

# Tomographic Reconstruction and Efficient Rendering of Refractive Gas Flows

Ivo Ihrke, Kai Berger, Bradley Atcheson, Marcus Magnor and Wolfgang Heidrich

**Abstract** This chapter introduces techniques for the capture and efficient display of dynamic three-dimensional non-stationary gas flows. We describe a flexible Schlieren-tomographic system consisting of multiple consumer camcorders. A special choice of background pattern for Background Oriented Schlieren (BOS) imaging provides for flexibility in the experimental setup. Optical flow techniques are used to measure image space deflections due to heated air flows from arbitrary camera positions. A specially tailored sparse-view algebraic reconstruction algorithm is employed to tomographically recover a refractive index gradient field. After robust integration of these gradient fields, time-varying, fully three-dimensional refractive index fields are obtained. These can be rendered efficiently using a ray-casting style algorithm that is suitable for graphics hardware acceleration. Additional optical properties can be rendered within the same computational framework.

---

Ivo Ihrke  
Imager Lab, University of British Columbia, Vancouver, BC, Canada,  
e-mail: [ivoihrke@cs.ubc.ca](mailto:ivoihrke@cs.ubc.ca)

Kai Berger  
Computer Graphics Lab, TU Braunschweig, Braunschweig, Germany,  
e-mail: [berger@cg.cs.tu-bs.de](mailto:berger@cg.cs.tu-bs.de)

Bradley Atcheson  
Imager Lab, University of British Columbia, Vancouver, BC, Canada,  
e-mail: [atcheson@cs.ubc.ca](mailto:atcheson@cs.ubc.ca)

Marcus Magnor  
Computer Graphics Lab, TU Braunschweig, Braunschweig, Germany,  
e-mail: [magnor@cg.cs.tu-bs.de](mailto:magnor@cg.cs.tu-bs.de)

Wolfgang Heidrich  
Imager Lab, University of British Columbia, Vancouver, BC, Canada,  
e-mail: [heidrich@cs.ubc.ca](mailto:heidrich@cs.ubc.ca)

## 1 Overview

Schlieren imaging techniques have long been used as a means to detect minute variations in refractive index. They convert angular deflections of light rays into intensity variations upon a surface, which can be more easily observed. Early systems placed a refractive medium (often a candle plume) inside a setup comprised of precision lenses and/or mirrors. A beam of collimated light, when shone through the medium, would pass by a filter before being focused onto an imaging plane. Any ray deflected from its original parallel path would be attenuated by this filter, resulting in a correspondingly darker spot on the imaging plane. Further details can be found in Settles' book [18], the most complete reference on Schlieren topics available today.

Classic Schlieren setups suffer from their high cost and difficulty of calibration. Recently, a different approach based on digital image processing has been proposed. The "Background Oriented Schlieren" (BOS) method [10, 15, 16] requires only a high frequency background pattern, positioned behind the refracting volume, and a camera to observe the pattern. Deflections of light rays passing through the medium cause apparent distortion of the background, which can be accurately measured using computer vision techniques.

While the projected refractive index variations can usually be interpreted by a trained user, they do not describe the full three-dimensional structure of the medium. For this, we need to perform a tomographic reconstruction of it, using Schlieren images captured from multiple viewpoints. Previous Schlieren-tomographic systems were made possible with only a single camera by operating on rotationally symmetric media [4] (in which views from any angle would be equivalent) or else on stationary flows [17] (in which the apparatus could be rotated to obtain more viewpoints). In our work, we present the first method to capture and reconstruct non-symmetric, non-stationary, time-varying refractive index variations [6]. We also develop a method to efficiently render the effects produced by inhomogeneous refractive index fields.

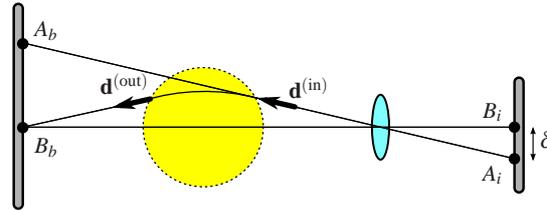
The chapter is structured as follows: we introduce the measurement setup for "Background Oriented Schlieren" capture of gas flows in Section 2 and describe their tomographic reconstruction in Section 3. In Section 4 we describe a light ray model which is used for the rendering of continuously varying refractive index fields such as the recovered gas flows. We conclude the chapter with our results and a discussion thereof in Section 5.

## 2 Background Oriented Schlieren Imaging

Any medium of inhomogeneous refractive index will cause light rays to bend as they traverse through it, leading to a characteristic *jittering* distortion of the background. BOS uses digital image processing techniques to convert this distortion directly into a measurement of the ray's total amount of deflection along its path.

Unlike with traditional Schlieren configurations, the BOS method has very moderate hardware requirements, which in turn makes multiple-viewpoint acquisition quite feasible. The basic principle, assuming a simple geometric optics model, is illustrated in Figure 1. A high frequency background pattern is viewed through a lens (implying that depth of field is of concern). Under normal circumstances, the ray travels straight, leading to point  $A_b$  being imaged at  $A_i$ . However, introducing a refractive medium into the scan volume causes the ray to bend, according to the well-known laws of optics. This causes an apparent distortion of the background pattern, with  $B_b$  now being imaged at  $A_i$ . *Optical flow* algorithms can be used to compare this distorted image with a known reference image in order to obtain the displacement vector field [8, 11, 14].

**Fig. 1** 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,  $B_b$  appears at  $A_i$  and we see a virtual displacement of  $\delta$ .



Note that the amount of distortion is proportional to the distance between the background and the scan volume, which should be maximised in order to ease the optical flow computation. However, to remain in focus, the scan volume should also be positioned as close to the background as possible. We compromise by placing it half-way between the camera and background (which are 5m apart). The camera array consists of 16 synchronised Sony HDR-SR7 camcorders, equipped with 400mm (35mm equivalent) lenses, positioned in an arc of almost  $180^\circ$  around the measurement volume.

Optical flow algorithms perform best when the image contains high frequency texture throughout. A random noise pattern will suffice, but could cause problems when its resolution differs greatly from that of the camera. For example, a noise pattern drawn from a Gaussian distribution will have an intensity histogram like that in the bottom left of Figure 2. Moving the pattern further away from the camera causes large patches of it to be imaged to each pixel in the camera, in effect blurring the whole image. The intensity histogram therefore becomes much narrower (see bottom right of Figure 2), which makes distinguishing between bright and dark spots more difficult, and hence the optical flow performance degrades considerably.

We solve this problem by using a multi-scale noise pattern, such as Wavelet Noise [9]. It is a sum of multiple independent noise functions with non-overlapping frequency spectra. This guarantees that the image will contain sufficient high-contrast detail at any scale. Histograms for the multi-scale noise pattern before and after scaling are shown in Figure 2, demonstrating the difference when compared to Gaussian noise.

We evaluated four different optical flow algorithms using a synthetically-warped noise pattern. For each algorithm we identified the single key parameter that most affects its results, and varied it across an empirically determined range. Figure 3 shows the average vector difference (endpoint) errors for the resulting vector fields, divided into three groups representing low, medium and high amounts of distortion. Within each group, the bars are grouped into four sets of three bars, representing the four algorithms and three of the selected parameter values. The green bars represent a variational algorithm by Brox et. al. [8] which produces good optical flow results in general, but tends to oversmooth this high frequency data. The MatPIV [19] toolbox (red) has been extensively used in earlier BOS work, but we found that significantly better results could be obtained with either the Horn-Schunck [11] (blue) or Lucas-Kanade [14] (brown) algorithms. The latter’s lower sensitivity to parameter choice led to us choosing it for our later experiments.

When performing tomographic reconstruction, the intersection of the view frusta from all the cameras defines the reconstruction volume. However, the refractive medium may only occupy a small region inside this volume, and so we detect the visual hull [13] of the medium and exclude all the empty space around it from the equation system. This both speeds up the solution, and improves on the quality of the results.

To obtain a 3D visual hull we first require a binary mask for each 2D optical flow vector field, classifying pixels as either in- or outside the medium. To see why such a mask cannot be obtained by simply thresholding the vector field, consider the trivial case of a ray propagating parallel to the refractive index gradient. Such a ray will not be deflected at all, yet it should clearly be classified as inside the medium. A solution to this problem is to first integrate the optical flow vector field by solving the associated Poisson equation, and then thresholding the resultant heightfield.

As input to the tomography algorithm we require 3D deflection vectors, but are only able to capture 2D projections of these via optical flow. These can be transformed into 3D worldspace via the camera calibration matrices to obtain  $\mathbf{d}^{(\text{in})}$  and an approximation  $\hat{\mathbf{d}}^{(\text{out})}$  to  $\mathbf{d}^{(\text{out})}$ .

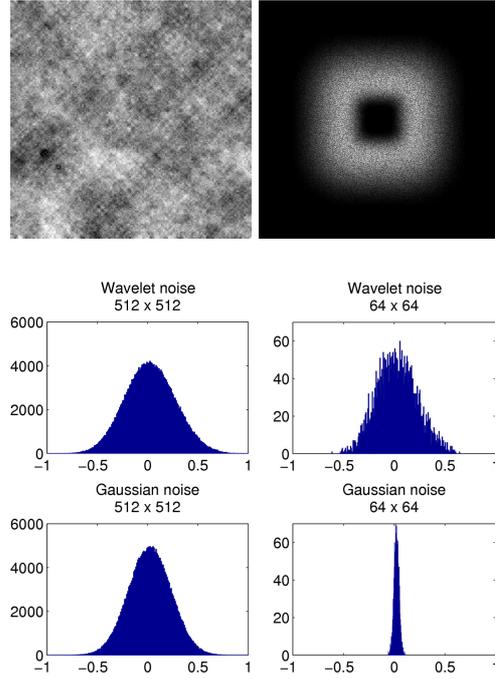
### 3 Tomographic Reconstruction

The 2D displacement vectors that have been computed by the optical flow algorithm are projections of 3D ray deflections, which are related to line integrals of the 3D refractive index gradients. We derive the relation between ray deflections and the refractive index field from the *ray equation of geometric optics*, which describes light propagation in optically inhomogeneous media as

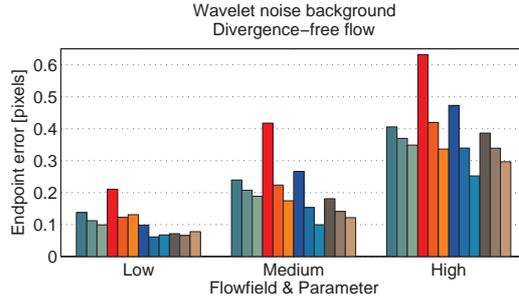
$$\frac{d}{ds} \left( n \frac{d\mathbf{x}}{ds} \right) = \nabla n \quad (1)$$

where  $n$  denotes the refractive index field,  $\mathbf{x}$  is the position of a mass-less particle travelling along the ray, and  $ds$  is the differential path along the curved ray. Re-

**Fig. 2** **Top left:** final wavelet noise image; **Top right:** frequency spectrum of one of the noise bands; **Middle row:** the intensity histogram of a wavelet noise pattern remains Gaussian-shaped, even if the pattern is downsampled; **Bottom row:** if a Gaussian noise pattern is downsampled, its dynamic range decreases considerably. Figures reproduced from [5].



**Fig. 3** The endpoint errors of different optical flow algorithms for a synthetic flow on a wavelet noise pattern. Figure reproduced from [5].



formulating Equation 1 as a first-order ODE system and integrating leads to the following equation, which relates 3D ray deflections to the gradient of the index field:

$$\mathbf{d}^{(\text{out})} - \mathbf{d}^{(\text{in})} = \int_c \nabla n ds. \quad (2)$$

Here  $\mathbf{d}^{(\text{in})}$  denotes the incoming ray direction and  $\mathbf{d}^{(\text{out})}$  denotes the outgoing ray direction with respect to the measurement volume for a ray path  $c$ .

We discretise the vector-valued function  $\nabla n$  using normalised scalar basis functions  $\phi_i$ . Insertion in Equation 2 yields

$$\sum_i \mathbf{n}_i \int_c \phi_i ds = \mathbf{d}^{(\text{out})} - \mathbf{d}^{(\text{in})} \quad (3)$$

where the vector-valued coefficients  $\mathbf{n}_i$  parametrise the components of the unknown gradient. Reformulating Equation 3 as three linear systems of equations in each of the vector components leads to

$$\mathbf{S} \mathbf{n}_{(x,y,z)} = \mathbf{d}^{(\text{out})}_{(x,y,z)} - \mathbf{d}^{(\text{in})}_{(x,y,z)}. \quad (4)$$

Here  $(x,y,z)$  denote the individual linear systems, and  $\mathbf{S}$  is the system matrix having the following structure:

$$\mathbf{S} = \begin{pmatrix} \int_{c_1} \phi_1 ds & \cdots & \int_{c_1} \phi_{n_b} ds \\ \vdots & \ddots & \vdots \\ \int_{c_{n_p}} \phi_1 ds & \cdots & \int_{c_{n_p}} \phi_{n_b} ds \end{pmatrix} \quad (5)$$

where  $n_b$  denotes the number of basis functions  $\phi_i$  and  $n_p$  expresses the total number of deflection measurements in all cameras. We approximate the integrals of the individual matrix entries by Riemann sums over the single basis functions. The curved rays are approximated by straight rays. This is known as paraxial approximation, and is justified by the fact that the deviation between straight and curved rays is small. Simulations of our measurement setup have shown that the difference does not exceed 0.1 mm for measurement volumes of about 20 cm<sup>3</sup>.

We employ localised radially symmetric basis functions to obtain a sparse linear system that can be solved by standard numerical algebra techniques. For the actual computation we only use the basis functions inside the visual hull [13] of the refractive index fields. The visual hull serves as a regularizer on the shape of the volume and minimizes projection artifacts [12].

The tomographic reconstruction results in a three-dimensional gradient field  $\widehat{\nabla n}$  of the refractive index. It is then integrated by solving a Poisson equation. Note however, that the set of gradient vectors is not consistent in general. To deal with gradient inconsistency we resort to anisotropically weighted Poisson integration [3]

$$\nabla \cdot (\mathbf{D} \nabla n) = \nabla \cdot (\mathbf{D} \widehat{\nabla n}) \quad (6)$$

where  $\mathbf{D}$  is a gradient-weighting diffusion tensor. For standard Poisson integration the tensor can be chosen as  $\mathbf{D} = 1$ . In our case we employ a tensor which prefers gradients from similar iso-surfaces of the underlying refractive index field, while weakening the influence of gradients orthogonal to them. For details on the definition and use of diffusion tensors refer to [2, 6, 20].

Equation 6 is again discretised within the visual hull, only. The boundary values are set to the refractive index of air, resulting in a sparse and positive definite linear system which is solved using a Jacobi-preconditioned Conjugate Gradient method [7].

## 4 Continous Refraction Rendering

Once three-dimensional models of gas flows have been obtained by the techniques described above, they can be placed into virtual environments. Their optical parameters can also be changed. For example, it is possible to add scattering and absorbing smoke particles even though smoke has not been present during the capture. To achieve this, we first derive a mathematical description of the underlying image formation model, which is suitable for rendering continuously refracting objects, including advanced optical properties like scattering, emission and absorption. We then discretize and simplify the model to map it efficiently onto modern graphics hardware.

The radiance recorded by a camera is accumulated along a curved ray passing through the pixel and the refracting volume. The image formation process can be described by the following equation:

$$L(c) = \int_c L_c(\mathbf{x}, \mathbf{d}) \alpha(s, c) ds + L_{bg} \alpha(s_\infty, c) \quad (7)$$

where  $L_c$  expresses the radiance at a particular point in space  $\mathbf{x}$ , in a particular direction  $\mathbf{d}$  due to the combined effect of emission and scattering,  $\alpha(s, c) = \exp(-\int_0^s \sigma_t \circ c(t) dt)$  denotes the absorption of light along the ray at a certain distance  $s$  and  $L_{bg}$  denotes background radiance. The radiance  $L_c(\mathbf{x}, \mathbf{d})$  can be written in more detail by summing the out-scattered and emitted radiance:

$$L_c(\mathbf{x}, \mathbf{d}) = \hat{\omega} L_s(\mathbf{x}, \mathbf{d}) + L_e(\mathbf{x}, \mathbf{d}). \quad (8)$$

Here  $L_s$  is the out-scattered and  $L_e$  the locally emitted radiance, while  $\hat{\omega}$  expresses the albedo of the material. Introducing the scattering phase function  $p$ , the scattering term  $L_s$  can be formulated as an integral over the sphere  $\Omega$  of incoming light directions

$$L_s(\mathbf{x}, \mathbf{d}) = \int_\Omega p(\mathbf{x}, \mathbf{d}, \omega) L(\mathbf{x}, \omega) d\omega = \int_\Omega p(\mathbf{x}, \mathbf{d}, \omega) dE_\omega. \quad (9)$$

This general formulation is clearly too involved to be computed quickly on graphics hardware. Therefore we simplify it by assuming that the scene is illuminated by a finite number of light sources, and that each point in the scene receives light by a finite number of incoming rays only. Applying these assumptions to our continuous model,  $L_s(\mathbf{x}, \mathbf{d})$  can be discretised and expressed as a sum over the incoming light rays:

$$L_s(\mathbf{x}, \mathbf{d}) = \sum_j p(\mathbf{x}, \mathbf{d}, \mathbf{l}_j) \Delta E_{\omega_j}. \quad (10)$$

This allows us to precompute the irradiance values for all incoming light directions. Afterwards, Equation 7 can be computed on the fly by performing irradiance lookups into an additional three-dimensional texture.

Applying this discussion to our acquired gas flows, we see that we do not acquire emission, absorption or scattering properties. Thus,  $L_c$  is simply zero. Because

the attenuation coefficient  $\sigma_r$  is zero throughout the volume, the absorption factor  $\alpha(s_\infty, c)$  equals one and we obtain

$$L(c) = \int_c 0 ds + L_{bg} \alpha(s_\infty, c) = L_{bg}. \quad (11)$$

The volume rendering of continuous refraction is a simple lookup of the background radiance. Additional optical properties like emission, absorption and scattering coefficients can be synthetically added to increase the visual appeal of the renderings. Since Equation 7 describes the forward model of volumetric light transport it could also serve as a basis for more advanced reconstruction schemes.

## 5 Results

We evaluated the accuracy of our reconstruction algorithm using simulated data experiments. We started by assessing the numerical error of the integration step, adding complexity to the synthetic experiments until the full system error as well as the errors introduced by the single processing steps could be determined.

Using a static, three-dimensional fuel injection data set [1], we first computed ground truth gradient vectors. By performing anisotropically weighted Poisson integration on the synthetic gradients we established a lower bound on the expected error of 42.15 dB peak signal-to-noise ratio (PSNR), or 0.78% root mean square (RMS) error.

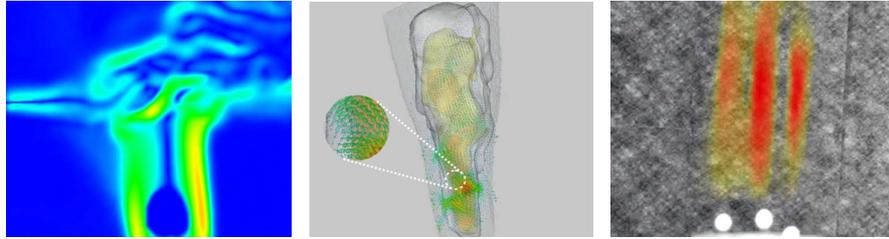
In a second experiment we simulated the 3D deflection vectors  $\mathbf{d}^{(\text{out})} - \mathbf{d}^{(\text{in})}$  by tracing curved rays according to Equation 1. A tomographic reconstruction was then performed on the simulated measurement data. The results of this experiment are shown in the first row of Table 1. In reality, only projections of this 3D deflection vector into the image planes of the cameras can be measured. The influence of this approximation was assessed in a third experiment (second row, Table 1). Finally, we distorted a wavelet background noise pattern using the projected, simulated deflection vectors. We applied optical flow computation to the distorted images to assess the influence of the optical flow on overall reconstruction accuracy (third row, Table 1). As can be seen from the results, optical flow computation introduces the largest error, followed by the approximation of the 3D deflection vector. However, using only 16 cameras, detailed reconstructions with only  $\approx 1\%$  RMS error can be achieved.

We performed a number of measurements on real-world gas flows. Figure 4 (left) shows the resulting deflections from an air jet dispersing a candle plume measured with the BOS technique, and (middle) visualizes the 3D gradients of a flow caused by a gas burner. The gradient vectors were recovered using approximately 150,000 basis functions and 700,000 pixel measurements for each time frame. Another experiment shows that our system is able to clearly separate spatially distinct features. The reconstructed plumes of three tea lights are overlaid onto one of the input images in Figure 4 (right).

Simulation errors for tomographic reconstruction						
Experiment	8 Cameras		16 Cameras		32 Cameras	
	PSNR	RMS	PSNR	RMS	PSNR	RMS
Ground truth $\mathbf{d}^{(out)}$	40.55	0.94%	41.29	0.86%	41.39	0.85%
Approximate $\hat{\mathbf{d}}^{(out)}$	40.01	1.00%	40.69	0.92%	40.76	0.92%
Optical flow	39.29	1.09%	39.84	1.02%	39.88	1.01%

**Table 1** Measured simulation errors for the tomographic reconstruction. The root mean square (RMS) error and the peak signal-to-noise ratio (PSNR) are compared for camera different setups.

We also modified the optical properties of the measured gas flows, adding absorption and scattering properties in regions of high refractive index gradient. In Figure 5, the refractive index volume of the gas burner flow is attached to an oil lamp and a candle data set is placed above a virtual candle. The light rays passing through the volume are distorted and cause the typical deflections of heated air flows. The wall relief is distorted as well as the content of the image on the opposite wall.



**Fig. 4** **Left:** the displacement magnitude of a candle plume's interaction with a jet of compressed air; **Middle:** the 3D refractive index gradients of a flow from a gas burner and their integrated refractive index field; **Right:** the reconstructed plumes of 3 tea lights rendered into one of the input camera's view. Figures reproduced from [6].



**Fig. 5** The reconstructed gas volume is placed into a virtual scene. The volume is attached to an oil lamp and distorts the relief on the wall.

**Acknowledgements** Ivo Ihrke was supported by a Feodor-Lynen Fellowship of the Humboldt Foundation, Germany. Part of this work was supported by the National Sciences and Engineering Research Council of Canada (NSERC), and the German Research Foundation (DFG) under grant MA2555/5-1.

## References

1. (2000) Fuel injection volumetric data set. <http://www.volvis.org>, DFG SFB 382
2. Agrawal A, Chellappa R, Raskar R (2005) An Algebraic Approach to Surface Reconstruction from Gradient Fields. In: Proceedings of ICCV, vol 1, pp 174–181
3. Agrawal A, Raskar R, Chellappa R (2006) What Is the Range of Surface Reconstructions from a Gradient Field? Proceedings of ECCV 3951:578–591
4. Agrawal AK, Albers BW, Griffin DW (1990) Abel Inversion of Deflectometric Measurements in Dynamic Flows. Applied Optics 38(15):3394–3398
5. Atcheson B, Heidrich W, Ihrke I (2008) An Evaluation of Optical Flow Algorithms for Background Oriented Schlieren Imaging. Experiments in Fluids p in print
6. 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 (Proc of SIGGRAPH Asia) 27(5):article 132
7. Barrett R, Berry M, Chan T, Demmel J, Donato J, Dongarra J, Eijkhout V, Pozo R, Romine C, van der Vorst H (1994) Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods. Society for Industrial and Applied Mathematics
8. Brox T, Bruhn A, Papenberg N (2004) High Accuracy Optical Flow Estimation Based on a Theory for Warping. Proceedings of ECCV pp 25–36
9. Cook R, DeRose T (2005) Wavelet Noise. Proceedings of ACM SIGGRAPH 24(3):803–811
10. Dalziel S, Hughes G, Sutherland B (2000) Whole-field Density Measurements by 'Synthetic Schlieren'. Experiments in Fluids 28(4):322–335
11. Horn B, Schunck B (1981) Determining Optical Flow. Artificial Intelligence 17(1-3):185–203
12. Ihrke I, Magnor M (2004) Image-Based Tomographic Reconstruction of Flames. In: ACM Siggraph / Eurographics Symposium Proceedings, Symposium on Computer Animation, pp 367–375
13. Laurentini A (1994) The Visual Hull Concept for Silhouette-based Image Understanding. IEEE Transactions on Pattern Analysis and Machine Intelligence 16(2):150–162
14. Lucas B, Kanade T (1981) An Iterative Image Registration Technique with an Application to Stereo Vision. In: International Joint Conference on Artificial Intelligence, vol 3, pp 674–679
15. Meier G (2002) Computerized Background-Oriented Schlieren. Experiments in Fluids 33(1):181–187
16. Richard H, Raffel M (2001) Principle and Applications of the Background Oriented Schlieren (BOS) Method. Measurement Science and Technology 12(9):1576–1585
17. Schwarz A (1996) Multi-tomographic Flame Analysis With a Schlieren Apparatus. Measurement Science and Technology 7(3):406–13
18. Settles G (2001) Schlieren and Shadowgraph Techniques: Visualizing Phenomena in Transparent Media. Springer
19. Svein JK (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>
20. Weickert J (1998) Anisotropic Diffusion in Image Processing. Teubner