

# Simultaneous Image Registration and Monocular Volumetric Reconstruction of a Fluid Flow

Florent Brunet<sup>1234</sup>

<http://www.brnt.eu>

Emmanuel Cid<sup>123</sup>

[cid@imft.fr](mailto:cid@imft.fr)

Adrien Bartoli<sup>4</sup>

<http://isit.u-clermont1.fr/~ab>

<sup>1</sup> INPT, UPS, IMFT (Institut de Mécanique des Fluides de Toulouse), Université de Toulouse, Allée Camille Soula, F-31400 Toulouse, France

<sup>2</sup> CNRS, IMFT, F-31400 Toulouse, France

<sup>3</sup> Fédération de recherche Fermat FR 3089, Toulouse, France

<sup>4</sup> Clermont Université, Université d'Auvergne, ISIT, BP 10448, F-63000 Clermont-Ferrand, France

---

## Abstract

We propose to combine image registration and volumetric reconstruction from a monocular video of a draining off Hele-Shaw cell filled with water. A Hele-Shaw cell is a tank whose depth is small (*e.g.* 1 mm) compared to the other dimensions (*e.g.*  $400 \times 800 \text{ mm}^2$ ). We use a technique known as molecular tagging which consists in marking by *photobleaching* a pattern in the fluid and then tracking its deformations. The evolution of the pattern is filmed with a camera whose principal axis coincides with the depth of the cell. The velocity of the fluid along this direction is not constant. Consequently, tracking the pattern cannot be achieved with classical methods because what is observed is the integration of the marked molecules over the entire depth of the cell. The proposed approach is built on top of classical direct image registration in which we incorporate a volumetric image formation model. It allows us to accurately measure the motion and the velocity profiles for the entire volume (including the depth of the cell) which is something usually hard to achieve. The results we obtain are consistent with the theoretical hydrodynamic behaviour for this flow which is known as the *laminar Poiseuille flow*.

## 1 Introduction

**General context.** This paper applies computer vision techniques to experiments in fluid mechanics. We propose a new method based on imagery to densely and accurately measure the Laminar Poiseuille Flow (LPF). The LPF is the theoretical model that describes the flow of a liquid in a draining off Hele-Shaw (HS) cell. Figure 1 gives an overview of the experimental setup. We propose a new algorithm combining direct image registration (DIR) and monocular volumetric reconstruction to track a pattern tagged into the liquid at a molecular level. Our




Figure 1: In this paper, we propose a new approach to register 2D images of a draining-off Hele-Shaw (HS) cell while simultaneously recovering the 3D motion of the fluid from a single video. The proposed approach relies on direct image registration (DIR). It uses as few physical assumptions as possible so as to experimentally validate the theoretical flow model. The interest of this work is twofold. In computer vision, it is a new algorithm for monocular volumetric reconstruction. In fluid mechanics, it allows us to experimentally measure the Laminar Poiseuille Flow (LPF) which is the expected theoretical model for this flow.

approach allows us to make measurements with a high spatial and temporal resolutions (up to 250 volumes per second with  $200 \times 200 \times 40$  voxels). Previous experimental measurements of this type of flow has principally been obtained for small regions of interest (around  $1 \text{ mm}^3$  using *e.g.* Micro Particle Image Velocimetry ( $\mu\text{PIV}$ ) [1]) or without taking into account the motion in the depth of the cell (using *e.g.* classical PIV [2]). Besides, our approach allows us to measure the development of the LPF whereas previous techniques often consider the flow when it has reached its steady state.

**Contribution.** Obtaining such measurements raises several strong computer vision issues. Indeed, our goal is to not use the theoretical flow model as a prior, but rather to obtain the flow using more general and weaker hypotheses, so as to be able to validate the theoretical flow model *a posteriori*. We thus meet the problem of modelling the 3D motion of a liquid in a HS cell and of making a volumetric reconstruction of an homogeneous moving liquid from a single point of view. We propose a set of simple but physically realistic hypotheses. In particular, we consider that the flow is *lamellar* which means that the liquid acts as if it were made of an infinite number of rigid layers of infinitesimal depth. These layers are parallel to the widest surface and move independently from each other. We approximate this model with a finite number of layers. We also make assumptions on the layers motion: the motion of the liquid is symmetric with respect to the centre of the cell, the motion is faster at the centre of the cell than near the cell surfaces, the motion of the layers can be described as downward translations and layers are not allowed to go up. We finally consider a volumetric image formation model based on the fact that pixel intensities are an integral measurement of the concentration of tracer in the corresponding volume of the HS cell. With all these elements, we propose a variant of DIR [3] adapted to the case of the superposed layers. The beauty of our approach is that even though it solves a complex fluid flow estimation problem, it only requires one to solve a single constrained least-squares optimization problem to obtain the 3D motion of the fluid for all the video sequence.

Note also that as a side contribution, we propose an image interpolation algorithm based on regularized splines after having shown that bilinear and bicubic interpolation introduce a systematic bias. This is important for the requirement of measurements' accuracy.

**Related work.** 3D reconstruction of a transparent volume have been extensively reviewed in [9]. Although many of the methods in computer vision for 3D reconstruction are dedicated to opaque objects, several methods have been proposed for volumetric phenomenon.

Tomography [6] is one of the most obvious approach. Similarly to our case, the intensity of a pixel correspond to an integral measurement of the volume over the corresponding line of sight. Our approach differs from tomography in the sense that we only use a single camera (mainly for practical reasons among which the problems of light source(s) reflection with the glass surface of the HS cell and the fact that high speed camera are expensive). Tomography from a single view point has been investigated for reconstructing astronomical objects such as planetary objects [10]. As in our case, [10] partly relies on the symmetry of the objects to reconstruct but, unlike us, they only reconstruct static scenes.

These previous approaches are formulated as inverse problems (with all the drawbacks of ill-conditioned problems due to noise or missing data). It is possible to make direct measurements with a proper experimental setup. For instance, [6] reconstructs clouds of smoke using a laser sheet which is swept across the volume. The laser sheet is filmed from an orthogonal point of view and the observed pixels are linked to the concentration of smoke. This kind of approach is not possible in our experiment because what we aim at observing is the inner motion of a completely homogeneous liquid without addition of tracer particles. Beside the temporal resolution of such a system would not be sufficient enough for our needs.

Similarly to our contribution, the authors of [9] propose a variant of classical optical flow estimation from images that incorporates a volumetric model of a gas in a microchannel. However, in [9], the LPF model is used as a strong prior.

**Notation.** Scalars are denoted in italics (*e.g.*  $x$ ), matrices in sans-serif (*e.g.*  $M$ ), images in capital letters (*e.g.*  $I$ ), and vectors in bold (*e.g.*  $\mathbf{p}$ ). The vectors are row vectors.

**Paper organization.** We describe the experimental setup and molecular tagging based on photobleaching in §2. Our contribution is explained in §3. The results are presented in §4. In §5, we conclude and discuss the possible implications of our work in fluid mechanics.

## 2 Experimental Setup

**The Hele-Shaw (HS) cell.** The HS cell we use for this paper is an almost bi-dimensional cell constituted of two parallel plates of glass ( $400 \times 800 \text{ mm}^2$ ) separated by a 1 mm wide gap. A picture of this cell is given in figure 2. The cell is filled with water and a downward translation is generated by draining off the cell with a valve located at the bottom middle of the cell. The motion of the liquid is observed from a fronto-parallel point of view, *i.e.* the principal axis of the camera is orthogonal to the surface of the cell. The acquisition is done with a fast camera (Photron RS3000) at a resolution of  $1024 \times 1024$  pixels with an effective intensity depth of 10 bits per pixel. The frequency of the acquisitions is limited by the sensitivity of the camera sensor compared to the luminous intensity of the observed phenomenon. Given the light source we use (see next paragraph), we were able to capture 250 frames per seconds. This is sufficient for the phenomenon we want to observe since the velocity of the fastest part of the draining liquid is expected to be less than 1 pixel per frame.

**Molecular tagging based on photobleaching.** In order to observe the fluid motion, we tag the liquid at a molecular level with a technique known as *photobleaching* [8]. It has the advantage of being truly non-invasive contrarily to other techniques such as PIV where

particles are added. Such particles may have a proper motion independent of the fluid and, moreover, they may adhere to the wall of the cell. We now give the general principle of photobleaching. Figure 2 illustrates the setup. A tracer, the *fluorescein*, is uniformly mixed to the water (at a concentration of  $10^{-4}$  g.L $^{-1}$ ). When the molecules of fluorescein are excited with an electromagnetic radiation of given wavelength, it becomes fluorescent *i.e.* it emits photons (which have a wavelength longer than the one used to excite the fluorescein which allows us to filter out the light coming from the fluorescein itself from the light of the green laser). We use a laser beam with wavelength of  $\lambda_1 = 527$  nm ('green laser') to excite the fluorescein. A spherical lens is placed in front of this laser to make the beam divergent so that an entire region of the cell is excited. A pattern is 'printed' by *photobleaching* in the liquid when it is perfectly still. For this purpose, we use another laser beam with wavelength  $\lambda_2 = 488$  nm ('blue laser') which has the property of inhibiting the fluorescence effect of the tracer when high level of energy are used. A cylindrical divergent lens is placed in front of the blue laser to obtain a beam shaped like a vertical layer. An opaque mask is positioned between this beam and the cell to create the pattern. Since a high level of energy is required to 'bleach' the fluorescein, it was not possible to simply use a divergent spherical lens to get a conic beam and then mark an entire two-dimensional pattern. Instead, we concentrated the energy of the blue laser on a single layer and, then, moved this layer horizontally with an automated arm to obtain a two-dimensional pattern in the liquid. Each position has to be illuminated by 1W of blue laser light during 15 seconds in order to bleach the liquid.




Figure 2: (a) Scheme of our setup (top view). (b) Front view of the set-up.

**Image formation model.** The intensity of the pixels we observe is linked to the concentration of non-bleached fluorescein. The important thing here is that the intensity of a pixel corresponds exactly to this concentration integrated on the small volume of the cell that corresponds to the pixel. This is illustrated in figure 3. In the sequel of this paper, we neglect the perspective effects and approximate the camera model by an affine projection. This makes sense because the depth of the cell is small compared to the distance from the camera to the cell and because the bleached pattern is located at the centre of the acquired images.

### 3 Simultaneous Registration and Volumetric Reconstruction

We start this section with a short reminder on DIR since this technique is used as a foundation for our contribution. At the same time, we illustrate why classical image registration is not usable for tracking the fluid motion. We then detail the core of our contribution *i.e.*




Figure 3: (a) The intensity value of a pixel is proportional to the quantity of non-inhibited fluorescein in the volume of the cell corresponding to that pixel. (b) Example of acquired image: the tagged pattern is the set of aligned dark marks at the centre of the image.

the algorithm able to simultaneously track the moving liquid and reconstruct the observed volume. Finally, we briefly mention a point of detail which is nonetheless important to get accurate results: image interpolation.

### 3.1 Reminder on Direct Image Registration (DIR)

Image registration is the problem of finding the geometric transformation between two images. In DIR [13], this transformation is determined directly from the intensities of the pixels. Let  $R$  be a reference image (*i.e.* in our case, an image of the bleached pattern before any motion) and let  $I_i$  be the  $i$ th image of the video sequence (for  $i = 1, \dots, m$ ). If we naively consider that the motion is a simple global downward translation, then aligning an input image  $I_i$  to the reference  $R$  with DIR is formulated as:

$$\min_t \sum_{\mathbf{q} \in \Omega} (I_i(\mathbf{q}) - R(\mathbf{q} - (0, t)))^2, \quad (1)$$

where  $\Omega$  is the *region of interest*, *i.e.* a portion of the input image  $I_i$  that contains the whole bleached pattern. Equation (1) is a non linear least-squares optimization problem easily solved using, for instance, a Gauß-Newton-like algorithm [14].

The formulation of (1) is valid only if the Brightness Constancy Assumption (BCA) is satisfied *i.e.* whenever the colours of corresponding pixels are identical. In our case, the intensity of a pixel is the integral of the corresponding volume of the cell. Since the motion of the fluid along the  $z$ -direction is not uniform, there is no direct correspondence between the pixel of the reference image and of an image of the input sequence. Consequently, the motion of the marked pattern cannot be explained with a simple translation (see figure 4).





Figure 4: In our experiment, a simple tracking of the bleached pattern is not possible because the brightness constancy assumption is not satisfied. This stems from the fact that the variation of the motion of the liquid across the depth of the cell is not negligible.

Figure 5: In our model, the intensity of a pixel is obtained by integrating the fluorescein concentration carried by the layers. A downward translation (red arrows) is associated to each one of the layers for each one of the images of the video sequence. These translations are the unknown of our problem. They are determined using our algorithm built upon DIR.



## 3.2 The proposed approach

**Motion model and image formation model.** The idea of our approach lies in the model we use for the motion and for the formation of the images. As explained before, the flow we consider is lamellar. We thus consider that the volume is divided into  $n$  layers parallel to the surface of the cell. The concentration of unbleached fluorescein contained in each one of these layers is given by dividing the reference image by the number of layers (we consider that before any motion appears the concentration of bleached and unbleached fluorescein is uniform along the  $z$ -direction). Let us note  $L_j$  the  $j$ th layer ( $j = 1, \dots, n$ ). We have that:

$$L_j(\mathbf{q}) = \frac{1}{n}R(\mathbf{q}) \quad \text{and} \quad R(\mathbf{q}) = \sum_{j=1}^n L_j(\mathbf{q}) \quad (2)$$

Still considering the lamellar hypothesis for the flow, the above defined layers will slide independently from each other with downward translations (see figure 5). A translation parameter  $t_{i,j}$  is associated to the  $j$ th layer for the  $i$ th input image of the video sequence ( $i = 1, \dots, m$ ). Given this model of the volume, the image formation model is given by integrating the concentrations carried by all the translated layers and we have:

$$I_i(\mathbf{q}) \approx \sum_{j=1}^n L_j(\mathbf{q} - (0, t_{i,j})). \quad (3)$$

**Additional hypotheses.** It is possible to make a few additional assumptions on the translations. These assumptions are quite simple and mainly based on the geometry of the experimental setup and on the fact that the only applying force is gravitational.

- **Symmetry.** The flow is symmetric with respect to the centre of the cell in the  $z$ -direction. This constraint is naturally enforced by modelling only half the cell (by convention,  $t_{i,n}$  is the innermost layer while  $t_{i,1}$  is the one closest to the wall).
- **Positivity.** All the translations are downward translations, *i.e.*  $t_{i,j} \geq 0$ .
- **Temporal consistency.** A layer never goes upward, *i.e.*,  $t_{i+1,j} \geq t_{i,j}, \forall i = 1, \dots, m-1, \forall j = 1, \dots, n$ .
- **Spatial consistency.** The inner layers are faster than the layers close to the surface of the cell, *i.e.*  $t_{i,j+1} \geq t_{i,j}, \forall i = 1, \dots, m, \forall j = 1, \dots, n-1$ .

**Final optimization problem.** Our approach is formulated as an optimization problem derived from the basic formulation (1) of DIR. The main difference is that the image formation model based on the integration of the volumetric layers replaces the reference image in the

initial formulation. The second difference is that inequality constraints are enforced on the parameters to find. Another difference is that we align the whole sequence at once in order to take into account the temporal consistency hypothesis. The final optimization problem is:

$$\min_{\mathbf{T}} \sum_{i=1}^m \sum_{\mathbf{q} \in \Omega} \left( I_i(\mathbf{q}) - \sum_{j=1}^n L_j(\mathbf{q} - (0, t_{i,j})) \right)^2$$

$$\text{subject to } \begin{cases} t_{i,j} \geq 0 & \forall i = 1, \dots, m, \forall j = 1, \dots, n \\ t_{i+1,j} \geq t_{i,j} & \forall i = 1, \dots, m-1, \forall j = 1, \dots, n \\ t_{i,j+1} \geq t_{i,j} & \forall i = 1, \dots, m, \forall j = 1, \dots, n-1 \end{cases} \quad (4)$$

where  $\mathbf{T}$  is an  $m \times n$  matrix containing all the translation parameters to determine. Problem (4) may be solved using for instance sequential quadratic programming. In practice, we used the `fmincon` function of `Matlab`.

### 3.3 Image Interpolation

In equation (4) the input images have to be evaluated at non integer locations. Besides, we know that the order of magnitude of the translation of the central layer between two successive images is about one pixel. Consequently, sub-pixel evaluation of the images must be performed with extreme care. Common algorithms such as bilinear or bicubic interpolation may introduce a bias. This is particularly true when the non-integer locations for the interpolation all have the same fractional part (which is the case for a simple vertical translation as we do in this paper). A simple experiment reported on figure 6 illustrates this phenomenon. For that, we take two successive images  $I_a$  and  $I_b$  at the beginning of a video sequence we acquired before any fluid's motion. Then, we look at the intensity difference between the first image  $I_a$  and the second image  $I_b$  translated along the  $y$ -axis by  $\Delta y$  pixels. Since  $I_a$  and  $I_b$  are identical (except for the noise), this difference should be minimal for  $\Delta y = 0$  pixel. Figure 6 tells us that with bilinear or bicubic interpolation,  $\Delta y = 0$  correspond to a local maximum and that two local minima appear on each side of  $\Delta y = 0$  for approximately  $\Delta y = \pm \frac{1}{3}$ .




Figure 6: The classical bilinear and bicubic interpolation algorithms introduce a bias in the cost function of DIR. Correct results are obtained with our approach based on regularized B-splines. Similar results have been obtained for hundreds of pairs of other images.

To solve this problem, we convert the images as continuous functions during a preprocessing step. For that purpose we fit a regularized tensor product B-spline [13] to the intensities of the images. The control points  $\mathbf{p}$  of the B-spline surface  $B$  approximating the image  $I$  are obtained by solving the following linear least-squares optimization problem [14]:


$$\min_{\mathbf{p}} \sum_{\mathbf{q} \in \Omega} (B(\mathbf{q}; \mathbf{p}) - I(\mathbf{q}))^2 + \lambda \iint_{\Omega} \left( \sum_{d=0}^2 \frac{\partial^2 B}{\partial x^d \partial y^{2-d}}(\mathbf{q}) \right)^2 d\mathbf{q}, \quad (5)$$

where the first term measures the closeness of the surface to the image data and the second term penalizes the bending energy of the fitted surface. The parameter  $\lambda$  controls the trade-off between these two aspects. In practice, it is empirically set to small value (*e.g.*  $10^{-6}$ ). If we do the same experiment as the one described previously and replace the interpolation by a simple evaluation of the B-spline, we see that the bias is drastically reduced (see figure 6).

## 4 Results

**Data.** We managed to do the acquisition of two video sequences with different patterns: dots aligned on an horizontal row and on a grid (see figure 7). The acquisition frequency was of 50 Hz but with an exposure time of  $\frac{1}{300}$  second. Classical background subtraction and normalization techniques have been applied to the raw data in order to correct the impurities of the HS cell and to compensate the variations of the illumination source.


Figure 7: Images extracted from the video sequences #1 (left) and #2 (right). These are images after correction (in particular, they are negative version of the original ones).



**Expected results for a LPF.** The underlying theory of the LPF predicts that in steady state, the velocity profile across the cell is parabolic. Between the start of the draining off and the steady state, there exists a development phase during which the velocity profile smoothly evolves from a flat line (*i.e.* same velocity for all the points between the surfaces of the cell) to the parabola of the steady state (see, for instance, [2, 2]).

**Checking our generative model.** We visually check our model by comparing the acquired data to the initial pattern warped using the translations computed with our approach. Figure 8 shows an example where the two images match which means that our approach is valid. Similar results were obtained for the other images of the video.

Figure 8: Comparison of the initial pattern warped using the translations computed with our approach and of the corresponding true image of the video.



**Computed translations and velocity profiles.** The computed translations and velocity profiles across the cell for different times are shown in figure 9. We can see that the computed motion is compatible with the theoretical LPF. During the initial steps of the LPF development phase, our approach seems to compute a velocity profile which is not completely flat as predicted by the theory. This is attributable to another law of fluid mechanics: the boundary layer phenomenon which states that the motion of a liquid at the immediate vicinity of a bounding surface is null. A comparison of the velocities computed with our approach to ground truth values is done in figure 10. A synthetic 3D representation of the

inner part of the cell is shown on figure 11.




Figure 9: Computed translations at different times (left) and corresponding velocities (right). Images #8 and #5 are taken when the LPF is developing whereas images #115 and #37 correspond to an established LPF. When the steady flow is established, a polynomial of degree 2 is fitted to the computed data: it shows that our results and the theoretical LPF are consistent. The difference of velocity profiles between videos #1 and #2 is most likely due to different temperature conditions (which impact the boundary layer phenomenon).




Figure 10: Velocities computed with our approach (averaged over the depth) compared to the ‘ground truth’ average velocity of the liquid (determined from the water/air interface). These results are obtained from video #1.




Figure 11: 3D representation of the reconstructed volume for video #1. The parabolic LPF is visible on the edge and its position is consistent with what is seen in front of the cell.

## 5 Discussion and perspectives

We have proposed a new algorithm that can track the apparent motion of a liquid in a draining off HS cell while simultaneously computing a full volumetric reconstruction of the inner part of the cell from a single point of view. We built our approach on top of DIR to which we added supplementary hypotheses to properly constrain the reconstruction problem. However, we have taken special care to use general enough assumptions on the physical flow model so as to make significant physical measurements. This allowed us to experimentally measure the LPF which is a new result in fluid mechanics for cell with this type of dimensions.

The technique we proposed may be useful to experimentally study other types of 3D motions in HS cells with a large field of view. This is an extremely promising perspective in fluid mechanics. Indeed, the principle of the superposed layers is quite generic and so is the volumetric image formation model we proposed. Our approach may be re-used by simply changing the deformation model (for instance by replacing the downward translation by local Euclidean transformations or even more complex deformation models). In particular, we are interested to measure the Lagrangian motion around a rising bubble in an HS cell. The proposed technique may also be useful to study the 3D effects of the well-known Saffman-Taylor instability in HS cell with a smaller gap than the one considered in this paper.

**Acknowledgement.** We would like to thank Emmanuella Bouche, Sébastien Cazin, Frédéric Risso, and Véronique Roig for their help on technical and fluid mechanics matters.

## References

- [1] Å. Björck. *Numerical Methods for Least Squares Problems*. 1996.
- [2] T. Caplow. URL [http://www.columbia.edu/cu/gsappp/BT/RESEARCH/Arch-atmos/u\\_double.gif](http://www.columbia.edu/cu/gsappp/BT/RESEARCH/Arch-atmos/u_double.gif).
- [3] P. Dierckx. *Curve and Surface Fitting with Splines*. Monographs on Numerical Analysis. Oxford University Press, 1993.
- [4] C. Garbe, K. Roetmann, V. Beushausen, and B. Jahne. An optical flow MTV based technique for measuring microfluidic flow in the presence of diffusion and Taylor dispersion. *Experiments in Fluids*, 44:439–450, 2008.
- [5] T. Hawkins, P. Einarsson, and P. Debevec. Acquisition of time-varying participating media. In *SIGGRAPH*, 2005.
- [6] G. T. Herman. *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*. Springer, 2nd edition, 2009. ISBN 185233617X, 9781852336172.
- [7] Ahmed Hifdi and Jaâfar Khalid Naciri. Étude analytique de perturbations de l'écoulement de Poiseuille dans un canal. *Comptes Rendus Mécanique*, 332(3):241–248, 2004.
- [8] S. Hosokawa and A. Tomiyama. Molecular tagging velocimetry based on photobleaching reaction and its application to flows around single fluid particles. *Multiphase Science and Technology*, 16:335–353, 2004.
- [9] I. Ihrke, K. Kutulakos, H. Lensch, M. Magnor, and W. Heidrich. State of the art in transparent and specular object reconstruction. In *STAR Proceedings of Eurographics*, pages 87–108, 2008.
- [10] M. Magnor, G. Kindlmann, C. Hansen, and N. Duric. Constrained inverse volume rendering for planetary nebulae. In *IEEE Visualization (Vis'04)*, pages 83–90, 2004.
- [11] M. Roudet, A. M. Billet, V. Roig, F. Risso, S. Cazin, and E. Cid. PIV method with volumetric lighting for measurements in narrow channel: Application to the flow around a rising bubble. In *International Symposium on Flow Visualization*, 2008.

- 
- [12] D. Sinton. Microscale flow visualization. *Microfluidics and Nanofluidics*, 1:2–21, 2004. ISSN 1613-4982.
- [13] R. Szeliski. Image alignment and stitching: A tutorial. *Foundations and Trends in Computer Graphics and Vision*, 2:1–104, 2006.