

Schlieren-Based Flow Imaging

by

Bradley Atcheson

B.Sc., University of the Witwatersrand, 2004

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

The Faculty of Graduate Studies

(Computer Science)

The University Of British Columbia

July, 2007

© Bradley Atcheson 2007

Abstract

A transparent medium of inhomogenous refractive index will cause light rays to bend as they pass through it, resulting in a visible distortion of the background. We present a simple 2D imaging method for measuring this distortion and then show how it can be used to visualise gas and liquid flows. Improvements to the existing Background Oriented Schlieren method for acquiring projected density gradients are made by placing a multi-scale wavelet noise pattern behind the flow and measuring distortions using a more reliable optical flow algorithm. Dynamic environment mattes can also be acquired, allowing us to render the flow into novel scenes. Capturing from multiple viewpoints then allows us to tomographically reconstruct 3D models of the temperature distribution of the fluid.

Contents

Abstract	ii
Contents	iii
List of Tables	v
List of Figures	vi
Acknowledgements	viii
1 Introduction	1
1.1 Motivation and Applications	3
1.2 Overview of the Method	6
1.3 Contribution	8
2 Schlieren Imaging	10
2.1 Optical Principle	10
2.1.1 Lens-Based Systems	11
2.2 Other Schlieren Configurations	14
2.2.1 Rainbow Schlieren	14
2.2.2 Grid-Based Systems	15
2.2.3 Background Oriented Schlieren	16
2.3 BOS and Shadowgraphy	20
3 Optical Flow	22
3.1 Basic Concepts	23

3.2	Block-Matching Methods	24
3.2.1	Spatial Correlation	25
3.3	Gradient-Based Methods	28
3.3.1	Lucas-Kanade	28
3.3.2	Horn-Schunck	29
3.4	Variational Methods	30
3.5	Schlieren-Specific Optical Flow	31
3.6	Wavelet Noise Background	32
3.7	Conclusion	34
4	Physical Setup	35
4.1	Acquisition Setups	35
4.1.1	Standard Setup	35
4.1.2	Camera Array	36
4.1.3	Lens Array	36
4.2	Angle Measurement	38
4.3	Implementation Notes	38
5	Results and Applications	40
5.1	Optical Flow Performance	40
5.2	Flow Visualisation	42
5.3	Environment Matting	46
5.4	Shadowgraphs	46
5.5	Tomographic Reconstruction	47
5.6	Conclusions	49
	Bibliography	51

List of Tables

5.1	Mean squared error of reconstructed fields.	41
5.2	Mean squared error after autocorrelation subtraction.	42

List of Figures

1.1	A vapour plume rises from burning alcohol.	3
1.2	Particle image velocimetry setup.	5
1.3	Interaction of candle plume with compressed air jet.	7
2.1	Standard high-level configuration of any Schlieren system.	11
2.2	A simple (idealised) lens-based Schlieren configuration.	11
2.3	Example Schlieren photographs.	13
2.4	Rainbow Schlieren filter.	15
2.5	Grid-based Schlieren system.	16
2.6	Virtual displacement caused by ray deflection.	17
2.7	Shadowgraph example.	20
3.1	Deformation applied to block-matching windows.	24
3.2	Cross-correlation matrix.	26
3.3	Fixed-pattern noise.	27
3.4	Horn-Schunck optical flow field for candle plume.	30
3.5	Random noise patterns.	33
3.6	Wavelet noise patterns.	34
4.1	Standard single camera setup and small candle.	36
4.2	Camera array setup and gas burner.	37
4.3	Lens array setup and water tank.	38
5.1	Synthetic optical flow tests.	41
5.2	Optical flow scaling test.	43

5.3	Candle plume results.	44
5.4	Captured images from plume interaction test.	44
5.5	Optical flow vectors from plume interaction test.	45
5.6	Corn syrup solution being injected into water.	45
5.7	Candle lantern environment matting example.	47
5.8	Synthetic shadowgraph example.	48
5.9	Tomography results.	50

Acknowledgements

Firstly, I'd like to thank everybody involved with this project for their comments, suggestions and proofreading. In particular, Wolfgang Heidrich, Ivo Ihrke, Derek Bradley and Marcus Magnor provided invaluable support and assistance.

To my colleagues in the Imager and PSM labs, my friends, new and old, my family back home, and Whistler, thank you for always being there and for helping me along the way.

I am also grateful to the Canadian Commonwealth Scholarship Programme for funding my studies, and to Kasi Metcalfe and Andrew Davidhazy for allowing me to use their images.

Chapter 1

Introduction

Humans have a natural tendency to want to understand the world around us. This classically involves making observations, devising theories and testing to see if predicted outcomes align with measured results. Physical phenomena and natural processes such as fire, cloud formation, plant growth, human behaviour etc. have all been studied in this way. In these cases, the subject is directly observable. Sometimes, however, we are only able to infer the existence of a process by its effects. Gravity itself is not visible, but through our observations of its interaction with objects of mass we know that it exists. The air around us, though at first glance stable, uninteresting, and *invisible*, is actually host to an extremely complex flow of gas. All around us, slight temperature and pressure variations in air and liquid currents, as well as slight imperfections in transparent solid materials create inhomogeneities in the refractive index of the medium. These are points at which the speed of light changes ever so slightly, causing rays to bend fractionally. Since they do not emit, absorb or reflect light, at a casual glance these inhomogeneities are entirely invisible, but to a careful observer they manifest themselves as a slight warping of the background seen through the media. Through their effects on the environment (*i.e.*, how they refract the light passing through them) we can learn their structure. With the right instruments, this invisible world can be observed, recorded, and subsequently illustrated for all to see in whatever manner we choose to render the information. The aim of this thesis is to present one such instrument and demonstrate its applications.

The deflection of light is commonly exploited by scanning and acquisition devices. Structured light range scanners project sheets of light onto solid, opaque objects, and measure the resultant large-scale deformations to the projected stripe(s) in order to

estimate scene depth [35]. At the other extreme, interferometry can be used to measure surface profiles at the nanometre scale by illuminating a target with coherent laser light and observing the resultant patterns of constructive and destructive wave interference [19]. We shall restrict ourselves to the class of media producing light deflections on a scale that is just barely visible to the naked eye. In this case the refractive indices of the medium and its surroundings differ by as little as 10^{-4} , resulting in a small yet still noticeable amount of refraction. Examples of this are the rising plume of heated air above a candle flame, the mixing of water and alcohol, and imperfectly formed panes of glass. This scale is large enough to be easily measured, yet small enough to have escaped much research interest in the past, and is useful in studying fluid flow. It is also large enough to be modelled using only geometric optics, allowing for a conceptually simple system.

The *Schlieren*¹ method has for centuries enabled scientists to observe the way in which light refracts when passing through these media. Although the presence of an inhomogeneity in the refractive index of the medium is often detectable without any aids, a cleverly-designed optical setup can be employed to produce a far superior view of it. Figure 1.1 shows the type of results that can be obtained. In this case, the flame itself is very small and no smoke is produced. In an ordinary photograph the only hint of the plume would be a slight distortion of the background, but even then we would not be able to discern nearly as much detail about its shape as is evident in the Schlieren photograph. Chapter 2 contains a more detailed discussion of the optical principle behind Schlieren imaging, and describes some of the classic instruments.

Although the images produced via Schlieren photography can be beautiful to look at, they are nonetheless only useful as a qualitative tool. It could be useful for example to notice the *presence* of a shock wave, but for a more thorough analysis, to *measure* the properties of the refracting medium, we require more quantitative information. Some progress has been made in quantitative Schlieren methods (Section 2.2.1), but these typically require extensive calibration and can be extremely difficult and tedious to set up, not to mention being rather expensive. A well-tuned Schlieren device is a

¹The German word for “streak”

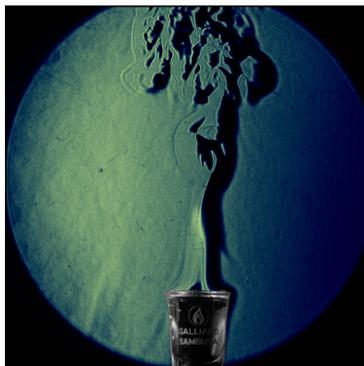


Figure 1.1: A vapour plume rises from burning alcohol. Photograph courtesy of Kasi Metcalfe.

remarkably elegant tool, and encompasses all of the complexity in an analogue optical setup. Once constructed and calibrated it requires almost no effort to take striking Schlieren photographs. However, the initial difficulty of construction is a significant barrier to wider adoption and development. The advent of enabling technologies such as high-resolution cameras and powerful desktop computers now allow us to shift the complexity from the optical setup to the post-processing side. Thus, *digital Schlieren* techniques, which perform heavy analysis on data obtained through very simple means, are now more attractive for the applications considered in this thesis (the only downside being the not-insignificant time and energy costs incurred in post-processing). One hopes that the loss of speed and elegance in having nature perform the computations for us directly in the optical equipment is made up for by the ease with which we are able to use these new digital methods.

1.1 Motivation and Applications

Schlieren photographs are useful because they allow us to *see* fluid flow. While the physics-based animation community searches for efficient ways to accurately simulate fluid flow [26], it would be useful to have ground truth data available against which the simulations could be verified. And if this data were acquired at the scale described above, we would have a natural benchmark against which to render flow that is just

physically accurate enough to *look right*.

The aeronautical industry attempts to directly measure air flow in wind tunnels to improve the performance of aerofoils and nozzles [38]. Aerodynamics, pipeline flow, even ink-jet print head design are all applications which require a good understanding of fluid flow near the object being designed. There are two main ways to go about measuring these flows:

- *inserting fixed probes into the flow*: the gas temperature around a nozzle can be directly measured with thin thermocouples. Unfortunately, the probes themselves impede the flow, and many are required in order to obtain high resolution data.
- *seeding the flow with markers*: when liquids interact, a dye can be used to follow the flow of one of them, but only until it diffuses across the full volume. For gases, tiny particles can be inserted into the flow. This is in fact the basis of a large branch of measurement science called *particle image velocimetry* (PIV) [17]. Since the particles have mass and thus inertia, they do not follow the gas flow exactly and thus should be as small as possible to minimise the discrepancy. These particles are illuminated by a pulsed sheet of laser light, and images are recorded by a camera, as shown in Figure 1.2. Analysing the motion of the particles over time reveals the fluid flow. Problems with this method stem from the fact that complex correction has to be applied to overcome the particle inertia issue, and more importantly, seeding a flow with a sufficiently dense cloud of particles is often difficult and sometimes impossible. Additionally, the algorithms used in most current PIV system for measuring the particle motion are rather naïve by modern computer vision standards. It would be beneficial to have a method that allows us to observe fluid flows directly, even in the cases where we are unable to insert marker particles.

A hybrid Schlieren-PIV system has been suggested by Jonassen et. al. [24] although it is only applicable to turbulent flow. By treating the tiny, turbulent eddies revealed in the Schlieren photographs as virtual particles, and running standard PIV software on the images, a crude flowfield can be visualised - but only for the regions of

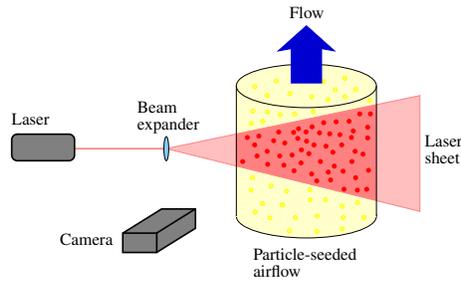


Figure 1.2: Particle image velocimetry setup.

high turbulence, and at a scale much larger than the internal dimensions of the eddies. The system described here works in cases of laminar flow and mild turbulence.

Aside from measuring fluid flows, a Schlieren-based acquisition setup has other potential applications, discussed further in Chapter 5:

Gas leak detection is not a trivial task since pipeline gases are often clear and odourless. However, these gases are generally optically dissimilar enough from air to be detectable via Schlieren methods. In this case we do not even require accurate flow data. Rather, it is simply the absence or presence of an inhomogeneity that we are interested in detecting. Traditional Schlieren devices are impractical in real-world environments, but as long as the background remains static our system is suitable.

Environment matting is the process of acquiring the properties (opacity, reflectivity, refractivity) of an optically active scene object and synthetically rendering it into new environments [49]. We capture the refraction in a controlled lab environment, and then render novel views (although our freedom to move and rotate the object is severely restricted) without ever requiring a full 3D model. A system capable of real-time acquisition allows for the capture of dynamic environment mattes. For example, the distortion caused by turbulent hot air above a fire can easily be acquired and used in visual effects applications. The necessity of having highly accurate distortion information for game and movie applications is debatable, but should one desire better quality results than those obtainable via simulations or random distortions, this method can be applied.

Tomographic reconstruction is one application where highly accurate distortion measurements *are* required. Under certain restrictions (which are met in our candle-flame class of refractive media) the total deflection of a light ray is a line integral of all the micro-deflections along the ray path. This is similar to the standard positron emission tomography problem where we record integrated absorption along the ray path. Therefore, with multiple cameras, we are able to record deflection data for dynamic media and reconstruct volumetric refractive index distributions of the scene using modified tomographic techniques, allowing us to build 3D models of the invisible field. Further discussion on this topic is provided in Chapter 5, but please note that this is the work of a separate project and does not form part of this thesis. Capturing the refractive index is the key point here. Rather than trying to obtain the flow directly via probes or PIV, we can acquire a physical quantity (index) which is related to the more interesting properties (temperature, pressure) via well-known physical laws. As Fu and Wu [14] note, “for a homogenous medium, the refractive index is a function of the thermodynamic state”.

1.2 Overview of the Method

The recently proposed background oriented [digital] Schlieren (BOS) technique forms the basis of our acquisition system. Full details on its workings are given in Section 2.2.3, but briefly, the method works by filming a static background through the refracting medium, and measuring per-pixel distortions in the background (see Figure 1.3). We use optical flow to do this measurement, and Chapter 3 describes the algorithm we use, as well as justification for why it is more suitable than the algorithms used in standard PIV systems. In addition, the optical flow can be aided by choosing a background with favourable properties. Section 3.6 justifies our decision to use wavelet noise for this purpose [10]. The goal of BOS is to produce a map, for every camera ray passing through the medium, specifying its magnitude and angle of deflection. This information is sufficient to enable the applications described above.

The output we get from this process is the 2D projected gradient of the refractive index field (Figure 1.3 right) and not the flow itself. The air itself is not moving in the

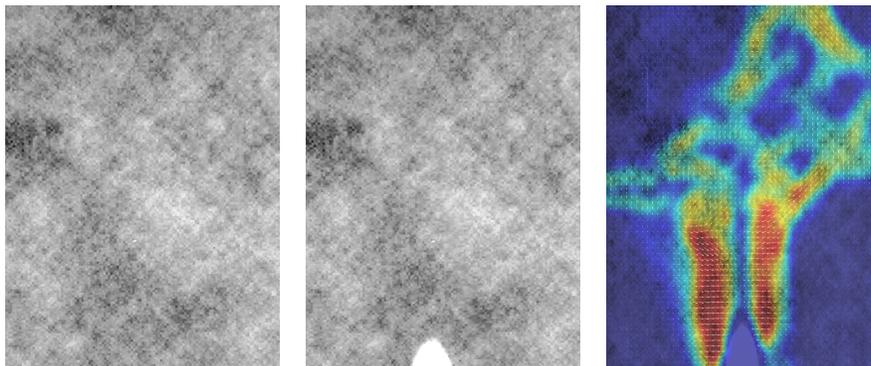


Figure 1.3: Reference background image (left) is distorted by the air above the candle flame (middle). Optical flow reveals the projected refractive index gradient (right). In this example we see the interaction of a candle plume being blown from the left by a can of compressed air.

direction of the arrows (in fact, it moves approximately orthogonally in this example), rather the arrows show the direction of deflection of the light rays passing through. However, an animated sequence of these images reveals very clearly the motion of the air.

The BOS method has many advantages, but does impose some limitations. In its favour, the remarkably simple setup is easy to construct, and the sensitivity and accuracy of the results are adequate for our purposes. By its nature, it allows for real-time acquisition of data, since all we require is a single image (per viewpoint) of the scene. An obvious benefit of this is that dynamic media are easily handled (a significant advantage in fact, as other scanning and tomography systems often require long acquisition times, forcing the subject to remain still), but not having anything more than a single image available forces us to make some assumptions about the system in order to interpret the results correctly:

- *transparency*: the medium should not absorb or scatter light. Optical flow algorithms attempt to match image patches in the background, and changes introduced when rays pass through different parts of the medium can cause this to break down. Ideally it will be perfectly transparent, but in practice a certain amount of uniform absorption can be tolerated (as the more robust optical flow

algorithms are insensitive to uniform changes in illumination).

- *smoothness*: the background distortion should be reasonably weak and smooth. The “objects” whose motion we wish to follow are essentially individual projected beams of light one pixel in size. These cannot be tracked via optical flow without a sufficient degree of spatial coherence.
- *no total internal reflection*: rays from the background must eventually strike the sensor if we are to hope to be able to determine their paths. No deflection measurements can be made in parts of the scene where total internal reflection occurs, and these regions will typically be smoothly interpolated. When dealing with complex glass shapes this will almost always be incorrect.
- *wavelength independence*: refraction has a dependence on wavelength which, for simplicity, is ignored here. The principle extends naturally to the general case, and colour sensors would allow for capturing spectral dispersion effects.

In addition to these assumptions, we cannot escape the fact that complete ray path information is not available from a single viewpoint. Although we can recover the start and end points (via camera calibration), the specific trajectory itself cannot be uniquely identified. Making further assumptions that the ray deflections are “small”, and the Schlieren volume itself is small relative to the spacing between the camera and the background, we can at least estimate the locations of additional points through which the ray must travel. Combining this measured and estimated information together makes the tomography possible, allowing us to understand the structure of the inhomogeneous field.

1.3 Contribution

The contribution of this thesis is that of providing a new acquisition tool to computer graphics. Inspired by the underlying Schlieren optical theory, we take a nascent acquisition setup (BOS) from the measurement science community and improve it through the use of better optical flow algorithms and specially designed background patterns.

While standard PIV processing software uses relatively weak optical flow algorithms to deal with difficult images, BOS allows us to select a background pattern that, when combined with a more reliable gradient-based optical flow algorithm, can reduce error by as much as 33% for the types of flow in which we are interested.

The use of Schlieren images has traditionally been restricted to studying static fluid interactions, but we recognise the similarities to PIV which make it possible to do more thorough dynamic fluid flow analysis - particularly when full 3D tomographic reconstruction is performed. In addition, the BOS setup serves as an alternative method for acquiring dynamic environment mattes of weakly-refracting transparent media with very little effort.

Chapter 2

Schlieren Imaging

A classical Schlieren imaging device can be thought of as a black box for converting angular deviations of light rays into intensity variations on an image plane. Many techniques exist for doing this, and this chapter begins with a simplified model of Toepler's classic lens-based system. A taxonomy of Schlieren instruments would be broken down into lens-and-mirror, grid, and background distortion-based systems. For completeness sake, a standard grid-based system is also described, and which prompts the conclusion that like lens-based designs, these are inappropriate for our applications. Background distortion emerges as a more suitable method, and background oriented Schlieren (BOS) is described in Section 2.2.3. Shadowgraphy is another imaging technique for viewing optical inhomogeneities, and despite having some similarities to Schlieren, it is quite distinct. An obvious difference is in the overall complexity of the method, with shadowgraphy being vastly simpler. Since BOS shares this trait, appears on the surface to be more like shadowgraphy, and is not clearly defined as being one or the other anywhere in the existing literature, the chapter concludes with an argument for why BOS is in fact a true Schlieren technique.

Note that it is possible to model Schlieren imaging in either the spatial domain using geometric ray optics (as done in this thesis) or in frequency space. Settles [42] has an excellent appendix on this for the interested reader, which describes how the Schlieren cutoff filter modifies the Fourier transform of the incident wavefront.

2.1 Optical Principle

All Schlieren techniques are based on the process illustrated in the schematic below. A light source emits rays of light which traverse the scan volume. If the volume is of

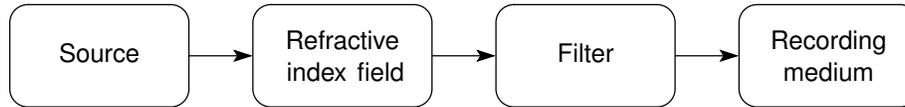


Figure 2.1: Standard high-level configuration of any Schlieren system. In some designs, the filter may lie before the field.

uniform refractive index, then the rays travel in straight lines. However, any inhomogeneities in the index will result in bending of the rays. A filter is chosen to modify the rays in some way depending on how they are bent. Finally, the recording medium captures an image of the modified light field and we can interpret this to *see* the refracting element.

2.1.1 Lens-Based Systems

The classic lens-based Schlieren method is illustrated in Figure 2.2. This design, the first described in detail by Settles [42], requires parallel light rays to pass through the field, so a lens is used to collimate the output of a point source. This is often the source of greatest expense (high quality lenses become extremely expensive beyond 100mm in diameter) and so parabolic mirrors are often substituted. The size of the beam of parallel light determines the maximum size of the scan volume that can be imaged.

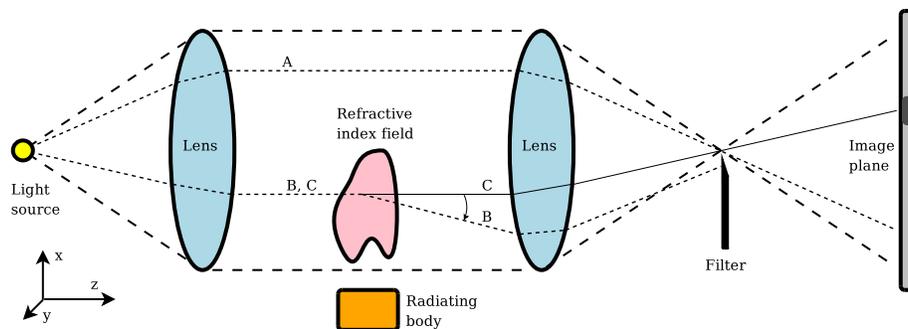


Figure 2.2: A simple (idealised) lens-based Schlieren configuration. In practice an extended light source has to be used in place of the point source, but save for an additional focusing lens and slightly more complex ray paths, the same principle applies.

Assume firstly that no refractive element is present, and consider the paths of rays *A* and *C*. They travel undisturbed through the scan volume until they encounter the second lens which focuses them back down to a point (actually, as near to a point as is possible given the diffraction limit). This point lies just above the fully-opaque filter, which in practice is often a razor blade. An image plane here would show a single spot, the image of the source. To produce a Schlieren image the recording plane is moved back, so that an inverted image of the scan volume is projected, as is clear from the ray paths.

Now consider what happens when refractions occur. Introducing a radiating body to heat the air inside the scan volume would cause its refractive index to change. Ray *B* (which would ordinarily follow the same path as *C*) is now deflected slightly (*very* slightly - the diagram is exaggerated for clarity). Instead of being focused down and passing above the filter, it now becomes obstructed and does not strike the recording medium. A dark spot thus occurs where *B* would have struck the image plane, had it not been filtered out.

This diagram is the key to understanding how it is possible to optically convert angular variations in light rays into intensity variations in an image. Of course, the overly simplified example here suggests a binary interpretation - either light reaches the sensor or it does not. In practice however, we use (small) area light sources which focus down not to a point, but to some finite area at the filter plane, and then a third lens focuses this to a point on the recording plane. The filter would be calibrated so that it blocks half of the light passing through (*i.e.*, it occludes half of the image of the extended light source) and then any deviations causing the beam to shift up (so that the filter blocks less of it) or down, would cause brighter, or respectively darker, spots to appear on the image plane. Figure 2.3 below shows some examples of typical results.

Constructing a small lens-based Schlieren device in a laboratory setting is quite possible for a few hundred dollars. Ideally we would like more flexibility in the tool, and the following disadvantages make it inadequate for our purposes:

- *expense*: as mentioned above, high quality lenses and mirrors are extremely ex-

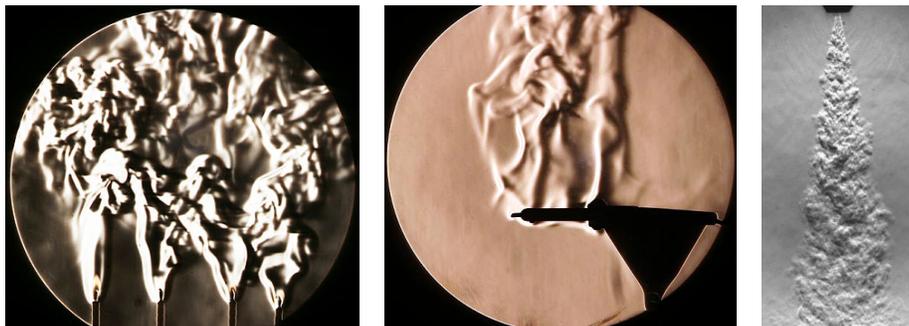


Figure 2.3: Example Schlieren photographs. Left: plumes above 4 candle flames. Center: plume from hot iron (photographs courtesy of Andrew Davidhazy). Right: turbulent helium jet (image reproduced from Volpe [46]).

pensive, so such systems are limited to studying small subjects. A laser and beam expander could also be used as a source of collimated light, but dangerously powerful lasers would be required to illuminate a reasonably large scan volume.

- *calibration difficulty:* although glossed-over here, the art of positioning all the optical components can be extremely delicate, and precise reproducibility is almost impossible after any reconfiguration. Due to diffraction effects, a razor blade is also not an ideal filter, and graduated-opacity filters produce better results [30]. Deciding on how sharp to make the transition, where exactly to position it, and the size of the source is the result of a trade-off between sensitivity and resolution. Obviously, once the beam is deflected so far that the entirety of it either passes over or is occluded by the filter, we can no longer see any further variation in the image intensity. Supporting a wider range of deviations requires mapping more slight intensity variations into the same image range, and we lose the ability to distinguish minor changes. Settles [42] discusses all of these issues and more in far more detail. In addition, image intensities are not easily mapped to absolute deflection angles. One way to accomplish this would be to place an object of known geometry (a high-quality lens, for example) into the scene, so that intensities in the image can be compared to a reference point.

- *anisotropy*: in the example in Figure 2.2, only vertical deviations produce any change in intensity. Any deflection of the beam parallel to the filter's edge does not alter the amount of light passing, and thus there is a strong dependence on the angle of deviation. Notice the similarity of the example photographs to images filtered with horizontal (vertical in the case of the helium jet) first-derivative edge detecting kernels. Of course, the filter can be rotated, but each image can only capture deflections in a single dimension. This highlights one of the main problems with classical Schlieren devices - by recording only intensity variations we lose directional information. More complex filter design alleviates this problem somewhat, at the expense of more complicated calibration.

2.2 Other Schlieren Configurations

Many ingenious variations of the original Schlieren method have been devised, only a scant few of which are mentioned in this thesis. This section begins with one of the most important small modifications to have been made, and then moves on to cover some radically different methods. Each of these solves one particular problem with the classic lens-based configuration, and BOS combines multiple advantages to provide a tool suitable to our purposes.

2.2.1 Rainbow Schlieren

Howes [21] replaced the knife-edge filter with a radial rainbow filter (Figure 2.4, centre). Colour provides an extra dimension to Schlieren images, allowing one to see not only the magnitude of a deflection, but its direction as well. Removing the sharp transition in the filter has the added benefit of reducing diffraction effects. Although it solves the anisotropy problem, one is still forced to deal with the same calibration and setup difficulties as with most other Schlieren methods.

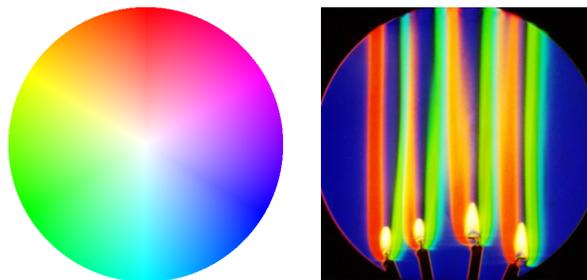


Figure 2.4: Left: the knife edge can be replaced by a colour wheel filter, having increasing opacity towards the edges. Both components of a deviation can then be identified, as hue maps the direction, and intensity maps the magnitude. Right: Schlieren photograph taken with a (different) colour filter, courtesy of Andrew Davidhazy.

2.2.2 Grid-Based Systems

Schardin [37] described a method applicable to very large field-of-view (even outdoor) settings. Forgoing collimating optics entirely, it replaces the point source with a background grid of alternating light and dark lines (for example, a metal plate with narrow horizontal stripes cut out of it). A piece of photographic film is placed in the filter plane, and an image of the background recorded. This is then developed, and a negative placed back *in exactly the same position as the original film*. An imaging lens and recording image plane are also included, as per Figure 2.5.

When rays can only travel in straight lines, any light from the bright portions of the background are blocked by the opaque parts of the filter, and so the image is dark. By placing a Schlieren object into the field of view, rays that pass through the transparent parts of the filter will be bent so as to strike the bright parts of the background.

This technique allows for very large subjects to be imaged. In fact, the Penn State University's full-scale Schlieren system [42] uses a 4×5 m background grid (retro-reflective material is used for the bright portions). However, the need to re-develop a filter each time an adjustment is made and the precise positioning it requires make using the device quite cumbersome. It is also limited to detecting only one-dimensional deflections.

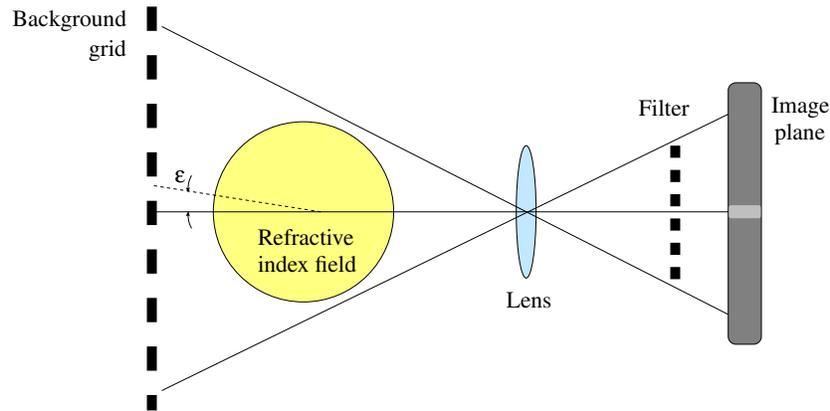


Figure 2.5: Grid-based Schlieren system. All rays are occluded by either the background grid or the complementary filter, unless refraction allows the ray to pass through.

2.2.3 Background Oriented Schlieren

Settles [41] described a number of Schlieren methods based on background distortion. By observing a distant light/dark background through an inhomogeneous field, density variations can be made quite visible. The hot air above a car's roof and the shock wave over a banked aircraft's wing both cause refractions visible to the naked eye when the background contains sufficient structure (fence or clouds, respectively). However, these distortions are only visible at the points where the regular structure in the background is disrupted, and a human viewer infers the existence of a continuous distortion field. Background oriented Schlieren (BOS), yet another distinct Schlieren technique, allows us to measure the distortion across the entire image.

Meier [29] patented BOS in 1999 and since then a handful of publications have used it. On the surface it appears quite different from the Schlieren techniques discussed above, but is nonetheless based on the same principle. The light source is replaced with a complex background pattern, and there is no longer a need for collimated light which makes the method far more attractive for large-scale, practical applications. Richard and Raffel [34] have demonstrated it in an outdoor setting, capturing the vortices formed by the rotor blades of a full-size helicopter. Schlieren lenses, mir-

rors and cutoff filters are removed entirely (although it can be argued that the iris of the observer's eye or camera plays the role of the filter) and all that is required is to capture an ordinary image of the background. Figure 2.6 illustrates the basic idea. Any deviation in ray paths due to refraction will manifest themselves as a corresponding distortion in the image. Assuming that we have a reference image of the undisturbed background, and are able to measure these image-space distortions, the deflection angle can be inferred. This image deformation is exactly the optical flow problem from computer vision, and there is potential to improve on previous BOS work by taking advantage of advances in optical flow research.

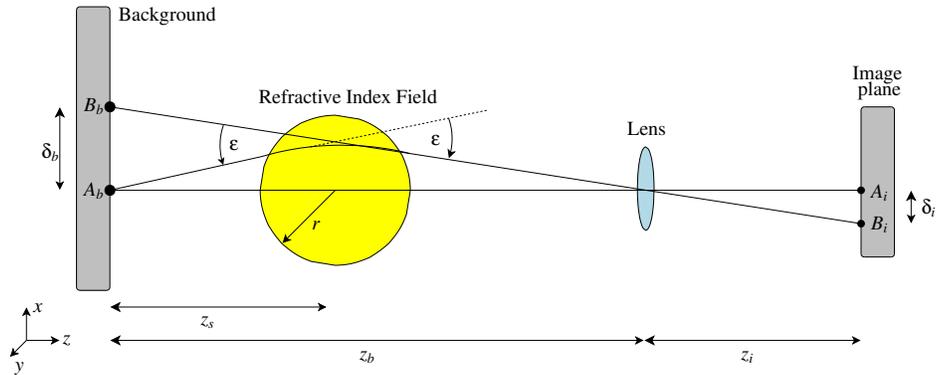


Figure 2.6: Virtual displacement caused by ray deflection. Under normal circumstances, points A_b and B_b on the background are imaged at A_i and B_i respectively. However, when refraction takes place, A_b appears at B_i and we see a virtual displacement of δ .

BOS combines the two-dimensional deflection sensing ability of rainbow Schlieren with the large scale of grid-based methods, and is the simplest technique in terms of equipment setup. It is not an entirely perfect solution, but its limitations are minor, and workarounds are discussed in Chapter 4. Other Schlieren methods always produce an image that is perfectly in focus, but with BOS we are back to a traditional projective imaging setup. Due to the finite depth of field, and the complicated and unknown lensing effects of the refractive media, it is impossible to keep both it and the background in sharp focus. The quality of the results is also constrained by the quality of the sensor (resolution and noise primarily) and of the optical flow algorithm used. However this

can also be viewed as an advantage. Both software and the sensor are easily replaceable in this setup, and with advances still being made in both of those areas, we can expect results to improve over time.

Elsinga et. al. [12] performed a thorough comparison of rainbow Schlieren and BOS using the airflow around a wedge of known geometry, and for which ground truth flow could be theoretically computed. They found both methods to be of comparable quality, giving results within 3% of the ground truth. BOS was found to have a superior dynamic range (it is able to capture both large and small deviations equally well) while rainbow Schlieren had to be tuned to work within a limited range. They did find the resolution of BOS to be lower, but this can be attributed to the optical flow algorithm used in the experiment. As the next chapter shows, higher resolution is achievable with better algorithms.

Determining Deflection Angles

The BOS displacement vectors can be used directly in visualisations, if projective scaling is not an issue. Rendering the magnitudes of these vectors according to a colour ramp is sufficient to qualitatively observe most flows. However, these 2D image-space vectors can also be mapped to deflection angles.

Ray deflections in Schlieren images are typically less than 200 arc seconds. As a consequence of this, the ray itself does not deviate much from the straight-line undeflected path, and making the *paraxial approximation* is standard practice in the Schlieren literature. Consider the ray deflection illustrated in Figure 2.6. Under the paraxial assumption, the refraction angle ϵ is so small that we can make the following approximations [45]:

$$\tan \epsilon \approx \epsilon \quad (2.1)$$

$$\tan \epsilon \approx \delta_b / z_s \quad (2.2)$$

From similar triangles we get that

$$\frac{\delta_b}{z_b} = \frac{\delta_i}{z_i} \quad (2.3)$$

and noting that z_i is the focal distance of the camera lens, we derive an expression for the deflection angle

$$\begin{aligned}\varepsilon &\approx \delta_b/z_s & (2.4) \\ &= \frac{\delta_i z_b}{z_s f} & (2.5)\end{aligned}$$

We do make two simplifying assumptions in the above, which allow us to treat the Schlieren field as a point instead of a volume:

- the surrounding medium is uniform, so no refraction occurs outside of the medium being imaged
- the overall size of the field is small relative to the distance between camera and background (*i.e.*, $r \ll z_b$)

Now the question arises as to what exactly this ε represents. After all, the ray does not make a single turn at one point and travel straight otherwise. Rather, it bends in some unknown fashion all throughout the field, which means that the ε we obtain is actually the integral of all the angular deflections the ray makes as it travels along its path. More so, it is related to the line integral of the gradient of the refractive index field. This fact can be derived from Fermat's principle [16] which states that a ray of light will follow a path corresponding to the shortest time of arrival (recall that light slows down when encountering a more optically dense material, and so will bend to find the quickest route through the media). The physical equation relating the time τ that it takes for a photon to traverse a point (x, y, z) with refractive index n is called the eikonal equation [6]

$$(\nabla\tau(x, y, z))^2 = n(x, y, z)^2 \quad (2.6)$$

Based on this, Gutierrez et. al. [18] derive the following equation governing the motion of a photon as it travels along an optical axis s

$$\frac{\partial}{\partial s} \left(n \frac{\partial x}{\partial s} \right) = \nabla n \quad (2.7)$$

We rearrange this into our 2D coordinate frame and then integrate once to get

$$\frac{\partial^2 x}{\partial z^2} = \frac{1}{n} \frac{\partial n}{\partial x} \quad (2.8)$$

$$\frac{\partial x}{\partial z} = \frac{1}{n} \int \frac{\partial n}{\partial x} dz \quad (2.9)$$

The quantity on the left corresponds to the deflection angle, and so we arrive at the final equation in the vertical dimension x

$$\varepsilon = \frac{1}{n_0} \int_{z_s-r}^{z_s+r} \frac{\partial n}{\partial x} dz \quad (2.10)$$

where n_0 is the refractive index of the surrounding medium, and a similar result holds for the horizontal dimension y .

2.3 BOS and Shadowgraphy

Schlieren imaging is very closely related to the principle of shadowgraphy. The latter is a far more simple technique, and requires nothing more than a light source (not necessarily a point, although it helps), a refractive field, and an image plane as illustrated in the figure below. Light is simply redistributed from one point to another, resulting in a pattern of light and dark. Shadowgraphs are easily cast outdoors by the sun.

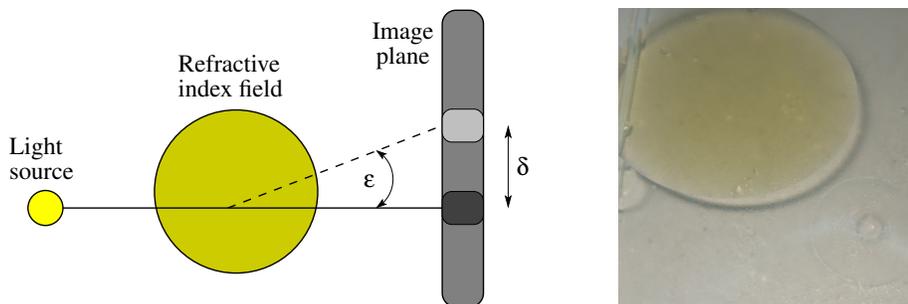


Figure 2.7: Left: a shadowgraph is quite literally the shadow of a refractive element, but instead of blocking rays entirely, it shifts them to a different part of the image causing that area to brighten. Right: example of a shadowgraph caused by a drop of oil floating on shallow water in a plastic bowl. Notice the bright area just inside of the dark shadow ring.

We know from Section 2.2.3 that Schlieren images respond to the gradient of the refractive index field. Shadowgraphs do not. If the gradient field were equal to a constant value everywhere, light rays would all shift together uniformly. In a Schlieren image this would produce a uniform change in the intensity at each point in the image, as the same fraction of each deflected ray is being blocked by the filter. However, in a shadowgraph, the entire (uniformly bright) image simply shifts, resulting in no overall change. As Settles [42] notes, Schlieren responds to the first derivative $\epsilon \propto \partial n / \partial x$ of the refractive index field, while shadowgraphy responds to the second derivative $\partial^2 n / \partial x^2$.

For applications such as ours, the temptation may exist to avoid the complexity of Schlieren photography entirely, to record shadowgraphs instead and then perform a double integration to find the deflection angles. Unfortunately, as Weinberg [47] notes, this is not possible in general due to superposition. With a nontrivial field in place, rays do not traverse unique parts of the scan volume, and so infinitely many solutions to the ray geometry could produce any given shadow.

A desirable property of most Schlieren method is that they produce images that are in focus. Shadowgraphy on the other hand, is subject to blur, as increasing the distance between the field and the image plane causes the shadow to spread out over a larger area. The most significant difference between the two techniques however, is in the complexity of the setup. Shadowgraphy requires almost nothing in the way of equipment and in fact bears striking resemblance to BOS when it comes to setup (and focus issues). One might then wonder whether or not BOS qualifies as a true Schlieren method. Settles [42] chooses the cutoff filter as the defining characteristic of a true Schlieren method. He does acknowledge that the iris of the aperture plays the role of the filter in certain Schlieren designs, although the lack of an obvious “knife edge” in BOS can be confusing (others implicitly class it as a Schlieren method). In the author’s opinion, it is more correct to classify the method based on the measured quantity. BOS responds to the first derivative, and so despite its massively simplified setup, it is indeed a true Schlieren method.

Chapter 3

Optical Flow

When observed from a fixed viewpoint, motion in a scene typically translates into corresponding motion of image intensities across the sensor. Therefore, analysing the motion in the image can lead to an understanding of how the scene is changing. Strictly speaking this is not always the case (for example, a rotating rotationally symmetric object will look the same in every frame as long as the lighting remains constant) but for our scenario such difficulties do not arise. Optical flow is a necessary component of many computer vision systems, and is still an active area of research. The fluid imaging community has thus far only begun to employ modern optical flow algorithms [11] and many particle image velocimetry systems (via which previous BOS studies are implemented) are still using modified versions of block matching [36].

This chapter begins with a brief review of the main principles in computing optical flow and then describes three main classes of algorithm. The focus here is on how they apply to our background distortion problem; it is by no means intended to provide an extensive coverage. Firstly, the prior BOS literature makes use of block matching algorithms, which tend to work reasonably well in practice, although with the necessary modifications they become highly complex and suffer from reliability and resolution problems. More sophisticated gradient-based approaches are described next, and we find that these provide a good balance between quality of results, ease of implementation and execution speed. Note that speed is an important issue here because we deal with video sequences, often from multiple cameras, and processing times of more than a few seconds per frame become intolerable. After a cursory discussion of a state-of-the-art variational approach, some more Schlieren-specific algorithms are considered. While at first it might appear that these specialized methods would be better suited

to our applications, experience shows that a solid general-purpose optical flow algorithm works well for BOS data. Finally, the carefully selected background pattern is described, showing how it can aid the optical flow by providing high frequency details at multiple scales.

3.1 Basic Concepts

Given two images, $I_0(x, y)$ and $I_1(x, y)$, the optical flow is defined as the 2D vector field $\vec{u}(x, y)$ describing the translation of each pixel from I_0 to I_1 . As per the well-known *brightness constancy* assumption, the intensity of each pixel is assumed to remain constant from frame to frame. This can be expressed as

$$I(x, y, t_0) = I(x + \delta x, y + \delta y, t + \delta t) \quad (3.1)$$

where $(\delta x, \delta y)$ represents the pixel's displacement over the time interval δt . In practice, uniform changes in intensity as well as minor local variations caused by gradual lighting changes can be tolerated reasonably well. Shadows and specular highlights (when the scene contains moving lights) are examples of common phenomena that violate brightness constancy. Fortunately the BOS setup satisfies it perfectly (aside from camera noise). In fact, it is almost an ideal case for optical flow, as we are able to guarantee that the background contains both high frequency detail and a high dynamic range. This ensures that each pixel is substantially different from those around it and we are thus better able to identify it in subsequent frames.

Optical flow is an ill-posed problem, and is usually regularised via a constraint on the smoothness of \vec{u} . An assumption of spatial (or spatio-temporal) coherence is made, based on the observation that it is highly unlikely for pixels belonging to real-world scene objects to move entirely independently of one another. BOS sequences do not contain typical foreground objects, but because we focus only on phenomena that are not highly turbulent, the tracked features (regions of the background) tend to behave like reasonably large moving objects. The refractive index fields we limit ourselves to are generally smooth, and produce spatially coherent distortions. To handle shock waves and sharp boundaries, an algorithm capable of finding piecewise smooth fields

should be used.

Note that we have made no allowances for attenuation or scattering. The scan volume and Schlieren medium are assumed to be perfectly transparent. If this were not the case, the brightness constancy assumption would not hold. Scanning the volume (perhaps by projecting multiple backgrounds) could provide a way to relax the constraint, but being restricted to a single image per time step precludes us from doing this.

3.2 Block-Matching Methods

The most naïve way to compute optical flow is to divide the image into a grid of small blocks (windows) and then to independently estimate the motion of each one. The assumption is that since each block represents such a small image patch, that it will exhibit primarily translational movement, and that everything else can be ignored. The validity of this assumption is highly dubious in practice, as typical window sizes of 16×16 pixels do in fact show (nonuniform) scaling and shearing for many of the media we are concerned with. Adding an explicit window deformation step significantly improves the results [13]. Deformation is applied in an iterative fashion, starting with square, regularly spaced windows. At each step, the current estimate of the vector field suggests a warping of the image, which is used to re-orient and reshape the windows. Figure 3.1 below illustrates the idea.

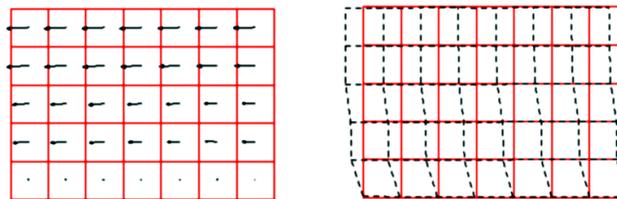


Figure 3.1: Deformation applied to block-matching windows. The flow suggested in the first (left) iteration suggest a shearing motion, and blocks are deformed accordingly in the subsequent (right) iteration (image reproduced from Scarano [36]).

Since we obtain only one vector per block, the window size should be minimised to give higher resolution in the output vector field. However the tradeoff is that each

vector is then less reliable, since fewer pixels are available in smaller windows to estimate the motion. Some users choose to overlap large windows to raise the resolution, but because each vector is an average of the motions of all pixels in the corresponding window, this results in neighbouring vectors that are highly correlated. A common solution is to use a hierarchical approach, beginning with large windows and successively refining them down to the smallest size.

Because blocks are processed completely independently of one another, there is no inherent enforcement of local or global smoothness. A rejection-filtering step is required to remove the spurious vectors *i.e.*, those that differ by more than a fixed threshold from their local neighbourhood's median (medians are more robust to outliers than means). It is also wise to perform a global filtering step, rejecting vectors that differ by too much from the global median. This helps to remove any large, but still spatially coherent, errors in the vector field. Removing these gross errors leaves holes which need to be filled in, typically via linear interpolation. Finally, a Gaussian or bilateral filter should be used to smooth the data.

3.2.1 Spatial Correlation

There is a limited range of motion that each block can undergo. If we know that no block can move more by more than M pixels over a time step, then a $(2M + 1) \times (2M + 1)$ matrix of *similarity scores* can be constructed, where the (i, j) th element of this matrix represents the likelihood of the block having translated by i pixels horizontally, and j pixels vertically (the indices ranging from $-M$ to M). This is an exhaustive evaluation of every possible translation that could have occurred, and once the matrix is constructed we need to search for the peak value which represents the overall translation of the block. In some cases there could be multiple peaks, or no one single element significantly higher than the surrounding noise, and selecting the wrong value in these cases is what leads to errors (spurious vectors) in the output.

Computing all the similarity scores is an expensive convolution process, which can be accelerated by working in the frequency domain. However the most common approach is spatial cross-correlation. Given an $N \times N$ image patch $f(x, y)$ (considered

to be periodic here for implementation purposes), and a shifted version $g(x, y) = f(x + u, y + v)$ of the same image, the goal is to evaluate how similar these two blocks are when overlaid. Finding a perfect match essentially means that we have recovered the offset (u, v) .

Correlation returns a measure of the similarity of two signals, and is defined (in the real, discrete case) as [27]:

$$(f * g)(x, y) = \sum_{i=0}^M \sum_{j=0}^N f(i, j)g(i - x, j - y) \quad (3.2)$$

This is sensitive to even uniform changes in intensity, and so it is commonplace to use instead the *normalised cross-correlation coefficient*

$$\gamma_{f,g}(x, y) = \frac{\sum_{i=0}^M \sum_{j=0}^N f(i, j)g(i - x, j - y)}{\sqrt{\sum_{i=0}^M \sum_{j=0}^N (f(i, j) - \bar{f})^2 \sum_{i=0}^M \sum_{j=0}^N (g(i - x, j - y) - \bar{g})^2}} \quad (3.3)$$

where \bar{f} and \bar{g} are the mean intensities over the windows.

Computing correlation scores for every possible (real-valued) translation is infeasible, so we sample at only the integer translations to produce a matrix like that plotted in Figure 3.2 below. Finding the location of the true peak, to subpixel precision is

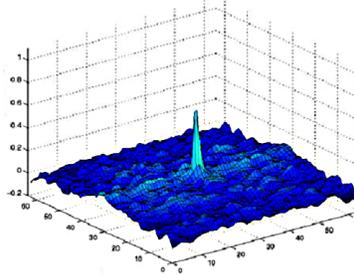


Figure 3.2: Matrix of cross-correlation coefficients, sampled at discrete integer locations (image reproduced from Scarano [36]).

done by fitting a curve through the maximum pixel and its neighbours. This is standard practice in PIV software, and three-point Gaussians are typically used [43]. Given consecutive data values p_{i-1} , p_i and p_{i+1} , with $p_i > p_{i-1}, p_{i+1}$, the location of the peak of an interpolating Gaussian, relative to p_i 's position is given by

$$x = \frac{\ln p_{i-1} - \ln p_{i+1}}{2(\ln p_{i-1} - 2 \ln p_i + \ln p_{i+1})} \quad (3.4)$$

Evaluation

One side-effect of the subpixel fitting procedure is that it produces a fixed-pattern background noise (see Figure 3.3). In regions of no Schlieren disturbance at all we would expect zero displacement. Slight camera noise means that this is not quite the case in practice, however, independent of noise, we still find some non-zero values in the correlation matrix next to the peak. And because a shift by one pixel to the left is unlikely to produce the exact same correlation coefficient as a shift one pixel to the right, an interpolating curve is going to have a peak that lies to one side of the central pixel. As a contrived example, when a window shifts by exactly 2 pixels, the correlation scores for 1 and 3 pixel shifts would most likely be distinct, non-zero values, resulting in an interpolated peak that does not correspond to an exact two pixel shift. Because these values don't change from frame to frame, we see an essentially random, but still temporally consistent fixed noise pattern. An effective way to remove this is to first compute the autocorrelation function of the reference windows with themselves and then to subtract it from the output vector field.

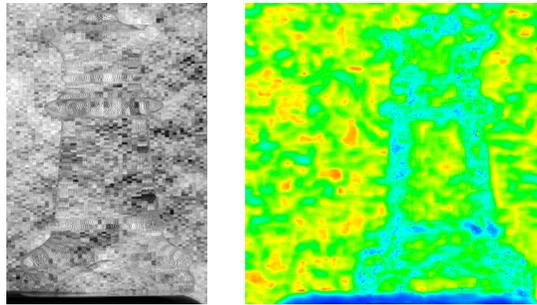


Figure 3.3: Left: glass chess piece. Right: colour-coded magnitude of optical flow vectors. The noise pattern outside of the chess piece is not caused by camera noise - it remains constant across frames.

Although extremely expensive, correlation is well-suited to parallel hardware implementation. We used Cg shaders on an Nvidia GeForce 7600 GPU to compute normalised cross-correlation coefficients. Despite being so hardware-friendly, implementation in real-time was still not possible when dealing with large images (256×256 and more). This is due to the fact that CPU \leftrightarrow GPU transfers are still relatively expensive,

and some operations (curve fitting) had to be performed on the CPU.

As previously mentioned, block-matching methods inherently assume constant linear translation of each window. Experiments with BOS images show that this is clearly not the case, and the warping that occurs under nonlinear motion can be quite significant. Window-deformation is used in PIV software to handle this, although being forced to use block-matching in that case makes it somewhat of a necessity. The input images are generally black-and-white, and contain insufficient detail to work well with the gradient-based optical flow techniques described in the next section. However, in the author's opinion, applying window-deformation to BOS data is overly complicated, inelegant, and ultimately inappropriate. Rather than tweaking a weaker algorithm, better results can be obtained more easily through more sophisticated optical flow techniques.

3.3 Gradient-Based Methods

Gradient-based methods have long been the stalwart of optical flow. They are based on the Taylor series expansion of the change at each pixel

$$I(x + \delta x, y + \delta y, t + \delta t) = I(x, y, t) + \frac{\partial I}{\partial x} \delta x + \frac{\partial I}{\partial y} \delta y + \frac{\partial I}{\partial t} \delta t + \dots \quad (3.5)$$

where the higher order terms are insignificant enough to be ignored. Taking the brightness constancy Equation 3.1 into account and dividing throughout by δt we get that

$$\frac{\partial I}{\partial x} \frac{\delta x}{\delta t} + \frac{\partial I}{\partial y} \frac{\delta y}{\delta t} + \frac{\partial I}{\partial t} = 0 \quad (3.6)$$

The $\partial I/\partial x$, $\partial I/\partial y$ and $\partial I/\partial t$ terms are just image derivatives. For each pixel then, we have one equation in two unknowns. Additional assumptions are thus required to get a unique solution.

3.3.1 Lucas-Kanade

To solve the under-determined system, the Lucas-Kanade algorithm makes an assumption similar to that of block-matching *i.e.*, that motion is constant over a small window

of neighbouring pixels [28]. Each pixel (i, j) within the window centered at (x, y) therefore contributes an equation of the form

$$\frac{\partial I(i, j)}{\partial x} \frac{\delta x}{\delta t} + \frac{\partial I(i, j)}{\partial y} \frac{\delta y}{\delta t} + \frac{\partial I(i, j)}{\partial t} = 0 \quad (3.7)$$

Solving this now overdetermined system for $\delta x, \delta y$ at (x, y) can then be done in a least-squares sense.

As with any gradient-based method, this only works for small motions of less than one pixel (image gradients across multiple pixels for random dot backgrounds are not useful). To support larger motions, the algorithm is implemented in a hierarchical fashion by downsampling the image a number of times, computing the flow at a coarse level and prolongating this up to finer levels. This acts as an initialisation, from which only a small differential motion takes place at the next level.

Evaluation

Due to its popularity, many very efficient implementations of the Lucas-Kanade algorithm exist. We tested with one from the Intel OpenCV library.

It is essential to have high detail in all parts of the image. Homogenous regions contain zero gradients which make it impossible to solve Equations 3.7 for nontrivial values of δx and δy . Fortunately BOS images contain detail everywhere and so are ideally suited to gradient-based methods.

The algorithm contains a slight nod to spatial coherence, in terms of the constant window motion assumption, but in practice it did not produce very smooth results. Camera noise makes it very difficult to get accurate local flow estimates and so heavy filtering was required to smooth the data.

3.3.2 Horn-Schunck

The Horn-Schunck algorithm takes an alternative approach to solving the under-determined System 3.6. A regularising term is instead added to enforce global smoothness [20].

The idea is to minimise the function

$$f = \sum_{i=0}^{N-1} \left(\frac{\partial I}{\partial x} \delta x + \frac{\partial I}{\partial y} \delta y + \alpha (|\nabla \frac{\partial I}{\partial x}|^2 + |\nabla \frac{\partial I}{\partial y}|^2) \right) \quad (3.8)$$

over all N pixels in the image. For simplicity, we have assumed that $\delta t = 1$. Smoothness is enforced by ensuring that the gradients of the vector flowfield are kept small, and the parameter α controls the weight of this penalty term (*i.e.*, how smooth we would like the output to be).

Evaluation

The Horn-Schunck algorithm performed very well in practice. Flowfields for BOS data are predominantly smooth and this was well reflected in our tests. Unlike with Lucas-Kanade, there is no sharing of information across a local neighbourhood and so camera noise did produce some high frequency noise even in the right-hand flow field of Figure 3.4. A post-smoothing step was able to remove this. Runtimes were

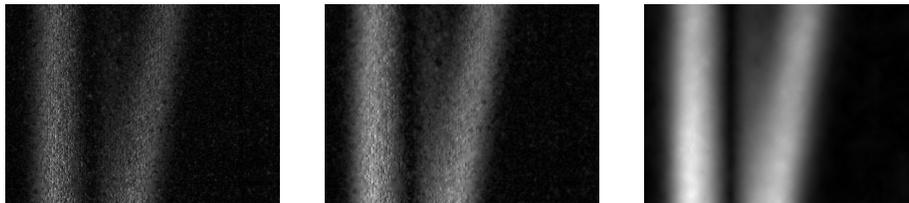


Figure 3.4: Magnitude of flow vectors for a candle plume. α increases from left to right, affecting the smoothness of the field.

acceptable, on the order of tens of seconds per frame. The biggest disadvantage is that it tends to over-smooth across edges. Since our experiments lacked any shock wave examples this was not a serious issue and Horn-Schunck worked well. In cases where sharp edges are likely to occur, it would be advisable to use instead the Black-Anandan algorithm, which uses robust statistics to fit piecewise smooth flowfields that do not smooth over edges [5].

3.4 Variational Methods

In order to deal with more difficult real-world scenes, with occluding objects and changing lighting, other approaches to optical flow have been developed. A current state of the art, variational algorithm proposed by Papenberg et. al. [31] produces some

of the best known results for hard scenes at the time of writing. Although it does make use of the same brightness constancy assumption as other methods, it does not immediately linearise this constraint via a Taylor series as the pure gradient-based approaches do. By avoiding this it is far better able to handle complex motions. Additional constraints on smoothness are also imposed and a nonlinear energy functional derived (although the numerical solution procedure itself requires linearisation). The results are impressive, and recent work on a multigrid implementation [7] suggests that in the future this will be the ideal algorithm for, amongst other things, BOS data. However, at present the algorithm is quite complex and no public implementation is available. The (non-multigrid) method is also known to be at least an order of magnitude slower than Horn-Schunck, making it unsuitable for our needs.

3.5 Schlieren-Specific Optical Flow

Some authors have proposed optical flow algorithms designed specifically for dealing with (traditional, lens-type) Schlieren images. The problem in this case is that these photographs typically have a very low dynamic range, and are quite smooth (recall that they are sensitive to deflections in only one dimension). Arnaud et. al. [3] use Horn-Schunck to compute optical flow for these images, but suggest a modified energy function to optimise, which includes a term specifically related to the divergence and curl of the underlying fluid flow. They reason that given the lack of contrast in the images, additional constraints guided by our knowledge of fluid flow can be used to obtain a better solution. However, it is important to distinguish between the 2D *optical* flow field, and the underlying 3D *physical* flow field. The former is a projection of the latter, and fluid motion concepts do not necessarily map directly onto optical flow fields.

Jonassen et. al. [24] suggested a hybrid Schlieren-PIV system, intended only for highly turbulent flows (see for example, Figure 2.3 right). Capturing the fine micro-turbulence structure is too difficult, but by treating the tiny turbulent eddies as virtual particles that remain unchanged from frame to frame (a reasonable assumption over *very* short time steps), they track these eddies using optical flow (block matching) to

produce a coarse flowfield.

To find the background distortion, Agarwal et. al. [1] use a method specifically designed for refractive media. They attempt to solve for the flow by projecting a sequence of structured backgrounds. This idea is quite common in scanning and environment matting techniques, but it does not support real-time acquisition and is thus not suitable for our needs.

Chuang et. al. [9] mention a real-time variant of their environment matte acquisition technique. For each input frame they obtain a per-pixel warping function by searching for a pixel with matching colour in the reference background frame. A colour ramp (slice through the RGB colour cube) is used for the static background. This process is extremely sensitive to camera noise, and coherence is not directly enforced, so spatio-temporal smoothing has to be applied as a post-process. They use Perona and Malik's edge-preserving smoothing operator [33] and report that smoothing just the output vector field produces better results than smoothing the input images instead. Since most colour cameras use Bayer filters, the sampling density of the chroma channel is lowered. Because no neighbour information is used to compute the flow vector at each pixel, and local smoothing is not an integral part of the process, this method is best suited to data with sharp boundaries (e.g. glass objects). For smooth gas flows, it makes sense to take advantage of values at neighbouring pixels.

3.6 Wavelet Noise Background

Optical flow requires images with high detail content in order to work effectively. The flow cannot be found in untextured regions, and must instead be smoothly interpolated there from the boundaries. Due to the physical constraints of seeding a flow with particles, PIV software is forced to work with images that are predominantly uniform black but covered in white speckles. Frequency content is therefore mostly zero, apart from the very high frequency particle edges. Previous work on BOS has tried to simulate the PIV input by splattering white paint on a grey background [34]. However this ignores the fact that we can design the background to aid the optical flow algorithm. While not strictly necessary, it does improve results in most cases.

Colour backgrounds are avoided because our optical flow algorithms typically work on single-channel images, and because the chromatic sampling rate of most cameras is lower than that of the luminosity sampling rate due to the Bayer filters. The most obvious high frequency background to use would be a random dot greyscale pattern. However this can give rise to aliasing problems. Whenever we try to represent the image using too few samples, spurious low frequency signals can appear [16]. There are two instances in which this downsampling occurs in our setup:

- *image pyramids*: the optical flow algorithm operates on a Gaussian pyramid of the image in order to deal with displacements of more than one pixel
- *background placement*: the background and camera can be moved according to the desired scan volume size, and to optimise sensitivity. We would prefer not to have to print multiple backgrounds, one for each possible separation distance.

Assuming that correct lowpass filters are used to prevent aliasing, we still face the issue of loss of detail at lower resolutions. Figure 3.5 shows the effect of downsampling a purely random pattern. The dynamic range is reduced until eventually it becomes a uniform 50% grey, against which no distortions can be measured.

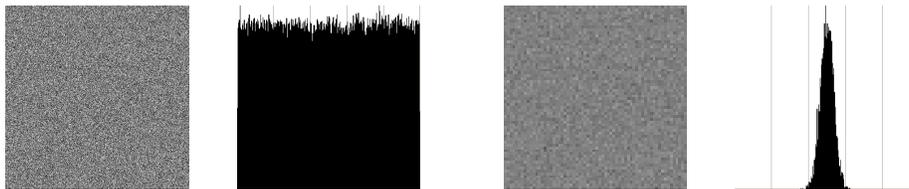


Figure 3.5: Left and center left: 512×512 uniformly distributed random noise pattern and its flat histogram. Center right and right: the same image downsampled to 64×64 has a greatly compressed dynamic range.

This problem is solved by using wavelet noise [10] for the background, shown in Figure 3.6. It contains high frequency detail at all scales. To generate it, an image pyramid of uniformly distributed random noise patterns is produced, with each level double the resolution of the preceding level. Then a separate bandpass filter is applied to each, ensuring that they contain non-overlapping frequency ranges (as a side-effect,

the histogram takes on a normal distribution). When the images are added together, the result can be used for BOS at any scale since it will always show sufficient detail.

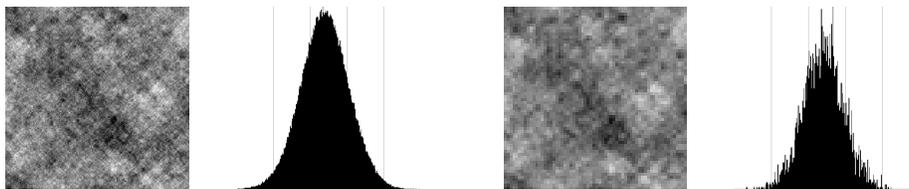


Figure 3.6: Left and center left: 512×512 wavelet noise pattern and its normal histogram. Center right and right: the same image downsampled to 64×64 maintains a wide dynamic range.

3.7 Conclusion

At some point a trade-off has to be made. There are a wide range of optical flow algorithms available. Some are tailored to specific input cases, and the quality of results and time required for computation differs greatly. It is important to bear in mind where the limiting factor of the overall system lies. In our BOS for tomography (Section 5.5) application for example, since we have so few views available, and so much 3D smoothing has to be applied to the data, accurate high resolution optical flow fields are simply not necessary. In other areas, a desire for high-resolution, high-framerate acquisition places constraints on the types of cameras we are able to use, and these tend to suffer from more noise than what we would hope for. Ultimately, while it may not always produce perfect results, we found the Horn-Schunck algorithm to be the best compromise, and to work well with the wavelet noise pattern.

Chapter 4

Physical Setup

This chapter describes our physical measurement setups for single-view and multi-view acquisition. Multiple problems encountered with the lens-array, as well as various practical implementation issues are discussed, along with our solutions.

4.1 Acquisition Setups

4.1.1 Standard Setup

Our default BOS acquisition setup consists of a backlit background and a stationary camera. For high resolution, high framerate, low noise images we use a monochrome Prosilica 1.5 megapixel C-mount camera, shown in Figure 4.1. The overall sensitivity is easily adjustable by moving the background plane. Smaller distortions produce larger image-space translations as the background is moved further away. In practice, the setup works well under a wide range of configurations. We obtained good results with candle flames positioned from 5cm to 1m away from the lens.

In addition to the scientific CCD camera, we tested with a recent-model consumer video camera and found it to produce good results. The Sony HDR-SR1 records in high definition 1440x1080 and uses MPEG4 compression. Despite noise levels higher than those of the Prosilica camera and the lossy compression, it allowed the distortion in a candle plume to be seen without difficulty.

Backlighting is not strictly necessary for the noise pattern. Unlike other Schlieren setups, BOS does not require collimated or narrow-band illumination. It is important to prevent any moving shadows or specular highlights from falling on the background. If a reference background cannot be captured, it can be approximated as an average

over all the frames. With a long enough sequence, this technique is very effective, and suggests that real-world (static) backgrounds can also be used.



Figure 4.1: Standard single camera setup and small candle.

4.1.2 Camera Array

For the tomographic application we require multiple simultaneous views of the scan volume. An array of 8 Imperx MDC-1004 video cameras with 1 megapixel resolution was configured in a half-ring for this. The automatic white balance on these cameras proved to be inaccurate, showing significant Bayer pattern artefacts. To correct for this we extracted just the green channel and ran the optical flow at half-resolution (distortions from the burner were still clearly visible). High BOS resolution was not critical in this application since, due to the low number of viewpoints, heavy 3D smoothing had to be applied to the volume data as part of the regularisation.

4.1.3 Lens Array

Since a camera array is very expensive, as well as being tedious to configure and operate, we desired another solution that would require only one camera. Georgiev et. al. [15] used a large external lens array (as opposed to placing microlenses over the sensor itself) to capture from multiple closely-spaced viewpoints. Since this inherently involves trading off resolution against the number of views, we would need a higher resolution sensor than that on the Prosilica camera and therefore used a 6 megapixel

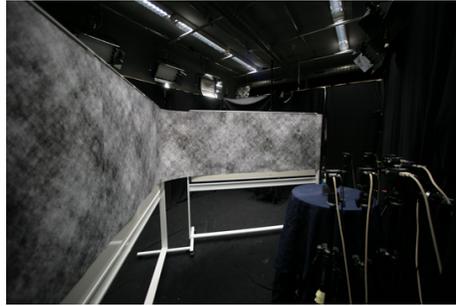


Figure 4.2: Camera array setup and gas burner.

Canon EOS D60 camera. Our lens array consists of a 6×5 grid of 15mm lenses in a close hexagonal packing, which fit just within the field of view of the camera and which have overlapping fields of view. There was some wastage of sensor area due to the plastic mounting, and because the edges of the lenslets were not in focus. Each lense therefore gave us an approximately 250×250 pixel area to analyze. Theoretically this would be sufficient, but in practice no useful results could be obtained for the following reasons:

- *focus*: even with the SLR set to the widest possible depth of field, we could not simultaneously get the centers and the edges of the lenslets into sharp focus.
- *chromatic aberration*: the lenslets suffered from significant dispersion, which introduced a radial blur after conversion to greyscale. To accommodate for this, we placed a green filter over the camera and ignored the red and blue channels. This effectively halved our resolution though, to the point that optical flow could not be accurately measured.
- *camera shake*: lack of a remote trigger forced us to touch the camera itself, introducing a constant offset to the flow fields (although this could be corrected for by subtracting that amount afterwards). In addition, the SLR mirror flip-up causes some shake.



Figure 4.3: Lens array setup and water tank.

4.2 Angle Measurement

Figure 2.6 shows how we approximate the deflection angle as δ_b/z_s (assuming ϵ is small, A_b, B_b and the centre of the refractive index field form a right-angled triangle). Converting optical flow vectors to angles involves an assumption that a single refraction occurs at one point. Since we do not yet know enough about the 3D structure of the index field, this point is assumed to lie on a plane passing through the centre of the field, parallel to the background). Since the field is often rather wide, we strive to move the background as far back as possible to make the assumption valid.

4.3 Implementation Notes

In practice the BOS setup was very easy to use and robust to minor adjustments. Experience shows that these are the most important implementation issues to overcome:

Noise from the sensor degrades the quality of the optical flow results. The cameras we used were all of high quality and so did not suffer from major noise problems, although a slight Gaussian blur did help to remove what there was. Hoping to construct a larger camera array, we tested with cheap and low-quality consumer cameras (webcams and compact digital cameras, at 640×480 resolution) to no avail. High noise levels and MPEG compression combined to mask the distortion fields.

Focus is not an issue in traditional lens-and-mirror Schlieren setups due to the orthographic projection. With BOS however, we need to focus both on the background plane and on the refractive index field. Due to the arbitrary lensing caused by the field, it can be difficult to bring it into focus (this is only a real problem when dealing with strong refractions from solids, discussed below). Our solution is to maximise the depth of field by using as small a camera aperture as possible, and using bright studio lights to allow images to be captured without high gain.

Solids are more difficult to image using the BOS method than fluids, because they tend to contain both sharp edges and significant warping due to very high refractive index gradients. Schlieren imaging is far better suited to cases where the field's refractive index differs from the surrounding medium in only the 3rd or 4th decimal place. Glass has an index of approximately 1.5 (for reference, air has an index of approximately 1.0, depending upon the temperature and pressure) and the distortions for complex objects like figurines and thin-stemmed wine glasses are simply too severe to measure using optical flow. Some environment matting techniques could be used [9] to acquire the distortion field, but it is also possible to adjust the index of the surrounding medium by immersing the glass object in another fluid [8]. Trifonov et. al. [44] required an exact index match to eliminate refractions. A potassium thiocyanate solution has an index high enough to match glass, but is toxic and requires careful mixing to reach the correct concentration. We were able to use harmless vegetable oil (although non-water soluble solutions are more difficult to clean up afterwards) with an index of approximately 1.47 to reduce the refractions in simple glass objects.

Chapter 5

Results and Applications

Once the displacement vector field has been acquired, a number of applications become possible. After showing that the optical flow results are of reasonable accuracy, this chapter describes the applications with which we have experimented, as well as speculating on the possibility for future improvement in the area of tomographic reconstruction.

5.1 Optical Flow Performance

To evaluate the performance of the optical flow algorithms we used a synthetically generated 2D random distortion field. It was used to warp two background images, and we then attempted to recover the original field, measuring the mean squared error of the differences. In addition to the wavelet noise image (Figure 3.6 left), we used the synthetic “PIV-like” image (Figure 5.1 top right) with energy across all frequencies as background images. As mentioned in Chapter 1, PIV is another method for imaging fluid flows via optical flow, and is limited by the number of physical particles that can be inserted into the flow. Our test image represented a best-case scenario, with an extremely high particle density. Nonetheless, better reconstructions were produced with the wavelet noise background. The bottom row of Figure 5.1 shows the reconstructed fields for both backgrounds when using the Horn-Schunck algorithm (with identical parameters). Table 5.1 lists the mean squared error values for these tests, demonstrating that we can improve results by up to 33% through control of the background.

We also found that in general, Horn-Schunck produced better results than Lucas-Kande, which in turn produced better results than block-matching. The only anomaly

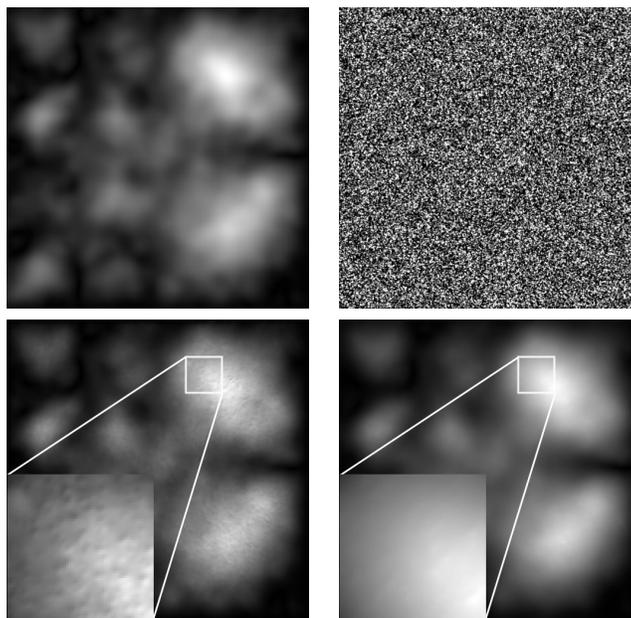


Figure 5.1: Synthetic optical flow test images. Top left: synthetic 2D field magnitudes. Top right: simulated PIV background, with as many particles as could practically be seeded into an airflow. Bottom left: reconstructed flow field from PIV background (Horn-Schunck). Bottom right: reconstructed flow field from wavelet noise background (Horn-Schunck).

was the extremely high error of 1.05 pixels for block matching with the wavelet noise background. This is as a result of the subpixel fitting problem mentioned in Section 3.2.1. The PIV background contains relatively large areas of black or white, which result in low correlation scores unless blocks are perfectly aligned. The wavelet noise image on the other hand, has less inter-pixel variation so correlation peaks are less distinct. Subtracting the autocorrelation field (the vector field we obtain by matching the background against itself) reduces this error dramatically, as shown in Table 5.2

MSE	Block-matching	Lucas-Kanade	Horn-Schunck
PIV	0.0554	0.0468	0.0353
Wavelet	1.0514	0.0456	0.0238

Table 5.1: Mean squared error (in pixels) of reconstructed fields for each algorithm with both PIV and wavelet noise background images.

MSE	Direct	Subtract autocorrelation
PIV	0.0554	0.0504
Wavelet	1.0514	0.1133

Table 5.2: Mean squared error after autocorrelation subtraction.

(note that the gradient-based algorithms returned identically zero autocorrelation fields for both background images). However, the average error of approximately 0.1 pixels is still relatively high, and in this case the wavelet noise proves to be a hindrance. The optimal procedure for recovering smooth fields with BOS is therefore to use the Horn-Schunck algorithm with wavelet noise as a background.

Lens effects in refractive index fields introduce minor local changes in scale, so we tested with another synthetically generated flowfield containing minification and magnification factors of up to 2X. The smooth synthetic field consists of vectors pointing away from the center, increasing in magnitude towards the bottom right. In the middle there is effectively no motion, and pure anisotropic magnification increases along the main horizontal and vertical axes, while isotropic scaling occurs along the diagonals. The largest displacement is by 80 pixels. As shown in Figure 5.2, the fields were recovered right up to the edges where scaling was most pronounced. Processing time is acceptable at under 30 seconds per frame (512×512) and average angular error is less than 7%. For comparison, Barron et. al. [4] report an average angular error of 9.78% for Horn-Schunck on the standard Yosemite sequence. Our case is made easier by having a smoother flow field, no discontinuities, and higher frequency detail in the image.

5.2 Flow Visualisation

The most obvious application for BOS is to simply view the flowfield directly. Figure 5.3 is a magnitude plot of two frames from a candle plume sequence (compare this to the candles in Figure 2.3). The lower sections of the plumes have the same shape - refractions in towards the centre, with no distortion along the central vertical axis. This is because a candle plume is cylindrical, and hotter than the air around it. The

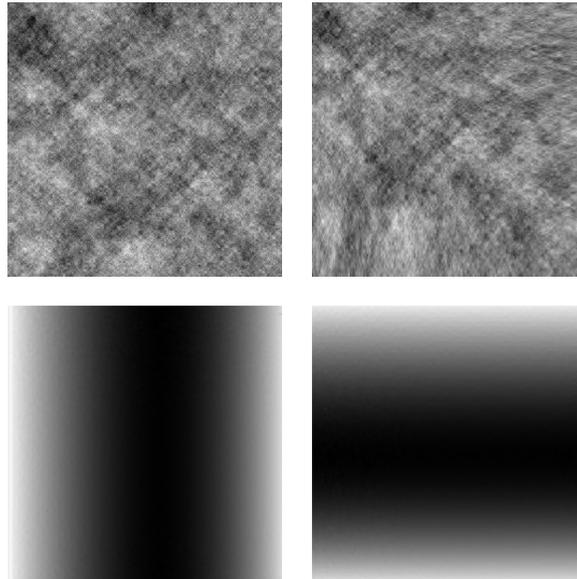


Figure 5.2: Optical flow scaling test. Top row: reference and distorted noise patterns. Minification occurs in the top left, while the bottom right is magnified. Bottom row: absolute values of horizontal and vertical vector components.

well-known Gladstone-Dale law

$$n - 1 = k\rho \quad (5.1)$$

relates refractive index n and gas density ρ via a constant ($k = 0.23 \text{ cm}^3/\text{g}$ for normal atmospheric conditions). As the gas heats up it becomes less dense. The right hand diagram in Figure 5.3 illustrates how a ray striking the plume will bend away from the boundary's normal on entering the plume, and towards it on exit. The background will therefore appear to move in towards the axis when the plume is present.

The upper regions of these images show non-laminar flow, where the gas has cooled enough to be driven more by random motion in the surrounding air currents than by its own upward force. An individual particle inside the flow would move upwards through the plume, which is certainly not in the same direction as the optical flow vectors are pointing. Schlieren photographs do not show flow directly like PIV tests do, and we need to interpret them correctly. Fu and Wu [14] ran optical flow directly on Schlieren photographs to get flow velocity, and we could do the same by applying another stage of



Figure 5.3: Candle plume results. Left and middle: magnitude plots of optical flow vectors from two consecutive images in the sequence. Right: refraction in a cylindrical lens makes the object at *A* appear where *B* did previously.

the optical flow algorithm to our magnitude plots. This only works if there is sufficient structure in the images to clearly show the motion. Still images make interpretation more difficult, but under animation the motion is obvious.

Figures 5.4 and 5.5 show the results of another test, with a candle plume being blown from the side by a can of compressed air. The jet of air is both colder and of higher density than the surrounding air, and so has a higher refractive index, causing the optical flow vectors to point outwards from the central axis. The distorted image

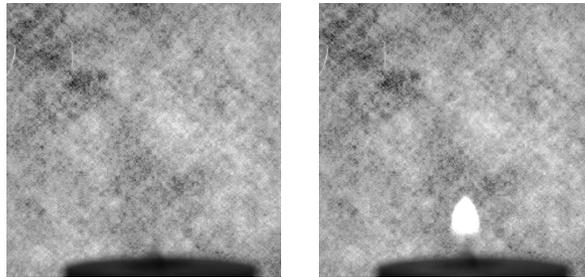


Figure 5.4: Captured images from plume interaction test.

contains a significant degree of blur in the compressed air jet. Because of its high speed and turbulence, we necessarily capture a time-averaged image given that the Prosilica camera was recording at 15fps. Despite this difficulty, the overall shape is still well captured, although we only see a smooth field and lose the high frequency turbulence.

Liquids can produce much the same degree of refraction as gasses. We injected a dilute solution of corn syrup into distilled water to produce the sequence in Figure 5.6. The results are over-smoothed and less clear when compared to the gas cases, since we

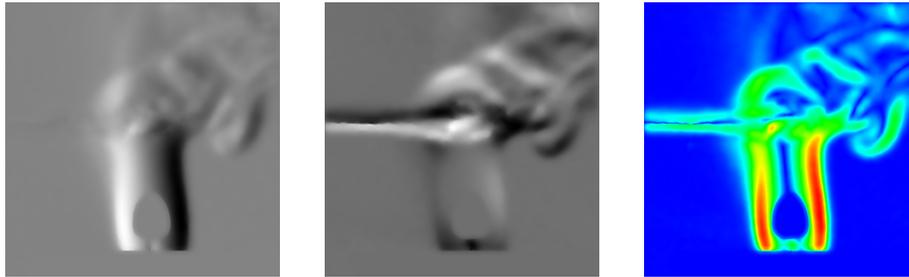


Figure 5.5: Optical flow vectors from plume interaction test. Left: horizontal components. Middle: vertical components. Right: magnitudes.

have interference introduced by the plastic container walls and dirt particles suspended in the water.

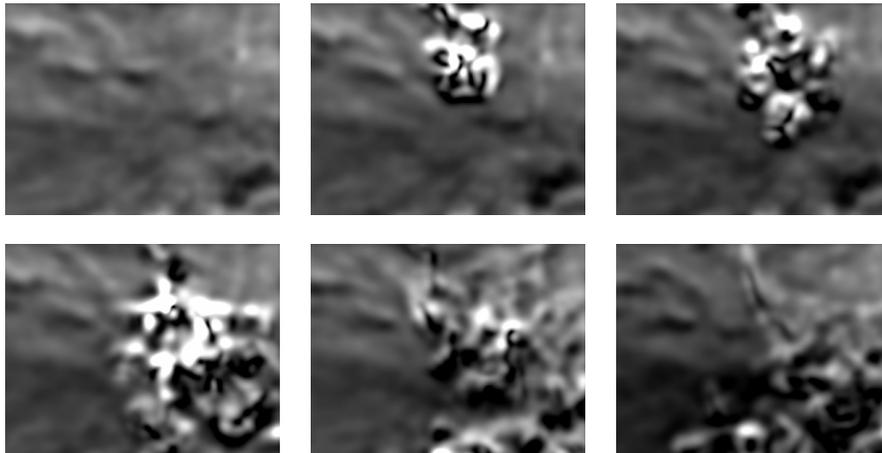


Figure 5.6: Corn syrup solution being injected into water.

Applied in an outdoor setting, we can also imagine this method being used for gas leak detection. Settles [40] proposes using the background grid method (Section 2.2.2) to see gas leaks from outdoor pipelines. BOS allows for natural (static) backgrounds to be used instead. Even if a suitable background is not present in the environment, printing a large wavelet noise pattern would be easier than producing the exact complementary grids as required by that method.

5.3 Environment Matting

Environment matting was first described by Zongker et. al. [49] and later extended by Chuang et. al. [9] as a method for capturing the complex lighting reflections and refractions of a translucent object with a simple 2D data structure and allowing it to be re-rendered in a new environment, maintaining the correct interaction with its new surroundings. Although a complete environment matte describes multiple properties (opacity, reflectance, refraction) we are concerned here with just the refraction, since typical Schlieren subjects are perfectly transparent.

An environment matte is usually captured by projecting a sequence of structured backgrounds behind the object and filming it. Various background patterns, from sweeping lines [9, 49] to natural images [48] to wavelets [32] have been used to obtain the mapping between sensor and background pixels. Subjects of interest for rendering are typically glass with many sharp edges and rough surfaces causing beams of light to spread over large areas. Environment mattes therefore consist of mappings to *regions* of the background (typically Gaussian blobs), whereas with Schlieren subjects we obtain a 1-1 mapping mapping to individual background pixels. In this respect BOS operates in a manner similar to the real-time acquisition method of Chuang et. al. [9].

Figure 5.7 shows the results of capturing a distortion field in one setting and applying it to another scene, making it look as if the optically active element were actually there.

5.4 Shadowgraphs

As described in Section 2.3, a shadowgraph is cast when light rays are deflected away from one part of a surface onto another by an inhomogeneous refractive index field. The resulting patterns of light and dark are the “shadow” of the field. We can cast synthetic 2D shadowgraphs via photon mapping [23]. Virtual particles (photons) are fired into the scene, with their paths being determined by the refractive index field. A simple raytracing simulation is performed by shifting each particle according to the distortion field. We splat a small Gaussian blob at the point where each photon

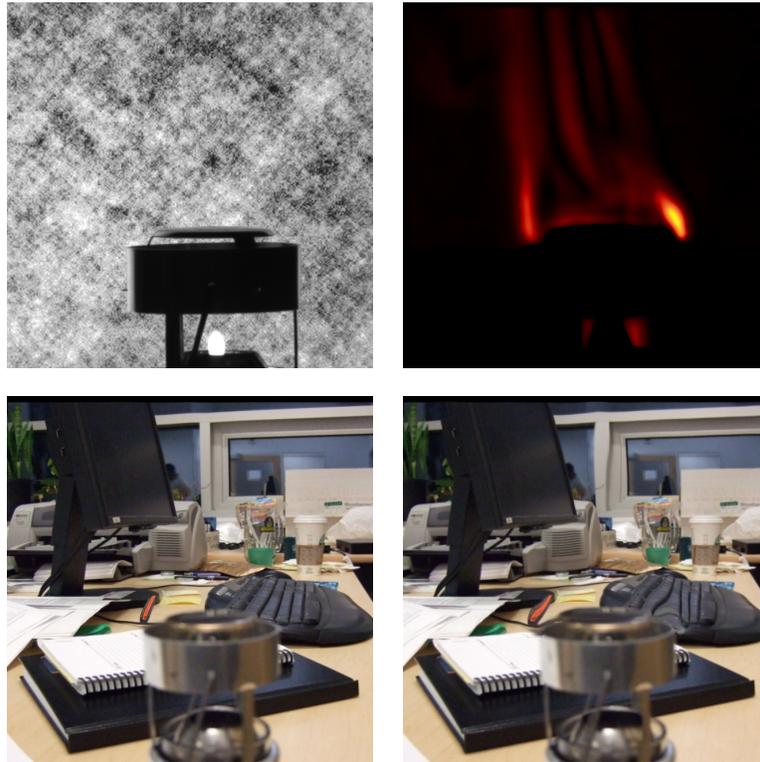


Figure 5.7: Candle lantern environment matting example. Top left: captured image of lantern in acquisition lab. Top right: magnitudes of optical flow vectors. Bottom left: novel background image (the lantern is physically in the scene and is not burning). Bottom right: environment matte composited onto novel background. Notice the distortion (exaggerated for clarity) in the keyboard and windowsill.

strikes the surface. Once a sufficient number of these have been cast, the final density estimation step is performed. Summing the normalised contributions from all nearby photon splats gives us an estimate of how much light is arriving at each scene point. The results, in Figure 5.8 look convincingly realistic under animation.

5.5 Tomographic Reconstruction

While this thesis has concentrated on 2D acquisition, we note that it is possible to use this information for 3D recovery of refractive index distribution fields. As part of a

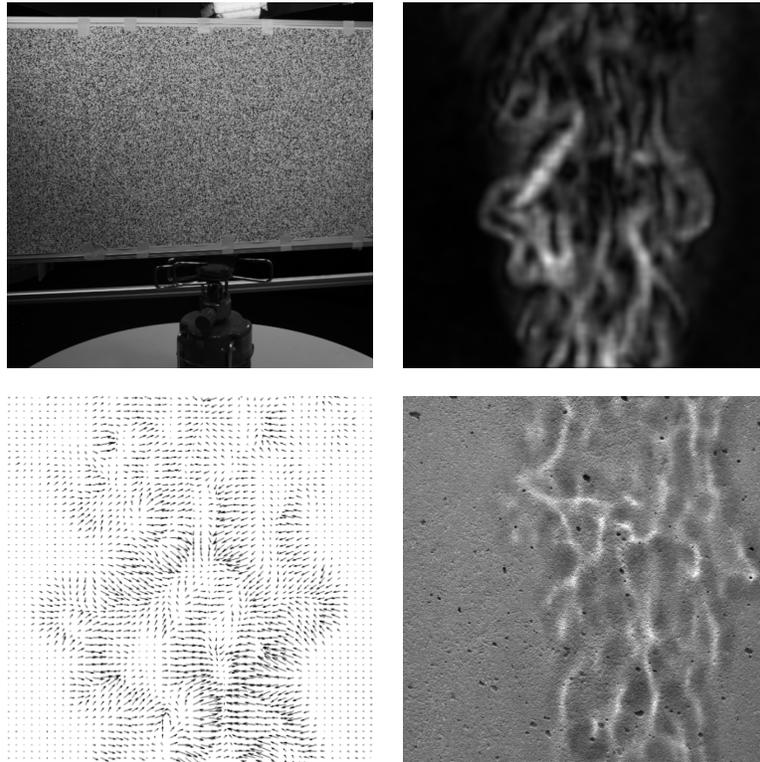


Figure 5.8: Synthetic shadowgraph example. Top left: captured image of gas burner in acquisition lab. Top right: magnitudes of optical flow vectors. Bottom left: vector field. Bottom right: Result of casting shadow onto novel background.

separate project, Ihrke [22] devised the reconstruction algorithm outlined in this section. The results shown here were computed using the flowfields we acquired through our BOS setup.

With a sufficient number of viewpoints, we can perform tomographic reconstruction of the Schlieren volumes by applying standard algorithms such as algebraic reconstruction (ART) [25]. Equation 2.10 shows how the optical flow vectors correspond directly to a line integral of the deflections along the ray paths. We capture the integrated deflections through BOS, and attempt to recover the refractive index field. Under the paraxial approximation, reconstruction via standard methods has been performed by a few researchers in the past. Schwarz [39] constructed a device with 20 pairs of

Schlieren mirrors arranged in a ring around a flame. With a single camera and a set of rotating mirrors he was able to capture a sequence of images from each of the view directions. Aside from the expense, the major disadvantage of this design is that dynamic media cannot be used. Since each view is captured at a different time, the flame could not be disturbed, ensuring that its thermal field remained constant throughout the acquisition.

The other limitation often imposed in Schlieren tomography is to consider only axially-symmetric media. Dynamic media such as flames and exhaust jets can then be captured, as long as they remain symmetric at all times. In this way, only a single viewpoint need be used, since we can assume it will look the same from every other direction. Agrawal et. al. [2] reconstructed symmetric fields using the Abel transform.

BOS makes it possible to capture multiple views of dynamic, large-scale, non-symmetric media, thanks to the simplicity and (relatively) low cost of the acquisition devices. A camera array was used to capture 8 views of a turbulent gas burner and an undisturbed candle plume, from which we reconstructed the volumes in Figure 5.9. Aside from being the first non-invasive method to capture time-varying, non-symmetric, inhomogeneous refractive index fields, we also removed the limitations of making the paraxial approximation in the reconstruction phase. Under stronger refractions, rays can be deflected outside of the voxels lying along their original paths. If this happens, ART will back-project any residual error into the wrong voxels, producing incorrect reconstructions. We employ an iterative process, alternately reconstructing the index field and then updating the light paths by raytracing. This ensures that after convergence the correct voxels are being updated. In this framework, much stronger refractions (such as those encountered in liquid mixing) can be accommodated. Our preliminary results indicate that the method shows much promise, although with only 8 views available we are not yet able to capture high frequency detail.

5.6 Conclusions

Although a very simple technique, BOS makes possible a number of interesting applications. Some of these, such as environment matting are possible through multiple

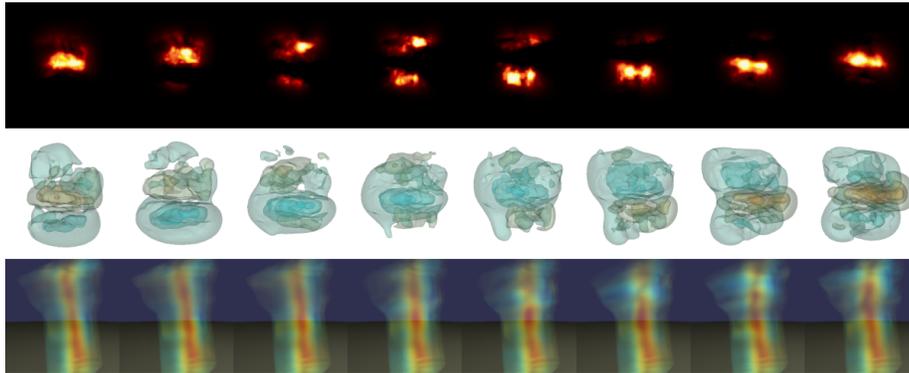


Figure 5.9: Tomography results. Top row: maximum intensity projections of gas burner refractive index field reconstructions. Middle row: isosurface renderings of the same sequence. Bottom row: volume renderings of laminar candle plume reconstruction.

methods, but the simplicity of BOS makes them more accessible. Other, such as tomography of general transparent gases, are not practical via any other method. All allow for subjective visualisation of a certain phenomenon. Objective analysis at a crude level is possible for flow visualisation (in that we can state where turbulence occurs, measure sizes and angles of plumes and so forth) but to obtain true PIV-like flowfields a more complete 3D reconstruction of the refractive index field is required. Moving forward, we would like to see a BOS-based system using an array of reasonably cheap cameras, able to build volumetric reconstructions of dynamic fluid interactions for such applications as wind-tunnel testing and numerical fluid simulation verification. With the increasing quality of image sensors, and a more sophisticated calibration system in place to better estimate the true world-space ray paths, such a system will hopefully become a reality.

Bibliography

- [1] S. Agarwal, S. Mallick, D. Kriegman, and S. Belongie. On refractive optical flow. *Lecture Notes in Computer Science*, (3022):483–494, 2004.
- [2] A. Agrawal, B. Albers, and D. Griffen. Abel inversion of deflectometric measurements in dynamic flows. *Applied Optics*, 38(15):3394–3398, May 1999.
- [3] E. Arnaud, E. Mémin, R. Sosa, and G. Artana. A fluid motion estimator for schlieren image velocimetry. *Lecture Notes in Computer Science*, (3951):198–210, 2006.
- [4] J.L. Barron, D.J. Fleet, and S.S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, February 1994.
- [5] M.J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding*, 63(1):75–104, 1996.
- [6] M. Born and E. Wolf. *Principles of Optics*. Cambridge University Press, 7th edition, 1999.
- [7] A. Bruhn, J. Weickert, T. Kohlberger, and C. Schnörr. A multigrid platform for real-time motion computation with discontinuity-preserving variational methods. *70(3):257–277*, 2006.
- [8] R. Budwig. Refractive index matching methods for liquid flow investigations. *Experiments in Fluids*, 17(5):350–355, 1994.

-
- [9] Y.Y. Chuang, D.E. Zongker, J. Hindorff, B. Curless, and D.H. Salesin. Environment matting extensions: Towards higher accuracy and real-time capture. In *Proceedings of ACM SIGGRAPH*, pages 121–130, August 2000.
- [10] R.L. Cook and T. DeRose. Wavelet noise. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 24(3):803–811, July 2005.
- [11] T. Corpetti, D. Heitz, G. Arroyo, E. Mémin, and A. Santa-Cruz. Fluid experimental flow estimation based on an optical flow scheme. *Experiments in Fluids*, 40(1):80–97, January 2006.
- [12] G.E. Elsinga, B.W. van Oudheusden, F. Scarano, and D.W. Watt. Assessment and application of quantitative schlieren methods: Calibrated color schlieren and background oriented schlieren. *Experiments in Fluids*, 36(2):309–325, 2004.
- [13] D. Di Florio, D. Di Felice, and G.P. Romano. Windowing, re-shaping and re-orientation interrogation windows in particle image velocimetry for the investigation of shear flows. *Measurement Science and Technology*, 13(7):953–962, 2002.
- [14] S. Fu and Y. Wu. Detection of velocity distribution of a flow field using sequences of schlieren images. *Optical Engineering*, 40(8):1661–1666, August 2001.
- [15] T. Georgiev, K.C. Zheng, B. Curless, D. Salesin, S. Nayar, and C. Intwala. Spatio-angular resolution tradeoff in integral photography. In *Proceedings of Eurographics Symposium on Rendering*, 2006.
- [16] A. Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufman, San Francisco, California, 1995.
- [17] I. Grant. Particle image velocimetry: A review. *Journal of Mechanical Engineering Science*, 211(1):55–76, 1997.
- [18] D. Gutierrez, A. Munoz, O. Anson, and F.J. Seron. Non-linear volume photon mapping. In *Proceedings of EGSR*, pages 291–300, 2005.
- [19] P. Hariharan. Optical interferometry. *Reports on Progress in Physics*, 54:339–390, 1990.

-
- [20] B.K.P. Horn and B.G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [21] W.L. Howes. Rainbow schlieren. Technical Report TP-2166, NASA, May 1983.
- [22] I. Ihrke. *Reconstruction and Rendering of Time-Varying Natural Phenomena*. PhD thesis, Universität des Saarlandes, 2007.
- [23] H.W. Jensen. Global illumination using photon maps. In *Proceedings of the 7th Eurographics Workshop on Rendering Techniques*, pages 21–30, Porto, Portugal, 1996.
- [24] D.R. Jonasses, G.S. Settles, and M.D. Tronosky. Schlieren piv for turbulent flows. *Optics and Lasers in Engineering*, 44:190–207, 2006.
- [25] A.C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. Society of Industrial and Applied Mathematics, 2001.
- [26] B.M. Klingner, B.E. Feldman, N. Chentanez, and J.F. O’Brien. Fluid animation with dynamic meshes. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 25(3):820–825, July 2006.
- [27] J.P. Lewis. Fast normalized cross-correlation. In *Vision Interface*, pages 120–123. Canadian Image Processing and Pattern Recognition Society, 1995.
- [28] B.D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of DARPA Image Understanding Workshop*, pages 121–130, 1981.
- [29] G.E.A. Meier. Hintergrund schlierenmeverfahren, 1999. Deutsche Patentanmeldung. Patent number DE 199 42 856 A1.
- [30] R.J. North. Schlieren systems using graded filters. Technical Report ARC Report 15099, British Aeronautical Research Council, 1952.
- [31] N. Papenberg, A. Bruhn, T. Brox, S. Didas, and J. Weickert. Highly accurate optic flow computation with theoretically justified warping. *International Journal of Computer Vision*, 67(2):141–158, 2006.

-
- [32] P. Peers and P. Dutré. Wavelet environment matting. In *Proceedings of the 14th Eurographics Workshop on Rendering*, pages 157–166, Leuven, Belgium, 2003.
- [33] P. Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- [34] H. Richard and M. Raffel. Principle and applications of the background oriented schlieren (bos) method. *Measurement Science and Technology*, 12:1576–1585, 2001.
- [35] S. Rusinkiewicz, O. Hall-Holt, and M. Levoy. Real-time 3d model acquisition. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 21(3):438–446, July 2002.
- [36] F. Scarano. Iterative image deformation methods in piv. *Measurement Science and Technology*, 13(1):R1–R19, 2002.
- [37] H. Schardin. Die schlierenverfahren und ihre anwendungen. *Ergebnisse der Exakten Naturwissenschaften*, 20:303–439, 1942. English translation published as NASA-TT-F-12731, April 1970.
- [38] S.P. Schneider. Laminar-flow design for a mach-6 quiet-flow wind tunnel nozzle. *Current Science*, 79(6):790–799, September 2000.
- [39] A. Schwarz. Multi-tomographic flame analysis with a schlieren apparatus. *Measurement Science and Technology*, 7:406–413, 1996.
- [40] G.S. Settles. Imaging gas leaks using schlieren optics. *American Society of Heating, Refrigerating and Air-Conditioning Engineers Journal*, pages 19–26, 1997.
- [41] G.S. Settles. Schlieren and shadowgraph imaging in the great outdoors. In *Proceedings of the 2nd Pacific Symposium on Flow Visualization and Image Processing*, Honolulu, HI, May 1999.
- [42] G.S. Settles. *Schlieren and Shadowgraph Techniques*. Springer, Berlin, 2001.

-
- [43] M. Thomas, S. Misra, C. Kambhamettu, and J.T. Kirby. A robust motion estimation algorithm for piv. *Measurement Science and Technology*, 16(3):865–877, 2005.
- [44] B. Trifonov, D. Bradley, and W. Heidrich. Tomographic reconstruction of transparent objects. In *Proceedings of Eurographics Symposium on Rendering*, 2006.
- [45] L. Venkatakrisnan and G.E.A. Meier. Density measurements using the background oriented schlieren technique. *Experiments in Fluids*, 37:237–247, 2004.
- [46] J.A. Volpe and G.S. Settles. Laser-induced gas breakdown as a light source for schlieren and shadowgraph particle image velocimetry. *Optical Engineering*, 45(8), August 2006.
- [47] F.J. Weinberg. *Optics of Flames: Including Methods for the Study of Refractive Index Fields in Combustion and Aerodynamics*. Butterworths, London, UK, 1963.
- [48] Y. Wexler, A.W. Fitzgibbon, and A. Zisserman. Image-based environment matting. In *Proceedings of the 13th Eurographics Workshop on Rendering*, pages 279–290, Pisa, Italy, 2002.
- [49] D.E. Zongker, D.M. Werner, B. Curless, and D.H. Salesin. Environment matting and compositing. In *Proceedings of ACM SIGGRAPH*, pages 205–214, August 1999.