GasSight: A Portable Gas Imaging Camera

Angus Grant MacGruer

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

June 2025

Abstract

Most gases do not absorb visible light, meaning they are invisible to the naked eye. Methane is one such invisible gas. According to NASA[1] methane is the second most impactful greenhouse gas, but its short lifespan in the atmosphere means that reducing atmospheric methane could bring climate benefit on a geologically rapid timescale. Whilst natural sources of methane gas are common, much of the atmospheric methane is from anthropogenic sources, for example via its use in natural gas fuel. Leaks of methane along pipelines often occur, but are challenging to detect. Many optical techniques exist for the imaging and hence localisation of gas source, but each have various strengths and weaknesses in terms of detection sensitivity, detection range, ease of use and frame rate. In this thesis a video-rate gas imaging system is introduced that features a compact size, an infrared frame rate of around 15 fps, and a moderate detection range of up to 16 m. The system makes use of flood illumination at 1653 nm to probe an absorption line of methane, allowing a full scene to be investigated whilst maintaining eye-safety. The device was designed to be compact with a low mass of 3.15 kg, making it suitable as a drone payload for aerial pipeline surveillance. The research described throughout this thesis covers the work developing this prototype system based upon an earlier concept demonstrator, including the improvements made to performance and the successful drone mounted field trials. An improved further iteration of the system was designed to overcome the shortcomings of the prototype system, specifically in terms of sunlight rejection. The main benefit of this iterated system are its increased infrared frame rate of 100 fps, while maintaining comparable detection range and identical form factor. This thesis presents a gas imaging device that satisfies the need of close range detection, with portability and high frame rate. This device could work in tandem

[1] https://climate.nasa.gov/vital-signs/methane/?intent=121

with lower frame rate systems capable of detection ranges exceeding 100 m in order to increase methane monitoring deployment worldwide and detect and reduce emissions.

Contents

List of Figures

List of Tables

Acknowledgements

There are many people I must thank for helping me both in and out of the lab.

I would like to begin by thanking Miles and Graham for their constant guidance and help throughout. Miles for providing me with ideas and goals all the way through my PhD and Graham for helping me realise them and teaching me so many things including, but not limited to: 3D printing, LabView, soldering (though, I was never very good at this), PCB design, paper writing (jury's out), product demonstration, and a disdain for organised fun.

This work would not be possible without Steven and Peter sacrificing their evenings to help me complete field trials, so thank you! I learnt a lot about data collection during these evenings and I think it produced some of the more interesting results presented in this thesis. I also want to thank Steven for motivating me to go out and get the data, help with scientific writing and providing me with lots of background physics explanations, I learnt a lot from working with you. I additionally thank Peter for his seemingly endless time for debugging, answering questions about literally anything, whistling into Shazam and more. That time was surely to the detriment of you but was certainly to the benefit of me.

I'd like to specifically acknowledge (but let's stop short of a full thank you) my three labmates Osian, Wei and Robert, in which this work was completed despite their best efforts. If impromptu music quizzes, spinning tops, basketball, puzzle cubes and Cinematrix could count towards my PhD I'd have graduated a long time ago.

I would also like to thank Sara and Mark. Sara for so effectively onboarding me into the group as she does for every new recruit, and Mark for always making sure that lunches were at least interesting (and because I know they will both read this section. Hi Sara and Mark!).

Thank you to the Optics Group for their support throughout my PhD and for making me feel

welcome!

Thank you to Lewis and Craig, my two oldest friends. Our regular calls discussing our latest scrapes were always the highlight of any week. I greatly anticipate our next culinary misadventure or Elden Hang whenever the shift pattern allows (assuming, of course, it is not a Go Day!). As I write this I know there is an update in the launcher not being downloaded, so get on it.

Thank you to Kieran, my more successful undergraduate counterpart and friend. You encouraged me to take this PhD and I'm not sure I would have done it otherwise, so this work wouldn't exist without you. Our PhD journeys both began by an organisational fluke in the literature review project. I'm not sure why I got single-pixel imaging and you didn't, but I think it worked out for both of us in the end.

A huge thank you to Amy for all of her advice and support throughout. I'd not have met you without doing this PhD, and that outweighs any degree.

Finally, I'd like to thank my family. Thank you to my Uncle Dods for making me laugh and for his liberal use of the sticker function on his phone. Thank you to my Granny for her love and for always checking in and keeping me updated on what's going on at home. Thank you to my Mam and Dad, for their love and support and for always being there to listen to me. Thank you to my cat Ted, you aren't very nice, but you are quite photogenic. I'd lastly like to thank my sister and flatmate, Raine, who has been with me the whole way through my PhD. You completed a whole undergraduate degree and got a first in this time, and I'm very proud of you and what you've achieved. I know you'll go on to do great things.

A final and quick special thanks to Peter, Mark, Graham, Steven, Emma, Chris, Amy and Miles for thesis feedback!

I would additionally like to acknowledge that his project was jointly funded by M Squared Lasers Limited and the UK National Quantum Hub for Imaging (QuantIC: EP/T00097X/1).

Declaration of Authorship

I hereby declare that this thesis is the result of my own work, except where explicit reference is made to the work of others, and has not been presented in any previous application for a degree at this or any other institution.

Angus Grant MacGruer

Publications

1. MacGruer, A.G., Johnson, S.D., Mekhail, S.P. et al. Characterising the performance of a drone-mounted real-time methane imaging system. Sci Rep **15**, 8787 (2025). https://doi.org/10.1038/s41598-025-93186-z

2. Publication in preparation: MacGruer, A.G., Padgett, M.J., and Gibson G.M. A Methane Gas Detection System Using An Optical Switching Method

Conference Proceedings and Presentations

1. A.G. MacGruer, G.M. Gibson, K. Nutt, and M.J. Padgett, "Real-time imaging of methane gas from a UAV mounted system" Proceedings Volume 12735, Remote Sensing Technologies and Applications in Urban Environments VIII; 127350J (2023) `https://doi.org/10.1117/12.2678412` Presented at SPIE Remote Sensing, Amsterdam, Netherlands. Presented 4th September 2023.

2. A.G. MacGruer, S.D. Johnson, K. Nutt, M.J. Padgett, and G.M. Gibson, "Drone-based gas leak detection system for use in industry" Proceedings Volume 12795, Quantum Technology: Driving Commercialisation of an Enabling Science IV; 1279503 (2023) `https://doi.org/10.1117/12.2691068` Presented at SPIE Photonex, Glasgow, UK. Presented 25th October 2023.

Summer Schools and Training

1. International Summer School in Quantum Technologies (August 2023) - A quantum technologies summer school based in Birmingham

2. Quantum Leap (November 2023- February 2024) - A Quantic organised entrepreneurship and innovation training programme.

Abbreviations

CMOS	Complementary Metal Oxide Semiconductor
DAC	Digital to Analogue Convertor
DFB	Distributed Feedback
DIAL	Differential Absorption Lidar
DMD	Digital Micro-mirror Device
FID	Flame Ionisation Device
FOV	Field of View
FPA	Focal Plane Array
FPS	Frames per Second
FSR	Free spectral range
FWHM	Full-Width Half-Maximum
LWIR	Long Wave Infrared
MEMS	Micro-Electromechanical System
MWIR	Mid Wave Infrared
OEM	Original Equipment Manufacturer
OGI	Optical Gas Imaging
PID	Proportional-Integral-Derivative
ppm	Parts per Million
ppm.m	Parts per Million Metre
ppb	Parts per Billion
ROI	Region Of Interest
SDK	Software Development Kit
SWIR	Short Wave Infrared

TDLAS	Tunable Diode Laser Absorption Spectroscopy
TDLidar	Tunable Diode Lidar
TTL	Transistor-Transistor logic
UAV	Unmanned Aerial Vehicle

Figure 1: **GasSight takes flight!** Proof of concept image for drone flight taken during first few months of my PhD. Miles to be digitally removed at a later date.

Chapter 1

Introduction

What if you were able to see something others cannot? This idea is the aim of novel imaging techniques. Being able to see something live is an incredibly powerful tool and something that is often taken for granted. Images and videos can enrich understanding of a situation and environment by providing context that could have been missing when conveyed via another medium. However, there are some things that cannot be seen with the human eye. By developing novel imaging techniques to see the invisible, visual context can be added to previously abstract scenarios. These scenarios range from the ubiquitous such as X-rays, for detecting internal bodily damage and for transmission microscopy, to the specific such as constructing 3D images through optical fibres in order to perform minimally invasive endoscopy [1–3].

One such scenario where imaging the invisible is incredibly useful is in that of gas detection. Many gases do not absorb visible light and as such cannot be seen without the use of a novel imaging technique. Being able to image gases can provide a huge depth of information that stems across environmental studies to industry. By being able to quickly and efficiently identify the presence of a gas can add context to a site of ecological interest, and improve safety by mitigating leaks of hazardous gases, such as explosive gases or potential asphyxiants in industry [4–6]. The work presented in this thesis demonstrates such a device which can image methane gas at video-rates.

1.1 Methane Gas

Methane has many properties that make effective and easy detection a priority. As a greenhouse gas that is 86× more effective at absorbing and re-emitting radiation than carbon dioxide over a 20 year period and 28× greater over a 100 year period, and so it is clear that the impact of methane leakage into the atmosphere on climate change is substantial [7, 8]. The fact that methane's 12 year lifetime in the atmosphere is comparatively short, compared to the centuries of carbon dioxide, underlines the fact that tackling methane emission could have a profound short-term impact in climate change prevention [7,9]. The issue with methane gas infrastructure is not the long-term impact, but the sharp short-term damage that methane can produce if leaks are not properly monitored [9]. The difficulties in monitoring methane leaks stems from the fact that methane is invisible to the human eye and odourless, requiring an indicator chemical to be added to the domestic gas supply [8]. Methane gas is also highly flammable with a lower flammability limit at atmospheric pressure of around 5% and an upper flammability limit of around 15%, though this is pressure dependent [10]. Methane gas leaks additionally cause a reduction of oxygen levels in the surrounding area, posing a direct threat to workers [11].

Whilst methane releases can be natural, such as from seepage or geothermal sources, there is no denying the anthropogenic origin of much of the world's methane [12]. The proof of the extent of anthropogenic emission lies in the fact that methane concentrations trapped in ice from before the industrial revolution are far less than those trapped afterwards [12]. One reason for the large amount of methane leakage is the use of natural gas as a fuel source [13]. The appeal of utilising natural gas, as opposed to other fossil fuels, is that less carbon dioxide is released, however this is not to say that methane release does not require careful monitoring [13]. The use of natural gas has been proposed as a low carbon "bridge" fuel allowing reduced carbon emissions whilst a cleaner source is found and more widely adopted [13]. Methane leaks can effect everyday life, an example being the fact that annually the city of Boston suffers thousands of gas leaks [14]. Leaks into an urban environment can have adverse effects on air quality and can in extreme cases lead to explosions due to the flammability of methane [14]. It was estimated that in 2011 as much as 2.4% of natural gas produced in United States is released

into the environment, this has an equivalent impact as 151 Tg of carbon dioxide [15]. In 2011 it was noted that many of the pipelines utilised are around, or over, 50 years old [9]. Due to the high pressures used, the flow rate of gas leaks can be high [16]. Methane leaks can also be costly with around 5 billion dollars worth of gas lost annually [17]. Landfill sites are a leading source of anthropogenic methane and carbon dioxide [18]. A third of the UK's methane is estimated to originate from landfill, with 2013 predictions estimating 667 kilo tonnes being emitted from these sites [19]. Subterranean pockets of methane gas can be pierced by drills during the search for oil, and fracking is considered to have accelerated natural methane escape [9, 12]. Additionally coal mining is a primary source of methane gas emission, at one point estimated to make up 8% of human-caused methane emission as pockets of trapped methane are ruptured and released during mining [20]. As methane is so widespread in society, the requirement to effectively and safely remedy leaks is clear. The fact that methane is also such a powerful greenhouse gas is what allows it to be detected, due to its ability to absorb infrared light.

Methane gas exhibits strong absorption at 3.3 μm, which is located in the mid wave infrared region (MWIR) of the electromagnetic spectrum, this is known as its fundamental vibrational mode [21, 22]. Whilst the absorption here is strong, and offers great promise for detection capabilities, it is difficult to create a system that could image at this wavelength due to the challenge and cost associated with the relevant detectors and sources [23, 24]. Such difficulties are the reason that many gas detection systems instead make use of the comparability less strong absorption lines at around 1.65 μm where the technology is more developed and affordable [21]. Absorption lines around 1.65 μm are known as the overtone band or the second harmonic [22, 25]. A comparison between the relative absorption strengths at 3.3 μm and 1.65 μm is presented in Figure 1.1, highlighting almost three orders of magnitude difference in absorption strength. Due to the great promise of detecting methane more sensitively in the MWIR, as a result of the greater absorption, a wide array of research is conducted in order to realise this. One shortcoming of detecting methane around 3.3 μm is that it is very close to a water absorption line, and care must be taken to avoid confusion and false detection [22, 26].

The cause of these absorption peaks are known as ro-vibrational energy levels, corresponding

to the degrees of freedom of both vibration and rotation of the bonds in the molecule [27]. In the case of methane, this is specifically the carbon and hydrogen bonds [27]. The particular energy levels causing the absorption peaks in Figure 1.1 (a) are modes known as C-H stretch modes and are cause by vibrational movements of the carbon and hydrogen on the bonds moving closer and further apart from one another, in this case specifically what is known as the ν_3 band, which is the degenerate stretching band [27, 28].The vibrational stretching requires a specific amount of energy which can be provided by specific wavelengths of light corresponding to the peaks seen in Figure 1.1 [29]. The peaks in Figure 1.1 (b) are, as mentioned before, the second harmonic of these absorption peaks [29].

The widths of the peaks themselves can vary and are largely dictated by two phenomena – pressure broadening and Doppler broadening. Broadening can cause peaks to overlap, both for peaks of the same gas and peaks from other gases [30]. Pressure broadening is caused by collisions between molecules, which is more likely to happen at higher pressures where this effect is more pronounced [31,32]. The collisions act to alter the energy levels of the molecule in question meaning that it could accept higher or lower wavelengths of light effectively broadening the range of absorbed wavelengths, with higher pressures causing increased widening [32, 33]. The shape of the widened response becomes Lorentzian in general, and in some cases even the central wavelength can shift [31, 34]. Doppler broadening on the other hand is more dependent upon temperature and the speed at which the molecules randomly move [32]. The direction of travel of the molecule can cause blue- or red-shifting of light, effectively broadening its response to light in a Gaussian response [32, 35]. This effect is far more dominant at higher temperatures [32, 35].

To estimate the width of the peak of interest in this project at 1.653 μm one can replot Figure 1.1, specifically looking in the region of the peak. Fig 1.2 shows the peak of interest plotted now in terms of GHz where for a mole fraction of 1, a temperature of 300 K and a pressure of 1 atm, an approximate order of magnitude full width at half maximum (FWHM) would be roughly 5 GHz.

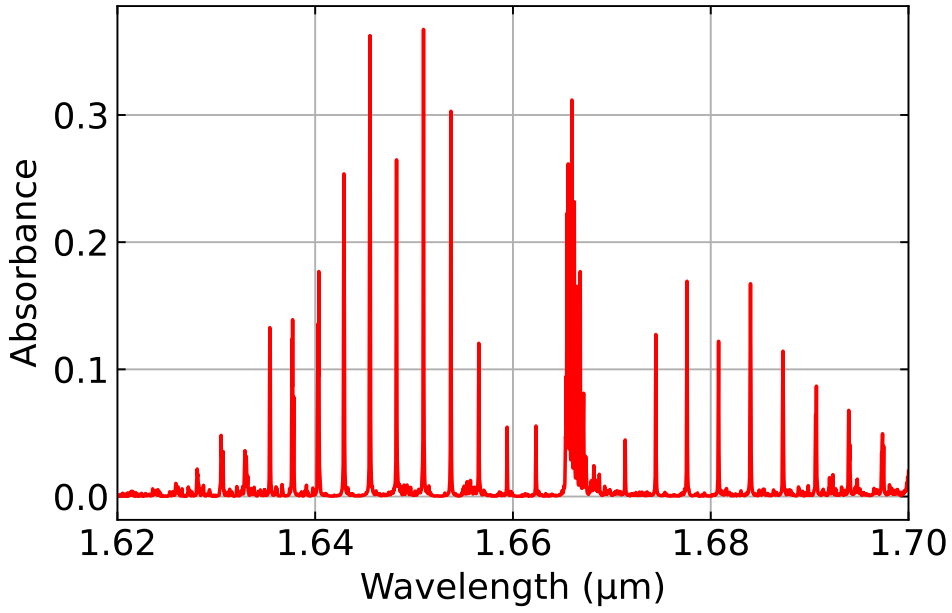

(a)

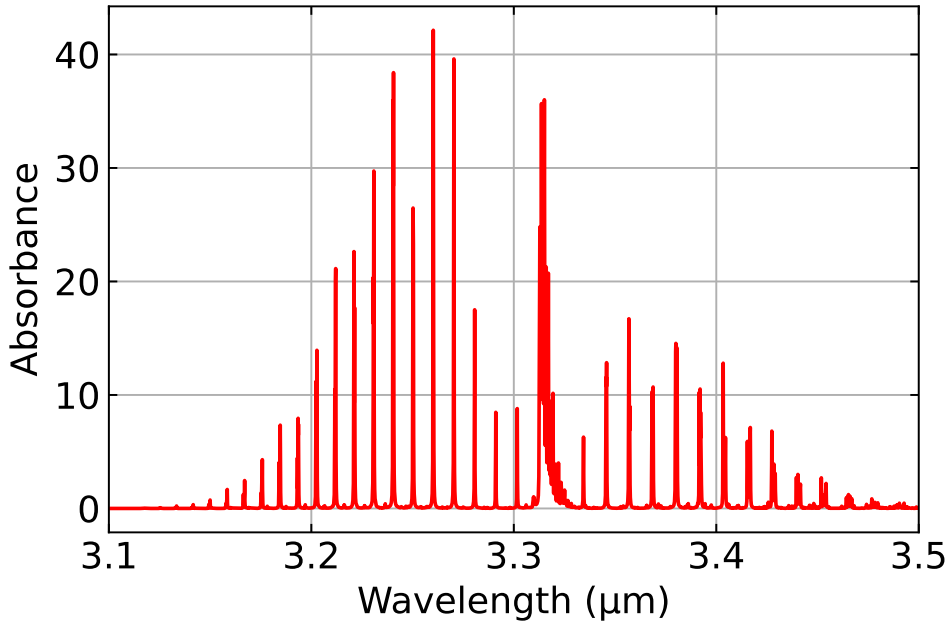

(b)

Figure 1.1: **Plots comparing the relative absorbance of wavelengths in (a) the SWIR region and (b) the MWIR region for** 1 cm **of pure methane gas with a mole fraction of 1 at a temperature of** 300 K **and a pressure of 1 atmosphere.** Plots are generated using data from the HITRAN database [36] obtained via the SpectraPlot toolset available at `https://www.spectraplot.com/absorption` [37].

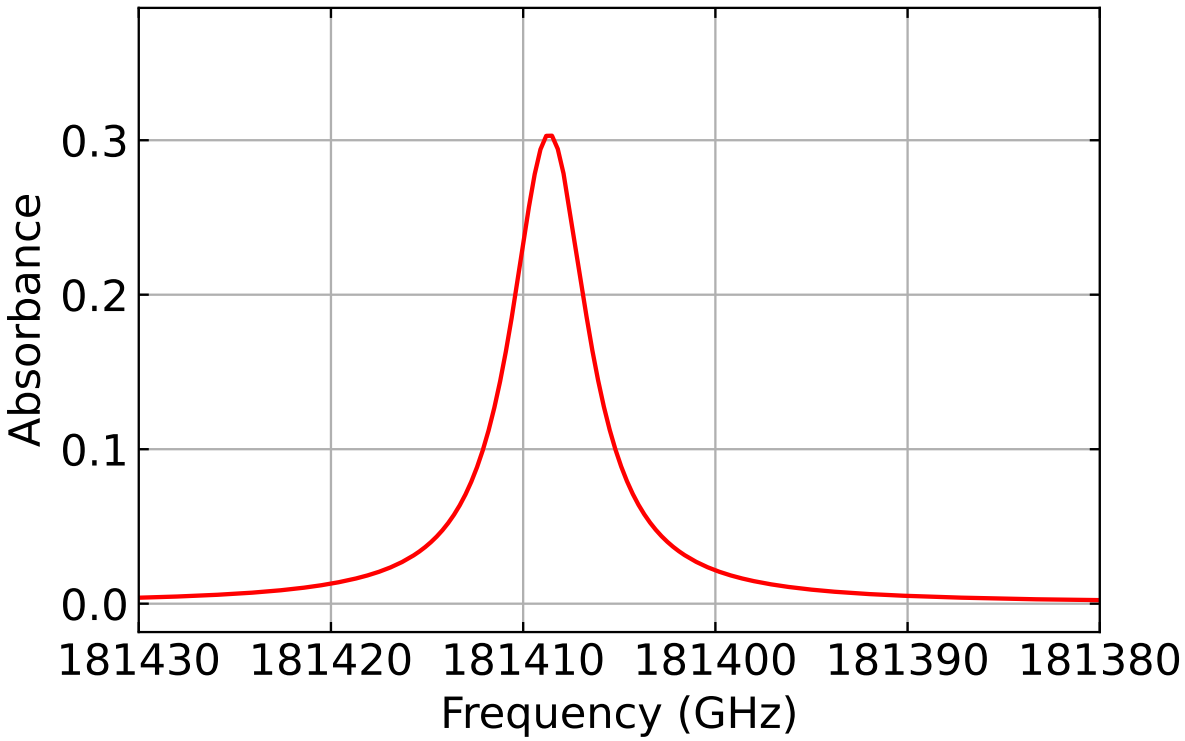

Figure 1.2: **A closer look at the absorption feature of interest in this project, in units of GHz.** This is plot was made for 1 cm of pure methane gas (with a mole fraction of 1) at a temperature of 300 K and a pressure of 1 atmosphere. Plots are generated using data from the HITRAN database [36] obtained via the SpectraPlot toolset available at `https://www.spectraplot.com/absorption` [37].

1.2 Gas Detection Techniques

Unfortunately, methodologies for estimating the level of methane present in the atmosphere either naturally or by human impact are regularly grossly inaccurate [38, 39]. It is noted that some of the Environmental Protection Agency's estimates of the amount of gas leaked in the United States of America do not match with other predictions, seemingly underestimating the extent [38]. In order to combat this inaccuracy, and to quantify the true extent of methane leaks, new approaches have been proposed in order to measure and monitor leaks. Two such approaches for extrapolating the amount of methane released by a facility or leak are the "bottom up" and "top down" approaches [38]. The bottom up approach uses data measured accurately from a leak in order to extrapolate and scale this to a much wider scope [38]. The top down approach takes an airborne ambient measurement and scales backwards via extrapolation [38]. Sometimes these top down and bottom up surveys can disagree, highlighting a level of inherent uncertainty [39].

There exists a wide variety of gas detection methodologies, each utilising different detection techniques. These techniques differ the information provided, leading to some being useful for specific niches. A few key techniques will be briefly summarised in the following.

1.2.1 Non-Optical Methods

The most common and accessible technique for detecting gas leaks is that of sniffers and other detectors which only detect gas at a single point [40]. These devices are often very accurate, but suffer from the fact that the devices must be manually moved throughout a region of suspected leak by an individual, with the output reading interpreted to ascertain leak presence [41]. An example of such a detector would be a flame ionisation device (FID) which works by using electrodes and a hydrogen flame to burn gases entering the system, burning some molecules will produce ions and produce a current which scales linearly with gas concentration [42]. These devices are particularly sensitive to hydrocarbons, and because of their inherently low noise can accurately measure the concentration of hydrocarbons present [42, 43]. These systems are widely used due to their portability [44, 45].

Electrochemical detectors have also been widely used and established [46]. These detectors oxidise the methane molecules between two electrodes, generating a current proportional to the gas concentration [8,46]. The performance of these electrochemical sensors can vary depending on the medium under which the oxidisation happens, with liquid media being more prone to leakage over time and solid media being more temperature dependent [8]. Acoustic sensors are also reported which can detect leaks by listening to the acoustic signals produced as the gas leaks [47]. An example of such a detector reported in the literature is a microphone array capable of recording gas leaks of very low flow rate, albeit only at a range of 20 cm. Some other examples of low cost, low information techniques involve using sniffer dogs or soap bubble based detection [48].

1.2.2 Tunable Diode Laser Absorption Spectroscopy

Tunable Diode Laser Absorption Spectroscopy (TDLAS) makes use of a tunable laser diode, usually that of a distributed feedback (DFB) laser [49]. This laser and its respective detector is used to investigate the absorption profile of a gas sample by detecting transmitted light and comparing it to a reference gas of equivalent path length [49, 50]. Generally the absorption profile of the gas is characterised by passing a saw-tooth wavelength response through the samples, the system also needs to be regularly calibrated against a reference cell [51]. A system that simply uses on and off absorption wavelengths as opposed to sweeping is also demonstrated [52]. Alternatives to using a cell do exist, such as use of a delay line or passing the laser light through the environment [50, 53]. This system utilises a differential measurement to compare transmission through the sample cell to that of the reference [51]. The system is not an imaging device but does provide accurate readings for concentration of desired gas present in the sample, with detection of around 0.1-1 parts per million (ppm) reported [54, 55]. Tactical use of light sources can mean one system can investigate a single source for presence of multiple gases, with one lab based system able to check a sample for water vapour, carbon dioxide and methane simultaneously [56]. These systems have also been reported in both handheld and unmanned aerial vehicle (UAV) mounted versions [54, 57]. The main drawback of this system is that, while being able to infer presence and concentration of gas accurately, it makes pinpointing the source of a

leak very difficult and involved - due to the lack of an imaging component [40].

1.2.3 Photoacoustic Spectroscopy

Photoacoustic spectroscopy is a very similar technique to TDLAS, in that it utilises a narrow beam-width laser and a cell of gas [58]. One uses a high power laser to excite a gas sample cell, and then listening for sounds caused by the pressure changes in the cell as a result of the temperature change that absorption and de-excitation produces [58, 59]. Whilst commonly the lasers used are that of carbon dioxide or carbon monoxide lasers, it has been demonstrated that certain less power demanding diode lasers can be effectively utilised [59]. Systems capable of detecting concentrations as low as 20 parts per billion (ppb) for acetylene and 6 ppb for ammonia are reported [59, 60]. Whilst the appeal of the system is its optical simplicity, compact size and low cost microphone detectors, it has been reported that, in comparison to some TDLAS techniques, some systems can detect as low as two order of magnitudes lower concentrations for certain gases with overtone absorption in the SWIR such as methane and ammonia [50, 59]. The system is also very vulnerable to external vibrations, with a large focus on photoacoustic spectroscopy research revolving around increasing a systems robustness in field settings [61].

1.2.4 Differential Absorption Lidar

Differential Absorption Lidar (DIAL), is a Lidar technique used to infer the presence of gas from great range. DIAL systems utilise differential back-scattered detection of on and off absorption illumination to detect gases, however, in this case a pulsed laser is utilised [62, 63]. The two wavelengths selected for probing a source must behave similarly in the environment [63]. The system is very sensitive and can provide range and concentration measurements with the correct calibration [63]. DIAL is not inherently an imaging system, though a picture of a gas distribution can be built up using scanning techniques [63]. Whilst many factors impact the systems performance, wind speed can have the greatest impact, in order to quantify this, a device to measure wind speed is often used in conjunction to the detector [63]. These systems are ideal for scanning full facilities from long range, with both aeroplane and space based measurements

reported in the literature [64, 65]. While this system can build up pictures of a facility's gas emission from range, it is not accessible and usable at shorter ranges, making leak source detection difficult. Another key difficulty is the requirement for an experienced operator and the large size of the system, with an example being a gas imaging system installed in a lorry [40, 63].

1.2.5 Tunable Diode Lidar

Tunable Diode Lidar (TDLidar) is a relatively new technique that utilises similar methods to TDLAS and DIAL, as well as time correlated single photon counting in order to characterise a scene [40]. The aim behind this system is to combine the accurate concentration measurements of TDLAS with the ranging and scanning capabilities of DIAL to build up an image of a scene and locate sources of gas leaks. The system makes use of the sweeping approach to probe an absorption profile via TDLAS to clearly identify the concentration of gas present. The single photon correlation counting allows depth to be interrogated based on degree of back scatter, which allows a 3D image of the scene to be built up [40]. The movement of the laser probe is achieved using a Risley prism pair, which is a technique that allows one to scan a scene by angling infrared light [66]. Currently the system is positioned as a stationary monitoring system, but advancements will hopefully allow for more handheld deployment [40]. The system also struggles at shorter ranges, and the time it takes to build up an image is on the order of a minute [40].

1.2.6 Optical Gas Imaging

Optical gas imaging (OGI) is an infrared technique used for directly imaging gas which does not use active illumination, but instead monitors infrared thermal images to infer the presence of gas from temperature changes, commonly in the MWIR and long wave infrared (LWIR) region [67, 68]. Because the thermal cameras at these wavelengths generally respond to a wide spectrum of wavelengths, the camera is commonly filtered around the absorption wavelength of the desired gas [68]. The thermal image will show gas plumes as black clouds if the camera is filtered to the absorption wavelength of the gas, or white clouds if filtered to emission wavelength

of a gas [69]. The signal must be above the noise floor in order to register a detection, however this is strongly dependant upon temperature [70,71]. In order to control the temperature, systems require extensive cooling, also, a difference in temperature between background and gas can begin to mask the signal, additionally some of the filters used in such devices also require cooling [71, 72]. By swapping the filter on the system one can utilise the same device to visualise a wide variety of gases with relative ease [68].

Wind speed and sunlight levels impact the accuracy and usability of these devices as they can both effectively mask or reduce the signal [72]. As a result of the additional cooling required, either to the detector or the filter, the overall cost of these devices can be high [40]. Despite the versatility of the system to image a wide array of different gas species, one drawback is that in some cases the filters exhibit overlap between different gas absorption or emission lines [68]. The lack of filter specificity can mean that identifying the difference between two gas species can be very difficult, this has a knock on effect of making it challenging for the system to be modified to quantify the amount of gas present in a scene [73]. The primary barriers to entry for this technology are the dependence on the external environment, cost of cooling and usability, with an element of the users own judgement being required for operation and assessment of gas leak severity [40, 69, 72].

1.2.7 Non-linear Interferometry

One of the more novel techniques is that of non-linear interferometry, where the gas can be probed in the MWIR but detected at a more accessible wavelength [26, 74–76]. An initial example of this utilised a pump laser of 1.064 μm and an idler laser of 3.22 μm in order to generate a detectable gas signal at 1.59 μm by using stimulated parametric down conversion [76]. Further iteration of this process led to a system in which the pump is now a 671 nm laser, allowing for detection at a wavelength of 848 nm. Light with a wavelength of 848 nm can be detected on more conventional silicon infrared cameras with lower noise signals and greater quantum efficiency [26]. While these systems are not yet capable of stand-off imaging, the ability to detect concentrations as low as 140 ppb has been demonstrated [26].

1.2.8 Raman Spectroscopy

Raman spectroscopy is another gas detection technique that can be used to quantify gas leaks. This process works by illuminating a gas sample or source with laser light and collecting the scattered light [77]. The scattering of the light in the gas will have a known effect on the returning wavelength [77]. This effect is molecule specific meaning the presence and amount of the gas can be inferred [77]. Raman Spectroscopy is beneficial because it does not require the use of a back scatter target behind the gas source in order to detect it, unlike many of the other techniques mentioned previously [78]. Raman spectroscopy has been demonstrated to be able to detect gases that are challenging to detect via other means, such as hydrogen [78, 79]. These hydrogen detections have been reported at ranges of up to 100 m [79]. The key issue with gas detection via raman spectroscopy is the relatively low signal to noise (SNR), especially in conditions of sunlight, meaning the optical system can become quite optically complex when attempting to filter out unwanted light [77–79]

Each of these techniques are summarised and presented in terms of pros, cons and imaging component in Table 1.1.

1.3 Other Gases

Whilst methane is the primary gas of interest for this research, there are other gases in which detection would be of interest to different groups. One such gas is carbon monoxide, which similarly to methane, is a difficult gas to detect due to its lack of colour and odour [80]. Detecting carbon monoxide is important because of its widespread nature and its potentially life-threatening impact upon health [80]. Optical detection is possible for carbon monoxide with techniques such as quartz enhanced photo-acoustic spectroscopy [81]. The wavelength of laser utilised for this was 4.55 μm which may not be accessible with a laser diode, however other, more accessible, wavelengths are possible. Another important gas to be able to detect is carbon dioxide, one of the most prevalent and damaging greenhouse gases [7]. Carbon dioxide can be detected via similar techniques making use of laser diodes at a wavelength of 4.2–4.3 μm but other, less strongly absorbed, wavelengths can be utilised [82]. Hydrogen, which shows promise

Detection Technique	**Imaging?**	**Pros**	**Cons**
Flame Ionisation	No	Low noise, low cost, portable, high ease of use, accuracy of concentration measurement	Must be manually swept across a region, only works at a point
TDLAS	No	Can detect low concentrations of gas, can be modified to look at different gases using one device, can be portable	Can measure concentration but is difficult to find the source of the leak
Photoacoustic Spectroscopy	No	Can detect very low gas concentrations, low cost detectors, small size	Vulnerable to external noise, can often require powerful laser sources, difficult to pinpoint leak source
DIAL	Can be via scanning	Long range detection, can provide range data as well as concentration, can observe a whole facility at once	Large size, difficulty of use, cannot be used at short range due to saturation
TDLidar	Yes	Can provide depth information as well as concentration, small size, provides imaging component	Takes around a minute to build up an image, struggles with saturation at short range, currently only stationary
OGI	Yes	Passive imaging system that detects gas via thermal profile, high ease of use, can detect gas leak sources in real time	Requires expensive filtering and is dependent upon external factor such as wind speed and temperature
Non-linear Interferometry	No	Can detect the MWIR absorption using a cheaper SWIR or even visible camera, can benefit from greater quantum efficiency and lower noise	Not yet proven for stand off detection, optically complex
Raman Spectroscopy	No	Can detect typically challenging to detect gases such as hydrogen. Does not require a backscatter target and is not gas specific. Demonstrated to work at ranges of 100 m	Often dealing with low signal to noise which needs to increased.

Table 1.1: **Table summarising the pros and cons of the discussed detection mechanisms.** The presence of an imaging component is also clearly highlighted.

as a coolant in power generators amongst other applications, is another gas where there seems to be much interest in video-rate imaging [83]. There are difficulties in reaching a wavelength that hydrogen gas can absorb, especially in a handheld device, so instead one can utilise other gas imaging via use of a tracer [83]. Raman detection of Hydrogen is reported with ultraviolet light at around 350 nm, which can be a challenging wavelength to work with due to safety concerns [79]. A carbon dioxide tracer can be introduced within the system and be detected from potential hydrogen leak sources by a gas imaging device [83]. Hydrogen gas can be potentially volatile so detection from a safe distance is of great importance [83].

Acetylene gas is another example of a gas where the demand for leak monitoring is high. Acetylene is a gas that is commonly used in welding, and is mixed with oxygen in order to achieve this [84]. The issue with using acetylene gas in industry is that it is highly explosive and, similarly to methane, it is invisible at visible wavelengths making detection difficult [85,86]. Acetylene does, however, absorb light at 1.53 μm, meaning that techniques such as photoacoustic spectroscopy and TDLAS can be used to detect its presence in an environment [87, 88]. As the wavelength of 1.53 μm has a close proximity to the 1.65 μm line of methane, it offers the potential to modify current systems in order to detect acetylene. Additionally to the absorption line at 1.53 μm, acetylene also more strongly absorbs around 13.7 μm, which can be challenging to access with current active illumination technologies [89].

Sulfur Hexafluoride (SF6) is another example of a greenhouse gas that requires careful monitoring. In terms of environmental impact SF6 has a greenhouse effect 24,300$\times$ greater than carbon dioxide and will persist in the atmosphere for 3,200 years [90, 91]. SF6 is most commonly used in the electronics industry where its properties as an inert gas with no associated risk of explosion or toxicity make it well suited for use as electrical insulation [92]. The caveat to the wide usage of SF6 in industry, however, is that under certain circumstances SF6 can produce toxic and dangerous substances, which can pose a serious health risk to those exposed [93]. SF6 features a strong absorption line located at around 10.5 μm, which is firmly in the LWIR band, making it difficult to image effectively using active illumination techniques [94,95]. An example imaging system uses a wide-band infrared camera and a narrow filter to passively image the gas at the 10.5 μm band [95]. Due to its potential to have a significant climate impact and potential

Gas	Detection Wavelength
Methane	3.3 μm and 1.65 μm
Carbon Monoxide	4.55 μm
Carbon Dioxide	4.2 μm - 4.3 μm
Hydrogen	0.35 μm
Acetylene	1.53 μm and 13.7 μm
Sulfur Hexafluoride	10.5 μm

Table 1.2: **Table summarising different gas species and their common detection wavelengths.**

health implications, widespread leak prevention is a priority [92].

A summary of the detection wavelengths of the discussed gases is presented in Table 1.2.

Overall, it is clear that having reviewed the topic that there is scope for an easy to use accurate gas imaging system that can quickly and reliably build up an image of a gas leak. The pressures on the environment posed by methane make leak remediation a key concern, but current technologies can be difficult to use, provide non-visual data that then must be processed by an expert, or be prohibitorily expensive. By designing a portable, lower cost gas imaging system it is hoped that gas leak detection could be simplified or even automated. If the detection mechanism is kept general, there is a additionally scope to modify the system for use in the detection of other gases.

1.4 Walkthrough

The structure of the thesis will be summarised in the following, with brief outlines of the chapters to follow.

Chapter 2 - This chapter will include an extensive literature review of the gas sensing research conducted by the optics group in Glasgow. The system that makes up the basis for this thesis, known as GasSight, will be introduced as it existed at the beginning of the PhD, this will primarily be based upon the published paper by Nutt et al. [41]. Aspects such as the diode tuning and SNR will be investigated in depth. The chapter will also summarise the years of research and different iterations of the technology that have existed over the years within the optics group. This summary aims to place the work conducted into context and allow a basis of

comparison.

Chapter 3 - This chapter will introduce the issue of movement noise, and the limitations of using the system as a moving device. Much of the processing pipeline and each individual technique implemented in order to make the system more robust will be introduced and explained. The key techniques discussed are timeline changes, correlation, persistency, erosion and dilation, and temporal smoothing. The contents of this chapter are a much expanded version of information included in the SPIE conference proceeding "Real-time imaging of methane gas from a UAV mounted system" [96].

Chapter 4 - This chapter will introduce the topic sunlight sensitivity within the GasSight. The impact of sunlight in imaging systems will be introduced before a focus on work conducted into filters and why it is not so straightforward to use narrow filters. A possible workaround for the limitations of using narrow filters is investigated. The contents of this chapter is again touched upon in the SPIE conference proceeding "Real-time imaging of methane gas from a UAV mounted system" [96].

Chapter 5 - This chapter includes the data that made up the paper "Characterising the performance of a drone-mounted real-time methane imaging system" [97]. This chapter is a culmination of the work introduced in the previous two chapters and looks at the real world drone trials conducted. Also tested was the systems robustness to different surface textures, with a focus on common types. Parts of this work is additionally mentioned in the SPIE papers "Drone-based gas leak detection system for use in industry" [98] and "Real-time imaging of methane gas from a UAV mounted system" [96].

Chapter 6 - Having introduced much of the software solutions to the problems encountered with the system, this chapter aims to introduce all the hardware changes made to the system. Whilst initially discussing the relatively few hardware changes made to the system, this chapter primarily acts as an introduction to the next version of the system and will feature an in depth breakdown into how it operates and the changes made to the previous system to achieve this, as well as the relevant justifications for the choices made.

Chapter 7 - This chapter will consist of a full comparison of the two gas imaging systems in order to establish if the decisions made when designing the system were successful in overcom-

ing the issues presented in the previous chapters. The key aspects of comparison will be that of frame rate, range, sunlight resistance, and polarisation resilience. By comparing these criteria it should be possible to conclude whether the construction of the new system was a success.

Chapter 8 - This brief chapter highlights the results of a collaboration conducted towards the end of the research period. This chapter discusses the key areas of success for the GasSight and where it was identified to succeed and areas that require improvements and additionally includes information about the future of the system..

Chapter 9 - This chapter acts as a final summary of all the work conducted.

Chapter 2

Device Introduction and Evolution

2.1 Introduction

The system that this thesis focuses upon, known as GasSight, is the latest research output from almost 2 decades of video-rate methane imaging research conducted by the Optics Group at the University of Glasgow. Before discussing the research completed with the system and into building the new system, these devices should be placed in context by the work completed prior and the expertise and knowledge that has paved the way for this device to be realised in the form that it has. Firstly, for clarity and ease of understanding the GasSight system and how it detects and indicates the presence of gas signals will be explained, this section will specifically discuss the GasSight as it functioned at the start of this PhD, to place the edits and modifications into context. The initial demonstration unit of this system was developed by K. Nutt, G. Gibson and M. Padgett. Following this a brief summary of the previous research leading up to and including the 2020 paper by Nutt et al. [41] will be presented.

2.2 GasSight

The GasSight device is a video-rate methane imaging camera that displays a live image feed of a methane gas leak atop a feed obtained by a visible camera. The camera consists of 4 primary components: a visible camera (Logitech c270), a short wave infrared (SWIR) camera (Allied

Vision Goldeye G-008) and two SWIR diodes corresponding to wavelengths in the region of 1653nm (Eblana Photonics). These diodes are tuned to correspond to a methane absorption band at 1653.7nm. The system is controlled centrally by an embedded Lattepanda mini-PC, which runs a user interface developed in LabView (National Instruments).

2.2.1 Diode Tuning

The GasSight device makes extensive use of diode tuning in order to access the absorption features of the methane gas and subsequently detect it. Diode tuning is a complex field related to laser spectroscopy, and takes advantage of inherent properties of all lasers in order to control the output power or wavelength. The type of laser will largely dictate how one can tune, the flexibility of the parameters or the possible wavelengths that are accessible. The diodes used here are known as distributed feedback (DFB) lasers. DFB lasers have narrow line-widths, meaning they can be used to generate very specific wavelengths of light with very little spread in wavelength, for example the diodes used in the GasSight have a line-width of 2 MHz [99]. This specificity of wavelength and narrow line-width is achieved via the inclusion of specified structure within the diode cavity itself, which can cause Bragg scattering [100]. If one simplifies the diodes laser cavity into two mirrors separated by a distance of, L and that the cavity is made of a medium with refractive index, n, then the laser light at wavelength, λ_0 must satisfy the following relationship in order to form a standing wave within the cavity:

$$\frac{2Ln}{\lambda} = m \tag{2.1}$$

For m an integer. It can also be seen from this equation why one is able to tune a diode using temperature as changes in temperature can cause changes in the cavity length, it can also have an effect upon the refractive index, which results in different wavelengths achieving resonance [101]. The diode can also be tuned by changing the current, the current has an impact primarily upon the charge carrier density which will primarily effect the refractive index of media [101]. It thus follows that another wavelength that would satisfy this condition would occur at $m-1$,

following:

$$\frac{2Ln}{\lambda + \Delta\lambda_{FSR}} = m - 1 \tag{2.2}$$

Where $\Delta\lambda_{FSR}$ is the required wavelength change to satisfy the relationship known as the free spectral range (FSR). The FSR is the difference between accepted wavelength modes within the cavity and can determine how far one can tune a diode, among other things. In reality a laser will also feature a gain curve, which is a property of the diode structure that can dictate which of the wavelength modes within the cavity are dominant [102, 103]. Which wavelength is promoted by the gain curve can change based on external factors such as temperature or pressure which effect the bandgap of the diode media [102, 103]. When a diode wavelength tunes outside of the gain curve another wavelength mode can be favoured rather abruptly, for example the next wavelength mode only one FSR away, this is known as mode hopping [102, 103]. Mode hopping can prevent some wavelengths of light being accessible and can prevent the diodes from tuning too far [101]. DFB lasers are noted for allowing a moderate tuning range of a few nanometres without mode hopping [101]. Following on from Equation 2.2, it can be rewritten such that:

$$\frac{2Ln}{\lambda + \Delta\lambda_{FSR}} = \frac{2Ln}{\lambda} - 1 \tag{2.3}$$

It thus follows that $\Delta\lambda_{FSR}$ can be expressed as:

$$\Delta\lambda_{FSR} = \frac{\lambda^2}{2Ln - \lambda} \tag{2.4}$$

An estimate of the FSR of the diodes used in the GasSight can be estimated, though it must be noted that the manufacturer data was not forth giving for specific data on cavity length or cavity medium, as such typical values were found in literature in order to provide an estimate. It is suggested that for diodes emitting at around 1650 nm have typical cavity lengths that range from 100 μm to 1000 μm, and are typically made from InGaAsP ($n = 3.2$) or InP ($n = 3.5$) [104]. Using these values and a wavelength of 1653.7 nm in Equation 2.4, one can obtain values in a range from 0.39 nm to 4.28 nm, which can be converted to frequency via:

$$\Delta f = \frac{c}{\lambda^2}\Delta\lambda \tag{2.5}$$

For Δf the width in terms of frequency, c the speed of light, λ is the central wavelength, here taken to be 1653.7 nm, and $\Delta\lambda$ is the change in wavelength, or in this instance the FSR. This gives a frequency range of $\approx$ 42.8 GHz to 470 GHz, which is comfortably greater than the bandwidth of around 5 GHz suggested by Figure 1.2.

All this is to say that diode tuning is very complex,and depends on a wide range of factors, as such it can be very difficult to accurately estimate the wavelength of a laser output without manufacturers data, an expensive spectrometer or, as in this case, a gas cell. In the following the implemented tuning mechanism for the diodes will be discussed.

The diodes are tuned to the absorption wavelength of methane by controlling their current and temperature, these are performed using two original equipment manufacturer (OEM) control chips (ThorLabs MLD203CLN and ThorLabs MTD415T, respectively). By accessing the Proportional-Integral-Derivative (PID) control settings for these diodes on the temperature control chip one can set the temperature, and speed of the temperature response. With these parameters controlled, an Arduino and a Digital to Analogue Convertor (DAC) can be used to supply a control voltage to the current modulation pin of the control chip. Both temperature and current can be controlled to set the wavelength and laser power. By varying this control signal, the diodes can be tuned to the methane absorption line, this is conducted using a simple set-up of a gas cell and a SWIR photo-detector in order to confirm the diodes were correctly tuned. With the diodes correctly tuned, it is also important to know a nearby voltage that corresponds to a diode wavelength that is not absorbed by the gas. By selecting a nearby voltage the overall settling time, defined by the PID settings, can be kept to a minimum. A change in current is used in order to more quickly change the diode wavelength, between on and off absorption. The change in current is set to the minimum value that switches the wavelength in order to minimise the difference in laser output power between the two tuned wavelengths.

A simple description of the process of PID tuning a system is as follows. The operation of the tuning occurs upon the error, $E(t)$, which is defined as the difference between the reference

value set by an operator and the actual output of the system, this has a time dependence because over time the size of the error will change as a result of the PID control and noise. Essentially it acts as a way of measuring how close the output is to the desired value. Proportional, or P, tuning takes the form of the following [105],

$$o(t) = k_P E(t), \tag{2.6}$$

where $o(t)$, is the output of the tuning and k_p is the proportional term. Integral, or I, tuning takes the form of,

$$o(t) = k_I \int^t E(\tau) d\tau, \tag{2.7}$$

where k_I is the Integral term controlling how strongly this tuning impacts the output. Finally, the Derivative, or D, tuning takes the form of,

$$o(t) = k_D \frac{dE}{dt}, \tag{2.8}$$

where k_D is the Derivative term which in turn controls the strength of this type of tuning. In a PID system all three types of tuning will act together as follows,

$$o(t) = k_P E(t) + k_I \int^t E(\tau) d\tau + k_D \frac{dE}{dt}. \tag{2.9}$$

Changing the strength of each of these tuning parameters will change elements of the response of the signal over time. A careful mix of tuning strengths must be employed to ensure that the tuning won't undershoot or overshoot and to ensure that in the case of over or undershooting that the system settles quickly. The ideal scenario here would be a response that, when undergoing change would be able to correct and settle quickly at the target value. A quick response to change in this scenario allows for the diodes to tune on and off resonance quicker. As this is temperature tuning, one direction will always settle fast than the other as it is generally quicker to heat an element than it is to cool it, but this is dependent largely on the heating system in place.

The issue with the approach of merely selecting a nearby wavelength is that it is quite difficult to accurately know what the offset wavelength corresponds to, however an estimate can be made. The information that is known is the supplied voltage and temperature settings to achieve both the on- and off-absorption wavelengths for both respective diodes. Also supplied by the manufacturer are approximate tuning parameters relating differences in temperature and current to differences in wavelength, assuming a smooth tuning with no mode hopping as described above. In the case of these diodes it can be stated that mode hopping is likely not occurring for the setting used, as a clear sweeping of the peak can be achieved and observed by varying the voltage and hence current and monitoring the amount of absorption through the cell. This conclusion is also supported by the calculations of the FSR suggesting mode hopping is unlikely. As temperature remains constant, with the diodes set to 21.9 °C and 26.4 °C, the change in current will dictate how far the diodes are tuned, but to know this the voltage change must be converted to a current change. By measuring the output current from the driver as the input voltage changes in the voltage vicinity of the settings used, the relationship between input voltage and output current is approximately linear appropriating the relationship of voltage across a resistor valued at 12.7 Ω. Using the simple relationship:

$$I = \frac{V}{R} \tag{2.10}$$

For I the output drive current, R the approximated resistance, and V the input voltage one can now convert between input voltage and output current. By then finding the difference between the drive currents for the respective on and off resonance wavelengths the change in current can be found and related to the change in wavelength via the manufacturers current tuning parameter. In the case of the Eblana diodes used here, the worst case scenario is reported as $15\,\mathrm{pm\,mA^{-1}}$ (as reported on `https://eblanaphotonics.com/products/1654-dm/`). So for the two diodes used in the GasSight system, the corresponding current differences are 5.5 mA and 3.9 mA respectively. Using the tuning parameter these correspond to wavelength differences of 82.5 pm and 58.5 pm respectively. In order to convert these wavelength changes to frequency changes Equation 2.5 can be used. This provides respective frequency tunings of 9.09 GHz

and 6.45 GHz. The size of these offsets is consistent with the width of the absorption peak, as previously presented in Figure 1.2. The differences in the size of the tuning widths for the two diodes could be due to a variety of factors including manufacturing differences resulting in different tuning parameters, or human error in the tuning process. It again must be iterated that due to the simplifications made in calculating these offsets that they are merely approximate.

2.2.2 Imaging

These tuned diodes are used as a form of active illumination. The diodes pass through 30 mm focal length lenses in order to collimate the output light before passing through a 20°square diffuser (Thorlabs ED1-S20-MD). The diffusers are used in order to illuminate a much higher area, allowing more of the scene to be investigated for methane as well as providing a flat field illumination. A scene can be illuminated by the diodes and the back scattered light collected through the aperture. The aperture features a cold mirror which separates the light into its SWIR and visible components, passing the SWIR light and reflecting the visible light. The back reflected SWIR light passes the mirror towards the SWIR camera, which is equipped with a 1653 nm filter with a full width at half maximum (FWHM) of 18.8 nm and a 25 mm focal length lens. The camera used in the GasSight system is an AlliedVision Goldeye G-008 camera, which has a resolution of 320×256 pixels. The lens used is a Kowa LM25HC-SW lens which as well as having a focal length of 25 mm, has an aperture radius of 12.7 mm. The field of view (FOV) of the system is approximately 20°, in order to match the spread of the diode diffusers. A 14-bit image is captured with the SWIR camera and saved to a memory location, A, using LabView, the exposure time of the camera is 33 ms. Next the diode state is changed to be the opposite to the previous exposure, if the previous frame was tuned to be absorbed by methane then the subsequent frame will be tuned to be off absorption and vice versa. A second image is taken and saved to a separate memory location, B. With two images saved, one featuring absorption and the other not, the presence of gas can be inferred from the difference between the two. In the case of this imaging regime the direction of subtraction must be the same direction, irrespective of whether A or B features absorption or not. The frame with absorption is always subtracted from the frame without absorption. This directionality is selected such that the regions in which

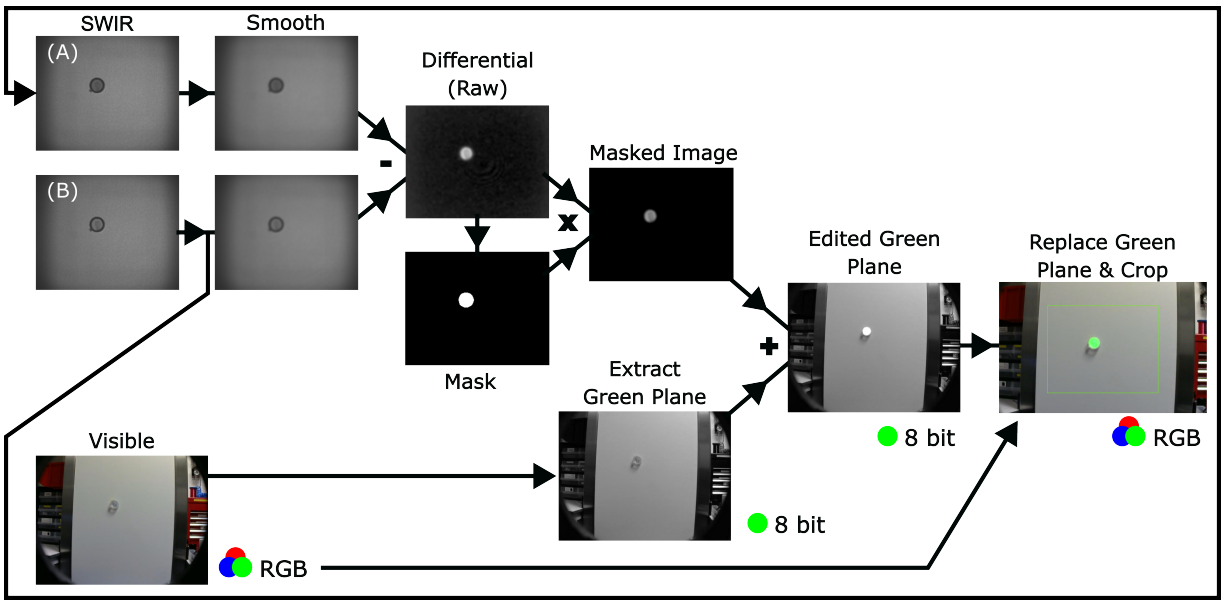

Figure 2.1: **A flow diagram of the image processing in the initial version of the system.** Here a SWIR image is taken, smoothed and subtracted from the previous iterations smoothed SWIR image, the direction of subtraction between A and B will change with each iteration. This differential image is masked, with filters applied to the mask before adding this signal to the green plane of the visible image that was taken at the beginning of the loop. This combination of visible and SWIR images reveals the gas location in the scene.

gas is present appear as a positive signal in the differential image. The differential image is then processed further. Whilst the SWIR images are being captured, the visible camera also captures an image. The final step of the imaging process is to move the image in location B to location A and repeat the process, making sure to swap the direction of subtraction. A simplified flow diagram of this process presented in Figure 2.1, highlighting the path of the SWIR and visible images in the process.

2.2.3 SNR Modelling

In order to give some idea of the performance of the system, a rough estimation of the amount of light reaching the system can be estimated relative to the camera noise of the system. In terms of the noise sources relevant to this observation there are three primary sources; shot noise, dark current and readout noise. First is the shot noise σ_s, this is defined as the variance in the number of electrons read by the signal, N, which if one assumes Poissonian statistics hold for

the incoming photons on the detector, can be defined as [106]:

$$\sigma_s = \sqrt{N} \tag{2.11}$$

The second source is the dark noise, σ_d which will vary from detector to detector and is a result of variations in the amount of electrons readout by the detector as a result of temperature [106]. Electrons can be released as a result of thermal excitations and so can be captured by the detector even in the absence of incoming photons [106]. This variance can be quantified by using manufacture estimated data for the dark current, DC, which has units of electrons per pixel per second, meaning that the dark noise can be calculated via [106]:

$$\sigma_d = \sqrt{DCt} \tag{2.12}$$

Where t is the exposure time of the camera. The final noise source considered is the readout noise, σ_r these are photons that are added to the signal electrons during the readout process, and are once again inherent to the detector [106]. To obtain the total noise, σ_{tot} each of the three noise sources can be combined as follows:

$$\sigma_{tot} = \sqrt{\sigma_r^2 + \sigma_d^2 + \sigma_s^2} \tag{2.13}$$

To begin the estimation, one must first calculate how much desired light is hitting the detector pixels. This will begin with the power leaving the two diodes, P, which combined at range will be equivalent to approximately 20 mW, as each diode outputs approximately 10 mW of power. Next is the fact that the light leaving the diodes is being diffused using a 20°square diffuser, which approximately matches the FOV of the camera. This means that to estimate the power per unit area at a range D one would use:

$$\rho = \frac{P}{(2D\tan(10^\circ))^2} \tag{2.14}$$

For ρ, the power density in units of $\mathrm{W\,m^{-2}}$. Important to remember is the reflectivity of

a surface, S, which will dictate what fraction of the light incident upon the surface will return to the detector. Next to consider is the geometric loss, G, of the imaging system, which is a measure of how much reflected from the surface will be incident upon the lens. This loss is proportional to the area of the lens, which is in turn proportional to the lens aperture radius, r, squared. The geometric loss is also proportional to the area of the illuminated region, which is itself proportional to R^2. This gives an approximate geometric loss of:

$$G = \frac{r^2}{R^2} \tag{2.15}$$

There is also a magnification gain, M, to consider in the imaging system, which is a result of the focal length of the lens in the system magnifying what is imaged, this is the ratio of the distance to the image R and the focal length f:

$$M = \frac{R}{f} \tag{2.16}$$

By combining these values one can then quantify the power density being imaged onto the sensor. To quantify the incident power per pixel, P_p, one must quantify the area of the pixel, A, this can be obtained from the data sheet and in the case of the Goldeye used here is equivalent to $(30 \times 10^{-6})^2 = 9 \times 10^{-10}\text{m}^2$. This means that the power per pixel can be estimated using:

$$P_p = \rho SAGM \tag{2.17}$$

What is more of interest however is the number of incident electrons in the detector, as that is how the noise is generally defined. To convert from incident power to number of electrons one must first know the quantum efficiency of the detector, Q, which dictates what percentage of incoming photons will produce an electron. The full response of the Goldeye camera is presented in Figure 2.2, but for the wavelength used here at $1.65\,\mu$m this is around 0.65. Next one must know the energy of each photon, E, so that a number of photons can be estimated, this can be found according to:

$$E = \frac{hc}{\lambda} \tag{2.18}$$

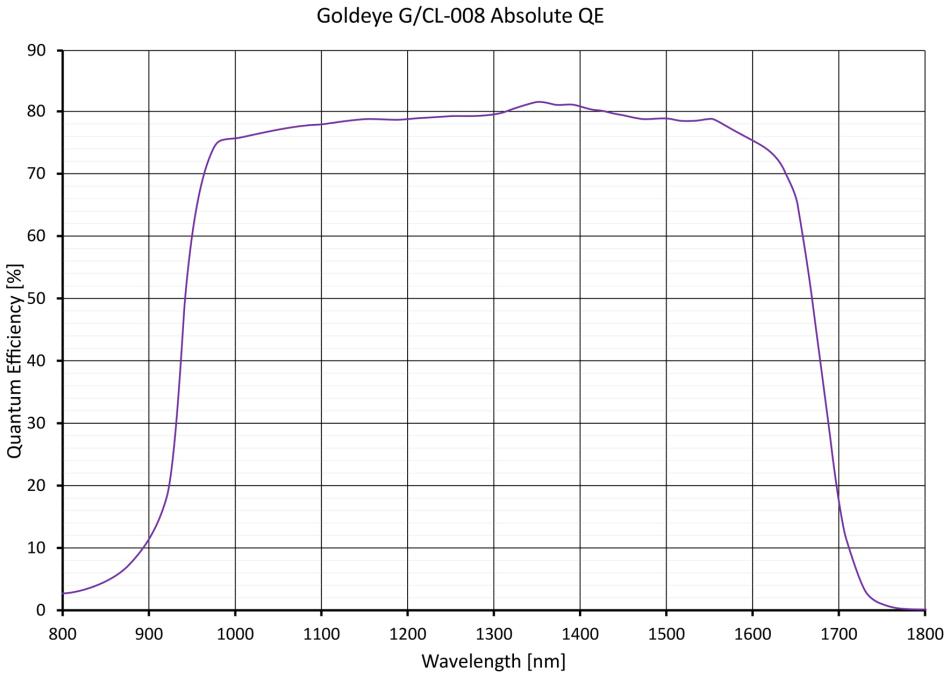

Figure 2.2: **The quantum efficiency of the Goldeye G-008 camera used in the system.** This figure was created by Allied Vision Technologies GmbH and is reproduced from their website with permission. The original figure can be found at: `https://www.alliedvision.com/en/camera-selector/detail/goldeye/g-008-swir-tec1/`.

Where h is Planck's constant, c is the speed of light and λ is the wavelength of the photon. The total energy incident at the pixel will be the product of the power per pixel and the exposure time of the camera, t, which was 33 ms here initially. Therefore the number of electrons, N, and subsequently the signal can be defined as:

$$N = \frac{P_P Q t}{E} \tag{2.19}$$

In order to define the noise, data can be obtained from the Goldeye data sheet, this provides a dark current value, D_C of 880×10^3 electrons per second. The readout noise, σ_r, of the system is difficult to quantify from the data provided, but appears to be 1600 electrons per pixel based upon what is available. The shot noise of this system is equivalent to the square root of the signal. This means that the SNR of the system can be defined by:

$$SNR = \frac{N}{\sqrt{(\sigma_r^2 + \sqrt{D_C t}^2 + \sqrt{N}^2)}} \tag{2.20}$$

Graphs comparing the relationship between SNR and range, reflectivity, camera aperture size and exposure time are included in Figure 2.3. These models suggest that the maximum range of the system, using an exposure time of 33 ms, a reflectivity of 0.85, full diode power and aperture radius, is approximately 4 m before the SNR drops below 1. These also show that in general increasing the aperture radius, reflectivity and diode power increase the SNR as would be expected as more light will enter the system. Increasing the exposure time is also of key interest because whilst it does increase the SNR, it additionally increases the noise, though the signal scales linearly while the noise scales to the power of $\frac{1}{2}$, as a result of the shot noise and dark current. This means that increased SNR can be achieved by increasing the exposure time without fear of the noise increasing too quickly.

The models in Figure 2.3 show quite a low SNR overall, especially compared to what is observed in reality. There are likely a few reasons for this. The first reason is that the dark current and readout noise values seem unusually high, it is possible this is incorrect but it is used in the way provided by the manufacturer. Of note is is the finding that when testing another camera, the Raptor Owl 640 N, it was evaluated to perform very similarly to the Goldeye in a brief lab based test of the two devices. This camera quotes a readout noise of one 18 electrons per pixel (`https://www.raptorphotonics.com/products/owl-640-n/`), however this is a low noise variation of the camera which is not a quoted feature of the Goldeye, as such to make this comparison fairer, the readout noise from the Owl 640 II can be used as it is higher at 50 electrons per pixel (`https://www.raptorphotonics.com/products/owl-640-ii/`). Important to note is that the pixel sizes for the Owl cameras are half that of the Goldeye, assuming the readout electronics scale in size, and thus readout noise, with linear size, this means one can double this to estimate the readout noise of the Goldeye as being 100 electrons per pixel. Secondly is the fact that the data here is the raw unprocessed data without further post processing steps such as averaging spatially by filtering or temporal averaging which will improve the SNR.

Additionally, of note is the fact that in this system the definition of signal and noise is slightly different to what is discussed in this section. Going forward the signal is the differential signal generated per pixel by taking the difference between on and off absorption images. This means

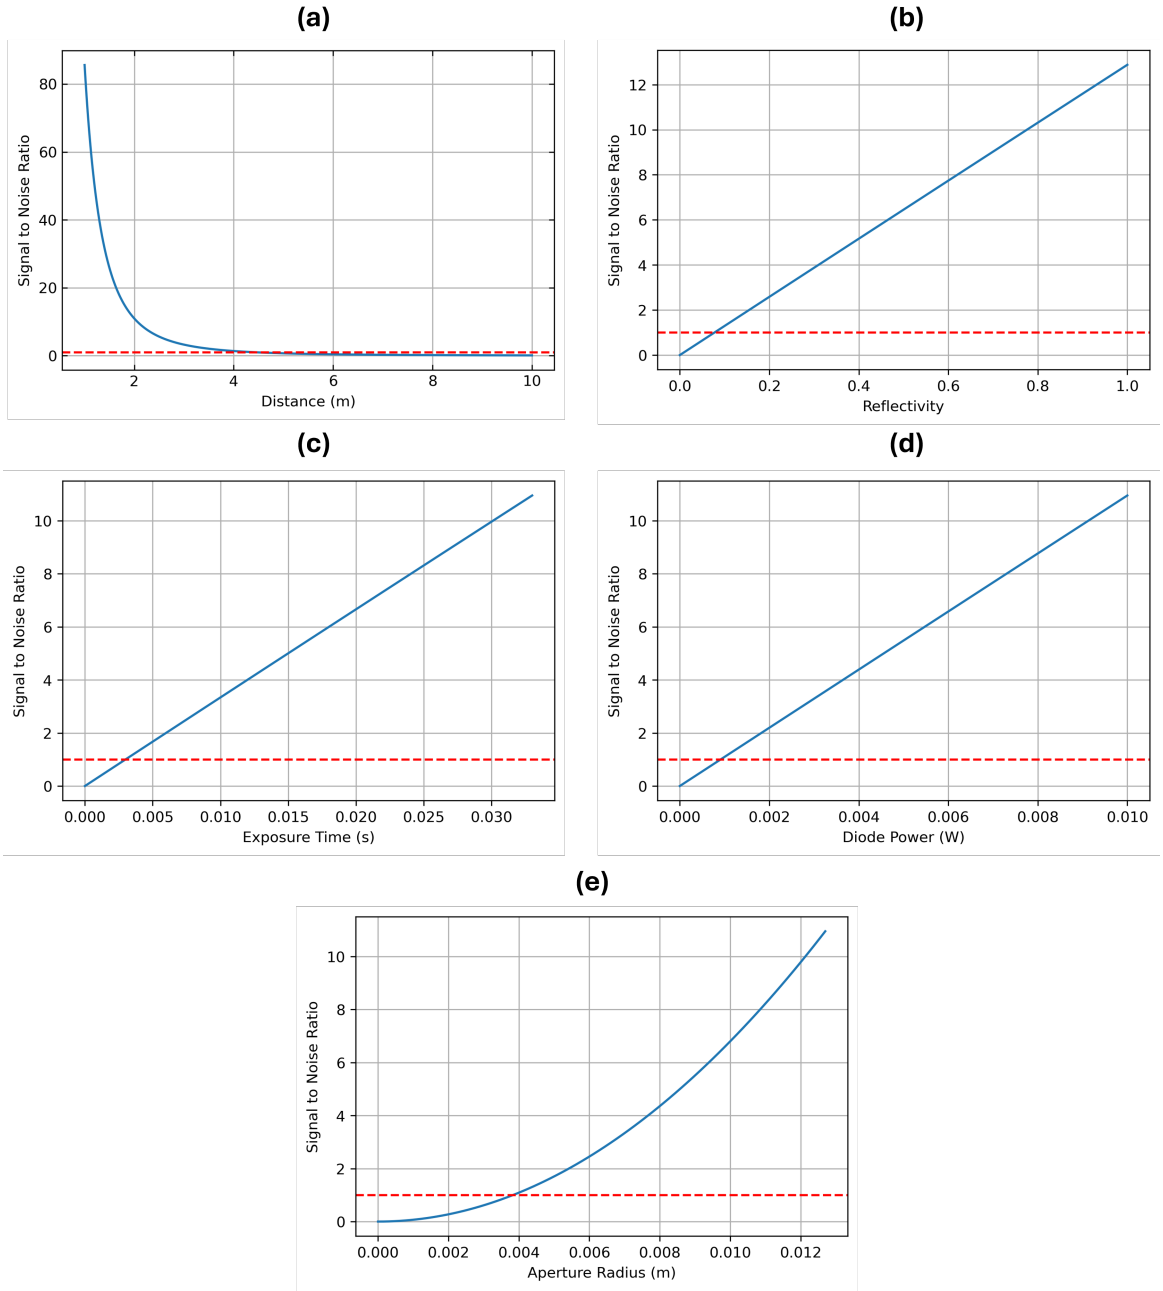

Figure 2.3: **Figures displaying the effect of varying (a) distance, (b) surface reflectivity, (c) exposure time, (d) diode power, and (e) aperture radius upon the signal to noise ratio.** A red line is included in order to indicate a SNR threshold of 1. In figure (a) a reflectivity of 0.85, an aperture radius of 12.7 mm, individual diode power of 10 mW and exposure time of 33 ms were used along with a varying distance. In figure (b) a distance of 2 m, an aperture radius of 12.7 mm, individual diode power of 10 mW and exposure time of 33 ms were used along with a varying reflectivity. In figure (c) a distance of 2 m, reflectivity of 0.85, individual diode power of 10 mW and exposure time of 33 ms, were used along with a varying aperture radius. In figure (d) a distance of 2 m, reflectivity of 0.85, an aperture radius of 12.7 mm, and exposure time of 33 ms, were used along with a varying diode power. In figure (e) a distance of 2 m, reflectivity of 0.85, an aperture radius of 12.7 mm,and an individual diode power of 10 mW, were used along with a varying exposure time.

that following on from the discussions above the actual signal would be the difference between an on and off absorption image pixel. A rough estimate of this signal can be calculated by using the data in Figure 1.2. This suggests that if the pixel is imaging methane gas of length 1 cm at 300 K, a mole fraction of 1 and a pressure of 1 atmosphere will have an absorbance of approximately 0.3. Absorbance, a, is defined by the following relationship [37]:

$$a = -\ln \frac{I}{I_0} \tag{2.21}$$

For I_0 the incident light intensity and I the resulting light intensity after absorption. As the system works in backscatter in reality this will be of the light as it passes through the gas twice, so rearranging the equation and accounting for the two passes gives:

$$I = I_o e^{-2a} \tag{2.22}$$

By using the equations above an estimate of the number of electrons detected at a range of 2 m, with an exposure time of 33 ms, as well as full sized aperture, full powered diodes and a reflectivity of 0.85 is approximately 17687. This means in the absorption frame, approximately 9707 electrons will be detected. Translating this to a difference in electrons detected this would be roughly 7980 electrons as the measured signal. It is also because of the differential nature of the signal definition that the noise in subsequent experiments is usually defined in terms of variance from frame to frame as opposed to detector noise, which is largely constant between frames.

2.2.4 Image Processing

The first step in the image processing is to register the images, which is achieved through calculating the positional difference between the two frames, this is obtained by utilising phase correlation. Phase correlation is a Fourier method for finding the offset between two images. Following the quantification and resolving of offsets between the images, each image is smoothed. The images undergo two smoothing algorithms, firstly by being convolved by a Gaussian signal

of size 7×7, this acts to spread the signal of each pixel over a 7×7, with a Gaussian distribution. The second convolution is with an 5×5 array of constants, selected to be 0.016, which then spreads this signal more evenly over a 5×5 pixel region. These smoothing algorithms act to improve continuity and smoothness of the gas signal, but also will average out issues caused by ocular speckle in the system. Speckle can result in regions of relative light or dark, so these convolutions act to average over the regions to make a more consistently bright image, the smoothing kernels must be larger than the regions of speckle for this to work. With the images now smoothed the difference is taken, ensuring the correct direction of subtraction, as discussed previously. In general, the issues caused by speckle here are very small and are dealt with entirely by the smoothing. The differential image is then passed through a threshold, where any pixel value that is negative is set equal to 0 and all positive signals pass unchanged. This approach decreases the noise of the differential signal greatly. Next a binary mask of the signal is created, the mask has value 1 where there is a signal and 0 where there is not. The mask is used as the basis of a persistency filter wherein a signal must be present in multiple successive frames in order to pass. In this instance a very weak persistency filter is used where signal must be in the current frame and the directly preceding frame to pass. This filtering is achieved by multiplying the mask of the current iteration by the mask of the previous iteration. This filter is effective for two key reasons, firstly it is effective at reducing random noise in the system as this is rarely persistent between subsequent frames, but it also fixes an issue associated with movement in the frame. When moving the camera the differential image generates false signals at the edge of the frame as a result of the change in signal due to the illumination profile caused by the diffusers. As a result of the differing order of subtraction in the imaging regime these false signals, which resemble solid bars, swap sides of the frame with each iteration, this is shown in Figure 2.4. This swapping in side of the false signal is due to the difference in temporal direction in which the subtraction is conducted i.e A-B subtracts the frame taken first in the imaging timeline from the image second in the timeline, whereas B-A is the opposite of this. The thresholded image is also added to the thresholded image from the previous iteration before being multiplied by the now filtered persistency mask, producing the gas signal image.

The next steps involve priming the gas signal to be added to the visible feed. This process

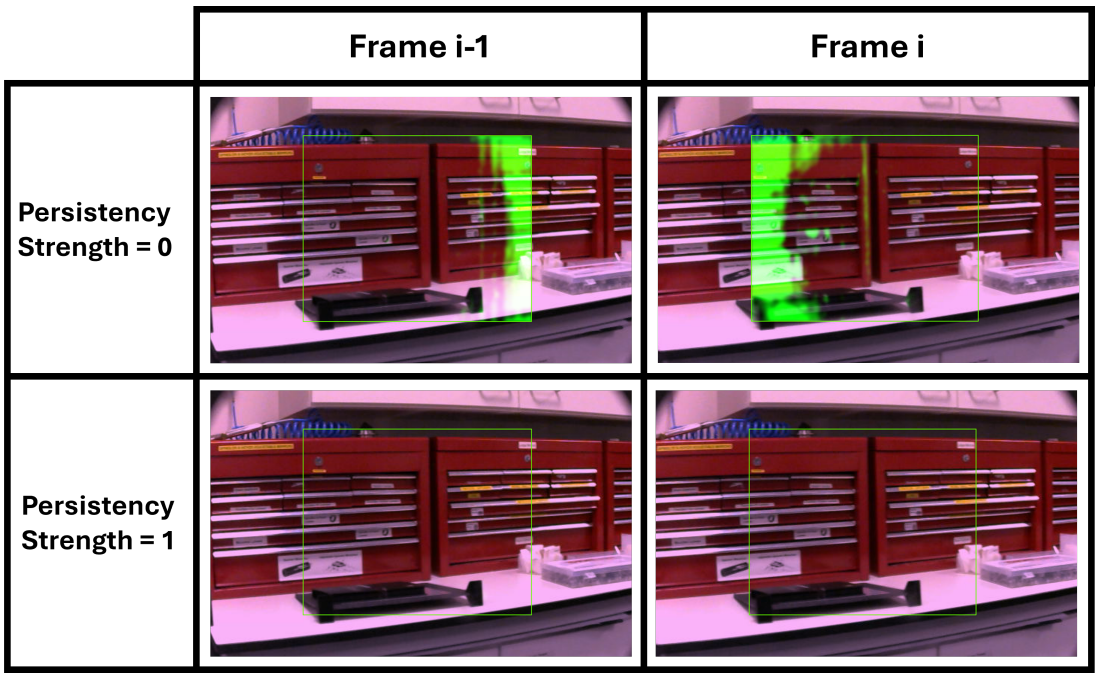

Figure 2.4: **Comparison of the presence of false signal bars on either side of the frame as a result of direction of subtraction swapping, and the impact of applying a weak persistency filter.**Here it can clearly be seen that the false signal bar swaps side based on the fact that the direction of subtraction temporally changes each iteration. Additionally, it is evident that by applying this persistency filter the bands of noise effectively disappear.

first starts by multiplying the gas signal by a sensitivity integer, this is a value that can be varied whilst the system is running and acts primarily as a gain for the gas signal. The signal is then smoothed again by convolving with a Gaussian signal, however it is normalised by dividing by the sum of the Gaussian kernel. Meanwhile the 8-bit green channel is extracted from the visible image and multiplied by 0.7, this is to prevent the channel becoming saturated too quickly when the gas signal is added and also acts to enhance the gas signal relative to the rest of the scene. The pixels that overlap with the SWIR camera pixels are extracted and saved as a subset of the visible camera data. A copy of the gas image is now divided by 16382 (i.e. 2^{14}), such that the values only range between 0 and 1, these values are subtracted each from 1 and multiplied by the attenuated green channel subset, priming the relevant pixels to have the gas signal added to them. Separately, the original gas image is converted to an 8 bit number by dividing by 64, this is what is added to the now primed green channel subset and rounded such that no single precision floats are added to the 8-bit channel. The green subset, now containing the gas signal is placed back into the attenuated full green image, which is then used to replace the visible images green

channel - this is what is displayed on the GasSight, revealing the gas location. For additional clarity a green border is added to the visible feed that highlights the region of interest (ROI) in which visible pixels overlap with SWIR pixels, essentially highlighting the active region of gas detection in the frame. Whilst there is a correlation between the saturation of the green channel with concentration of gas detected, it is not calibrated in such a way that it can accurately reveal this concentration value and acts more as a binary detection of the presence or lack of gas.

2.2.5 Size, Weight, Power and Cost

The overall size of the completed system had dimensions of $19.4\,\text{mm} \times 12.4\,\text{mm} \times 28.3\,\text{mm}$. The mass of the system was additionally measured to be 3.15 kg. The system is powered centrally with a 14.4 V, 5.2 A h lithium ion battery, which can be recharged using a 24 V power supply. The battery gives GasSight approximately 2 hours of battery life, suggesting a current draw of approximately 2.5 A. GasSight is also equipped with a HDMI port as well as a USB socket, this mean it can be plugged into a monitor and controlled like a PC. An image of the system as well as a schematic representation of presented in Figure 2.5. The cost of the system is primarily a result of the SWIR camera and the diodes. The SWIR camera cost approximately £9,000, and the diodes were around £1,000 each. The rest of the components are a mix of 3D printed and off the shelf components in order to keep additional costs low. It can therefore be approximated that the total cost is around £12,000, however as SWIR camera technology develops this cost will decrease with the camera.

With regards to the associated ocular hazard of the system, it was designed such that the laser power output by the system is of class 1, or eye safe. Class 1 was a key consideration of the system in order for it to be used as safely and widely as possible. As the system works in the SWIR the ocular hazard is lower than those in other bands, with an advised upwards band of 10 mW for a collimated beam to be considered class 1. Whilst the individual diodes are capable of exceeding this, and on their own are classified as class 3R the use case of this system requires a diffuse beam, which by the time it is accessible by a power meter will read less than 5 mW due to the impact of the distance between the diode and the first collimating lens and the external diffuser spreading the beam out further. The diodes themselves are inaccessible to the user as

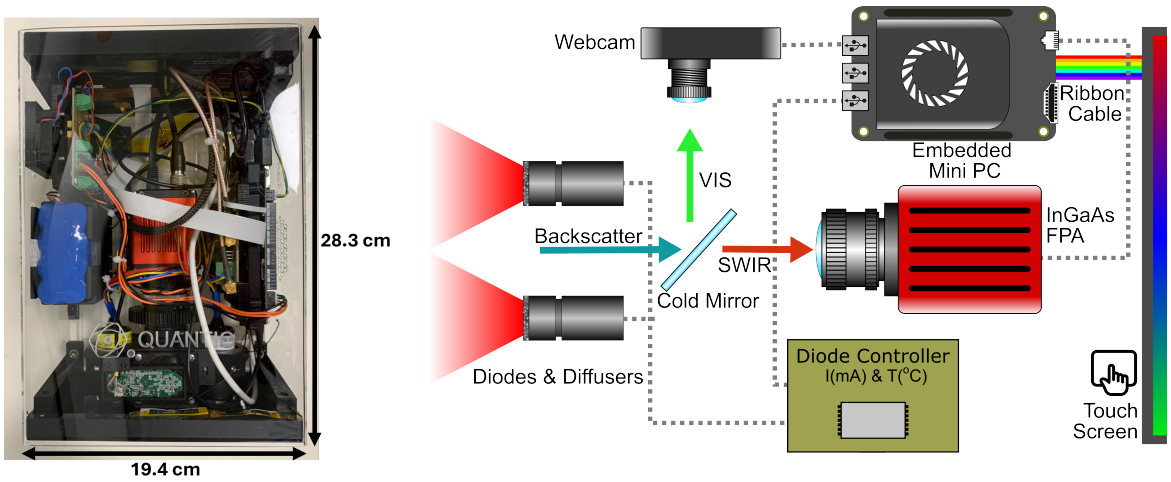

(a) (b)

Figure 2.5: **Image showing (a) a top down view of the GasSight system and (b) a schematic diagram of the individual main components.** Figure adapted under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from the paper "Characterising the performance of a drone-mounted real-time methane imaging system" [97], where scale bars have been added and the image has been cropped.

they are sunken internally into the system, additionally the diffusers are also inaccessible to a user and cannot be unscrewed without opening the system

2.3 Previous Work

Before the development of the GasSight, gas detection research within the Optics group at Glasgow had been extensively covered. Research relating to gas detection in the group dates back to 1998, with work conducted in the paper "Continuous-wave optical parametric oscillator based on periodically poled KTiOPO4 and its application to spectroscopy" by Gibson et al. [107]. This research demonstrated a device capable of detecting the absorption spectrum of methane around 1649nm, whilst not a imaging system, it does highlight the use of methane cells and absorption spectroscopy in the group before the more direct focus on imaging systems.

An ethane detection system was also developed in "A field-portable, laser-diode spectrometer for the ultra-sensitive detection of hydrocarbon gases" by Gibson et al. [108]. This system is a dual ethane methane detection system which utilises 3.34 μm lead-salt laser to investigate a cell for the presence of both gases by scanning across a wavelength region. The cell is pumped

with nitrogen to obtain a reference before drawing in atmospheric gas to investigate. A reference frame is taken every 5 minutes to ensure that the background reference is relevant. The system also requires liquid nitrogen temperature control, and a 2 kW power supply meaning that this system must be transported with an appropriate power supply of suitable size. The system does not feature an imaging component but does detect very low concentrations of gas with observations averaged over 1 minute being capable of detecting up to 50ppt of ethane gas. The ethane detection capabilities of this system have uses in medical science which has been explored further in separate studies [109–111].

This system was adapted, in collaboration with Shell, and is discussed in "Oil and gas prospecting by ultra-sensitive optical gas detection with inverse gas dispersion modelling" by Hirst et al. to create 2D images (maps) of hydrocarbon emissions [112]. This technique is again looking for ethane as opposed to methane due to ethane having a lower overall background concentration such that detecting ethane may be more indicative of a non-natural source, in this case an oil or gas reservoir. The device used for this detection is a version of the lead-salt laser system as described above but used with wind data to create a map of the gas distribution. To improve accuracy of wind speed and atmospheric measurements the inlet and wind measure were extended to an altitude of 5 m. The maps are made by way of an inversion process and utilise an x and y wind direction measurement and a concentration measurement. By using this wind data in addition to GPS location data accurate maps of the hydrocarbons could be reconstructed.

The final detector version of the gas sensor was discussed in "An open-path, hand-held laser system for the detection of methane gas" by van Well et al. [113]. This version could be back pack mounted and could measure gas concentrations over a path length of 1m. It is of note for utilising a tunable laser diode and for swapping to the 1.6 μm absorption line, both featured that remain in the more modern iterations of the system, additionally it featured a handheld component and was portable. This system was then adapted into an imaging system in “Imaging of methane gas using a scanning, open-path laser system” by Gibson et al. [114]. This design utilised an InGaAs detector, an amplifier, a Fresnel lens, a charge coupled device, a laser diode and a temperature controller. This system illuminated a scene with the laser and detected the back-scatter to discover information about the nature of the medium the light had

traversed. The images obtained with this method are fairly low resolution but this can be partially remedied via up-scaling and resampling. A higher resolution was not pursued in favour of a fast processing speed to achieve higher frame rates. A key discovery of this paper was that utilising back-scattering is sufficient, outside on uneven and varied surfaces, to successfully image the gas. A similar scanning system is possible at around $3.27\,\mu$m, which results in an stronger absorption than at $1.67\,\mu$m [115]. The absorption strength of methane at this wavelength was reported to be greater than four times that of water vapour and carbon dioxide, making it an ideal candidate [116]. The output laser from the pump in this scenario has a power of around 50 mW, which is approximately $5-10\times$ times greater than what is achieved by the laser diodes in the first GasSight system, this allowed for greater contrast between on and off absorption frames and helped to reduce overall noise in the system [115]. The wavelength was accessed using a pump-enhanced optical parametric oscillator which is considerably more expensive than a laser diode, similarly the detector in this system exhibited more noise and as such had to be thermoelectrically cooled, which is not a consideration in other systems [115].

Another detector created by the optics group, utilising similar techniques, was the single pixel camera detector presented in "Real-time imaging of methane gas leaks using a single-pixel camera" by Gibson et al. [117]. Single pixel imaging can allow one to achieve high frame rate video. Imaging in this scenario made use of structured illumination, which involved the use of a spatial light modulator, in this case a digital micromirror device (DMD). The DMD could be programmed to illuminate a scene with a specific pattern such as Hadamard patterns. The DMD took a finite time to swap between such patterns and this set limits upon the maximum possible frame rate one could achieve. Compressed sampling could also be implemented in the image processing, this allows an image to be reconstructed from less data points than there are pixels in the image, by methods such as a priori frame weighting. This system also made use of differential imaging, meaning that it imaged one frame at the absorption length of methane and the next at a close non-absorbing wavelength. The difference between these two images highlighted the gas, which could be overlaid onto a visible image of the scene. The key components that were required to achieve this system are as follows: A DMD, a complementary metal oxide semiconductor (CMOS), a photo-diode, a laser diode, a projector lens and a collection

lens. This technique permitted the gas to be imaged with a resolution of 16×16 pixels, this could be improved again by using smoothing techniques. The reason for using a photodiode and single pixel imaging, as opposed to a SWIR camera, was due to the high cost of the technology at the time, with a photodiode being a more cost effective option. This design behaved similarly to previous designs which utilised four separate detectors, for red, green, blue and infrared light respectively, as well as a hot mirror and x-prism to split the light into the relevant constituents [118] This system, whilst unable to detect gas, was able to simultaneously image in SWIR and visible and the use of a hot mirror is reminiscent of the introduction of a cold mirror to a later GasSight revision.

As described above, the most recent iteration of this system prior to the beginning of this PhD is the GasSight, as presented in "Developing a portable gas imaging camera using highly tunable active-illumination and computer vision" by Nutt et al. [41]. By this point in time the cost of the SWIR cameras were low enough that it could replace the single pixel technique. The general operation of the system is not dissimilar to that of the previous iterations in that it utilises differential imaging. The design of the system is demonstrated in Figure 2.6, which highlights 4 ports on the front of the system, relating to the two illumination sources as well as the SWIR and visible feeds. The introduction of two diodes as opposed to one is to increase return signal and illumination area of the scene. The primary difference between this and the single pixel variant is the fact that it is now a handheld device, that can be moved and used far more easily in a wider variety of scenarios. Another benefit is that a higher resolution gas image can be obtained based solely upon the camera utilised, which in this case is the Goldeye G-008 focal plane array (FPA), which offers a 320×256 pixel resolution. This version attains a lower frame rate than what is achieved by the single pixel variant with 10 fps compared to approximately 25 fps. This decreased frame rate presents an opportunity for improvement. The gas camera also now has to take into account movement and how this will affect the difference between subsequent frames, the solution for this is phase correlation. This phase correlation approach has limits in how far and how quick it can work however, and may be a road towards improvement. The lens system within the camera can also be manually focused allowing it to be utilised at a variety of distances from the gas source. This system was envisioned as being used in conjunction with a drone, such

that pipeline leakages could be monitored remotely.

A figure comparing the form factor of the camera systems over time are shown in figure 2.6, additionally a timeline summarising the key advancements of the research within the Optics group over time is included in Figure 2.7.

In between the publishing of the 2020 paper and the beginning of the PhD, some further edits were made to the system. The primary alteration made to the system was the introduction of the previously mentioned cold mirror. This cold mirror was introduced in the prior systems SWIR aperture and acts to pass SWIR light but reflect visible light. The mirror was positioned such that incoming visible light is reflected upwards, at a right angle to the axis of entry. The visible cameras aperture at the front of the system is instead removed, and the visible camera re positioned to point downwards, thus collecting the incoming, reflected visible light. The action of removing the visible aperture and instead sharing the aperture between both cameras acts to make the two systems co-linear. This co-linearity acts to prevent issues pertaining to parallax and ensures that the SWIR feed and the visible feed are exactly the same in terms of position and imaging angle. The cold mirror must be placed before either lens system, meaning that, unlike the exposed lenses apparent in Figure 2.6, the lens systems are now sunken into the shell and inaccessible to an outside user. Whilst this is a minor change it does mean that the lenses are protected within the system and are less prone to scratches or scrapes, making the system more robust for use in the field.

2.4 Conclusion

Overall, gas imaging projects have been a constant source of interest for the Optics Group at the University of Glasgow. The GasSight system was introduced in addition to a summary of how the system functions, providing context for the rest of this thesis. All experiments throughout the course of this research were built upon, or a modification of, this device, so it was vitally important to introduce it at this point. In introducing the device, important parameters such as how far the diodes tune and an estimation of the range were discussed which will add context to decisions made throughout this thesis. Whilst it was important to introduce the starting point for

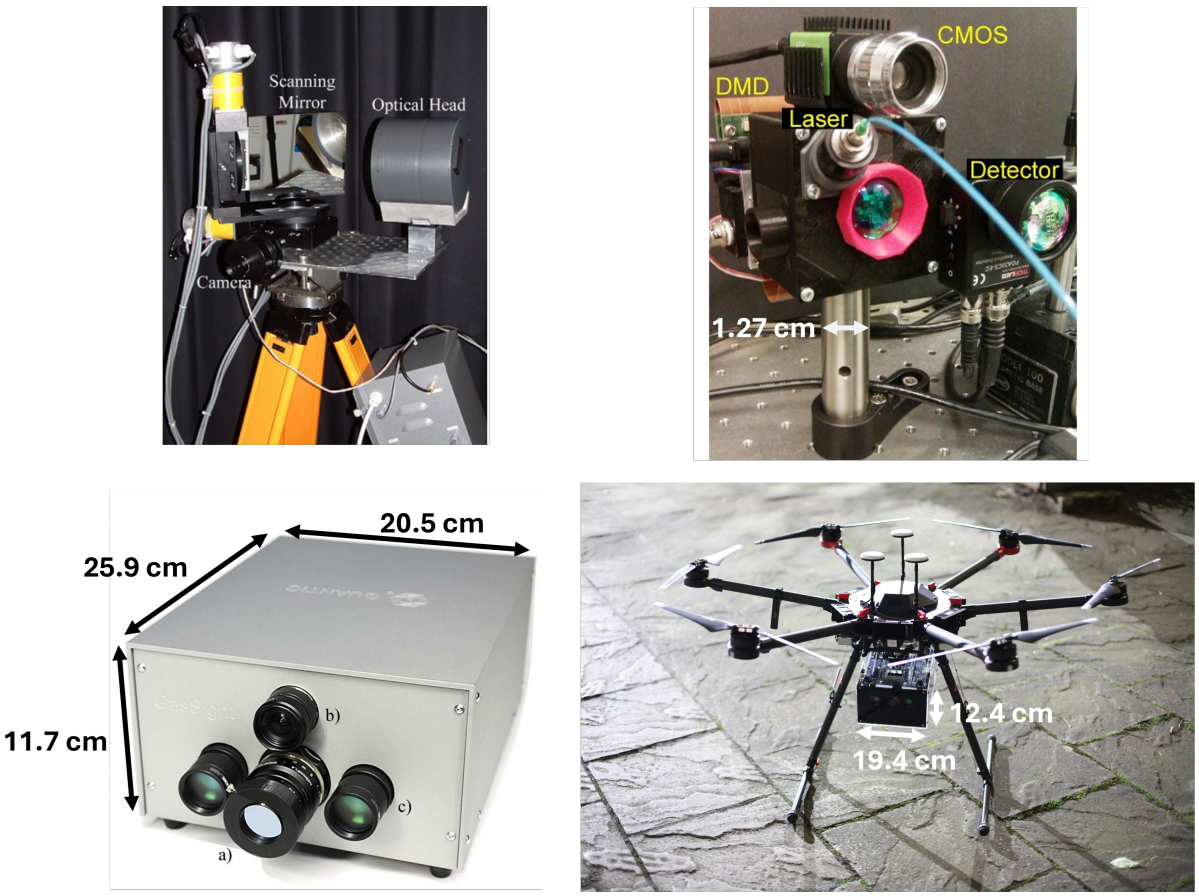

Figure 2.6: **Figure representing the design evolution of the gas camera.** The top left image is the scanning system, highlighting the scanning mirror and optical head [114]. The top right image represents the single pixel design, this highlights a now decreased size as well as the DMD, CMOS, detector and laser [117]. Image on the bottom left highlights a more portable version of the gas camera, (a) refers to the infrared lens system for the focal plane array, (b) refers to the webcam used for the colour images and (c) refers to one of the two laser diodes which are within lens tubes [41]. The bottom right image highlights the latest drone-mounted gas camera revision, which combines the lens systems for the optical and infrared signals by splitting the incoming light using a mirror. This version also has an internal computer and touch display, meaning the whole system is fully self-contained and easy to control. Where possible a sense of scale was provided for size comparison, no accurate measurements for the top left device could be found retroactively, but of note is the inclusion of a 150 mm Fresnel lens within the optical head providing a minimum radius for this component [114]. Top left image reproduced under a creative commons license (`https://creativecommons.org/licenses/by-nc-sa/3.0/deed.en`) from "Imaging of methane gas using a scanning, open-path laser system" [114]. Top right image adapted under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from "Real-time imaging of methane gas leaks using a single-pixel camera" [117], where a scale bar has been added. Bottom left image adapted under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from "Developing a portable gas imaging camera using highly tunable active-illumination and computer vision" [41], where scale bars have been added.

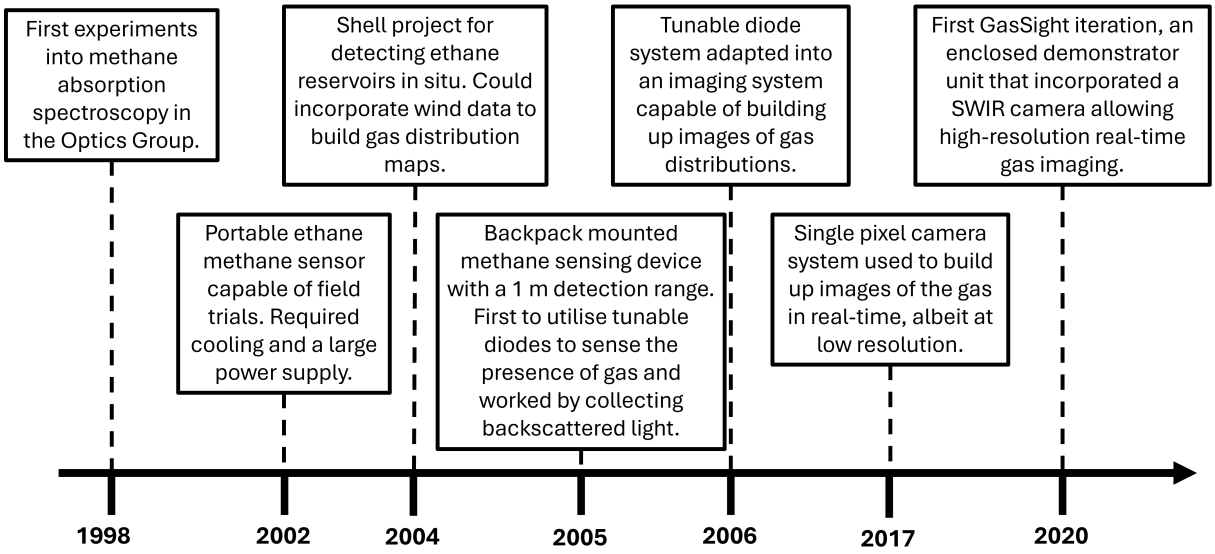

Figure 2.7: **A timeline summarising the key advancements of the gas detection systems researched within the Optics Group at the University of Glasgow.**

this research, it was similarly as important to place the device in context against the many years of research culminating to the work presented here. It is clear that gas detection research by this group has spanned many methodologies and use cases, such as by tracking plume movement and wind speed to identify sources of leaks, to single pixel cameras and even into health care with breath analysis. There have been experiments with making the systems handheld, making the systems fast, and making the systems robust, all of which lead up to the GasSight device which aims to capitalise on all three of these design challenges. How this system has been updated, the use cases, and further design alterations will be the focus of the remainder of this thesis.

Chapter 3

Movement Suppression

3.1 Introduction

For the practical usability of GasSight in real-world settings, outside of controlled lab conditions, it must be relatively robust to movement. The differential nature of the gas detection mechanism means that differences between two subsequent frames can be flagged as signal as they bypass some of the noise suppression filters. These false signals add a degree of uncertainty to the system, which becomes a key issue in terms of usability, and adds a degree of risk to measurements as these signals need to be accurately differentiated from real gas signal so as to not misinterpret the severity of a leak. In this chapter, the steps to strengthen the system against movement will be introduced and the results of their implementation presented. The relative strengths and weaknesses of the different approaches are highlighted. Some of the explored techniques were not implemented into the system but are important to consider as potential solutions to future problems. Parts of the discussion in this chapter was published in the conference proceedings: "Real-time imaging of methane gas from a UAV mounted system" [96].

3.2 Changes to the Image Processing

The most effective way to reduce the motion artefacts is to ensure that the two frames being subtracted from one another are as similar as possible. Artefacts arise due to differences between

the frames. These can be very simple differences and mainly result from new positioning of objects within the frame, either due to something in the frame moving or the camera itself. The size of the artefact is directly proportional to the extent of the difference between the two frames. By having two images that are taken closer together temporally, one can minimise the differences between frames, this can be achieved via multiple routes.

In an earlier version of GasSight, the SWIR camera used a software for-loop to take one SWIR image with an exposure time of 33 ms per iteration. The downtime between images was set primarily by the time required to process the images, there was a single loop in which the SWIR images were processed and overlaid upon the visible image. This loop was set to take 100 ms, or upon failing to be completed in 100 ms, the loop instead was instructed to take 200 ms as a result of the strict timing parameters set in LabView. Within the processing loop, the first operation performed was the tuning of the diodes, with odd/even loop iterations switching between on/off resonance. The SWIR and visible imaging were controlled in separate loops which constantly captured new images and saved them to a respective SWIR or visible image memory location. The latest image saved to the memory location was the image used in the next iteration of the processing loop. It was difficult to guarantee when the latest image in the processing loop was taken relative to the other as there could be as much as a 200 ms difference between them depending on when in time the processing loop accesses the memory location. By having extended periods of downtime, it was evident that differences between the two frames would be inevitable.

In order to reduce the length of system downtime, it was decided to change the overall image processing algorithm such that two SWIR images are taken in the same loop in sequence with one another. In order to adjust for the fact that capturing two images per iteration will slow the process down, the loop was instead split into two separate loops, one taking SWIR images and generating differential images, and the other capturing visible images and handling the false colouring and overlay of gas signal. An updated flow diagram of this is highlighted in Figure 3.1, which highlights the subtle difference between image processing sequences in Figure 2.1, with the addition of a timeline. In order to expedite the imaging pipeline further, it was decided to decrease the exposure time of the SWIR camera from 33 ms, to 20 ms. While this may have a

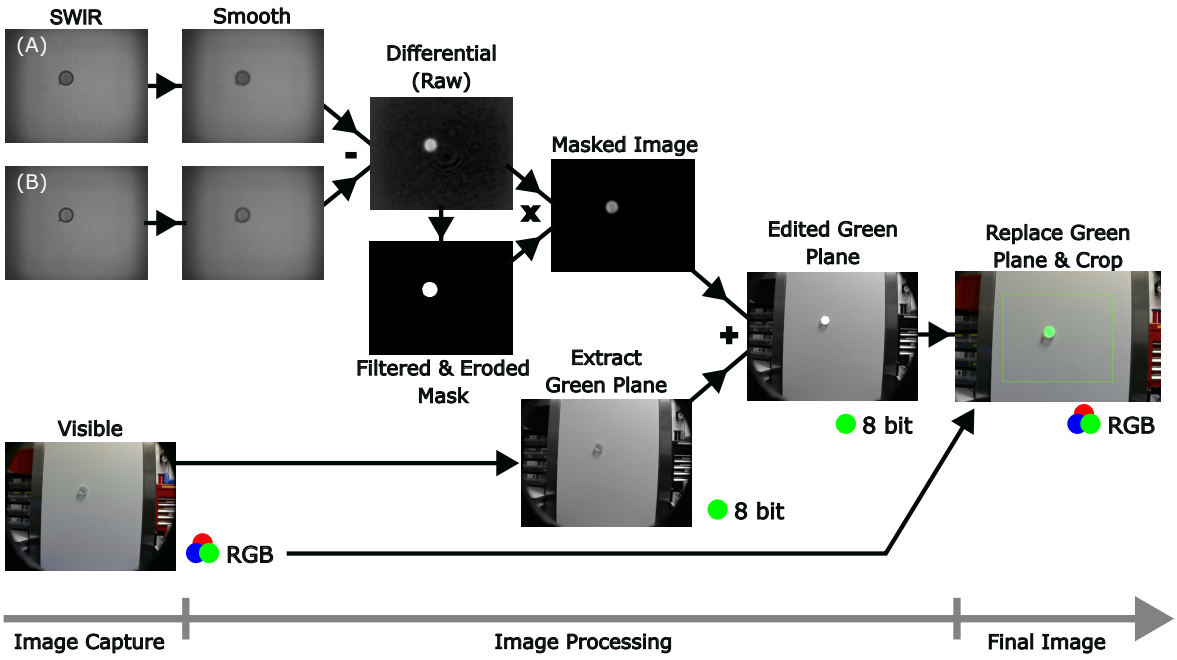

Figure 3.1: **An updated flow diagram depicting the image processing of the GasSight system after changing when the SWIR images are taken.** The main differences from the previous flow diagram shown in Figure 2.1, is that the two SWIR images are taken in quick succession, meaning Image B is never reused as Image A in the next loop iteration. Below, a timeline is present to more clearly highlight when the processes are happening. In the previous version of the image processing, there was less constraint on when the images were taken. Figure reproduced under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from "Characterising the performance of a drone-mounted real-time methane imaging system" [97].

minor impact on the SNR of the system, the decrease in movement noise, as a result of reduced downtime between images, was expected to be an appropriate trade off for future measurements that require the system to move.

To acquire the images in quick succession, one must still tune the diodes on and off resonance. Tuning occurs between the two SWIR captures, however, an issue related to this imaging regime is the tuning time of the diodes. As the diodes take a finite time to tune on and off, this sets a fundamental hard limit upon how fast the system can be run in such a way, as if the diode is not fully tuned when the image is taken then signal strength is decreased. The different tuning times can be observed in Figure 3.2, which shows the finite time required to swap between the on and off states as a result of the temperature tuning. The work around for this is to introduce wait times into the system that allow the diode to finish tuning in-between images. The differences in tuning times t_2 and t_4 are noticeable because of the fact that times t_1 and t_3 are so similar. The reason for this difference in tuning time is likely due to a combination of three factors. the first is that the two diodes do not tune over the same size voltage difference with one tuning over a difference of 0.07 V and the other 0.05 V, this will directly impact the amount of current across the diode which in turn can cause additional temperature changes that could increase settling time. Next is the fact that the two diodes are set at different operating temperatures, with one operating at 21.9 °C and the other 26.4 °C, these differences in temperature could impact the settling time as a diode operating at a temperature further away from room temperature may have to perform more cooling or heating to reach its goal. These temperatures were selected based on manufacturers data of where to find the methane absorption line. Finally, the largest reason for the difference in settling time will be batch to batch differences and imperfections in the manufacturing process this is, for example, why the diodes tune at different temperatures or drive currents.. In order to further increase the efficiency of the process, tuning is only performed once per iteration and the direction of subtraction swapped with each iteration. Swapping the direction of subtraction ensures that only a single wait is required in each iteration of the loop. However, although the need for two separate wait times within each loop, an additional issue was introduced requiring consideration. As the tuning is in part temperature dependent, the time taken to tune in each direction is not the same, meaning that for maximum efficiency the wait

times could differ for odd and even iterations of the loop. Overall, the average iteration time of this imaging loop is approximately 70 ms, resulting in a differential SWIR image frame rate of approximately 15 fps. Whilst, there is only a small change in the length of time each iteration takes when compared to the previous imaging timeline, the two SWIR images are taken far closer together in time, with some improvement to the overall frame rate of the system from 10 fps to 15 fps.

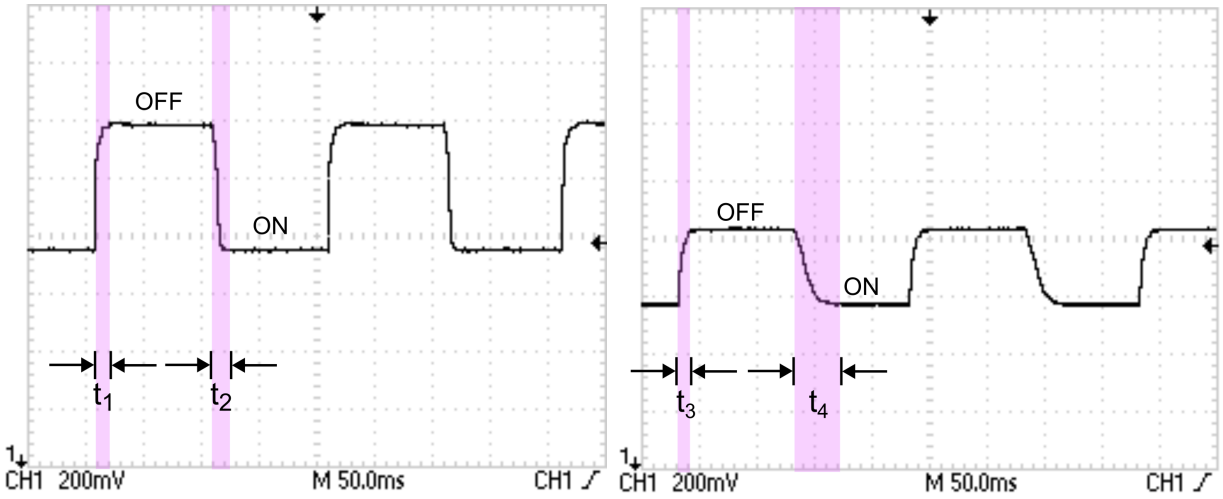

(a) (b)

Figure 3.2: **Images highlighting the tuning times for (a) the left diode of the device and (b) the right diode.** Here it can be seen that the finite tuning time of the diodes as they tune on and off absorption, showing that both directions do not take equally the same amount of time to settle between tunings. Each division along the x axis is 50 ms It can be observed in (a) that tuning to the off absorption wavelength, t_1, takes approximately 10 ms, where as the time taken to tune to the on absorption wavelength, t_2 takes roughly 15 ms. In (b) an off absorption tuning time, t_3, of approximately 10 ms can be observed, with an on absorption tuning time, t_4 of as much as 35 ms recorded.

Figure 3.3 highlights the impact of varying the wait time, that allows the diode to fully tune, upon differential signal strength and uncertainty. This figure was made by following the same process that will be described in greater depth in Section 5.5, where the average differential signal pixel value is calculated over the extent of the signal object over 100 captured frames. The uncertainty is the overall standard deviation of this value over the same 100 frames. It can be seen that as the wait time increases from 0 ms to 10 ms, the differential signal begins to increase and the variance clearly decreases. This decrease in variance is a direct result of the fact that, as the wait increases the camera captures less of the transition between on and off states and

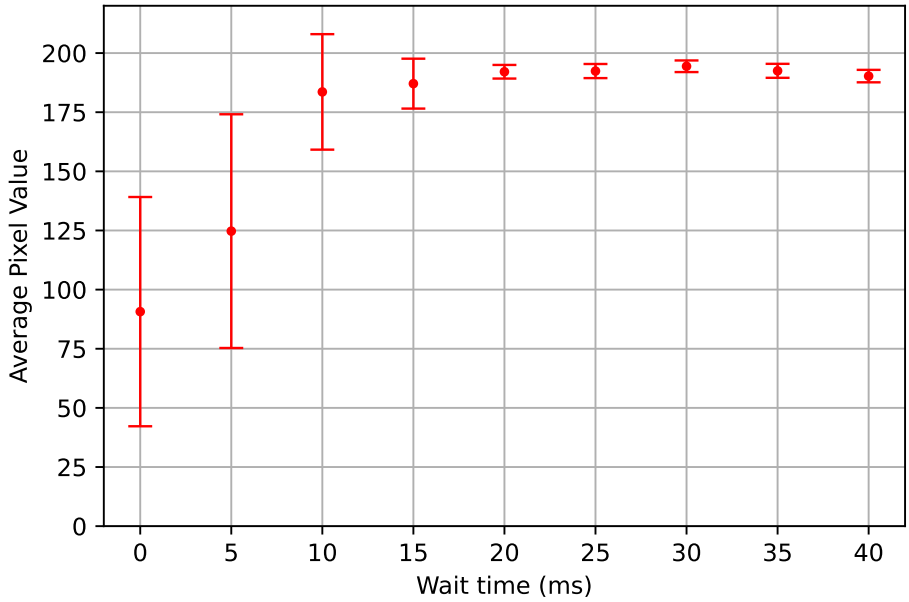

Figure 3.3: **The effects of increasing wait time to allow diodes to tune between consecutive images on differential signal strength and variance.** This data was generated by looking at a single gas cell of approximately 10,000 ppm.m at close range and performing the full imaging pipeline to the captured images, allowing average differential signal strength and standard deviation over the full set of images to be calculated, over the size of the signal region.

as such, a smaller percentage of the capture is spent observing the diode mid-tuning. The fact that the settling time differs from on to off is highlighted by the initial rapid change in average signal value and variation from 0 ms to 10 ms, followed by a shallower change from 10 ms to 20 ms. The difference in tuning time for each direction of tuning was highlighted in Figure 3.2, from which it can be observed that in one direction there is a tuning time of approximately 10 ms and the direction taking approximately 20 ms, matching the trend in Figure 3.3. This difference follows the idea that an optimal wait time in one direction will be obtained quicker than the other. From figure 3.3, it can be seen that, in general, adding a wait of 20 ms is enough to provide a relatively stable signal, however, in reality this means adding a wait time equivalent to another full exposure time, wasting about a third of the total available photons. As a result a wait time of 10 ms is used as this provides almost equivalent signal but with the caveat of greater variation. In general the true limitation on the imaging frame rate is the diode tuning time.

3.3 Image Correlation

The most effective way to remove movement noise from the system is to not move the system at all and to image only stationary objects, this would ensure no positional discrepancy between objects in the frame and, as a result, no false signals. A stationary system works well in a lab setting but is very limited in use case externally, so the camera needs to be able to truck and pan in order to be portable. A solution to this problem would be to use the position of the objects in the frame to calculate by how much the camera has moved and use this to correct the positioning of the objects in the two consecutive frames relative to one another. A technique that allows one to quantify movement is correlation.

As mentioned in Chapter 2, the correlation technique used initially in the gas camera was phase correlation. Phase correlation allows the Fourier components of the two images to be investigated in order to find the offset between the two. The equation for this is [119, 120],

$$P = \mathscr{F}^{-1}\left(\frac{\mathscr{F}(I_a) \circ \mathscr{F}(I_b)^*}{|\mathscr{F}(I_a) \circ \mathscr{F}(I_b)^*|}\right). \tag{3.1}$$

Here, I_a is one SWIR image, labelled A, I_b is the other SWIR image, labelled B, and P is the resultant image. Within the resulting image, there will exist a peak at the pixel corresponding to the horizontal and vertical offset. For example a peak at location (4,5) suggests an x-offset of 4 pixels and a y-offset of 5 pixels. Once this offset has been calculated, the images can be moved such that they overlap perfectly. This offset is implemented by shifting one of the images by the calculated amount, this shift is achieved by cycling the rows and columns of the image around. The issue with this cyclical swapping is that at the edges of the image the pixel values become incorrect. For example, for an image with M columns, if the image requires a shift of 5, then the pixel values in column 1 will be shifted to column $M-4$. After shifting, the images are cropped by the size of the offsets, in order to prevent edge effects. Any minor discrepancies can also be filtered out by the persistency filter, introduced in Chapter 2, that is applied to the image.

Whilst image correlation is an established technique, it was identified as being a potential frame rate bottle-neck in the post processing of the system because of the time required to calculate the Fourier transforms [120–123]. In order to prevent this, two similar functions to

phase correlation were implemented which does not require the use of Equation 3.1 and removes the need for Fourier transforms but instead checks for a location of minimum difference by calculating the difference between the two images with offsets applied to the original image, in a kind of cross correlation method. These two different approaches, referred to as full area cross correlation and simplified cross correlation, are introduced in the following.

3.3.1 Full Area Cross Correlation

The first of these methods is the full area cross correlation, which investigates a small area of the images for a likely offset. Whilst it would be much slower to move the full first image over the entirety of the second image pixel by pixel, if one was to define a maximum possible offset, N, then only a small number of pixel values need to be investigated. There is an inherent aspect of risk in including a maximum offset, as if the offset is greater than the limit then the operation will fail, and as a result there must be careful consideration as to what this value must be such that it is not over or underestimated. A low maximum offset between 5 and 10 pixels is usually used, this is primarily to ensure the process is quick and, in the case of failure, the incurred noise will likely be filtered out by subsequent filters introduced in this chapter. A visual summary of the full area cross correlation method is presented in Figure 3.4, which highlights the changing size of the images, as well as the location of the image. The process involved in calculating the offset is as follows: firstly, the reference image, to which the other image will be compared, is cropped by the size of the maximum offset, this was performed in order to essentially offset the image, this is show at the top of Figure 3.4. This crop is applied such that it crops from the left and from the top of the image. The second image, referred to as the move image, is then cropped iteratively within a for loop until it has been cropped by the amount of $2N+1$ in each direction, creating $(2N+1)^2$ cropped images varying in size, which is demonstrated as the second step in Figure 3.4. Each cropped move image is subtracted from the reference image, the disparity in size is accounted for by LabView which will automatically crop the larger image to match the size of the smaller image. As a result of the the reference image being already cropped by N, this means that the iteration i=0 acts as having the move image in offset position -N, thus allowing the negative direction of offset to be investigated. The +1 in the $2N+1$ originates from

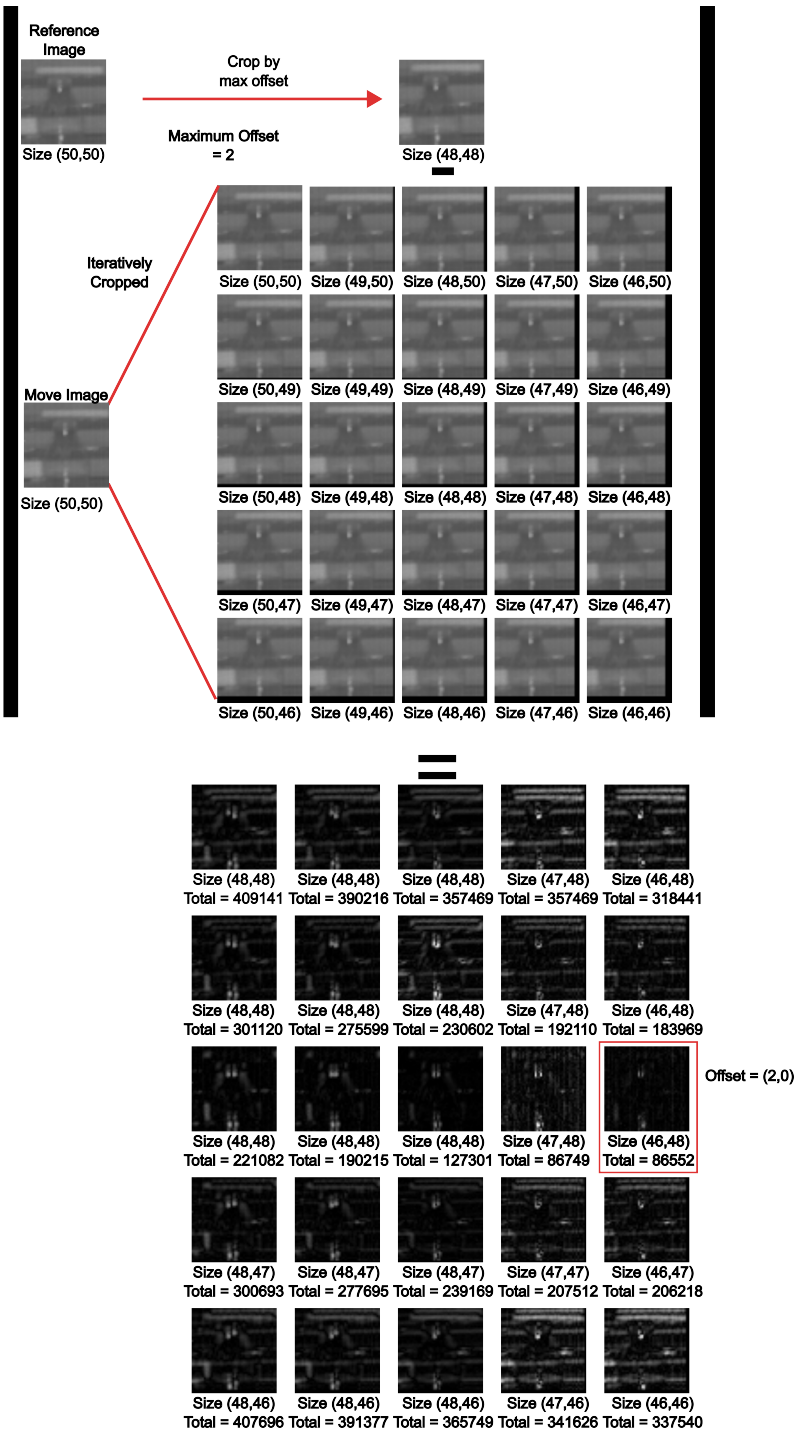

Figure 3.4: **Diagram showing the process for calculating the offset between a reference image and a move image.** The reference image is firstly cropped by the size of the maximum offset, N, which here is 2. Cropping is performed in order to allow both positive and negative offsets to be investigated. The move image is iteratively cropped by up to $(2N+1)$ for all crop permutations in each direction, and these cropped move images are subtracted from the cropped reference image. Any size mismatch is handled by LabView, due to the in-built properties of the subtract function, which crops the larger image to match the size of the smaller image. The magnitude of each differential image is taken and the pixel sum then calculated, and compiled to generate a 2D array of sums. The minima position relative to the centre of the array is identified as being the offset in x and y. In this case a positive offset in the x direction of 2 pixels, with no detected offset in the y direction.

the need to additionally include a calculation for where the offset is zero. For example if the maximum offset was 3, the process would investigate offsets in position -3, -2, -1, 0, 1, 2, and 3 which is 7 points, or $(2 \times 3) + 1$. The magnitude of each differential image is taken and the sum of the pixels is calculated, and an array of totals of dimension $(2N + 1)^2$ is generated, as shown in the bottom array in Figure 3.4. The addition of the magnitude before summing allows the difference between the reference and move images to be quantified as any difference, regardless of direction, will suggest an offset. A lower sum means that the images are more similar and thus in the array the minima position will correspond to the offset that generates the least movement noise, and the maximal level of overlap. One extra step that must be taken is that because of the initial cropping of the reference frame the maximum offset is subtracted from the coordinates of the minima in order to calculate the offsets.

One potential issue that stems from this process is the fact that the images are made smaller with each iteration. This was implemented for two reasons, firstly it was quicker than having the pixels sequentially shifted around and secondly, swapping the pixel locations around may have dominated the differential sum masking the offset. One worry could stem from the fact that, because the image has been cropped, there are less pixels to sum over and as such less signal in general which could cause the local minima to erroneously be identified elsewhere. In general, however, this was found not to be the case as the introduction of false signal and noise from poorly overlapping images greatly increases the sum, and will increase as the offset becomes a poorer fit, this means that the point of optimal correlation remains clear despite the decreased size. Further to this issue, would be the case in which the two images are rendered very different by the presence of gas, causing less return photons in the on-absorption image and thus a darker image, or a scene that is relatively featureless or repetitive. In both of these cases this type of correlation will provide less accurate offsets or fail completely. In scenarios such as these the maximum offset acts to dampen the overall impact this can have by preventing the calculation of an offset so large it introduces further movement artefacts.

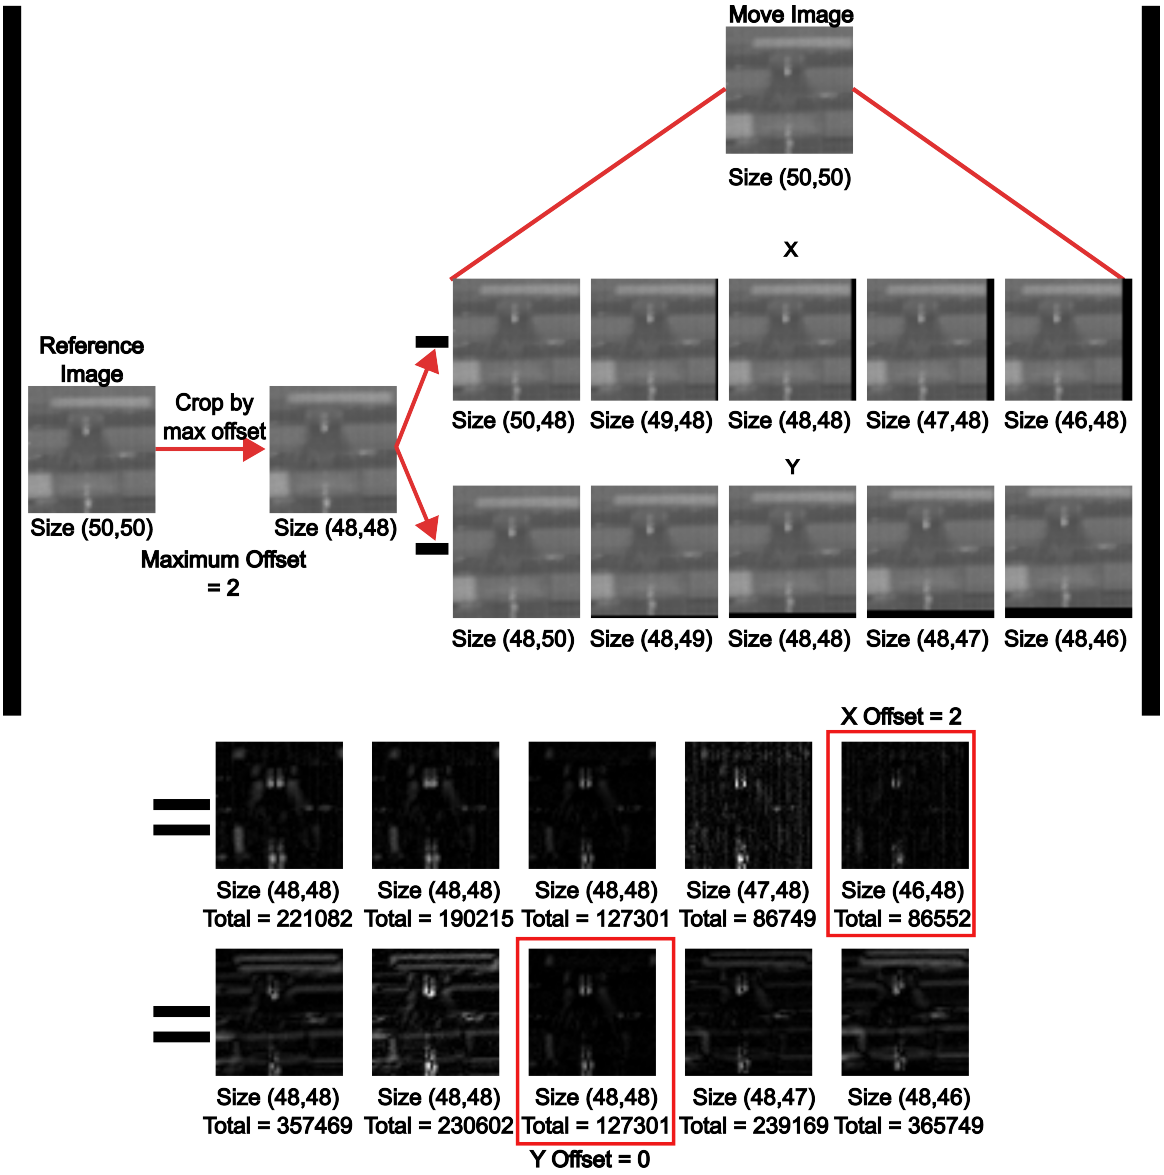

Figure 3.5: **Diagram showing the simplified process for calculating the offset between a reference image and a move image.** The reference image is firstly cropped by the size of the maximum offset, which here is 2, allowing both positive and negative offsets to be investigated. The move image is iteratively cropped by up to $(2N+1)$ in each direction, and these are subtracted from the cropped reference image. However, not all permutations are explored here, just $(2N+1)$ different points in the x direction and the y direction. Any size mismatch is handled by LabView subtraction function to crop the larger image to match the size of the smaller image. The magnitude of each differential image is taken and the pixel sum then calculated, and 2 1D arrays can be built up, one for each direction. The minima position relative to the centre of each array is identified as being the offset in x and y respectively, in this case a positive offset in the x direction of 2 pixels, with no detected offset in the y direction.

3.3.2 Simplified Cross Correlation

It was quickly discovered that this subtraction method, even with such a small number of iterations, is not particularly much faster than phase correlation. Instead a new, simplified strategy was devised. This simplified process is highlighted in Figure 3.5, highlighting its difference and relative simplicity when compared to the process outlined in Figure 3.4. If one moves the image only either horizontally and vertically, and looks for the respective minima in each direction, then a relatively accurate offset can be quickly found on a timescale of approximately 10 ms for a maximum offset of 10 pixels, greatly reduced from the 50 ms and above required for the full area approach. This means that the processing downtime can be brought down by this approach, increasing the overall frame rate of the differential image generation. The system is however still limited by the pre-selected offset range meaning that the system would only be able to move roughly at a maximum of one pixel per millisecond. Instead of checking each individual potential offset the system now checks for the x offset and y offset separately. This method, instead of scaling with $(2N+1)^2$, now scales with $4N+2$ which means that the total number of operations is less and one can investigate a wider region for a similar length of computing time. This can be thought of as a simplified correlation technique, which is much less time consuming than the more complete correlation described previously. As an example, for a maximum offset of 3 the simplified method only investigates 14 points, as opposed to the 49 points before. The x and y offset calculation can also be performed in two separate for-loops meaning they can be run in parallel further improving the overall efficiency of the system.

In order to compare the efficiency of the two correlation methods a test was conducted. SWIR images were saved whilst the system is moved on its tripod first to the right ($\rightarrow$), then upwards ($\uparrow$) and then approximately back to the origin via a diagonal path ($\swarrow$). The timing of the simplified correlation technique compared to the more full technique and using a maximum offset of 6, performed on this data set can be seen in Figure 3.6. Additionally, Figure 3.7 shows the two techniques retain a similar level of accuracy, as there is generally a high level of agreement between the calculated offsets, bar a minor discrepancy in the x direction. A further timing comparison between the computation time of phase correlation and this simplified cross correlation can be seen in Figure 3.8, where phase correlation is demonstrated to take longer

per calculation than the simplified correlation with a maximum offset of 20. It is clear that this methodology is much faster than phase correlation overall.

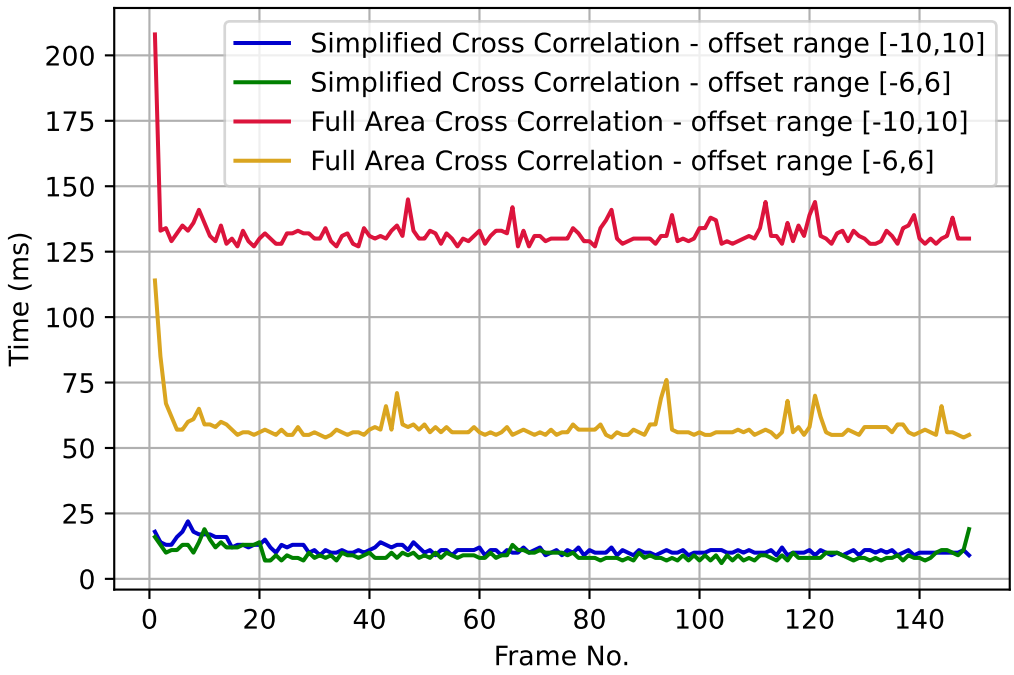

Figure 3.6: **Comparison of time needed to complete full correlation and simplified correlation.** The first frame was recorded to take $\approx 6 \times 10^6$ ms, which is a result of a LabView timing issue and so these timings were discarded, it can additionally be observed in the sluggishness of the initial few frames that were not cropped. The full area correlation refers to the technique in which the images are scanned over the reference image in an area defined as $(2N+1)^2$, for N the maximum offset, whereas the simplified technique only investigates the horizontal and then vertical offsets with only $2 \times (2N+1)$ points investigated.

Having concluded that there is only minor disagreement between the data obtained for the simplified cross correlation, it could be compared to the performance of phase correlation. The offsets for the 150 saved frames are plotted in Figure 3.9. From this graph it can be seen that, in general there is a level of agreement between the two methods, at least within one pixel offset, which can be filtered out with further processing if incorrect. In the y offset data it can be observed that it appears to detect more movement than the phase correlation does, this was an issue experienced when using the phase correlation in general, it often did not detect movement. A few possible reasons for this lack of sensitivity can include the low return signal to the SWIR camera, or an incorrect focus resulting in a lack of features to correlate, the simplified correlation

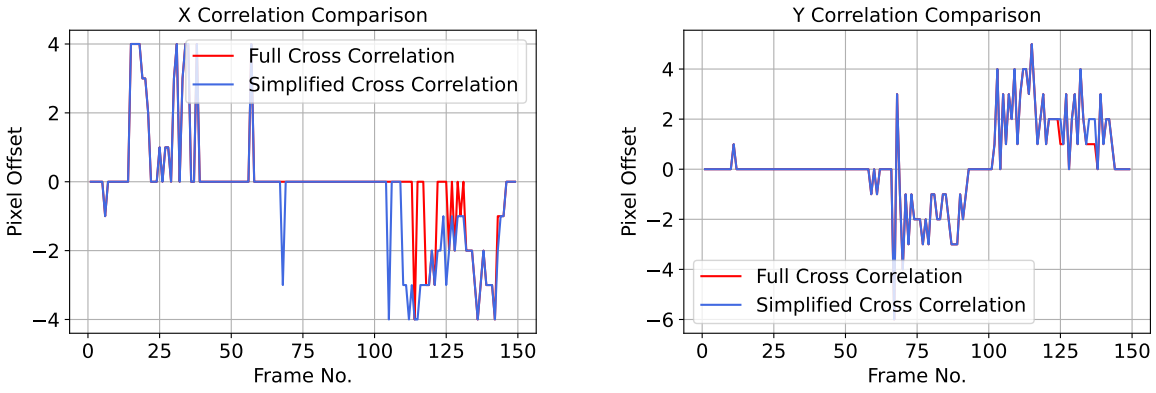

(a) (b)

Figure 3.7: **Comparison of the offsets calculated using the full correlation and the simplified correlation for (a) the x direction and (b) the y direction.** It can be seen that to within a few pixels, the faster simplified procedure of moving the image only horizontally and then vertically matches very closely to the more complete system where the image is more fully scanned across. Any inaccuracies in offset of this size can be filtered out by other imaging techniques later in the post processing. A maximum offset of 6 is implemented here and throughout.

appears to be able to overcome these issues.

3.3.3 Visible or SWIR channel

When correlating the question was asked as to which feed would be best to perform the correlation on; the visible images or the SWIR images? The images being overlapped are the SWIR images so it makes intuitive sense to try and use the SWIR data in order to quantify movement. Additionally, the SWIR images are mostly always guaranteed some level of visibility due to the active SWIR illumination, this provides diminishing returns for further away sources, but in low light situations, such as at night, the visible camera cannot correlate without additional lighting. The visible images however, have a much larger FOV providing more information to correlate, while the pixel sizes may not be exact matches, there is already information in the program that maps the SWIR pixels to the visible pixels so these values can easily be converted in case of difference. Another issue with the visible channel is that the edges of the visible images are aperture stopped and may have to be cropped, otherwise the constant nature of the vignetting could override the correlation calculation, this is shown in Figure 3.10. Correlating the visible camera is only an effective solution in the scenario wherein the visible images is taken concurrently with the SWIR images. This was effective in the previous imaging timeline where the

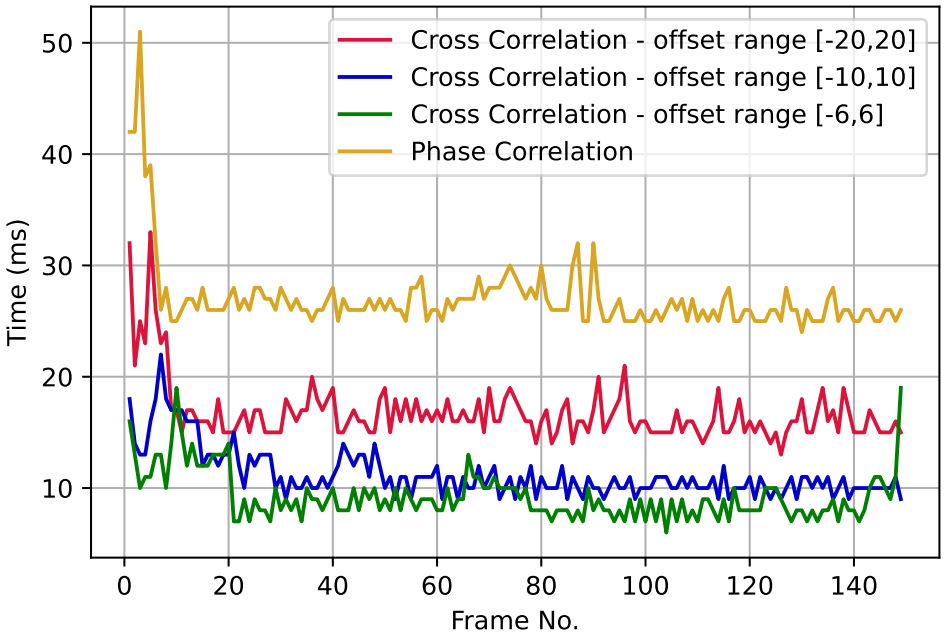

Figure 3.8: **Comparison of timings for different correlation techniques.** Phase correlation is performed on the whole image, and the cross correlation is performed within the defined pixel offset ranges. There is an initial sluggishness in starting, which is a common issue in timing processes in LabView. The first frame was recorded to take $\approx 6 \times 10^6$ ms, which is an additional LabView timing issue and so these times were discarded. It can clearly be observed that the simplified cross correlation technique provides a quicker alternative to phase correlation for reasonably sized maximum offsets.

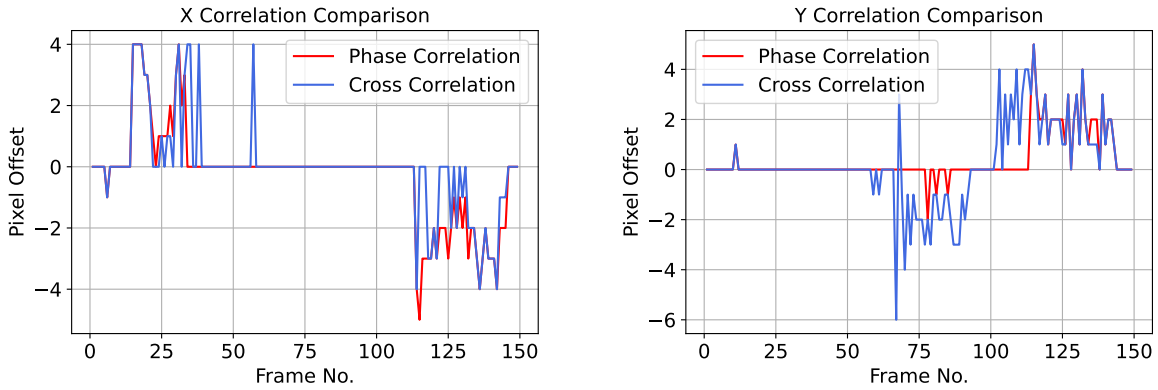

(a) (b)

Figure 3.9: **Comparison between phase correlation and the simplified cross correlation for (a) the x direction and (b) the y direction.** The movement of the system was initially to the right, followed by an upwards movement, before returning to the origin via a diagonal movement down and to the left.

visible camera and SWIR camera were set to have the same frame rate and ran continuously, however in introducing the timeline where two SWIR images are taken in quick succession this implements a need to take an additional visible image per loop, which has the potential to reduce the overall frame rate of the system. Performing the simplified correlation technique to the same data set as before, but this time including the visible frames provides the data shown in Figure 3.11.

Figure 3.10: **Example frame showing the vignetting caused by the outer aperture of the system.**

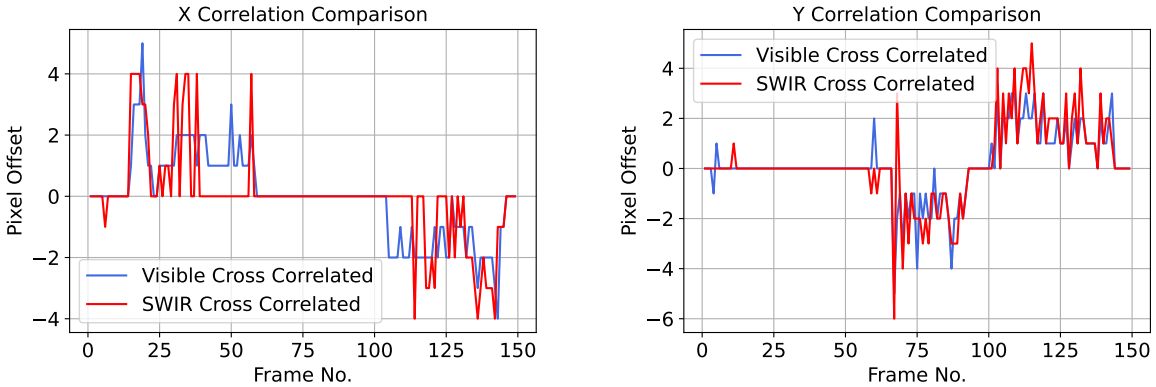

(a) (b)

Figure 3.11: **Comparison between cross correlation performed on the visible and SWIR feeds of the camera for (a) the x direction and (b) the y direction.** The movement of the system was initially to the right, followed by an upwards movement, before returning to the origin via a diagonal movement down and to the left. Of note here is that for completeness the visible image is not cropped and the vignetting is still present.

In general there is quite a high level of agreement between correlations performed on the two channels, despite their different FOVs and pixel numbers. Whilst in some imaging regimes it could have been desirable to utilise the visible feed to calculate offsets, it was instead decided that in the new imaging timeline there is not room to additionally take two visible images each iteration as this would slow down the system. As such the visible channel is almost entirely used for the live feed.

3.4 Persistency

In the initial GasSight program there existed a set of commands that created a binary mask of the gas signal after the step in which the negative values are removed. Non-zero regions, which should ideally represent the presence of gas, are assigned a value of one regardless of magnitude, and regions of zero signal are kept as such. This binary mask is then multiplied by the binary mask generated on the previous iteration of the loop, before being applied to the original image. The purpose of these steps was to prevent a common issue that stems from the fact that in the original version of the GasSight program the order in which the images are subtracted temporally, swaps with each loop. For example if image A which is taken off absorption, is taken before image B which is taken on absorption then, when subtracting A-B to reveal the gas signal, the order of subtraction is such that the temporally later image is taken from the temporally earlier image. After this subtraction image B is then saved into the position A and a new off absorption image is taken and saved in position B. Now, in order to generate a gas signal B-A must be computed, which temporally is backwards to the previous loop iteration where, in this case, the earlier image is taken from the later one. One consequence of this swapping of temporal direction is that when moving the camera the false signal that arises at the edges of the frame ends up swapping sides each and every frame. For a sharp movement to the left for example, large bars of signal are present on either side of the frame. By performing the masking step outlined above, there is effectively a limit in place that says that a gas signal must be present in two consecutive frames in order to count as a real signal. This eradicates these edge signals because by definition of their origin, they cannot exists in multiple frames in a row,

as was demonstrated previously in Figure 2.4.

When examining methods for how to suppress further movement noise this small section of code was re-examined. It was thought that, as this code removes things that are not persistent between two consecutive frames, it could it be harnessed in order to filter out movement artefacts by extending its effect beyond two consecutive frames. The code was changed to include a variable number of consecutive frames, dictating that a signal must be present within 3 , 4 or 5 frames to pass, for example. The persistency filter can have a destructive impact on actual differential gas signal, especially in a moving frame, or whilst the gas is moving, as it cannot differentiate between noise and signal. A persistency strength of 5 (i.e. there are six consecutive frames involved, the current one and the previous 5 frames) is used in all future investigations.

Persistency has different effects on a signal, and can affect a moving signal in a stationary frame, a stationary signal in a moving frame, and a moving signal in a moving frame in similar ways. If a gas signal is moving within the frame then it can result in a loss of sensitivity at the edges of the movement. An example of this can be seen in Figure 3.12, wherein a circular signal is squeezed in the direction of motion of the camera. If the signal is moving slowly then the impact of this is far less than a fast moving signal as there will be an increased level of overlap in consecutive frames. It is entirely possible that, for demanding levels of persistency, a fast moving signal can become invisible to the system. It is for this reason that the user must be careful with how strong the persistency setting is. For example, if viewing a plume it could be such that the more amorphous, transient edges of the plume may be lost to the filter but the more consistent origin point of the plume is more likely to pass the filter. One benefit to this is that movements external to the gas signal moving in the frame are far less likely to be recorded as signals.

Persistency filtering is especially harsh upon signal moving in the frame. This harshness is dependent upon the filter strength but will ultimately impact signal that moves or evolves in a period of time slower than the total time taken to take the impacted images. In the worst cases it can result in the system becoming fully insensitive to signals moving in the frame at certain speeds. The speed at which something must be moving before it is blocked by the persistency filter depends largely upon the frame rate of the imaging camera, a higher frame rate means that

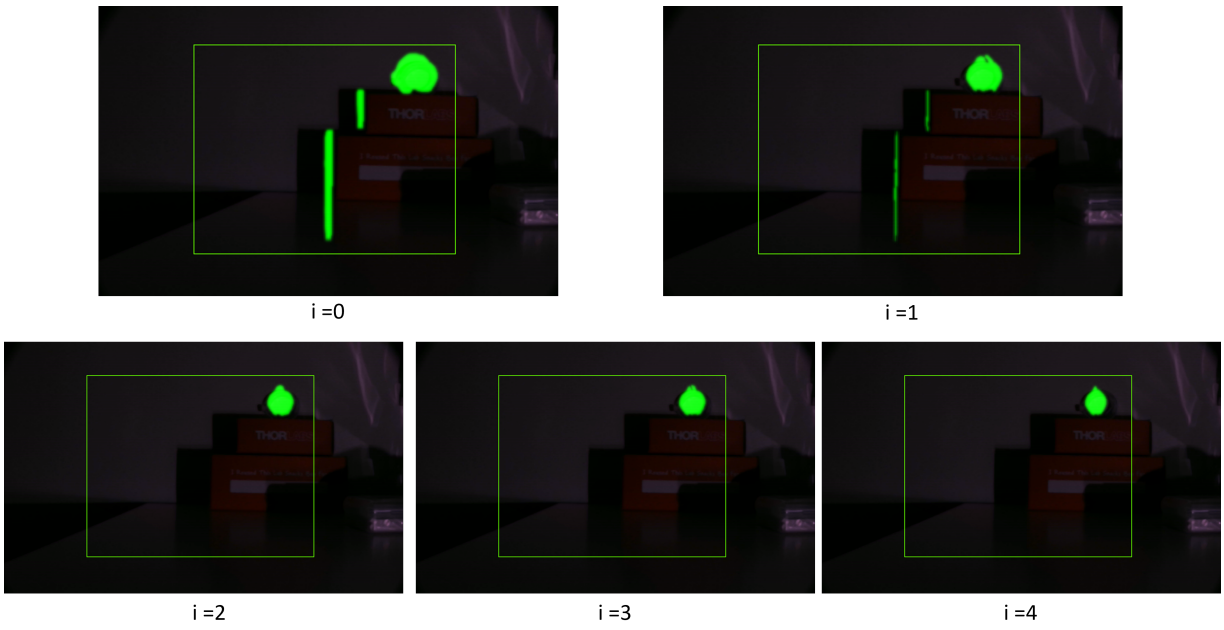

Figure 3.12: **Example of the effect of increasing strength of persistency filter.** Persistency strength increased by increasing the number of previous frames , i, included in the calculation. The first image, where $i = 0$, has no reliance on previous frames and displays a large and noticeable edge effect on the side of the boxes. By steadily increasing i, it is noticed that the edge effects disappear at $i = 2$, however the size of the signal itself continues to decrease in extent along the axis of movement as more frames are taken into account.

more images are taken in the same amount of time so a persistency filter will now apply over a shorter timescale than before. This means that a previously strong filter will become weaker, but for the case of a higher frequency it is also likely that there will be less differences frame to frame as a result of movement, so the impact of the persistency filter will be lessened further. The persistency filter will also be impacted by the number of pixels in the image as well as the field of view as these will dictate the position and size of the moving object in each image and could mean, for example, in a movement captured by a higher resolution camera that there is no overlap of individual pixels meaning there is no persistence. This means that for a higher resolution camera there needs to be a weaker persistency filter.

Overall, whilst this method can be a very effective noise filter, its strength and harsh nature means that it must be used sparingly to avoid signal loss. In conjunction with other less aggressive filtering types, this can slot into the system as a slight "clean up" filter that can get rid of a large amount of random fluctuating noise.

3.5 Erosion & Dilation

Erosion and dilation is a widely used technique in image processing to de-noise and clean up processed images [124]. Erosion acts to decrease the spatial size of a signal, whereas dilation acts to increase the spatial size of a signal. By using both techniques in conjunction one can decrease small signals away with an erode and recover the data lost on the signal by dilating, but because the erode removed the small noise signals they do not return in the dilation. This can be used essentially as an additional clean-up step after noise signals have been filtered and decreased in extent by other means.

The process of eroding and dilating is very flexible and revolves around the definition of a kernel called a structuring element. A structuring element is a binary array made up of 1s and 0s. This structuring element can have almost any shape one desires, and different shapes will provide different effects, such as increasing or decreasing the harshness of the filter or by helping define the expected shapes of signals. In general erosion and dilation is performed upon binary images. The erosion and dilation function used is the inbuilt LabView function works by first defining what type of process one would like to perform on the image, in this case the option "open" is selected. An open is an erosion followed by a dilation, other orders are available, as well as just either an erosion and dilation. A number of iterations can also be added, meaning that many erosion and dilations can be performed in quick succession. Multiple iterations can act to clean up the image further, but for the purposes of this device and to maintain frame rate, the process is only ever performed once. The structuring constant then passes over the binary image, centring on each image pixel. The structuring element must be defined to have a size of odd number $\times$ odd number, this is to ensure that there is a central pixel. This central pixel is the point at which the element is overlaid onto each image pixel, allowing the vicinity of that pixel to be probed. For an erosion action the action of the element is an OR operation wherein if a 1 in the element overlaps a 0 in the binary image, the central pixel is set to zero, if there is no 0 in the nearest neighbours then the central pixel value is unchanged. For a dilation the OR operation acts such that if the element overlaps with a nearest neighbour pixel containing a 1 the central pixel is set to 1. At the edges of the image LabView requires the introduction

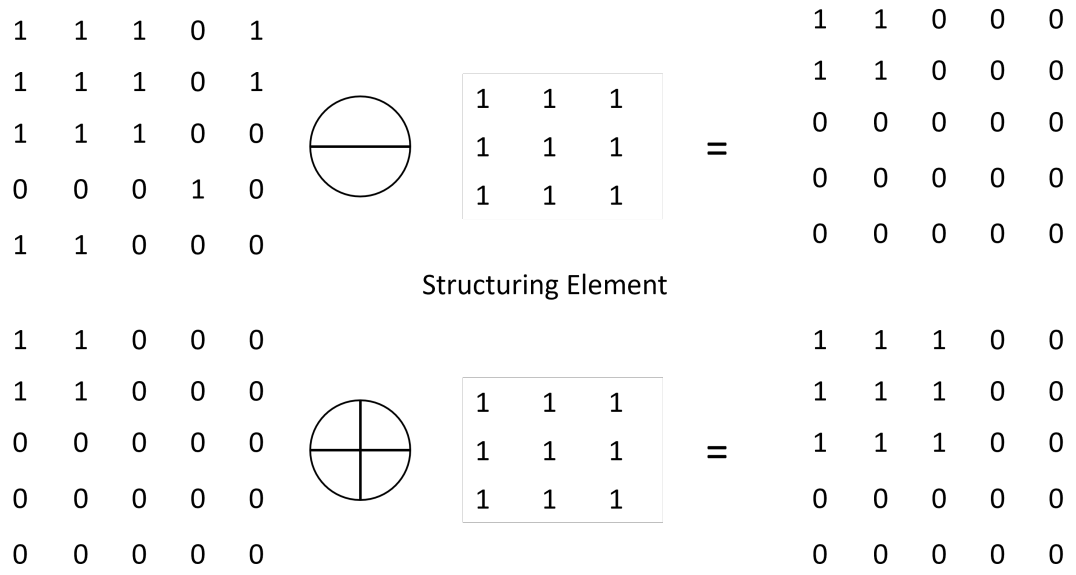

Figure 3.13: **Simplified examples of the erosion and dilation process.** These example uses the same structuring element, which is a 3×3 array of 1s. The top calculation shows an erosion where the structuring element is selected such that the 1s that are not connected to the small cluster in the top left of the image are eroded away. The bottom calculation shows the result of a dilation on the end product of the erosion via the same structuring element. This demonstrates that the 1s that were removed by the erode outside of the cluster will not return upon dilation.

of a temporary border, which by default consists of 1s. A brief and very simple example of an erosion and dilation is presented in Figure 3.13.

In the system, the erosion and dilation is applied to the binary mask created for the signal position after the persistency filter has taken effect. The erosion and dilation occurs after the persistency filter, which already decreases the noise in the image. This means that a weaker structuring element is required, therefore, having minimal impact upon signal shape. The fact that the erosion is applied directly to a binary mask made this easy to implement as the signal has already been converted to binary. In terms of selecting the structuring element this involved some trial and error. The aim is to have minimal impact on the signal itself, so smaller structuring elements are inherently better, but they must also be effective at reducing noise. There was much experimentation with the size and shape of the structuring element throughout, for example using square structuring elements consisting only of 1s, and various fills of circles with 0s in the corner. An example of the effect of structuring element choice for circular elements is presented in Figure 3.14. In general Figure 3.14 highlights the fact that this step is really only used for tidying the image up and has no great effect on the overall gas signal.

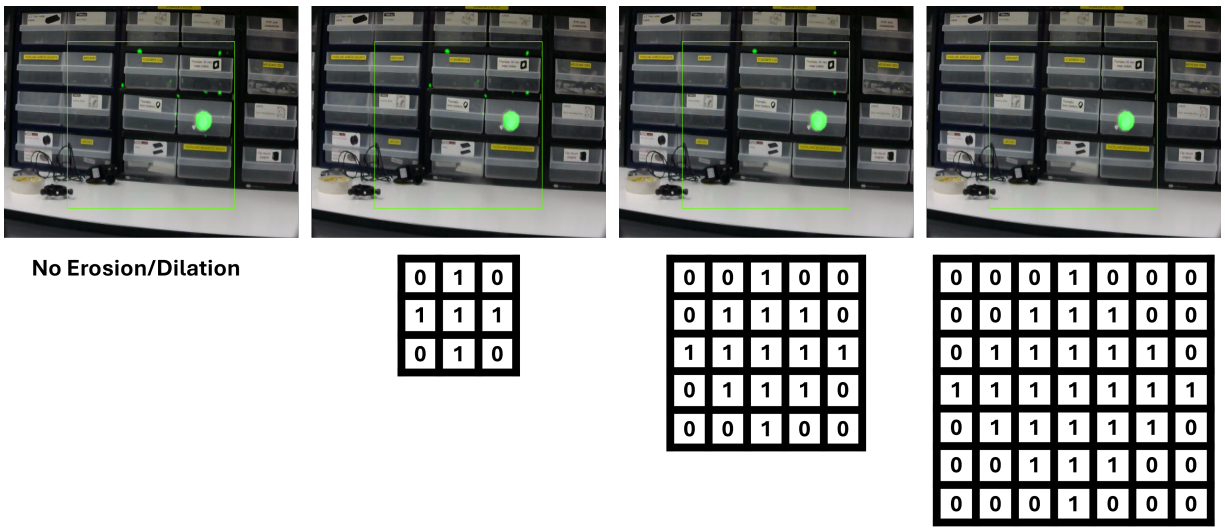

Figure 3.14: **The effect of increasing the size of structuring element upon de-noising and signal strength.** It can be seen that as the extent of the element increases, from 3×3 pixels to 7×7 pixels, the small regions of noise begin to disappear. The SWIR image size that the erosion and dilation operates upon is 320×256 pixels. As the size of the structuring element never exceeds the size of the gas signal, the gas signal remains largely unaltered.

3.6 Temporal Smoothing

The final method to reduce the effect of the motion artefacts is that of temporal smoothing. This is a technique that, in essence behaves similarly to that of persistency, as it effects signals based on when in time they were generated. The way that it functions is that the latest generated signal will make up a defined proportion of the overall signal. The equation describing such a filter would look like the following,

$$S_N = (R * I_N) + ((1 - R) * S_{N-1}) \tag{3.2}$$

where, S_N is the total signal for iteration N, S_{N-1} is the total gas signal calculated through the same equation for iteration $N-1$, I_N is the image signal taken during the current frame, and R is the recursive strength which dictates the percentage of how much the current image signal contributes to the current signal. This means that depending on the recursive strength the information recorded in previous frames could have a heavy weighting upon the latest frame. What this can mean is that small variations present in only the most recent frame can, with the correct recursive strength, essentially be filtered out of the final image. A high recursive filter

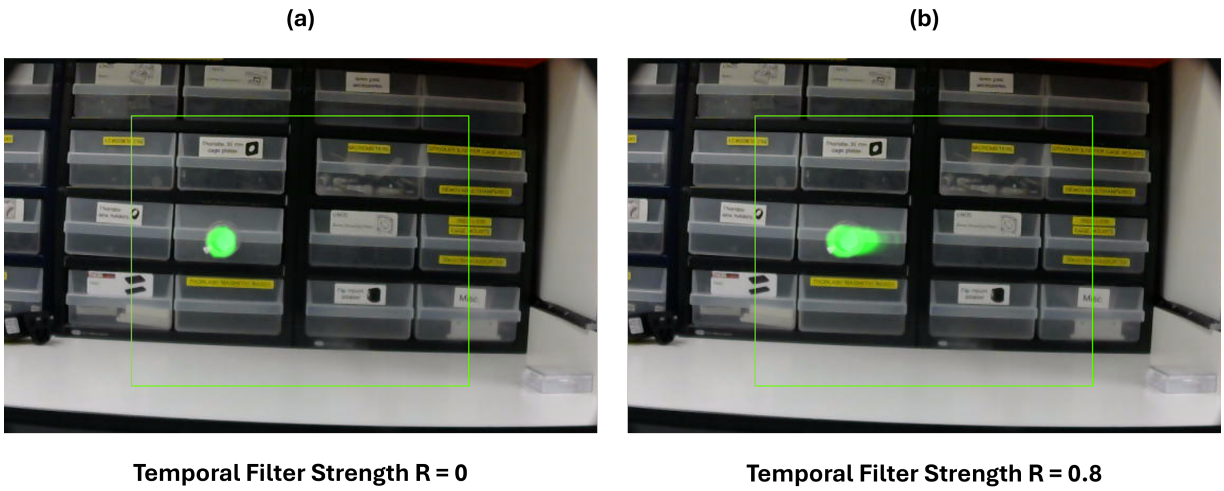

Figure 3.15: **An example of the trailing quality that can be added when using a filter with high recursive strength.**

means that the latest signal is dominant, and in general this means that the signal updates much more quickly but also means that transient false signals are more likely to pass through. A low recursive strength means that the final signal is more consistent of data gathered in the past and this can mean that small false signals such as edge effects can be washed out relative to a gas signal but the slowness to update can result in quickly moving signals to be similarly reduced. By implementing this filter into the system it was hoped it would act to counteract the loss of signal size caused by the persistency filter by allowing previous gas signals to appear in the frame, after being weighted appropriately.

What this often ends up looking like in motion, is essentially a trail that follows a signal, as seen in Figure 3.15 (b), which is not present in (a). The effects of the temporal smoothing can be fairly variable, with the potential to smooth out small examples of noise. In the worst case scenario, there could be a situation in which the false signals are consistent enough between frames that they end up persisting longer than they would without the filter in place. An example may be that the edge effects are so consistent that they adopt the trail also making them even more of a distraction from the real gas signal.

In a faster system, the temporal filter acts to essentially compile and average a memory of all the previous frames. This can allow for a more asynchronous approach to imaging where the differential imaging can be banked into the temporal filter and the SWIR camera can contin-

uously image whilst a slower visible imaging and post processing loop just accesses the latest temporally smoothed image. This cannot currently be implemented in the system due to the elongated tuning times of the diodes limiting the frame rate.

3.7 Conclusion

Overall, there were a number of software avenues of exploration available for improving the systems resilience to movement. This process allowed for a range of experimentation, and it was difficult to truly make the system robust to movement. Of the techniques explored, the image processing timeline changes, introduced in section 3.2, probably had the greatest impact of all. Changing where in the timeline the images are taken permanently decreases the downtime between consecutive images, thus automatically decreasing the extent of the differences, caused by movement, between the two images and decreasing the size of any artefacts caused by motion. With this change, it placed weaker constraints on the required strength of subsequent de-noising techniques, meaning that these weaker filters would also have a much lesser impact on the actual gas signal reducing the potential for a loss of sensitivity.

Image correlation, as discussed in Section 3.3, was the technique that was explored the most, and was the primary technique used to account for movement in the system before the start of this PhD research. Many techniques were tried with the primary aim of maintaining the frame rate by avoiding computationally slow processes. This led to the introduction of the quick, but less accurate, cross correlation approach that replaces the slightly more accurate, but slower, phase correlation. Performing correlation upon the visible channel was also investigated but dismissed due to the additional need to take a second visible frame.

Introducing a harsher persistency filter was also a largely effective step but had to be used carefully. Persistency can be very effective at removing transient noise signals and moving or narrow edge signals but has quite a large destructive potential on non-noise signals. A careful use of this, usually defining a signal present in 6 consecutive frames as being a real signal, was found to be very effective at increasing robustness to movement. The main downside of this approach is that the system is very susceptible to sudden jerky movements and in this case

almost all signal, noise or not, disappears from the frame.

Erosion, dilation and temporal smoothing, discussed in Sections 3.5 and 3.6, were introduced in order to pull back on the dependence on persistency, by allowing these processes to clean up the images with minimal impact on gas signal. Erosion and dilation in particular were used in quite a small capacity in later works, but acts to be a very effective clean up step for reducing the presence of noise, especially as it is implemented after the persistency step, where noise sizes are generally decreased. Temporal smoothing was found to be effective at accounting for the loss of positional data caused by the persistency filter. The temporal filter could introduce its own shortcomings such as the potential to actually enhance noise in the final image. It was decided that temporal smoothing is better suited in a faster, more continuous system.

In conclusion, many techniques were explored in great detail in this chapter, with the resolution being that in order to effectively de-noise the system and allow it to function as a system that can be freely moved and trucked requires a careful balance of each of the techniques. This balance can vary for the scenario that the system is required to be used in, but in general the movement sensitivity is hugely dependent upon the frame rate of the system and the tuning time of the diodes, which is ultimately the fundamental frame rate limit of the system.

Chapter 4

Spectral Filtering for Daytime Operation

4.1 Introduction

The content of this chapter is discussed in the conference proceedings "Drone-based gas leak detection system for use in industry" [98] and "Real-time imaging of methane gas from a UAV mounted system" [96].

Following the improvements made to the motion reduction, another key issue had to be tackled - sunlight. As part of the trialling of the system, an observation of a gas cell was set up outside on a bright, sunny day. It was quickly identified that despite a relatively short range the system could not detect the gas signature within the cell. The image was also noted to be significantly noisier than when investigated inside. By comparing this to the performance of the system during the night it was clear that sunlight was causing the issue. If the system can only detect gas leaks in the absence of sunlight then its likelihood of widespread deployment is significantly reduced. An example of the difference in ambient SWIR light levels indoor versus outside is shown in Figure 4.1.

As a result of this significant impasse, the difficulty of overcoming the sunlight induced lack of sensitivity became a large focus for the research, with both hardware and software changes explored. The method used to attempt to remove this ambient sunlight background was to filter the SWIR camera. This chapter discusses these filters and introduces the associated benefits and detriments of using them by modelling and implementing them. A technique to overcome the

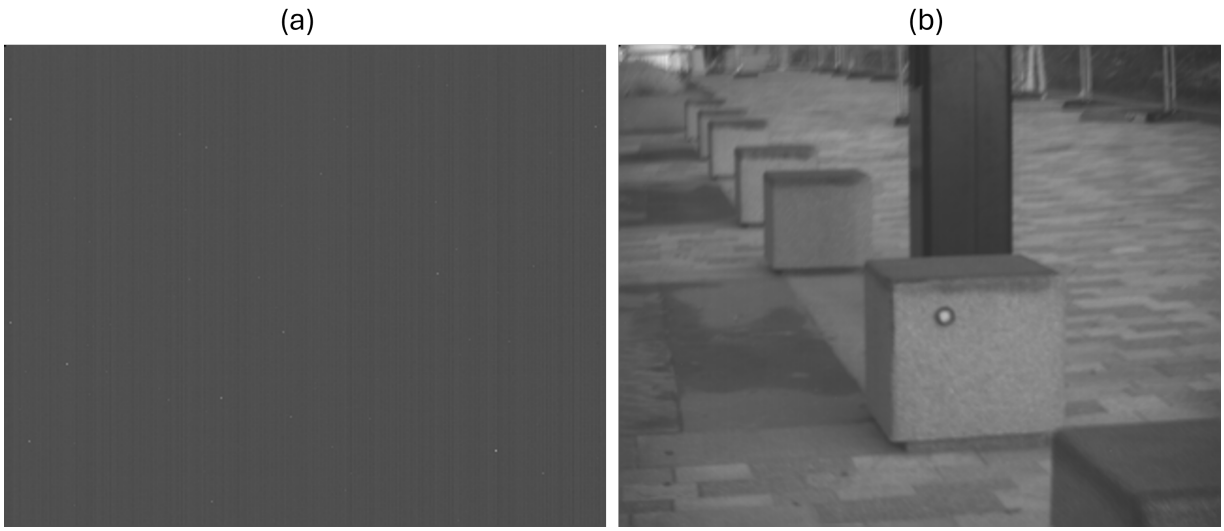

Figure 4.1: **Example frame taken on the SWIR camera with the same settings inside and outside in the absence of SWIR diode illumination.** The SWIR camera features a spectral bandpass filter centred on 1653 nm with a FWHM of 18.8 nm. It can be seen that without a source of SWIR light, the inside frame (a) only contains sensor noise. On the other hand the SWIR camera can clearly see the full frame outside in (b). This contrast reveals that a high level of sunlight can enter the SWIR camera via the filter.

loss of information introduced as a result of adding filters is also introduced and tested.

4.2 Background

In order to explore how best to filter sunlight, an understanding of how filter choice, sunlight and temperature can each affect the performance of the system is introduced.

4.2.1 Filtering

The primary issue that affects the system is that the desired wavelength (1.65 μm) exists in an atmospheric window, meaning that there will be a high background level of light due to sunlight [125–127]. The Goldeye SWIR FPA used has a peak quantum efficiency from around 1 μm to 1.65 μm with a sharp fall-off at longer and shorter wavelengths, this means that the wavelength used in the system exists on the down trend of the systems efficiency, providing an approximate quantum efficiency of 65% and means that any choice that decreases the amount of light reaching the sensor must be carefully considered. The quantum efficiency of the system is presented in Figure 2.2. It is additionally important to note that the well capacity of this camera

is 5 Me^{-1} at the Gain 0 setting used throughout this research. Additionally, the dynamic range is reported to be 70 dB, in the Gain 0 setting. The exposure time of the system varies for many of these experiments but is generally set to be 20 ms, unless otherwise stated.

The most effective way to reduce sunlight entering the system is to filter the camera so that only light at 1.65 μm enters the camera. The filtering method used in this system is a dielectric bandpass filter, which aims to pass only desired wavelengths of light through. Dielectric bandpass filters work via interference and consist of a series of dielectric materials of different refractive index deposited on top of each other [128]. The aim of these layers is to reflect unwanted wavelengths of light such that they destructively interfere, whilst passing desired wavelengths with minimal loss. Varying the design of the layers can alter the filter response. For example by changing layer thickness, and thus path length, one can alter the transmittance, similarly changing the number of dielectric layers impacts the range of allowed wavelengths the filter will pass [129]. Different types of filters such as long pass, short pass and bandpass filters with differing transmission widths can all be engineered by altering the structure and quantity of the dielectric coatings. All filters used throughout this chapter are dielectric bandpass filters.

4.2.2 Sunlight

As evident from Figure 4.1, unwanted sunlight signal is entering the system via the filter as a result of the atmospheric window coinciding with the imaging wavelength. Strong summer sunlight levels are enough to saturate the SWIR camera. In some cases this saturation can cause blooming where saturated pixels can overflow into neighbouring pixels, though it is not evident that this is happening in this system [130]. There will be environmental changes that vary the intensity of the background sunlight, such as cloud cover and atmospheric turbulence [131,132]. For cloud cover, the type of cloud can also dramatically impact the level of SWIR blocking especially relative to visible light, for example cumulus clouds have a stronger blocking effect than cirrus clouds [133]. The height of the sun in the sky will also impact the intensity of sunlight, as the sun approaches the horizon, the intensity decreases [133].

4.2.3 Temperature

Varying temperature can also have an impact on the systems performance. Temperature can alter the absorption profile of a gas, for example the methane absorption spectrum can significantly vary in shape and absorption strength, with absorption lines absorbing more strongly or weakly with temperature. As the strength of absorption is temperature dependent, this must be considered when detecting methane over an extended temperature range [134, 135]. It is reported that over the range of temperatures from 300 K to 1000 K, there would be an approximately 16× decrease in device sensitivity as a result of the absorption strengths temperature dependence [136]. Whilst this decrease is obviously a large change, it is not representative of actual working parameters for the device. Over a range of 290 K to 350 K one would experience a 21% decrease in line intensity [135]. More typically the system would work around 300 K, slightly lower if working at night as often necessitated by other sunlight related issues. These observations follow on from previous discussions of Doppler an pressure broadening, where higher temperatures act to increase the Doppler broadening effect and decrease the relative absorption experienced. Plots showing the effect of temperature changes on absorption are included in Figure 4.2. From this figure it can be observed clearly that as temperature increases, absorbance decreases. Additionally, included for completeness in Figure 4.2 is the effect of changing pressure as in a sealed gas cell the temperature will impact the overall pressure. It can be observed that for the values provided relative to the temperature changes the pressure broadening effect is much greater suggesting that for the situations observed here, pressure broadening is the dominant effect.

The uncontrollable issues, outlined above, make trying to account for sunlight, particularly in how its variance impacts the captured images, very difficult. This difficulty means that the approach taken was instead to decrease the amount of sunlight detected as opposed to trying to overcome it at its current level.

4.3 Saturation and Variation

Upon reviewing data obtained during sunny days, where the gas in the cell could not be detected, it was clear that two issues caused a reduction in sensitivity. Firstly, the frames were significantly

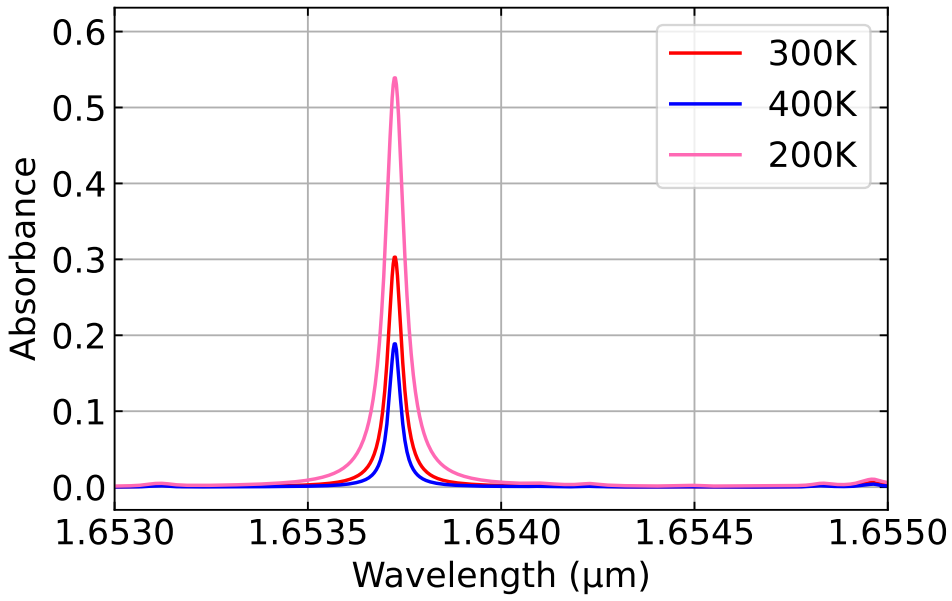

(a)

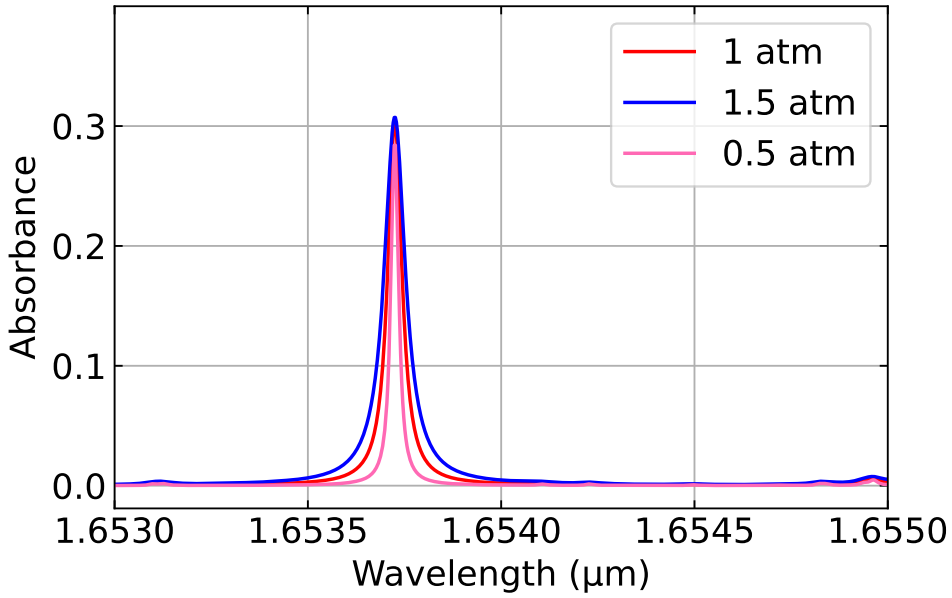

(b)

Figure 4.2: **[Comparison of the effect of increasing temperature (top) and pressure (bottom) upon absorbance of a gas species.** In the top image the gas is methane gas with a mole fraction of 1 at and a pressure of 1 atmosphere and varying temperature. The bottom image the gas is methane gas with a mole fraction of 0.01 at a temperature of 300 K and varying pressure. Plots are generated using data from the HITRAN database [36] obtained via the SpectraPlot toolset available at `https://www.spectraplot.com/absorption` [37].

saturated, despite the fact the system features a bandpass filter with a full width at half maximum (FWHM) of 18.8 nm centred upon 1653 nm, a plot of the filter response is included in Figure 4.3. It was clear that this filter was not capable of filtering enough of the sunlight.

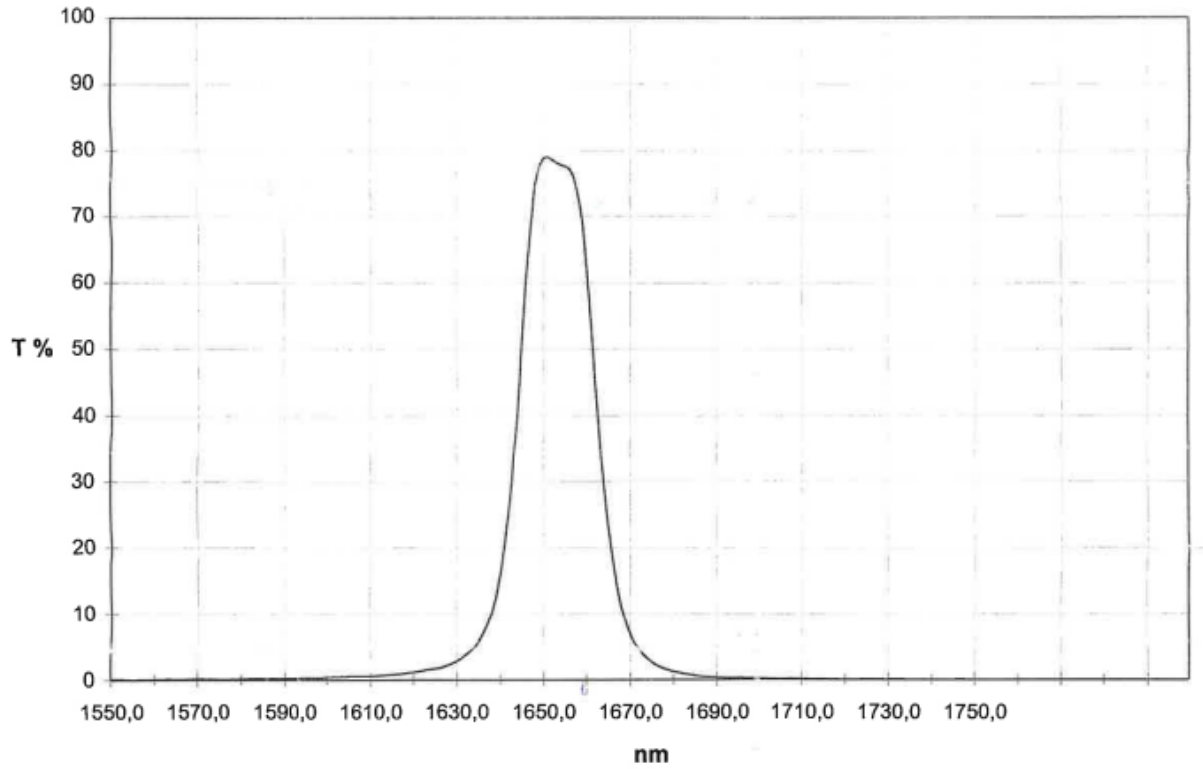

Figure 4.3: **A scan of the specification sheet for the initial filter used in the system.** The filter has a FWHM of 18.8 nm, centred upon 1653 nm. The maximum transmission of the system is 79%. Specification sheet and filter created by Spectrogon, specification sheet reproduced here with permission.

Secondly, the high levels of noise signal in the captured frame appear to be a result of large variations in sunlight level and brightness over short timescales. Much like the false motion detection signals reported in Chapter 3, this noise is a result in differences between the two subsequent frames that are not related to the presence of gas. These variations are hard to predict and similarly hard to reduce for the system at hand, for reasons similar to that of the motion suppression. A higher frame rate system will lead to a more efficient system that is more resilient to these fluctuations. Due to the quick and seemingly random nature of the noise, the previously very effective persistency filter does not aid much in reducing their effect.

4.4 Filters

Filters are defined primarily by two main features, the central wavelength, λ_0, and the width of the filter, usually quantified by its FWHM. In terms of filtering the system the most simple

solution to detecting too much ambient light is would be to use a filter with a narrower FWHM. Having a narrower FWHM centred on the desired 1653.7 nm would mean less light entering the system at undesirable wavelengths, whilst maintaining the same level of gas signature. Using a narrower FWHM filter means that less sunlight will reach the sensor without the need to alter the exposure time of the SWIR camera. The primary drawbacks of introducing reduced FWHM filters are their increased manufacturing cost and FOV reduction caused by the reduced angular acceptance of the filter. The increased cost comes about from the requirement to make custom optics, that require expensive coatings. Narrow filters also come with the risk of not being effective at the desired wavelength, this is a result of the manufacturing process wherein an exact central wavelength and FWHM are rarely exactly attained. For example for the filter seen in Figure 4.3, there is an uncertainty on the specifications sheet of ± 2 nm in central wavelength and ± 3 nm in FWHM. If a filter is narrow to the extent that its FWHM is of comparable size to the uncertainty on the FWHM then it follows that in some cases this can actually result in a filter that is either wider than anticipated, or more crucially not centred on the desired wavelength. This offset in centring can mean that the expected transmission for the desired wavelength is actually much lower than desired or, in the worst cases, blocked by the filter. In these scenarios the filter cannot be used as effectively as a wider filter. The high cost and long development time of the films and optics required for custom filters means that experimenting with this can be a costly and time consuming effort, which can lead to little or no benefit in the long run.

4.5 Irradiance Approximation

By following a similar methodology to that in Section 2.2.3, one can compare the amount of sunlight passing through the filter and compare it to the amount of diode light. To do this a value of the solar irradiance, I, at 1653 nm must be found, this is approximately $0.23\,\mathrm{W\,m^{-2}\,nm^{-1}}$, and an extract of the blackbody irradiance response is included in Figure 4.4 [137]. By using this value of solar irradiance one can compare the power density of sunlight, ρ_s to that of the diodes at a specified distance, where the diode power density, ρ, is calculated again using Equation 2.14. To begin, one must calculate the power density of the sunlight, this can be done by multiplying

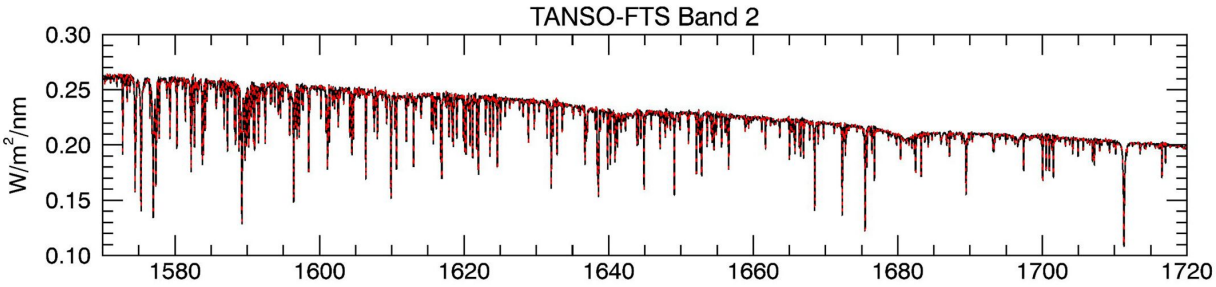

Figure 4.4: **Figure showing the solar irradiance in the wavelength region of interest.** This figure is adapted under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from the paper "The TSIS-1 Hybrid Solar Reference Spectrum" [137]. The figure has been cropped to only include the plots relevant to the wavelength of interest.

the solar irradiance by the width of the filter in nm, Δf:

$$\rho_s = I\Delta f \tag{4.1}$$

For the filter used here, Δf is 18.8 nm, giving a power density of approximately $4.32\,\mathrm{W\,m^{-2}}$. Now, calculating the diode power density using Equation 2.14, at a distance of 2 m, using two 10 mW diodes gives a power density of $0.04\,\mathrm{W\,m^{-2}}$. This means that at a range of 2 m, there is approximately $107.6\times$ more sunlight than diode light. A graphs highlighting the relationship between the diode to sunlight ratio as function of distance is highlighted in Figure 4.5. At the range of 2 m, the $107.6\times(\approx 100\times)$ ratio of sunlight to diode power becomes an issue because it likely saturates the camera. The camera has a dynamic range of 70 dB but is also a 14-bit camera and so can only register a maximum pixel value of 16383, so at a range of 2 m the camera registers an approximate average of 2500 pixel counts per per pixel from the diodes in the absence of sunlight through the filter. If sunlight were to enter this scene the registered pixel value would add roughly 250,000 extra pixel counts per pixel which would saturate the camera easily. As such, a narrower filter would linearly decrease the amount of sunlight entering the system to a more manageable level. What is really challenging here is the variance of the sunlight from frame to frame which is very difficult to quantify as it depends on a wide variety of uncontrollable factors. Instead the focus should be on reducing the sunlight such that it does not saturate the detector and allows the differential gas signal to be seen

To decrease the filter width such that there is a 1:1 ratio of sunlight to diode light at 2 m,

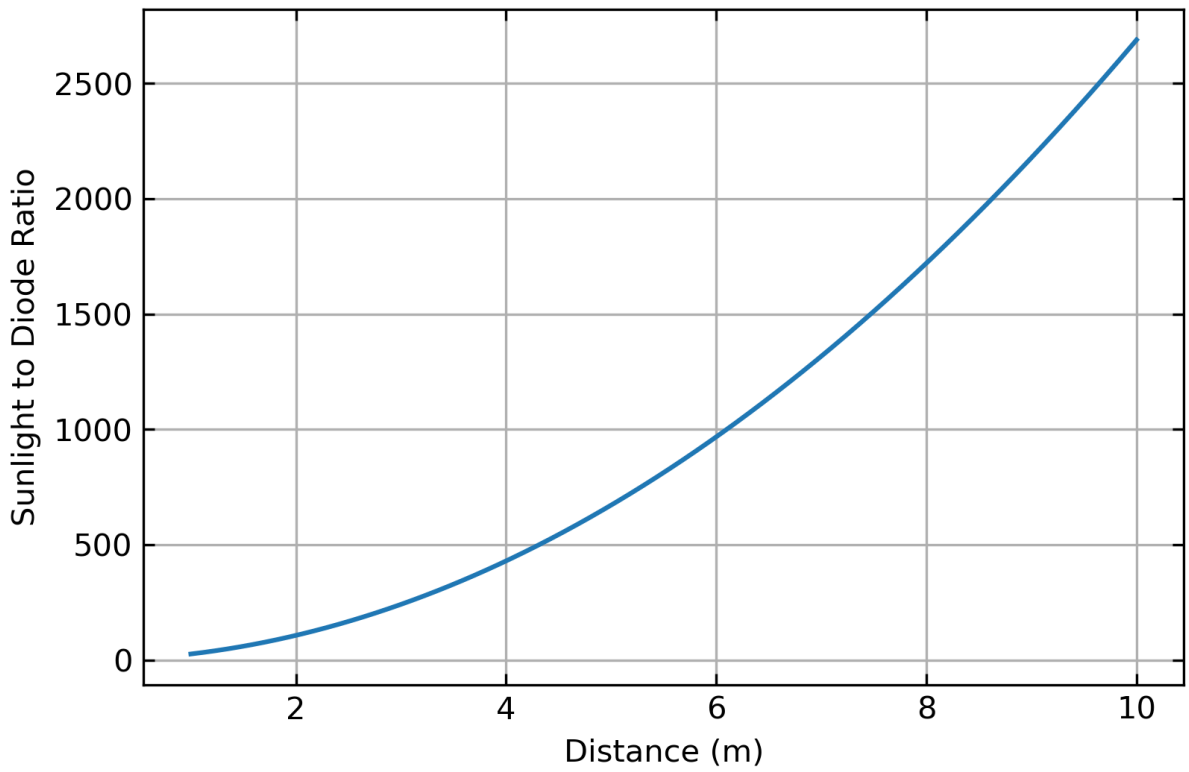

Figure 4.5: **Figure showing the ratio of the power density of sunlight to the diodes as a function of distance.** Here the diodes used are $2\times$ 10 mW diodes.

one would need a filter width of 0.188 nm, which would equate via Equation 2.5 to a width of 20.7 GHz, meaning it would still be wide enough to contain the diode tuning shifts as calculated in Chapter 2.2.1, however assuming a Gaussian response, the transmission of the off wavelength may be significantly reduced, if instead a 5:1 ratio was explored, looking for a filter width of around 0.94 nm $\approx$ 1 nm $\approx$ 110 GHz should work just as well without such a stark contrast in on and off transmission..

4.6 Filter Models

Models of potential ideal filters were made by assuming the filter response to be approximately Gaussian in nature. Model filters with differing FWHM, transmission amplitude, refractive index and central wavelength are all investigated in terms of response to light at a desired wavelength, for a wide array of incident angles. To simulate the angular dependence of transmission $T(\theta)$ as a Gaussian response one must first understand how the incident angle of incoming light

θ impacts the wavelength passing through the filter, $\lambda(\theta)$,

$$\lambda(\theta) = \lambda_0 \sqrt{1 - \left(\frac{\sin\theta}{n_{eff}}\right)^2}, \tag{4.2}$$

where λ_0 is the central wavelength of the filter, and n_{eff} is the effective refractive index of the filter as a whole. The transmission response for the filter as a function of θ is as follows,

$$T(\theta) = T_0 \exp\left(-\frac{(\lambda(\theta) - \lambda_{diode})^2}{2\sigma^2}\right), \tag{4.3}$$

where T_0 is the maximum transmission of the response, λ_{diode} is the illumination wavelength and σ is the standard deviation of the Gaussian response, which is related to the FWHM of the filter via,

$$\sigma = \frac{FWHM}{\sqrt{8\ln 2}}. \tag{4.4}$$

By using these three key equations and varying the parameters, a range of possible filter responses could be modelled.

The first parameter that can varied is the refractive index of the filter. Refractive index is often less considered as it is mostly defined in the manufacturing process and cannot be easily changed. Most filters take the form of glass, and as such the refractive index will not vary much from this. The impact of the refractive index is such that the angular response of the filter as a function of angle varies according to Equation 4.2, which in turn impacts the angular transmittance according to Equation 4.3. A lower refractive index results in a reduced FOV. A heightened version of the effect of this is presented in Figure 4.6.

The second parameter to change is the FWHM of the filter. The main issue outside of uncertainties that come with custom filter manufacturing is that the current detection protocol of the gas detection camera is very susceptible to reductions in FWHM. As the system collects back scattered light from an active illumination source, it thus follows that the light entering the filter will enter from a variety of angles. This wide variety of entry angles means that, in terms of Equation 4.2 light entering at more extreme angles can be shifted outwith the transmitted wavelengths of the filter, this is most notable at the edges of the frame. In reality, the effect this

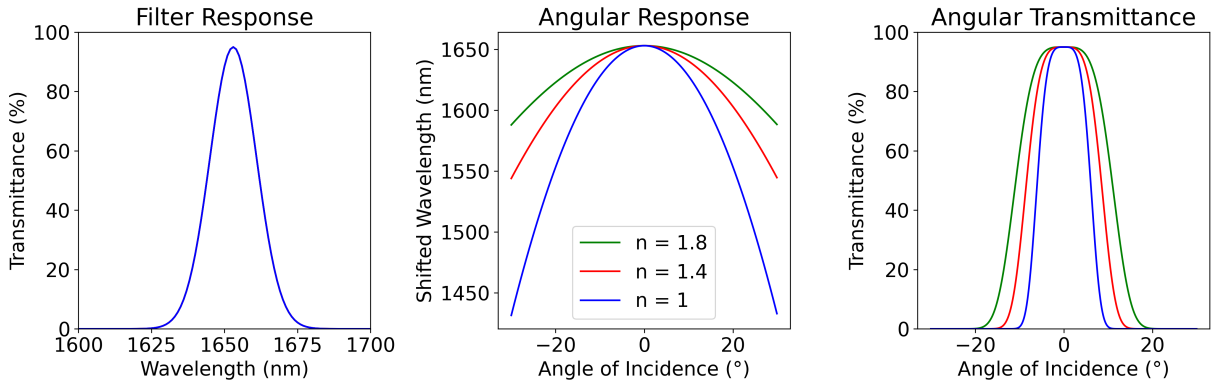

Figure 4.6: **Effect of varying refractive index on filter angular response.** Parameters used for plotting are as follows: a peak transmission of 95%, a FWHM of 18.8 nm, an illumination source at 1653 nm, and a central filter wavelength of 1653 nm. Plots vary refractive index, n. The difference between filters with refractive indexes of 1.8, 1.4 and 1 is presented.

has on an imaging system is that light entering at these extreme angles is blocked by the filter, despite being at the desired wavelength of illumination, and this reduces the amount of light detected overall. In the image this results in a reduction of the FOV of the system. A reduction of FOV means that, in effect, the imaging system is able to see less, reducing the active area of image. An example of a reduction in FOV is shown in Figure 4.7, which compares the resultant FOV of a system with a filter FWHM of 18.8 nm and a system with a filter FWHM of 12 nm. Narrower filters allow for less of an angular acceptance and wider filters allow for more angular acceptance. This reduced angular acceptance clearly demonstrates the dangers of using increasingly narrow filters upon the FOV of an imaging system. The FWHM has no impact upon the angular wavelength shifting response of the filter. Widening the filter past a certain threshold will have less of an effect upon the FOV as opposed to decreasing it. This can clearly be seen in Figure 4.8, where increasing the FWHM from 18.8 nm to 30 nm has a far smaller impact on FOV than decreasing the FWHM to 4 nm.

Another key concern is that, because the system utilises a dual wavelength exposure mechanism, the filter response will differ for the two wavelengths. It thus follows that by making sure these two wavelengths are as similar as possible to prevent differing FOVs. In scenarios in which the FOV changes between the two wavelengths in a differential imaging system then the size difference in illumination profile detected will flag as a differential signal. The shape of the differential signal will appear as a ring, an example of this is presented in Figure 4.9.

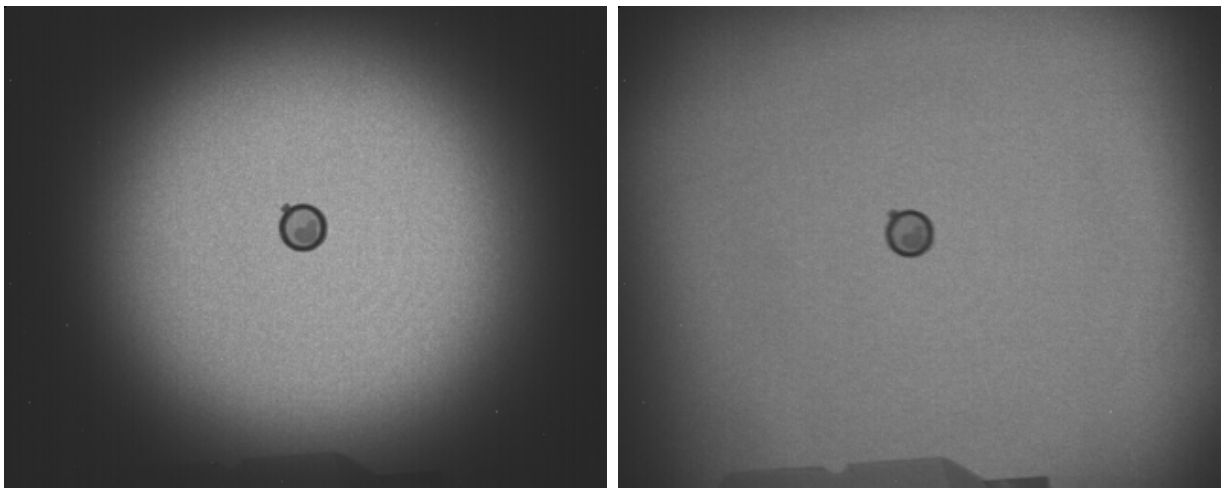

(a) (b)

Figure 4.7: **Comparison of the effects of filter FWHM on FOV, with (a) a filter with FWHM of 12 nm and (b) a filter with FWHM of 18.8 nm.** A scene, consisting primarily of white card and a gas cell of an approximate concentration of 100,000 ppm.m, was illuminated using the internal diodes. The two images were taken from the approximately the same distance with the same wavelength of illumination, only the filter has varied between the two images.

This reduction in FOV brought about by the filter becomes a trade-off between sunlight blocking and active imaging area. A filter with a narrower FWHM will result in a reduced active area on which to detect gas, but could strongly alleviate the impact of sunlight. In order to understand this relationship, models for the filters were developed and experiments conducted with a range of spectral filters to optimise the final detection scheme.

Finally, of interest is the effect of varying the central wavelength of the filter on the angular transmission capabilities. By shifting to a central wavelength to a shorter wavelength to the illumination light, the effect is that there is a decrease in overall transmission with a slight decrease in overall FOV, this can be seen in Figure 4.10. This response is because of the shifting impact of Equation 4.2, where the light entering the system suffers from an overall lower transmission due to the shifting of the transmission curve. It is clear that from this data that using a shorter central wavelength has no benefit to the system, and will actively decrease the FOV and peak transmission of the system. Conversely, if the central wavelength of the filter is instead longer than the illumination wavelength as in Figure 4.10, a very different response is observed. Additionally, with each of these models there is concern that this only looks at the on absorption wavelength for methane, with no consideration of the off wavelength. Based on the estimations made in

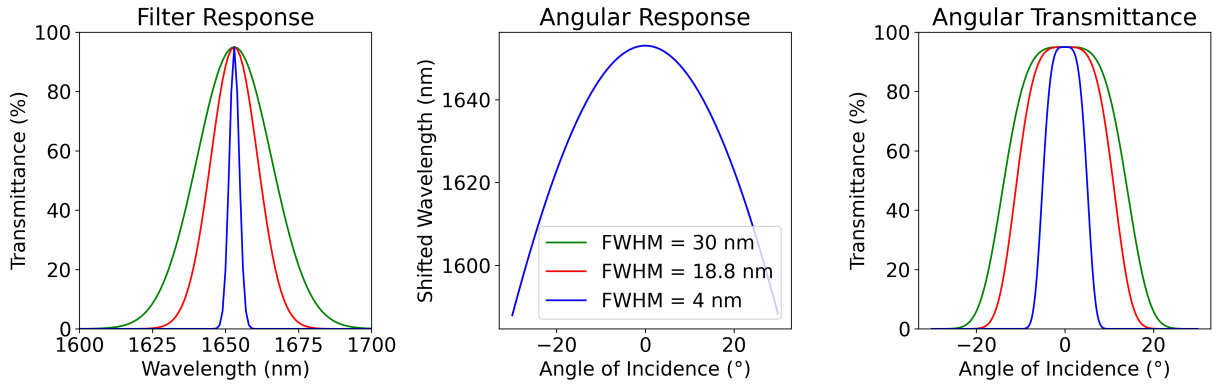

Figure 4.8: **Effect of varying FWHM on filter angular response.** Parameters used for plotting are as follows: a peak transmission of 95%, a refractive index of 1.8, an illumination source at 1653 nm, and a central filter wavelength of 1653 nm. Plots vary FWHM. The difference between filters with a FWHM of 30 nm, 18.8 nm, and 4 nm is presented.

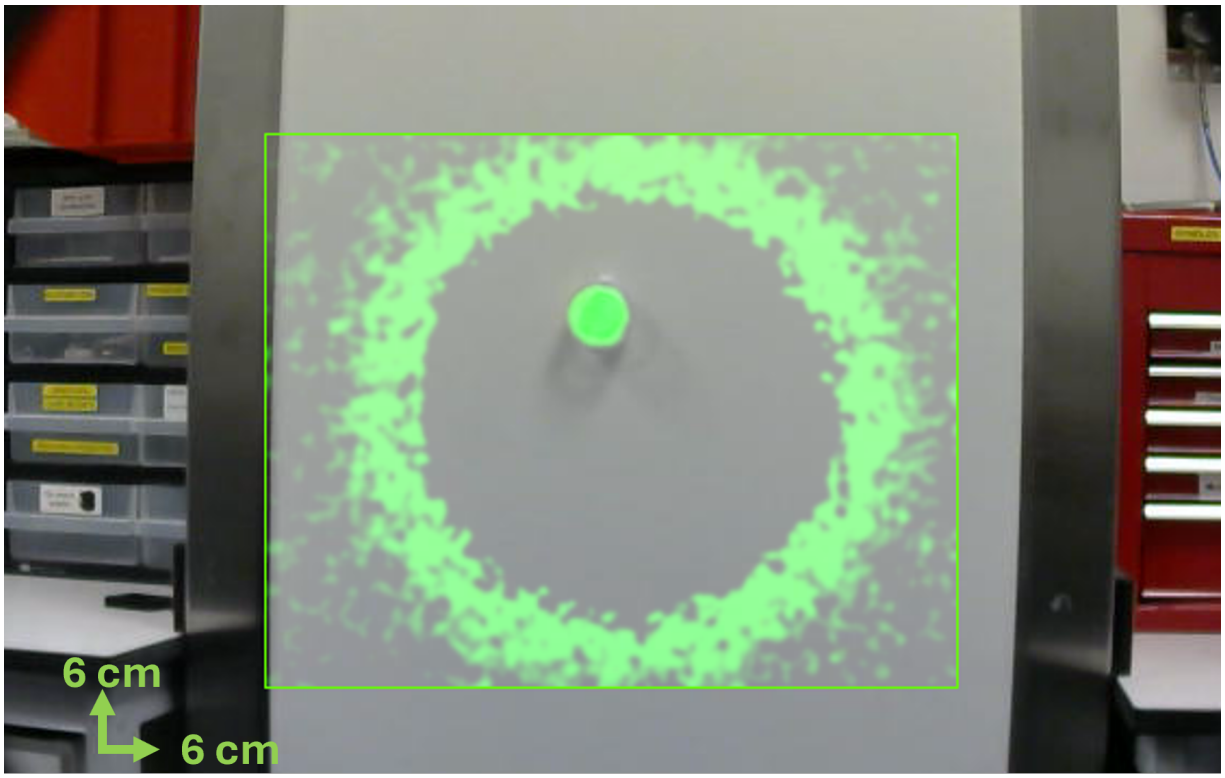

Figure 4.9: **Example of the ring visible when imaging using the** 12 nm **filter.**

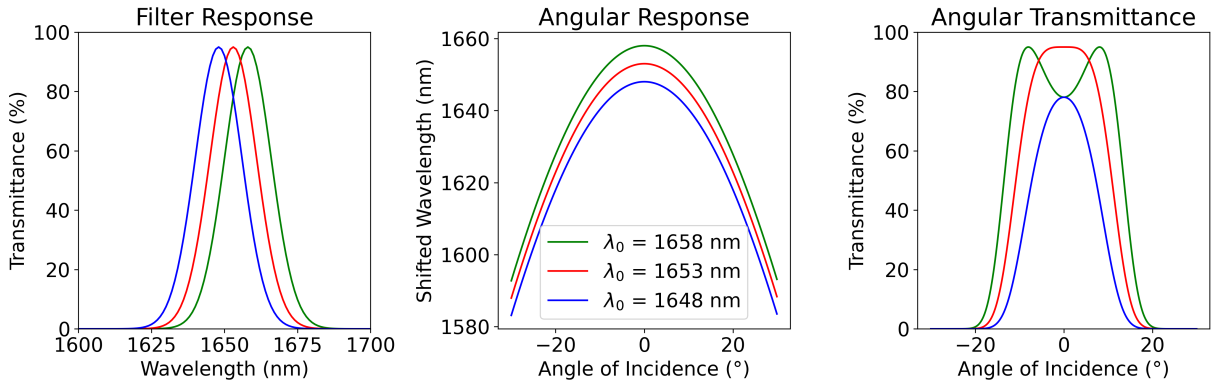

Figure 4.10: **Effect of varying central wavelength on filter angular response.** Parameters used for plotting are as follows: a peak transmission of 95%, a FWHM of 18.8 nm, a refractive index of 1.8, and an illumination source at 1653 nm. Plots vary central wavelength of the filter, λ_0. The difference between filters with λ_0 of 1658 nm, 1653 nm, and 1648 nm is presented.

Chapter 2.2.1, the differences in wavelength is on the order of tens of pico-meters, meaning that the transmission difference compared to the on absorption wavelength should be negligible. The shifted central wavelength results in a lobe like effect where two distinct peaks appear, made up from the two peaks caused by the fact that light entering at angles greater than the normal will have a higher transmission. This shifting also causes a dip in transmissible light at normal incidence, however it can be used to artificially increase the FOV of a system by separating the two peaks out more by increasing the central wavelength, though one must account for the central dip, which itself will become wider the further the central wavelength is from the desired wavelength of detection. Whilst the longer central wavelength solution shows great promise, it was unable to be explored experimentally due to difficulties and cost involved in manufacturing a filter that would match the desired specifications.

Similar effects can also be achieved by tilting the filter away from the incoming light, as this will effectively alter the angle of incidence upon the filter. There was no desire to investigate this effect due to the idea that implementing a tilted filter would be challenging to implement, and the ability to fine tune this within the system could add an extra level of complexity. Whilst there are benefits to filter tilting, such as amore favourable sunlight to diode ratio which itself lends itself to a higher SNR, there were fears that reducing the FOV would be seen to potential users as a downgrade from what the system had previously been presented as being capable of.

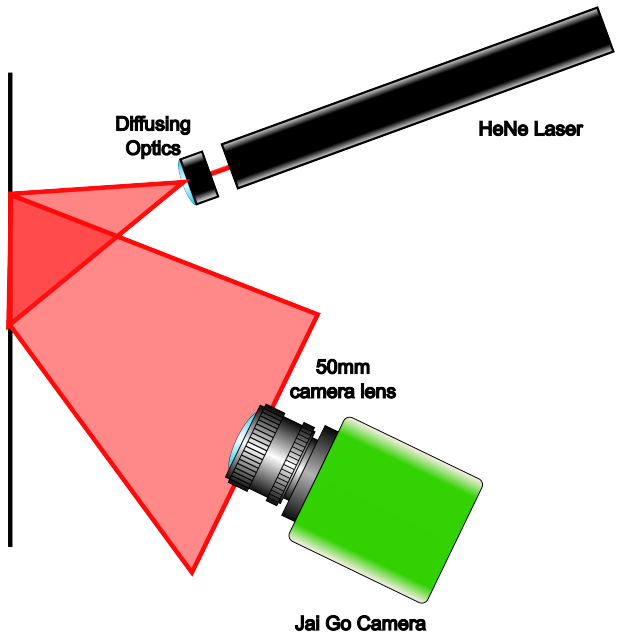

Figure 4.11: **The basic HeNe based filter testing apparatus.** Here the HeNe is diffused in order to maximise coverage. This light is reflected off the wall and recorded by the visible camera. This apparatus was then monitor the effects of using different filters in front of the camera.

4.7 Filter Experiments

In order to verify and experiment with the parameters of these models in a real setting, a small optical set-up was created. This setup utilised a Helium Neon laser emitting at 632.8 nm, and a visible camera (Jai Go). The HeNe was diffused with the aid of a 100× microscope objective lens, and a 20° square diffuser in order to illuminate a wide area of a wall. Additional tubing was also introduced after the diffuser in order to maintain eye safety. The illuminated wall was then imaged using the visible camera, this is presented in Figure 4.11.

This apparatus was constructed in order to investigate the impact of etendue, which is a conserved property relating the solid angle (which for a circular beam is related to the numerical aperture) to the physical area of light, but more generally describes the spread of the light [138]. As a property, etendue is conserved within a lens system meaning the product of the physical and numerical aperture entering a perfect, lossless system must equal the product exiting the same system. By taking advantage of this conservation a very simple lens system was created using two lenses with a 2 : 1 ratio of focal lengths, separated by the sum of their respective focal lengths. By introducing these two lenses one can alter the angle of incidence of light entering

the system as it leaves the second lens. This apparatus is shown in Figure 4.12. This is important because the light will be incident upon the filter at this new shallower angle and thus be more accepting of the light. The way that this works is that essentially the image is magnified by the lens system. By the conservation of etendue it can be said that approximately, the product of the physical size of the beam and its angular spread is approximately conserved following:

$$R_1\theta_1 \propto R_2\theta_2 \tag{4.5}$$

For $R_{1,2}$ being the radius of the physical size of the beam in states 1 and 2 and $\theta_{1,2}$, being the angular spread of beam in states 1 and 2. It thus follows that as the size of a beam of light increases its angular spread must decrease. This will follow a relationship similar to:

$$\frac{R_1}{R_2} \propto \frac{\theta_2}{\theta_1} \tag{4.6}$$

This means that the beam at lens 2 post magnification will be larger but will have a much shallower angular spread. This decrease in angular spread has the cost of decreasing the overall FOV of the system, but it is hoped that it will not decrease the FOV more than which would be lost due to the filter blocking the incoming light. The true trade off here is that because the beam size is now larger then new larger optics, including the filter, would be required after the lens system. This lens system is more favourable to simply altering the lens because of the position of the filter being in front of the lens. It is possible different experiments and gains could be achieved by placing the filter after the lens, but due to space limitations between the lens and sensor this was not considered.

To test this theory two lenses were selected, in this case both with 50 mm diameters, with respective focal lengths of 75 mm and 150 mm which were separated by the sum of the focal lengths, which is 225 mm. A diameter of 50 mm was selected such that the aperture of the lens would not be the limiting factor on the FOV, if the lenses acted to reduce the FOV individually there would be no advantage to implementing them into the system. These lenses were mounted within cage plates so they could be adjusted along a metal rod apparatus. The main idea here is that light incident on the outermost lens, which is the lens with the shorter focal length, at angle

θ_1, will leave the second lens with a shallower angle of travel θ_2, relative to the normal. This conservation means that if the angle of incidence is decreased at the point of hitting the filter, then the wavelength will be shifted less according to Equation 4.2, and as such some of the FOV will be recovered, the lens system is shown in Figure 4.12. By using the diffused HeNe laser set-up as well as the visible camera, with an attached 50 mm focal length lens, the effect on FOV was examined. A filter could be placed in two key positions; externally - wherein the filter is placed before the 75 mm focal length lens, which in turn effectively means that the impact of the two lens system is ignored and the FOV should be reduced as in the gas imaging system, or internally - wherein the filter is placed between the 150 mm focal length lens and the camera lens meaning that the filter will interact with the light at the reduced angle of incidence. The filters used are centred on 632.8 nm with FWHMs of 3 nm (ThorLabs FL632.8-3) and 1 nm (ThorLabs FL632.8-1) respectively. These lenses were 1 inch in diameter so are not perfect given that the size increase of the beam, but were initially used to provide evidence of the recovery of FOV. An experiment was assembled such that these filters could be placed in the external and internal positions and the outcomes compared. The outcome of the experiment is highlighted within Figure 4.13, wherein the FOV of the camera is severely hampered in the case where the filters are placed externally to the lens system. It is immediately evident by comparing the two external images that the FOV is decreased substantially, with the FOV when the 1 nm filter is in place being noticeably smaller. It is however more difficult to pinpoint a decrease in FOV for the system when the filters are placed internally. Whilst the images are certainly darker for both filters in comparison to the image when there is no filter, as a result of the filter cutting out all light that is not at, or around, 632.8 nm, it is difficult to identify if there is any noticeable decrease in FOV in addition to this darkening. It was shown that the two lens system can effectively be used in conjunction with a narrow filter and benefit from the filtering action without any extreme FOV loss.

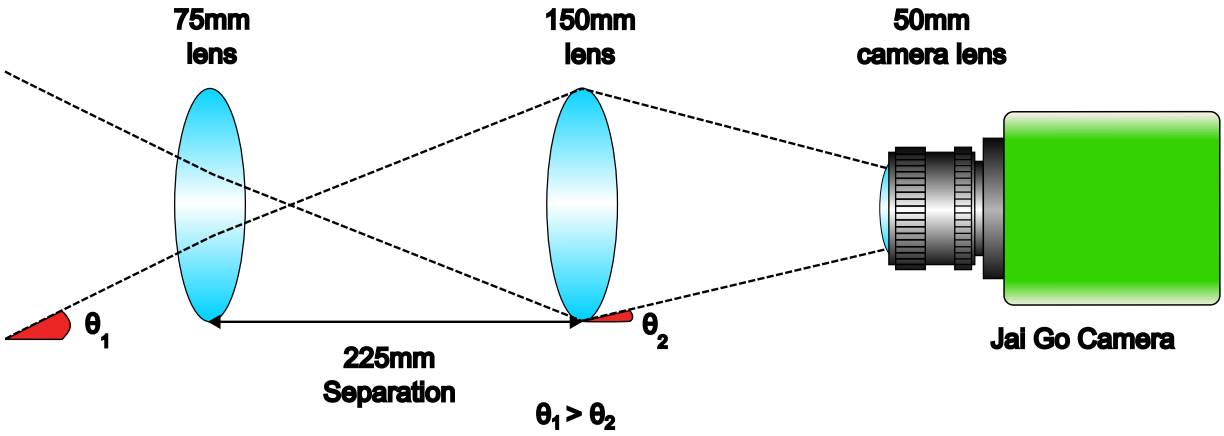

Figure 4.12: **Schematic demonstrating beam path through the dual lens system.** The result of passing through the lenses is a swallowing of the incident angle upon the camera lens. Figure adapted from the SPIE conference proceeding "Real-time imaging of methane gas from a UAV mounted system" [96], where increased clarity has been added to each component and the change in angle.

4.8 System Changes

Despite having demonstrated that it is possible to recover FOV lost to narrower filters, it was decided that the approach is quite cumbersome to adopt into the current system. The size of the dual lens system, as well as the difficulty in procuring larger diameter filters, would necessitate changes to the gas imaging systems form factor and mass, which must be kept minimised for portability, although these improvements may be incorporated into future device iterations. It was also found that obtaining the filter with a FWHM of 1 nm required to achieve the desired 5:1 sunlight to diode ratio may be challenging to manufacture as well as potentially very limiting in terms of FOV.

Attention was turned to potential hardware and software implementations which could improve performance in direct sunlight. The most simple approach is to somehow gain an advantage over the sunlight noise, either by raising the signal strength or diminishing the sunlight. If one was to reduce the exposure time of the system, the improvement would be twofold, as the system could run faster and be less susceptible to changes in the sunlight background, and having the shutter open for a decreased amount of time would in turn reduce the amount of sunlight being detected by default. The trade-off here is that in reducing the exposure time also reduces the amount of useful signal detected proportionally, essentially gaining no real advantage but

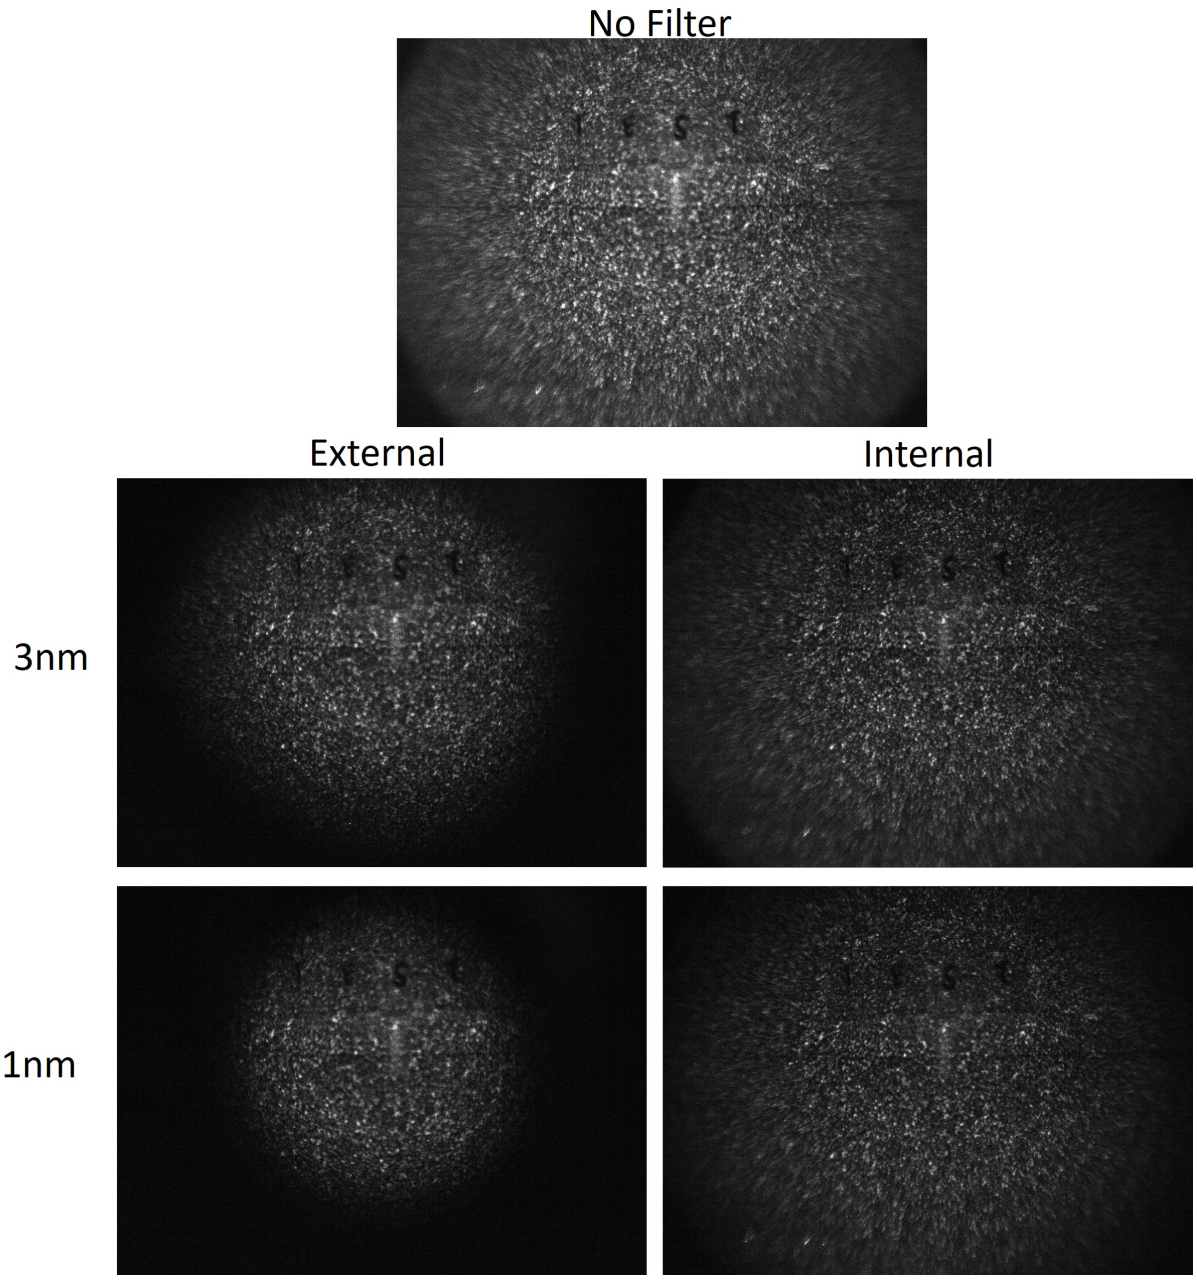

Figure 4.13: **Example images taken using the two lens system in conjunction with two filters.** These filters had FWHM of 3 nm and 1 nm respectively. The topmost image, included for reference, depicts what the system images with no filter present. The filters are then placed in the external and internal positions within the system and the obtained images are depicted as such. A high level of optical speckle is observed in these images due to the narrow line width of the HeNe laser used. Figure adapted from the SPIE conference proceeding "Real-time imaging of methane gas from a UAV mounted system" [96], where the layout has been changed from horizontal to vertical for clarity.

also decreasing the overall detection range capabilities of the system.

Another change that can be made to the system would be to use higher power diodes which should return a higher differential signal from the gas. A higher return signal could result in an increased amplitude, which could allow for an easier and clearer way of differentiating real signal from sunlight. One issue that can arise is the sunlight causes the camera to become saturated, then the additional power will not aid the system overall. The extra power can increase the functional range of the system as it ensures that a higher power level returns to the camera as a result of the inverse square law.

While in isolation neither of these changes appear to provide much of an advantage against sunlight saturation, the combination of the two might. If one increases the power level by a factor of two and decreases the exposure time by a factor of two then the same amount of light enters the camera as before. The advantage of performing such a balance is that one can decrease the exposure time as a result, thereby reducing the amount of sunlight entering the system per exposure. Decreasing the exposure time in addition to increasing the diode power will not result in an increase in system range but it will instead add extra benefit to the movement artefacts as discussed before. There is potential for less variation between frame, something that also poses a problem in high light level conditions. It was thus concluded that attempting to create a new system that can balance power and frame rate would be the best practical approach for minimising sunlight noise. It follows from the discussion of the desirable filter width target of 5:1 sunlight to diode power that this could be achieved by a $\approx 20\times$ decrease in filter width, it stands to reason that a similar ratio could be achieved via a similar increase in diode power and decrease in exposure time. Also to consider is the fact that the model used to come to the $20\times$ increase does not consider post processing techniques that can be used to further enhance SNR and potentially win out over sunlight with a far less extreme increase in power. Increasing the power by a factor of 20 would also raise concerns of cost, manufacturability and eye safety, meaning that any increase in diode power must balance these three concerns on top of the benefits over sunlight.

4.9 Conclusion

In this section a few practical solutions to filtering sunlight were presented. The problem of sunlight saturation and variance is very complex, and overall it proved difficult to adapt the system to overcome this. By modelling the spectral response of sunlight in relation to the diode power, it was decided that the best way to combat this was to decrease the amount of sunlight entering the system where possible as opposed to trying to account for it. The decision to not account for sunlight at video-rates was due to the wide array of contributing factors such as temperature and cloud cover. Initially, the solution was to block an increased amount of sunlight by using a filter with a decreased FWHM. It was discovered that a lower FWHM resulted in a diminished FOV, which had to be accounted for. Whilst a solution was found that allowed for the FOV to be regained using a dual lens apparatus, it was deemed too bulky to be effectively integrated into the system. Instead it was decided that in order to overcome these issues a new system would need to be designed that could run at a higher frame rate with more powerful diodes and a reduced exposure time.

Chapter 5

Drone Based Gas Detection and Surface Characterisation

5.1 Introduction

In this chapter the data obtained when flying the system on the drone is presented and discussed. Additionally, the systems ability to accurately detect a gas signal against a variety of surfaces was also explored in order to give an indication as to how robust the system could be to the end user. The bulk of this work is covered in the paper "Characterising the performance of a drone-mounted real-time methane imaging system". Data was taken for this chapter with the assistance of Steven Johnson and Simon Mekhail, whilst technical and construction support was provided by Graham Gibson. Overall supervision and project direction was provided by Miles Padgett.

5.2 Background

As introduced in Chapter 1, gas leaks can pose harmful to the environment, costly to the industries controlling the underlying infrastructure and pose a significant danger, with the worst cases leading to explosions causing harm to individuals or property. Therefore the promise of drone based gas detection for remote sensing would have significant benefits to personal

safety [139]. Being able to operate unmanned gas detectors provides hope that, by decreasing the human risk and increasing the ease of leak detection, routine inspections will become far more commonplace [139]. Helicopters pose an interesting option for gas monitoring due to their high speed and versatility, and for detector systems that work over an increased range they are ideal [140, 141]. The shortcoming of helicopter based detection is primarily related to their inability to access tight or difficult to access regions, for example forested areas or a gas plant [140]. Measurements have also been demonstrated using both small and large winged aircraft, with larger crafts able to host larger more sophisticated detector technology [142, 143]. Another key factor to consider is cost, the cost of using a helicopter or winged aircraft is much higher than purchasing and using a drone [140, 144, 145]. As an example, a single field trial using an aeroplane based detector was estimated to cost $50,000 compared to the approximate £6,000 total price for the drone used throughout this work [146]. Of note is the fact that the $50,000 total also included additional costs such as set up, access and staff, costs of which could be avoided with drone usage [146].

Satellite based leak detection is also widely used in research. The main drawback of satellite based detection is that it is rarely suited for small scale leaks, due to the low resolution of images which can be on the kilometre scale, though finer resolutions on the scale of tens of metres are also reported [147, 148]. Satellite detections have, however, found success in discovering large scale leaks that were not previously reported, such as one reported in Ohio [149].

Compared to a helicopter, drones have greater mobility and can access smaller spaces, giving increased access to gas leak sources. Similarly, the rotors produce less downwards flow which can disperse the gas leak being observed, this is known as down draft [140, 150]. Drones are also a useful tool in less built up areas, where access with cars is limited or where general population density is lower [151]. Drones can also be programmed to follow the exact same flight path each deployment, meaning the system can be easily automated and optimised, this is one of the primary features that could enable more regular facility checks [152]. This automation can even potentially lead to less expense in terms of insurance and cost incurred as a result of leaks [147]. Drone swarms are also used to sweep larger areas far more quickly, though is more expensive as it requires many drones and many detectors [153]. A different type of detection, regime called

"pushbroom" sensing can also be employed using drones and satellites [154, 155]. Such devices build up images by imaging a full sensor line at a time, as the system moves, either as the drone moves or whilst in orbit [154,155]. These systems can provide very high resolution images when the data for multiple lines are combined, but can be challenging to use in practice as accurate position data is required in order to correctly stitch together the images continuously, this issue can also be exacerbated by the movement of the drone [154].

Drone deployment is, however, not without difficulty. Drones themselves suffer from very short flight times meaning that, relative to other options, they present a low active detection time [140]. Another issue related to drone use are the rotors and the down draft they produce. Down draft can directly impact the target gas, potentially dispersing it out of frame or below the detection threshold [156]. Detector position can be selected to avoid impact from rotor disturbance, with imaging systems commonly attached to a gimbal below the UAV [157, 158]. Despite this, it is likely that a detection mechanism that can detect from a suitable distance, such as the laser based system utilised in this project, should be able to circumvent this issue because the air current will not directly impact the gas stream [156]. Other forms of gas detection have been trialled including those that gather air samples during the flight for later analysis [159]. The elevation from the source of the leak is also important - as the gas rises it becomes increasingly disperse making it harder to detect [156]. In the other direction, if the drone is too close to the ground the downwards stream spreading along the ground can completely displace the gas [156]. A further issue of the process is the fact that the drone is constantly in motion - even whilst in hover mode [139]. The constant movement must be accounted for in some way, with methods such as phase correlation being a common solution [139]. Whilst the targets used in the following are enclosed gas sources, in real leaks scenarios wind speed will greatly impact the visibility of the gas, as high winds will disperse the gas over a much shorter range [160]. The gas will however still be most concentrated at the source of the leak meaning that if the system can detect the cells from a moderate range, such as around 5 m it should still be able to detect the leak source [160].

5.3 Pre-Flight Preparations

Before the system was flown it needed to be attached to the drone in a safe way and the parameters of the flight carefully considered. The model of drone used was a DJI Matrice 600 Pro, this is a larger model of drone capable of lifting off with a payload of up to 5 kg. With the maximum payload the total flight time was predicted to be approximately 15 minutes, whereas with no payload the maximum was estimated as 30 minutes. The mass of the GasSight system is 3.15 kg, placing it safely under the maximum take-off mass for a payload. The mounting mechanism itself consisted of a custom mount that connects to the load-bearing support beams on the drone. A metal plate is placed atop the beams in order to distribute the weight over both sides of the drone. This plate has 3D printed screw-on clips that secure the plate to the beams. In the centre of the plate there is a standard camera ball joint with a tripod plate attachment. This plate attachment connects to the same shape of plate as attached to the base of the GasSight allowing for quick and easy mounting. The plate attachment features an additional lock to prevent the plate from loosening from the grip throughout the duration of the flight. The ball joint adds a degree of flexibility to the angle at which the GasSight can be pointed. Whilst this has to be moved manually when the system is grounded, it does ensure that the system can face downwards ensuring light emitted is back-scattered towards the detector. The tilt is fixed during each drone flight, but due to difficulties in replicating the exact angle each time the system is mounted it is likely that it differed from campaign to campaign. In general, the angle used was fairly shallow such that the camera could see ahead of itself whilst still being able to backscatter light off of the ground.

In terms of backscattering there are a few considerations to be made here. The scattering nature of the output light is fully dependent on the nature of the surface it scatters from. In the most extreme case a Lambertian scattering surface assumes that all light scattered off a surface is scattered evenly in all directions, whereas a specular reflective surface would scatter the light all in one direction depending on the incident angle of the incoming light [161, 162]. Different surfaces will have different extents of diffusivity, with some scattering more or less light over a range of angles [161]. The diffusivity of a surface is generally related to its roughness,

essentially describing how flat it is to the incoming light or whether it is made up of smaller elements [163, 164]. The relative roughness of a surface, and thus how diffusively it scatters the light, can vary depending on the wavelength used and the coarseness of the surface [163, 164]. Many sophisticated approaches for modelling surfaces and their likely scattering profiles exist, but were not utilised in the work conducted here [165, 166]. An experiment into the relationship between surface type and return signal is performed to indicate how large a factor this plays for the typical viewing angles and surfaces used with the device.

The drone also features a connecting HDMI cable which allows the screen of the system to be streamed to the mobile device used to control the system. The drone also features a range of up to 5 km, though it is unlikely it would be able to travel that far in the limited flight time afforded by the batteries. The drone is easy to keep track of as it provides real time positional coordinates, however in all trials conducted the drone never left line of sight. Another serious consideration of this particular system is that it was not waterproof, meaning it could not be flown in wet conditions. The system was also only capable of streaming up to a maximum frame rate of 60 fps, meaning that if the system was running at a rate faster than this some data loss could occur, this was not the case here however. Overall, the key limitations of the drone mounted system lies in its finite battery life and limited usability in adverse weather. Whilst in theory it could be used over large ranges providing extensive coverage, this was deemed unlikely to be possible given the limitations, nor was it required for the field trials conducted here. The system mounting is shown in Figure 5.1.

In order to test if the system would work effectively with the GasSight load, a dummy payload was made. This dummy payload made use of the casing and perspex breadboard of a previous system iteration with optical posts screwed into the breadboard to ensure a mass slightly greater than that of the real payload. This dummy system was locked onto drone mount and flown in a way that was envisioned to be similar to how the actual system would be used, which involves flying at an altitude of around 3 m and slowly hovering forwards as shown in Figure 5.2. In addition to these predicted flight paths, slightly higher and slightly faster flights were performed in order to ensure that the system would remain safely affixed for all the predicted flight parameters experienced during a real flight. It was thus concluded that the real system

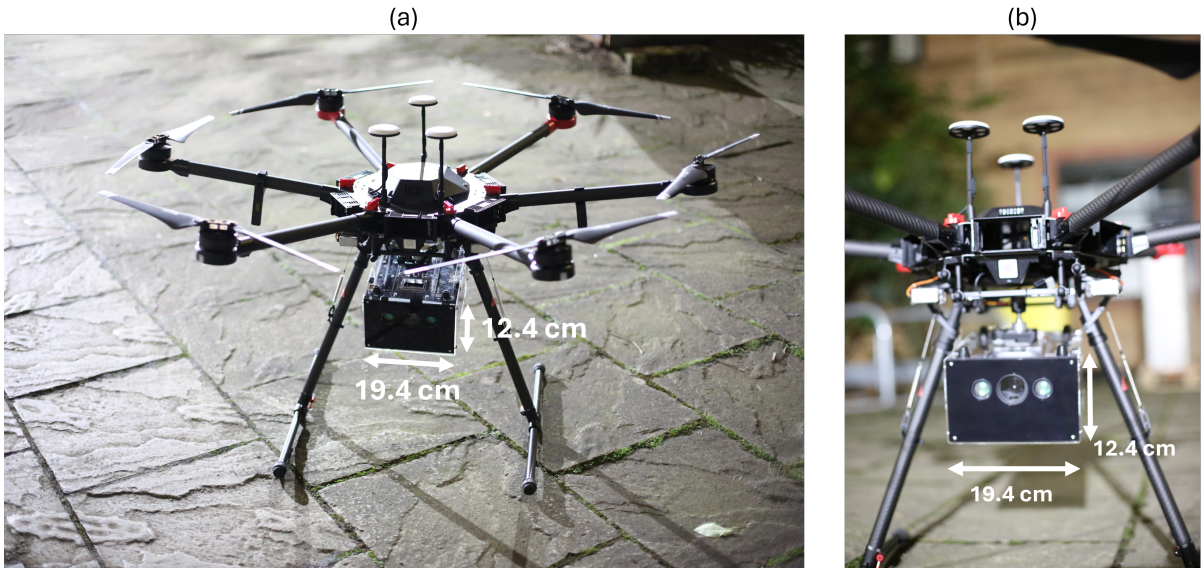

Figure 5.1: **The combined drone and gas detection device, highlighting the mounting below the chassis of the drone.** (a) demonstrates the full extent of the combined system, whilst (b) demonstrates the specifics of the mounting.

could be safely mounted onto the underside of the drone.

5.4 Drone Campaigns

Over the course of the work conducted with the drone, three distinct campaigns were conducted. Whilst the gas system's components and settings were kept consistent between these trials other factors such as time of day and other equipment brought along differed in order to try and maximise effectiveness. The three campaigns each highlighted separate challenges that had to be overcome.

5.4.1 Campaign One

The first flight was conducted during the day, the GasSight was affixed to the bottom and multiple flights were conducted. A key success was the ability to test the streaming of the camera feed to the phone controlling the drone. This test allowed live and unbroken access to the gas detection feed. As these flights took place during the day sunlight related issues proved problematic and resulted in a reduction in sensitivity. A ground sheet was brought to offer a more uniform reflective back scatter surface, additionally a co-operative back scatter board was brought, with

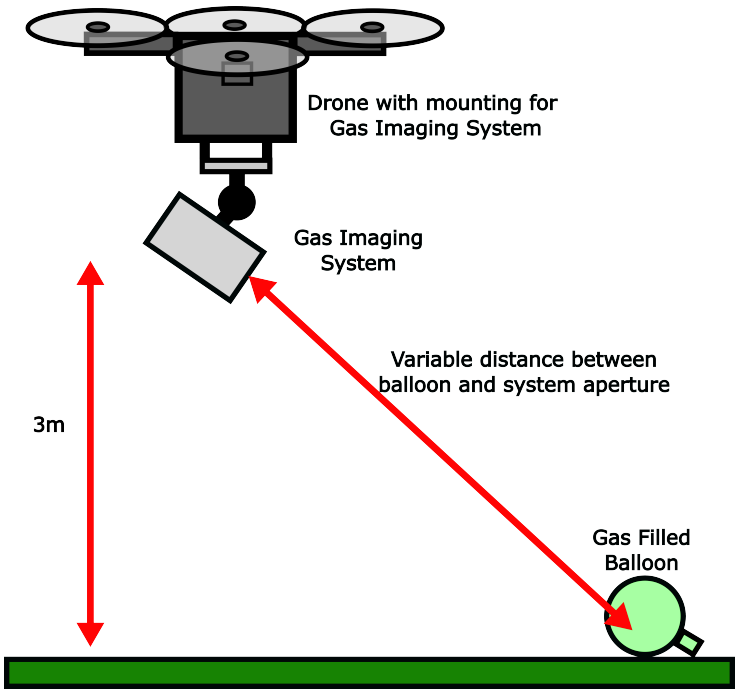

Figure 5.2: **A Diagram illustrating how the gas imaging system was used when mounted to the drone.**

gas balloons attached. These boards would act as a best case scenario, where despite other external factors gas should be able to be detected when in front of the co-operative back scatter surface. The balloons used were made of plastic and were clear both in the visible and in the infrared. Whilst a full spectroscopic analysis was not conducted as to the infrared absorption properties of the balloon material, reference balloons containing nitrogen were routinely used to demonstrate that this impact would be at best negligible as no differential signal arises from these nitrogen balloons. A similar rationale was used to justify using the glass gas cells used earlier, with a reference nitrogen cell used in tandem with some measurements. The balloons were 20 cm thick and filled with methane and as such contained a high concentration of approximately 200,000 ppm.m at the centre. In general, balloons and cells were used to avoid releasing methane into the atmosphere and to produce replicable experiments, though it must be noted that such measurements will not be representative of real gas leaks in the environment. Observations were taken using grass, the ground sheet and the co-operative back scatter board as the back scatter targets.

The data from this flight unfortunately was overcome with issues related to sunlight and as such much of the feed was dominated by noise. The gas signature could only be reliably

detected via the use of the co-operative back scatter board. This meant that ground sheet and gas were unable to return a gas signal. It was challenging to conclude whether or not the system was significantly hindered by the motion of the drone during this campaign because the availability of a clear signal against the noise was infrequent. Example frames highlighting the high noise levels and low detection capability of the data run are presented in Figure 5.3. The horizontal lines of noise are odd, but appear to be caused by image registration issues caused by a failing correlation in the y direction. This can be prominently seen on the top of the board in figures (d) and (e), the other lines appear to be from differences in the grass background exacerbated by the increased offsets attributed by the cross correlation.

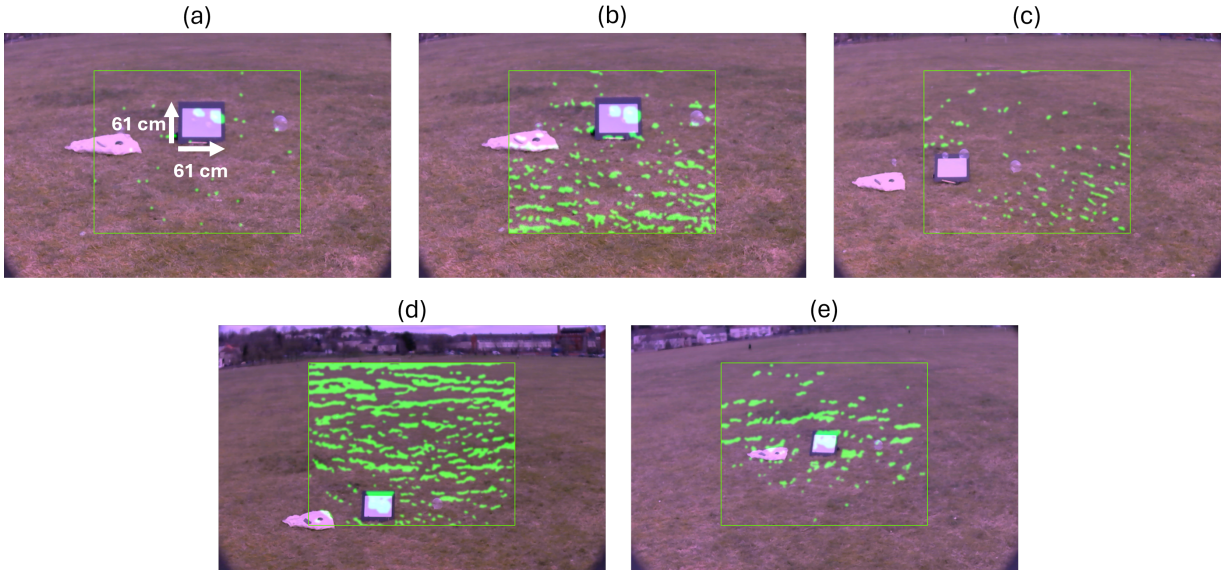

Figure 5.3: **A variety of example frames captured during the course of the multiple test flights.** These images were taken over the course of a variety of flights at different but similar distances from the balloons. A persistency filter strength of 2 is utilised in each. All 5 examples highlight issues encountered during the exposures. Frame (a) highlights a relatively low noise example where the balloons mounted to the co-operative back scatter board are clearly visible, but the balloon directly to the right is not detected. There is a large amount of false signal present in these frames, likely a combined result of saturation, sunlight noises, and a failing correlation calculation. Images (b), (d) and (e) all demonstrate frames with high noise levels but are still capable of detecting the board mounted balloons, whereas the frame in image (c) cannot, owing to the fact the wind has blown the balloons above the board. For a sense of scale, the dimensions of the backing board is provided in (a).

5.4.2 Campaign Two

In order to tackle the sunlight issues the next set of flights were conducted after sunset. The general technique employed here was kept the same as the daytime flights. Overall, the performance was noted to be much better, and the balloons were capable of being detected clearly against the grass without the need for the co-operative back scatter board. The downside, however of these test flights was the limited visibility afforded to the visible camera, which made locating the gas signals trickier. While this has no real impact upon the detection mechanism it does add a level of uncertainty to the system which could be met poorly by end users and may need to be avoided. One such solution would be add a toggle-able LED spotlight onto the system to act as a spotlight. One would need to ensure this light source has no emission in the SWIR however to prevent additional noise or loss of sensitivity. Some example frames from this campaign are presented in Figure 5.4. There was no noticeable impact from the drone vibrations upon the systems ability to detect the presence of gas consistently frame to frame. The extent to which the vibrations move the system are not enough to pass the various motion filters in place and as such do not cause any increased occurrence of noise artefacts. One issue pertaining to movement however, is that in instances where the drone is in movement this can cause some issue in relation to that of the persistency filter. The impact of this depends on the speed of the drone, so as a result many of the images shown in the figures are taken with the drone in a stationary hover mode. Figure 5.4(a) is an exception to this and features the effect of an abrupt stop of the drone in direction of travel towards the balloon, resulting in a harsh jerking motion backwards caused by the slight swing of the camera below the drone. In order to account for the higher degree of movement involved with the drone based detection a relatively high persistency strength (as introduced in Chapter 3.4) was used, specifically that of 5, requiring something to be in 6 subsequent frames to pass. This harsh filter is very noticeable in instances in which the drone moves forwards, backwards or pans left or right as the gas signal is temporarily filtered out as a result of the sudden, jerky nature of the motion as in Figure 5.4. In future iterations of the system it would be hoped that a higher imaging frame rate could result in a system less reliant on persistency and as such less sensitive to the jerky movements of the drone, outside of the standard hover configuration.

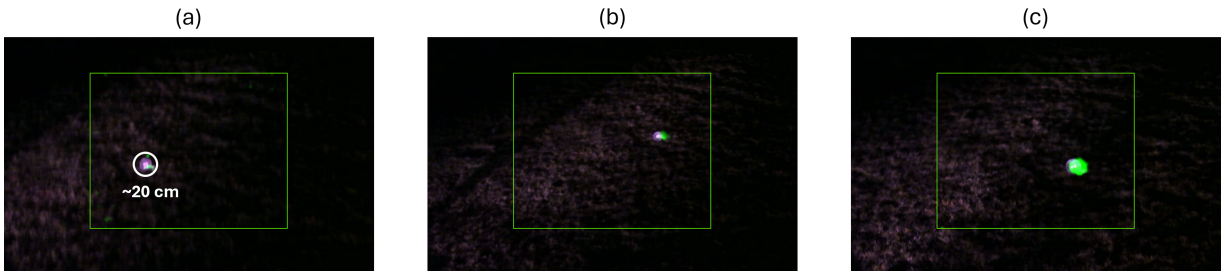

Figure 5.4: **Three example frames taken during the night time drone flights, all taken with a persistency strength of 2.** Note that a torch has been used to illuminate the scene for clarity. (a) demonstrates a frame in which the gas signal has been almost entirely lost due to a sudden jerky motion from the drone. (b) highlights a frame from slightly further away, there is an offset between the visible and SWIR camera here that appears to be as a result of viewing angle. (c) demonstrates an ideal frame taken from a closer range where the gas signal clearly overlaps with the position of the balloon. To provide scale, the approximate diameter of the balloon is provided.

5.4.3 Campaign Three

Having learnt from the first two campaigns, it was decided that the best mix of low sunlight and visible camera visibility would be to image during sunset and into twilight. This time a light meter (RS Pro ILM-201L) was also used in order to gauge how signal changes as the sunlight levels decrease. Given the higher visibility and experience at using the drone more data was taken on this trip, and footage obtained was much clearer, with the position of the balloon much easier to pinpoint in a frame. This data consisted of a range of flights taken as the sunset progressed at differing ranges as well as images and videos of the system in flight. An example frame, as well as an image of the drone in flight are included in Figure 5.5. These images are all taken at brightness levels measured to be less than 400 Lux and approaching levels of sub 10 Lux as the night progresses - which is consistent with sunset and significantly less than the approximately 10,000 Lux of sunlight where the data was taken [167].

Thorough investigations into which settings work best were conducted over the course of the three trips but in general it was found that a stable clear signal could be maintained for a gas balloon against grass for approximately 10 m horizontally and 3 m vertically. The effects of the setting Sun can be examined firstly, as demonstrated in Figure 5.6. As the Sun sets it can also be noted that the processed images become less noisy, this is a combination of a decrease in sunlight related issues as well as a decrease in an overall wind breeze that caused grass movement, which

may have led to movement artefacts to occur. Images taken around the same time as Figure 5.6 (B) give the best insight into the performance of the system on the drone, this is because there is a perfect amount of sunlight to see things clearly on the visible camera, but not such that noise related problems are encountered with the SWIR camera. It is also clear that there appears to be no issue of down draft interfering with the detection of gases at this range. The gas is enclosed within a balloon and so direct interaction between gas flow and down draft was not examined, however, the balloons were only secured to the ground using a tent peg through a non filled tip, meaning that they could be easily moved by airflow. The balloons themselves were not disturbed by the down draft in the flights conducted here.

For the data presented in Figures 5.5 and 5.6 a persistency level of 5 was used for each and every trial. As highlighted in the images, the gas balloons were of an approximate radius of 20 cm and were approximately 10 m horizontally from the drone. In the case of the images in Figure 5.6, these were taken when the drone was in a hover mode and the only movement of the camera would be small perturbations caused external factors such as by air current, for example. Figure 5.5, instead features a frame from a video in which the drone has flown towards the balloon at a constant altitude of 3 m. The speed of is difficult to quantify and is likely not constant given it was directly user controlled but was flown at such a slow pace that in the source video, as the balloons enter the detectors range the methane balloon is clearly highlighted despite the high levels of persistency. The still frame included in the figure occurs after the drone has stopped its transit and has returned to a stationary hover.

5.5 Surface Characterisation

Whilst the drone data was effective at showing that the system could successfully be integrated together the only surface available was grass. This was primarily due to limitations on where the drone could be flown, and as such the range of surfaces that could be accessed from above were limited. It was still believed that, in order to be convincing as a device that could be used in a wide variety of scenarios, the system should be able to detect gas using a wide variety of back scatter surfaces.

Figure 5.5: **Two images taken during the third set of flights.** Left image is a capture of the system mid flight, flying at an altitude of approximately 3 m. The left image is an example frame from the system wherein a methane balloon can clearly be distinguished from 3 other Nitrogen filled balloons. Here the gas is highlighted in red as opposed to green to improve contrast against the grass. To provide scale, the approximate diameter of the balloon is provided. Figure adapted under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from "Characterising the performance of a drone-mounted real-time methane imaging system" [97]. Figure was edited to include the scale of the balloon.

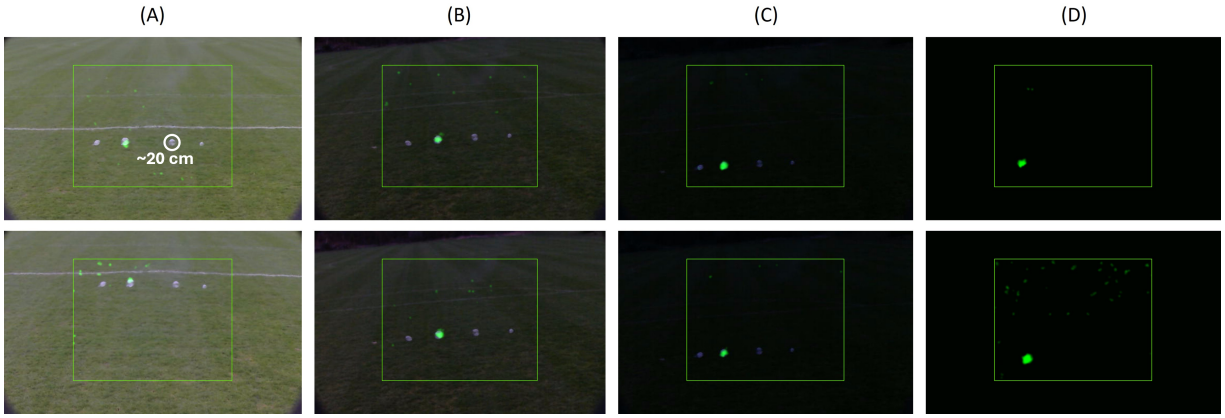

Figure 5.6: **Processed images taken by the drone mounted system as the sun sets.** The progression of time goes from (A) to (D) where it is clearly visible from the visible camera feed that the available sunlight decreases as the frames become progressively darker. The lux values measured for these images are as follows: (A) was 250 Lux, (B) was 80 Lux, (C) was 35 Lux and (D) was 6 Lux. To provide scale, the approximate diameter of the balloon is provided. Figure adapted from the SPIE conference proceeding "Real-time imaging of methane gas from a UAV mounted system" [96]. The figure was edited to include the approximate balloon diameter.

It was decided that a range of common surfaces should be investigated, such as paving, grass and tarmac. In order to mimic the height and perspective of the drone a tall tripod was used, this height also facilitates a long range due to the fact it allows us to look down onto a gas, giving a larger back-scatter surface compared with looking at the horizon where there is less likelihood of return signal. The tall tripod was set such that the imaging device was placed at a height of 2.12 m, whilst this is not as high as the height maintained by the drone it does allow a horizontal range exceeding 10 m to be investigated. Another consideration that was taken into account was the fact that the investigations would have to be conducted at night for the same reasons as before, the light meter was taken in order to ensure some level on consistency between the light levels of the readings.

The six surfaces that were investigated were grass, tarmac, aluminium, steel, paving stone, and card. The initial five surfaces were selected due to the fact that they are common surface types one would encounter in urban areas, where the system is envisaged to be used by an end user. Additionally to these common surfaces, a surface made up of 18% reflective grey card was used as this represents a common and specific reference in imaging. This card was introduced in order to present a potentially replicable surface for control purposes, its properties in the SWIR were not investigated, but given that it is commonly used a photography standard it was deemed suitable as a control. The surfaces were investigated such that a balloon was placed a measured distance from the aperture of the system, and then a sequence of 400 pairs of images would be obtained for each distance in order to post process at a later date. Firstly, grass was investigated with a balloon, at ranges from 2.094 m to 11.083 m in increments of approximately 1 m, next the paving was investigated similarly for a range of values 2.132 m to 11.046 m again in increments of 1 m. The balloons used here were of the same type as used in the drone experiments, they contained pure methane and corresponded to approximately 200,000 parts per million metre (ppm.m) at the centre of the balloon. These two data sets were obtained in a single night, however due to the length of time required to set up and capture the image sets for these two surfaces a different approach was taken for the second night. As these measurements were taken in an environment with visible light pollution, the light meter readings were taken in approximately the same location each time.

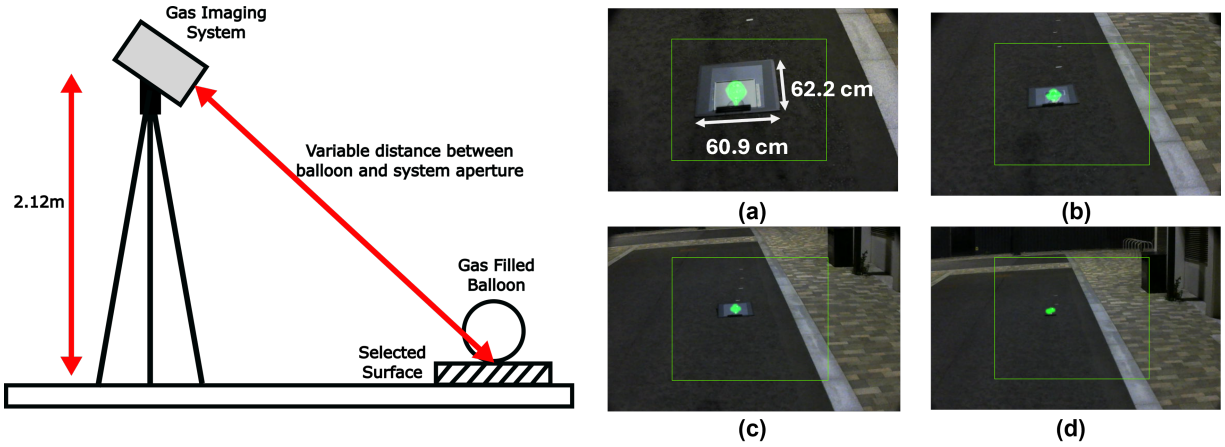

Figure 5.7: **The surface imaging apparatus.** Simple figure of the imaging set-up used for surface presented on left-hand panel Right-hand panel displays processed images taken for (a) steel at a distance of 3.1 m, (b) aluminium at a distance of 5 m, (c) 18% reflective photography calibration card at a distance of 7.1 m, and (d) tarmac at a distance of 9 m. All distances are quoted in terms of diagonal distance to the aperture of the camera. Here a gain of 0.7 was applied to the visible camera feed to increase gas signal clarity. Figure reproduced under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from "Characterising the performance of a drone-mounted real-time methane imaging system" [97].

The remaining four surfaces were observed in the second night. 400 images were again taken for these data sets, but this time the distances were pre-marked out with masking tape at distances ranging from 2.171 m to 9.905 m. By marking out the target locations before the imaging campaign, the data could be gathered more efficiently as the card, aluminium, and steel were individual sheets and could be slotted underneath the balloon, allowing for the surface to be quickly swapped, this process is shown in Figure 5.7, of note is the fact that the x-axis here is plotted in terms of the inverse square of the distance. This x-axis was selected to give a linear response to the data. Whilst initially the sheets were laid flat underneath the balloon, at further distances a slight tilt was added to each of the portable surfaces in order to fully ensure that the light that passes through the balloon will reach the desired back scatter media. This may have been an error in hindsight as it changes the angular dependence of the diffuse scattering as discussed previously. This could explain the noticeable jump in signal seen in the second night data in Figure 5.7, though of note is the fact that this jump is present in the tarmac data as well, meaning it is more likely to be a result of the tilt of the system changing instead of the surfaces. A future experiment with tighter control over the angle of illumination may provide greater insight into the differences in the specular reflection between the surfaces.

Once all the data for each of the surfaces had been gathered the data was processed in order to directly compare the return signals obtained from the different surfaces. The process of this largely follows the image processing timeline as outlined previously in Chapter 3, but with some subtle edits. Firstly, as always, the direction of subtraction is key, with the correct order of an on absorption frame being subtracted from the off absorption frame. In the imaging timeline used to capture the 400 image data sets there is a control designed to swap the order of which the images are taken such that if iteration i takes an on absorption followed by an off absorption frame, iteration $i+1$ will feature an off absorption followed by an on absorption, as has been outlined before. These differential signals are thresholded and then masked, this mask is passed through a persistency filter of 5, and eroded and dilated. The nature of the erosion and dilation structuring element in this instance is 3×3 array with zeros in the corners. The mask is then applied back to the differential SWIR image. In order to allow the persistency filter to take proper effect, the first 5 frames of the 400 are not included in the overall processing and instead act as buffers for this filter. Instead of converting this masked signal from 16-bit to 8-bit as usual, the masked signal is used to quantify the signal strength, as converting to the 8-bit signal could cause a loss in sensitivity. For each masked differential signal, the average signal strength over the extent of the gas signal is calculated by dividing the sum pixel values within the masked signal by the number of active pixels in the mask. This acts as an average signal strength for a reading. In order to quantify the variance on the signal a standard deviation is calculated for the average values and is used as the error bounds on the measurement.

By obtaining these average values one can then compare the return signal between the different surfaces and evaluate if there was a marked difference in their ability to return gas signal. The data is plotted and presented within Figure 5.8. What can be seen from these graphs is that while performance does vary between each of the surfaces, each and every one of them returns an identifiable signal from the system within the examined range of the system. Surfaces such as steel and paving are clearly more reflective than grass, as demonstrated by the higher overall signal return for these surfaces at close range, but grass still returns an adequate signal at the far distances that allows for gas identification against noise.

Additionally included in Figure 5.9 is the signal to noise ratio (SNR) of these data points as

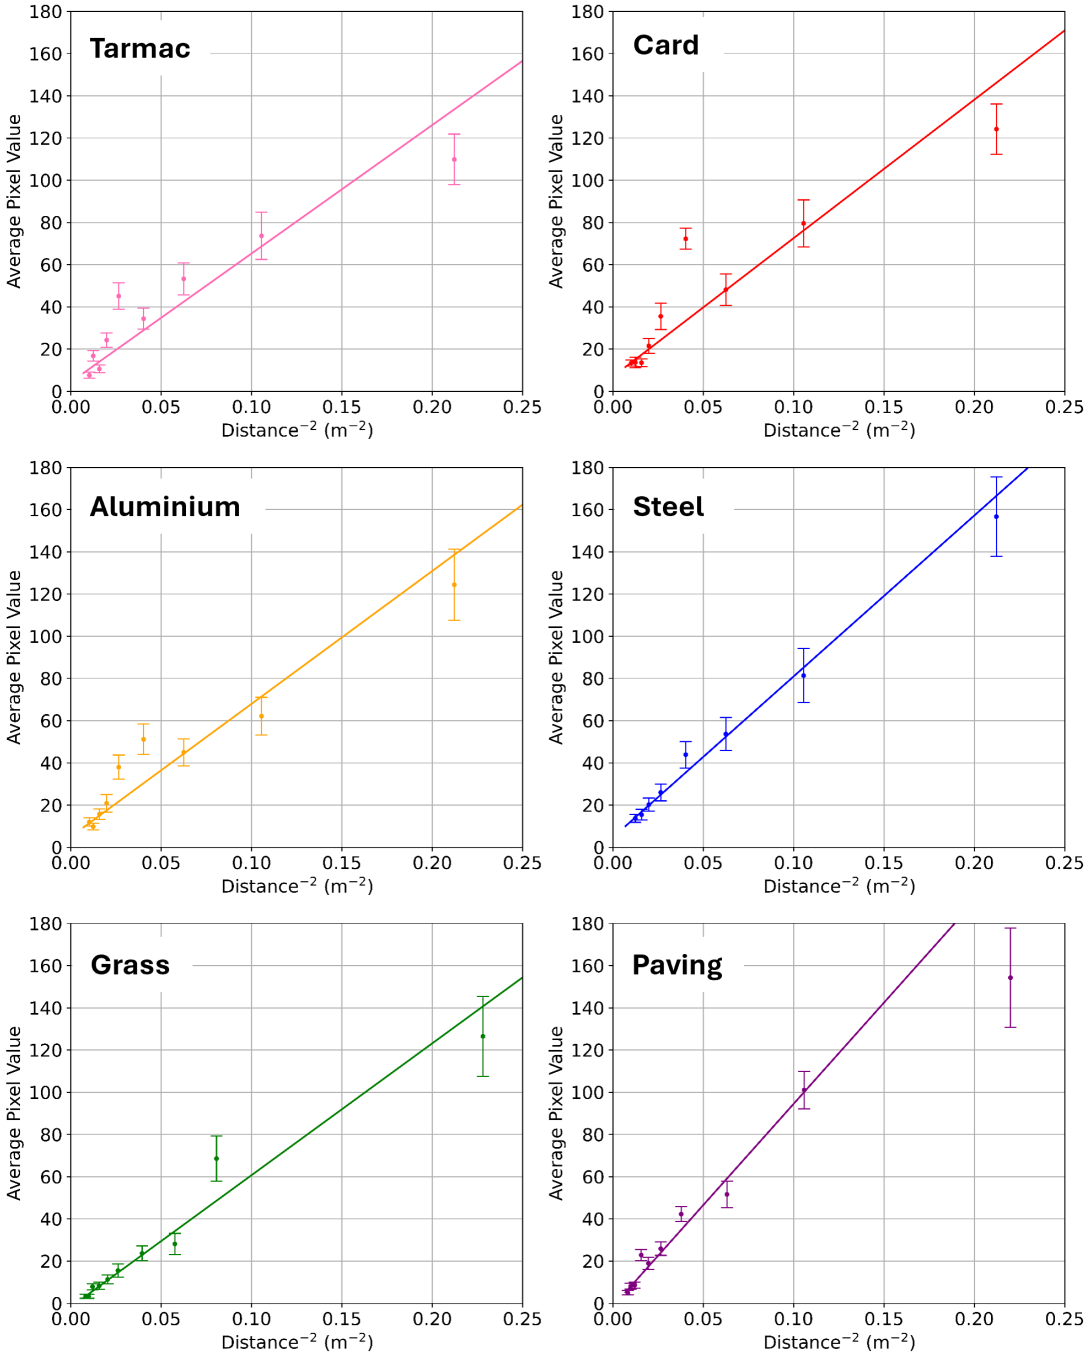

Figure 5.8: **Graph showing the overall trend in signal strength with increasing distance from gas balloon, using a simple linear fit to the data.** The x-axis here has distance to the power of -2, this was selected to incorporate the $\frac{1}{d^2}$ relationship with signal strength and distance. Figure adapted under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from "Characterising the performance of a drone-mounted real-time methane imaging system" [97]. Figure format edited to a 2×3 layout for clarity.

a function of distance. Whilst there is no overall clear trend presented in this data there are a few key takeaways. Firstly, is the fact that paving appears to provide the overall highest SNR suggesting that it could narrowly be the most favourable surface, but in general they all behave fairly similarly. Secondly, is the fact that a lack of a clear trend may indicate that the definition of noise in this experiment may be incomplete and may need to include further distance dependent considerations as discussed in Chapter 2.

Further to these surface measurements, an additional measurement was taken in order to effectively characterise the systems range. These tests were conducted using a glass gas cell filled with a concentration of 5000 ppm.m placed at one end of a hall and the system set to collect frames over a range of distances. This is a lower concentration than the balloons and cells used previously, by up to a factor of 20×. The overall procedure follows that of the surface characterisation. The results of this data run is shown in Figure 5.10. What this shows is that the system, in the most favourable scenarios is capable of detection from ranges exceeding 15 m.

5.6 Conclusion

In conclusion, the system has been demonstrated to work effectively in tandem with a drone for aerial observations of gas signals, albeit with promising results only possible at dusk with the hardware and setting used at the time. Mounting to a drone is desirable due to the higher level of flexibility and automation that it provides. It was shown that the impact of rotor based vibrations had no substantial impact on the ability to accurately locate the gas position within the frame. The system integrates into a pre-existing drone easily using a custom mount, which provides a strong and robust connection, securing the system safely. Whilst sunlight continues to be an issue, the performance of the system at night highlights the improvements made to the motion filtering by being able to maintain a gas signature whilst the drone is in hover mode. In addition to proving that the system can integrate with a drone, the detection response against a variety of surfaces was explored. It was discovered that the system overall can detect gas signals at desirable ranges for a wide variety of common surfaces. Of note however is that lack of testing of wet surfaces, which have different reflective properties. Anecdotally the system

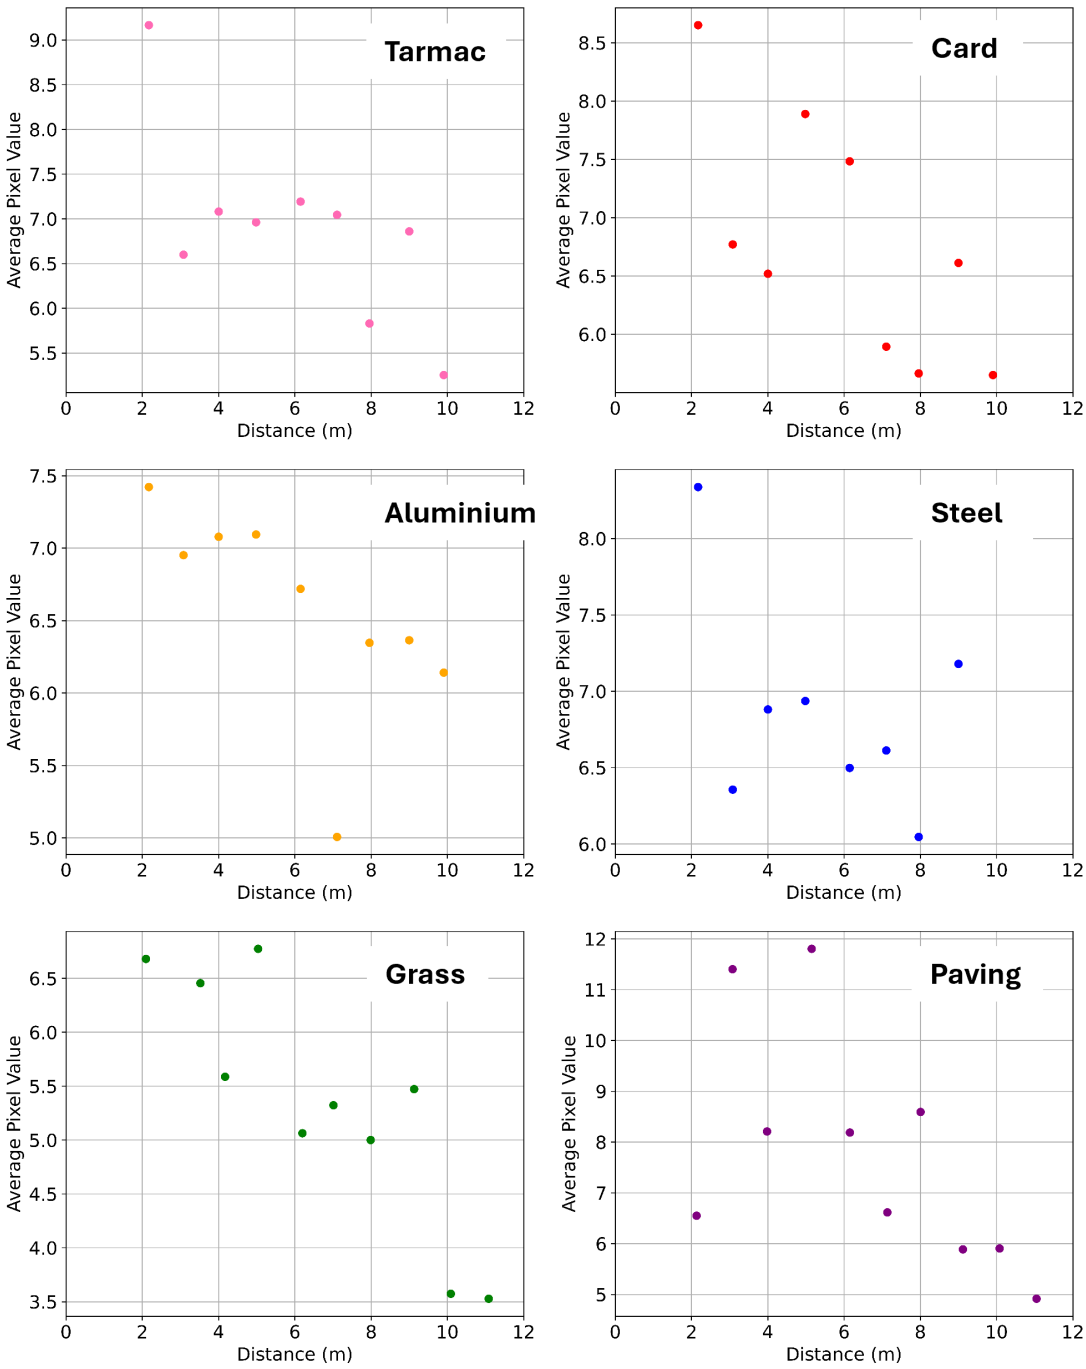

Figure 5.9: **The gas signal to noise ratio as a function of distance for the different surfaces of interest.** It can be observed that, with the exception of paving, many of the surfaces provide very similar responses.

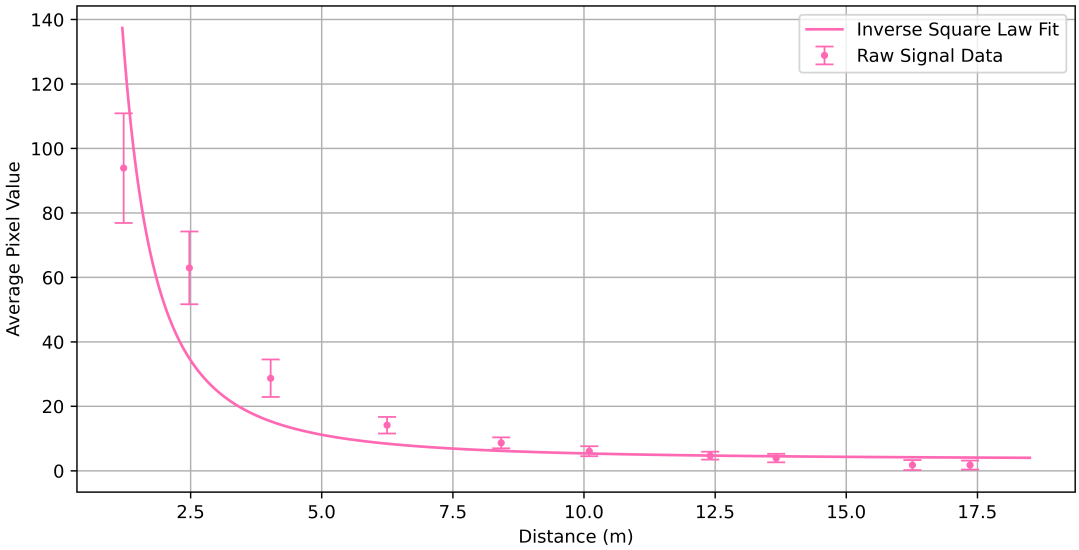

Figure 5.10: **A graph demonstrating the return signal from a gas cell of 5000ppm.m as a function of distance.** Ranges between 16.3 m and 17.3 m are shown as returning signals of approximately equal size to their standard deviation, resulting in an intermittent flicker in the image. A clear $\frac{1}{d^2}$ trend is highlighted from the data. Figure reproduced under a creative commons license (`https://creativecommons.org/licenses/by/4.0/`) from "Characterising the performance of a drone-mounted real-time methane imaging system" [97].

has been shown to work in such conditions such as on wet paving slabs at close range, but a deeper study into how this may affect the range was not conducted. Finally, the range was tested for a cell of of concentration 5000 ppm.m showing reliable performance up to a range of 16 m. The results of the surface tests and the drone mounting are promising results in terms of highlighting the robust nature of the system. It is hoped that, given the opportunity to perform a real world demonstration, the system would be able to detect gas leaks along a pipeline whilst drone mounted.

Chapter 6

Device Upgrades and Redesigns

6.1 Introduction

This chapter summarises the upgrades made to the system hardware as well as the process behind building the next iteration of the GasSight system. The new system acts as a full re-evaluation of the components of the older system in such a way as to utilise the data obtained so far to improve performance. The rational behind the differing parts will be presented. The upgrades made to the system hardware, consist mainly to changes to the laser diodes. These upgrades, in conjunction with the software changes made previously, allowed the performance of the system to be maximised. The next iteration of the GasSight will be introduced, explaining how it works differently and what the key components are. This version was conceptualised to combat the motion and sunlight related issues introduced previously. The aim of the new design revolves around increasing the overall frame rate and efficiency of the system, with many of changes made specifically to delivery of the SWIR diode flood illumination.

6.2 Device Upgrades

A primary hardware upgrade made to the existing system was to update the diodes, internal to the system, for a higher power illumination. This was achieved by using diodes with a higher power at a lower temperature, with not dissimilar drive current. All 4 diodes involved in this process

are Eblana Photonics EP1653-7-DM-TP39-01. Excerpts from the data sheets, highlighting the increased power for similar drive, are presented in Figure 6.1. The aim was to be able to eke out more range from the system with the increased power, with no additional drive current required. The driver board within the system was also replaced with a newer design, that was slightly more compact, reducing the space taken within the system. The diodes were tuned using the customisable PID settings within the temperature control chips on the OEM driver board. The PID settings were optimised for wavelength shifting with minimal temperature overshoot, and thus reduced tuning time. As before, an arduino was also used to trigger the diodes on and off absorption by changing the output voltage from the DAC.

In general, hardware changes made to this version of the system were kept to a minimum. This was to allow for the creation of a new system, whilst also experimenting with the software on the previous system, as well as perform field trials, at the same time.

6.3 New Design

A new version of the GasSight system was worked upon consistently throughout the period of this PhD. All progress conducted on building the new system was performed in parallel with work conducted with the old design, this was to ensure that use and demonstration of the older design could continue uninterrupted. The primary motivation behind the construction of a new system was to use the data gathered so far with the previous version and use it to tackle issues related to range, frame rate and sunlight noise. The main avenue for achieving this was to run the system far more efficiently by decreasing the amount of downtime in the system where the camera is closed, whilst trading off as little in performance and sensitivity as possible. The first idea was to purchase a new type of diode in order to maximise power output.

6.3.1 SWIR Camera

Before tackling the new diodes, an important investigation into the SWIR camera was conducted. There were initial plans to upgrade the SWIR camera in the system in order to improve overall performance, which would be compared to the current Goldeye camera. The camera se-

Parameter	Value	Units
T_{sub}	35	°C
I_{th}	25.33	mA
SE	0.12	W/A
V_f at I_{op}	1.32	V

Table 1: Key LIV parameters

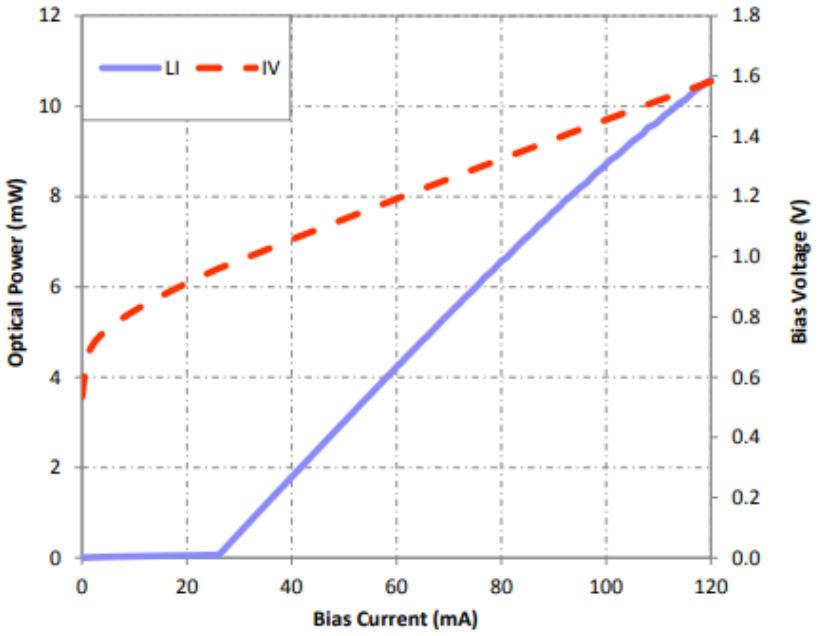

(a)

Parameter	Value	Units
T_{sub}	28	°C
I_{th}	19.18	mA
SE	0.16	W/A
V_f at I_{op}	1.38	V

Table 1: Key LIV parameters

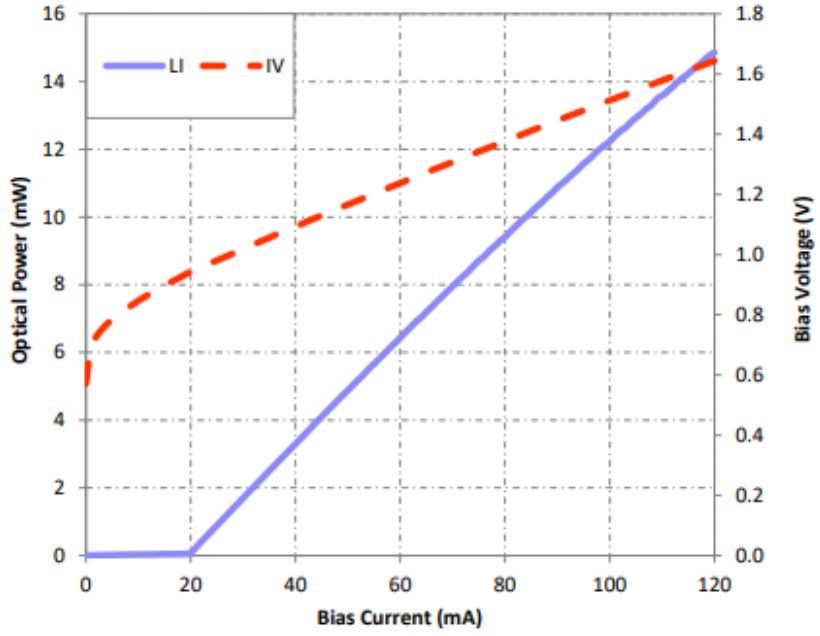

(b)

Figure 6.1: **Data sheet extracts highlighting the power increase from a previous diode (a) to a new diode (b) for similar current input.** The original data sheets were created by Eblana Photonics and are reproduced here with permission.

lected was a FirstLight C-RED 2 Lite, which offered a similar form factor with higher resolution camera with 640×512 pixels with a size of 15 μm. It was hoped that by increasing the resolution one could overcome some issues with range wherein objects can often be extent limited such as small cells or plumes. The trade-off here is that, as the output light from the diodes remains the same, the amount of light per pixel is less, potentially limiting overall range, highlighting the importance of photon efficiency over resolution. The new model also promises a high frame rate, reaching up to 600 fps which would facilitate an overall frame rate increase to the system. A comparison of a frame from the C-RED 2 Lite and the Goldeye G-008 is present in Figure 6.2. Another reason that this new SWIR camera was not used further was that the software development kit (SDK) was too computationally demanding. Running the camera drew far too much computational power that it resulted in quite a laggy performance especially for the LattePanda micro-PCs used in the system. It was decided that the higher resolution was not worth pursuing if it meant that a small compact PC could not be used as it could limit the handheld nature of the system. Additionally, the increased resolution decreasing the light per pixel was deemed not worth it in conjunction with the difficulty in operation. A more in depth comparison between the cameras is included in Table 6.1, comparing further features such as size and noise. Overall, it was concluded that the Goldeye camera used previously was reasonably optimal for the system, despite newer systems being available.

6.3.2 Visible Camera

A very minor change made to the system was the inclusion of a higher resolution visible camera. The new camera (Logitech Brio 100 webcam) has a resolution of 1920×1080, which is an improvement when compared to the previous cameras' (Logitech c270 webcam) 640×480. The reason for the change is primarily aesthetic, the lower resolution of the previous camera was deemed at times to be too low which can sometimes produce quite a fuzzy feed. The new camera, though not a large improvement, offers a clearer image without compromising on the small form factor that can be easily implemented within the system. The output of the two visible cameras for the same scene is demonstrated in Figure 6.3. Both cameras do suffer from aperture stopping as a result of the outer casing of the system, this is not an issue because the

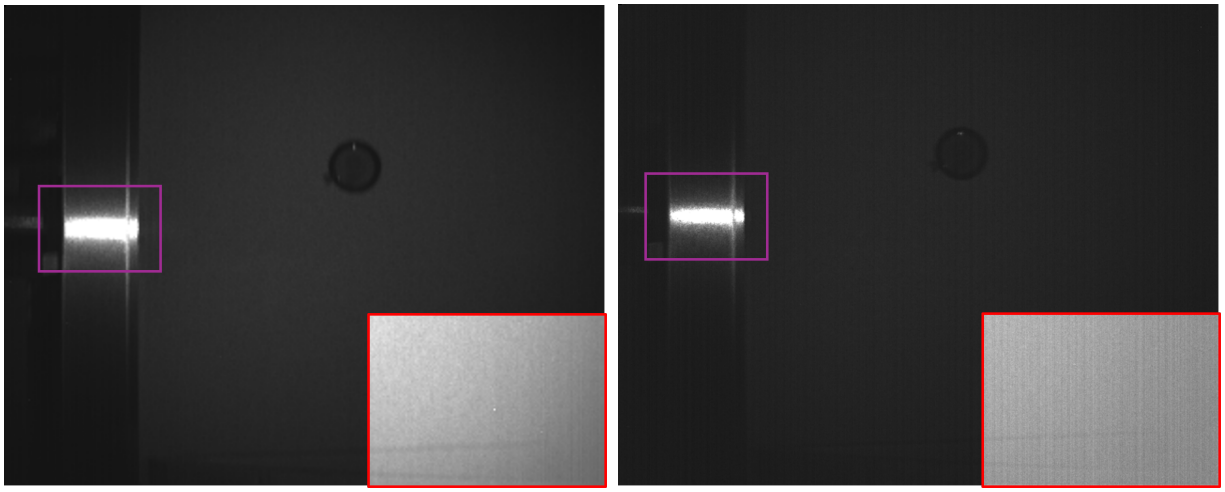

(a) (b)

Figure 6.2: **Example frames comparing the performance of the (a) Goldeye G-008 camera and (b) the C-RED 2 Lite camera.** What can be observed firstly is that the Goldeye camera produces a visibly brighter image, this is due to the lower resolution, this is most prominent in the size and extent of the specular reflection off the metal bar, highlighted in each image by a purple square, it is smaller in the C-RED 2 Lite image. The FOV of both cameras is essentially identical, this is because of the fact that the overall sensor size is the same despite the differences in resolution as a result of the difference in pixel size. Finally, the C-RED 2 Lite exhibits a much higher degree of structure in the images, caused potentially by underlying readout or pixel issues causing persistent stripes of higher or lower signal to be seen in the image. To highlight this a region of 10× gain is added to the images, as highlighted by the red square, where this structure is more prominent in the C-RED 2 Lite image. This noise should cancel out in the differential images.

Specification	**Goldeye G-008**	**FirstLight C-Red 2 Lite**
Resolution	320×256	640×512
Pixel Size	$30\,\mu\text{m} \times 30\,\mu\text{m}$	$15\,\mu\text{m} \times 15\,\mu\text{m}$
Quantum Efficiency at $1.65\,\mu\text{m}$	$\approx 65\%$	$\approx 65\%$
Sensor Type	InGaAs	InGaAs
Connector Type	Gig-E	USB 3.1
Gain Modes	0, 1	Low, Medium, High
Dark Noise	$1.6\ \text{k}e^{-1}$ (Gain0), $210e^{-1}$ (Gain1)	$35e^{-1}$ (High Gain)
Dark Current	$880\ \text{k}e^{-1}s^{-1}$	$24\ \text{k}e^{-1}s^{-1}$
Max FPS at full frame	344 fps	600 fps
Dimensions	$78\,\text{mm} \times 55\,\text{mm} \times 55\,\text{mm}$	$78.1\,\text{mm} \times 65\,\text{mm} \times 65\,\text{mm}$
Mass	340 g	460 g
Approximate Cost	£9,000	£25,000

Table 6.1: **Table comparing key specifications of the two SWIR camera options.** Data was obtained from the respective data sheets available at `https://www.alliedvision.com/en/camera-selector/detail/goldeye/g-008-swir-tec1/` and `https://andor.oxinst.com/products/c-red-series/c-red-2-lite`, respectively.

FOV in the non-aperture stopped region still far exceeds that of the SWIR camera and thus there is no demand to increase the size of the window. In order to make the SWIR image more accurately overlap with the pixels of the visible image, the SWIR camera pixels are re-binned in such a way that it can be scaled over the extent of visible picture.

(a) (b)

Figure 6.3: **Comparison between the (a) old and (b) new visible webcams.** Whilst there issues in terms of the outer casing of the device causing aperture stopping, it can be noted that camera (b) offers a sharper image and a larger FOV.

Different cameras were investigated such as a Sandberg 1080P Saver. The Sandberg camera also features a resolution of 1920×1080 pixels, but was not used due to its less reliable initialisation within the LabView software when compared to the Logitech cameras.

6.3.3 Laser Diodes

As discussed previously the diodes used within the original GasSight system output approximately 10 mW. It was discovered that diodes with approximately an order of magnitude higher power could be purchased (Aerodiode 1650LD-4). The diodes received had powers of 55 mW and 70 mW and had a different form factor as they were fibre coupled, this meant that they had a greater degree of flexibility for placement within the system housing. The only requirement was that the diodes were housed in such a way as to minimise damage to the fibres themselves. A table comparing additional features of the two diode types is included in Table 6.2.

The new diodes were more powerful than the previously used diodes and this in turn allows the exposure time to be reduced to achieve similar signal levels than before. The diodes are driven using proprietary laser driver OEM boards (Aerodiode CCS-CW-std-OEM) which can

Specification	**Eblana**	**Aerodiode**
Threshold Current (mA)	20-25	20-30
Spectral Width (kHz)	2000	200-500
Side mode suppression ratio (dB)	30-40	35-50
Wavelength shift with temperature ($\mathrm{nm\,^{\circ}C^{-1}}$)	0.07-0.1	0.08
Wavelength shift with current ($\mathrm{pm\,mA^{-1}}$)	10-15	3

Table 6.2: **Table comparing key specifications of the two SWIR diodes.** Data was obtained from the respective data sheets available at `https://eblanaphotonics.com/products/1654-dm/` and `https://www.aerodiode.com/product/1650-nm-laser-diode/`, respectively.

be integrated easily into the LabView program by using the provided SDK. Despite the higher power output benefits, these diodes are still limited by the switching time. The diode tuning time was significantly slower than that of the previous lower power diodes, with Figure 6.4, highlighting as long as around 100 ms in the direction of off to on absorption for both diodes. The primary reason for this increased tuning time, in addition to batch to batch differences, will most likely be due to the much higher currents these diodes are tuned at, which are on the order of 500 mA compared to the approx 100 mA that the previous diodes required. Much higher currents can contribute to higher temperatures, which could cause the increased settling times. This long and finite switching time introduces an unfortunate downtime between the two consecutive images as before, meaning that, despite the potential to lower the exposure time, there can only be a comparatively less impressive increase in frame rate. The downtime between consecutive frames increases motion artefacts, and so would require strong filters to prevent false signals and present a clear gas distribution. The downtime can also be viewed from a perspective of wasting photons, if in order to generate one final false coloured visible image it takes two 6 ms exposures, a 50 ms down time for tuning the and a further 15 ms processing time for adding the SWIR data to the visible channel, then out of a total process time of 77 ms the shutter is only open for 12 ms meaning that $\approx 85\%$ of the SWIR photons are wasted. Whilst some of this can be recovered via parallelisation of the imaging process, the tuning time places hard limits on the system frame rate and photon waste. It was clear that to truly optimise the system a new approach to switching the illumination wavelength was required. These new Aerodiode diodes

were used in this new iteration of the system over continuing to use the Eblana Photonics diodes due to the higher optical power provided, despite the slower tuning time.

6.3.4 Fibre Switcher

In order to minimise movement issues a new solution was required, this led to the introduction of a fibre switcher (ThorLabs OSW12-1310-SM - MEMS 1x2 Fibre Optic Switch). These allow one to swap between two light paths by using the three ports in the system, as seen in Figure 6.5. Having the new diodes already fibre coupled made this an attractive option. The switch is a small trigger-able micro-electromechanical system (MEMS) mirror system that is designed to redirect light through the system and out of one of the ports with a quoted switching time of approximately 1 ms.

The switch can be used for gas detection by initially tuning the two diodes separately. The 55 mW diode was first tuned to the absorption wavelength of methane, whilst the 70 mW diode was instead tuned to a nearby wavelength and driven such that it outputs at a similar power to the tuned 55 mW diode. To estimate the wavelength of this off absorption wavelength a similar process to that performed in Chapter 2.2.1 can be followed, if the same approximations are made. The new diode drivers can be directly interfaced with and provided exact output currents without the requirement for converting input voltage as before. The diode is considered tuned to the absorption line at a temperature of 48.5 °C and a drive current of 319 mA, whereas the off absorption has a drive current of 306 mA, a difference of 13 mA. The aerodiode diodes have a quoted current tuning parameter of $0.003\,\mathrm{nm\,mA^{-1}}$ (as found at `https://www.aerodiode.com/product/1650-nm-laser-diode/`). This results in a wavelength tuning of 0.039 nm or 4.29 GHz. The reason for tuning the 55 mW and not the 70 mW diode was because it means that if, once tuned, the more powerful diode emitted at a power that the lower power diode could not reach there would not be a large difference between the brightness of the on and off frames. A potential difference in brightness as a result of this on/off power differential would introduce false signals that would need to be accounted for with a threshold - potentially decreasing the sensitivity of the system greatly. Once initially tuned these diodes were then not dynamically tuned as before, so once tuned they emit only at the one

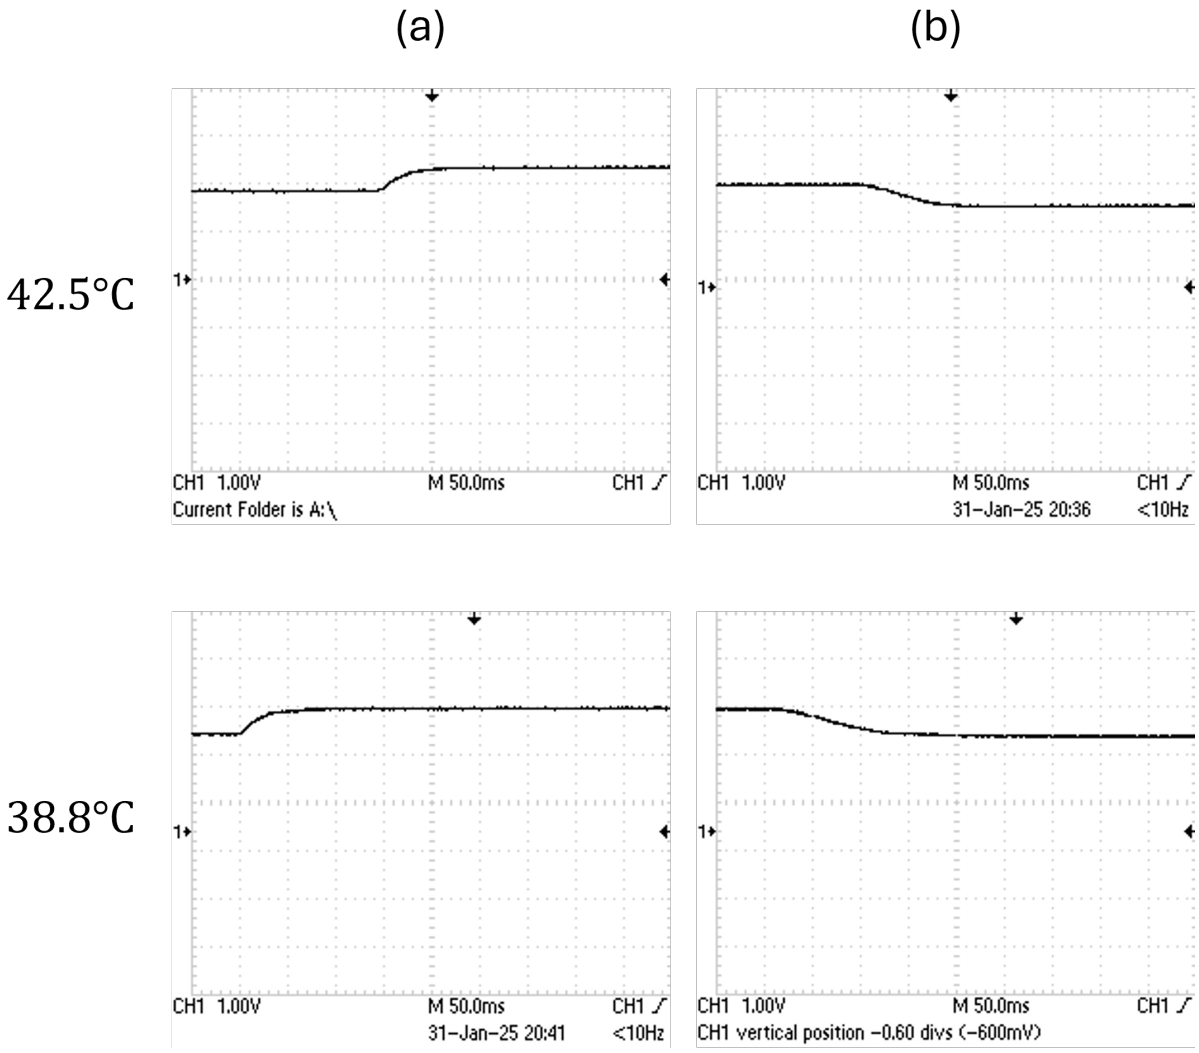

Figure 6.4: **A demonstration of the tuning time required by the two new diodes, for (a) tuning on to off resonance and (b) tuning off to on resonance.** It can be clearly observed that for both the diode running at 42.5°C and the diode running at 38.8°C behave very similarly in terms of tuning response. Each division along the x-axis in these plots corresponds to 50 ms. These plots show approximately a 50 ms tuning time from on to off resonance and an approximately 100 ms tuning time required for off to on resonance.

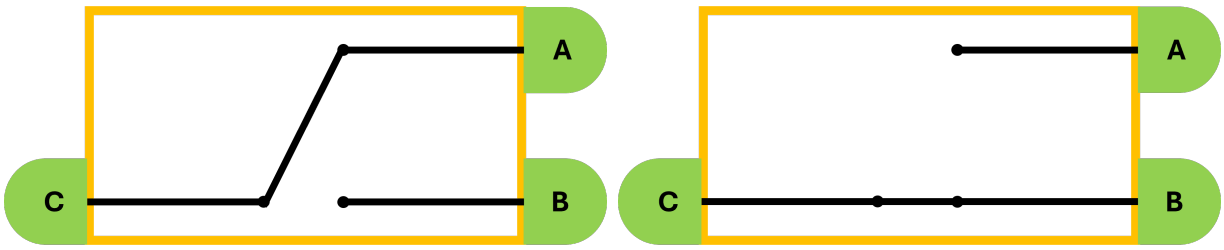

(a) (b)

Figure 6.5: **Figure indicating the two states of the optical switch, (a) that passes light from input A to input C, and (b) that passes the light from input B to input C.** An electronic trigger signal causes the state to change.

wavelength. The respective fibres were then coupled into the switch allowing the mirror to act as a selector between the on and off absorption wavelengths. This process means that the limiting factor was now the switching time and not the tuning time.

The optical switch can be triggered using a pulsed 5 V signal. The board level version of the switch swapped the mirror position upon receiving a falling edge pulse. The camera can be used to provide this trigger. There are many ways in which the camera can send a trigger, in this case the transistor-transistor logic (TTL) pin on the camera was used to send the signal to the switch. There are a variety of trigger types that the camera can send but the one initially utilised here was an inverted exposure trigger - which was a signal that was low when the camera was active and high when it was not. The trigger will fail to work if the exposure time and frame rate of the camera are incompatible, i.e. if the exposure time is less than the set frame rate allows the images will be taken with the correct exposure time but there is a higher risk of the trigger activating twice in a quick succession, causing the direction of subtraction between consecutive frames to get out of sync. This change in direction results in an inverted gas image, i.e. where the regions of gas are clear but elsewhere is green, an example of what this can look like is shown in Figure 6.6. In general, and unless otherwise stated, a compatible frame rate of 100 fps and exposure time of 8 ms were used, effectively adding a 2 ms downtime between each consecutive image capture.

Similarly to the issue of the loss of synchronisation between mirror position and subtraction direction was the fact that the initial starting position of the mirror was simply the last position of the mirror before stopping. In order to overcome this and to maintain the correct directionality

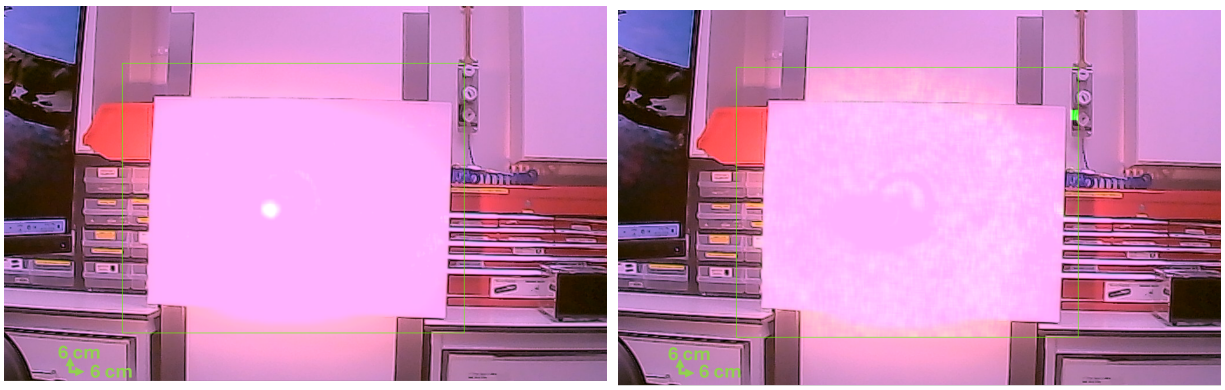

(a) (b)

Figure 6.6: **Comparison between camera response for when the direction of subtraction is (a) the correct direction and (b) the inverse direction.** When the direction of subtraction is incorrect only the noise that was thresholded out is observed and there is no gas signature.

throughout the imaging process the direction of subtraction was instead defined by the slight difference in brightness between the on and off absorption images. The brightness difference is very minor between the two states, but is consistent enough that it could be detected and utilised effectively here even in scenarios where the camera was in transit.

The old system illuminated symmetrically with similar diodes on each side, this allowed an even illumination profile which does not misalign with distance. For this even profile to be achieved in the new system, the light exiting the optical switch needed to be split and coupled into the two lens tubes on either side of the aperture as before. The splitting and coupling was implemented using a 50:50, 1×2 fibre coupler (ThorLabs TW1650R5A1) between the switch and lens tubes. This coupler allows equal ratios of the on/off resonance light to be coupled into the lens tubes on either side of the aperture mimicking the parallax solution of before. A schematic of the optical components of the system are presented in Figure 6.7.

6.3.5 Parallelisation

The overall triggering system streamlines the imaging pipeline substantially, because the switch timing is completely controlled by the camera, the code does not acknowledge the presence of the switch in any way. As a result of the camera triggering the switch there can be a higher level of parallelisation implemented than before, reducing photon waste and increasing system frame rate. A producer-consumer type algorithm can be implemented such that the producer loop

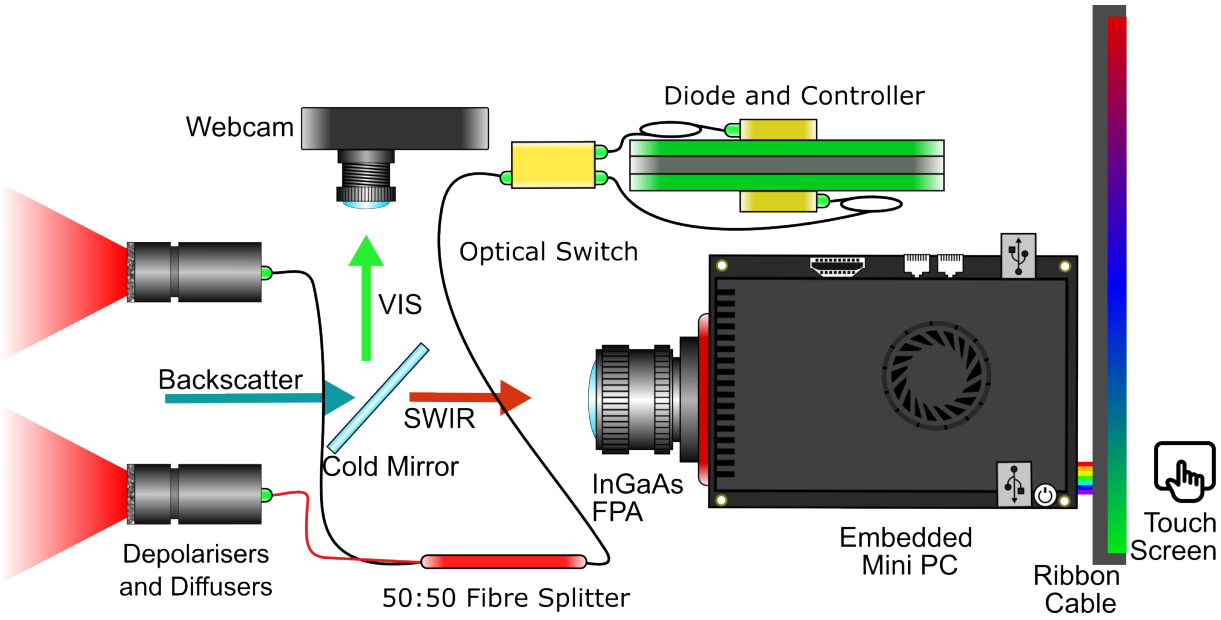

Figure 6.7: **The SWIR optical components of the redesigned system.** The diode tubing containing the collimating lens and diffusers, as well as the SWIR FPA maintain their original positions, with the depolarisers added to the external side of the diffuser in each tube. The diode boards, cooling boards, fibre splitter and optical switch, as well as the path of the fibre, highlight the pathway of the diode light through the system. Also highlights the position of the Mini PC and touch screen, the mini PC has the new position of being placed above the FPA. The cold mirror is also brought over, for figure simplicity the fibres appear to pass through the beam path, however in reality the beam bath is carefully enclosed.

produces unprocessed differential SWIR images and the consumer loop takes visible images as well as processing the SWIR images in order to be overlaid.

In order to fully incorporate this new pipeline, it was decided that it may be necessary to implement a new, more capable mini-PC. The mini-PC that was selected was a newer model of the LattePanda, known as the LattePanda Sigma. This upgrade provided a much more capable CPU, upgrading from 2 cores to 12. The increase in number of cores led to a higher degree of flexibility and allows a more consistent processing time to be maintained by the system at all times. The processing power offered by the increased specifications of this mini-PC are what allows the implementation of the new algorithm. The new model also offers both an extra USB-C port and an extra USB-A slot. The main drawback related to changing to this new PC is the increased size. The previous system, the LattePanda Alpha 864s, was 115 mm × 78 mm, whereas the Sigma is 146 mm × 102 mm. In order to facilitate this increase in size the placement of the PC needed to be reconsidered. The overall solution to this issue was to make use of the empty vertical space that was present in the previous system by designing a 3D printed table-like mount for the PC to sit above the SWIR camera. The raised height of the LattePanda allowed for the system battery to be moved underneath, granting a higher degree of flexibility in regards to part placement throughout. The Sigma model is also substantially heavier increasing to 1.8 kg from the 0.7 kg of the Alpha model. The system also requires much more power to run, requiring between 12 V and 20 V to operate, compared to the between 7.4 V and 15 V of the alpha model. These trade offs in weight and power consumption were considered satisfactory given the increased computational power afforded.

The producer loop acts far more simply than the previous systems capture loop, but mirrors the original imaging algorithm written by the previous PhD student as summarised in Figure 2.1. The loop captures a SWIR image and saves it to a memory location A, this image then has the SWIR image taken in the previous iteration, which is saved in the memory location B, subtracted from it, which in turn is saved in memory location D. The image in memory location B is then overwritten with the image in memory location A and the capture loop begins again. The direction of subtraction is swapped each iteration to maintain the directionality of the gas signal, which as mentioned previously is dictated by the slight brightness imbalance between

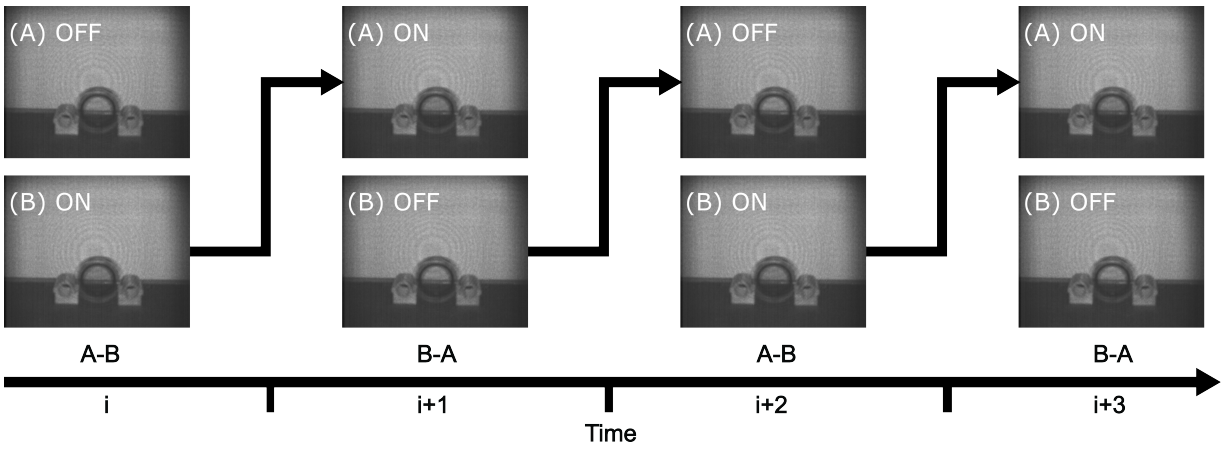

Figure 6.8: **A diagram showing the path of SWIR images in the imaging algorithm of the new system.** Also highlighted is the changing direction of on and off absorption states of the images and the requirement to change the direction of subtraction each iteration.

the two states. The movement of images with respect to memory location, absorption state and direction of subtraction is summarised in Figure 6.8. The image in memory location D is first low pass filtered using a kernel of size 11×11 and then converted to an array. This array is then passed into the visible capture and image processing loop.

As the image processing loop takes longer to run than the SWIR capture and subtraction loop, it only processes the differential image present in the memory location, D, at the beginning of each loop, meaning that it misses many frames that are generated during the course of an iteration. One of the larger changes made to the operation of the system is the reintroduction of the temporal filtering, which aims to counter the loss of these frames. Missing frames leads to photon loss, and matching the iteration time of the two loops would introduce considerable downtime to the SWIR capture loop, again leading to photon loss. Implementing temporal filtering means that a "memory" of the frames missed by the slower loop is maintained and carries into the image processing. Temporal filtering, as explained in Section 3.6, means that the latest frame contributes to only a fraction of the overall signal for that iteration. As a result of the increased frame rate, the temporal filter can be used far more effectively than in the previous system and the most current frame could make up a comparatively lower percentage of the current frame, with fractions as low as 0.01 giving reliable results.

The introduction of the temporal smoothing truly unlocks the potential of the parallelisation of the processing. By thinking of this from a point of view of photon loss, this imaging technique

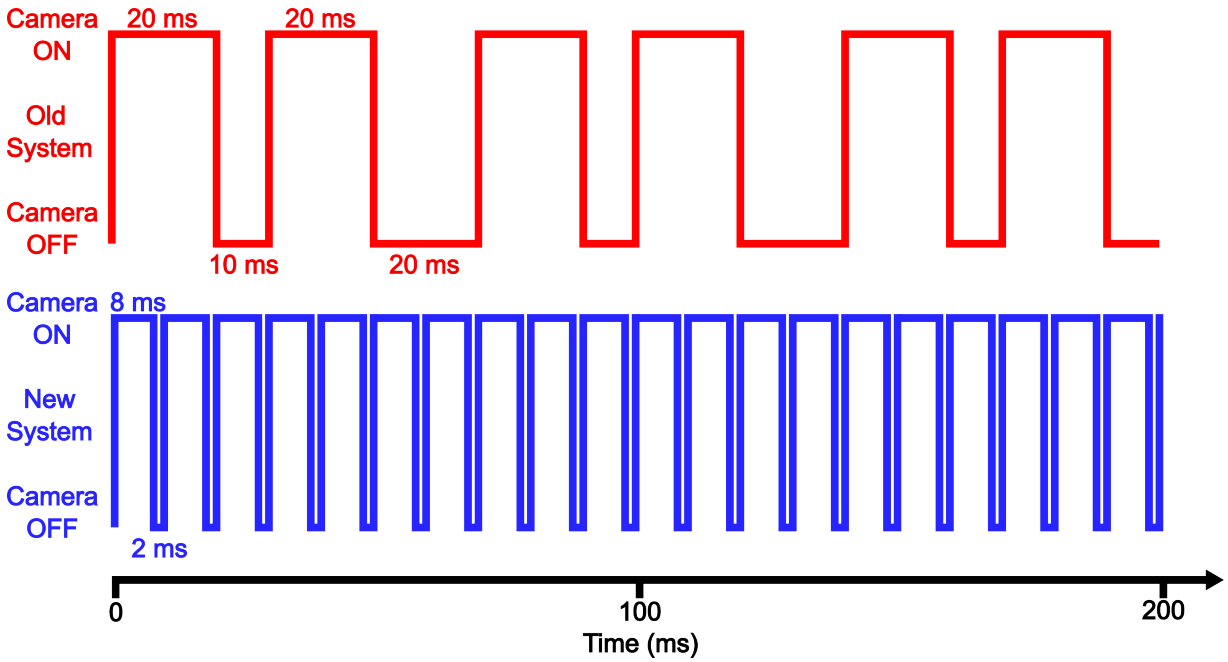

Figure 6.9: **A comparison between the overall active time of the SWIR camera between the new and old system.** It can be seen that over the course of 200 ms, the old system will take only 6 SWIR images, whereas the new system will take 20. The SWIR camera is only capturing for 57% of the time in the old iteration, whereas the new system is capturing for 80%.

means that there is far less downtime in the system on the whole, with the only example of this currently is in the small window introduced in order to properly space the trigger response. If a frame rate of 10 fps is used with an exposure time of 8 ms then the camera will be active 80% of the time, increasing it's potential photon budget. In the previous system a worst case example of this would be a loop that takes at least 70 ms with the SWIR camera only active for two 20 ms exposures, this means the camera is collecting light around 57% of the time, meaning the new method provides a roughly $1.4\times$ increase in potential signal capture. A comparison of the active time of the SWIR camera in both systems is shown in Figure 6.9. The visible capture and production loop is limited by the frame rate of the visible camera which is set to 30 fps, depending on the processing the overall iteration time can vary but is generally found to be approximately 35 ms. As the visible capture and processing loop acts as the consumer and produces the combined visible and SWIR images, the speed of this loop is what dictates the speed of the feed on the camera display. This means that the feed will show ≈ 30 fps but will, as a result of the temporal filtering, contain 100 fps worth of gas signature.

6.3.6 Size, Weight, Power and Cost

Another key design challenge considered when designing the new system was to maintain the compact size and low mass. Size was maintained to ensure easy integration with the tripod or the drone used with the previous system. In order to achieve this, it was decided to maintain the same dimensions as the previous system, with the same sized base plate, faceplates, touch screen and lid, each edited slightly to ensure proper placement of the new components and to allow sufficient cooling. The 14.4 V, 5.2 A h battery used is also the same, and can be charged using the same 24 V charger as before. Due to the introduction of new components the battery life of the new system is now 1.5 h, suggesting a current draw of approximately 3.5 A. The majority of the 3D printed internal mountings had to be redesigned in order to maintain this form factor, often requiring inventive placements and considerations for fibres. The result was a new system that has the same dimensions as the previous system, and has a mass of 3.878 kg, which is approximately 0.7 kg more than the previous system. This mass is still safely below the 5 kg maximum payload size of the drone. The only change to the overall form factor is that due to the inclusion of the extra depolarising optics in the lens tubes, discussed later in Chapter 7.5. These extra optics cause the tubes to poke out from the face plate by an extra centimetre, though it is possible with some internal rearrangement to further sink these into the system to prevent this. Images of the fully assembled system are shown in Figure 6.10. In terms of cost the new version uses the same camera and the majority of the same off the shelf and 3D printed components. The additional costs come from the new diodes and the fibre switcher, with the new diodes costing £2,500 each and the fibre switcher costing approximately £1,000. Additional costs include the new driver boards, which cost £1,500 each, meaning that the new system costs approximately £7,000 more.

As with the previous system, the complete optical system was made to comply with the an ocular hazard classification of class 1. Whilst the diodes are now capable of much more power output, and as such are now classified as class 3B lasers, they still do not reach cross the upper threshold of equivalent to a collimated 10 mW beam, this is again largely due to the diffusers. The diodes and diffusers are once again inaccessible without opening the system and in this case the diffuser is placed behind two depolarisers adding extra length in which to

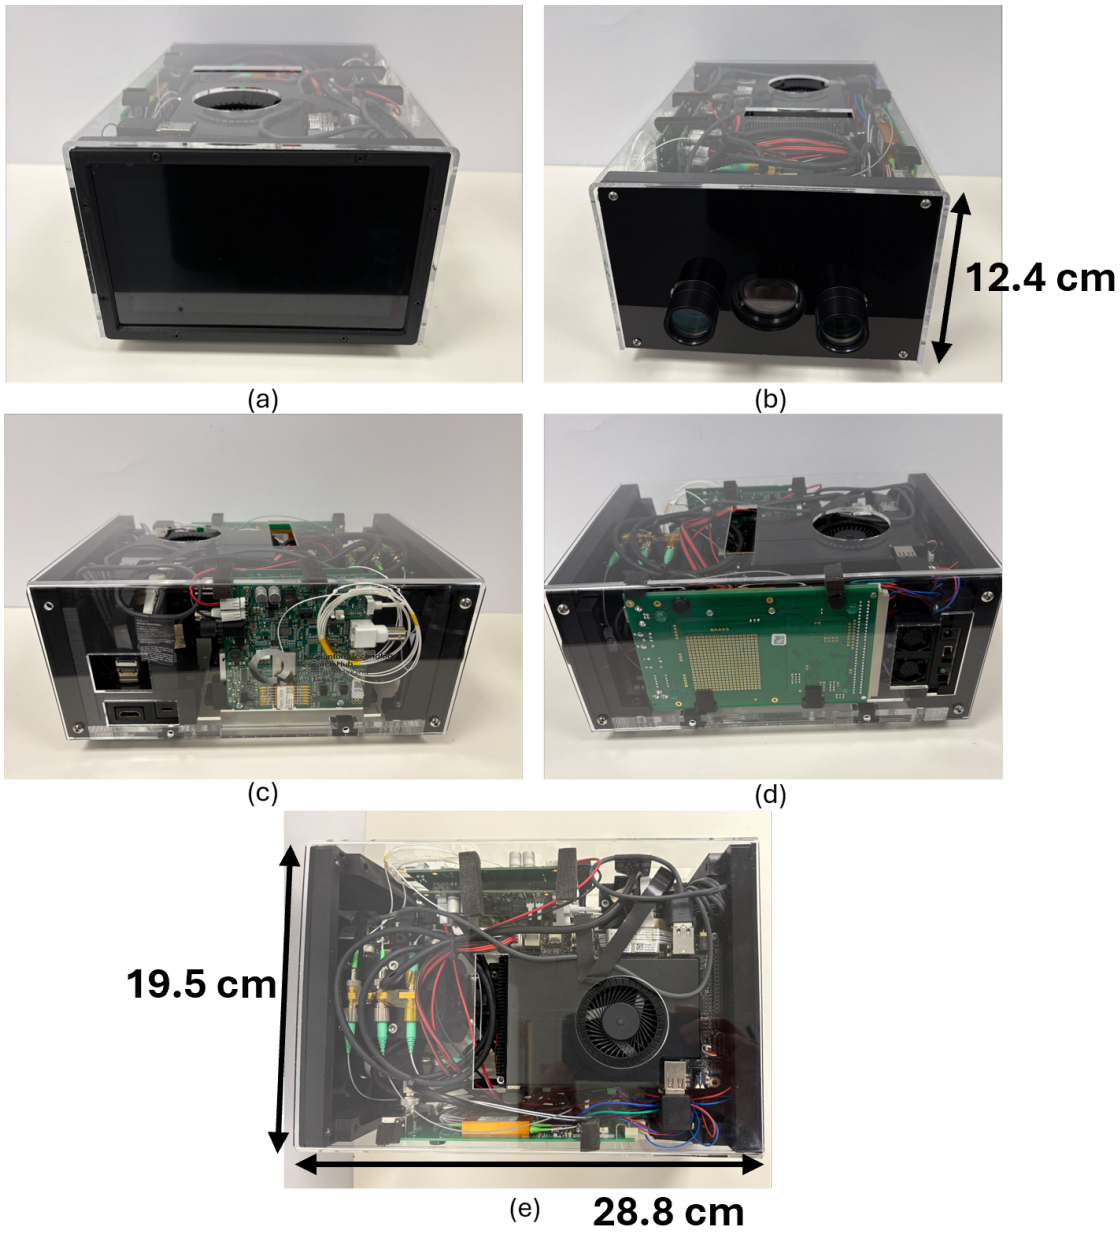

Figure 6.10: **Images of the fully assembled new gas detection system from (a) behind, (b) the front, (c) the right side, (d) the left side and (e) above.** This image reveals details such at the large switch board and on/off switch visible at the side of the system in (d), the HDMI port, a fan vent, the fibre coupled diodes and their driver boards in (c), the touch screen in (a), the diode outputs and collection optics in (b) and the micro PC and heat vents in (e).

diffuse the beam further, reducing ocular hazard. At the diffuser exit a power meter reading read 9.53 mW, placing it very close to the upper limit of the safety classification, albeit this is without considering the impact of the depolarisers. The system was verified and signed off as class 1 by the laser safety officer.

6.4 Example Data

In order to demonstrate the performance of the system, it was set up to film the ignition of natural gas from a gas hob. A small amount of natural gas was released before ignition. Domestic gas pressure is approximately 21 mbar and the range at which the system detected this was 1.27 m [168]. Three frames from this test are presented in Figure 6.11, which shows the hob before releasing gas, during gas release and after the hob has been ignited. It can be seen that gas is clearly detected from the hob at this range. The settings used in the reconstruction of this image were a temporal smoothing fraction of 0.01, a smoothing kernel of size 12×12, a persistency filter strength of 2, a 5×5 circular erosion kernel with zeroes in the corners (as in Figure 3.14), and a re-binned SWIR image size of 616×464. A black card backing sheet is used in order to offer a uniform back scatter. These results highlight the systems sensitivity, even within a domestic setting. It must be noted that during observation the synchronisation of the direction of subtraction was lost a couple of times resulting potentially in a loss of sensitivity in a few frames, this was manually corrected in post processing.

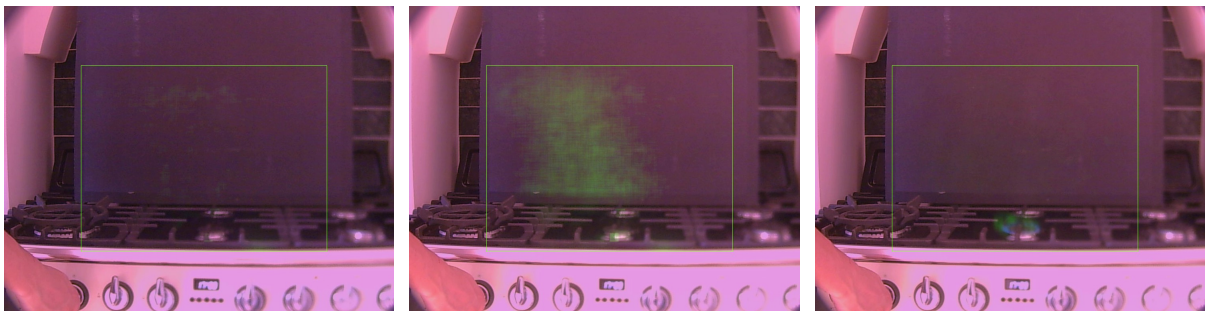

(a) (b) (c)

Figure 6.11: **Example frames taken during the observation of the gas hob, (a) was taken before gas release, (b) during release and (c) after ignition.** The gas signature is clearly observed in (b) and is not detected after ignition, highlighting the fact that the concentration of gas has been dispersed quickly after ignition.

6.5 Conclusion

Overall, this chapter highlighted and discussed the reasoning behind the hardware improvements that were made during the course of the PhD. Whilst relatively minor changes were made to the existing system by replacing the diodes with slightly more powerful ones, the more important decisions relate to the building of the new system. The new system was designed and built with the shortcomings of the previous iteration in mind with the fundamental limits outlined in the previous chapters, specifically relating to the maximum frame rate of the system, were front of mind from the outset. With this design the requirement for long waits in the image timeline could be overcome, allowing for more continuous imaging overall. Continuous imaging allowed for the reintroduction of temporal filtering in order to achieve effects that were previously less effective. Whilst much of the overarching design was reused, such as the overall form factor and footprint, many of the new parts introduced have increased sizes and different geometries. These new shapes meant that new mounting designs and layouts were experimented with in order to achieve the target size. The completed system is as such very similar in size and weight, but with improvements in power and frame rate. Finally, the system was demonstrated to work at a moderate range to detect gas release from a domestic gas hob, providing insight into its sensitivity and capabilities.

Chapter 7

Device Comparisons

7.1 Introduction

Having completed the redesign of the gas camera it came time to compare the two systems in order to quantify the improvements achieved. There are three primary categories that the system was designed to improve upon; detection range, frame rate and sunlight rejection. An advantage in any of these categories could potentially greatly improve the usability of the system as it would be able to work over longer distances, with higher flexibility of movement, location and time of observation. A further point of comparison is that of polarisation sensitivity, which as construction neared completion became apparent as a difference between the capabilities of the two systems. An image comparing the overall appearance of the two systems is present in Figure 7.1.

7.2 Frame Rate

One of the primary advantages of the new system was its increased frame rate. As the system was built to tackle sunlight noise, lowering the exposure time was of great concern from the outset. The more powerful diodes are what allowed for an increase in frame rate, as this enabled a higher amount of light captured per unit time than before. More light means that the system can achieve similar levels of signal for much lower exposure times. As the laser power was

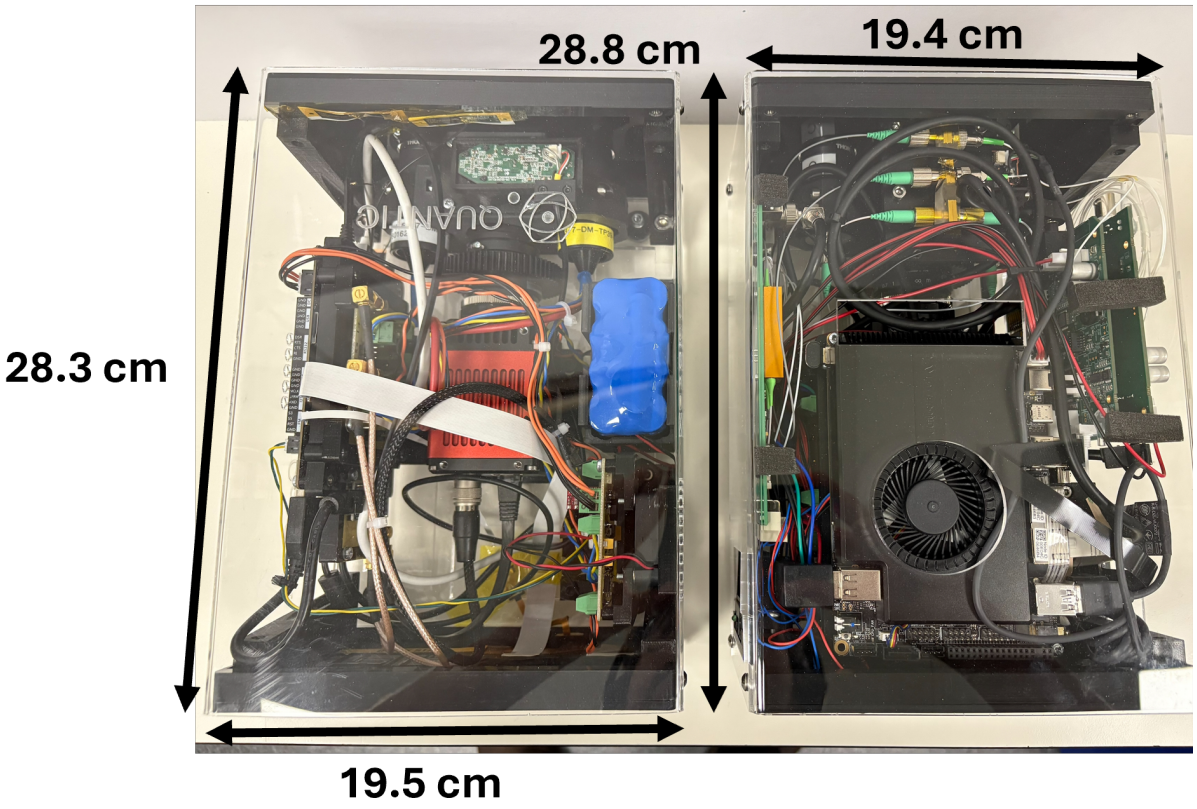

Figure 7.1: **Image showing the two gas imaging devices with the older system on the left and the newer system on the right.** Whilst many of the individual components within the system differ the overall form factors are the same. The overall size and footprint of the two systems are the same.

approximately 50 mW at each diode before being 50:50 split, it means that the diodes were at least 25 mW at each exit point. This power increase of around 2.5×, meant that the exposure time could be approximately halved, as long as the decrease in exposure time is proportional to the increase in power there would be no loss of range experienced by the system overall. A benefit of decreasing the exposure time was that the image to image downtime will be decreased and, as a result, the differences present between the two images will be diminished thereby reducing the effects of camera motion, or movement within the scene.

7.2.1 Triggering

The improved frame rate of the system was mainly achieved by the introduction of the MEMS switcher. The new laser diodes were more difficult to tune on and off absorption with there being no reliable way to do this via the LabView SDK. As seen in Section 6.3.3, it can be seen that in the worst case scenario the tuning time for the new diodes was as much as 100 ms which is slower than the previous system. By using the switcher one can effectively decrease this downtime to just 1 ms, as seen in Figure 7.2. While this 1 ms switch is fast, there was no requirement to add a wait in the system for this to trigger. As explained when introducing the switch in Section 6.3.4, it could be triggered using the SWIR camera. The ability to trigger meant that the camera could change the switch state of the mirror based on when it was exposing.

As mentioned in Section 6.3.4, because of the triggering mechanism there had to be a compatible frame rate and exposure time combination that prevented the trigger signal from being missed or recorded twice, and as such throwing the synchronisation of the system out. The potential to lose synchronisation does set another fundamental limit on the system as there had to be enough exposure time to include the dip whilst also giving enough space temporally between triggers. Too high a frame rate caused synchronisation issues, whereas too low an exposure time can mean that signal was lost in the switch. In this case too high was found via trial and error to be exceeding 8 ms, as this causes indistinct trigger peaks that could often be missed. Too low here again would be anything less than what provides a stable trigger signal with no missed or false triggering. For example 6 ms would be considered too low for a 10 fps system if 8 ms provided a stable response as it could be seen to waste 2 ms of potential data collection. The trigger

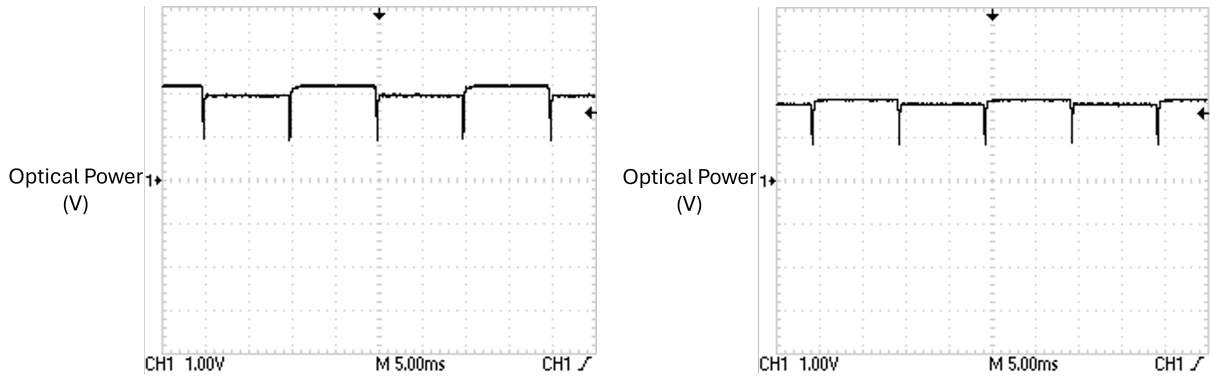

(a) (b)

Figure 7.2: **Images highlighting the mirror switching time for (a) the left diode of the new device and (b) the right diode.** The x-axis here denotes one division as 5 ms. The y axis is the optical power, in volts, measured from the diode as detected on a SWIR photodiode detector. The switching time here is very quick taking approximately 1 ms, as quoted on the product. Also of note, and as to be expected, there is no difference between either sides diode in terms of switching time. The large dips here are a result of the mirror being in transit between reflecting the response of the two diodes and acts to block light temporarily.

signal for a frame rate of 100 fps and an exposure time of 8 ms is shown in Figure 7.3. The 2 ms downtime gave enough space to separate the trigger signals, and acted to prevent overlapping signals, thus improving synchronisation.

The form of triggering highlighted in figure 7.3 was implemented initially and unchanged for much of the construction and testing. This triggering regime was settled upon early in the camera development, before it was clear that the trigger outputs needed to be so regularly spaced with a compatible frame rate. Upon revisiting the concept, it was noticed later in the development that there was a more optimal way to run the trigger. In Figure 7.3 it could be seen that the falling edge occurs when the camera begins its exposure, this meant that for the first few moment of the exposure the mirror was actually in transit and either the wrong wavelength or no light is being produced by the diodes. As seen in Figure 7.2, it can be seen that, relative to the overall length of the maximum 10 ms period allowed by the frame rate, the mirror transit only accounted for a very small fraction of this time. The 8 ms exposure time, however was slightly shorter than the 10 ms window of the frame rate, meaning that, marginally, the transit time will take over a larger fraction of the camera's active time. In the worst case the switching taking 1 ms will account for one eighth of the exposure time. To counter this loss, the newly introduced 2 ms downtime to allow the triggers to be spaced can be utilised. By non-inverting the signal as demonstrated

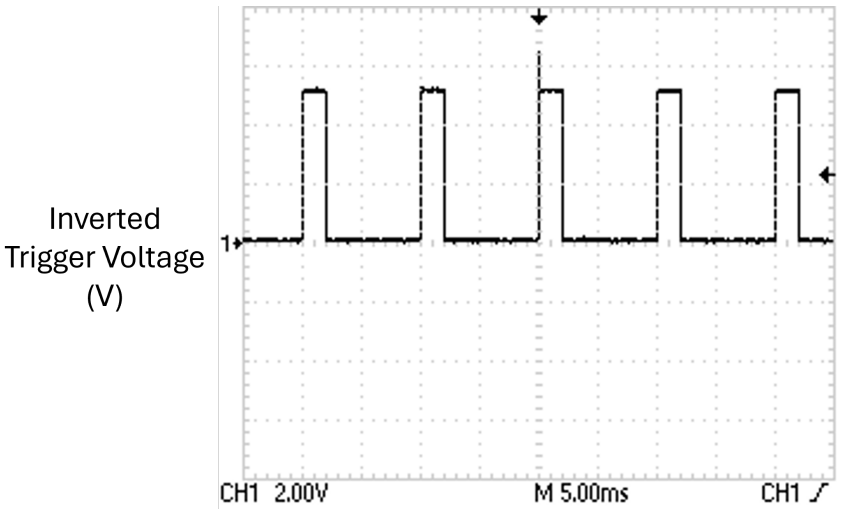

Figure 7.3: **Initial inverted trigger signal from camera used for triggering the optical switch.** This trigger signal features a 0 V signal when the camera is exposing, a 5 V signal when the camera has stopped. This means that the falling edge signal occurs at the beginning of the exposure. The frame rate of the camera is 10 fps, and the exposure time is 8 ms. One division along the x-axis is 5 ms.

in Figure 7.4 the falling edge can instead fall at the end of the exposure. The placement of the falling edge meant that the mirror transit time could be fully enclosed within the camera downtime and be in the correct state for the beginning of the next exposure. By using this imaging regime light was used more efficiently, as light could be captured for the full duration of the image acquisition.

The effects of this change on the system was first explored in terms of overall brightness of the images taken. In order to test the brightness, images were taken in the same location using both trigger approaches. The pixels in the raw SWIR images were summed over and the average for the 435 frames taken was calculated. By comparing the average total brightness of the two frames it was calculated that the images taken with the normal, non-inverted trigger as in Figure 7.4, were approximately 1% brighter. Whilst the images were brighter, and confirm that the non-inverted trigger is correct, the improvement in brightness is fairly minimal. From the perspective of photon waste, however the non-inverted trigger is preferable but the overall impact upon the data is unlikely to make much of a difference and will not invalidate any data taken before the swap in trigger direction.

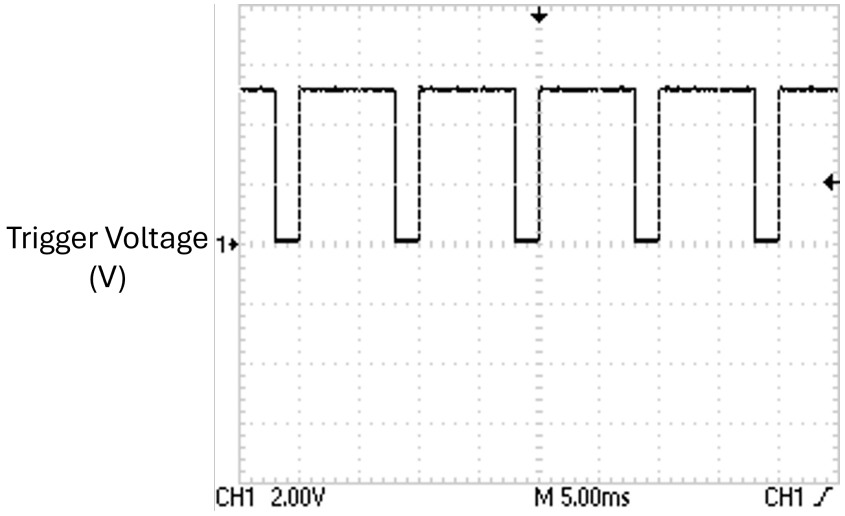

Figure 7.4: **Final non-inverted camera trigger signal used to change the optical switch position.** This trigger signal features a 5 V signal when the camera is exposing, a 0 V signal when the camera has stopped. Here the falling edge signal occurs after the camera exposure has finished, meaning the switch can occur in the camera downtime. The frame rate of the camera is 10 fps, and the exposure time is 8 ms. One division along the x-axis is 5 ms.

7.2.2 Compact Switch

During development, it was convenient to utilise the development board for controlling the MEMS switcher. The development board within the system had a large size, taking up almost a whole side of the device. On the board the only component required to run the system was the switch itself, the rest was merely a simplified interface and contains monitoring information that was useful in set up but less so once the system was enclosed and running. The large space designated to fibre winding was used effectively, however. The switch could be purchased individually and run on a custom PCB as the switch only consisted of 3 pins. One was power in and one was ground, with the third being the trigger line, meaning that the camera trigger was still applicable to control the switch. The trigger behaved differently on the switch itself as opposed to the board and instead would be in one state for a high signal (5 V) and the other state for a low signal (0 V), this means that the previous trigger would not work here as one exposure consists of the same low and high signal each time, meaning that no flip will occur between subsequent frames.

The solution to the lack of a switch could be the introduction of a flip-flop logic circuit. The most likely candidate circuit for this would something similar to a "divide by 2" type flip-

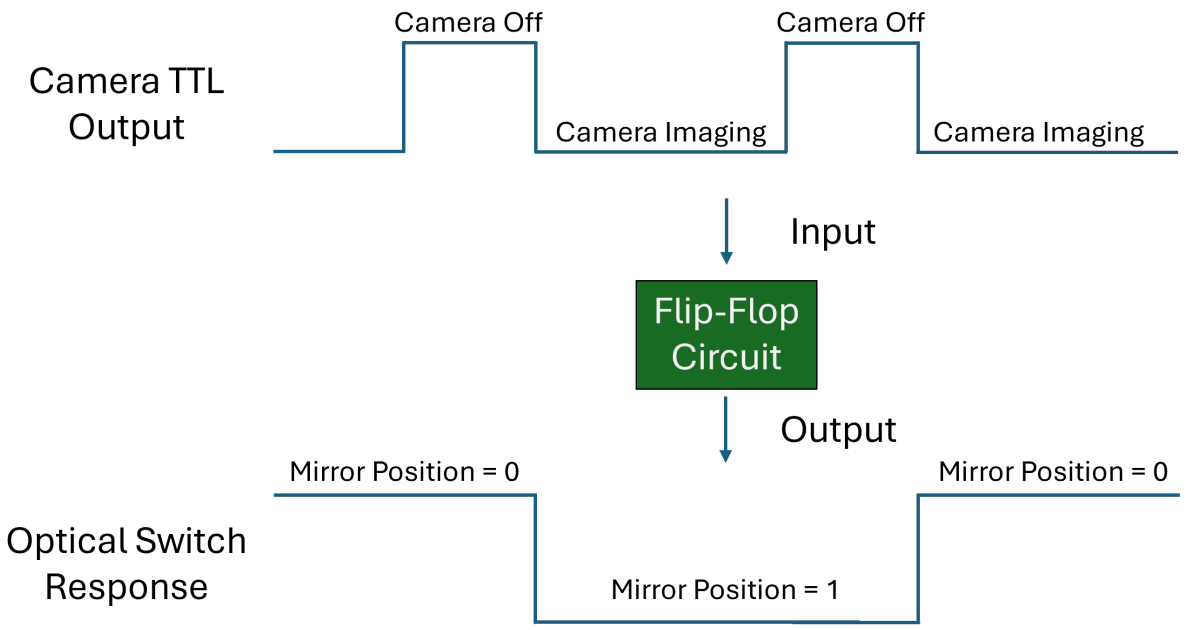

Figure 7.5: **A simplified example of how a potential divide by two flip-flop circuit would convert the TTL output signal to a format readable by the by the optical switch module.**

flop [169, 170]. This kind of flip-flop essentially will toggle an output from high to low upon reception of a rising edge, for a regular repeating signal, with frequency f, this means that the output would also be a square wave, but with frequency $\frac{f}{2}$, hence the name "divide by two". This flip flop would then convert the trigger input into a signal that means the mirror position changes when required, with each iteration. This level of control also has the added benefit of ensuring the known position of the mirror at any one time, because one can ensure for example that the input signal is low or high initially. An example of how this circuit will work is shown in Figure 7.5.

7.3 Range

In order to compare the range of the systems the two cameras were used in order to image a glass gas cell of known concentration from an increasing distance. This cell was made up in order to contain a concentration with an upper limit of 5000 ppm.m. By using the same signal and noise calculation program used in Section 5.5, relative signals could be found for these balloons. It was shown during data collection for the paper "Characterising the performance of a drone-mounted real-time methane imaging system" [97] and presented in Figure 5.10, that

for the older iteration of the system working with an exposure time of 20 ms the gas cell still returned an intermittent signal at ranges as high as 17.3 m, though it must be noted that the noise and signal at such a range are comparable in size and as such the signal flickers in and out at this range. The size of the cell on the SWIR camera feed here is also quite small, though seemingly not such that this is the limiting factor in the detection criteria.

In order to compare the range capabilities of the new system, the cell was once again used and the two devices were tested in the corridor outside of the laboratory. Sets of images were taken for 6 distances which had already been concluded to be suitable for detection by the older system as in Figure 5.10. Each gas camera was mounted on a tripod and placed at a distance from the cell, which was measured using a handheld laser range measuring device. After each run of data the two cameras were swapped so that the distance for both was the same. The aim of this investigation was to see if the range of the new system was greater or less than the older system by observing whether the return signal at each of these distances was more or less than the old systems response.

Fundamentally, there are quite a few differences between the two systems, many of which have already been discussed, but in order to process the saved data some changes were made. Firstly was the change to the imaging algorithm that reuses the previous iterations SWIR image as seen in Figure 6.8, as well as the system being inherently noisier and as such utilising a larger 7×7 erosion kernel. This increased noise may be the result of speckle. Whilst in the previous version of the system the same laser was tuned on and off, in this case two different lasers are used which may have different speckle patterns which can be flagged as different between the images. The fact that such speckle related differences pass the smoothing suggest that the regions of speckle are larger in size, hence the larger erosion kernel. The increase in size may be due to the fibre coupled nature of the diodes, as this can introduce speckle. The speckle introduced by the fibre can also vary as the bends in the fibre move or change. In future the speckle could be overcome by using a larger smoothing kernel, as in Figure 6.11, but this may cause issues for signals of the same size as the speckle being washed out in the worst cases. The most obvious difference was the fact that the exposure time of the new system was only 8 ms as compared to the 20 ms of the old system. The results of this investigation are shown in Figure

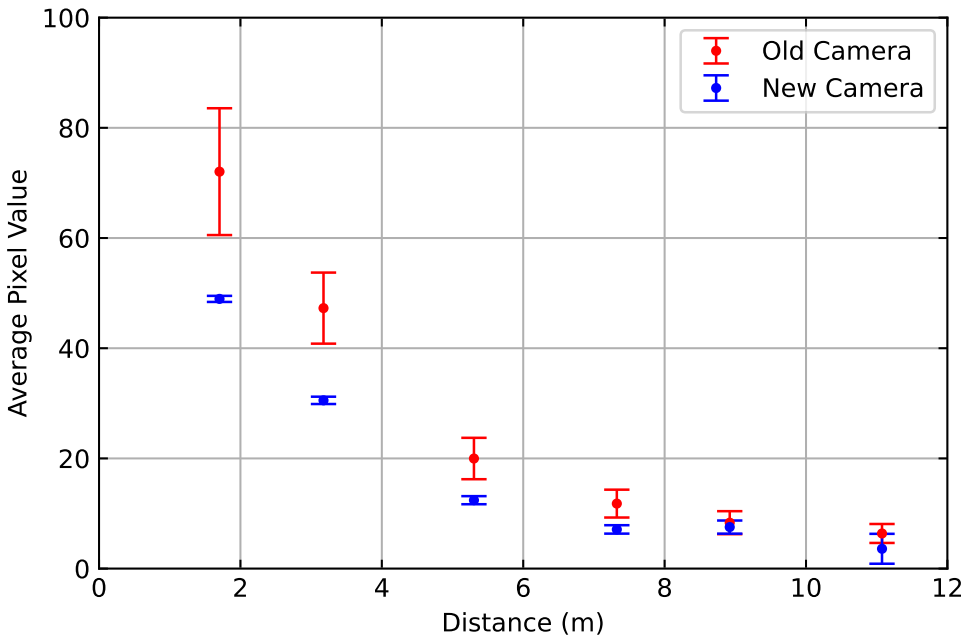

Figure 7.6: **Comparison of the return signal between the old and new camera systems as a function of distance.** It is clear that whilst comparable, the return signal calculated for the newer device is on the whole, lower than that of the older version. This lower signal suggests that the overall range of the newer system is less. Another interesting point to note is that the uncertainty of the newer system is much less, and in facts gets smaller with increasing range, which is the exact opposite of what is observed in the data from the older device.

7.6.

The first takeaway from this data is that, on the whole, the new system receives less return signal than the old system. Whilst at first this is surprising, it must be balanced with the fact that, despite the large increase in output power, the exposure time is less than half of what the old system is capable of. This comparison can be taken further by comparing the SNR of the two systems as a function of distance as calculated in 5.5, this can be seen more clearly in Figure 7.7. From Figure 7.7, it can be seen that according to these models the range should be very similar despite the differences in exposure time and diode power, with a comparable range of approx 4 m before reaching an SNR of 1. The figure, however, also supports the idea of the older system having an higher SNR with an SNR exceeding 80 at a distance of 1 meter, compared to the new camera having an SNR just shy of 60 at the same range. This signal could be recovered by increasing the exposure time but this could introduce more noise into the system as well as a trade off in terms of system frame rate. The new system also begins

to lose the signal around the 11 m mark, as the gas signal flashes in and out, as it did for the older system around 17 m, suggesting that the range is slightly less, though still above 10 m. This flashing and lack of consistent signal could also be observed in Figure 7.6 from the size of the deviation relative to the signal magnitude, as the uncertainty bar approaches the same size as the signal. The increased noise in the new system also may act to lower the average signal somewhat, but the low standard deviation would suggest that this impact is low. Finally, what was perhaps most noticeable was the fact that the deviation on the return signal for the new system was by far less than that of the old system, even despite the noisier images. The primary reason for this appears to stem from the fact that the size of the differential signal does not vary much frame to frame. This consistency could stem from two sources. Firstly, was the increased frame rate of the system meaning that, because more frames are captured in a short time, environmental changes have less of an overall impact, though it would be envisioned that this would only have a very minor impact. It is more likely that this is a result of the removal of tuning times from the system, meaning that the swapping was less likely to capture light tuned to the wrong wavelength or capturing light during the transition from on to off resonance. This improvement in stability may also suggest that, as shown in Figure 3.3, that the old system was using a suboptimal wait time, something that has since been verified. So, whilst the range results show that the new system lacks the overall sensitivity of the old system with the current settings, it does highlight that the new systems switching mechanism results in far more consistent on/off images, improving the reliability of the measurements.

7.4 Sunlight Rejection

One of the key reasons for building the new system was to overcome the sunlight sensitivity of the previous system by avoiding saturation of the SWIR camera. The pathway to increasing the sunlight rejection is closely tied to that of the increased frame rate of the system. By using lower exposure times the camera should accept less sunlight, but because the laser power is now higher, a similar level of signal can be maintained, providing methane sensitivity but with less saturation.

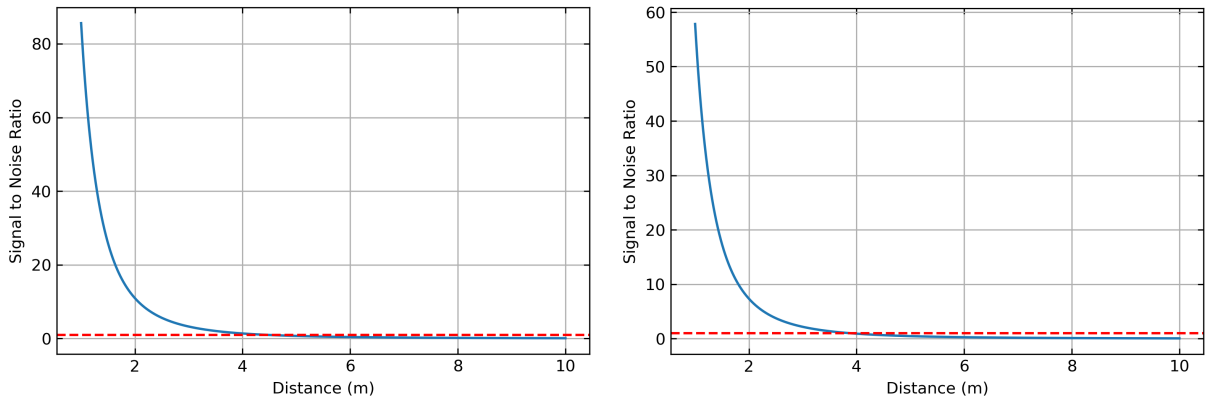

(a) (b)

Figure 7.7: **Comparison between the SNR as a function of distance for (a) the old system and (b) the new system.** A red line indicating a SNR of 1 is included for clarity. Here the distance ranges from 1 m to 10 m. In (a) a reflectivity of 0.85, an aperture radius of 12.7 mm, individual diode power of 10 mW and exposure time of 33 ms were used. In (b) a reflectivity of 0.85, an aperture radius of 12.7 mm, individual diode power of 27.5 mW and exposure time of 8 ms were used.

7.4.1 Old vs. New

It can be very difficult to compare the two systems in terms of sunlight, one has no control over the weather and the brightness of the Sun as a result in order to compare the sunlight performance of the system, a new approach was devised. By taking a standard Thorlabs power meter with the 700 nm to 1800 nm (`https://www.thorlabs.com/item/S122C`) head attachment outside and tuning it to 1653 nm, a measure of a sunlight level could be established. The ambient sunlight level could then be simulated within the lab using a theatrical stage lamp. The stage lamp illuminates broadly using a wide array of wavelengths including SWIR, and hence a significant component of light around 1653 nm, meaning that both systems can observe the illumination. This is a potentially flawed measurement in hindsight because the optical head, whilst only being sensitive within the same region as the internal SWIR camera, is not used as in conjunction with a filter to ensure that only light at 1653 nm is being accurately recorded, meaning that other wavelengths of light in the response range is being counted additionally. The theatre lamp, whilst not matching the spectral response of sunlight, will cover the full spectral response of the optical head on the power meter meaning that a similarly wide response to sunlight is experienced [171]. Whilst the accuracy of this procedure as a true approximation of

sunlight is unlikely, what is established is a strong and constant source of light containing light at 1653 nm. The systems were then positioned such that the incoming backscattered power from the stage lamp, measured on the same power meter, approximately matched that of the outside reading. It could then be observed which of the two systems fared better under these high levels of ambient light, raw SWIR images taken on both systems are presented in Figure 7.8, highlighting the difference in degree of saturation experienced.

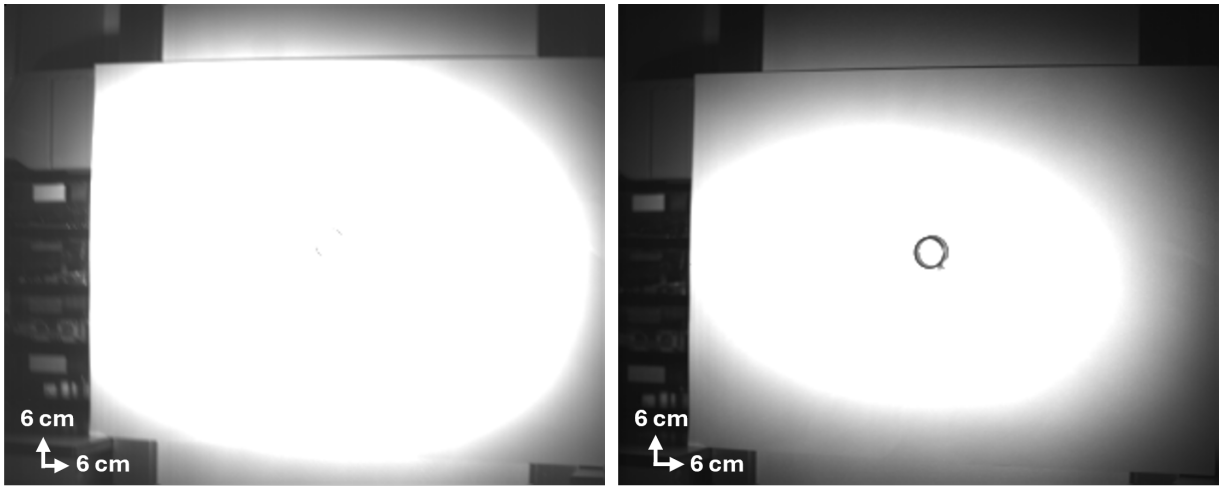

(a) (b)

Figure 7.8: **Comparison between raw SWIR frames taken on (a) the older system and (b) the new system.** The degree of saturation is far more evident in (a), as a result of the higher exposure time, however in (b) it is possible to make out the location of the methane cell.

Immediately it could be seen that, there was a lesser degree of saturation when using the new device in comparison to that of the old, due to the lower exposure time. It could also be noted that, in looking at the images side by side, the regions not directly illuminated by the spotlight, and thus only illuminated by the device's active illumination, highlights a similar return signal, despite differing exposure times. The sunlight sampled for this experiment was that of a bright winter day, meaning that this could potentially correspond to a particularly challenging scenario for the devices to overcome. The new system is close to saturation during this exposure, as such, it could be possible to scale the exposure time based on external ambient brightness, however one must be careful with implementation because this will have impact on maximum range of the system.

Upon processing the images, as presented in Figure 7.9, gas can clearly in the frames cap-

tured by the new system. In terms of the new system, the visible feed is just a singular frame captured independently of the SWIR capture. The exposure time of the visible camera in the new system is evidently higher than that of the old system's, causing it to be a bit more saturated. Because the green channel is attenuated one can still add to the visible feed and see gas despite this saturation, though this should in general be avoided due to the need to relate the gas signal to a real world location during field trials. By looking at a processed differential frame for the new system, as seen in Figure 7.10, it is clear that, because the full shape of the cell cannot be seen and only a crescent is visible, that there is still some degree of saturation experienced over the extent of the gas signal location. In reality, this experiment features a particularly challenging scenario as the backdrop is a white sheet whereas it is more likely a user would view more realistic targets with less reflectivity, however it is also a possibility that days with brighter sunlight signal could be experienced causing saturation even on these less reflective surfaces.

(a) (b)

Figure 7.9: **Comparison between processed gas detection frames taken on (a) the older system and (b) the new system.** In (a) there is no gas signal detected, it is also relatively noise free as a result of the frame being mostly saturated. In (b) there is a clear gas signal, however there is a small amount of noise signal present on the outside of the area where the saturation occurs.

From this comparison it is clear that the new system is far more capable at dealing with higher levels of ambient light. Whilst it fared much better here than the previous system it did not exhibit perfect results with some noisiness as well as being close to saturation in some areas of the SWIR image where gas is located. It is proof that the higher power diodes in addition to a lower exposure time is an avenue to overcoming sunlight saturation.

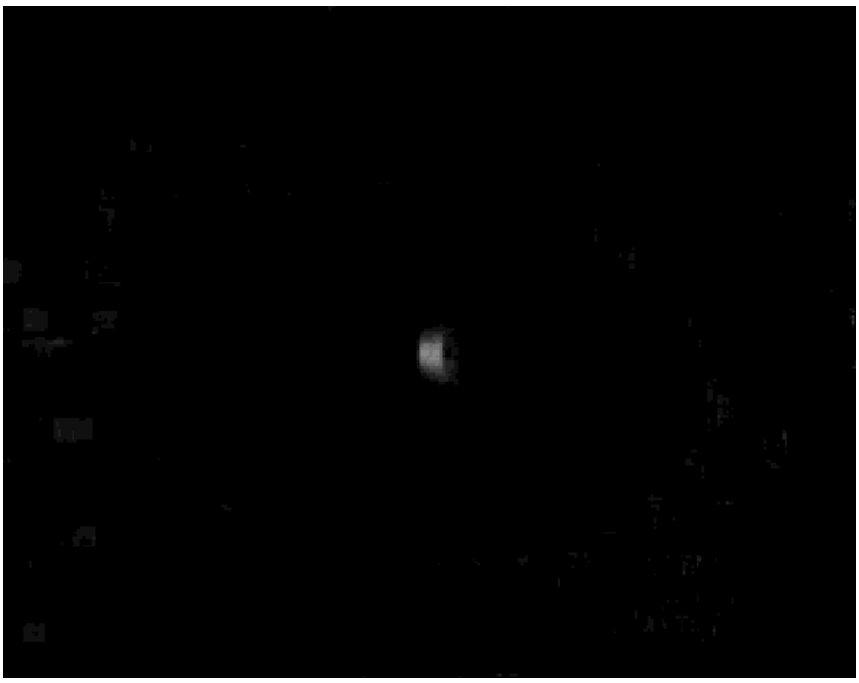

Figure 7.10: **Processed raw differential SWIR image from new system.** In this frame, which has been smoothed, resampled in size and had a threshold applied. It can be seen that the full, circular extent of the gas cell cannot be directly observed suggesting a degree of saturation.

7.4.2 Decreasing the Exposure Time

Having demonstrated the operating parameters for the new system fare better at decreasing the amount of ambient light saturation than the previous system, attention was turned to trying to find the optimal exposure time for gas detection in these conditions. As discussed, it was theorised that by decreasing the exposure time further one could potentially win out over the ambient light levels in terms of return gas signal, but it could potentially reach a point of diminishing returns where the low exposure time actually causes a reduced return signal to be obtained.

In order to test this hypothesis the new system was used once again with the stage light, set up with a power meter to record approximately 74 μW at the point of the cameras aperture, this matches a power meter reading for sunlight taken outside for 1.653 μm. A glass gas cell of approximately 10,000 ppm.m was then placed against a white background at a distance of 2.5 m away from the aperture, in the direct light of the stage light. SWIR images were then taken for varying exposure times, from 10 ms to 2 ms in intervals of 2 ms. Of note here is that, whilst exposure time varied, the frame rate of the camera remained locked at 10 fps, as the system remained stationary it is unlikely that this will result in increased motion artefacts. Roughly 400

frames were taken per exposure, but in processing only the first 100 were selected for use, this was to avoid any instances in which the double trigger activated and as such would cause a small reduction is signal strength. The differential images are filtered using a low pass filter of size 9×9, and the visible overlay images follow the same production pipeline as all images taken on the system, the visible image here is again simply a single visible frame taken in isolation, which given this is a static image should be valid for all exposures. Additionally an average signal over 100 frames was calculated to quantify performance, using the same algorithm as introduced in Chapter 7.3. The results of this investigation are presented in Figure 7.11.

From Figure 7.11 it can be seen that saturation increases with increasing exposure time, which is to be expected. A saturation increase in general means that the signal begins to become undetectable, and as can be seen from the 8 ms and 10 ms frames featuring additional false detections. One caveat to this is in the comparison between the 2 ms and 4 ms where is can be seen that the gas signal for the 2 ms is the lesser of the two despite the lower levels of saturation overall. This result highlights the fact that for this gas cell at this particular observation distance, then 4 ms is an ideal exposure time. The 4 ms exposure time provides the highest average pixel value of 18.9 counts and any increase or decrease in exposure time results in a decrease in signal, with 2 ms providing an average pixel value of 10, and 6 ms providing an average pixel value of 13.1. As mentioned this reduction in exposure time will have an effect on range, if one were to move closer to the gas source the 2 ms exposure may begin to prove more valuable than the 4 ms as the saturation increases.

7.5 Polarisation Sensitivity

Unfortunately, one new issue introduced in the design of the new system was that of polarisation disparity. Because different diodes illuminate the on and off absorption frames, if the polarisation of the light emitted by these diodes does not match then disparities between in-frame features that are sensitive to polarisation can be highlighted as a false gas distribution. These false signals were particularly noticeable on surfaces with a high degree of specular reflection, such as smooth metals or some glasses.

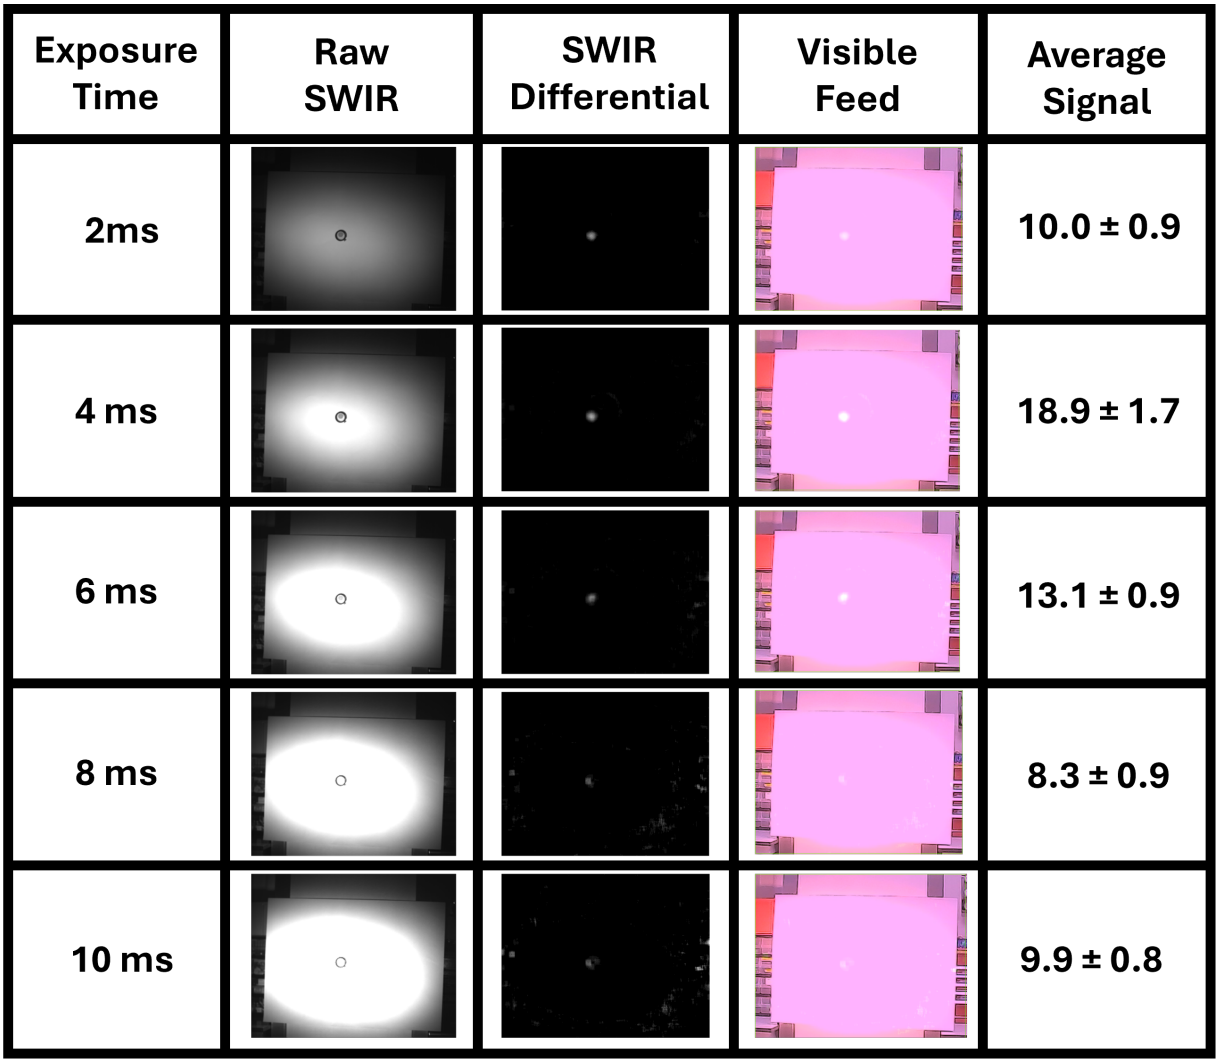

Figure 7.11: **Effects on increasing exposure time in the presence of high ambient light levels.** It can be seen that as the exposure time increases so too does the level of saturation in the raw SWIR feed. Whilst, in general this causes a loss of sensitivity, which can be gleaned from the weakening of the signal from 4 ms to 10 ms it can also be seen from the comparison between the images taken at 2 ms and 4 ms, wherein the shorter exposure time image has a less prominent signal as a result of receiving less signal.

One solution that was attempted was to include depolarising optics into the lens tubing after the diffuser. These are the extra optics mentioned before that cause the lens tubing to poke out of the faceplate, as well as being referred to in Figure 6.7. Two types of SWIR depolariser were used in conjunction, one made of quartz (ThorLabs DPU-25-C) the other a liquid crystal polymer (ThorLabs DPP25-C). The idea of including these was to scramble the polarisation of the light in such a way that any polarisation differences could be minimised. In reality, while this does help, it does not entirely remove the issue from the system.

In order to completely solve this problem the system would need to be rebuilt with this issue in mind. The most straightforward way to tackle this would be by introducing polarisation maintaining fibre throughout. To maintain this would mean that all fibre components would need to be replaced throughout the system. The diodes themselves would have to be replaced with versions coupled into polarisation maintaining fibre, as well as the mirror switch and fibre splitter.

7.6 Conclusion

In conclusion, the new system, which was built to address the key shortcomings of the previous system, is successful in improving the performance of the system in terms of frame rate and ambient light rejection. Increasing the SWIR frame rate from approximately 15 fps to 100 fps is by far the greatest success of the modifications. This frame rate boost was facilitated primarily by the introduction of the mirror switch to eliminate the tuning time limit that existed previously. By adding in the ability to camera trigger the switch and increasing the degree to which the system utilises parallelisation allow one to more fully take advantage of these frame rate increases as well as increasing the overall photon budget in the system. Another great success is the increase in the systems sunlight rejection. It was demonstrated, albeit in a lab setting, that the system beats the older system primarily as a result of the decreased exposure time. It is still clear that the system is sensitive to sunlight but, despite using a very similar filter to the old system, should only capture half of the sunlight with enough return signal to still detect the gas signal. The system also shows a greater degree of on-off stability by exhibiting far less

variance in the processed differential images frame to frame. Where the system still has room for improvement is in terms of its polarisation sensitivity and its range. The lack of polarisation control in the system means that there was a sensitivity to surfaces where polarisation response was a concern. This sensitivity was evident in the lab by the fact that the on and off diodes produced different sizes of specular reflection when reflecting off of metal surfaces, causing a persistent differential signal. Whilst depolarisers were used to decrease the extent of this effect the real solution to this would be to further redesign the system such that polarisation is controlled and maintained throughout. Finally, the new system exhibited a slightly lower total range. This range reduction was most likely due to the choice of exposure time, meaning that the system wasn't as sensitive as it could be, however, this must be seen as a potentially minor trade off in comparison to the clear benefits in frame rate, consistency and sunlight rejection, which could be hampered with a larger exposure time. Another way to recover range would be to increase laser power, a potential solution that must be implemented in accordance to laser safety regulations as well as also taking polarisation into account.

Chapter 8

Proposed Improvements and Future Work

8.1 Introduction

The final investigation conducted as part of the PhD was to interact with industry to gauge how this technology could end up as a commercially viable product. This research was performed in conjunction with QLM Technology LTD, who design a long range lidar system for accurately imaging and quantifying gas leaks over large distances. In order to evaluate the system the then newly completed system, introduced in Chapter 6, was brought down to the QLM research lab in order to test the systems capabilities. The relative strengths of the system as well as the key areas of improvement were explored and discussed.

8.2 Successes

The key successes of the successes of the system were identified as being the range, frame rate and ambient light rejection.

8.2.1 Range

The ranges that the GasSight system was capable of working at were seen as a success in terms of being a complementary technology to more sophisticated long range systems. As demonstrated in Figure 7.6, for the outlined settings both iterations are capable of accurate detection of a gas

cell of 5000 ppm.m of methane for ranges up to approximately 11 m. Longer range systems, such as those produced by QLM, whilst very sensitive with ranges up to and exceeding 100 m, can suffer at closer observations distances due to saturation, with few results reported below 10 m [40]. By having a portable device that works for short distances one can fine tune and cover the regions that these longer range systems cannot.

8.2.2 Frame rate

The frame rate of the system was also praised. As introduced previously, the SWIR feed and differential imaging runs at 100 fps (10 ms per iteration), whereas the visible capture and processing loop runs at $\approx$ 30 fps (35 ms per iteration), this slower feed is what is displayed on the camera interface. By having a 30 fps feed it is possible to track a leak on a moving system, which can be challenging for imaging regimes that require scanning or many data points to build up an image of the gas scene. The QLM system for example takes on order of 100 s to generate each image, meaning that the 30 fps visible refresh feed shows a considerable speed increase [40].

8.2.3 Ambient Light Rejection

Another key success was the systems rejection of high ambient light levels. During a visit to the company research lab metal gas cells with glass windows containing methane concentrations of 1082 ppm.m and 2615 ppm.m were placed outside in direct sunlight. The system was able to detect the gas cells with almost the same ease as indoors, albeit with an increased level of false detection. The settings used here were as before, with a SWIR frame rate of 100 fps, and a SWIR exposure time of 8 ms. Additionally the gas cells did require a backdrop, in this case a chequerboard card, to allow for detection. The results of this test can be observed in Figure 8.1. As discussed in earlier chapters the previous gas imaging system struggled with sunlight, causing issues in the detection of the gas, so being able to demonstrate rejection in a real world setting, as opposed to the simulated sunlight of Figure 7.9, was deemed highly successful. Whilst the success of the ambient light rejection was initially put down to the decreased exposure time and higher power diodes, it is possible however that the shade caused by the metal casing on the cells

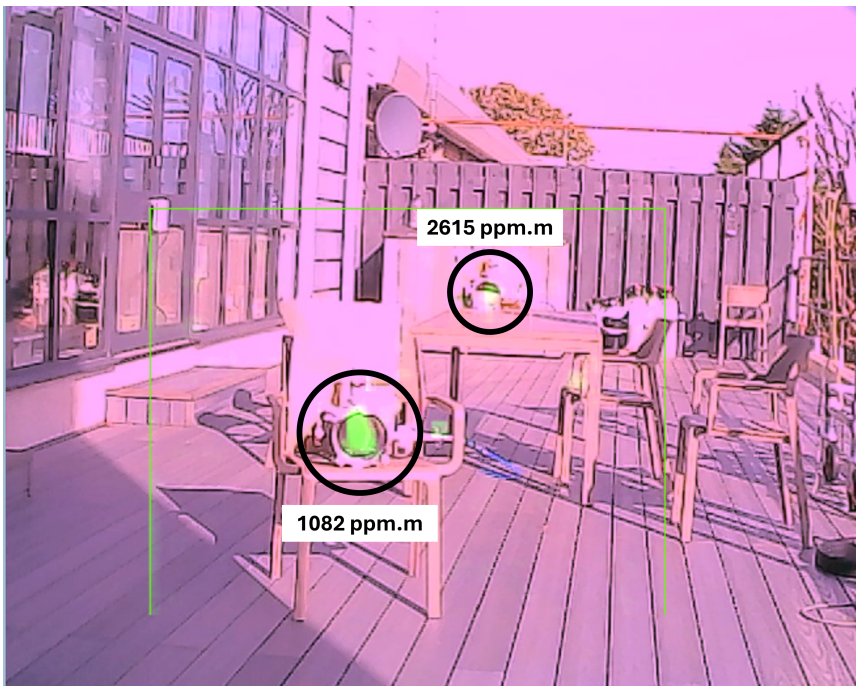

Figure 8.1: **Example frame taken in bright ambient conditions, with two gas cells visible.** The concentration of the closer cell was 1082 ppm.m and the further cell was 2615 ppm.m. Both cells were provided cardboard backgrounds. The cells were also positioned such that they were directly in the sunlight. While the image is quite noisy and there are a few false artefacts due to glare, the system can identify the gas locations.

provided protection from the sunlight and may have artificially aided the results by reducing the amount of sunlight in the path of the gas. This is something that would need to be explored in the future.

8.3 Necessary Improvements

Perhaps more importantly than identifying what the system already does well, are the changes that must be made to the system in order to be a commercially viable product.

8.3.1 Tuning

A point of discussion was the fact that the diodes are manually tuned only periodically. Whilst manual tuning does ensure that the system is tuned to the correct wavelength, it is also prone to drifting if not regularly checked. Drifting would cause the accuracy of the system to vary over time as the slightly offset wavelength will be absorbed slightly less strongly by the methane

until, in the worst case scenario, it is no longer absorbed. This would lead to a system that only provided false negative readings, which could lead to consequences such as missing leaks during costly surveys or even worse explosion and health risks. Such a scenario would be far from ideal for a product on the market which requires a high level of consistency with each use. There are multiple methods one could implement in order to ensure that the systems tuning does not drift over time. One such method may be to include a gas reference cell, which is used to automatically tune the diodes in order to ascertain the settings that provide maximal absorption. Such a calibration could be implemented upon boot up of the system. It could also possible to implement an external tuning apparatus containing a gas cell, with a program which could allow an end user to tune the system directly, with minimal need to manually probe a source. A suggestion from industry was that this could be a dock type system that automatically recalibrates the system using an external cell whilst charging. This approach would not involve adding components internally, allowing the form factor to remain the same, but could be very complex to implement and only allows calibration with the dock. An internal cell would allow for calibration to be performed whenever, but would require an internal redesign of the system where light can be passed through the cell and detected, either on a camera or a small photodiode.

8.3.2 Detection Threshold

As relayed by the company, the minimum size of leak that a gas detection device needs to be able to detect as outlined by the US government is 1 ton per year [40]. To test if the gas camera could achieve this goal, QLM provided a gas cell designed to challenge the system based on the detection capabilities of other systems such as those produced by FLIR [172]. This concentration was 300 ppm.m, which was a factor of 3 less concentrated than what was detected with the other two cells in Figure 8.1. The gas camera was unfortunately unable to detect this cell at any range with an exposure time of 8 ms meaning a sensitivity increase is required. In order to increase the sensitivity threshold would require more light collected back in the system, this could potentially be achieved in a few different ways, for example more laser power or a higher exposure time. Increasing the exposure time is the easier of the two options but opens up problems in regards to movement issues and sunlight as the time between subsequent frames

increases as well as increasing the likelihood of saturation. Increasing the laser power also adds its own challenges, mainly the difficulty in procuring a diode with a power output greater than 100 mW. A work around for this would be to introduce more diodes to the system to build the power that way. Adding diodes would have an increased cost, but also the driving solution may need to be reassessed such that the system is not too bulky and that the battery can support such a power draw, though this could be avoided by running the diodes in series in order to prevent this increased draw from being unmanageable. By looking at the findings of Chapter 5.5 the SNR, and thus detection threshold could be increased in additional ways to the two avenues already explored. Firstly, increasing the collection area of the system, perhaps by increasing the aperture radius via the introduction of a larger lens could lead to an increase in SNR, but would have to be performed carefully so as to not increase the system footprint or mass. The system could also make better use of its temporal smoothing abilities to average over a greater number of images and pixels to increase the SNR, though this must not come at the expense of the frame rate. Finally, the system could make use of a lower noise camera such as the Raptor Owl or FirstLight camera mentioned previously. A lower dark current, and readout noise is particularly of benefit at lower exposure times where the readout noise is more comparable in size to the signal and the dark noise scaling is more similar. The issue with swapping the camera again comes with size and mass concerns, but also usability and computing cost, an issue experienced with the FirstLight camera. Increasing detection sensitivity should be the first issue that is tackled in the near future. Given there is a minimum detection threshold for a system to be recognised as a solution, it stands to reason that it should be a key goal of any system to reach this level before pursuing potential commercialisation.

8.3.3 Concentration Readings

Another key recommendation was that the system be able to tell one how much gas is present as opposed to just where the gas is. Adding such capabilities would involve altering the system in order to provide accurate concentration readings. The system would have to be far more controlled than it currently is, especially in regards to the laser. Such a modification could go in step with the tuning of the diodes as slightly drifted diodes will cause the system to give inaccurate

values. To properly quantify the gas concentration will also require knowledge of exactly what the off wavelength of the system is. By knowing the exact absorption profiles for the two wavelengths it should be possible to assign concentration values to the gas signal detected. Providing a gas measurement was highlighted as vital for the system moving forwards, the information as to the amount of gas is very valuable as it reduces the requirement for end users to make assessments on leak severity based upon gas images [40, 69]. By having a combination of accurate concentration reading as well as a visual component could mean that this device could act as quite a compelling companion device to the long range QLM system with no loss of accuracy between the two systems.

Additionally, there is potential here for machine learning to be integrated to expedite this process as in [173]. Here one may be able to train the device on a wide range of data sets and different gas concentrations and match what is observed to these. Whilst machine learning could be an ideal solution there are huge technical challenges that must be overcome. The main issue with a machine learning approach is that, because gas is amorphous in shape and size the learnsets required to characterise such a variable system would have to be large. Learnsets with such large size can run quite slowly which could in turn slow the frame rate of the system to below video-rate, which could be a huge knock to the systems strengths.

8.3.4 False Detections

As described throughout this work, certain scenarios will cause false gas detections in the system. One origin of these false signals is from the sensitivity to polarisation. This sensitivity is brought about by the fact that polarisation is not maintained throughout the system and the outputs of the two laser diodes will differ between frames. The movement issues stem from the movement related issues described throughout this thesis, though in general these were easier to filter out than the glare related noise.

It was made clear that any form of false detection would be looked upon poorly in industry. Removing false detection is desirable because remedial actions for gas leaks can often be costly so a high degree of certainty is desired to prevent unnecessary cost. In order to reduce these false signals there is a great deal of work that needs to be undertaken. The movement

issues are somewhat unavoidable in this differential imaging regime where the two images are taken at different times and not concurrently. By increasing the frame rate of the system and employing filtering in both time and space these false signals can be diminished but care has to be taken to ensure this has no impact on the systems potential sensitivity and provide incorrect readings. There are possible solutions where the two images are taken simultaneously on the same detector chip by splitting the two wavelengths equally among half the pixels. This dual imaging would eliminate any background noise issues as well as movement artefacts. A dual imaging system would require many changes to the overall system, the first of which would be a more broadband light source that contains both the desired on and off absorption wavelengths. Secondly the light entering the system would need to be split into the constituent wavelengths, this could be achieved using a dichroic. The anticipated issues that arise with the dichroic is that it would have to have a very narrow wavelength edge in order to facilitate the fact that the two wavelengths used are so close together, which may be impossible. It could mean selecting an off wavelength further away from the absorption band but this would have to be investigated more completely, for example by ensuring that the intensity of the light at both wavelengths matches. The resolving of two non-overlapping images onto the detector is also challenging and would require careful optical setups which may increase the size and complexity of the handheld system. Similar effects can be obtained by using Wollaston prisms [174].

An interference based GasSight design featuring a setup similar to a homodyne telescope and based upon research conducted within the group was briefly entertained as a work around for sunlight [175]. The working principle would be that the two differential images could be taken at the same time and thus avoid movement noise. This idea was dismissed due to difficulty and limited usage in the field. The main drawback, aside from the increased complexity, was the rigidity of the range of the system. The requirement to have a fixed delay line length, paired with a potentially narrow bandwidth of laser light meant that industry use such as attached to a UAV would be far more difficult and limited in use.

The polarisation based issue has been less explored so far, but potential solutions have been devised. The most obvious solution to removing this issue is to introduce control over the polarisation from the beginning. To introduce such a level of control would involve replacing all

fibre coupled components with polarisation maintaining fibre. Adding this control will mean that the polarisation of the system will be more predictable. The diodes would also require replacing as currently they are only coupled into single mode fibres, though a polarisation maintaining option does exist.

Other such solutions involve rapidly changing the bend in the fibre in order to randomise the polarisation, much like the depolarisers attempt to do. It is hoped that by depolarising in such a way that the change in polarisation is far more dynamic then over the course of an exposure the randomised polarisation will hopefully balance out between the two frames in such a way that the glare related issues are diminished. Integrating such a process could again place tight constraints upon the size, weight and form factor of the system.

Overall, there are plenty of potential options available in order to attempt to mitigate false detection. The feasibility of some of these techniques are limited however, with potentially great overhaul required in order to achieve these improvements.

8.4 Conclusion

In conclusion, while the systems created so far are functional in such a way that they can generate interest, they are far from ready for commercialisation and deployment within the field. Whilst the system shows a lot of promise in regards to its frame rate, sunlight rejection and range it has many limitations that would need to be overcome fully before the system could be taken to market. These improvements are all substantial in that they would require changes be made to the system that could compromise the form factor and mass of the system. The end goal would be to create a system that can provide accurate concentration readings of gas, in addition to generating an imaging component. The system has to also be made far more robust to false detection, with almost no room for error, whilst the system is quite good at mitigating these currently, new techniques may need to be put in place in order to fully eradicate uncertainty in a measurement.

Chapter 9

Conclusions

The key aim for this thesis was to develop a system that can image invisible gases in a real world environment. The gas in question was methane, a dangerous greenhouse gas with serious climate impact. In order to develop such a system required looking back to the previous iterations and similar technologies before using this knowledge to move forward and refine the current system before then using this experience to build a new system that overcomes any hurdles found during this process. This thesis has presented that journey by acknowledging the past and documenting the work conducted to refine and further develop the system.

Work in the group has, over the years, undergone a fair few revisions, from backpack mounted to a single pixel camera. All of these systems had unique benefits, but also shortcomings. Some were able to image, some were able to accurately measure concentration, some were low cost, and some were portable. The culmination of this work was the GasSight system which was portable, relatively cost effective and easy to operate. This iteration of the GasSight combined the learning of the previous versions and was the springboard for the work presented throughout this thesis.

The journey of this research continues with this existing system and found that the image analysis was very sensitive to movement. This sensitivity meant that it could not be trucked or moved without the addition of high levels of noise. In order to tackle this sensitivity to movement, many motion compensation techniques were experimented with in such a way that the system remained operating at a high frame rate. This meant that traditional solutions such

as phase correlations were identified as bottle necks, slowing the system. It was from here that the move to a different imaging sequence where two SWIR images are taken per loop was introduced. By changing where the images are taken in the algorithm the downtime between images was reduced from 100 ms to that of the tuning time of the diodes, which was around the order of 15 ms, a $\approx 7\times$ decrease in downtime. This, in addition to the introduction of harsher persistency, erosion and dilation and quick and simple cross correlations resulted in a system that, whilst not taking full advantage of the photon budget, could be moved around without too much loss of sensitivity and, crucially, without adding movement artefacts.

Following the implementation of the improved movement correction was the unfortunate discovery that the system had difficulties in direct sunlight. Upon taking the system into the summer sunlight it was found that almost all the sensitivity was lost. This started a study into how one could overcome this issue quickly and easily by changing the properties of the internal spectral filter, designed to block this unwanted light. What followed was an in depth study into the behaviour of the filter in the system specifically in regards to its impact upon the field of view. What was discovered was that, in order to effectively filter the light, one would require a narrow filter which in turn would significantly reduce the value of the FOV. A proposed solution, that used a large aperture dual lens system to allow the integration of narrow bandwidth filters while maintaining a suitable FOV, was trialled but deemed cumbersome in practice. Instead it was decided that, in order to overcome this, a new system would need to be made that addresses the problem from the outset by employing much shorter exposure times.

With the limitations of the system in mind, outside trials continued and drone mounting began. The decision was made early on to perform these trials as the Sun sets and into the night in order to balance visibility with low sunlight levels. These trials reveal that the system is sensitive enough to work on a drone from a range of 10 m horizontally and 3 m vertically, but at the same time robust enough not to suffer due to rotor vibrations. Further to this, an investigation was conducted into the return signal strength as a function of distance for a range of different surfaces. What was found was that, up to a range of around 10 m, a wide range of surfaces were able to return enough signal to identify a gas filled balloon of approximately 100,000 ppm.m. Further to this, the overall range of the system was tested for a cell of concentration 5000 ppm.m

and found to have a range of up to 16 m. These studies were key in characterising the systems performance in real world scenarios.

The findings of the work up until this point was the fact that many of the cameras issues stem from its inherent lack of frame rate and processing speed. At this point much was done to increase overall frame rate of the system, so a new system focusing on maximising this further was designed. The fundamental frame rate limit of the system was the tuning time of the diodes, this was eliminated by the introduction of a camera triggered MEMS mirror switch. This increased frame rate lent itself to a higher degree of parallelisation of the SWIR and visible image processing, less movement noise, and much improved photon budget. The new, more powerful diodes allow for further frame rate increases by allowing a much lower exposure time to be used, greatly decreasing the time length of the differential imaging, achieving a SWIR frame rate of 100 fps. A key trade off of this system is its increased cost of around £7,000, as a result of increased costs related to the diodes. Whilst this cost is high it is still inline with comparable OGI systems from companies such as FLIR which sell systems that cost from around £18,000 to £25,000[1]. It is hoped that the cost of the system will decrease over time, especially as SWIR cameras are developed into more commercial products driving the price down further.

In addition to the frame rate benefit of the switching and triggering, it also results in a far less varied differential signal response. This increase in frame rate greatly reduces motion artefacts and also provides clear wins in terms of sunlight resilience enabled by the significantly lower exposure time. This allows the system to overcome the sunlight more easily whilst still being able to detect gas. The decrease in exposure time comes with a minor decrease in the systems range capabilities, something that could be recovered with higher exposure times, perhaps at the cost of some movement and sunlight resilience. Despite the successes elsewhere, the system does falter in terms of its sensitivity to polarisation, specifically due to the varying intensity of backscattered light from specular reflective surfaces such as metals, which differ for on and off absorption states due to differences in polarisation between the two diodes. Despite attempts to minimise this impact by introducing depolarisers, the true solution would be instead to design a

[1]Based upon example listings from January 2026: https://www.tester.co.uk/teledyne-flir-gf77a-fixed-gas-find-uncooled-ir-thermal-camera, https://www.tester.co.uk/flir-teledyne-gf77-gas-find-uncooled-ir-thermal-imaging-camera-choice-of-lens

system that incorporates polarisation control throughout.

As a final experiment the new system was demonstrated to a gas detection company, QLM Technologies, in order to understand what improvements would be required if the system was to be commercialised and how it differs in use case to their device. The system was deemed successful in terms of its range, frame rate and ability to work in sunlight. The system was demonstrated to work up to 11 m with an overall final imaging frame rate of 30 fps, whereas the QLM system works up to ranges of 100 m with an imaging times of around 100 s. This comparison shows that the system could work effectively as a companion product to QLMs existing long range scanning devices, as a short range real time imager. Despite this, in order to make the system commercially viable, key future work was identified. The system would need to be able to provide concentration readings and this would involve careful improvements to the tuning, as well as incorporating some level of calibration. The system also needs to be able to detect significantly lower concentrations of gas, for example 300 ppm.m, in order to be competitive and the likelihood of false detection would need to be fully removed. These improvements are attainable but need to be carefully implemented as they can exacerbate one another.

To conclude, the work conducted in the course of this thesis brings an existing system further towards being a tool that could be reliably used to detect and mitigate serious methane gas leaks. Reaching this point involved improving the robustness of the system to a wide variety of complex scenarios such as drone mounting, movement and sunlight. Throughout this work the system has been repeatedly demonstrated to be far more resilient than before. These improvements are most evident when you consider that initially the system was a stationary lab based detector that ended up being attached to a drone and successfully trialled.

Bibliography

[1] Xiangyu Ou, Xue Chen, Xianning Xu, Lili Xie, Xiaofeng Chen, Zhongzhu Hong, Hua Bai, Xiaowang Liu, Qiushui Chen, Lin Li, and Huanghao Yang. Recent development in X-ray imaging technology: Future and challenges. *Research*, 2021, 2021.

[2] Chuntian Cao, Michael F. Toney, Tsun-Kong Sham, Ross Harder, Paul R. Shearing, Xianghui Xiao, and Jiajun Wang. Emerging X-ray imaging technologies for energy materials. *Materials Today*, 34:132–147, 2020.

[3] Daan Stellinga, David B. Phillips, Simon Peter Mekhail, Adam Selyem, Sergey Turtaev, Tomáš Čižmár, and Miles J. Padgett. Time-of-flight 3D imaging through multimode optical fibers. *Science*, 374(6573):1395–1399, 2021.

[4] Torgrim Log, Wegar Bjerkeli Pedersen, and Heike Moumets. Optical gas imaging (OGI) as a moderator for interdisciplinary cooperation, reduced emissions and increased safety. *Energies*, 12(8), 2019.

[5] Lydia H. Kifner, Aram J. K. Calhoun, Stephen A. Norton, Kristine E. Hoffmann, and Aria Amirbahman. Methane and carbon dioxide dynamics within four vernal pools in Maine, USA. *Biogeochemistry*, 139(3):275–291, August 2018.

[6] Sheel Bansal, Olivia F. Johnson, Jacob Meier, and Xiaoyan Zhu. Vegetation affects timing and location of wetland methane emissions. *Journal of Geophysical Research: Biogeosciences*, 125(9):e2020JG005777, 2020. e2020JG005777 2020JG005777.

[7] Joshua F. Dean. Old methane and modern climate change. *Science*, 367:846–848, 2020.

[8] Tahani Aldhafeeri, Manh-Kien Tran, Reid Vrolyk, Michael Pope, and Michael Fowler. A review of methane gas detection sensors: Recent developments and future perspectives. *Inventions*, 5(3), 2020.

[9] Robert W. Howarth. A bridge to nowhere: Methane emissions and the greenhouse gas footprint of natural gas. *Energy Science and Engineering*, 2:47–60, 2014.

[10] Kenneth L Cashdollar, Isaac A. Zlochower, Gregory M Green, Richard A Thomas, and Martin Hertzberg. Flammability of methane, propane, and hydrogen gases. *Journal of Loss Prevention in the Process Industries*, 13(3):327–340, 2000.

[11] Nitish Kumar and Himanshu Gupta. Chapter 18 - Methane: Risk assessment, environmental, and health hazard. In Jaspal Singh, R.D. Kaushik, and Malvika Chawla, editors, *Hazardous Gases*, pages 225–238. Academic Press, 2021.

[12] Benjamin Hmiel, V. V. Petrenko, M. N. Dyonisius, C. Buizert, A. M. Smith, P. F. Place, C. Harth, R. Beaudette, Q. Hua, B. Yang, I. Vimont, S. E. Michel, J. P. Severinghaus, D. Etheridge, T. Bromley, J. Schmitt, X. Faïn, R. F. Weiss, and E. Dlugokencky. Preindustrial 14CH4 indicates greater anthropogenic fossil CH4 emissions. *Nature*, 578:409–412, 2020.

[13] Xiaochun Zhang, Nathan P. Myhrvold, Zeke Hausfather, and Ken Caldeira. Climate benefits of natural gas as a bridge fuel and potential delay of near-zero energy systems. *Applied Energy*, 167:317–322, 2016.

[14] Nathan G. Phillips, Robert Ackley, Eric R. Crosson, Adrian Down, Lucy R. Hutyra, Max Brondfield, Jonathan D. Karr, Kaiguang Zhao, and Robert B. Jackson. Mapping urban pipeline leaks: Methane leaks across Boston. *Environmental Pollution*, 173:1–4, 2013.

[15] Robert B. Jackson, Adrian Down, Nathan G. Phillips, Robert C. Ackley, Charles W. Cook, Desiree L. Plata, and Kaiguang Zhao. Natural Gas Pipeline Leaks Across Washington, DC. *Environmental Science & Technology*, 48(3):2051–2058, February 2014. Publisher: American Chemical Society.

[16] Levi M. Golston, Nicholas F. Aubut, Michael B. Frish, Shuting Yang, Robert W. Talbot, Christopher Gretencord, James McSpiritt, and Mark A. Zondlo. Natural gas fugitive leak detection using an unmanned aerial vehicle: Localization and quantification of emission rate. *Atmosphere*, 9(9), 2018.

[17] Sleight Halley, Lok kun Tsui, and Fernando Garzon. Combined mixed potential electrochemical sensors and artificial neural networks for the quantificationand identification of methane in natural gas emissions monitoring. *Journal of The Electrochemical Society*, 168:097506, 9 2021.

[18] X.F. Lou and J. Nair. The impact of landfilling and composting on greenhouse gas emissions – a review. *Bioresource Technology*, 100(16):3792–3798, 2009. Selected papers from the International Conference on Technologies and Strategic Management of Sustainable Biosystems.

[19] Grant Allen, Peter Hollingsworth, Khristopher Kabbabe, Joseph R. Pitt, Mohammed I. Mead, Samuel Illingworth, Gareth Roberts, Mark Bourn, Dudley E. Shallcross, and Carl J. Percival. The development and trial of an unmanned aerial system for the measurement of methane flux from landfill and greenhouse gas emission hotspots. *Waste Management*, 87:883–892, 2019.

[20] C. Özgen Karacan, Felicia A. Ruiz, Michael Cotè, and Sally Phipps. Coal mine methane: A review of capture and utilization practices with benefits to mining safety and to greenhouse gas reduction. *International Journal of Coal Geology*, 86(2):121–156, 2011.

[21] Takaya Iseki, Tai Hideo, and Kiyoshi Kimura. A portable remote methane sensor using a tunable diode laser. *Measurement Science and Technology*, 11:594–602, 2000.

[22] Nicholas M. Davis, Daniel Francis, Jane Hodgkinson, and Ralph P. Tatam. Compact methane sensor using an integrating sphere and interband cascade laser at 3313 nm. *Sensors and Actuators B: Chemical*, 389:133866, 2023.

[23] Chengzhi Xie, Mohsin Aziz, Vincenzo Pusino, Ata Khalid, Matthew Steer, Iain G. Thayne, Marc Sorel, and David R. S. Cumming. Single-chip, mid-infrared array for room temperature video rate imaging. *Optica*, 4(12):1498–1502, Dec 2017.

[24] Shuhao Wu, Khue Tian Lai, Vincenzo Pusino, and David R. S. Cumming. High-quality factor mid-wave infrared metasurface for methane detection applications. In Duncan L. Hickman, Helge Bürsing, Ove Steinvall, and Philip J. Soan, editors, *Electro-Optical and Infrared Systems: Technology and Applications XXI*, volume 13200, page 132000F. International Society for Optics and Photonics, SPIE, 2024.

[25] A. Kannath, J. Hodgkinson, R. G. Gillard, R. J. Riley, and R. P. Tatam. A VCSEL based system for on-site monitoring of low level methane emission. In James K. Guenter and Chun Lei, editors, *Vertical-Cavity Surface-Emitting Lasers XV*, volume 7952, page 79520F. International Society for Optics and Photonics, SPIE, 2011.

[26] Jinghan Dong, Arthur C. Cardoso, Haichen Zhou, Jingrui Zhang, Weijie Nie, Alex S. Clark, and John G. Rarity. Methane sensing via unbalanced nonlinear interferometry using a CMOS camera and undetected mid-infrared light. *Applied Physics Letters*, 126(6):061104, 02 2025.

[27] Kyriaki Kefala, Vincent Boudon, Sergei N. Yurchenko, and Jonathan Tennyson. Empirical rovibrational energy levels for methane. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 316:108897, 2024.

[28] Francesco Mazza, Ona Thornquist, Leonardo Castellanos, Thomas Butterworth, Cyril Richard, Vincent Boudon, and Alexis Bohlin. The rovibrational $\nu 2$ mode spectrum of methane investigated by ultrabroadband coherent Raman spectroscopy. *The Journal of Chemical Physics*, 158(9):094201, March 2023. _eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0138803/18043022/094201_1_5.0138803.pdf.

[29] F. Menard-Bourcin, C. Boursier, L. Doyennette, and J. Menard. Rotational and Vibrational Relaxation of Methane Excited to $2\nu 3$ in CH4/H2 and CH4/He Mixtures at 296 and 193 K from Double-Resonance Measurements. *The Journal of Physical Chemistry A*, 109(14):3111–3119, April 2005. Publisher: American Chemical Society.

[30] V. Malathy Devi, D. Chris Benner, Keeyoon Sung, Timothy J. Crawford, Shanshan Yu, Linda R. Brown, Mary Ann H. Smith, Arlan W. Mantz, Vincent Boudon, and Syed Ismail. Self- and air-broadened line shapes in the $2\nu 3$ p and r branches of 12ch4. *Journal of Molecular Spectroscopy*, 315:114–136, 2015. Spectroscopy with Synchrotron Radiation.

[31] Jean-Michel Hartmann, Ha Tran, Raymond Armante, Christian Boulet, Alain Campargue, François Forget, Livio Gianfrani, Iouli Gordon, Sandrine Guerlet, Magnus Gustafsson, Joseph T. Hodges, Samir Kassi, Daniel Lisak, Franck Thibault, and Geoffrey C. Toon. Recent advances in collisional effects on spectra of molecular gases and their practical consequences. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 213:178–227, 2018.

[32] Erfan Chenshen, Juntao Tan, Bin Wang, Erlong Jiang, Nan Zhao, Shaofeng Zheng, Zeren Luo, Jiaming Li, and Qingmao Zhang. A review of spectral broadening in laser atomic absorption spectroscopy. *J. Anal. At. Spectrom.*, 40:1902–1915, 2025.

[33] Henry Margenau. Pressure shift and broadening of spectral lines. *Phys. Rev.*, 40:387–408, May 1932.

[34] Darrell E. Burch, Edgar B. Singleton, and Dudley Williams. Absorption line broadening in the infrared. *Appl. Opt.*, 1(3):359–363, May 1962.

[35] Mu Li, Jie Gao, Yongqi Wu, Rui Zhu, Qiang Gao, and Yungang Zhang. Study on temperature effects of broadband absorption spectroscopy based on doppler broadening combined with a multilinear shape superposition. *Opt. Lett.*, 48(22):5923–5926, Nov 2023.

[36] I.E. Gordon, L.S. Rothman, R.J. Hargreaves, R. Hashemi, E.V. Karlovets, F.M. Skinner, E.K. Conway, C. Hill, R.V. Kochanov, Y. Tan, P. Wcisło, A.A. Finenko, K. Nelson, P.F. Bernath, M. Birk, V. Boudon, A. Campargue, K.V. Chance, A. Coustenis, B.J. Drouin, J.–M. Flaud, R.R. Gamache, J.T. Hodges, D. Jacquemart, E.J. Mlawer, A.V. Nikitin, V.I. Perevalov, M. Rotger, J. Tennyson, G.C. Toon, H. Tran, V.G. Tyuterev, E.M. Adkins, A. Baker, A. Barbe, E. Canè, A.G. Császár, A. Dudaryonok, O. Egorov, A.J. Fleisher, H. Fleurbaey, A. Foltynowicz, T. Furtenbacher, J.J. Harrison, J.–M. Hartmann, V.–M. Horneman, X. Huang, T. Karman, J. Karns, S. Kassi, I. Kleiner, V. Kofman, F. Kwabia–Tchana, N.N. Lavrentieva, T.J. Lee, D.A. Long, A.A. Lukashevskaya, O.M. Lyulin, V.Yu. Makhnev, W. Matt, S.T. Massie, M. Melosso, S.N. Mikhailenko, D. Mondelain, H.S.P. Müller, O.V. Naumenko, A. Perrin, O.L. Polyansky, E. Raddaoui, P.L. Raston, Z.D. Reed, M. Rey, C. Richard, R. Tóbiás, I. Sadiek, D.W. Schwenke, E. Starikova, K. Sung, F. Tamassia, S.A. Tashkun, J. Vander Auwera, I.A. Vasilenko, A.A. Vigasin, G.L. Villanueva, B. Vispoel, G. Wagner, A. Yachmenev, and S.N. Yurchenko. The hitran2020 molecular spectroscopic database. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 277:107949, 2022.

[37] Christopher S. Goldenstein, Victor A. Miller, R. Mitchell Spearrin, and Christopher L. Strand. Spectraplot.com: Integrated spectroscopic modeling of atomic and molecular gases. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 200:249–257, 2017.

[38] Ramón A. Alvarez, Daniel Zavala-Araiza, David R. Lyon, David T. Allen, Zachary R. Barkley, Adam R. Brandt, Kenneth J. Davis, Scott C. Herndon, Daniel J. Jacob, Anna Karion, Eric A. Kort, Brian K. Lamb, Thomas Lauvaux, Joannes D. Maasakkers, Anthony J. Marchese, Mark Omara, Stephen W. Pacala, Jeff Peischl, Allen L. Robinson, Paul B. Shepson, Colm Sweeney, Amy Townsend-Small, Steven C. Wofsy, and Steven P. Hamburg. Assessment of methane emissions from the U.S. oil and gas supply chain. *Science*, 361:186–188, 2018.

[39] A. R. Brandt, G. A. Heath, E. A. Kort, F. O'Sullivan, G. Pétron, S. M. Jordaan, P. Tans, J. Wilcox, A. M. Gopstein, D. Arent, S. Wofsy, N. J. Brown, R. Bradley, G. D. Stucky, D. Eardley, and R. Harriss. Methane leaks from north american natural gas systems. *Science*, 343(6172):733–735, 2014.

[40] James Titchener, Doug Millington-Smith, Chris Goldsack, George Harrison, Alexander Dunning, Xiao Ai, and Murray Reed. Single photon lidar gas imagers for practical and widespread continuous methane monitoring. *Applied Energy*, 306:118086, 2022.

[41] Kyle J. Nutt, Nils Hempler, Gareth T. Maker, Graeme P. A. Malcolm, Miles J. Padgett, and Graham M. Gibson. Developing a portable gas imaging camera using highly tunable active-illumination and computer vision. *Optics Express*, 28(13):18566, 2020.

[42] Torkil Holm. Aspects of the mechanism of the flame ionization detector. *Journal of Chromatography A*, 842(1):221–227, 1999.

[43] J. N. Driscoll and Jennifer Maclachlan. Monitoring Methane from Fracking Operations: Leak Detection and Environmental Measurements. *ACS Energy Letters*, 2(10):2436–2439, October 2017. Publisher: American Chemical Society.

[44] Derek Hollenbeck, Demitrius Zulevic, and Yangquan Chen. Advanced leak detection and quantification of methane emissions using suas. *Drones*, 5(4), 2021.

[45] Patrick V. Riquetti, Jayne I. Fletcher, and Colin D. Minty. Aerial Surveillance for Gas and Liquid Hydrocarbon Pipelines Using a Flame Ionization Detector (FID). In *IPC1996*, pages 681–687, Volume 2: Design, Construction, and Operation Innovations; Compression and Pump Technology; SCADA, Automation, and Measurement; System Simulation; Geotechnical and Environmental, June 1996.

[46] Takaaki Otagawa, Solomon Zaromb, and Joseph R. Stetter. A room-temperature electrochemical sensor and instrument for monitoring methane. *Sensors and Actuators*, 8(1):65–88, 1985.

[47] Jian Li, Yulin Li, Xinjing Huang, Jiahao Ren, Hao Feng, Yu Zhang, and Xiaoxia Yang. High-sensitivity gas leak detection sensor based on a compact microphone array. *Measurement*, 174:109017, 2021.

[48] Pal Stefan Murvay and Ioan Silea. A survey on gas leak detection and localization techniques. *Journal of Loss Prevention in the Process Industries*, 25:966–973, 2012.

[49] Qixin He, Chuantao Zheng, Huifang Liu, Bin Li, Yiding Wang, and Frank K. Tittel. A near-infrared acetylene detection system based on a 1.534 μm tunable diode laser and a miniature gas chamber. *Infrared Physics and Technology*, 75:93–99, 2016.

[50] Chuan-Tao Zheng, Jian-Qiang Huang, Wei-Lin Ye, Mo Lv, Jing-Min Dang, Tian-Shu Cao, Chen Chen, and Yi-Ding Wang. Demonstration of a portable near-infrared CH4 detection sensor based on tunable diode laser absorption spectroscopy. *Infrared Physics & Technology*, 61:306–312, 2013.

[51] Kan Rui-Feng, Dong Feng-Zhong, Zhang Yu-Jun, Liu Jian-Guo, Liu Cheng, Wang Min, Gao Shan-Hu, and Chen Jun. Influence of laser intensity in second-harmonic detection with tunable diode laser multi-pass absorption spectroscopy. *Chinese Physics*, 14(9):1904, sep 2005.

[52] Liu Jun, Tan Qiulin, Zhang Wendong, Xue Chenyang, Guo Tao, and Xiong Jijun. Miniature low-power IR monitor for methane detection. *Measurement*, 44(5):823–831, 2011.

[53] Javad Shemshad, Saiied Mostafa Aminossadati, and Mehmet Siddik Kizil. A review of developments in near infrared methane detection based on tunable diode laser. *Sensors and Actuators B: Chemical*, 171-172:77–92, 2012.

[54] Kun Liu, Lei Wang, Tu Tan, Guishi Wang, Weijun Zhang, Weidong Chen, and Xiaoming Gao. Highly sensitive detection of methane by near-infrared laser absorption spectroscopy using a compact dense-pattern multipass cell. *Sensors and Actuators, B: Chemical*, 220:1000–1005, 6 2015.

[55] Chuan-Tao Zheng, Wei-Lin Ye, Jian-Qiang Huang, Tian-Shu Cao, Mo Lv, Jing-Min Dang, and Yi-Ding Wang. Performance improvement of a near-infrared CH4 detection device using wavelet-denoising-assisted wavelength modulation technique. *Sensors and Actuators B: Chemical*, 190:249–258, 2014.

[56] U. Gustafsson, J. Sandsten, and S. Svanberg. Simultaneous detection of methane, oxygen and water vapour utilising near-infrared diode lasers in conjunction with difference-frequency generation. *Applied Physics B: Lasers and Optics*, 71:853–857, 2000.

[57] Xin Yu, Rui-Hong Lv, Fang Song, Chuan-Tao Zheng, and Yi-Ding Wang. Pocket-sized nondispersive infrared methane detection device using two-parameter temperature compensation. *Spectroscopy Letters*, 47(1):30–37, 2014.

[58] D.U. Schramm, M.S. Sthel, M.G. da Silva, L.O. Carneiro, A.J.S. Junior, A.P. Souza, and H. Vargas. Application of laser photoacoustic spectroscopy for the analysis of gas samples emitted by diesel engines. *Infrared Physics & Technology*, 44(4):263–269, 2003.

[59] Michael E. Webber, Michael Pushkarsky, and C. Kumar N. Patel. Fiber-amplifier-enhanced photoacoustic spectroscopy with near-infrared tunable diode lasers. *Appl. Opt.*, 42(12):2119–2126, 2003.

[60] Ke Chen, Min Guo, Shuai Liu, Bo Zhang, Hong Deng, Yonghao Zheng, Yewei Chen, Chen Luo, Li Tao, Mingqi Lou, and Qingxu Yu. Fiber-optic photoacoustic sensor for remote monitoring of gas micro-leakage. *Optics Express*, 27:4648, 2 2019.

[61] Jane Hodgkinson and Ralph P Tatam. Optical gas sensing: a review. *Measurement Science and Technology*, 24(1):012004, Nov 2012.

[62] E.V. Browell. Differential absorption lidar sensing of ozone. *Proceedings of the IEEE*, 77(3):419–432, 1989.

[63] Fabrizio Innocenti, Rod Robinson, Tom Gardiner, Andrew Finlayson, and Andy Connor. Differential absorption lidar (DIAL) measurements of landfill methane emissions. *Remote Sensing*, 9(9):953, 2017.

[64] E.V. Browell, S. Ismail, and W.B. Grant. Differential absorption lidar (DIAL) measurements from air and space. *Applied Physics B: Lasers and Optics*, 67(4):399–410, October 1998.

[65] Axel Amediek, Gerhard Ehret, Andreas Fix, Martin Wirth, Christian Büdenbender, Mathieu Quatrevalet, Christoph Kiemle, and Christoph Gerbig. Charm-f - a new airborne integrated-path differential-absorption lidar for carbon dioxide and methane observations: measurement performance and quantification of strong point source emissions. *Appl. Opt.*, 56(18):5182–5197, Jun 2017.

[66] Anhu Li, Xingsheng Liu, Jianfeng Sun, and Zhiyong Lu. Risley-prism-based multi-beam scanning lidar for high-resolution three-dimensional imaging. *Optics and Lasers in Engineering*, 150, 3 2022.

[67] N. Maistry, E. Cox, and R. A. Schutz. The investigation of different image processing techniques to improve the visibility of volatile gas plumes in industry using optical

gas imaging. In Monuko du Plessis, editor, *Fifth Conference on Sensors, MEMS, and Electro-Optic Systems*, volume 11043, page 1104319. International Society for Optics and Photonics, SPIE, 2019.

[68] Thomas Strahl, Johannes Herbst, Armin Lambrecht, Eric Maier, Jonas Steinebrunner, and Jürgen Wöllenstein. Methane leak detection by tunable laser spectroscopy and mid-infrared imaging. *Applied optics*, 60:C68–C75, 2021.

[69] Jingfan Wang, Lyne P. Tchapmi, Arvind P. Ravikumar, Mike McGuire, Clay S. Bell, Daniel Zimmerle, Silvio Savarese, and Adam R. Brandt. Machine vision for natural gas methane emissions detection using an infrared camera. *Applied Energy*, 257:113998, 2020.

[70] Arvind P. Ravikumar, Jingfan Wang, and Adam R. Brandt. Are optical gas imaging technologies effective for methane leak detection? *Environmental Science and Technology*, 51:718–724, 2017.

[71] Robert Olbrycht and Marcin Kałuża. Optical gas imaging with uncooled thermal imaging camera - impact of warm filters and elevated background temperature. *IEEE Transactions on Industrial Electronics*, 67(11):9824–9832, 2020.

[72] Daniel Zimmerle, Timothy Vaughn, Clay Bell, Kristine Bennett, Parik Deshmukh, and Eben Thoma. Detection limits of optical gas imaging for natural gas leak detection in realistic controlled conditions. *Environmental Science & Technology*, 54(18):11506–11514, 2020. PMID: 32786569.

[73] Nathan Hagen. Survey of autonomous gas leak detection and quantification with snapshot infrared spectral imaging. *Journal of Optics*, 22(10):103001, sep 2020.

[74] Paul Kaufmann, Helen M Chrzanowski, Aron Vanselow, and Sven Ramelow. Mid-IR spectroscopy with NIR grating spectrometers. *Opt. Express*, 30(4):5926–5936, Feb 2022.

[75] Dmitry A. Kalashnikov, Anna V. Paterova, Sergei P. Kulik, and Leonid A. Krivitsky. Infrared spectroscopy with visible light. *Nature Photonics*, 10(2):98–101, February 2016.

[76] Arthur C. Cardoso, Jinghan Dong, Haichen Zhou, Siddarth K. Joshi, and John G. Rarity. Methane sensing in the mid-IR using short wave IR photon counting detectors via non-linear interferometry. *Opt. Continuum*, 3(5):823–832, May 2024.

[77] Deblina Majumder, Ronak Janani, Alex Scrimshire, Alex Stone, William Brooks, Chris Holcroft, Rob Werner, Spencer Green, Natalie Wheeler, and Paul A Bingham. Raman gas sensing technology: A new horizon? *Sensors and Actuators Reports*, 9:100311, 2025.

[78] David J. M. Stothard, Roman Spesyvtsev, John Leck, Rory Pringle, Carolyn O'Dwyer, Craig Hunter, and Stuart Bennett. Molecule-specific, stand-off airborne substance detection with Deep-UV excited, range-resolved, single-photon "Quantum" Raman spectroscopy: Towards an optical "tricorder" with molecular LIDAR. In Jason A. Guicheteau, Christopher R. Howle, and Tanya L. Myers, editors, *Chemical, Biological, Radiological, Nuclear, and Explosives (CBRNE) Sensing XXVI*, volume 13478, page 134780J. International Society for Optics and Photonics, SPIE, 2025.

[79] M. Hollmann, A. Walter, and F. Duschek. Evaluation of molecular hydrogen detection from distance by raman scattering. *Journal of Instrumentation*, 20(07):C07048, jul 2025.

[80] Turja Nandy, Ronald A. Coutu, and Cristinel Ababei. Carbon monoxide sensing technologies for next-generation cyber-physical systems. *Sensors (Switzerland)*, 18, 10 2018.

[81] A. A. Kosterev, Y. A. Bakhirkin, and F. K. Tittel. Ultrasensitive gas detection by quartz-enhanced photoacoustic spectroscopy in the fundamental molecular absorption bands region. *Applied Physics B: Lasers and Optics*, 80:133–138, 1 2005.

[82] S. E. Aleksandrov, G. A. Gavrilov, A. A. Kapralov, B. A. Matveev, M. A. Remennyi, and G. Yu Sotnikova. InAsSb diode optical pairs for real-time carbon dioxide sensors. *Technical Physics*, 63:1390–1395, 9 2018.

[83] O'Neill C. Improving power generation safety with advanced leak detection technology. *FLIR Application Note*, 2020.

[84] Rudra Pratap Singh, Subodh Kumar, Somil Dubey, and Aman Singh. A review on working and applications of oxy-acetylene gas welding. *Materials Today: Proceedings*, 38:34–39, 2021. 2nd International Conference on Future Learning Aspects of Mechanical Engineering.

[85] Ming Dong, Chuantao Zheng, Dan Yao, Guoqiang Zhong, Shuzhuo Miao, Weilin Ye, Yiding Wang, and Frank K. Tittel. Double-range near-infrared acetylene detection using a dual spot-ring herriott cell (DSR-HC). *Opt. Express*, 26(9):12081–12091, Apr 2018.

[86] Juliet Attah, Latifatu Mohammed, Andrew Nyamful, Paulina Donkor, Anita Asamoah, Mohammed Nafiu Zainudeen, John Adjah, Charles K. Klutse, Sylvester Attakorah Birikorang, Frederick Agyemang, and Owiredu Gyampo. Oxy-hydrogen gas as a sustainable fuel for the welding industry: Alternative for oxy-acetylene gas. *Cleaner Energy Systems*, 9:100160, 2024.

[87] Anirban Roy, Arup Lal Chakraborty, and Kenneth T. V. Grattan. A 1531.52 nm tunable diode laser-based system for high-sensitivity remote monitoring of acetylene and ammonia leaks in industrial plants. In Kehar Singh, A. K. Gupta, Sudhir Khare, Nimish Dixit, and Kamal Pant, editors, *ICOL-2019*, pages 379–382, Singapore, 2021. Springer Singapore.

[88] Y. Cao, W. Jin, H. L. Ho, L. Qi, and Y. H. Yang. Acetylene detection based on diode laser QEPAS: combined wavelength and residual amplitude modulation. *Applied Physics B*, 109(2):359–366, November 2012.

[89] Utsav KC, Ehson F. Nasir, and Aamir Farooq. A mid-infrared absorption diagnostic for acetylene detection. *Applied Physics B*, 120(2):223–232, August 2015.

[90] Shengyao Shi, Yi Li, Zhaolun Cui, Yongxu Yan, Xiaoxing Zhang, Ju Tang, and Song Xiao. Recent advances in degradation of the most potent industrial greenhouse gas sulfur hexafluoride. *Chemical Engineering Journal*, 470:144166, 2023.

[91] P. G. Simmonds, M. Rigby, A. J. Manning, S. Park, K. M. Stanley, A. McCulloch, S. Henne, F. Graziosi, M. Maione, J. Arduini, S. Reimann, M. K. Vollmer, J. Mühle,

S. O'Doherty, D. Young, P. B. Krummel, P. J. Fraser, R. F. Weiss, P. K. Salameh, C. M. Harth, M.-K. Park, H. Park, T. Arnold, C. Rennick, L. P. Steele, B. Mitrevski, R. H. J. Wang, and R. G. Prinn. The increasing atmospheric burden of the greenhouse gas sulfur hexafluoride (SF_6). *Atmospheric Chemistry and Physics*, 20(12):7271–7290, 2020.

[92] L.G. Christophorou, J.K. Olthoff, and R.J. Van Brunt. Sulfur hexafluoride and the electric power industry. *IEEE Electrical Insulation Magazine*, 13(5):20–24, 1997.

[93] Constantine T. Dervos and Panayota Vassiliou. Sulfur hexafluoride (SF6): Global environmental effects and toxic byproduct formation. *Journal of the Air and Waste Management Association*, 50:137–141, 2000.

[94] Jeremy J. Harrison. New infrared absorption cross sections for the infrared limb sounding of sulfur hexafluoride (SF6). *Journal of Quantitative Spectroscopy and Radiative Transfer*, 254:107202, 2020.

[95] Xu Zhang, Weiqi Jin, Pan Yuan, Chao Qin, Hongchen Wang, Ji Chen, and Xuan Jia. Research on passive wide-band uncooled infrared imaging detection technology for gas leakage. In Juan Liu, Baohua Jia, Xincheng Yao, Yongtian Wang, and Takanori Nomura, editors, *2019 International Conference on Optical Instruments and Technology: Optical Systems and Modern Optoelectronic Instruments*, volume 11434, page 114340H. International Society for Optics and Photonics, SPIE, 2020.

[96] Angus G. MacGruer, Graham M. Gibson, Kyle J. Nutt, and Miles J. Padgett. Real-time imaging of methane gas from a UAV mounted system. In *Remote Sensing Technologies and Applications in Urban Environments VIII*, volume 12735, page 127350J. International Society for Optics and Photonics, SPIE, 2023.

[97] Angus G. MacGruer, Steven D. Johnson, Simon P. Mekhail, Kyle J. Nutt, Miles J. Padgett, and Graham M. Gibson. Characterising the performance of a drone-mounted real-time methane imaging system. *Scientific Reports*, 15(1):8787, March 2025.

[98] Angus G. MacGruer, Steven D. Johnson, Kyle J. Nutt, Miles J. Padgett, and Graham M. Gibson. Drone-based gas leak detection system for use in industry. In *Quantum Technology: Driving Commercialisation of an Enabling Science IV*, volume 12795, page 1279503. International Society for Optics and Photonics, SPIE, 2023.

[99] Sergei K. Turitsyn, Sergey A. Babin, Dmitry V. Churkin, Ilya D. Vatnik, Maxim Nikulin, and Evgenii V. Podivilov. Random distributed feedback fibre lasers. *Physics Reports*, 542(2):133–193, 2014. Random Distributed Feedback Fibre Lasers.

[100] Michalis N. Zervas. Chapter 1 - advances in fiber distributed-feedback lasers. In Ivan P. Kaminow, Tingye Li, and Alan E. Willner, editors, *Optical Fiber Telecommunications (Sixth Edition)*, Optics and Photonics, pages 1–24. Academic Press, Boston, sixth edition edition, 2013.

[101] Mark G Allen. Diode laser absorption sensors for gas-dynamic and combustion flows. *Measurement Science and Technology*, 9(4):545, apr 1998.

[102] W. Trzeciakowski, A. Bercha, F. Dybala, R. Bohdan, P. Adamiec, and O. Mariani. Pressure and temperature tuning of laser diodes. *physica status solidi (b)*, 244(1):179–186, 2007.

[103] N. Hashizume and H. Nasu. Mode hopping control and lasing wavelength stabilization of fiber grating lasers. *Furukawa Electric*, pages 7–10, 04 2001.

[104] V.P. Duraev, A.A. Marmalyuk, and A.V. Petrovskiy. Tunable laser diodes for the 1250–1650nm spectral range. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, 66(4):846–848, 2007.

[105] Michael A Johnson and Mohammad H Moradi. *PID control*. Springer, 2005.

[106] Suhavi Kaur, Zhe F. Tang, and David R. McMillen. A framework to enhance the signal-to-noise ratio for quantitative fluorescence microscopy. *PLoS One*, 20(9):18, 09 2025. Copyright - © 2025 Kaur et al. This is an open access article distributed under the terms of the Creative Commons Attribution License: http://creativecommons.org/licenses/by/4.0/ (the "License"), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. Notwithstanding the ProQuest Terms and Conditions, you may use this content in accordance with the terms of the License; Last updated - 2025-09-05; SubjectsTermNotLitGenreText - Environmental.

[107] G. M. Gibson, M. Ebrahimzadeh, M. J. Padgett, and M. H. Dunn. Continuous-wave optical parametric oscillator based on periodically poled KTiOPO4 and its application to spectroscopy. *Opt. Lett.*, 24(6):397–399, Mar 1999.

[108] Stephen Monk Graham Gibson and Miles Padgett. A field-portable, laser-diode spectrometer for the ultra-sensitive detection of hydrocarbon gases. *Journal of Modern Optics*, 49(5-6):769–776, 2002.

[109] Kenneth D. Skeldon, Claire Patterson, Cathy A. Wyse, Graham M. Gibson, Miles J. Padgett, Chris Longbottom, and Lesley C. McMillan. The potential offered by real-time, high-sensitivity monitoring of ethane in breath and some pilot studies using optical spectroscopy. *Journal of Optics A: Pure and Applied Optics*, 7(6):S376–S384, June 2005.

[110] Kenneth D. Skeldon, Graham M. Gibson, Cathy A. Wyse, Lesley C. McMillan, Steve D. Monk, Chris Longbottom, and Miles J. Padgett. Development of high-resolution real-time sub-ppb ethane spectroscopy and some pilot studies in life science. *Appl. Opt.*, 44(22):4712–4721, Aug 2005.

[111] Jo Etienne Abela, Kenneth Skeldon, Robert Stuart, and Miles Padgett. Exhaled ethane concentration in patients with cancer of the upper gastrointestinal tract - a proof of concept study. *Bioscience trends*, 3:110–4, 06 2009.

[112] Bill Hirst, Graham Gibson, Steve Gillespie, Ian Archibald, O. Podlaha, Kenneth Skeldon, Johannes Courtial, Steve Monk, and Miles Padgett. Oil and gas prospecting by ultra-sensitive optical gas detection with inverse gas dispersion modelling. *Geophysical Research Letters*, 31:L12115, 06 2004.

[113] Ben Well, Stuart Murray, Jane Hodgkinson, Russ Pride, Rainer Strzoda, Graham Gibson, and Miles Padgett. An open-path, hand-held laser system for the detection of methane gas. *Journal of Optics A Pure and Applied Optics*, 7:S420–S424, 05 2005.

[114] Graham Gibson, Ben Van Well, Jane Hodgkinson, Russ Pride, Rainer Strzoda, Stuart Murray, Steve Bishton, and Miles Padgett. Imaging of methane gas using a scanning, open-path laser system. *New Journal of Physics*, 8, 2 2006.

[115] David J. M. Stothard, Malcolm H. Dunn, and Cameron F. Rae. Hyperspectral imaging of gases with a continuous-wave pump-enhanced optical parametric oscillator. *Optics Express*, 12:947, 2004.

[116] Fei Wang, Shuhai Jia, Yonglin Wang, and Zhenhua Tang. Recent developments in modulation spectroscopy for methane detection based on tunable diode laser. *Applied Sciences*, 9(14):2816, July 2019.

[117] Graham M. Gibson, Baoqing Sun, Matthew P. Edgar, David B. Phillips, Nils Hempler, Gareth T. Maker, Graeme P. A. Malcolm, and Miles J. Padgett. Real-time imaging of methane gas leaks using a single-pixel camera. *Optics Express*, 25:2998, 2 2017.

[118] Matthew P. Edgar, Graham M. Gibson, Richard W. Bowman, Baoqing Sun, Neal Radwell, Kevin J. Mitchell, Stephen S. Welsh, and Miles J. Padgett. Simultaneous real-time visible and infrared video with single-pixel detectors. *Scientific Reports*, 5, 5 2015.

[119] H. Foroosh, J.B. Zerubia, and M. Berthod. Extension of phase correlation to subpixel registration. *IEEE Transactions on Image Processing*, 11(3):188–200, 2002.

[120] Kaaren May and Nicholas Krouglicof. Moving target detection for sense and avoid using regional phase correlation. In *2013 IEEE International Conference on Robotics and Automation*, pages 4767–4772, 2013.

[121] Hui Wang, Jing Zhao, Jiwen Zhao, Fei Dong, Zhenbao Pan, and Yinyi Feng. Position detection method of linear motor mover based on extended phase correlation algorithm. *IET Science, Measurement & Technology*, 11(7):921–928, 2017.

[122] Zhelin Yu, Lidong Zhu, and Guoyu Lu. An improved phase correlation method for stop detection of autonomous driving. *IEEE Access*, 8:77972–77986, 2020.

[123] James W. Cooley, Peter A. W. Lewis, and Peter D. Welch. The fast fourier transform and its applications. *IEEE Transactions on Education*, 12(1):27–34, 1969.

[124] K A M Said and A B Jambek. Analysis of image processing using morphological erosion and dilation. *Journal of Physics: Conference Series*, 2071(1):012033, oct 2021.

[125] Brandon Stark, Matthew McGee, and YangQuan Chen. Short wave infrared (SWIR) imaging systems using small unmanned aerial systems (sUAS). In *2015 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 495–501, 2015.

[126] Tristan M. Goss, Henning Fourie, and Johan W. Viljoen. SWIR sensor design considerations. In Monuko du Plessis, editor, *Fifth Conference on Sensors, MEMS, and Electro-Optic Systems*, volume 11043, page 110430B. International Society for Optics and Photonics, SPIE, 2019.

[127] Marc P. Hansen and Douglas S. Malchow. Overview of SWIR detectors, cameras, and applications. In Vladimir P. Vavilov and Douglas D. Burleigh, editors, *Thermosense XXX*, volume 6939, page 69390I. International Society for Optics and Photonics, SPIE, 2008.

[128] M. Kajikawa and I. Kataoka. A design method of optical bandpass filters. *Journal of Lightwave Technology*, 15(9):1720–1727, 1997.

[129] Arup Kumar Ghosh, Gautam Garai, and Subhankar Bhattacharjee. Distributed adaptive differential evolution algorithm for mathematical functions optimization and development of optical band pass filter. *Optical and Quantum Electronics*, 57(1):74, December 2024.

[130] A. Donval, T. Fisher, G. Blecher, and M. Oron. Dynamic Sunlight Filter (DSF): a passive way to increase the dynamic range in visible and SWIR cameras. In Bjørn F. Andresen, Gabor F. Fulop, and Paul R. Norton, editors, *Infrared Technology and Applications XXXVI*, volume 7660, page 766024. International Society for Optics and Photonics, SPIE, 2010.

[131] G. Pfister, R. L. McKenzie, J. B. Liley, A. Thomas, B. W. Forgan, and C. N. Long. Cloud coverage based on all-sky imaging and its impact on surface solar irradiance. *Journal of Applied Meteorology*, 42(10):1421 – 1434, 2003.

[132] Linshen Yao, Yang Shang, and Dan Fu. Influence of atmospheric turbulence on optical measurement: a case report and review of literature. In Weimin Bao and Yueguang Lv, editors, *Selected Papers of the Photoelectronic Technology Committee Conferences held November 2015*, volume 9796, page 979604. International Society for Optics and Photonics, SPIE, 2016.

[133] Lindsey Wiley, Richard Cavanaugh, Joshua Follansbee, Derek Burrell, Robert Grimming, Rich Pimpinella, Jeff Voss, Orges Furxhi, and Ronald Driggers. Comparison of reflective band (Vis, NIR, SWIR, eSWIR) performance in daytime reduced illumination conditions. *Appl. Opt.*, 62(31):8316–8326, Nov 2023.

[134] L. Wang, S. Kassi, A.W. Liu, S.M. Hu, and A. Campargue. High sensitivity absorption spectroscopy of methane at 80K in the 1.58μm transparency window: Temperature dependence and importance of the CH3D contribution. *Journal of Molecular Spectroscopy*, 261(1):41–52, 2010.

[135] Buyue Duan, Yan Lv, Zhiguo Wang, Di Wang, Dong Li, and Bin Yao. Correlation of absorption spectrum properties of methane with coupling variables between temperature and concentration. *Optik*, 185:82–87, 2019.

[136] V. Nagali, S. I. Chou, D. S. Baer, R. K. Hanson, and J. Segall. Tunable diode-laser absorption measurements of methane at elevated temperatures. *Appl. Opt.*, 35(21):4026–4032, Jul 1996.

[137] O. M. Coddington, E. C. Richard, D. Harber, P. Pilewskie, T. N. Woods, K. Chance, X. Liu, and K. Sun. The tsis-1 hybrid solar reference spectrum. *Geophysical Research Letters*, 48(12):e2020GL091709, 2021. e2020GL091709 2020GL091709.

[138] Scott A. Lerner and Brett Dahlgrenn. Etendue and optical system design. In Roland Winston and Pablo Benítez, editors, *Nonimaging Optics and Efficient Illumination Systems III*, volume 6338, page 633801. International Society for Optics and Photonics, SPIE, 2006.

[139] Parham Nooralishahi, Fernando López, and Xavier Maldague. A drone-enabled approach for gas leak detection using optical flow analysis. *Applied Sciences*, 11(4), 2021.

[140] Bara J. Emran, Dwayne D. Tannant, and Homayoun Najjaran. Low-altitude aerial methane concentration mapping. *Remote Sensing*, 9(8):823, 2017.

[141] Kouki Ikuta, Noboru Yoshikane, Nilesh Vasa, Yuji Oki, Mitsuo Maeda, Michihiro Uchiumi, Youichirou Tsumura, Jun Nakagawa, and Noriyuki Kawada. Differential absorption lidar at 1.67 μm for remote sensing of methane leakage. *Japanese Journal of Applied Physics*, 38(Part 1, No. 1A):110–114, 1999.

[142] Rebecca V. Hiller, Bruno Neininger, Dominik Brunner, Christoph Gerbig, Daniel Bretscher, Thomas Künzle, Nina Buchmann, and Werner Eugster. Aircraft-based CH4 flux estimates for validation of emissions from an agriculturally dominated area in Switzerland. *Journal of Geophysical Research: Atmospheres*, 119(8):4874–4887, 2014.

[143] S. J. O'Shea, S. J.-B. Bauguitte, M. W. Gallagher, D. Lowry, and C. J. Percival. Development of a cavity-enhanced absorption spectrometer for airborne measurements of CH_4 and CO_2. *Atmospheric Measurement Techniques*, 6(5):1095–1109, 2013.

[144] Thomas E. Barchyn, Chris H. Hugenholtz, Stephen Myshak, and Jim Bauer. A uav-based system for detecting natural gas leaks. *Journal of Unmanned Vehicle Systems*, 6(1):18–30, 2018.

[145] E. Förster, H. Huntrieser, M. Lichtenstern, F. Pätzold, L. Bretschneider, A. Schlerf, S. Bollmann, A. Lampert, J. Nęcki, P. Jagoda, J. Swolkień, D. Pasternak, R. A. Field, and A. Roiger. A helicopter-based mass balance approach for quantifying methane emissions from industrial activities, applied for coal mine ventilation shafts in poland. *Atmospheric Measurement Techniques*, 18(23):7153–7176, 2025.

[146] Evan D. Sherwin, Yuanlei Chen, Arvind P. Ravikumar, and Adam R. Brandt. Single-blind test of airplane-based hyperspectral methane detection via controlled releases. *Elementa: Science of the Anthropocene*, 9(1):00063, March 2021. _eprint: https://online.ucpress.edu/elementa/article-pdf/9/1/00063/459354/elementa.2021.00063.pdf.

[147] E. G. Nisbet, R. E. Fisher, D. Lowry, J. L. France, G. Allen, S. Bakkaloglu, T. J. Broderick, M. Cain, M. Coleman, J. Fernandez, G. Forster, P. T. Griffiths, C. P. Iverach, B. F. J. Kelly, M. R. Manning, P. B. R. Nisbet-Jones, J. A. Pyle, A. Townsend-Small, A. al Shalaan, N. Warwick, and G. Zazzeri. Methane mitigation: Methods to reduce emissions, on the path to the paris agreement. *Reviews of Geophysics*, 58(1):e2019RG000675, 2020. e2019RG000675 2019RG000675.

[148] D. J. Jacob, A. J. Turner, J. D. Maasakkers, J. Sheng, K. Sun, X. Liu, K. Chance, I. Aben, J. McKeever, and C. Frankenberg. Satellite observations of atmospheric methane and their value for quantifying methane emissions. *Atmospheric Chemistry and Physics*, 16(22):14371–14396, 2016.

[149] Sudhanshu Pandey, Ritesh Gautam, Sander Houweling, Hugo Denier van der Gon, Pankaj Sadavarte, Tobias Borsdorff, Otto Hasekamp, Jochen Landgraf, Paul Tol, Tim van Kempen, Ruud Hoogeveen, Richard van Hees, Steven P. Hamburg, Joannes D. Maasakkers, and Ilse Aben. Satellite observations reveal extreme methane leakage from a natural gas well blowout. *Proceedings of the National Academy of Sciences*, 116(52):26376–26381, 2019.

[150] Levi M. Golston, Lei Tao, Caroline Brosy, Klaus Schäfer, Benjamin Wolf, James McSpiritt, Bernhard Buchholz, Dana R. Caulton, Da Pan, Mark A. Zondlo, David Yoel,

Harald Kunstmann, and Marty McGregor. Lightweight mid-infrared methane sensor for unmanned aerial systems. *Applied Physics B: Lasers and Optics*, 123, 6 2017.

[151] Oscar Alvear, Nicola Roberto Zema, Enrico Natalizio, and Carlos T. Calafate. Using UAV-based systems to monitor air pollution in areas with poor accessibility. *Journal of Advanced Transportation*, 2017:1–14, 2017.

[152] Tommasofrancesco Villa, Felipe Gonzalez, Branka Miljevic, Zoran D. Ristovski, and Lidia Morawska. An overview of small unmanned aerial vehicles for air quality measurements: Present applications and future prospectives. *Sensors (Switzerland)*, 16, 7 2016.

[153] Maurizio Rossi, Davide Brunelli, Andrea Adami, Leandro Lorenzelli, Fabio Menna, and Fabio Remondino. Gas-drone: Portable gas sensing system on UAVs for gas leakage localization. In *SENSORS, 2014 IEEE*, pages 1431–1434, 2014.

[154] J. Pablo Arroyo-Mora, Margaret Kalacska, Deep Inamdar, Raymond Soffer, Oliver Lucanus, Janine Gorman, Tomas Naprstek, Erica Skye Schaaf, Gabriela Ifimov, Kathryn Elmer, and George Leblanc. Implementation of a uav–hyperspectral pushbroom imager for ecological monitoring. *Drones*, 3(1), 2019.

[155] E. Armandillo, D. Stepanova, D. Rees, I. Ramirez, I. Mishin, I. Chemezov, H. Leuschner, L. Fries, F. Galassi, and M. Elhag. GHG monitoring from CubeSat using compact microLiDAR and pushbroom spectrometer. In Max Petrozzi-Ilstad, editor, *Small Satellites Systems and Services Symposium (4S 2024)*, volume 13546, page 135462Y. International Society for Optics and Photonics, SPIE, 2025.

[156] Ryohei Sato, Kento Tanaka, Hanako Ishida, Saki Koguchi, Jane Pauline Ramos Ramirez, Haruka Matsukura, and Hiroshi Ishida. Detection of gas drifting near the ground by drone hovering over: Using airflow generated by two connected quadcopters. *Sensors*, 20(5):1397, 2020.

[157] Javier Burgués and Santiago Marco. Environmental chemical sensing using small drones: A review. *Science of the Total Environment*, 748, 12 2020.

[158] Josefa Wivou, Lanka Udawatta, Ali Alshehhi, Ebrahim Alzaabi, Ahmed Albeloshi, and Saeed Alfalasi. Air quality monitoring for sustainable systems via drone based technology. In *2016 IEEE International Conference on Information and Automation for Sustainability (ICIAfS)*, pages 1–5, 2016.

[159] T. Andersen, B. Scheeren, W. Peters, and H. Chen. A UAV-based active aircore system for measurements of greenhouse gases. *Atmospheric Measurement Techniques*, 11(5):2683–2699, 2018.

[160] J Hodgkinson, Q Shan, and R D Pride. Detection of a simulated gas leak in a wind tunnel. *Measurement Science and Technology*, 17(6):1586, may 2006.

[161] Virginio Cantoni, Stefano Levialdi, and Bertrand Zavidovique. 2 - visual cues. In Virginio Cantoni, Stefano Levialdi, and Bertrand Zavidovique, editors, *3C Vision*, pages 19–113. Elsevier, London, 2011.

[162] Colin Ware. Chapter two - the environment, optics, resolution, and the display. In Colin Ware, editor, *Information Visualization (Third Edition)*, Interactive Technologies, pages 31–68. Morgan Kaufmann, Boston, third edition edition, 2013.

[163] H. E. Bennett and J. O. Porteus. Relation between surface roughness and specular reflectance at normal incidence. *J. Opt. Soc. Am.*, 51(2):123–129, Feb 1961.

[164] L. C. Wye, H. A. Zebker, and R. D. Lorenz. Smoothness of titan's ontario lacus: Constraints from cassini radar specular reflection data. *Geophysical Research Letters*, 36(16), 2009.

[165] Vittorio Degli-Esposti, Franco Fuschini, Enrico M. Vitucci, and Gabriele Falciasecca. Measurement and modelling of scattering from buildings. *IEEE Transactions on Antennas and Propagation*, 55(1):143–153, 2007.

[166] H. Ragheb and E.R. Hancock. Lambertian reflectance correction for rough and shiny surfaces. In *Proceedings. International Conference on Image Processing*, volume 2, pages II–II, 2002.

[167] A. Richards and M. Hübner. A new radiometric unit of measure to characterize SWIR illumination. In Gerald C. Holst and Keith A. Krapels, editors, *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXVIII*, volume 10178, page 101780C. International Society for Optics and Photonics, SPIE, 2017.

[168] Daniel R. Jones, Waheed A. Al-Masry, and Charles W. Dunnill. Hydrogen-enriched natural gas as a domestic fuel: an analysis based on flash-back and blow-off limits for domestic natural gas appliances within the uk. *Sustainable Energy Fuels*, 2:710–723, 2018.

[169] J.R. Gibson. *Electronic Logic Circuits - Third Edition*. Arnold, London, 1992.

[170] P. Horowitz and W. Hill. *Art of Electronics - Second Edition*. Cambridge University Press, Cambridge, 1989.

[171] Ricardo Araguillin, Ángel Méndez, José González, and Ćesar Costa-Vera. Comparative evaluation of wavelength-scanning otto and kretschmann configurations of spr biosensors for low analyte concentration measurement. *Journal of Physics: Conference Series*, 2796(1):012009, jul 2024.

[172] Flir: Optical gas imaging: https://www.flir.co.uk/instruments/optical-gas-imaging/. Date accessed: 17/05/2025.

[173] Jingfan Wang, Jingwei Ji, Arvind P. Ravikumar, Silvio Savarese, and Adam R. Brandt. Videogasnet: Deep learning for natural gas methane leak classification using an infrared camera. *Energy*, 238:121516, 2022.

[174] Julien Fade, Swapnesh Panigrahi, Anthony Carré, Ludovic Frein, Cyril Hamel, Fabien Bretenaker, Hema Ramachandran, and Mehdi Alouini. Long-range polarimetric imaging through fog. *Applied Optics*, 53:3854, 6 2014.

[175] O. Wolley, S. Mekhail, P.-A. Moreau, T. Gregory, G. Gibson, G. Leuchs, and M. J. Padgett. Near single-photon imaging in the shortwave infrared using homodyne detection. *Proceedings of the National Academy of Sciences*, 120(10):e2216678120, 2023.