

Bradley Atcheson · Wolfgang Heidrich · Ivo Ihrke

# An evaluation of optical flow algorithms for background oriented schlieren imaging

Preprint (accepted: Sep 2008)

**Abstract** The background oriented schlieren method (BOS) allows for accurate flow measurements with a simple experimental configuration. To estimate per-pixel displacement vectors between two images, BOS systems traditionally borrow window-based algorithms from particle image velocimetry.

In this paper, we evaluate the performance of more recent optical flow methods in BOS settings. We also analyze the impact of different background patterns, suggesting the use of a pattern with detail at many scales.

Experiments with both synthetic and real datasets show that the performance of BOS systems can be significantly improved through a combination of optical flow algorithms and multiscale background.

**Keywords** Background Oriented Schlieren Imaging · Optical Flow

---

## 1 Introduction

Background oriented schlieren imaging (BOS) has in recent years evolved as an attractive option for fluid imaging (e.g. Meier, 2002; Jensen et al, 2005). BOS systems are inexpensive and easy to set up, and produce quantitative ray deflection results. As such, they are ideal not only for 2D imaging, but also for tomographic setups (Raffel et al, 2000; Venkatakrishnan and Meier, 2004; Goldhahn and Seume, 2007; Atcheson et al, 2008).

BOS systems record per-pixel ray deflections by comparing an undistorted view of a background with a distorted view, in which the background is viewed through a refracting flow. Typically, rectangular neighborhoods of pixels in the two images are compared using a cross-correlation metric developed for particle imaging velocimetry (Westerweel, 1997). If matching neighborhoods are found, their relative

position in the two images defines a 2D displacement vector representing the refraction of light due to the variation of refractive index in the flow.

This estimation of displacement vectors from an image pair is very similar to a problem that has received a lot of attention in the computer vision community for many years. Computer vision researchers are interested in *optical flow*, that is, the apparent motion of objects in 2D images from one time step to another (see Section 2 for a more detailed description of the problem). After initial work with cross-correlation methods similar to those used in PIV and BOS, more accurate gradient-based and variational methods have been developed for this problem over the past two decades.

One difference between the typical optical flow setting and BOS is, however, that optical flow is usually applied to relatively simple motions, such as camera movement or rigid motions of individual objects. The resulting displacement field is typically very smooth with a few discontinuities at object edges. In contrast, BOS datasets usually do not have sharp discontinuities, but contain other high frequency variations, especially when measuring turbulent flow.

In this paper, we therefore analyze the performance of optical flow methods when applied to datasets that are more typical of BOS problems. We specifically compare the standard window matching algorithm to both gradient-based and variational optical flow methods. Although a vast number of variations have been proposed on both approaches over the years, a number of representative methods can be chosen for each category. The algorithms by Lucas and Kanade (1981) and Horn and Schunck (1981), as well as their multi-resolution, multi-scale variants represent the classic gradient-based approaches. Similarly, the algorithm by Brox et al (2004) is one of the most popular variational approaches. In our experiments, we test these algorithms on both real BOS datasets as well as synthetic ones, the latter for comparisons with ground truth data.

Another, previously unexplored, avenue for improvement lies with the control we have over the background pattern. While in PIV the pattern is given through the distribution of particles in a flow, for BOS we can choose any background image. In particular, we can choose a pattern

that has a dense set of distinguishable features across a wide range of scales. Specifically, we analyze the performance of the *wavelet noise* pattern (Cook and DeRose, 2005), to see if it can help in estimating a dense set of pixel correspondences between the images, while providing a large degree of flexibility in the optical setup.

It should be noted that recent research in PIV imaging has led to algorithms similar to those used in optical flow research, such as the one by Ruhnau et al (2005), which proposes a multi-scale and multi-resolution variant of the Horn and Schunck algorithm similar to one of the algorithms we evaluate. Other recent PIV algorithms make use of domain-specific assumptions such as incompressibility of flow (Corpetti et al, 2006; Ruhnau and Schnörr, 2007). These assumptions are not valid for BOS, and hence these more specialized algorithms cannot be applied in this setting.

In the following, we first give an introduction to the optical flow problem, and discuss the individual algorithms used in our evaluation (Section 2). We then describe a new background pattern with detail across a wide range of scales (Section 3), before we present the experiments on synthetic and real data (Sections 4 and 5, respectively).

## 2 Optical Flow

Given two images  $I(x, y, t_0)$  and  $I(x, y, t_1)$ , the *optical flow* (e.g. Davies, 2004) is defined as the 2D vector field describing the *apparent motion* of each pixel between images of a dynamic 3D scene taken at different points in time. This apparent motion is computed using the assumption of *brightness constancy*, that is, the assumption that pixel intensity for corresponding 3D points is the same in both images:  $I(x, y, t) = I(x + \delta x, y + \delta y, t + \delta t)$ . (1)

This problem is very similar to the pixel correspondence problem solved in BOS methods, although the characteristics of the apparent motion field differ from those of the typical optical flow field in computer vision. BOS optical flow generally does not exhibit sharp discontinuities, but can contain high frequency features distributed over a large area, especially when BOS is used to image turbulent flows. By comparison, for optical flow in 3D scenes, brightness constancy is violated when the scene contains shadows, specular highlights, and foreground objects that occlude the background. Much of the complexity that exists within modern optical flow algorithms is designed to deal with these cases, as well as with homogeneous image regions by assuming smooth flow fields with occasional discontinuities.

Despite these differences, the basic optical flow framework is an appropriate model for the refractions in BOS imaging. In particular, brightness constancy (Equation 1) holds for typical flows of interest, as long as the BOS setup is well engineered. In particular, brightness constancy assumes

- transparent, non-scattering flows with only small density gradients to minimize internal reflections and dispersion effects. Refraction angles are usually in the order of a small fraction of a degree.

- stable lighting from sources such as incandescent or metal halide lamps. Flickering sources such as fluorescent lights are to be avoided.
- approximately diffuse material for the background pattern to eliminate angular variations of radiance levels for a point on the pattern as observed from different directions. Due to the small magnitude of deflections, this assumption only places very minor constraints on the material used as a background.
- lastly, a small camera aperture, such that refocusing effects from the volumetric refractions can be neglected. Note that small apertures are already required in BOS systems in order to keep both the volume and patterned background in focus at the same time (see, e.g. Goldhahn and Seume, 2007).

All these assumptions are consistent with the state of the art in BOS imaging setups (Venkatakrishnan and Meier, 2004; Goldhahn and Seume, 2007). In the remainder of this work, we assume a BOS system optimized for the above considerations. Our work analyzes the best choice of optical flow algorithm for extracting information about light refraction from images acquired with such a setup.

In the following, we first review the block-matching algorithms commonly applied to BOS and PIV problems, and then summarize some gradient-based and variational methods developed in computer vision.

### 2.1 Block-Matching Algorithms

PIV processing software typically makes use of spatial correlation in order to detect the motion of small rectangular windows from frame to frame. As shown by Elsinga et al (2004), the same software may be used to process BOS data. For each overlapping window, the method computes the correlation score for each integer translational movement. These computations can be executed in the frequency domain for efficiency reasons. Integer-precision is insufficient for many applications, and so a subpixel accurate peak in the correlation matrix is located by fitting a curve through the points around the maximum value and then finding its peak. The computational cost for this algorithm is high, but can be mitigated with specialized, or parallel computing hardware.

Each output vector (ideally) represents the average motion of all pixels across the associated window. In order to maintain stability, these windows are often quite large ( $8 \times 8$  pixels and higher), and the assumption of uniform purely translational motion across such a large image region is often not valid. Deforming the windows has been shown to produce significantly improved results (Scarano, 2002), but at the cost of much increased complexity and execution time. Windows should be as small as possible to prevent excessive smoothing due to averaging over the windows. However, small windows may not contain sufficient image detail in order to produce a reliable peak in the correlation matrix. Some errors can be filtered out in post-processing, but one technique that does generally improve results is to apply the

algorithm in an iterative fashion, first using large windows, and then using the resulting vector fields to pre-shift smaller windows on successive iterations (Scarano and Riethmuller, 1999).

## 2.2 Gradient-Based Algorithms

Gradient-based methods are based on the first degree Taylor series expansion of the change at each pixel

$$I(x + \delta x, y + \delta y, t + \delta t) = I(x, y, t) + \frac{\partial I}{\partial x} \delta x + \frac{\partial I}{\partial y} \delta y + \frac{\partial I}{\partial t} \delta t + \dots, \quad (2)$$

where the higher order terms are ignored. Taking Equation 1 into account and dividing throughout by  $\delta t$  we get that

$$\frac{\partial I}{\partial x} \frac{\delta x}{\delta t} + \frac{\partial I}{\partial y} \frac{\delta y}{\delta t} + \frac{\partial I}{\partial t} = 0. \quad (3)$$

The  $\partial I/\partial x$ ,  $\partial I/\partial y$  and  $\partial I/\partial t$  terms are easily computable image derivatives. For each pixel, we then have one equation in two unknowns,  $\delta x$  and  $\delta y$ . Additional constraints are thus required to get a unique solution, and they are usually provided by neighboring pixels.

### 2.2.1 Lucas-Kanade

Equation 3 can be solved at a given point by considering all the pixels  $(i, j)$  within a fixed-size window around that point and constructing the following system of equations (Lucas and Kanade, 1981):

$$\left( \frac{\partial I}{\partial x} \frac{\delta x}{\delta t} + \frac{\partial I}{\partial y} \frac{\delta y}{\delta t} + \frac{\partial I}{\partial t} \right) \Big|_{i,j,t} = 0 \quad (4)$$

Since the system is now overdetermined,  $\delta x$  and  $\delta y$  are determined using a least-squares optimization, which also improves robustness under noise.

### 2.2.2 Horn-Schunck

Horn and Schunck (1981) take an alternative approach to solving the aforementioned underdetermined system, by adding a regularizing term in order to enforce global smoothness. The idea is to minimize the function

$$\int_{\Omega} \left( \frac{\partial I}{\partial x} \delta x + \frac{\partial I}{\partial y} \delta y + \frac{\partial I}{\partial t} + \alpha \left( |\nabla \delta x|^2 + |\nabla \delta y|^2 \right) \right) dx dy, \quad (5)$$

which combines Equation 3 (with  $\delta t = 1$  for convenience) with a weighted smoothness term based on the magnitudes of the vector field gradients. The direct incorporation of global information removes the necessity to post-process the vector field to remove outliers due to incorrect matches.

### 2.2.3 Multi-resolution and Multi-scale Extensions

In images with high frequency content, gradient information is only useful at a very small scale, and displacements of more than one pixel cannot be measured. The standard approach to dealing with this problem is to use a *multi-resolution* coarse-to-fine algorithm. An image pyramid is constructed by repeatedly downsampling the image by a factor of two. The optical flow can then be found on the smallest image in the pyramid, and is used to unwarp the next smallest image by computing  $I^{(n)}(x - \delta x(x, y, t), y - \delta y(x, y, t), t)$  where  $\delta x$  and  $\delta y$  are upsampled and scaled from the previous level. Interpolation is used for the fractional pixel locations. This process is then iterated until reaching the original image resolution (Anandan, 1989).

Block-matching algorithms use an iterative process, starting with larger interrogation window sizes, to deal with large displacements. With a large window overlap, this can significantly increase the execution time. Image pyramids serve the same purpose, but allow for using a constant window size on successively smaller images, with much lower computational cost (Adelson et al, 1984).

In many cases, the performance of multi-resolution algorithms can be further improved by controlling the spatial frequency content more finely than with power-of-two image pyramids. This is usually achieved by filtering the individual pyramid levels with a low-pass filter to generate different *scales*. An example of such a *multi-scale* algorithm is the work by Ruhnau et al (2005), who apply it to the Horn-Schunck framework.

The disadvantage of both multi-resolution and multi-scale methods is that some image detail is lost. In an image consisting of uniform or normally distributed noise, repeated downsampling and/or filtering results in an image that converges to a uniform grey. The lack of detail in those images makes it difficult to compute optical flow. We show in Section 3 that this problem can be avoided by making use of a multi-scale noise pattern.

## 2.3 Variational Algorithms

Recent work on optical flow has turned towards variational approaches, which seek to minimize an energy functional over the image domain. Whereas earlier gradient-based approaches linearize the optical flow equation in the problem formulation, Papenberg et al (2006) extend an algorithm proposed by Brox et al (2004), and perform a non-linear optimization. They report greatly improved results on standard optical flow test datasets as a result. Their method, expressed below in equation form, is based on a generalization of the brightness constancy assumption to include constancy of other linear functions  $\mathcal{L}_i$  (*i.e.*, identity, gradient, Laplacian) of the image intensity. The norms of these weighted data terms  $D_i$  are minimized across the whole image. The penalty function  $\Psi$ , which asymptotically approaches some maximum value, is applied to each term to prevent outliers

from skewing the result, and a smoothness term is included that is similar to that of the Horn-Schunck method.

$$D_i = \|\mathcal{L}_i(I(x + \delta x, y + \delta y, t + \delta t)) - \mathcal{L}_i(I(x, y, t))\|_2^2 \quad (6)$$

$$E = \int_{\Omega} \left( \Psi \left( \sum_i \alpha_i D_i \right) + \alpha \Psi(|\nabla \delta x|^2 + |\nabla \delta y|^2) \right) dx dy \quad (7)$$

## 2.4 Algorithm Parameters

The accuracy of optical flow estimation is strongly influenced not only by the choice of algorithm, but also by choice of parameter values for a given algorithm.

In the case of the block-matching and Lucas-Kanade algorithms, window size ( $w$ ) is the key parameter that balances stability with spatial resolution. Larger windows provide a higher degree of confidence in the vector field, but do not reveal fine-scale detail. The goal is therefore to make  $w$  as small as possible while maintaining an acceptable degree of stability.

The Horn-Schunck and Brox algorithms are chiefly controlled by the smoothness parameter  $\alpha$ , which damps high frequency spurious errors, as well as turbulent flow. Therefore,  $\alpha$  should be made as small as possible to prevent over-smoothing, but beyond a certain (data-dependent) threshold, the effect of the regularizing smoothness term is diminished to the point that Equations 5 and 7 again become underdetermined.

The Brox algorithm contains an additional set of parameters  $\alpha_i$  controlling the respective weights of the data terms  $D_i$ . Our implementation of this algorithm contains only the identity and gradient operators  $\mathcal{L}_0$  and  $\mathcal{L}_1$ , and since the gradient information in a noisy image can be very high frequency, we fixed  $\alpha_0$  and  $\alpha_1$  to values experimentally determined to produce the best results (0.9 and 0.1, respectively), and only analyze the impact of the remaining parameter  $\alpha$ .

## 3 Wavelet Noise

All optical flow methods rely on inhomogeneous, high-frequency image detail for the estimation of displacement. With PIV, a flow is seeded with particles in an attempt to indirectly control this image detail. For outdoor experiments, Kindler et al (2007) have shown that natural scenes are often suitable for BOS, while for laboratory experiments, random noise functions are often used as background (Richard and Raffel, 2001). We argue that whenever direct control of the background pattern is possible, a multi-scale pattern, such as that described in this section, can significantly improve the accuracy of results.

For most random noise patterns, such as uniform or Gaussian noise, the intensity histogram is affected by changes in scale due to averaging of noise values over

pixel sensors. In particular, downsampled uniform noise obtains a Gaussian intensity profile, and downsampled Gaussian noise has a smaller variance than the original (see Figure 2), resulting in a loss of detail and contrast. While the normalized cross correlation used in optical flow algorithms can partially compensate for this contrast reduction, quantization and random noise in the camera sensor result in significantly degraded optical flow estimates even for small levels of downsampling. Upsampled noise patterns, on the other hand, produce homogeneous color regions that do not provide enough detail for determining dense optical flow. Therefore, random noise patterns work best if the resolution of the pattern and that of the camera sensor are matched. This restriction, which proves cumbersome in practical setups, can be avoided with multi-scale patterns such as wavelet noise (Cook and DeRose, 2005), which have discernible high frequency content at all scales.

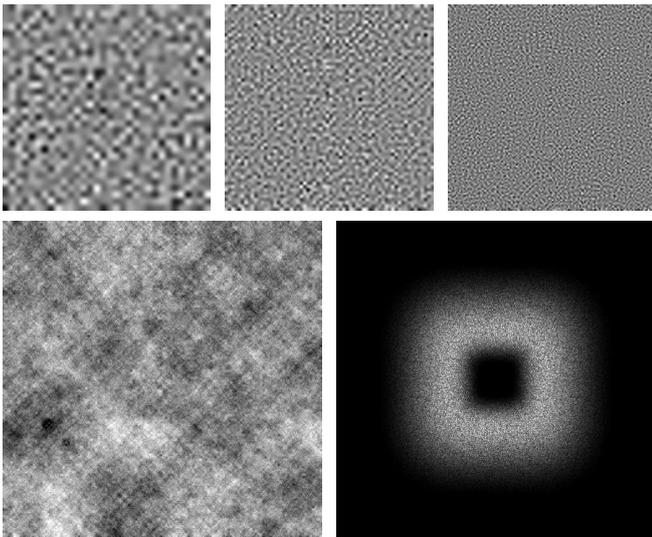
A wavelet noise pattern of size  $2^k \times 2^k$  is the sum of individual band-limited noise functions generated at resolutions of  $2^j \times 2^j$  for  $j \leq k$ . Each individual function is constructed by initializing a  $2^j \times 2^j$  image with uniformly distributed random pixel values. This image is then downsampled by a factor of two, and upsampled again to its original resolution. The difference between those two images is then effectively a high-pass filtered image which, when upsampled to the final resolution  $2^k \times 2^k$ , contains frequencies in only a narrow band. Together, these bands form a partition of the entire frequency space. The top of Figure 1 shows three such noise functions for different resolutions  $j$ . The bottom right of the figure shows the spectrum of one of these levels, revealing its band-limited nature. The final wavelet noise pattern is created by upsampling all the images to the final resolution of  $2^k \times 2^k$  pixels, and summing them (bottom left of Figure 1).

There are two reasons for using a multi-scale noise pattern. Firstly, it allows for hierarchical optical flow algorithms to safely downsample the images to handle large displacements, without concern for loss of accuracy at those stages. Secondly, as a practical matter it is tedious to reprint a background pattern of the appropriate resolution for every experimental configuration. Having a single large multi-scale pattern affords a large degree of flexibility with respect to the geometric and optical setup, which simplifies experiments.

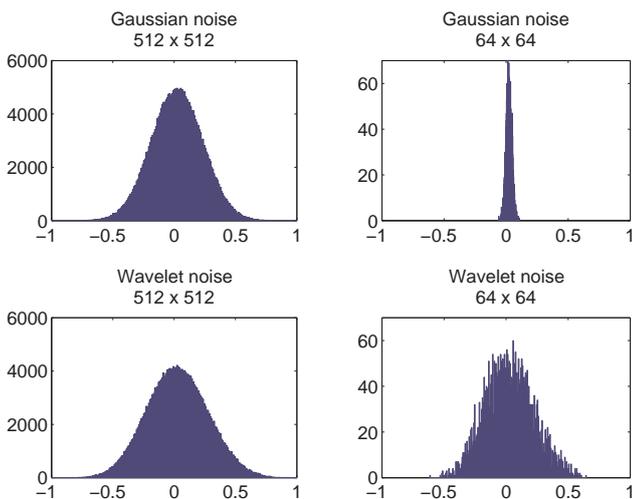
The problem of loss of detail on non-multi-scale patterns is illustrated in Figure 2. Viewing a high frequency, high dynamic range, normally distributed noise image from afar, or under minification, results in a flatter grey image with low dynamic range. Uniformly distributed noise, as well as many natural high frequency scenes, also exhibit this behavior, whereas the wavelet noise retains contrast and detail under scaling.

## 4 Synthetic Data Evaluations

In this section we present results from experiments with synthetic datasets. We focus on flows that exhibit typical characteristics of BOS data. Evaluations for data representative



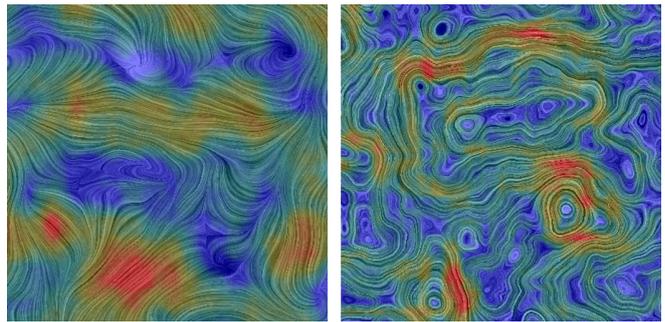
**Fig. 1** Wavelet noise. **Top row:** Three of the individual noise bands. **Bottom left:** the final noise image is the sum of eight such bands, equally weighted in this case. **Bottom right:** the frequency spectrum of the second band is confined to the range between one half and one quarter of the maximum frequency.



**Fig. 2** Background pattern intensity histograms show the effects of downsampling for two different types of noise.

of the apparent motion of 3D objects have been performed elsewhere (e.g. Barron et al, 1994).

Optical flow fields from BOS measurements can exhibit sources and sinks even if the underlying fluid flow is incompressible. At the same time, BOS optical flow may also contain rotational components, especially when imaging turbulent fluid flows. Both sources/sinks and rotational components may be expected to cause problems with the optical flow estimation, so to analyze the impact of both phenomena we created two types of datasets, illustrated in Figure 3 using Line Integral Convolution images (Cabral and Leedom, 1993). The first optical flow dataset was a simple 2D random



**Fig. 3** Random noise flowfield (left) and Curl-noise (right). Line-integral convolution is used to visualize the direction of flow, while the color-coding illustrates velocities.

noise function that was low-pass filtered to ensure varying degrees of smoothness. These random fields contain sources and sinks, but almost no rotational components. To investigate the effect of the latter, we also generated divergence-free vector fields containing many small vortices using the *curl noise* method by Bridson et al (2007).

Fields of each type were generated with varying levels of *turbulence*, where we simulate increased turbulence by allowing higher frequency variations in the field. The low turbulence fields represent what would typically be seen with a laminar flow such as an undisturbed candle plume (Figure 7), whereas high turbulence represents a flow more complex than that of the interaction depicted in Figure 8. More turbulent flows are expected to be more problematic for optical flow estimation since they require the use of smaller window sizes in both the block matching and Lucas-Kanade algorithms, as well as a lower weights for the regularization terms in the Horn-Schunck and Brox algorithms. For all algorithms, the robustness of the optical flow estimation is therefore expected to decrease with increasing turbulence. Note that because the divergence-free flow fields are constructed by taking gradients of the random flow fields, they exhibit higher frequency variations. Error magnitudes between both types of flow fields can therefore not be compared directly.

Three background patterns were chosen and subsequently warped by each vector field. Normally distributed noise was selected as a baseline in order to evaluate the newer wavelet noise pattern (Section 3). A sample high density particle image from a PIV experiment was also used, in order to determine whether or not the other optical flow algorithms could be carried over to PIV. In order to avoid the effects of sampling artifacts, we generated all images and fields at higher than necessary resolution, and downsampled before computing the displacements. Additive Gaussian noise with a standard deviation of 3 grey levels was added to all images to model an imperfect image sensor.

For each combination of background pattern and vector field, we computed the optical flow using each algorithm and varied the key parameter across an empirically determined range. For Horn-Schunck we used  $\alpha$  values of 0.06,

0.02 and 0.006. With smaller values the algorithm became unstable and severe errors appeared in the output. The Brox smoothness term was set to 6, 5 and 4. Values smaller than these resulted in very poor results, and larger ones caused the algorithm to crash in certain cases due to a singular matrix appearing in the Gauss-Seidel solver. For Lucas-Kanade, square windows of size 7, 5 and 3 pixels were used and for block-matching we used the open-source MatPIV toolbox (Sveen, 2004) with  $w = 32, 16$  and 8 pixels. In the hierarchical algorithms, image pyramids were restricted to three levels, and MatPIV used three iterations of successively decreasing window sizes, with a window overlap of  $7/8$  to maximize spatial resolution.

For the Horn-Schunck algorithm, we experimented both with the multi-scale, multi-resolution variant of the algorithm, and the multi-resolution-only variant. We consistently found the former method slightly superior to the latter, albeit at somewhat higher computational cost. The tables in the following only report the results from the multi-resolution, multi-scale algorithm.

For evaluation we cropped all results to remove the difficult border cases, and computed average angular and endpoint errors. Endpoint error is simply the magnitude of the difference between the ground truth  $\mathbf{g} = (g_u, g_v)$  and computed optical flow  $\mathbf{f} = (f_u, f_v)$  vectors

$$e = \|(f_u, f_v) - (g_u, g_v)\|_2. \quad (8)$$

#### 4.1 Results

The endpoint error results are shown in Figure 4. We found that angular errors follow the same trends as endpoint errors, and have omitted them for the sake of brevity. The left column of the graphs shows results for random flow, whereas the right column shows the results for divergence-free flow. From top to bottom, the rows correspond to wavelet noise, Gaussian noise, and a PIV background pattern, respectively. Within each graph, three groups of error bars represent the endpoint error for low, medium, and high turbulence. Each group contains the color coded error bars for different algorithms and parameter settings. In all cases, lower numbers are better. The results of less than 0.1 pixel error for the wavelet noise pattern with a high frequency random flow are encouraging (maximum deflection was around 3 pixels).

The gradient-based algorithms typically performed better than the block-matching and variational methods. Horn-Schunck was able to produce lower absolute errors than Lucas-Kanade in most cases, but proved to be more sensitive to parameter variations. The wavelet noise pattern resulted in generally lower errors than either of the other backgrounds. In some cases where large errors were produced with the Gaussian background, the same algorithm and parameters were able to produce a good result when using the wavelet background instead. We have observed that even with a conservatively chosen parameter, the combination of a gradient-based optical flow algorithm and the wavelet noise back-

ground pattern is often able to outperform MatPIV at its optimal parameter setting.

All algorithms can produce severe errors with the wrong choice of parameter values. The choice is usually between excessive smoothing (large windows or high weights for regularization terms), or low robustness (small windows or low regularization weights). For block-matching, we have used windows of at least  $8 \times 8$  pixels in size, below which the results become extremely unstable. Lucas-Kanade is stable for smaller window sizes such as  $5 \times 5$  or, in some settings,  $3 \times 3$ , allowing for a better reproduction of high-frequency detail. The continuous parameters of Brox and Horn-Schunck are more difficult to select, but with the right settings Horn-Schunck outperforms all other algorithms.

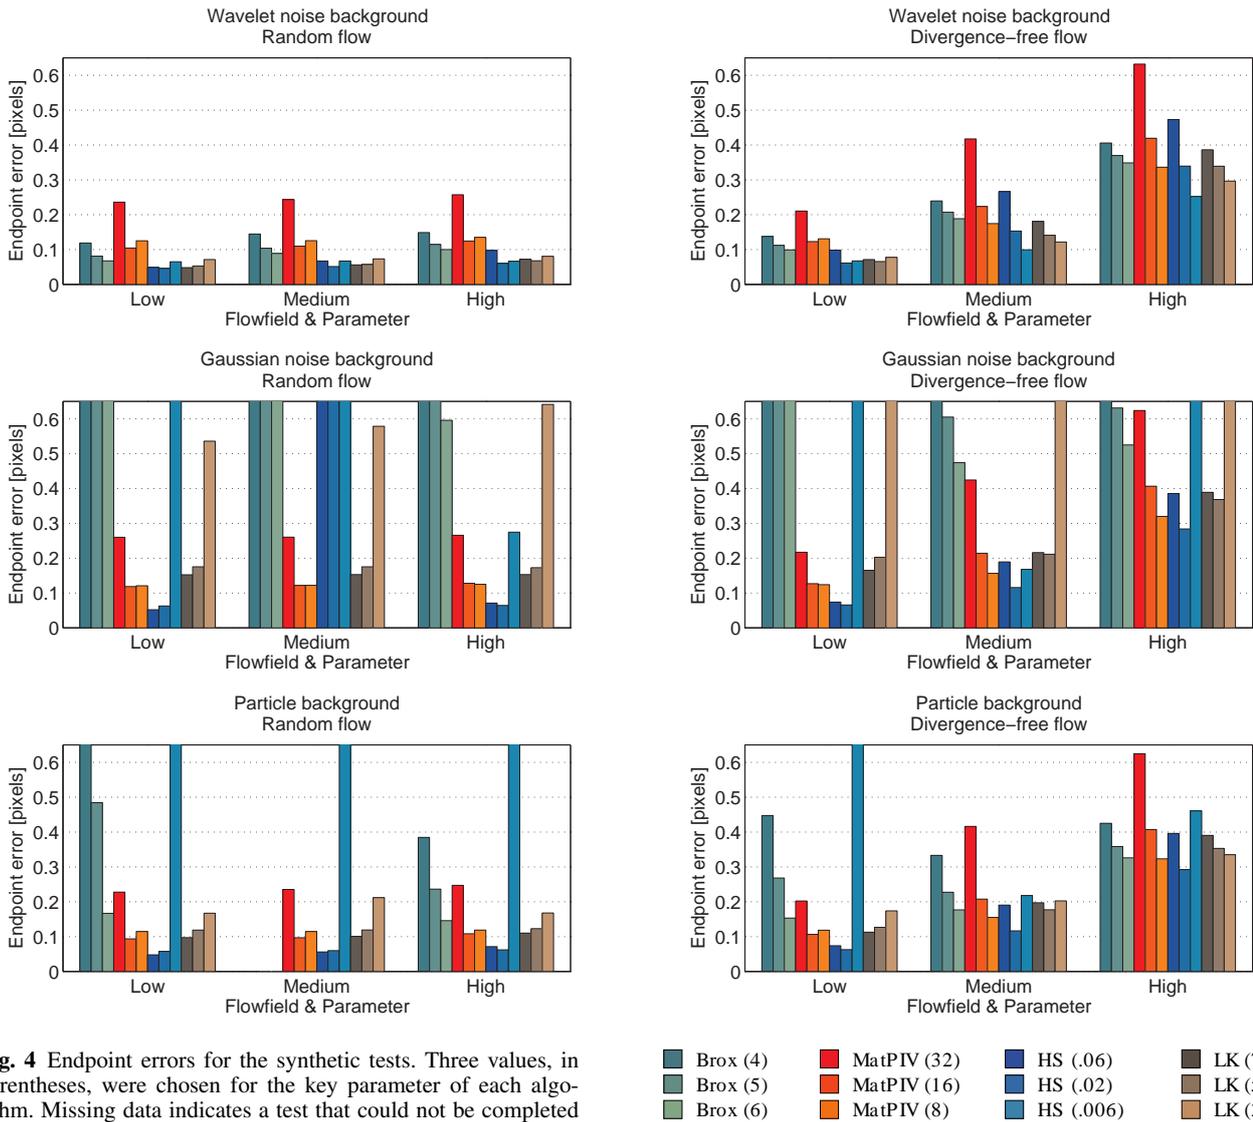
The PIV background returned results of similar quality to the Gaussian noise background, albeit with better stability. While the PIV image took on an approximately Gaussian intensity profile after downsampling, its dynamic range remained higher than that of the Gaussian pattern for the same amount of downsampling.

Timing results cannot be meaningfully compared, as both the MatPIV and Brox implementation were run in the interpreted Matlab environment rather than as optimized C code. Nevertheless, block-based cross-correlation is an inherently expensive operation, and algorithms that avoid this step can be made significantly faster. In the case of Horn-Schunck, we were able to process  $512 \times 512$  images in as little as 3 seconds on a standard desktop PC.

A *peak-locking* effect has been observed in PIV experiments (Westerweel, 2000). It arises when vectors are biased towards integer displacements. There is some evidence that this is due to the curve fitting procedure used to find the sub-pixel peak location, and various methods have been developed to mitigate it (Scarano, 2002). The histograms of displacement values in Figure 5 show that, for BOS, the effect is apparent under block-matching, whereas a gradient-based algorithm is largely immune.

#### 4.2 Scale tests

In order to test the efficacy of wavelet noise in allowing for resizing of the background, we conducted tests using identical algorithms and flow fields, varying only the background pattern and scale. Figure 6 shows a comparison of estimated displacement magnitudes for the Gaussian noise pattern (top) and the wavelet pattern (bottom), using Lucas-Kanade with  $w = 7$ . We chose this algorithm because it has the lowest parameter sensitivity. The left column shows the initial results at the original resolution of  $512 \times 512$ . Both fields are similar, but the Gaussian noise resulted in a few patches of spurious errors, which tend to be localized along contours of high gradient. On the right, we reduced the image sizes before computing the optical flow. Notice that the quality of the field produced with the Gaussian noise background decreased noticeably, whereas that produced with the wavelet noise pattern was largely unaffected.



**Fig. 4** Endpoint errors for the synthetic tests. Three values, in parentheses, were chosen for the key parameter of each algorithm. Missing data indicates a test that could not be completed due to instability.

## 5 Application to Acquired Data

While our synthetic tests show that the combination of gradient-based optical flow algorithms and a wavelet noise background is very effective, the question remains as to how accurately these results reflect true BOS data. To investigate this question, we conducted a qualitative evaluation on real-world datasets. Figure 7 shows the estimated horizontal displacements for a laminar candle plume. Figure 8 shows a more turbulent flow, where the candle plume is disturbed with a jet of compressed air. Finally, Figure 9 depicts the results of a flow from a candle lantern.

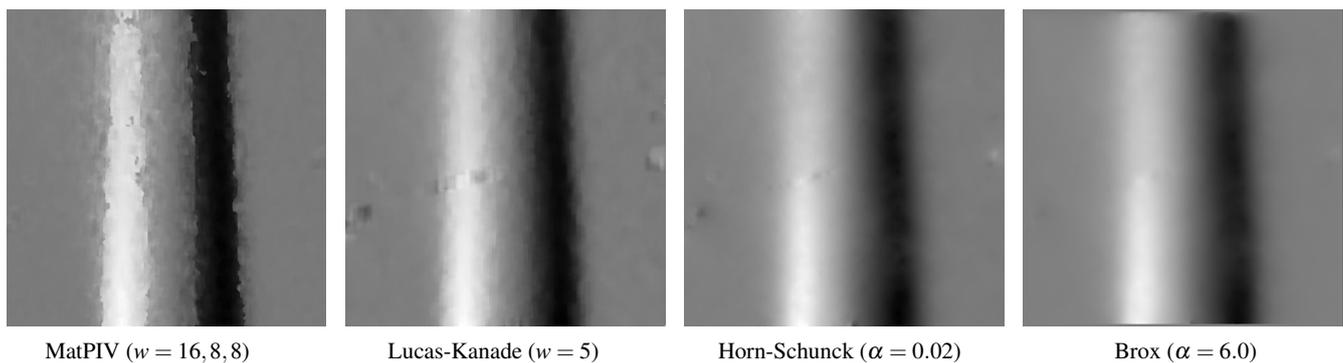
Peak locking is evident in the images generated with block-matching. Standard practice would be to low-pass filter these images, which would, however, sacrifice precision in turbulent regions. Due to the explicit regularization

term, Horn-Schunck produces smoother results than Lucas-Kanade. The Brox results are even smoother.

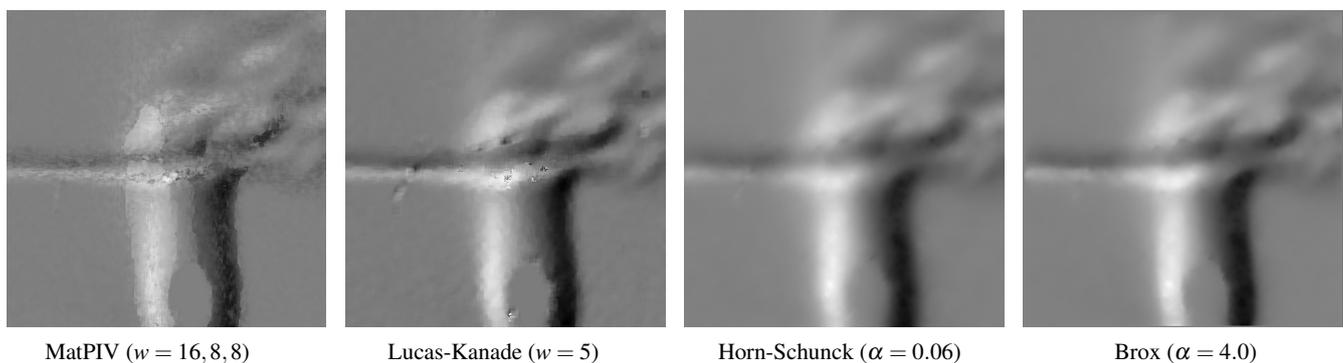
Finding optimal parameters for the different algorithms can be difficult without the luxury of a ground truth result. In the presented results, we estimated the best settings based on our experience with the synthetic data, and then refined the choice until the resulting image looked visually the best. Horn-Schunck and in particular Brox can be very sensitive to the setting of the smoothness parameter. By contrast, Lucas-Kanade is less sensitive to the window size, which we found we could always set to either  $3 \times 3$  or  $5 \times 5$  pixels.

## 6 Discussion and Conclusion

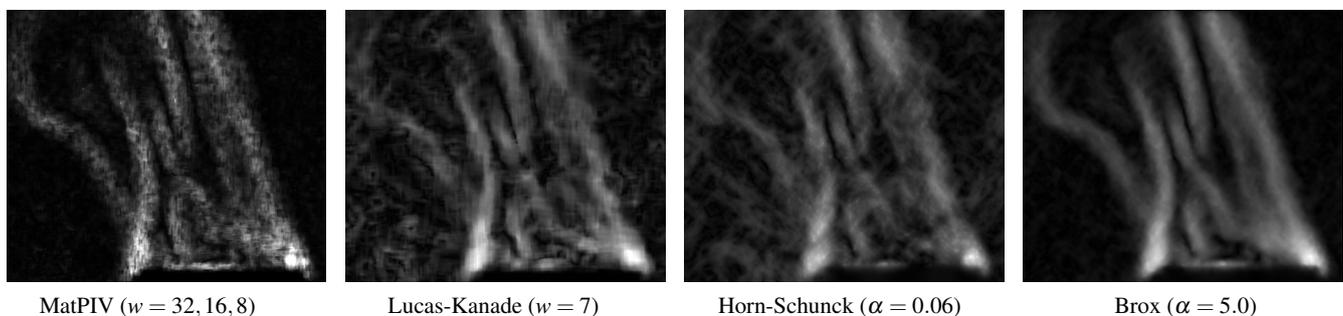
Our experiments confirm that BOS results can be improved upon by using more sophisticated algorithms than simple block matching. In particular, the multi-resolution,



**Fig. 7** Horizontal components of displacement vectors from a candle plume, imaged against the wavelet noise background.



**Fig. 8** Sum of horizontal and vertical vector components of a candle plume being disturbed by a jet of compressed air, imaged against the wavelet noise background.



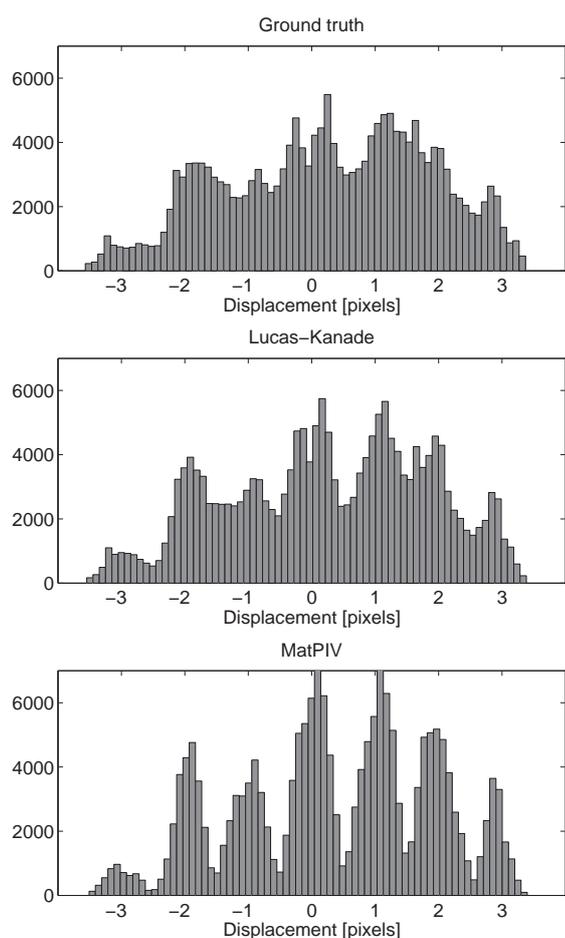
**Fig. 9** Magnitudes of displacement vectors of a candle lantern, imaged against the wavelet noise background.

multi-scale variants of gradient-based algorithms (Lucas and Kanade (1981) and Horn and Schunck (1981)) provide significantly better accuracy than block-matching algorithms, as shown in our comparisons with ground truth data. More recent variational algorithms such as Brox et al (2004) tend to excessively smoothen the higher-frequency content that is expected in BOS datasets.

In most cases we found that Horn-Schunck produces slightly lower error than Lucas-Kanade for the respective best parameter settings. The explicit regularization term in Horn-Schunck also provides visually superior results without the oversmoothing artifacts of the Brox algorithm. On

the other hand, Horn-Schunck is somewhat sensitive to the choice of parameter settings, and without ground-truth data, it may not be obvious how to set the weight of the regularization term for a particular dataset. By comparison, Lucas-Kanade only has one discrete-valued parameter, the size of the comparison window, which should be either  $3 \times 3$  or  $5 \times 5$  for best results.

Our experiments also show that a wavelet noise background does indeed further improve accuracy. Furthermore, experiments with scaling show that wavelet noise can, as expected, be used more easily across different scales, which simplifies the experimental setup for BOS measurements.



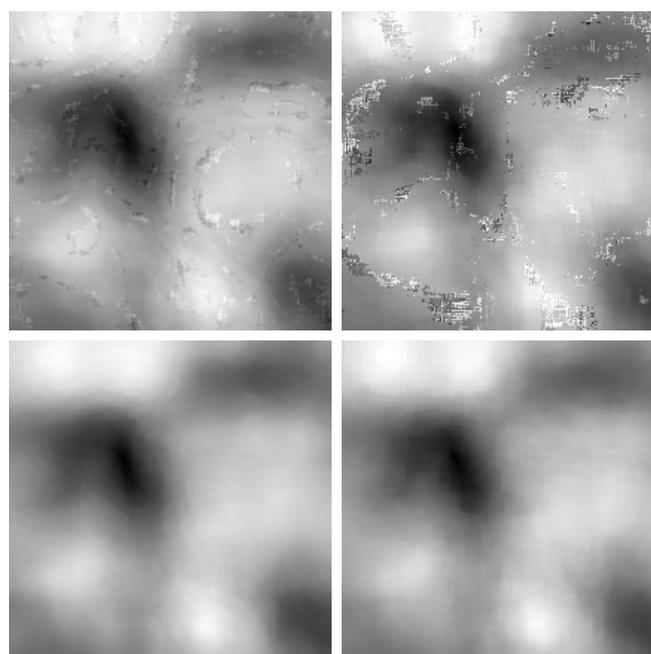
**Fig. 5** Histograms of vector horizontal components for the medium turbulence random flow, using a wavelet noise background.

Whenever one has control over the background in BOS experiments, the use of a wavelet noise background is a simple measure to increase accuracy.

On the other hand, Lucas-Kanade and Horn-Schunck improve results even without the use of a wavelet noise background, which suggests that these algorithms could also be attractive for particle imaging velocimetry with high density particle fields.

## References

- Adelson E, Anderson C, Bergen J, Burt P, Ogden J (1984) Pyramid methods in image processing. *RCA Engineer* 29(6):33–41
- Anandan P (1989) A computational framework and an algorithm for the measurement of visual motion. *IJCV* 2(3):283–310
- Atcheson B, Ihrke I, Heidrich W, Tevs A, Bradley D, Magnor M, Seidel HP (2008) Time-resolved 3D capture of non-stationary gas flows. *ACM Transactions on Graphics* 27(5)



**Fig. 6** Vector field magnitudes from the scaling tests. **Top row:** Gaussian noise. **Bottom row:** Wavelet noise. **Left column:** original resolution ( $512 \times 512$ ). **Right column:** downsampled to  $64 \times 64$ .

- Barron J, Fleet D, Beauchemin S (1994) Performance of optical flow techniques. *IJCV* 12(1):43–77
- Bridson R, Houriham J, Nordenstam M (2007) Curl-noise for procedural fluid flow. *ACM Transactions on Graphics* 26(3):46.1–3
- Brox T, Bruhn A, Papenbergh N, Weickert J (2004) High accuracy optical flow estimation based on a theory for warping. In: *Proc. 8<sup>th</sup> ECCV, Prague, Czech Republic*, pp 25–36
- Cabral B, Leedom LC (1993) Imaging vector fields using line integral convolution. In: *SIGGRAPH '93: Proc. 20<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, USA, pp 263–270
- Cook R, DeRose T (2005) Wavelet noise. *ACM Transactions on Graphics* 24(3):803–811
- Corpetti T, Heitz D, Arroyo G, Mémin E, Santa-Cruz A (2006) Fluid experimental flow estimation based on an optical-flow scheme. *Experiments in Fluids* 40(1):80–97
- Davies E (2004) *Machine vision: Theory, algorithms, practicalities*, Morgan Kaufmann, chap 18.2, pp 505–509
- Elsinga G, van Oudheusden B, Scarano F, Watt D (2004) Assessment and application of quantitative schlieren methods: Calibrated color schlieren and background oriented schlieren. *Experiments in Fluids* 36(2):309–325
- Goldhahn E, Seume J (2007) The background oriented schlieren technique: sensitivity, accuracy, resolution and application to a three-dimensional density field. *Experiments in Fluids* 43(2–3):241–249
- Horn B, Schunck B (1981) Determining optical flow. *Artificial Intelligence* 17:185–203

- 
- Jensen O, Kunsch J, Rösgen T (2005) Optical density and velocity measurements in cryogenic gas flows. *Experiments in Fluids* 39(1):48–55
- Kindler K, Goldhahn E, Leopold F, Raffel M (2007) Recent developments in background oriented schlieren methods for rotor blade tip vortex measurements. *Experiments in Fluids* 43(2–3):233–240
- Lucas B, Kanade T (1981) An iterative image registration technique with an application to stereo vision. In: Proc. 7<sup>th</sup> Intl Joint Conference on Artificial Intelligence, Vancouver, Canada, pp 674–679
- Meier G (2002) Computerized background-oriented schlieren. *Experiments in Fluids* 33(1):181–187
- Papenberg N, Bruhn A, Brox T, Didas S, Weickert J (2006) Highly accurate optic flow computation with theoretically justified warping. *IJCV* 67(2):141–158
- Raffel M, Richard H, Meier G (2000) On the applicability of background oriented optical tomography for large scale aerodynamic investigations. *Experiments in Fluids* 28(5):477–481
- Richard H, Raffel M (2001) Principle and applications of the background oriented schlieren (BOS) method. *Measurement Science and Technology* 12(9):1576–1585
- Ruhnau P, Schnörr C (2007) Optical stokes flow estimation: an imaging-based control approach. *Experiments in Fluids* 42(1):61–78
- Ruhnau P, Kohlberger T, Schnörr C, Nobach H (2005) Variational optical flow estimation for particle image velocimetry. *Experiments in Fluids* 38(1):21–32
- Scarano F (2002) Iterative image deformation methods in PIV. *Measurement Science and Technology* 13(1):R1–R19
- Scarano F, Riethmüller M (1999) Iterative multigrid approach in PIV processing with discrete window offset. *Experiments in Fluids* 26(6):513–523
- Sveen J (2004) An introduction to MatPIV v.1.6.1. Eprint no. 2, ISSN 0809-4403, Dept. of Mathematics, University of Oslo, <http://www.math.uio.no/~jks/matpiv>
- Venkatakrishnan L, Meier G (2004) Density measurements using the background oriented schlieren technique. *Experiments in Fluids* 37(2):237–247
- Westerweel J (1997) Fundamentals of digital particle image velocimetry. *Measurement Science and Technology* 8(12):1379–1392
- Westerweel J (2000) Theoretical analysis of the measurement precision in particle image velocimetry. *Experiments in Fluids* 29(7):S003–S012