connected pair of closed loops which are imaged into one another. One of the loops could successfully form in the half of the cloak containing the free cells. The second loop cannot occur as there are no neighboring free cells indicated by the red cross where the dark red ray has to meet the used cells. A similar scene can be seen in (b), but instead, an attempt was made to produce a closed-loop as shown in Fig. 3.19(b). To complete the loop, the ray would once again be required to travel into a used cell, indicated by a red cross. These regions are unavoidable regardless of trajectory choice simply through the constraint that the pair of R'_2 images must be crossed. The corresponding pair of R'_2 lines is indicated by the shade of grey used in the diagram.

A complete generalization may be made when considering a different limiting factor. Namely the available free *neighboring* cells. In any of the closed-loop cases, it has been shown that a ray must cross the R' line at least twice or in the case of a double loop 4 times. The location of these intersections is not random but strictly defined by the cloak imaging properties such that they are always on opposite sides of the cloak. In an odd base vertex cloak, this is still true but the images of R will be split across two type 2 cells on one side of the cloak. Therefore, modifying Fig. 3.17 by adding the image pair of the R plane in the relevant cells, Fig. 3.20 was made. Taking any point on one of these R' lines, it can be seen that either the point itself or the image of it only has a single free neighboring cell that does not lead to a dead end. Since a closed-loop trajectory must enter and leave a cell through two different neighboring cells, either the incident ray comes from the dead end, which is impossible, or the emerging ray would need to travel into the dead end. Therefore such a ray cannot exist. This forms a condition whereby a ray may have sufficient cells to form a closed-loop trajectory on one half of the cloak but not in the corresponding other half. This is illustrated in Fig. 3.20 where a “stuck” ray is shown. The

rule found for the 2D case must then be adjusted and should read:

5. If a ray crosses D it cannot have a real image that forms a closed-loop, which consequently also means it cannot have a real image anywhere. In the case of $n_{\text{even}} < 8$ or $n_{\text{odd}} < 7$ this is due to a lack of free cells while for any larger number of base vertices, n , due to a light ray dead-end condition.

Having shown, in some parts more easily than others, that the 2D rules found in section 3.3 can be directly extended to 3D, the focus can now be shifted to the consequences these rules impart.

Lens stars

A special kind of trajectory was identified whereby a light ray can follow a closed-loop trajectory entirely contained in cells of type 2. The lenses separating neighboring type 2 (and 3) cells form a structure known as a “lens stars” [74]. These consist of at least 3 ideal lenses arranged in a star shape sharing a common edge (or point in 2D). The principal points of the lenses have to follow the edge-imaging condition, which is automatically satisfied in the ideal-lens cloak. Considering only the lenses separating the different type 2 cells, it can be seen that these also share a common edge in the 3D case, therefore forming a lens star.

It should be noted that this is technically also satisfied in the type 3 cells, since their principal point lies on p_4 in Fig. 3.1(b). However, a further property of lens stars is that any ray trajectory passing through the lens star must lie along a conic section in a plane that contains the principal point. In Cells 2 this is satisfied for a number of planes that intersect the lenses in cells 2, with one being shown in Fig. 3.22 as a grey shaded triangle. Such a plane is not possible to construct in the type 3 cells as it cannot intersect all of the lenses that constitute the lens star. Therefore, no *real* lens star closed loop trajectories may occur in the type 3 cells.

Within a lens star, a light ray may follow a close-looped trajectory passing only through the lens-star lenses. Within the cloak, such a trajectory will therefore be entirely contained within cells of type 2. It is possible to show this with a sample trajectory in the topologically reduced, $n = 8$, lens cloak in Fig. 3.21 and a ray trace diagram within a wireframe model of an $n = 3$ ideal-lens cloak in Fig. 3.22, confirming the results.

For this special case of a closed-loop trajectory, the total number of cells used is n , regardless of the base vertex number parity. Therefore, by simply following rule 5 it may seem as though the limit is in fact less than previously stated. In fact, it would mean that the 3D generalization would require a cloak with less than 3 base vertices with is unphysical. Fortunately, there is a simple consideration that can be made to prove that no ray entering through the base lens can have an image that forms a lens star.

When a ray enters the base lens it must then leave the cloak through one of the type 3 cells. Since a lens star only exists within cells 2, it must use *all* of the type 2 cells. Therefore, cell




Figure 3.21: A schematic of a possible ray trace trajectory of a lens star within cells 2 of the ideal-lens cloak. In this case the trajectory can be seen to cross the R'_2 lines in two places across the centre, upholding the imaging conditions outlined. The cells the ray travelled through are greyed out to indicate that these have been used. It can be seen that a lens star trajectory will split the cloak fully, meaning any ray entering through the base lens into cell 1 cannot travel to the outside again.

1 and cells 3 are disconnected, meaning that a ray entering the base lens cannot have an image that forms a closed-loop trajectory within the lens star as it would then no longer be able to pass into cells 3. Ultimately, this consideration holds for any number of base vertices, upholding the previously imposed limit of $n_{\text{even}} < 8$ or $n_{\text{odd}} < 7$.

3.4 Omnidirectional cloaks

It was found that the main cloaking mechanism of the ideal-lens cloak is due to closed-loop trajectories. Therefore, any light source placed inside of the cloak must have some of the rays follow a closed-loop trajectory for the view directions over which it is cloaked. This was used, although through an understanding of the imaging properties, in [6] to create an omnidirectional ideal-lens cloak. As the name suggests, this cloak works for all viewing directions, which in turn must mean that a light source placed within such a cloak will have any emitted ray following along a closed-loop light-ray trajectory. This can be shown in the 2D raytracing case in Fig. 3.23 where an omnidirectional lens cloak was constructed and raytraced. The equivalent 3D realization is shown in [6], where the construction was carried out by misaligning the images of the light source placed within the cloak, such that there can not be any line of sight for which




Figure 3.22: A ray-traced lens star trajectory within an $n = 3$ base lens ideal-lens cloak. In (a), the lens star is identified with the grey triangle forming a possible plane in which the lens-star trajectory can occur. It can be seen that the lenses, shaded in cyan, intersect this plane as required for the formation of a lens star. To confirm the existence of closed-loop trajectories following lens stars, a light ray was placed onto the lens star plane within the cloak in (b). This was then ray traced through the scene, shown as a red line. The ray travels along a closed-loop intersecting each lens once before returning to the starting point as expected. Subsequent round trips follow the same trajectory. In (b), the lenses were not shown for clarity but all lenses were included in the ray trace of the scene.

the physical outer cloak and all of the images of the object (once due to the outer and once due to both cloaks) are aligned. Ultimately this results in the light source being cloaked from all outside positions. The same cloak may be constructed by considering how closed-loop light-ray trajectories are formed within the cloak.

Since the rules of the closed-loop trajectories are now understood, it is possible to describe why this cloak works and create a new omnidirectional ideal-lens cloak. Rule 2 states that the special line to consider is the R' line. Any ray which passes this line must always be close-looped. Considering the location of the R' line in the inner and outer cloak of Fig. 3.23 it can be shown that all the light rays must cross at least one of these R' lines. Therefore, there will be closed-loop light-ray trajectories in both the inner and outer cloaks, such that any rays that escape the inner cloak must be trapped by the outer cloak.

When designing the new, here called the *extreme* ideal-lens cloak, similar conditions to force rays across the R' surfaces need to occur. One way to ensure this happens is by considering only rays that start on the R' surface. Such rays must then be close-looped for all directions, indicating that even the very basic ideal-lens cloak can act as an omnidirectional cloak, but only for a line (or a surface in 3D), as demonstrated in Fig. 3.24. Since we wish to not only cloak a single line or surface but a whole volume of space, this only acts as a starting point for the new extreme ideal-lens cloak.

Rule 6 states that any object within cell 1 is only visible when seen along a line of sight that




Figure 3.23: A construction of the original omnidirectional lens cloak proposed in [6]. In (a), the basic construction and concept of an omnidirectional ideal-lens cloak is shown. An inner cloak is placed within an outer cloak such that it is imaged outside of the outer cloak. The image can be seen as a small dashed cyan triangle below the outer cloak. Within the inner cloak, a point, P , can be seen. As expected, the image of this point due to the outer cloak lies within the image of the inner cloak and is labeled as P'_o . The combination of the cloaks image this point further onto the point P'_{oi} . A red line connecting these two points can be seen to miss the outer cloak. A viewing position and direction along this line has a line of sight through both images of P , but never through the physical cloak. Therefore, it is fully cloaked from these directions. Since any other view direction does not contain a line of sight where the images of P align, P must be hidden for *all* view directions. This is confirmed using a ray trace simulation in (b). It can also be seen that this coincides with the condition of all rays having to cross the R'_2 lines, here given as black dashed lines.

intersects the base lens. Therefore, the opposite must also be true, namely any ray originating in cell 1 can only cross to the outside if it passes through the base lens. Taking two cloaks, an outer and an inner one, it is then possible to place the inner cloak such that its base lens coincides with the R' surface of the outer cloak. Therefore, any ray starting within the inner cloak can only reach the outer cloak through the base lens (as a consequence of rule 6). Any other ray will be trapped within the inner cloak. All the rays that make it out of the inner cloak pass through the base lens and must then also cross the R' surface in the outer cloak, trapping them in a closed-loop within the outer cloak. A schematic of an extreme ideal-lens cloak is shown in Fig. 3.25(a) for the 2D case, while a ray-traced version is shown in (b).

As the same rules apply in the 3D case, the extension into 3D is straightforward with the base lens now being placed onto the R' surface within the outer cloak. A wireframe model of the simplest case for a 3 base vertex cloak (which can be extended to any number of base vertices) is shown in Fig. 3.26(a) with a sphere placed within the inner cloak (for clarity, only one inner cloak and sphere are used, but this can be extended). Viewing the same scene from an angle where the sphere would typically be visible in (b), completes the setup. Now replacing the




Figure 3.24: The simplest version of an omnidirectional cloak found. The basic ideal-lens cloak, shown as a construction of cyan lenses, was found to act as an omnidirectional cloak for a specific line (or surface in 3D). This line is the R'_2 line shown here as a black dashed line in one of the type 2 cells. When a light source is placed onto this line and traced, all rays cross the R'_2 line and as a result follow a closed-loop trajectory. Therefore anything along this line will automatically be cloaked for all directions.




Figure 3.25: The setup for an *extreme* ideal-lens cloak. In (a), an inner cloak can be seen to be placed such that the base lens of the inner cloak overlaps the R'_2 line of the outer cloak here shown as a black dashed line. The resulting structure should therefore, according to the rules found, act as an omnidirectional cloak for any object placed within the cloaking area of the inner cloak. To confirm this, a light source was placed within the inner cloak in (b). The rays were traced and it was found that none of them escaped to the outside as hoped.

wireframe model with the appropriate ideal lenses and ray tracing the scene, it can be seen that the spheres disappear from view, demonstrated in Fig. 3.26(c). Of course, keeping the sphere in the same position and removing the inner cloak in (d) of Fig. 3.26 reveals the sphere confirming that it is indeed the combination of the two cloaks that hides the sphere.




Figure 3.26: The same extreme ideal-lens cloak shown in Fig. 3.25, but now generalized to 3D. In (a), a wireframe model of the setup can be seen where the base lens of the inner red cloak was aligned to lie on top of the R'_2 surface of the outer blue cloak. Within the inner cloak, a sphere is shown. The goal is to hide the sphere from a camera position where it should be visible when using just one cloak. One such position is shown in (b) where the same scene is viewed. When the wireframe model is now replaced with ideal lenses, and the scene is ray traced again, (c) is formed. In (c), it can be seen that the sphere is completely invisible and the overall cloak acts as a TO device with the space behind the cloak appearing undistorted as required. To prove that it was indeed the additional cloak that allowed for the omnidirectional invisibility, a ray-traced simulation of the same scene but removing the inner cloak is shown in (d). Here the sphere is once again visible as expected when only using the outer cloak.

3.5 Attractors

An attractor may be described as any state that a system evolves towards. A straightforward example is a simple damped pendulum swinging that will tend towards a state of zero movement due to friction. It is possible to create similar attractor-like properties within the ideal-lens cloak leading to light ray trapping.

Initially, a simple ray may be passed into the cloak through the base lens. This ray will, under normal circumstances, simply leave the cloak through one of the L_3 lenses. By placing an additional lens in the path of the ray, such that it is redirected through the R' line, it is possible to stop the ray from leaving the cloak. The process behind this is as follows. Initially, the light ray is redirected by the additional ideal thin (trapping) lens. Taking the angle of incidence on the lens as θ_0 and the angle of refraction after transmission through the lens as θ_1 , a formula can be derived from the vectorial imaging equation of an ideal-thin lens (given as (4) in [16]) to relate

the two angles as

$$\tan \theta_1 = \tan \theta_0 + \frac{r}{f}, \quad (3.2)$$

where r is the radial distance from the lens principal point to the ray incidence position on the lens and f is the focal length of the lens. This also holds in 3D provided the plane of incidence intersects the principal point. In those cases, we note that all angles, θ , will always be formed between the ray and a line connecting the incidence position on the lens and the lens principal point. The derivation can be found in a Mathematica document called *AttractorBehaviourCalculations*, provided in [1]. Starting a new ray at the lens intersection with a direction that crosses the R' line, we can assume it must follow a standard closed-loop trajectory in every cell due to the ideal-lens cloak construction conditions. Starting with the assumption that the additional lens is placed in cell 1, this means the ray must then travel back into cell 1 with the same direction, and along the same trajectory as the initial refracted ray leaving from the trapping lens. As a consequence it must then intersect the trapping lens again at the same position as in the first intersection and with an angle of incidence equal to the angle of refraction from the previous round trip. Due to the intersection, the direction will change once again resulting in a new angle of refraction, θ_2 , that follows the same relationship as shown in Eqn (3.2), with θ_0 becoming θ_1 and θ_1 turning into θ_2 . The process will repeat for every round trip, tending to a ray direction that is perpendicular to the trapping lens's optical axis, or parallel to the lens. This aligns with the definition of an attractor trajectory and can also be expressed mathematically whereby iterating Eqn (3.2) n times leads to

$$\tan \theta_n = \tan \theta_0 + \frac{nr}{f}, \quad (3.3)$$

where θ_n is the light ray angle of refraction (after n round trips) and θ_0 is the initial angle of incidence upon the trapping lens. Plotting this it can be seen that the system indeed evolved towards a single state proving the attractor-like behavior demonstrated in Fig. 3.27(a) with the equivalent ray tracing example shown in (b). It should be noted that this expression is generally true for any cell provided it lies within the cloaking region (shaded grey in Fig. 3.4), and has an appropriate trapping lens. A further consequence, and the reason the additional lens is referred to as a trapping lens, is that some light rays which continuously intersect R' after every round trip will be able to enter the cloak from the outside but never escape. Therefore, it may act as an ideal-lens light ray trapping mechanism. This concept spells doom for some other applications of these types of cloaks, which is explored further in section 4.3.

When placing multiple trapping lenses within the cloak, the system becomes more complex as the ray will no longer intersect the same point of each lens on every round trip. Through experimentation with the ray tracer (*RayPlay2D* in [3]) a ray will usually—often after multiple round trips—“choose” a single trapping lens to which it becomes attracted, ultimately leading

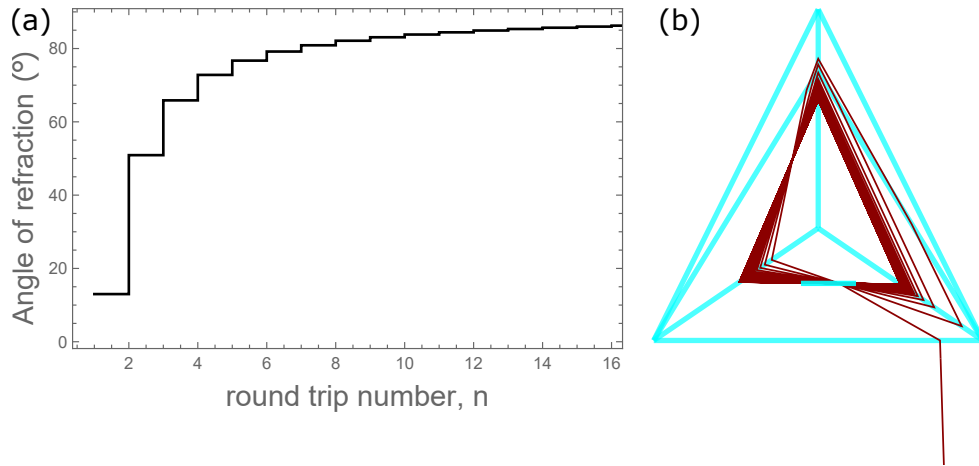


Figure 3.27: A graphical representation of the attractor-like behaviour shown through Eqn (3.3) in (a). It can be seen that with each round trip, shown as a step in the graph, the angle of refraction moves closer towards a final value of 90° which is also the angle obtained when taking the limit of $n \rightarrow \infty$. At 90° , the ray moves parallel to the lens or perpendicular to the optical axis. This can be seen in (b) where a red ray is incident on a trapping lens in cell 1. Upon each round trip, the ray direction has a larger component parallel to the lens, agreeing with the graphical results and the equation found.

to the familiar case where only a single trapping lens is used, as shown in Fig. 3.28.




Figure 3.28: An ideal-lens cloak with multiple ideal lenses placed within it. In (a), a simple schematic can be seen with each cyan line representing a lens. A single ray is then passed into the cloak in (b) and made to intersect all the lenses at least once. After subsequent round trips, it can be seen that the ray is fully attracted to just one lens (the left one) and no longer intersects the other lenses.

There is one notable exception that was found, namely when using two trapping lenses set up in a telescopic fashion. Configuring the telescope such that it magnifies results in a light ray travelling parallel to the optical axis to be shifted closer to the telescope's optical axis after crossing both lenses. After passing through the telescope, the ray will still be parallel to the

optical axis but will shift closer to it. If the initial ray is such that it narrowly misses the R' line and the shift is sufficient to yield an outgoing ray that does cross the R' line, the ray can be captured. Subsequent round trips will pass through the telescope again, causing the ray to shift towards the optical axis asymptotically while remaining captured, illustrated using a ray trace simulation in Fig. 3.29. Therefore, this yields an attractor system that traps light rays through two lenses. In theory, this can be expanded to more complex telescopes and/or multiple telescopes acting in series, yielding an even larger number of trapping lenses a light ray would cross on every round trip.




Figure 3.29: A setup of an ideal-lens cloak containing two additional lenses within, shown in (a). These extra lenses were now arranged such that they form a telescope, with the black dotted line representing the focal plane. The telescope was constructed to magnify the beam when passing left to right, making any incident ray parallel to the optical axis exit the second lens closer to the optical axis. When passing in a light ray at a specific direction, it can be made to run parallel to the telescope's optical axis upon reaching the telescope, as shown in (b). Due to telescopic properties, this means it must also leave the second lens parallel to the optical axis, and the process will begin anew upon the second round trip, with each iteration moving the ray closer to the optical axis. As can be seen in (b), this occurs until the ray appears to be traveling perfectly on the optical axis, always intersecting both lenses.

3.6 Discussion and conclusions

Building on the existing ideal-lens cloak concepts, the cloaking mechanism was explored further. As it turned out, the underlying principle originated from closed-loop light-ray trajectories. These are light rays that travel perpetually within the cloak, passing through each cell along the exact same path. The origin of the closed-loop trajectories may be tracked back to simple imaging properties of the ideal-lens cloak. Considering further imaging properties, a set of rules

was established for when closed-loop trajectories can occur, provided the base lens of the cloak is within the range dictated by Eqn (3.1).

Based on the rules, the initial question posed may be answered. When placing a light source within the cloak, any ray emitted from the source will follow a closed-loop trajectory provided it (or its backwards continuation) does not pass through the base lens of the cloak. Similarly, if it does pass through the base lens, it cannot follow a closed-loop trajectory. Therefore, an observer placed within the cloak, looking outward, will either see the outside space if they are looking through the base lens—which acts as a window out of (or into) the cloak—or alternatively see the back of their own head. The range over which this happens is indicated by the red trapped rays in the ray-traced Fig. 3.7, while the green rays can be seen to leave the cloak. In the case of an omnidirectional ideal-lens cloak [6] no rays may leave to the outside, therefore all light ray trajectories originating within the cloak will be close-looped. Hence an observer within the omnidirectional ideal-lens cloak would see the back of their head in all directions.

Most of the ray-traced simulations shown were for a 2D cloak. Utilizing some common imaging properties, the rules found were generalized to a 3D cloak. Since the 3D version does not restrict the number of base vertices of the cloak, some rules had to be adjusted and reinstated using a different approach. Ultimately, it was possible to generalize the rules to any 3D ideal-lens cloak configuration. This was then used to explore the possibility of a novel omnidirectional cloaking idea called the “Extreme ideal-lens cloak”. Once again, ray tracing simulations were shown to prove the functionality of such a cloak along with the general construction in section 3.4.

In the past, some constructions utilizing ideal-lens cloaks along with other optical component encountered unknown black regions where a ray appeared to enter the cloak but did not leave it. Some of these will be discussed further in chapter 4, particularly in section 4.3. Due to these dark spots, a further investigation was carried out utilizing our newfound rules along with a deeper understanding of the imaging properties within the ideal-lens cloak. It was found that under the right conditions, the simple combination of a cloak and lens can exhibit attractor-like properties. Specifically when directing a light ray from the outside along a trajectory that should be closed-looped according to our rules using a simple lens. Such a ray was found to repeatedly pass through the lens until it asymptotically approaches a fully stable closed-loop trajectory. This violates Liouville’s theorem, which leads to a conclusion that something was fundamentally wrong with the system. The components in question are the *ideal* lenses used throughout. One example where this becomes apparent is that the attractor equation (Eqn (3.2)) can be derived from the simple thin lens equations in combination with the closed-loop trajectory property in every cell of a lens TO device, but is no longer restricted to a paraxial use. A normal lens however, will only follow this when used paraxially. This means our ideal lenses are unphysical. Therefore simply basing our cloak on ideal-lenses should also then mean it is fundamentally unphysical. However, the ideal-lens cloak structure as a whole does not perform

a light ray transformation that could not be replicated with metamaterials. Therefore, it may not be as simple as initially thought. As it turns out, the simple addition of one (or more) idealized lenses to the ideal-lens cloak system causes it to no longer perform according to the loop and edge imaging conditions, and ultimately TO, which then makes it impossible to replicate using metamaterials.

All of these findings are exciting within their own scope, but ultimately they still concern themselves with an *ideal*-lens cloak which currently cannot be constructed. However, by simply using the ideal-lens cloak as a base, it may be possible to extract and generalize some of the rules found to other, more feasible, cloaks. Specifically, the shifty cloak [75], which utilizes a similar imaging idea to the ideal-lens cloak may follow some of the same rules. This is investigated further in chapter 4 which in general concerns itself with the make-up of the shifty cloak. There may also be some transferable ideas to more mathematically inspired cloaking mechanisms such as the non-euclidean cloak [76], where there may also be a line that cannot be crossed unless a closed-loop occurs similar to the R' lines found here. These investigations and generalizations are yet to be completed and may present future research opportunities, building on the concepts explored in this chapter. It is also clear that not all cloaking devices will follow the rules outlined here. Specifically, cloaks that do not inherently turn an object completely invisible, such as the original cloak [67] which utilizes the form invariance of Maxwell's equations to construct a TO device that effectively expands a single point into a finite size space. If an object is now placed within this space, an outside observer would see the object shrunk to the size of a single point which is infinitely small in nature therefore appearing invisible. If the concept is reversed and a point light source is considered at this position, there is no rule that dictates that all rays must remain within the cloak. Instead, it may be the case that such rays will leave to the outside. This subtle difference leads to the incompatibility between the idealized considerations and hence rules found in this chapter and some other types of invisibility cloaks.

Chapter 4

Shifty cloak

4.1 Introduction

The potential cloaking applications of some transformation optics (TO) devices were outlined in chapter 3, with a quick shift in focus towards lens TO devices such as the Rochester cloak [71] which, when generalized to 3D, gave rise to the ideal-lens cloak [6, 16]. The focus here remains on TO optics but with molding ideas from the lens TO cloaks into the currently more realizable (compared to the ideal-lens cloak) metamaterial cloaks. This is by no means the first time TO concepts have been generalized, with metamaterial TO devices being physically demonstrated using crystalin-equivalent TO devices. The carpet cloak [70], originally designed as a metamaterial cloak, has been realized using natural birefringent crystals [77, 78], a relatively low-tech solution to an existing cloak. Similar approaches using crystals have also been used to successfully realize other cloaking types such as the original Pendry cloak in [79, 80]. Ultimately, these wonderfully demonstrate the importance of finding alternative ways of recreating a known cloaking concept all within the confines of general TO.

In much the same way, the initial interest was to generalize the ideal-lens cloak from a lens TO construction into a metamaterial cloak that may, theoretically, be feasible to realize. The underlying cloaking principle of an ideal-lens cloak is, in the simplest case, only a device that images the interior space outside of the physical structure. In the ideal-lens cloak, this leads to further distortions, but the fundamental concept remains a simple *shift* of the interior space. Therefore, a TO device that simply shifts the interior to the outside can in principle act as a cloak in the same fashion as an ideal-lens cloak. Such a cloak may therefore simply be called a “shifty cloak”. The shifty cloak was originally proposed in [75] as a realizable alternative to the ideal-lens cloak. As a consequence, it is advisable to read chapter 3 first to fully understand some of the concepts explained here.

Consequently, due to the beautiful simplicity of the cloak, it comes as no surprise that similar devices have already been theorized [81–88] using wave optical approaches. These extend past optical applications and have been investigated in fields other than optics, with applications in

wireless energy transfers [89] and heat flow control [90,91]. In an effort to expand on the existing research, we not only studied these cloaks in 3D (something only [88] and [85] have done) but also investigated alternative applications of these shifty cloaks. One of the closely related applications to the ideal-lens cloak is the construction of an omnidirectional cloak consisting of two nested shifty cloaks. More exotic applications were also investigated by using a combination of shifty cloaks along with other optical components to create an optical wormhole (a device that provides an apparent optical shortcut to anything behind from a range of viewing positions, but which is invisible when viewed from certain other positions) [92], and a Janus device [93] which appears to look like different things when viewed from differing viewing positions.

This chapter is laid out as follows. A more detailed description of the operating principles and the theoretical construction of a shifty cloak is given in section 4.2. The shifty cloak was then used to create two devices, first an optical wormhole in section 4.3 and then a Janus device in section 4.4.

4.2 Shifty cloak

In its simplest form, a shifty cloak is a construction that images the interior space to the outside. This has been simulated several times using wave optical analysis and ray optics. In our work, the focus was mainly an extension into 3D using a purely ray optical approach.

A 2D schematic which can be extended to 3D later is shown in Fig. 4.1, where a simple box shape was chosen. Within this box, there is an enclosed space that acts as the interior of the cloak. Anything within this space will be imaged by the optical elements of the cloak indicated as a simple grey medium in the figure. The optical elements used can, in principle, be constructed of homogeneous slabs of material. Each slab is responsible for shifting light on transmission in a manner where the shift imposed by one slab is equal in magnitude but opposite in direction to the shift of the slab placed on the opposite side of the cloak. Therefore, the net result is that the two slabs cancel one another out, causing the component to be invisible from the outside. When choosing the appropriate slabs of material it is then possible to extend this to make a cloak that is invisible from all directions acting as a TO device. For this to work fully, each slab must image the interior to the same location. Therefore one set of slabs, opposite to one another, must image the interior to the same spot as the second slab pair. In the simplest case, this will be the case where all slabs do precisely nothing. To achieve more interesting situations, the interior of the cloak should be imaged outside of it resulting in a similar cloaking mechanism as a single ideal-lens cloak where certain viewing directions will not be able to see an object placed within the interior of the cloak. One of these is also shown in Fig. 4.1.

An extension to the same 3D configuration can be seen in Fig. 4.2. The slabs were designed to induce an affine shift in the form of a simple translation of the interior to the outside. Each slab, consisting of a homogeneous material, was designed with flat faces. This results in an affine




Figure 4.1: A schematic of the shifty cloak. The 4 grey-shaded regions represent homogeneous material slabs, each used to redirect light rays in such a manner that the interior region (white square) is imaged to the dashed region and thus "shifted" (in this case to the right). A green ball was placed within the inside region of the cloak, which for some viewing positions appears to be invisible. The image of the ball is also shown with a dashed outline. One such viewing position is shown below the cloak where the red-shaded regions represent the view seen. The darker red part of the view passes through the cloak but *around* the ball while the light red region never intersects the cloak and can therefore not see the image of the ball. To clarify how a ray transverses the cloak, a ray, shown as a dark red line was drawn traveling within within the dark red region before leaving the cloak such that it preserved the initial direction and projected position (shown by a red dashed line intersecting the side of the cloak).

transformation mapping between the object to image positions, which can be replicated exactly using glenses. To feasibly simulate this, the slabs were replaced by glenses with the exact same object-to-image mapping without a loss of generality. The whole scene was ray-traced using the `PolyhedralShiftyCloakVisualiser` [3], which confirmed the cloaking properties of the shifty cloak. The precise material parameters required can be calculated based on the Jacobian matrix of the desired transformation followed by the general TO procedure to determine the material permittivity and permeability tensors. The detailed calculations are outside the scope of this chapter and can be found in [75]. It was found, however, that at least one material slab had to be negatively refractive. Negative refraction has been shown to be achievable with metamaterials but still poses an issue for the feasibility of the shifty cloak. Namely, the bandwidth of resonance-based negative refractive index materials is often limited, and the physical size over which it can be achieved is typically small (due to the lossy nature of negative index materials) [85, 94]. However, with ongoing research interest in negative refraction, especially




Figure 4.2: A 3D model and raytraced simulation of the shifty cloak shown in Fig. 4.1 using the `PolyhedralShiftyCloakVisualiser`, available on [3]. All further simulations in this section were also carried out using this simulation file. In (a), a 3D render was used to demonstrate a green sphere placed within a wireframe model of a shifty cloak. The image of this ball, due to the shifty cloak, can be seen as a green-striped sphere to the right of the cloak. When disabling the wire frame and enabling the optical slabs the full shifty cloak effect comes into play. As demonstrated for the 2D case in Fig. 4.1, the sphere is now imaged to the striped sphere location, but since there is no direct line of sight from the camera that intersects both the image of the sphere and the cloak, the sphere is completely hidden. The only visible feature is a darker outline of the cloak due to simulated absorptions for illustration purposes.

non-resonance-based approaches [95], it may one day be possible to fully realize a broadband shifty cloak. With that prospect in mind, some abstract applications using shifty cloaks were considered.

The first application stems from the same imaging considerations as the ideal-lens cloak in chapter 3, consisting of a nested shifty cloak to hide an object within from all viewing directions. Here, a shifty cloak is placed within the interior space of another shifty cloak. The slabs of material were designed to image the interiors of the cloaks such that the images produced by the outer cloak and the combination of the two cloaks cannot lie along a line of sight. This is shown in Fig. 4.3 where a sphere was placed within the inner shifty cloak. The image position of the sphere is also shown to illustrate the cloaking mechanism along with a ray-traced confirmation that demonstrates the necessity of the inner shifty cloak.

4.3 Optical wormhole

One potential application involving shifty cloaks is an optical wormhole. In the simplest definitions of a wormhole, it is described as linking two distant parts of space through a “hole” in the space-time fabric. The hole then acts as a shortcut to the other side. A simple schematic shows this in Fig. 4.4 where two view positions are shown, one passing through the wormhole and one missing it. Based on the length of the path taken, it can be seen that the route through the wormhole is significantly shorter and any object seen through the wormhole should and any ob-

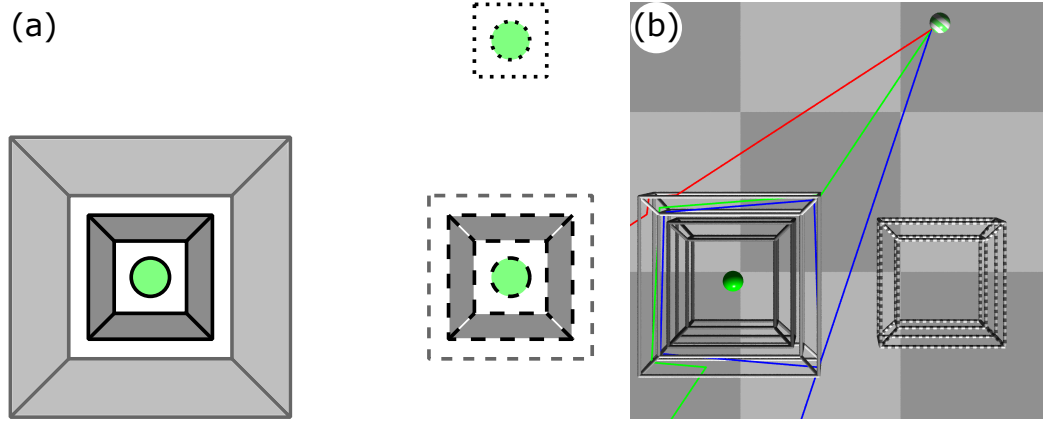


Figure 4.3: An omnidirectional cloak demonstration, (a), and simulation, (b), using two nested shifty cloaks. An inner shifty cloak was nested within an outer shifty cloak as shown in (a) outlined using solid lines. A sphere was placed within the inner cloak as the view object. When taking the first image position of the inner cloak and sphere due to the outer cloak alone, the dashed image is created to the right of the system. A vertical shift due to the inner cloak is then considered, imaging the sphere and interior space of the inner cloak to the dotted second image positions above the first image. Since there is no line of sight connecting the first and second image positions with the physical cloak, it should be fully invisible from all viewing positions. In (b), a top-down version of a ray-traced diagram of a similar scene is shown. Three light rays were shown originating from the second image position, a striped sphere, incident on the physical cloak. These indicate possible lines of sight intersecting the nested cloak structure. As made apparent, the rays do not intersect the first image and hence show that there exists no line of sight between the two images and the nested cloak. As expected, when tracing these rays through the cloaks, it can be seen that none of them pass through to the inner space of the cloaks.

ject seen through the wormhole should therefore also *appear* to be closer. Optically, this would correspond to a shortened optical path length, which is not what our optical wormhole does and would presumably require a real wormhole to accurately reproduce. Instead, an approximation is made to imitate a shortened path length by imaging the points on one side of the optical wormhole closer to the other side. This then yields an approximate image of the view seen through a wormhole, ignoring any warping effects that may actually occur in a real wormhole.

The initial goal was to image an object closer to an observer with magnification 1. One method for doing exactly this is by considering 4 lenses set up as shown in Fig. 4.5(a). The image through all 4 lenses is then shifted closer to an observer by the overall length of the 4 lens construction. The exact expression for this shift can be given by Δz if we assume a shift purely along the z -axis, where

$$\Delta z = 2(f_1 + f_2) \left(1 + \frac{f_2}{f_1} \right), \quad (4.1)$$

for lenses with an optical axis along the z direction and where the first lens (L_1) and fourth lens (L_4) have a focal length of f_1 and the second and third (L_2 and L_3 respectively) have a focal length of f_2 as shown in Fig. 4.5. This was derived by considering the image position due to a single ideal lens given by equation (3) in [72]. When applying the equation to 4 subse-




Figure 4.4: Two distant regions of a 2D space shown to be connected through a simplified wormhole model. The folded space, in this case, can be seen as a blue-tinted transparent surface embedded in 3D space. The wormhole is labeled and can be seen to connect two distant regions of the folded space. When considering two positions on the bottom half of the space (positions 1 and 2), viewing two distant positions on the top half (positions 1' and 2' respectively), it can be seen that the distance travelled by the red line connecting the number 1 points will be much shorter than the distance of the number 2 points. Optically, this means that point 1' will appear to be significantly closer to 1 than point 2' will be for 2, reducing the path length of the light taken through the wormhole. A way to replicate a similar effect purely ray optically would be to image 1' closer to 1 whereby the path length is not physically shortened, but it *appears* closer in proportion to the physical path length reduction. On the other hand, for some view positions, such as 2, this should not occur, and the wormhole should be invisible from the view. This ignores any possible distortions caused by a real wormhole.

quent lenses the overall image position for an object placed at (x, y, z) should be at $(x, y, z + \Delta z)$. Provided the focal lengths of both lens types are given as f_1 and f_2 , the separation between the lenses along their optical axis can be solved analytically for a unit lateral magnification scenario. Effectively this makes use of the fact that the overall image should not be shifted in the x or y component. Finally, it can then be calculated what the corresponding value of Δz should be. Fig. 4.5(b-d) prove this by first simulating a standard scene with a simple object –here performed by a “Tim” head model. The model is then placed behind 4 lenses and simulated. Simulating the same scene but now with the camera moved closer to the Tim head then yields the direct comparison. The distance the camera is moved is equivalent to the size of the full 4 lens system and can be seen to then produce the same image as through the 4 lens system.

Having found a simple lens setup that can image anything seen through it closer by Δz , one-half of the wormhole is completed, namely imaging something closer to an observer. The other half requires this to only work from some viewing positions and be completely invisible from all others. Using a pair of shifty cloaks this may theoretically be achieved. When considering the imaging properties that make the omnidirectional shifty cloak work in section 4.2 another nested cloak setup can be created which does not cloak the interior from all directions but instead for

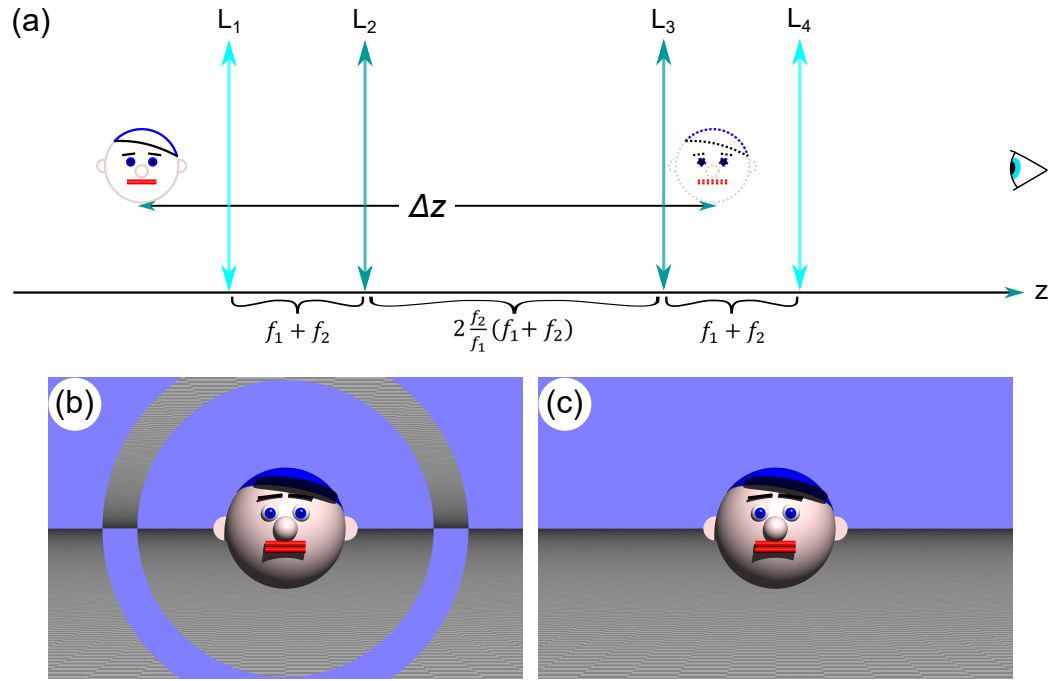


Figure 4.5: In (a), a demonstration of a 4 ideal lens setup that images an object closer to an observer, therefore making it appear larger. The distance by which the the object, a simple model of a Tim head, is imaged closer is given by Δz which is analytically expressed in Eqn (4.1) for an observer shown as an eye. The relative lens separations are shown as the sum, and scaled sum of their focal lengths ultimately creating two telescopes separated by some intermediate distance that is also dependent on their focal length. A raytraced simulation confirming this is shown in (b) and (c) created using the `IdealLensWormholeVisualiser` available at [3]. Here, (b) shows the view through a lens setup with all focal lengths set to 1 unit. This should be equivalent to a z shift of 8 units closer to the original camera position. Note that the simulation software uses arbitrary units where one floor tile represents a length of 1. When moving the camera 8 units closer in (c), and removing the lenses, it can be seen that the Tim head and background are of the same size as in (b). The outer ring in (b) represents regions of space where a ray does not intersect all 4 lenses.

just some. This was done by designing a nested cloak setup where the images of the interiors lie along a restricted line of sight, that is for some directions and positions a straight line may be drawn connecting all the images created. An example of such a construction can be seen in Fig. 4.6.

To determine the viewing directions from which the interior is visible, an additional similarity to the ideal-lens cloak is considered. Namely, the existence of a window (explained in section 3.3.1 Fig. 3.15) into the interior of the cloak. By considering the image of all the slab interfaces, it can be shown that any shifty cloak, that images the entirety of the interior to the outside, has a window which lies along the outer edge of the slab the image is shifted over. Schematically this is shown in Fig. 4.7(a). The reason for the window can be explored in terms of the images of the slab boundaries. When taking a light ray that passes through the interior in the *image* space it may pass through some of the slab interface images. The image of the ray must then pass




Figure 4.6: The setup for an optical wormhole. A nested shifty cloak construction is shown outlined by solid lines. Within the outer cloak, there is an inner cloak and two lenses, with the inner cloak containing an additional 2 lenses. Overall, each pair of lenses is separated by the sum of their focal lengths. The first images, due to the outer cloak only, are shown using dashed outlines, consisting of the inner cloak and all 4 lenses. The second images will be due to the combination of the outer and inner cloaks and are given by dotted outlines. The exact shift amount and hence image locations were constructed such that the separation between the first lens images and the second lens images was the required distance for the 4 lens construction shown in Fig. 4.5. This was done to allow for a slightly larger shift than simply the size of the inner cloak cavity. The light grey shaded region represents the view positions where a line of sight intersecting all image positions may occur further explained in Fig. 4.7.

through the same interfaces within the real cloak dictating how it propagates through the cloak. An extreme case example is shown in Fig. 4.7(a) with (b) and (c) confirming the lines of sight along which the object is visible using ray trace simulations. This is precisely the reason that restricts the possible viewing positions and directions to finally arrive at what is shown in Fig. 4.6.

By placing the 4 lens setup into the nested shifty cloak construction in Fig. 4.6, it should theoretically image anything seen through it closer, but when looking at the device along any line of sight for which the interior is invisible, it should appear to be fully invisible. A small modification to increase the observed shift size was carried out by placing only one lens pair into the innermost cloak and the second pair into the outer cloak, ultimately allowing for a larger variation in focal lengths and hence separation of the final image positions as shown in Fig. 4.6 due to the relative shift size of the cloaks. Raytraced simulations confirm this for first the on-axis view straight through the cloaks and then for a view from the side in Fig. 4.8(b) and (c). However, when viewing it for some position further off the optical axis—but still within the theoretically permitted region—a black region appears instead of the expected 4 lens system view demonstrated in Fig. 4.8(d).

These black regions correspond to the raytracer running out of ray object intersections which we call trace levels. These represent the number of permitted ray trace intersections before simply returning black for the specific pixel. In the simulation used this was set to 1000. Therefore,

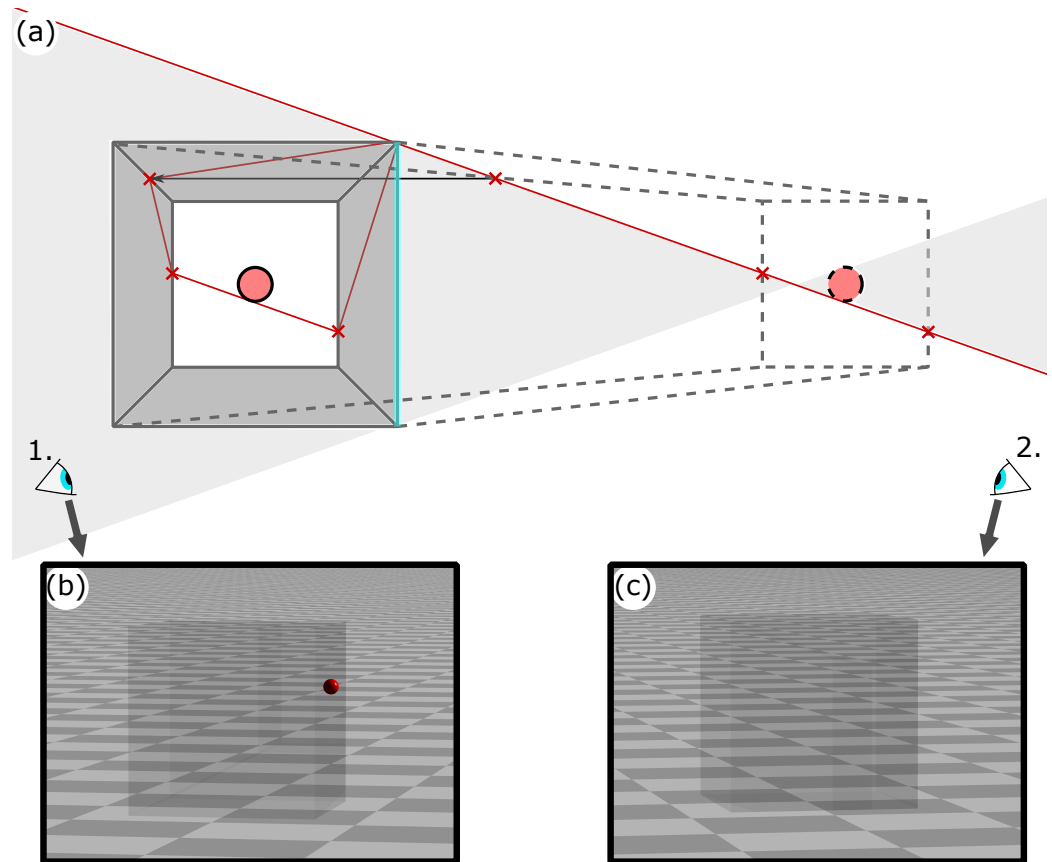


Figure 4.7: An illustration and demonstration of the restricted viewing directions of a shifty cloak. Similar to the ideal-lens cloak, a window into the shifty cloak may be considered to lie along the surface of the cloak over which the shift occurs. This is shown as a blue-shaded region on the right-hand side of the cloak in (a). The image positions of the interior are also shown in (a), giving the image of the slab intersections and interior as dashed lines. Using the intersections between the ray and the image positions, the trajectory through the cloak was constructed indicated as a red line, while the red crosses indicate the corresponding intersections. Two viewing positions are given as 1 and 2, with 1 being placed within a grey-shaded region representing all lines of sight that can intersect both the window and the view object (a red sphere). Unlike 1, 2 is placed outside the shaded region and therefore cannot intersect the window and the image of the red sphere. A ray-traced simulation confirming this is shown in (b) for view position 1 and (c) for position 2. In (b) the red sphere can therefore be seen while in (c) it cannot. Once again, for this, the `PolyhedralShiftyCloakVisualiser` was used.

this means that there were more than 1000 ray intersections for the wormhole construction. Regardless of the trace level, the result remained unchanged. The reason for this lies within the attractor trajectories explained in section 3.5. In a similar manner, it was considered that a ray may only pass through a slab or the interior cell once, and therefore if it were to follow a closed-loop trajectory, it would need to retrace the exact same trajectory within each part of the cloak. The simplest closed-loop that may form in the shifty cloak is therefore one that travels through the cloak cavity and two neighboring slabs of material crossing two slab cavity boundaries and a single slab-to-slab boundary. Considering that this mimics the behavior in the ideal-lens cloak,




Figure 4.8: Simulating the shifty cloak wormhole setup outlined in Fig. 4.6 using the `ShiftyCloakApplicationVisualiser` simulator available at [3] yields all images shown here. In (a), the cloaks and lenses were removed to see a small Tim head in the distance. This represents view 2 in Fig. 4.4 and acts as the control. The same scene is then viewed again but now with a shifty cloak wormhole setup, first the on-axis view in (b), then from a position from where the interior of the cloak is expected to be invisible in (c), and finally slightly off-axis in (d). The Tim head appears to be closer in (b) as desired while the lenses are fully invisible in (c) as desired. However, in (d) a dark region appears where the lenses should be. The extent of this region may be better understood when viewing the video in [7], which shows a simulated fly-around view of the shifty cloak wormhole.

it comes as no surprise that when adding just a single lens within the interior of the shifty cloak, it fully reproduces the attractor behavior of the ideal-lens cloak explored in the attractor section of section 3.5 and therefore follows both Eqns (3.2) and (3.3). Adding multiple lenses in the ideal-lens cloak has shown that a ray will in most cases be attracted to a single one until it runs approximately parallel to the lens yielding a complete closed-loop trajectory.

One way to confirm this result is by comparing it to an optical wormhole constructed using the ideal-lens cloak. Creating a similar nested ideal-lens cloak to restrict the viewing positions and placing 4 lenses inside yields a similar construction to the one seen in the shifty cloak wormhole. It should be noted that the transformation is no longer a uniform translation of the entire interior. Therefore, the separations between the lenses and their focal lengths had to be accounted for such that the final image of the lenses (the image due to both cloaks) was precisely the desired 4 lens setup. This calculation along with a detailed construction of the

ideal-lens wormhole are outside the scope of this chapter and have been shown previously in [8], a summary is also given in appendix A. The results agree with what can be seen in the shifty cloak wormhole with the equivalent attractor trajectories appearing for certain viewing directions, demonstrated in Fig. 4.9.




Figure 4.9: A simulation of an ideal-lens wormhole sharing some of the properties seen in the shifty cloak wormhole. Here a similar setup to Fig. 4.6 was used, with some differences in lens focal length and position. The exact details were previously covered in [8] and have been summarized in appendix A. The properties seen in (a) show a Tim head imaged closer to the observer in the same on-axis view used previously. To compare the size increase, a small portion of the original Tim head can be seen in (b) while also showing the same, but enlarged Tim head towards the center. Since this is now an off-axis view, a similar black region forms due to attractor-like properties. This comparison was made in the `IdealLensWormholeVisualiser`.

The overall shifty cloak optical wormhole construction therefore only performs as an optical wormhole in some parts. It fulfills the previously outlined properties of an optical wormhole for the on-axis view but not for all the theoretically permitted viewing positions. Instead for some it simply returns a ray-optical black hole (seen as a black spot) due to attractor trajectories. The wormhole is correctly hidden for the expected view positions outside of the theoretically allowed ones. This can be seen in a short video of the shifty cloak wormhole provided in [7] where all of the aforementioned scenarios are visible when moving the camera around the system.

4.4 Janus device

A Janus device is a construction named after the Roman god of the duality of the same name. It comprises a property in that the view seen through the device looks different from differing viewing positions. In the original optical Janus device [93], the two different viewing options were either a lens or a (transverse) beam shifter. For our purposes, we therefore want to create an optical device that contains at least two “modes” that show a different view when seen from different viewing positions.

The shifty cloak may be used to create this by using two shifty cloaks to image their interior to the same space. Therefore, depending on the viewing position and direction, one or both




Figure 4.10: The Janus device setup in (a), shows two shifty cloak wire frame models each containing a sphere. One of the spheres is red and solid while the other is green striped with each of them being imaged into the same space shown as a red and green striped sphere. Three camera positions are shown with each viewing the image from a different direction. For the camera placed at 1, some sample rays were drawn to indicate the expected view and show that the camera will see the red sphere only. The same could be done for all other camera positions which can be checked using (c-e). When viewing the setup from camera position 1, shown in (b), the drawn trajectories can be confirmed by showing that the red-green striped sphere only lies in a line of sight with the cloak containing the red sphere. (c-d) correspond to camera positions 1-3 respectively, where (c) confirms the expected outcome based on the rays shown in (a), where only the red sphere is visible. As the camera is moved around the cloak either one sphere, the other sphere or both can be seen at any one time. When both can be seen as shown in (d), half of each sphere is shown. lastly, (e) as expected shows only the green striped sphere. A video showing a round trip around the cloak is given in [7] to better understand these properties. All the simulations in this section were carried out using the simulations software named `ShiftyCloakApplicationVisualiser`.

of the interiors can be seen at the same time. The image, and hence the interior of the cloak, can be seen if there exists a direct line of sight between the shifty cloak window and the image intersecting the viewing position, as demonstrated in Fig. 4.7. A diagram of a Janus device constructed from shifty cloaks can be seen in Fig. 4.10, where, within the interior of two cloaks, two spheres and their image positions are shown. A pair of traced rays are also shown to indicate rays entering camera 1 and their interactions with the scene. A ray-traced version of the same setup is visible in Fig. 4.10(c-e), which confirms the expected results.

4.5 Discussion and conclusion

Barring the effects of attractor behaviors in the optical wormhole, both an optical wormhole and a Janus device were successfully constructed using shifty cloaks. The Janus device shown successfully fulfilled all the desired properties, and while the application is currently only a proof of concept, it is possible to exchange the interior object from simply a sphere to some optical component. Ultimately, this would then be able to perform any permitted transformation that does not redirect incident light rays along a closed-loop trajectory. In some form, one of these may be an optical wormhole, where from one direction it appears as bringing space closer to the observer while from another it looks invisible. In the chapter, the component inside the wormhole responsible for the translation was constructed using 4 internal lenses. There may be an alternative approach to this, which is yet to be investigated, which involves simply removing one of the slabs that form the inner shifty cloak. Since a single slab effectively shifts a light ray that enters into it, with the opposite slab undoing that shift, removing one of the pairs would induce the light-ray shift. If the on-axis slabs are chosen, this could then shift light rays closer to an observer in a similar way to the 4 lens setup. Of course, this then still requires a second cloak to dictate the directions from which the system will be seen but it may decrease the complexity of the system.

As described in section 4.3, the optical wormhole performs correctly on-axis and when viewed from a fully cloaked direction, but not when seen from an off-axis position where the full wormhole effect is still expected to occur. Instead, what can be seen is a black patch representing a region where the simulation traced the rays for more than 200 intersections without exiting the cloak. The only feasible way this may occur is when the rays become trapped within the shifty cloak. The same occurrence can be seen when constructing the wormhole using the previously investigated ideal-lens cloaks. It has been shown in section 3.5 that attractor-like properties can be found in the ideal-lens cloak when placing a lens inside, causing the light rays to be redirected to a closed-loop trajectory. Since the shifty cloak may equivalently be considered using glenses that follow the rules of the ideal-lens cloak construction –the loop-imaging and the edge-imaging condition– and utilize a similar underlying cloaking mechanism, it appears likely that shifty cloaks will also exhibit the same attractor properties as the ideal-lens

cloak. Unfortunately, this then also hints at the possibility that in order to remove the attractor behavior, the cloak can no longer function within the cloaking regime outlined in section 3.3, Eqn (3.1) for an ideal-lens cloak. Or perhaps even simpler, these components can only exist paraxially where a lens can act akin to an ideal-thin lens.

This result, derived from the ideal-lens cloak, shows that the theoretical study of a currently impossible-to-realize cloak may still shed light on some of the mysterious behaviors of a cloak grounded in reality. On the other hand, it also shows that the same shortcomings still apply for some of the potential applications of these cloaks. It should be considered that the light-ray trapping itself may be useful for certain purposes in resonator physics or as a means of potentially “catching” sunlight for energy production.

While the shifty cloak is closer to a realizable cloak when compared to the ideal-lens cloak, it should still be noted that the feasibility of producing such a cloak in its ideal form remains outside the current possibilities of commonly available optical materials. The main issue is the inevitable use of a slab that has to perform negative refraction which is not, as of now, possible for a material over a broad spectrum of wavelengths. Regardless, we believe that the use of a heavily idealized model in the ideal-lens cloak allowed for a thorough understanding of a, more realistic, shifty cloak. This extends past simply the underlying cloaking properties into some of the strange occurrences that may appear when combining the cloaks with other optical components. If it is possible to generalize the ideal-lens cloaks further to apply to a broader range of cloaks, it may be useful in a similar manner to what was shown in this chapter for understanding the closed-loop trajectories and attractor behaviors.

Final conclusions

This thesis explores two seemingly distinct but ultimately connected fields of study: medical spectacles and transformation optical cloaking. At the heart of the research lies a unifying approach, namely that carefully structured micro-optical elements can redirect light in unconventional ways. In the case of the view rotation spectacles this permitted a pseudo-rotation which, if you do not look too closely, can seem like the device rotates the view seen through it. However, a perfect image rotation should be optically forbidden when performed by a single continuous and thin holographic optical element. Since this thesis does not claim to circumvent physical laws, a trade-off was required to achieve the view rotation. This came in the form of a pixelated component, where each pixel performed an optically permitted transformation, a local shift, to give the impression of an overall rotation when considering the combined view through all the pixels. In all instances these claims were backed up by ray-trace simulations constructed in our open-source ray tracer.

For the second pair of medical spectacles, the adaptive lenses, a similar approach was used, where the image was segmented to permit a change in composite focal length in a way that is both compact and adjustable. In this case the segmented images consisted of spiral windings. In both medical spectacles the focus was on investigating how the view through the spectacles changes and how it could be improved using various optical techniques that once again involved a trade-off between improved view quality and further optical or mechanical complication. For a quantitative measure of the improved view, a Snellen chart was used for the view rotation spectacles, which showed a roughly tenfold improvement through pupil restrictions and derivative control. In the adaptive spectacles this came in the form of winding focusing, which allowed the integral image between neighboring windings to form a nearly seamless image, barring ghost image formation through misdirected light rays.

When investigating the ideal-lens cloak, several conditions for closed-loop light-ray trajectories were found, a phenomenon also present in other cloaks. These closed-loop trajectories permitted the construction of an attractor through the addition of a further ideal lens. According to Liouville's theorem the type of attractor demonstrated here is unphysical, this indicated that some assumption made must have also been unphysical. As it turned out, this was the ideal thin lens, which was assumed to work equally well outside the paraxial regime. We found that a specially assembled structure of these unphysical lenses may in fact create a component that

could, at least in theory, be replicated using metamaterials. The simple addition of one further ideal lens to form an attractor, however, breaks this possibility. Given that this again verges on a seemingly forbidden concept, it may be possible to apply some of the tricks used for the medical spectacles — such as a pixelated approach — to work around some of the unphysical nature of the ideal lenses. In fact, an approximation to an ideal lens, which is a special case of a glens, can already be made by using lenslet arrays arranged to form a Gabor superlens. This could serve as a future research direction for the ideal-lens cloak.

Alternatively, by taking the cloaking properties of the ideal-lens cloak and identifying material properties that could perform a similar type of imaging, the shifty cloak was devised. Since this was now a system inherently designed with feasibility in mind (albeit difficult due to negative refraction), applications previously theorized using the ideal-lens cloak were explored, such as an optical wormhole. These constructions again used ideal lenses and encountered similar attractor-like properties which, for the purposes of a wormhole, were not desired. We note that it may be possible to remove the ideal lenses from the system by simply removing one of the material slabs which should, at the least for some view directions reproduce the wormhole results.

In future, we hope to further expand our catalog of medical spectacles to provide treatment alternatives for several eye conditions. Testing and optimizing these for the best possible view quality and form factor will be natural further steps. The understanding we gained from investigating these devices beyond a purely theoretical construction may also aid in developing theoretical cloaking structures into more feasible forms, even if their usefulness remains limited.

In conclusion, this work has shown how structured micro-optical elements, whether wedges, spirals, lenslet arrays, or cloaks, expand our ability to manipulate light. Furthermore, it highlights both the practical promise and the theoretical depth of micro-optics, contributing toward a future where optical devices following these principles may simultaneously improve lives in everyday medical contexts and deepen our understanding of what is physically — or in some instances nearly physically — possible in light control.

Appendix A

Ideal lens wormhole

The underlying principle in the ideal lens wormhole is the exact same as for the shifty cloak wormhole described in section 4.3. It consists of 4 lenses surrounded by a nested cloak –in this case, two ideal lens cloaks– to control the directions for which it is visible. The arrangement for the ideal lens cloak is such that the image positions of the 4 internal lenses are such that they are spaced by the required separation demonstrated in Fig. 4.5. To achieve this, first, the desired image position of the 4-lens configuration is taken as a starting point, which we will call the *second* image. This is then imaged by the *image* of the base lens of the inner cloak (where the image is due to the base lens of the outer cloak), leading to the first image positions for the 4-lens configuration. The way this is calculated is by considering the imaging properties of an ideal thin lens given by

$$\mathbf{Q}' = \mathbf{Q} + \frac{(\mathbf{Q} - \mathbf{P}) \cdot \hat{\mathbf{a}}}{f - (\mathbf{Q} - \mathbf{P}) \cdot \hat{\mathbf{a}}}(\mathbf{Q} - \mathbf{P}), \quad (\text{A.1})$$

taken from [16], where \mathbf{Q}' is the image position, \mathbf{Q} is the object position, \mathbf{P} is the principal point of the lens, $\hat{\mathbf{a}}$ is the normalized optical axis direction vector, and f is the focal length of the lens. This first position is then imaged once again by the base lens of the outer cloak, yielding the final physical positions required in order for the second image to align with the desired 4-lens configuration. The only requirement is now that the physical positions of the 4 lenses must lie within the inner cloak, which can be checked. This is highlighted in Fig. A.1 where a schematic of the setup is shown, including the region where the wormhole should be visible, highlighted in grey.

In a similar manner, the focal lengths will also change whenever a lens is imaged. Since the final lens positions are such that they follow Fig. 4.5, this means the focal length of the physical lenses must be calculated in a similar manner to the position. Considering the focal lengths in the same order as the positions, that is, the focal length is first changed by the image of the inner cloak base lens and then by the outer cloak base lens. This change is given by the product of the magnification, due to the imaging lens, and the focal length of the lens that will be imaged. The final expression for each focal length can be calculated generally, but the expressions are




Figure A.1: A schematic of the ideal lens cloak wormhole. The solid lines represent the physical space objects, where the outer cloak is cyan colored and labelled as O , the inner cloak is blue and labelled as I , and the 4 lenses within are all green, labelled as L_{1-4} . The inner cloak and the lenses are then imaged to the outside as a first image shown as dashed lines, with the same labels but primed i.e. I' and L'_{1-4} . The first images (L'_{1-4}) of the 4 lenses are then imaged once again by I' , to their final image position given by L''_{1-4} . These are shown by dotted green lines, and can be seen to be separated by exactly the 4-lens relationship outlined in Fig. 4.5 as desired. The area of permitted viewing positions for which a view direction exists where the wormhole should be visible is given by a grey background.

too lengthy to be included here.

When adhering to the positions and focal lengths calculated, a similar wormhole to the shifty cloak wormhole can be simulated, which is the ray-traced setup shown in Fig. 4.9.

Appendix B

Ray tracing simulations and software

The ray tracing simulations used comprised the CPU Dr Tim [4] software and GPU-based web browser simulations, programmed using three.js [48]. As a third alternative, Unreal Engine 5 (UE5) [96] was briefly explored. A short summary describing each simulation package and its usage is given in this appendix.

B.1 Dr Tim

Dr Tim is a ray tracer written in Java for a CPU-based approach. Due to this, it suffers from some of the shortcomings that most CPU programs have when compared to GPU-based counterparts, namely, processing speed. Therefore, Dr Tim’s use is restricted to be interactive but not exactly real-time. Instead, each scene is rendered, where every pixel is calculated in a ray-casting approach.

Ray casting, as the name suggests, is the process of sending a ray outwards from the “camera” to intersect with the scene. When the ray eventually intersects an object that is neither transparent, refractive, or reflective, the ray is stopped, and the pixel it was sent from returns the color of the intersected object. This is done for every pixel. For intersections with any other object, the incident light ray direction at the intersection point is taken and changed according to the surface property. In order to approximate some optical components, phase holograms are used where the phase profile is selected to mimic that of the optical component it should represent. This makes the scene object intersection calculations significantly easier. Other idealized components also exist, such as ideal mirrors and lenses. Since the interactions are all based on ray optics only, most wave optical effects are not considered or have to be programmed in separately, such as certain diffraction effects.

By placing the relevant objects in front of a camera and setting up a backdrop scene (typically consisting of a Tim head or a lattice), most optical components can be simulated, provided the appropriate surface profile exists (or is created) within Dr Tim. Some objects have been integrated to be available in the fully interactive version of Dr Tim available at <https://>

`github.com/jkcuk/Dr-TIM/releases/tag/Interactive_TIM_v.1.0`. This along with further components may be found in the extensive (but potentially a bit messy) catalog available on the GitHub [97].

B.2 Web browser simulations

Similar to the Dr Tim simulations, the ray tracer developed for the web browser (hereafter referred to as the web app) is primarily based on ray optical interactions. Since it was written in `three.js`, it utilizes the GPU, allowing for significantly faster processing speeds. Since this is a new approach, there are several optical components that are yet to be integrated into the web app. Therefore, the use is limited throughout this thesis.

One of the main advantages of using the web app (specifically `three.js`) is the compatibility with virtual reality (VR) headsets. These could be of special interest for clinical studies, especially in conjunction with some of the spectacles outlined throughout this thesis. By using a VR headset, some of the expensive prototyping could be delayed until the effectiveness of the spectacles is more clearly confirmed using a relatively cheap and efficient simulation-based approach.

The ray tracing is constructed in the same way as with Dr Tim but with the added complexity of integrating the GPU. The result is a ray tracer that can simulate several scene intersections and interactions in real time. The currently available simulations can be found at `https://github.com/jkcuk`.

B.3 Unreal engine 5

The last ray tracer we considered and to some extent tested was developed by Epic Games and is called Unreal Engine 5 (UE5) [96], a free (with some commercial restrictions applying) downloadable program. As a potential entry barrier for UE5, we note that some of the ray tracing capabilities require a moderately powerful but importantly recent graphics card, equipped with ray tracing cores. This engine was primarily developed as a game engine but features some of the most advanced ray tracing capabilities in the field. The user interface allows for a simple and quick scene construction (we typically used an empty basic scene) along with the capability to create refractive material properties. The refractive properties used followed from an online tutorial available at `https://www.youtube.com/watch?v=UY9nf6HEGIM`.

A pipeline was developed whereby we started by identifying a refractive component we wanted to ray trace. The component was then calculated in Mathematica (for example the view rotator in chapter 1) and then turned into a 3D model using some CAD conversions, or directly modeled in a 3D CAD software. The 3D model was then exported into UE5 using the Datasmith plugin (available at `https://www.unrealengine.com/en-US/datasmith/plugin`




Figure B.1: An UE5 simulation of the refractive view rotator. In (a) a scene with a green bar tilted at 10° can be seen. The same scene is then viewed but after placing a view rotator in front of the camera in (b). It can be seen that the green bar is approximately rotated back to a horizontal orientation. The slight offset from being perfectly horizontal is due to the camera distance not being placed exactly where the view rotator was designed for it to lie. This could be fixed by creating another 3D model of the view rotator but for a further eye distance. It is for the same reason that the edges appear to be bending, with a similar effect clearly noticeable in Fig. 1.8, we conclude that the view rotator is accurately (to any noticeable degree) simulated in UE5.

ins). After importing the model, the refractive (or if desired reflective which we have not experimented with) material properties were set.

The scene could then be viewed in real time with an example scene of the refractive view rotator being visible in Fig. B.1. The curving on the edges and imperfect correction is due to the camera setup, which was not placed at the appropriate distance from the view rotator, but it can be seen that the component does what we expected it to do, namely “rotate” the green bar back to the roughly horizontal position when compared to the view without the component. (If this were to be repeated, it would be advisable to design the rotator with a further eye distance in mind to permit the spacing in UE5.) This was promising, but we noticed some undesired effects when attempting to create a more complicated optical setup. This second setup consisted of two lenslet arrays constructed such that shifting them relative to one another will change the magnification from the very basic unit magnification case in Fig. B.2(b) to a larger magnification in Fig. B.2(c). What can be seen however is that in (a), the view appears not magnified as intended but we notice some imperfections, but even more severely in (b), we can see that the view does not at all match with the expected view. When comparing this with a Dr Tim simulation we decided that it was likely an issue with UE5 (or maybe more fairly, an issue with how we used the UE5).

The issue could be from any of the following sources or a combination of them. Generally, the pipeline developed may not be the most streamlined, with each exporting and importing step




Figure B.2: three screenshots of an UE5 scene containing two lenslet arrays, forming a telescopic array, designed to change magnification when shifted by one period to the side. In (a), a view without the components can be seen of just a 3D model of a Tim head. (b) adds back the component but in the basic configuration where it should return a unit magnified view. Some gray stripes are visible which are baffles required for this construction and are there by design. Besides this, the view seen mostly aligns with the expected view in (a), barring some small double images in each telescope center. However, when shifting the component by one period in (c), the view seen is significantly deteriorated, past any expected levels when comparing it with a Dr Tim simulation. This might indicate some issues with using the UE5 in the way we have done here for accurate ray tracing.

creating small but in an optical sense noticeable imperfections. For example, when exporting a sphere made in Mathematica into a 3D CAD software, it does it by creating small triangulations along the surface consisting of a point cloud of connected vertices. This can then be turned into a solid model in the CAD software (which is required for the UE5 model). Exporting this solid model into the UE5 using the Datasmith import tool then re-triangulates the model based on some its own optimization process. This will again lead to further surface “defects” specifically for curved surfaces. Producing the model in UE5 directly may therefore yield better results, but we have not investigated the feasibility of doing so sufficiently.

We also note that, to us, the ray tracing performed by UE5, while impressive, is very much a black box. As such, any optimization procedures performed by the engine to increase the ray trace speed may come at a cost of physical accuracy. This is certainly the case in classical “ray traced” scenes in game engines where the ray tracing was simply done using surface textures that are far from physically accurate ray tracers but instead simply look nice. While the code to determine the ray trace accuracy of UE5 is available (<https://dev.epicgames.com/documentation/en-us/unreal-engine/downloading-source-code-in-unreal-engine>) the source code is extensive and beyond the scope of this small experiment. Instead we attempted to determine the accuracy by comparing the results with the Dr Tim simulations with moderate success.

A work around to all of these issues is to use a feature within UE5 that permits a user to work directly with some of the underlying code written in C++. As an extension, modifying the source code to ensure physically accurate ray tracing may also be an option, but one that would be closer to a computer science project rather than a physics one. By doing both, it may be possible to create even complex components directly in UE5 and ray trace them.

While we do not use the UE5 in this thesis, nor in any draft papers currently in preparation, we thought it may be of interest to some readers who want to experiment further with the UE5 as it could serve to be a powerful ray tracing tool when used appropriately. We have managed to show that at least for flat optical surfaces, it performs (mostly) physically accurate ray tracing. As a last note, we have not noticed any wave optical considerations such as diffractive effects. This is again something that may be possible to include when modifying the code directly.

Appendix C

Wave optics of the adaptive Fresnel lens

As a declaration, the following paper draft was not written entirely by the thesis author, but rather as a collaboration between the authors on the paper, with Dr Johannes Courtial being the main contributor. The contents of the paper draft are also subject to changes before and after submission to a journal and may therefore not be fully representative of the final version. As such, it mainly serves as an extension for any curious reader, past the purely ray optical regime of the adaptive Fresnel lens, into the wave optical world.

Wave optics of spiral adaptive Fresnel lenses

MAIK LOCHER¹, ELISE BLACKMORE¹, ZHAO-KAI WANG¹, AND JOHANNES COURTIAL^{1,*}

¹*School of Physics and Astronomy, College of Science and Engineering, University of Glasgow, Glasgow G12 8QQ, UK*

**Johannes.Courtial@glasgow.ac.uk*

Compiled May 24, 2025

We recently introduced adaptive Fresnel lenses (AFLs) based on complementary cylindrical lenses “bent into” spirals. Here we start to investigate the wave optics of such components. We calculate typical point-spread functions and investigate diffraction effects in the case of narrow windings. We design a modification that results in the phase profile for light that passes through corresponding windings being that of a parabolic lens, exactly, but only in the case of infinitesimally thin phase components. Our findings contribute to the improvement and understanding of spiral AFLs, and pave the way towards their practical realisation.

<http://dx.doi.org/10.1364/ao.XX.XXXXXX>

1. INTRODUCTION

Lenses are ubiquitous optical components. Adaptive lenses, in which the focal length is adjustable, have added versatility and found numerous applications [1]. There are many types of adaptive lenses [1], including designs based on liquid confined between elastic membranes and liquid crystals (which includes Ref. [2], which sounds very closely related to this work but uses quite different ideas). Of special significance to our work are designs in which two refractive components combine to redirect transmitted light like lenses whose focal length depends on the lateral offset between the components, described by Alvarez [3] and Lohmann [4] and referred to here as Alvarez-Lohmann lenses, and Bernet and Ritsch-Marte’s varifocal Fresnel lens [5], called BRM lens here for short, in which two diffractive components with phase profiles (Eqn (8) in Ref. [5])

$$\Phi_{\text{BRM},i} = \pm ar^2\varphi, \quad (1)$$

where $i \in \{1, 2\}$ is the number of the component, the + and – sign corresponds to component 1 and 2, respectively, a is a parameter and r and φ are polar coordinates on the surface of the (unrotated) components, combine, when rotated relative to each other by an angle $\Delta\theta$, into a parabolic lens with a $\Delta\theta$ -dependent focal length (Eqn (10) in Ref. [5])

$$F_{\text{BRM}}(\Delta\theta) = \frac{\pi}{a\Delta\theta\lambda}. \quad (2)$$

A recent paper [6] presented a design for novel adaptive Fresnel lenses (AFLs) which was further investigated in Ref. [7]. The main idea is that pairs of complementary cylindrical lenses with focal lengths $+f$ and $-f$ can redirect parallel bundles of light rays in proportion to their transverse offset (Fig. 1(a)); that both complementary cylindrical lenses can be manufactured such that they are “bent”, specifically so that the nodal lines (the lines

which do not redirect light rays intersecting them) form identical spirals; that rotation of the two resulting complementary-cylindrical-lens spirals by an angle $\Delta\theta$ results in a radial offset ΔR between the cylinder-axes (Fig. 1(b)); and that the setup can be modified such that the combination redirects transmitted light rays approximately like a simple Fresnel lens whose focal power, $P = 1/F$ (F is the focal length), changes with $\Delta\theta$ (Fig. 1(c, d)) as

$$P = p\Delta\theta, \quad (3)$$

where p is the ratio of focal power and rotation angle between the components (units diopters / ° or diopters / rad).

The last point requires that the radial offset ΔR between the nodal lines of cylindrical lenses at a radial distance from the centre R , the magnitude f of the focal lengths of those cylindrical lenses at radial distance R , and the focal length $F = 1/P$ of the resulting Fresnel lens are related through the *radial-offset equation*

$$\frac{\Delta R}{R} = \frac{f}{F}. \quad (4)$$

Ref. [6] studied the case of the nodal lines forming logarithmic spirals, resulting in a value of f that is independent of R . For all other spiral shapes, such as the Archimedean and hyperbolic spirals studied in Ref. [7], f depends on R , which means it depends on the position on the spiral. Table 1 lists the relationship $f(R)$ for different nodal-line spiral shapes.

Fig. 2(a) shows the radial phase profile of a logarithmic-spiral AFL constructed in this way, together with the quadratic radial phase profile of a parabolic lens with the same focal length as the AFL (Fig. 2(d)). What matters for the redirection of transmitted light rays is not the absolute value of the phase, but the phase *gradient*. Within each winding, the radial phase gradient of the curve shown in Fig. 2(a) is exactly constant (note that, for other spiral shapes, this is the case only approximately), unlike in

spiral type	R	f	A	C
logarithmic	$\exp(b\theta)$	f_1	$-1/(Rf_1)$	$R^2/(2f_1)$
Archimedean	$b\theta$	f_1/R	$-2/f_1$	$R^3/(3f_1)$
hyperbolic	$-1/(b\theta)$	$f_1 R$	0	R/f_1

Table 1. Basic expressions for the radial distance of the nodal line of the cylindrical-lens spiral in terms of the azimuthal parameter θ , $R(\theta)$; the cylindrical lens's focal length, f ; and the coefficients respectively multiplying the Alvarez-Lohmann-phase and azimuthal-phase-correction term, A and C . b is the winding parameter. The expressions for f , A and C depend only on R and the ratio of b and p , which we write as $f_1 = b/p$ (a quantity that happens to be the focal length of the cylindrical lens at radial distance $R = 1$). f , A and C are expressed in terms of R here, and can easily be expressed in terms of θ by replacing R with the corresponding expression for R in terms of θ . These formulae treat R , f , f_1 , A and C as dimensionless quantities but — as and when required — we can interpret R , f , and f_1 to have units meters and A and C to have units meters⁻². (b and θ are dimensionless.)

the parabolic lens, so while the AFL redirects the part of a light beam that passes through a particular winding like a lens of focal length F , it does not focus it. This is a shortcoming that can be overcome by adding to each winding a radial focusing power $1/F$, which can be achieved by adding to each winding of the components one part of a radial Alvarez-Lohmann lens [3, 4]. This *winding focusing* results in the radial phase profile shown in Fig. 2(b), which now has almost exactly the same phase gradient as the corresponding parabolic lens (Fig. 2(d)).

The spiral AFLs presented in Refs [6, 7] have a number of shortcomings. We call passage through corresponding windings *standard transmission*, and the components are designed to redirect light undergoing standard transmission. However, light can also undergo *winding hopping*, that is, pass through one winding of the first component and not through the corresponding winding of the second component, but a *different* (usually a neighbouring) winding of the second component (Fig. 1(d)). Winding hopping results in light being redirected incorrectly. In the case of thin, touching, components, light does not change transverse position upon transmission through the combination, and if the components are designed such that, for $\Delta\theta = 0$, any transverse position corresponds to the same winding on both components, winding hopping is minimised for $\Delta\theta = 0$, that is, when the focussing power of the adaptive Fresnel lens is zero (Fig. 1(c)). In other words, the design works best around zero focussing power. Another shortcoming is that the unwanted azimuthal phase gradient, unlike in a parabolic lens, is non-zero in the adaptive spiral Fresnel lens with $\Delta\theta \neq 0$, that is, whenever its focussing power is non-zero. Ref. [7] presents a modification that enables the Fresnel lenses to work best around non-zero focusing power, an alternative method of winding focusing by separating the components, and a web app [8] that allows the view through adaptive spiral Fresnel lenses to be experienced in augmented-reality (AR) style raytracing simulations. The raytracing simulations presented in Refs [6, 7] show that adaptive spiral Fresnel lenses work, at least ray-optically.

Here we perform initial wave-optical investigations of these devices. As this is easier to simulate, we consider each spiral AFL

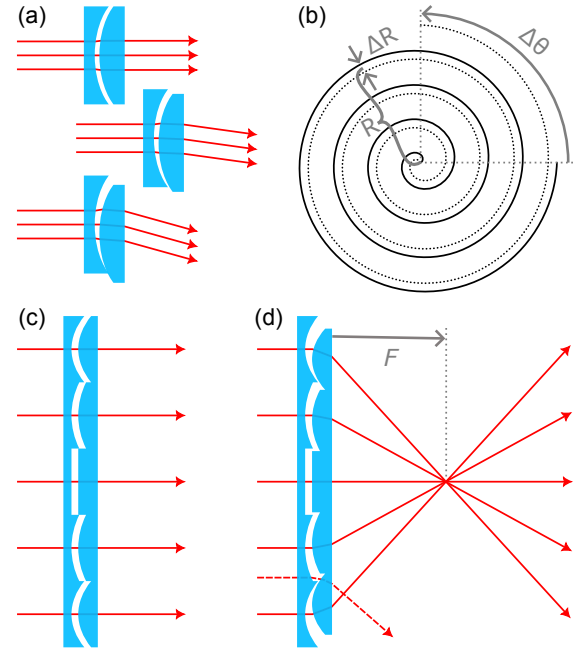


Fig. 1. Basic principle of the spiral AFL. (a) The direction change of light rays (red arrows) on transmission through complementary cylindrical lenses (cyan) depends on their lateral offset. (b) The radial separation, ΔR , between two rotated, but otherwise identical, spirals (here Archimedean spirals, one shown solid, the other dotted) depends on the relative rotation angle, $\Delta\theta$. (In general, although not in the example shown here, it also depends on the radial distance, R .) (c,d) Longitudinal, central, cross sections through two components comprising complementary cylindrical lenses whose nodal lines have been bent into identical spirals, shown for the case of unrotated (c) and rotated (d) components, which act approximately like lenses with focussing power 0 (c) and $1/F$ (d). The solid arrows indicate rays that pass through corresponding windings (we call this *standard transmission*); the dashed arrow shown in (d) shows an example of a ray that passes through one winding of the left component and a neighbouring winding of the right component, committing the sin of *winding hopping*.

component to be a thin, planar, phase hologram. In Sec. 2 we calculate the phase profiles of these spiral-AFL components and of the spiral AFLs themselves. In Sec. 3 we derive a modification of these phase profiles that is designed to remove the unwanted azimuthal phase gradient of spiral AFLs and which results in a close relationship to the BRM lens, which we discuss in Sec. 5. We study point-spread functions in Sec. 4. The effect of diffraction close to the centre of hyperbolic-spiral Fresnel lenses can be argued to lead to drastic blurring and to no blurring at all, and we settle this argument in Sec. 6. We conclude in Sec. 7.

2. PHASE PROFILE OF SPIRAL AFLS

The goal of adaptive-lens design is to create components that impart on transmitted light the same phase shifts — or at least the same phase gradients — as corresponding lenses of varying fo-




Fig. 2. Radial phase profile of a logarithmic-spiral AFL (a) without (dotted) and (b) with Alvarez-Lohmann (AL) winding focussing (dashed) and (c) with AL winding focussing and azimuthal phase compensation (solid), and (d) of a parabolic lens with the same focal length. In (a-c), the part of the curves that corresponds to one particular winding has been highlighted in red. The “steps” in (a,b) are actually steep phase gradients that correspond to winding hopping. The tiny, circled, lines in the bottom left quadrant are parts of the (discontinuous) curve (c) that correspond to winding hopping. The boundaries between windings in (a) and (b) are chosen such that the phase across the clear aperture of both components is continuous, which results in the boundaries being at different radial distances in the two cases. The curves were calculated for $\Delta\theta = 10^\circ$, $b = 0.1$, $p = 0.1$ diopters / rad, r plotted in the range $[0, 1 \text{ m}]$, the phase, Φ , is plotted in arbitrary units. For clarity, the curves are offset vertically.

cal lengths. Here we aim for our adaptive lens to correspond to a (thin, parabolic) lens of focal length F (or, equivalently, focussing power $P = 1/F$), which phase-shifts light of wavenumber k that passes through the lens at radial distance r by [9]

$$\Phi_L(r) = -\frac{kr^2}{2F} = -\frac{1}{2}kPr^2, \quad (5)$$

thereby adding a radial phase gradient

$$\frac{d\Phi_L(r)}{dr} = -kPr. \quad (6)$$

(Note that it is possible to design adaptive lenses that impart slightly different phase profiles, or indeed optical components that impart different phase profiles altogether. This is done in Ref. [5] and should also be possible for our AFLs.)

Our spiral AFLs comprise two components that each impart a spatially-dependent phase shift on transmitted light. We use a polar coordinate system in the plane of the component, and without loss of generality choose to restrict the values of the azimuthal coordinate φ to the range $(-\pi, \pi]$. This is convenient for calculating phase cross-sections numerically, as $(-\pi, \pi]$ is the range of values of the atan2 function [10]. We also define the *azimuthal parameter*, θ , as an azimuthal coordinate that is not restricted to the range $(-\pi, \pi]$ and which, for a position (r, φ) on the n th winding of the spiral, takes the value

$$\theta = \varphi + n2\pi. \quad (7)$$




Fig. 3. Structure and parametrisation of the components, shown here for components based on Archimedean spirals. The solid line is the nodal line (in the shape of an Archimedean spiral) of a cylindrical lens, the dashed line is the boundary between neighboring windings. The n th winding has been shaded (in red). φ is the azimuthal coordinate, restricted to the range $[-\pi, \pi]$; θ , an unrestricted azimuthal angle that is related to φ by the equation $\theta = \varphi + n2\pi$ for the n th winding, parametrises the spiral $R(\theta)$ that describes the nodal line. For one particular azimuthal angle φ ($-\pi \leq \varphi < \pi$), the radial coordinate (r) axis is shown, and the radial distances $R(\theta - 2\pi)$, $R(\theta)$, and $R(\theta + 2\pi)$ are marked on it.

It is this azimuthal parameter θ that we use in Table 1 to parametrise the spirals describing the nodal lines. Each value of θ corresponds to one particular azimuthal direction and one particular winding, specifically the azimuthal direction corresponding to azimuthal angle φ and the n th winding (Fig. 3). In combination, r and θ therefore describe a particular position (with polar coordinates (r, φ)) on a particular winding (namely the n th winding).

The phase shift imparted on transmitted light by component i ($i = 1, 2$) at (r, θ) can then be written as

$$\Phi_i(r, \theta) = \mp \frac{k(r - R(\theta))^2}{2f(\theta)} \pm \frac{A(\theta)k(r - R(\theta))^3}{6} \pm \frac{C(\theta)k}{2}, \quad (8)$$

where $R(\theta)$ and $f(\theta)$ are the radial coordinate of the nodal line and the focal length of the cylindrical lens corresponding to the azimuthal parameter θ ; $A(\theta)$ and $C(\theta)$ are θ -dependent coefficients; the upper and lower sign corresponds to $i = 1$ and $i = 2$, respectively; and where we have ignored any uniform phase offset. The expressions for $f(\theta)$ and $A(\theta)$ as well as the number n depend on the spiral shape, that is, on the expression for $R(\theta)$. Table 1 lists these expressions for three different spiral types, namely logarithmic, Archimedean, and hyperbolic spirals. The first, quadratic in $(r - R)$, term in Eqn (8) describes the “bent cylindrical lens”, more precisely the phase profile of a “radial lens” with radial focussing power $\pm 1/f(\theta)$ that is centred at radial coordinate $R(\theta)$. The second, cubic in $(r - R)$, term describes the two components (each sign describes one component) of a radial Alvarez-Lohmann-lens centred at radial coordinate $R(\theta)$ [3, 4]. The third term describes an azimuthal compensation. This azimuthal phase compensation is new, derived in Sec.




Fig. 4. Phase profile of a logarithmic-spiral AFL and its components. The images show grayscale representations (0=black, 2π = white) across a central square of sidelength 1 cm of the phase profiles of (a) component 1 and (b) component 2 on their own (calculated according to Eqn (8) with R , f and A according to the logarithmic-spiral-type row in Table 1) and (c) in combination (Eqn (9)), with areas in which winding hopping occurs shaded (in red). The inset in (c) shows the bottom right quadrant of the phase profile of a parabolic lens of focal length $F = 1$ m (Eqn (5)). The images were calculated using the open-source [11] wave-optics simulator *Young TIM* [12] with parameters $b = 0.1$, $p = 1$ diopter/ 60° , $\Delta\theta = 60^\circ$, $\lambda = 632.8$ nm. Like in all subsequent simulations, the winding boundary was chosen to be half-way between the nodal lines of neighbouring windings. Further details can be found in Supplement 1.

3, and in earlier sections we discuss our components without the azimuthal-phase-compensation term to represent our understanding prior to this paper. Section *Expressions for phase-profile terms* of Supplement 1 lists expressions for the individual terms in Eqn (8) for different spiral types. Fig. 4(a,b) shows grayscale representations of the functions Φ_1 and Φ_2 for an example of an adaptive logarithmic-spiral Fresnel lens (without azimuthal phase compensation).

The components are designed to be arranged such that they are perpendicular to the optical axis (and parallel to each other), touching, centred on the optical axis, and rotated (around the




Fig. 5. Phase profiles (central 1 cm \times 1 cm) of hyperbolic-spiral (a, c, e, i.e. left column) and Archimedean-spiral (b, d, f, i.e. right column) AFLs with $p = 1$ diopter/ 60° for rotation angles $\Delta\theta = 60^\circ$ (a, b; top row), 30° (c, d; middle row), and -30° (e, f; bottom row), resulting in focal lengths $F = 1/(p\Delta\theta) = 1$ m (a, b), 2 m (c, d), and -2 m (e, f). The insets in the images in the left column show the bottom right quadrant of the phase profile of a parabolic lens of the same focal length. In (a, b), the area where winding-hopping occurs has been shaded (in red). The hyperbolic-spiral phase profiles were calculated for $b = 10$, the Archimedean-spiral phase profiles for $b = 0.0002$. Supplement 1 contains further details.

optical axis) relative to each other by an azimuthal angle $\Delta\theta$. On transmission through both components, the phases due to the 1st and 2nd components, one rotated by $+\Delta\theta/2$, the other by $-\Delta\theta/2$, therefore add to

$$\Phi_{\text{AFL}}(r, \theta) = \Phi_1\left(r, \theta + \frac{\Delta\theta}{2}\right) + \Phi_2\left(r, \theta - \frac{\Delta\theta}{2}\right). \quad (9)$$

Fig. 4(c) shows the function Φ_{AFL} for an example of a logarithmic-spiral AFL, specifically the sum of the functions Φ_1 and Φ_2 shown in Fig. 4(a) and (b). Fig. 5 shows the phase profile of Archimedean-spiral and hyperbolic-spiral AFLs and different values of the rotation angle $\Delta\theta$. For comparison, the phase profiles of parabolic lenses of focal power P (given by Eqn (3)) are also shown.

The plots reveal differences between Φ_{AFL} and Φ_{L} , including the following:

1. The wrong radial phase gradient in those parts of the clear aperture that correspond to winding hopping (highlighted in Fig. 4(c)).
2. A lack of phase variation around the centre of the hyperbolic-spiral lens (Fig. 5(a, c, e)) and, to a lesser degree, of the logarithmic-spiral lens (Fig. 4(c)), caused by the steep phase gradients in winding-hopping areas limiting the phase values in precisely the same way in which the thickness discontinuities in a Fresnel lens limit its thickness.
3. A non-zero azimuthal component of the phase gradient, especially around the centre in the case of the Archimedean-spiral AFL (Fig. 5(b, d, f)), which leads to a distortion at the centre that is visible in the raytracing simulations shown in Ref. [7].

At the same time, there is also a clear similarity between Φ_{AFL} and Φ_{L} : in those parts of the clear aperture that correspond to transmission through corresponding windings, the radial phase gradients appear to be equal, at least approximately. Unless otherwise stated, it is those parts we consider in the following.

3. AZIMUTHAL PHASE COMPENSATION (APC)

Fig. 6(a) shows the phase profile of an Archimedean-spiral AFL. Two paths are highlighted to draw attention to two properties of the phase in those parts of the clear aperture that correspond to standard transmission:

1. On the spiral describing the nodal lines of the cylindrical lenses in the two components before rotation (the components are rotated by $\pm\Delta\theta/2$), i.e. on the line $R(\theta)$, the phase is constant, at least approximately.
2. Almost everywhere, specifically on the circle shown, the phase as a function of the azimuthal angle is a monotonic function. As we will see in Sec. 4, this results in darkness at the center of the point-spread function (PSF), which is undesirable for our purposes.

We discuss these below.

Whereas in the spiral AFL shown in Fig. 6(a) the phase on the spiral $R(\theta)$ is constant, on a parabolic lens with a finite focal length it is given by Eqn (5) with $r = R(\theta)$. This discrepancy can be eliminated by adding to the AFL the θ -dependent phase of a lens of $\Delta\theta$ -dependent focal power $P = p\Delta\theta$ (Eqn (3)) at radial distance $R(\theta)$, namely

$$-\frac{kp\Delta\theta R(\theta)^2}{2}. \quad (10)$$

The problem is that this corrective phase is proportional to $\Delta\theta$, but we can add precisely such a term (certainly for small angles $\Delta\theta$) to the AFL by adding an additional θ -dependent *azimuthal phase correction* (APC) $\pm\Phi_{\text{APC}}(\theta)$ to the two components, using the + sign for one component and the – sign for the other. For small rotation angles $\Delta\theta$ between them, the APCs then add to

$$\Phi_{\text{APC}}(\theta + \Delta\theta/2) - \Phi_{\text{APC}}(\theta - \Delta\theta/2) \approx \Phi'_{\text{APC}}(\theta)\Delta\theta, \quad (11)$$

where $\Phi'_{\text{APC}}(\theta)$ is the derivative of $\Phi_{\text{APC}}(\theta)$ with respect to θ , evaluated at θ . We want this to equal the expression given in Eqn (10), so

$$\Phi'_{\text{APC}}(\theta) = -\frac{kpR(\theta)^2}{2}, \quad (12)$$

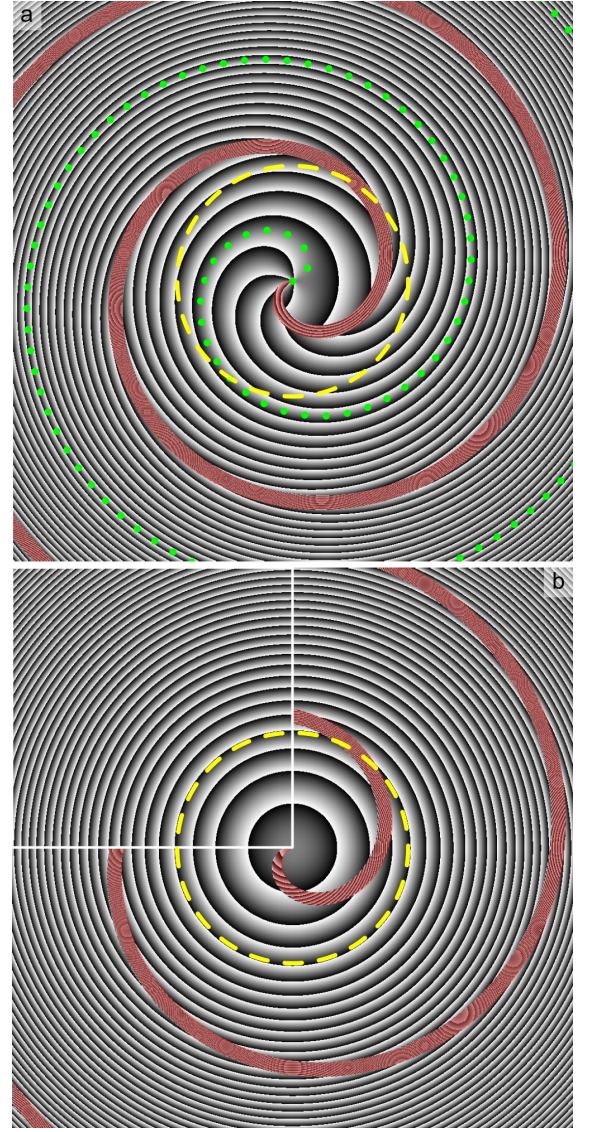


Fig. 6. Phase profile (central 1 cm × 1 cm) of an Archimedean-spiral AFL without (a) and with (b) azimuthal phase compensation (APC). The inset in (b) shows the parabolic phase cross-section of the upper left quadrant of a lens with the same focal length ($F = 1$ m). The dotted green line in (a) indicates the position of the nodal lines of the cylindrical lenses forming the two components before they were rotated by $+\Delta\theta/2$ and $-\Delta\theta/2$, respectively, $R(\theta)$, along which the phase can clearly be seen to be approximately constant. The azimuthal phase gradient on the dashed yellow circle can be seen to be positive (as the phase grows (black to white) counter-clockwise) in (a) and zero in (b). The figure was calculated for $b = 0.0005$, $p = 1$ diopter/ 30° , $\Delta\theta = 30^\circ$, resulting in $F = 1$ m. Supplement 1 lists further details.




Fig. 7. Example of the radial height profile, $h(r) = \Phi_1(r)/(k(n-1))$, of component 1, realised refractively, of different adaptive Fresnel lenses with $p = 1$ diopter/5°. (a) Archimedean-spiral AFL without APC (range ≈ 0.3 mm) and (b) with APC (drawn 10-fold vertically de-magnified; range ≈ 17.5 cm), and (c) BRM lens (range ≈ 1.4 cm). The curves were calculated for azimuthal angle $\varphi = \pi$. The refractive index of the material was $n = 1.5$.

or

$$\Phi_{\text{APC}}(\theta) = \int \Phi'_{\text{APC}}(\theta) d\theta \approx - \int \frac{kpR(\theta)^2}{2} d\theta. \quad (13)$$

Substituting the different expressions for $R(\theta)$ for the different spiral types, performing the integration, and ignoring (as usual) uniform phase offsets yields

$$\Phi_{\text{APC}}(\theta) = \frac{R(\theta)^n kp}{2nb}. \quad (14)$$

where $n = 1$ for the hyperbolic spirals, $n = 2$ for the logarithmic spirals, and $n = 3$ for the Archimedean spiral. As in Eqn (8), this additional phase term can alternatively be written in terms of the azimuthal-phase-compensation coefficients $C(\theta)$ (listed in Table 1) in the form

$$\Phi_{\text{APC}}(\theta) = \frac{k}{2} C(\theta), \quad (15)$$

which was already used in Eqn (8).

Perhaps surprisingly, in the case of hyperbolic-spiral and Archimedean-spiral AFLs the resulting component (in areas of the clear aperture where corresponding windings overlap) has the phase profile of a parabolic lens of focal length $P = p\Delta\theta$ (given by Eqn (3)), *exactly* (and in the case of logarithmic-spiral AFLs approximately). Fig. 6(b) shows an example.

In the case of diffractive components, addition of the APC results in the desired phase profile and represents a significant improvement of the PSF (Sec. 4). In the case of *refractive* components it can also significantly extend the range of the height profile of the components, in the example shown in Fig. 7 from ≈ 0.3 mm without APC to ≈ 17.5 cm with APC. This has a number of unwelcome (for our purposes) consequences, including




Fig. 8. Phase profile (1 cm \times 1 cm) of a three-part AFL. The figure was calculated for a combination of two Archimedean-spiral AFL, one with parameters $b = 0.0005$, $\Delta\theta = 30^\circ$, the other with $b = -0.0005$, $\Delta\theta = -30^\circ$, in both cases $p = 1$ diopter/60°. The focal length of each AFL was $F = 2$ m, that of the combination was 1 m. See Supplement 1 for further details.

increased thickness and the requirement for farther separation to enable the components to be rotated relative to each other, and while the components can be designed for non-zero separation between components [6], the fraction of winding-hopping light then depends on the incident light-ray direction and ultimately leads to a field-of-view limitation.

We briefly discuss a method to achieve partial azimuthal phase compensation, specifically the removal of the monotony of the (standard-transmission) phase as a function of the azimuthal angle that leads to the central intensity null in the PSF, that avoids the very significant height addition discussed above. The idea is to split the spiral AFL into two spiral AFLs such that the focussing powers of the two “half spiral AFLs” add up to that of the “full” spiral AFL and such that the (average) slope of the azimuthal phase gradient in the two half spiral AFLs have the opposite sign. Note that the two “inner” components can be combined into a single component which is then rotated relative to the two “outer” components, resulting in a three-component spiral AFL. Fig. 8 shows the phase profile behind an example of such a three-component spiral AFL. It can be seen that the phase pattern is symmetric with respect to mirroring on the horizontal axis, which implies that the azimuthal phase gradients at opposite positions relative to the horizontal axis have equal magnitudes but opposite signs. A number of shortcomings are clearly visible:

1. There are clear differences between the (standard-transmission) three-part spiral AFL and lens phase profiles.
2. The area that corresponds to winding hopping is almost twice that of a spiral SFL with the same value of the winding




Fig. 9. Power in the focal plane contained within radius r of the focal point for Archimedean-spiral AFLs (black curves, a-d) and, for comparison, a parabolic lens (gray curve, e) under Gaussian illumination. (a) without AL winding focussing and APC; (b) with AL winding and without APC; (c) with AL winding focussing and APC; (d) three-part AFL with AL winding focussing and without APC. All spiral-AFL components (a-d) and parabolic lens (e) were simulated as a thin phase holograms. Spiral-AFL parameters: $b = 0.000159$ (± 0.000159 in the case of the two spiral AFLs forming the three-component spiral AFL in (d)); $p = 1$ diopter/ 30° (a-c), $p = 0.5$ diopter/ 30° (d); $\Delta\theta = 30^\circ$ (a-c) and $\pm 30^\circ$ (d), in all cases resulting in a focal length $F = 1$ m. The components were placed in the waist plane of the illuminating Gaussian of waist size 3 mm and wavelength $\lambda = 632.8$ nm. The insets show the intensity profiles (and, in (a) only, also the phase profile) across the central 1 mm \times 1 mm in the focal plane. In the phase profile in (a), the positions of three charge-1 optical vortices near the centre are marked by red arrows. The simulations were performed using Young TIM [13]. Further details of the simulations can be found in Supplement 1.

parameter, b , and the same rotation angle, $\Delta\theta$.

In Sec. 4 we study the effect of these azimuthal-phase-compensation efforts on the PSF.

4. POINT-SPREAD FUNCTIONS

Fig. 9 allows a comparison of the point-spread functions (PSFs) of spiral AFLs whose components are constructed using only the first term in Eqn (8) (Fig. 9(a)), or the first two terms (the second term describes AL winding focussing; Fig. 9(b)), or all three terms (the third term describes APC; Fig. 9(c)). The PSFs without

(Fig. 9(a)) and with AL winding focussing (Fig. 9(b)) differ only very subtly in appearance, but a significant improvement due to AL winding focussing is clearly visible in the curve $P(r)$ that shows the power contained within a radius r from the focal point. Figs 9(a) and (b) were calculated without APC; adding APC (Fig. 9(c)) dramatically improves both the PSF's appearance and $P(r)$ curve, and both are very similar to their parabolic-lens counterparts (Fig. 9(e)).

APC adds an azimuthally dependent, monotonic, phase to the spiral AFL's phase profile, making the phase profile independent of the azimuthal angle. Without APC, a spiral AFL's phase profile therefore contains an azimuthally dependent, monotonic, phase (with the same magnitude as the APC but the opposite sign), which is reminiscent of a spiral phase plate [14], which is perhaps the simplest component with an azimuthally dependent, monotonic, phase profile (in fact, the phase profile is simply proportional to the azimuthal angle), and which is used to generate optical vortices, also known as phase singularities [15]. Three such phase singularities are clearly visible very close to the centre of the phase profile of the PSF in Fig. 9(a), and as optical vortices are intensity zeros they create a dark centre (Fig. 9(a,b)).

APC removes the azimuthal phase and with it the vortices, resulting in a PSF that is very similar to that of a parabolic lens (Fig. 9(e)), that is, essentially a diffraction-limited point. The difference between the PSFs in Fig. 9(c) and (e) is hard to spot in the intensity profiles but clearly visible in the $P(r)$ curves: while in the parabolic-lens PSF the central peak contains essentially all the power in the beam (which means that $P(r)$ reaches ≈ 1 at a relatively small value of r), in the case of the PSF of the spiral AFL with APC the central peak contains essentially only the power of any light that underwent standard refraction (minus any power that was diffracted outwards due to the spiral shape of the clear-aperture area corresponding to standard refraction acting essentially as an aperture, plus the power corresponding to any winding-hopping light "accidentally" refracted into the PSF's central peak).

Fig. 9(d) relates to the PSF of the three-part spiral AFL discussed in Sec. 3. It was constructed from two 2-part spiral AFLs such that the azimuthal phases due to the two cancel (on average). Like in the case of the 2-part spiral AFL with APC, this removes the vortices at the centre of the APC, resulting in a relatively narrow central peak in the PSF. Closer inspection also shows up a number of (more or less complete) rings surrounding the central peak, which show up very clearly in the $P(r)$ curve as (rounded) "steps".

APC clearly improves the PSF of spiral AFLs, provided the elements are diffractive. Aside from issues related to winding hopping, APC effectively makes diffractive spiral AFLs perfect, certainly in the case of Archimedean-spiral and hyperbolic-spiral lenses. However, because of the added thickness when realised refractively (see Fig. 7), APC is less interesting for our purposes; for this reason, it is not spiral AFLs *with* APC we investigate further below, but spiral AFLs *without* APC.

In the *Mini PSF catalog* section in Supplement 1 we show plots of the PSFs and $P(r)$ curves of spiral AFLs without APC for different spiral types and different values of the winding parameter, b . In the case of the Archimedean-spiral and logarithmic-spiral shapes, the $P(r)$ curve arguably has the best characteristics at some intermediate value of the winding parameter, b , whereas in the case of hyperbolic-spiral lenses the characteristics of the $P(r)$ curves improve as b increases. Fig. 10 shows what are, out of the values we tried, arguably the best $P(r)$ curves for the different spiral shapes. It suggests that hyperbolic-spiral AFLs




Fig. 10. Best (arguably) power-in-radius curves in the focal plane for the different spiral shapes under Gaussian illumination. The insets show a grayscale representation of the central $1\text{ cm} \times 1\text{ cm}$ of the phase profiles of the corresponding spiral AFLs with AL winding focussing (but without APC), which were calculated for parameters $b = 0.00016$ (Archimedean), $b = 0.08$ (logarithmic), and $b = 10^{10}$ (hyperbolic). In all cases $p = 1\text{ diopter}/30^\circ$ and $\Delta\theta = 30^\circ$, resulting in a focal length $F = 1\text{ m}$. The waist size of the illuminating Gaussian ($\lambda = 632.8\text{ nm}$) was 3 mm ; the spiral AFLs were placed in the plane of the beam waist. See Supplement 1 for further details of the simulations.

might have the best imaging characteristics.

5. RELATIONSHIP TO THE BRM LENS

Once all three terms in Eqn (8) — the cylindrical-lens term, the Alvarez-Lohmann term, and the azimuthal-phase-correction term — have been added together, what is the resulting phase profile? In the case of logarithmic-spiral AFL, for which the phase profile is that of the corresponding parabolic lens only approximately, the phase profile of the components is somewhat complicated, but in the case of the hyperbolic- and Archimedean-spiral AFL it becomes surprisingly simple: in the case of the hyperbolic-spiral AFL, the phase profile of component i is

$$\Phi_i(r, \theta) = \pm \frac{kr}{f_1} \pm \frac{kp}{2} r^2 \theta, \quad (16)$$

for the Archimedean-spiral AFL it is

$$\Phi_i(r, \theta) = \mp \frac{kr^3}{3f_1} \pm \frac{kp}{2} r^2 \theta. \quad (17)$$

In both cases, the first terms of the first and second component, by virtue of being independent of θ and differing only in sign, add up to zero and therefore do not contribute to the phase profile of the resulting AFL, Φ_{AFL} (Eqn (9)). The second, θ -dependent, terms are the same in both cases, and they are precisely the phase profiles of the components of the BRM lens,

$\Phi_{\text{BRM},i}$ (Eqn (1)), with $a = kp/2$. What we have therefore invented in a roundabout way is a simple variation on the BRM lens.

But the relationship between spiral AFLs and the BRM lens goes further: as can be seen in the *Expressions for phase-profile terms* section in Supplement 1, the phase profile of the components of a hyperbolic-spiral AFL without APC is

$$\pm \frac{1}{2} kp \theta \left(r + \frac{1}{b\theta} \right)^2, \quad (18)$$

and in the limit $b \rightarrow \infty$ this becomes

$$\pm \frac{1}{2} kp \theta r^2, \quad (19)$$

which is again the phase profile of the components of the BRM lens. In other words, in the limit $b \rightarrow \infty$ the hyperbolic-spiral AFL becomes the BRM lens. This, of course, explains why the power-in-radius characteristics of a hyperbolic-spiral AFL with a large b value, such as the one shown in Fig. 10, are so good.

Given the very close relationship between spiral AFLs and the BRM lens we expect the equivalent of winding hopping to occur in the BRM lens. Indeed it does, but with a simplification: whereas the area that corresponds to winding hopping in spiral AFLs is spiral-shaped, in the BRM lens it is in the shape of a sector. Bernet and Ritsch-Marte provide a way of removing this sector, namely clever scaling of the phase on each circle centred on the lens centre by a factor, chosen to be as close to 1 as possible, that scales the phase jump at the phase discontinuity to become an integer multiple of 2π , thereby eliminating the phase discontinuity. This approach can also be applied to hyperbolic-spiral and Archimedean-spiral AFLs with APC and to hyperbolic-spiral AFLs without APC but with $b \gg 1$, where it should work just as well as it does in the BRM lens, or to logarithmic-spiral AFLs with APC, where we expect it to work, but less well. But note that this scaling works best for phase cross-sections comprising an azimuthal phase ramp; spiral AFLs without APC do not have such an azimuthal phase ramp, and we therefore do not expect the scaling to work in those.

Given that our spiral AFLs with APC are arguably merely more complicated versions of the BRM lens, and that, in the limit $b \rightarrow \infty$, logarithmic-spiral AFLs are BRM lenses, should we keep investigating spiral AFLs? We believe the answer is “Yes”, and we discuss two reasons below.

First, spiral AFLs that are less closely related to the BRM lens are an alternative adaptive Fresnel-lens design with lower image quality than the BRM lens (see Sec. 4), but in many cases also a significantly thinner form factor when realised refractively. This is not always the case (reducing the clear aperture significantly reduces the height-profile range of the BRM lens more than that of an Archimedean-spiral AFL), but in the example shown in Fig. 7, calculated for parameters relevant to application in eye glasses, the range of the height profile of an Archimedean-spiral AFL (without APC) is $\approx 0.3\text{ mm}$ while that of an equivalent BRM lens is $\approx 1.4\text{ cm}$, so while the lens in the first case can be very thin, that in the second case is very thick.

Second, we could design the components such that APC is added only where it matters the most, thereby improving the imaging quality but restricting the thickness gain. One particularly promising modification involves multiplying the APC factor by a θ -dependent factor that varies between 0 and 1, chosen to take the value 0 where the optical quality is good even without APC and where the addition of the APC would add




Fig. 11. Periodic array (period w) of complementary (focal lengths $+f$ and $-f$), shifted relative to each other (by ΔR), pairs of cylindrical lenses. For clarity, one pair of complementary, shifted, cylindrical lenses is highlighted. We use such an array to describe the local structure of a hyperbolic-spiral AFL without AL winding focussing and without APC near its centre, in which case the cylindrical-lens arrays are periodic in the r direction and the cylinder axes are tangential to the φ direction.

significant thickness to the components (like far from the centre in an Archimedean-spiral AFL) and 1 where the APC significantly improves the optical quality without adding significant thickness (like close to the centre of an Archimedean-spiral AFL).

6. DIFFRACTION EFFECTS NEAR THE CENTER OF HYPERBOLIC-SPIRAL AFLS

In logarithmic- and hyperbolic-spiral AFLs, the winding width, defined as

$$w(\theta) = |R(\theta - \pi) - R(\theta + \pi)|, \quad (20)$$

asymptotically approaches zero in proportion to r and r^2 , respectively, as the centre is approached (i.e. as $r \rightarrow 0$). This causes practical challenges (for example, to align corresponding windings requires very precise alignment of the components) and it raises questions about diffraction.

We concentrate on hyperbolic-spiral AFLs without AL winding focussing and without APC. The reason is that the local phase modulation near the centre is approximately that behind complementary, but shifted, periodic arrays of straight, identical, cylindrical lenses as shown in Fig. 11 (see section *Calculation related to diffraction effects near the center of hyperbolic-spiral AFLs* in Supplement 1), and a mathematical description of the field in a plane immediately behind such an array, and of its far-field diffraction pattern, can be obtained relatively easily. The complex field in the far field is the product of a delta comb and an envelope, whereby the delta comb is due to the periodicity, an idealisation that implies that the array is infinite, which it is not in reality, but *locally* the infinite array and the actual array are almost identical. We can use a top-hat function to pick out the part where the two arrays are almost identical; after convolution with the Fourier transform of this top-hat function, which replaces the delta peaks with (relatively narrow) sinc functions, the complex field represents the far field due to the part of the array picked out by the top-hat function.

Fig. 12 shows the structure of the far-field intensity distribution in the radial direction — the radial power spectrum — behind two periodic arrays of cylindrical lenses, one comprising cylindrical lenses of focal length $+f$, the other comprising cylindrical lenses of focal length $-f$ and shifted relative to the former by ΔR in the r direction (see Fig. 11), when the combination is




Fig. 12. Structure of the radial wavenumber (k_r) power spectrum behind the cylindrical-lens array shown in Fig. 11 under uniform-plane-wave illumination with wavelength λ . The function is a (squared) Dirac comb of period λ/w , centred on the origin, with an envelope given by the function $\bar{u}_1^2(k_r)$ (dashed curve), derived in Supplement 1 (Eqn (S39)). The function $\bar{u}_1(k_r)$ comprises two terms proportional to sinc functions, respectively centred on $k_r/k = -\Delta R/f$ and corresponding to standard transmission and centred on $k_r/k = (w - \Delta R)/f$ and corresponding to winding hopping. (The visible modulation of the central peak of the peak centred at $k_r/k = (w - \Delta R)/f$ is due to the outer lobes of the sinc function centred at $-\Delta R/f$.) The graph was calculated for $w = 1$ mm, $\Delta R = 0.3$ mm, $f/k = 1.05 \times 10^{-8}$ m².

illuminated by a uniform plane wave of wavelength λ . This power spectrum comprises a comb of discrete radial wavenumbers modulated with an envelope whose two main peaks, due to standard transmission and winding hopping, respectively pick out spatial frequencies around $k_r/k = -\Delta R/f$ and around $k_r/k = (w - \Delta R)/f$.

We are mostly interested in the peak due to standard transmission. Here are a few of its properties and their consequences:

1. The peak is centred on $k_r/k = -\Delta R/f$, but according to the radial-offset equation (Eqn (4)) $\Delta R/f = R/F = RP$, and so the peak is centred on radial wavenumber $k_r = -kPR$. Wavenumber is phase gradient, and indeed the radial phase gradient added by a lens at radial distance R is (according to Eqn (6)) the same.
2. The width of the peak, $2\lambda/(w - |\Delta R|)$, is greater than, or equal to, two times the separation between neighbouring discrete wavenumbers in the comb, λ/w , which means that there are always ≥ 2 discrete wavenumbers present in the peak.
3. The spread in radial wavenumbers due to the width of the peak, $\Delta k_r/k = 2\lambda/(w - |\Delta R|)$, leads to radial angular blurring. We can roughly estimate this to be $\beta \approx \Delta k_r/k = 2\lambda/(w - |\Delta R|)$ which, for the case of small rotation angles $\Delta\theta$ (and thus radial offsets $|\Delta R| \ll w$) becomes $\beta \approx 2\lambda/w$. In a hyperbolic-spiral AFL, $w(R) \approx 2\pi bR^2$, so in a hyperbolic-spiral AFL with $b = 6$ (for which $w(5 \text{ mm}) \approx 1$ mm) we can expect blurring greater than $\beta = 1^\circ$ for $R \lesssim \sqrt{\lambda/(\pi b \beta)} \approx 1.3$ mm, and greater than $\beta = 1'$ (one arc minute, the human eye's angular resolution) for $R \lesssim 1$ cm.

The other peak, due to winding hopping, is normally undesired, as it corresponds to an unwanted additional image (or even more than one image, if light also “hops” more than one winding). It might therefore be considered advantageous if such light becomes evanescent and gets reflected, so that the additional image is no longer transmitted. If $\Delta R \ll w$, this happens when $R \gtrsim (2\pi p)^{-1}$; for example, in a hyperbolic-spiral AFL with $p = 1$ diopter/30°, this corresponds to $R \gtrsim 1.3$ cm.

7. SUMMARY AND CONCLUSIONS

We have presented an initial wave-optics investigation of spiral AFLs. The main motivation for designing spiral AFLs in the first place was the desire to build relatively thin adaptive lenses from *refractive* components, but to simplify our simulations we treated the components as thin (planar) *diffractive* components.

We found that the main difference in the phase profile of parabolic lenses and spiral AFLs is an azimuthal phase term. We introduced azimuthal phase correction (APC), a modification to the spiral-AFL components that adds this azimuthal phase term, thereby making the phase profile of spiral AFLs much more similar to that of parabolic lenses. Unfortunately, APC adds significant thickness to refractive components.

Our more surprising findings include the close links between spiral AFLs and the Bernet-Ritsch-Marte (BRM) lens. APC makes the azimuthal phase profiles of the components of Archimedean- and hyperbolic-spiral AFLs identical to those of BRM-lens components. (The radial phase profiles differ, but in the case of diffractive components this has no effect.) Most notably, in the limit $b \rightarrow \infty$, hyperbolic-spiral AFLs and BRM lenses are the same.

We calculated point-spread functions (PSFs) of a range of spiral AFLs. In the case of spiral AFLs that are more spiral-AFL-like than BRM-lens-like (Archimedean- and logarithmic-spiral AFLs and hyperbolic-spiral AFLs with “small” values of b , in all cases without APC), these contain optical vortices that widen the PSF. In the case of BRM-lens-like spiral AFLs, as well as BRM lenses, the PSF comprises what is essentially a diffraction-limited spot, formed by light undergoing standard refraction, and a background due to light undergoing winding hopping.

Our results show that the BRM lens is the better diffractive lens, but often spiral AFLs without APC can be significantly thinner. Spiral AFLs therefore offer an attractive compromise for applications in which thinness is required.

In future we plan to develop this work in a number of different directions.

First, we intend to build refractive prototypes of different spiral AFLs and characterise their performance.

Second, we want to investigate diffraction effects due to the individual components, starting from the mathematical expression for the far-field diffraction pattern behind an array of microlenses under uniform-plane-wave illumination (Eqn (9) in Ref. [16]). Such an investigation would concentrate on, but not be limited to, the possibility and consequences of light being evanescent *between* the components.

Third, we will generalise our design to allow greater separation between the two components, a requirement for being able to rotate the components relative to each other. A number of approaches have been discussed in Ref. [6].

Fourth, we plan to consider design variations, for example “Fresnelizing” (i.e. introducing steps into, and thereby reducing the range of) the thickness profile of spiral AFLs with APC,

ideally in a way that still allows the components to be rotated relative to each other in both directions.

Further down the list is relating the field behind spiral AFLs to relevant published results, for example the far-field diffraction pattern behind a slit shaped like an Archimedean spiral [17], the presence of a vortex behind a slit shaped like a single winding of an Archimedean spiral [18], and the field due to logarithmic spiral antennas [19].

REFERENCES

1. G. Li, “Adaptive lens,” *Prog. Opt.* **55**, 199–283 (2010).
2. M. A. Geday, M. C. no García, J. M. Otón, and X. Quintana, “Adaptive spiral diffractive lenses — lenses with a twist,” *Adv. Opt. Mat.* **8**, 2001199 (2020).
3. L. W. Alvarez, “Two-element variable-power spherical lens,” U. S. Patent 3,305,294 (1967).
4. A. W. Lohmann, “A new class of varifocal lenses,” *Appl. Opt.* **9**, 1669–1671 (1970).
5. S. Bernet and M. Ritsch-Marte, “Adjustable refractive power from diffractive moiré elements,” *Appl. Opt.* **47**, 3722–3730 (2008).
6. I. Armstrong, M. Locher, and J. Courtial, “Adaptive Fresnel lens: basic theory,” *J. Opt. Soc. Am. A* **42**, 211–220 (2025).
7. M. Locher, D. Wu, and J. Courtial, “Adaptive spiral Fresnel lens: generalisations, improvements, and augmented-reality simulations,” in preparation (2025).
8. University of Glasgow, “SpiralFresnelFrenzy,” <https://jkcuk.github.io/SpiralFresnelFrenzy/> (2024).
9. J. W. Goodman, *Introduction to Fourier Optics* (McGraw-Hill, New York, 1996), chap. 5.1.3, 2nd ed.
10. Wikipedia, “atan2,” <http://en.wikipedia.org/wiki/Atan2>.
11. “Young TIM source code,” <https://github.com/jkcuk/YoungTIM> (2023).
12. “Young TIM 1.1.0,” (2025).
13. S. Leavey and J. Courtial, “Young TIM: A wave-optics simulator with slightly special powers,” submitted for publication (2016).
14. S. S. R. Oemrawsingh, J. A. W. van Houwelingen, E. R. Eliel, J. P. Woerdman, E. J. K. Verstegen, J. G. Kloosterboer, and G. W. ’t Hooft, “Production and characterization of spiral phase plates for optical wavelengths,” *Appl. Opt.* **43**, 688–694 (2004).
15. M. Berry, “Singularities in waves and rays,” in *Les Houches Lecture Series Session XXXV*, R. Balian, M. Kléman, and J.-P. Poirier, eds. (North-Holland, Amsterdam, 1981), pp. 453–543.
16. T. F. Q. Iversen and S. G. Hanson, “Geometrical characterization of micro-optical cylindrical lens arrays using angular resolved diffraction,” *Meas. Sci. Technol.* **21**, 055301 (2010).
17. R. N. Bracewell and J. D. Villaseñor, “Fraunhofer diffraction by a spiral slit,” *J. Opt. Soc. Am.* **7**, 21–25 (1990).
18. V. Logachev, “Simulation of the vortex beams formation during diffraction by a thin spiral aperture,” in *2022 VIII International Conference on Information Technology and Nanotechnology (ITNT)*, (Samara, Russian Federation, 2022), pp. 1–5.
19. P. E. Mayes, “Frequency-independent antennas and broad-band derivatives thereof,” *Proc. IEEE* **80**, 103–112 (1992).

Bibliography

- [1] M. Locher, “Mathematica figure and calculation bundle.” <https://doi.org/10.6084/m9.figshare.29114723>, 5 2025. Figshare.
- [2] S. S. Welsh and J. Courtial, “View rotation with parallel ray-rotating windows,” *Optics Communications*, vol. 285, no. 24, pp. 4802–4806, 2012.
- [3] M. Locher, “Dr Tim based runnable ray-trace software bundle.” <https://doi.org/10.6084/m9.figshare.29109590>, 5 2025. Figshare.
- [4] S. Oxburgh, T. Tyc, and J. Courtial, “Dr tim: Ray-tracer tim, with additional specialist scientific capabilities,” *Computer Physics Communications*, vol. 185, pp. 1027–1037, 3 2014.
- [5] J. Courtial, “Spiralfresnelfrenzy.” <https://github.com/jkcuk/SpiralFresnelFrenzy>, 2024. GitHub repository.
- [6] J. Běĺín, T. Tyc, M. Grunwald, S. Oxburgh, E. N. Cowie, C. D. White, and J. Courtial, “Ideal-lens cloaks and new cloaking strategies,” *Optics Express*, vol. 27, p. 37327, 12 2019.
- [7] M. Locher, “Shifty cloak supporting videos.” <https://doi.org/10.6084/m9.figshare.29114714>, 5 2025. Figshare.
- [8] M. Locher, “Ideal lens wormhole,” 01 2021. Unpublished BSc dissertation, University of Glasgow.
- [9] K. Iga, Y. Kokubun, and M. Oikawa, “Fundamentals of microoptics,” 1984. New York: Academic Press.
- [10] L. Hornbeck, “Current status of the digital micromirror device (dmd) for projection television applications,” in *Proceedings of IEEE International Electron Devices Meeting*, pp. 381–384, 1993.
- [11] H. Zappe, *Fundamentals of Micro-Optics*, p. 3–16. Cambridge University Press, 2010.
- [12] D. Gabor, “A new microscopic principle,” *Nature*, vol. 161, pp. 777–778, 5 1948.

- [13] K. Bang, Y. Jo, M. Chae, and B. Lee, “Lenslet vr: Thin, flat and wide-fov virtual reality display using fresnel lens and lenslet array,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 5, pp. 2545–2554, 2021.
- [14] K.-Y. Huang, C.-F. Wang, D.-L. Shih, K.-P. Chang, Y.-P. Chang, C.-N. Liu, and W.-H. Cheng, “Scheme of flash lidar employing glass aspherical microlens array with large field of illumination for autonomous vehicles,” *Optics Express*, vol. 32, pp. 35854–35870, 9 2024.
- [15] C. Hembd-Sölner, R. F. Stevens, and M. C. Hutley, “Imaging properties of the gabor superlens,” *Journal of Optics A: Pure and Applied Optics*, vol. 1, p. 94, 1 1999.
- [16] J. Courtial, T. Tyc, J. Bělin, S. Oxburgh, G. Ferenczi, E. N. Cowie, and C. D. White, “Ray-optical transformation optics with ideal thin lenses makes omnidirectional lenses,” *Optics Express*, vol. 26, p. 17872, 7 2018.
- [17] J. Courtial and T. Tyc, “Generalized laws of refraction that can lead to wave-optically forbidden light-ray fields,” *Journal of the Optical Society of America A*, vol. 29, pp. 1407–1411, 7 2012.
- [18] A. C. Hamilton and J. Courtial, “Metamaterials for light rays: ray optics without wave-optical analog in the ray-optics limit,” *New Journal of Physics*, vol. 11, p. 013042, 1 2009.
- [19] G. K. V. Noorden and E. C. E. C. Campos, *Binocular vision and ocular motility : theory and management of strabismus*. Mosby, 2002.
- [20] Z. Georgievski, M. Sleep, and K. Koklanis, “Simulated torsional disparity disrupts horizontal fusion and stereopsis,” *Journal of AAPOS*, vol. 11, pp. 120–124, 4 2007.
- [21] S. R. Hatt, D. A. Leske, E. A. Bradley, S. R. Cole, and J. M. Holmes, “Comparison of quality-of-life instruments in adults with strabismus,” *American Journal of Ophthalmology*, vol. 148, pp. 558–562, 10 2009.
- [22] H. B. McBain, C. K. Au, J. Hancox, K. A. MacKenzie, D. G. Ezra, G. G. Adams, and S. P. Newman, “The impact of strabismus on quality of life in adults with and without diplopia: A systematic review,” *Survey of Ophthalmology*, vol. 59, pp. 185–191, 3 2014.
- [23] S. J. Woo, J. M. Seo, and J. M. Hwang, “Clinical characteristics of cyclodeviation,” *Eye*, vol. 19, pp. 873–878, 2005.
- [24] L. B. D. Lott, K. A. Kerber, P. P. Lee, D. L. Brown, and J. F. Burke, “Diplopia-related ambulatory and emergency department visits in the united states, 2003-2012,” *JAMA Ophthalmology*, vol. 135, pp. 1339–1344, 12 2017.

- [25] M. Harada and Y. Ito, “Surgical correction of cyclotropia,” *Japanese Journal of Ophthalmology*, vol. 8, pp. 88–96, 1964.
- [26] Önder Ayyıldız, F. M. Mutlu, M. Küçükevcilioğlu, G. Gökçe, and H. İbrahim Altınsoy, “Clinical features and surgical results in harada-ito surgery patients,” *Turkish Journal of Ophthalmology*, vol. 48, pp. 267–273, 10 2018.
- [27] D. A. Iliescu, C. M. Timaru, N. Alexe, E. Gosav, A. D. Simone, M. Batras, and C. Stefan, “Management of diplopia,” *Romanian Journal of Ophthalmology*, vol. 61, pp. 166–170, 9 2017.
- [28] Y. S. Bradfield, M. C. Struck, B. J. Kushner, D. E. Neely, D. A. Plager, and R. E. Gangnon, “Outcomes of harada-ito surgery for acquired torsional diplopia,” *Journal of AAPOS*, vol. 16, pp. 453–457, 2012.
- [29] S. Flodin, P. Karlsson, A. Rydberg, M. A. Grönlund, and T. Pansell, “Surgical outcome of graded harada-ito procedure in the treatment of torsional diplopia – a retrospective case study with long-term results,” *Strabismus*, vol. 30, pp. 8–17, 2022.
- [30] A. M. Miller, “Torsional diplopia,” *American Orthoptic Journal*, vol. 65, pp. 21–25, 2015.
- [31] J. Courtial, N. Chen, S. Ogilvie, B. C. Kirkpatrick, A. C. Hamilton, G. M. Gibson, T. Tyc, E. Logean, and T. Scharf, “Experimental demonstration of ray-rotation sheets,” *Journal of the Optical Society of America A*, vol. 35, p. 1160, 7 2018.
- [32] A. C. Hamilton, B. Sundar, J. Nelson, and J. Courtial, “Local light-ray rotation,” *Journal of Optics A: Pure and Applied Optics*, vol. 11, 2009.
- [33] D. W. Swift, “Image rotation devices — a comparative survey,” *Optics & Laser Technology*, vol. 4, pp. 175–188, 1972.
- [34] M. C. Hutley, R. Hunt, R. F. Stevens, and P. Savander, “The moire magnifier,” *Pure and Applied Optics: Journal of the European Optical Society Part A*, vol. 3, p. 133, 3 1994.
- [35] J. Courtial, “Standard and non-standard metarefraction with confocal lenslet arrays,” *Optics Communications*, vol. 282, pp. 2634–2641, 7 2009.
- [36] H. Snellen, *Probebuchstaben zur Bestimmung der Sehschärfe*. Utrecht, 1862.
- [37] E. R. Tkaczyk, “Vectorial laws of refraction and reflection using the cross product and dot product,” *Opt. Lett.*, vol. 37, pp. 972–974, 3 2012.
- [38] Wolfram Research Inc., “Mathematica, Version 14.1.” Champaign, IL, 2024.

- [39] M. Locher, “Supporting simulations for view rotator pixel and aperture size.” <https://doi.org/10.6084/m9.figshare.30009886>, 8 2025. Figshare.
- [40] D. A. Cronin, E. H. Hall, J. E. Goold, T. R. Hayes, and J. M. Henderson, “Eye movements in real-world scene photographs: General characteristics and effects of viewing task,” *Frontiers in Psychology*, vol. 10, 1 2020.
- [41] L. R. Young and D. Sheena, “Survey of eye movement recording methods,” *Behavior Research Methods & Instrumentation*, vol. 7, pp. 397–429, 1975.
- [42] C. Paterson, “Computer-generated diffractive optics with spiral phase dislocations,” *Imperial College London*, pp. 21–23, 1996. PhD thesis.
- [43] O. Reshef, M. P. DelMastro, K. K. Bearne, A. H. Alhulaymi, L. Giner, R. W. Boyd, and J. S. Lundeen, “An optic to replace space and its application towards ultra-thin imaging systems,” *Nature Communications*, vol. 12, 12 2021.
- [44] R. Lazar, J. Degen, A. S. Fiechter, A. Monticelli, and M. Spitschan, “Regulation of pupil size in natural vision across the human lifespan,” *Royal Society Open Science*, vol. 11, 6 2024.
- [45] P. Benitez, “Simultaneous multiple surface optical design method in three dimensions,” *Optical Engineering*, vol. 43, 7 2004.
- [46] E. Eaglesham, “Derivative control of novel generalised refractive optics,” 03 2024. Unpublished MSci dissertation, University of Glasgow.
- [47] K. Shastri, O. Reshef, R. W. Boyd, J. S. Lundeen, and F. Monticone, “To what extent can space be compressed? bandwidth limits of spaceplates,” *Optica*, vol. 9, pp. 738–745, 7 2022.
- [48] Mr.doob, “Three.js.” <https://github.com/mrdoob/three.js>, 3 2021. GitHub repository.
- [49] D. Pascolini and S. P. Mariotti, “Global estimates of visual impairment: 2010,” *British Journal of Ophthalmology*, vol. 96, pp. 614–618, 5 2012.
- [50] H. Hashemi, A. Fotouhi, A. Yekta, R. Pakzad, H. Ostadimoghaddam, and M. Khabazkhoob, “Global and regional estimates of prevalence of refractive errors: Systematic review and meta-analysis,” *Journal of Current Ophthalmology*, vol. 30, p. 3, 3 2017.
- [51] V. S. Gudlavalleti, K. P. Allagh, and A. S. Gudlavalleti, “Self-adjustable glasses in the developing world,” *Clinical Ophthalmology*, vol. 8, pp. 405–413, 2 2014.

- [52] T. C. Kriss and V. M. Kriss, “History of the operating microscope: from magnifying glass to microneurosurgery,” *Neurosurgery*, vol. 42, pp. 899–908, 4 1998.
- [53] The Editors of Encyclopaedia Britannica, “Fresnel lens,” *Encyclopedia Britannica*, 3 2025.
- [54] A. Fresnel, “Memoir upon a new system of lighthouse illumination,” *U.S. Lighthouse Society*, 7 1822.
- [55] T. R. Fricke, B. A. Holden, D. A. Wilson, G. Schlenther, K. S. Naidoo, S. Resnikoff, and K. D. Frick, “Global cost of correcting vision impairment from uncorrected refractive error,” *Bulletin of the World Health Organization*, vol. 90, pp. 728–738, 10 2012.
- [56] M. G. Douali and J. D. Silver, “Self-optimised vision correction with adaptive spectacle lenses in developing countries,” *Ophthalmic and Physiological Optics*, vol. 24, pp. 234–241, 2004.
- [57] N. Hasan, A. Banerjee, H. Kim, C. H. Mastrangelo, W. Tasman, E. A. Jaeger, L. Johnson, J. G. Buckley, A. J. Scally, and D. B. Elliott, “Tunable-focus lens for adaptive eyeglasses,” *Optics Express*, vol. 25, pp. 1221–1233, 1 2017.
- [58] C. W. Fowler and E. S. Pateras, “Liquid crystal lens review,” *Ophthalmic & physiological optics : the journal of the British College of Ophthalmic Opticians (Optometrists)*, vol. 10, pp. 186–194, 1990.
- [59] L. W. Alvarez and W. E. Humphrey, “Variable-power lens and system,” 1970. US Patent, US3507565A.
- [60] A. W. Lohmann, “A new class of varifocal lenses,” *Appl. Opt.*, vol. 9, pp. 1669–1671, 7 1970.
- [61] Optotune, “Focus tunable lenses.” <http://www.optotune.com/technology/focus-tunable-lenses>, 2016.
- [62] Adlens, “Adjustable lens glasses: How they work.” <https://adlens.com/how-it-works/>, 2017.
- [63] S. Bernet and M. Ritsch-Marte, “Adjustable refractive power from diffractive moiré elements,” *Appl. Opt.*, vol. 47, pp. 3722–3730, 7 2008.
- [64] I. Armstrong, M. Locher, and J. Courtial, “Adaptive fresnel lens: basic theory,” *Journal of the Optical Society of America A*, vol. 42, pp. 211–220, 2 2025.
- [65] L. Galinier, P. Renaud-Goud, J. Brusau, L. Kergadallan, J. Augereau, and B. Simon, “Spiral diopter: freeform lenses with enhanced multifocal behavior,” *Optica*, vol. 11, p. 238, 2 2024.

- [66] M. McCall, J. B. Pendry, V. Galdi, Y. Lai, S. A. Horsley, J. Li, J. Zhu, R. C. Mitchell-Thomas, O. Quevedo-Teruel, P. Tassin, V. Ginis, E. Martini, G. Minatti, S. Maci, M. Ebrahimpouri, Y. Hao, P. Kinsler, J. Gratus, J. M. Lukens, A. M. Weiner, U. Leonhardt, I. I. Smolyaninov, V. N. Smolyaninova, R. T. Thompson, M. Wegener, M. Kadic, and S. A. Cummer, “Roadmap on transformation optics,” *Journal of Optics (United Kingdom)*, vol. 20, 6 2018.
- [67] J. B. Pendry, D. Schurig, and D. R. Smith, “Controlling electromagnetic fields,” *Science*, vol. 312, pp. 1780–1782, 2006.
- [68] U. Leonhardt, “Optical conformal mapping,” *Science*, vol. 312, pp. 1777–1780, 6 2006.
- [69] D. Schurig, J. J. Mock, B. J. Justice, S. A. Cummer, J. B. Pendry, A. F. Starr, and D. R. Smith, “Metamaterial electromagnetic cloak at microwave frequencies,” *Science*, vol. 314, pp. 977–980, 11 2006.
- [70] J. Li and J. B. Pendry, “Hiding under the carpet: A new strategy for cloaking,” *Physical Review Letters*, vol. 101, 11 2008.
- [71] J. S. Choi and J. C. Howell, “Paraxial ray optics cloaking,” *Optics Express*, vol. 22, p. 29465, 12 2014.
- [72] G. J. Chaplain, G. Macauley, J. Běln, T. Tyc, E. N. Cowie, and J. Courtial, “Ray optics of generalized lenses,” *Journal of the Optical Society of America A*, vol. 33, p. 962, 5 2016.
- [73] T. Tyc, J. Běln, S. Oxburgh, C. D. White, E. N. Cowie, and J. Courtial, “Combinations of generalized lenses that satisfy the edge-imaging condition of transformation optics,” *Journal of the Optical Society of America A*, vol. 37, p. 305, 2 2020.
- [74] J. Běln, J. Courtial, and T. Tyc, “Lens stars and platonic lenses,” *Optics Express*, vol. 29, p. 42055, 12 2021.
- [75] J. Courtial, J. Běln, M. Soboňa, M. Locher, and T. Tyc, “Shifty invisibility cloaks,” *Optics Express*, vol. 32, pp. 11–25, 1 2024.
- [76] U. Leonhardt and T. Tyc, “Broadband invisibility by non-euclidean cloaking,” *Science*, vol. 323, pp. 110–112, 2009.
- [77] B. Zhang, Y. Luo, X. Liu, and G. Barbastathis, “Macroscopic invisibility cloak for visible light,” *Physical Review Letters*, vol. 106, 1 2011.
- [78] X. Chen, Y. Luo, J. Zhang, K. Jiang, J. B. Pendry, and S. Zhang, “Macroscopic invisibility cloaking of visible light,” *Nature Communications*, vol. 2, pp. 1–6, 2 2011.

- [79] H. Chen and B. Zheng, “Broadband polygonal invisibility cloak for visible light,” *Scientific Reports*, vol. 2, pp. 1–4, 2 2012.
- [80] H. Chen, B. Zheng, L. Shen, H. Wang, X. Zhang, N. I. Zheludev, and B. Zhang, “Ray-optics cloaking devices for large objects in incoherent natural light,” *Nature Communications*, vol. 4, pp. 1–6, 10 2013.
- [81] W. Li, J. Guan, Z. Sun, and W. Wang, “Shifting cloaks constructed with homogeneous materials,” *Computational Materials Science*, vol. 50, pp. 607–611, 2010.
- [82] C. Jiang and X. Zang, “Overlapped optics, illusion optics, and an external cloak based on shifting media,” *Journal of the Optical Society of America B*, vol. 28, pp. 1994–2000, 8 2011.
- [83] X. F. Zang, B. Cai, and Y. M. Zhu, “Shifting media for carpet cloaks, antiobject independent illusion optics, and a restoring device,” *Applied Optics*, vol. 52, pp. 1832–1837, 3 2013.
- [84] Y. Zhou, H. Chen, al, Z. J. Wong, Y. Wang, J. C. Miñano, P. Benítez, J. C. González, Y. Du, X. Zang, C. Shi, X. Ji, and Y. Zhu, “Shifting media induced super-resolution imaging,” *Journal of Optics*, vol. 17, p. 025606, 1 2015.
- [85] F. Sun, Y. Liu, and S. He, “True dynamic imaging and image composition by the optical translational projector,” *Journal of Optics*, vol. 18, p. 044012, 4 2016.
- [86] M. Huang, J. Yang, C. Yang, F. Mao, and T. Li, “Target illusion by shifting a distance,” *Optics Express*, vol. 26, pp. 24280–24293, 9 2018.
- [87] H. A. Madni, B. Zheng, M. Akhtar, F. Jaleel, S. Liu, A. Mahboob, S. Iqbal, W. X. Jiang, and T. J. Cui, “A bi-functional illusion device based on transformation optics,” *Journal of Optics*, vol. 21, p. 035104, 2 2019.
- [88] C. Yang, M. Huang, J. Yang, T. Li, T. Pu, and F. Mao, “General design of 3d piecewise homogeneous illusion devices with arbitrary shapes,” *Journal of Applied Physics*, vol. 128, 7 2020.
- [89] S. Li, F. Sun, D. An, and S. He, “Increasing efficiency of a wireless energy transfer system by spatial translational transformation,” *IEEE Transactions on Power Electronics*, vol. 33, pp. 3325–3332, 4 2018.
- [90] Y. Liu, F. Sun, S. He, W. Wang, L. Lin, X. Yang, J. Cui, C. Du, and X. Luo, “Novel thermal lens for remote heating/cooling designed with transformation optics,” *Optics Express*, vol. 24, pp. 5683–5692, 3 2016.

- [91] Y. Liu, K. Chao, F. Sun, S. Chen, H. Dai, and H. Chen, “Active thermal metasurfaces for remote heating/cooling by mimicking negative thermal conductivity,” *Advanced Materials*, vol. 35, p. 2210981, 7 2023.
- [92] A. Greenleaf, Y. Kurylev, M. Lassas, and G. Uhlmann, “Electromagnetic wormholes and virtual magnetic monopoles from metamaterials,” *Physical Review Letters*, vol. 99, 10 2007.
- [93] T. Zentgraf, J. Valentine, N. Tapia, J. Li, and X. Zhang, “An optical “janus” device for integrated photonics,” *Advanced Materials*, vol. 22, pp. 2561–2564, 2010.
- [94] D. R. Smith, W. J. Padilla, D. C. Vier, S. C. Nemat-Nasser, and S. Schultz, “Composite medium with simultaneously negative permeability and permittivity,” *Physical Review Letters*, vol. 84, p. 4184, 5 2000.
- [95] S. Bang, S. So, and J. Rho, “Realization of broadband negative refraction in visible range using vertically stacked hyperbolic metamaterials,” *Scientific Reports*, vol. 9, pp. 1–7, 10 2019.
- [96] Epic Games, “Unreal engine.” <https://www.unrealengine.com>, 2019.
- [97] J. Courtial, “Dr-tim.” <https://github.com/jkcuk/Dr-TIM>, 2019. GitHub repository.