Bayesian analysis for image-based motion estimation with application to environmental flows

Dónal Harkin[1], Naratip Santitissadeekorn[1], Christof Meile[2], Erik Bollt[3], and George Waldbusser[4]

[1]Department of Mathematics, University of Surrey, Guildford, UK

[2]Department of Marine Sciences, The University of Georgia, Athens, GA, USA

[3]Department of Mathematics and Computer Science, Clarkson University, Potsdam, NY, USA

[4]College of Earth, Ocean, and Atmospheric Sciences, Oregon State University, Corvallis, OR, USA

April 3, 2018

Abstract

We develop a Bayesian algorithm to determine the unobserved velocity field that governs the evolution of a scalar in a fluid, which is observed as a time series of images with pixel intensity reflecting tracer concentration. The mathematical model reflects the mass conservation equation for the dissolved tracer, which is simplified in our application to the so-called optical flow equation. The Bayesian algorithm provides an uncertainty estimate that measures the reliability of the estimated flow field, and a correlation structure between image pixels that gives more spatial information about the underlying dynamics. We use Gaussian error distributions for uncertainty estimation since they allow tractable computation in high dimensions, which tends to be the case for most real-world images. In synthetic data experiments, where the "true" velocity field is constructed, we investigate the Bayesian algorithm versus a Horn-Schunck optical flow method. We also estimate the flow for a real experimental image sequence in a marine surficial sediment, where flow is caused by the pumping activity of burrowing macrofauna, and analyse how the variance and divergence change over time.

1 Introduction

Quantification of flow magnitude and direction from repeat observations, such as image time series, have received considerable attention in many scientific fields [1–6]. Many image-based velocity approximations are based on the optical flow framework [3, 4, 6, 7], which is widely used to estimate the planar displacement vector: this represents the apparent motion at all pixels and maps one image to the next. Brightness constancy is assumed within this framework [5, 6, 8–10], which is for planar flow equivalent to a divergence-free assumption. The optimal estimation is achieved by solving the variational cost function associated with a regularisation constraint, which should be appropriately chosen for the specific application [6]. The Horn-Schunck regularisation term imposes a global smoothness constraint [8] and mitigates the "aperture" problem [11, 12]. The Euler-Lagrange equations associated with this cost function can be explicitly derived, providing a coupled PDE to be iteratively solved for the displacement vector (e.g. by the iterative Gauss-Seidel or successive over-relaxation solvers) [3, 5, 13]. There is a large body of research on motion estimation from images, and video analysis, based on the Horn-Schunck framework; a comprehensive survey of this is given in [11, 14, 15].

For a realistic flow model with "complicated" boundary conditions, the Euler-Lagrange equations and their solutions can be difficult to obtain: a computationally expensive direct optimisation of the cost function may have to be carried out instead [16]. Moreover, a smoothness constraint may not be suitable in many fluid applications where the divergence and vorticity fields become important (see [17] for a detailed discussion). These issues limit use of the optical flow framework and, in particular, the Horn-Schunck regularisation constraint in some complex fluid applications. In addition, many variational-based optical flow algorithms do not provide a systematic way to measure reliability of a recovered velocity field. This motivates a Bayesian approach similar to [18, 19], where the velocity field is probabilistically estimated by combining a model "prediction" simulated by a given fluid model with an actual observation (i.e. image sequences) [3, 15, 20]. It avoids the difficult task of deriving the Euler-Lagrange equations. More importantly, the Bayesian framework also provides an uncertainty quantification of the estimated flow field that can give an insight into the inherent uncertainty of the optical flow algorithm [2, 3].

We develop a Bayesian algorithm in this communication to identify fluid flow in a porous medium, assuming Gaussian error distributions to allow for tractable computation in high dimensions. Specifically, we quantify movement of porewater in marine sediments. The activity of macrofauna who live in the top 0.2 m of marine sediments is critically important for exchange of solutes between the water column and sediment [21, 22]. But quantification of biologically induced solute transport, referred to as bioirrigation, can be challenging. Bioirrigation rates can be: estimated from flow velocity microelectrodes, particle image velocimetry [23], conductive exchange experiments [24], other sophisticated imaging techniques [24, 25]; or derived from more commonly available vertical solute concentration profiles [26].

A thin aquarium (0.22 m height × 0.445 m width × 0.022 m depth) was set up to study the effect of macrofauna on porewater flow. The aquarium was filled with approximately 0.16 m of sediment overlain by seawater. A lugworm, *Arenicola marina*, was placed in this aquarium and allowed to acclimate. Lugworms are a few centimetres long, and are subsurface deposit-feeders who pump water through their burrows into the sediment (for example, see [27]). The biological flow induced by the lugworm's activity was visualised by adding fluorescein, a dissolved tracer, through a thin pipette to a location in the sediment close to where the burrowing animal resided. It is assumed that the flow induced by the animal is two-dimensional because the aquarium is thin; this flow is studied using a time series of images depicting spatio-temporal evolution of the tracer field, with intensity reflecting tracer concentration.

The Bayesian algorithm is developed in Section 2. Its applications are illustrated in Section 3 with synthetic data experiments, where the "true" solutions are known. The method is then used to approximate the velocity field for the experimental image sequence introduced in the previous paragraph.

2 Probabilistic model of optical flow

The temporal evolution of tracer concentration $C(x, y, t)$, under the action of diffusion with coefficient $D \in \mathbb{R}$ and advection driven by a flow velocity $v = [v_x, v_y] \in \mathbb{R}^2$, is governed by the Fokker-Planck equation:

$$(\phi C)_t = \nabla. (\phi D \nabla C - \phi v C) \,. \tag{2.1}$$

We simplify (2.1) and express it in the optical flow framework, assuming that porosity $\phi \in \mathbb{R}$ is spatio-temporally constant:

$$- f_t = \nabla f \cdot v, \tag{2.2}$$

where $f(x, y, t)$ is image intensity. The diffusion term in (2.1) is negligible when D is "sufficiently small", i.e. the Péclet number is greater than one; and the optical flow is divergence-free, so image intensity is "rigidly translated" in time. For each pixel, denoted by s, equation (2.2) locally imposes a linear constraint on the optical flow. The following discrete equation is written for an image with N pixels:

$$y = Hv, \tag{2.3}$$

where $y = -f_t \in \mathbb{R}^N$ and $v \in \mathbb{R}^{2N}$. The observation operator is defined as $H = [f_x, f_y] \in \mathbb{R}^{N \times 2N}$, where $f_x,\ f_y \in \mathbb{R}^{N\times N}$ are diagonal matrices. Additively perturbing (2.3) by a stochastic term gives:

$$y = Hv + \epsilon,$$

where $\epsilon \in \mathbb{R}^N$ has a Gaussian distribution $N(0, R)$ and R is a symmetric, positive semi-definite matrix. The likelihood function of v is then given by:

$$P(y|v) = N(Hv, R). \tag{2.4}$$

The maximum likelihood estimate (MLE) $\hat{v}$ is obtained by optimising (2.4); it is the same as the optimal least-squares solution:

$$\hat{v} = (H'R^{-1}H)^{-1}H'R^{-1}y. \tag{2.5}$$

Note that the prime symbol indicates transposition of a matrix or vector. The MLE is the normal flow solution $\hat{v}_n$ when $R = I_N$, i.e. the optical flow is given by (2.6) at each pixel:

$$\hat{v}_n = \left(-\frac{f_x f_t}{\|f_s\|^2}, -\frac{f_y f_t}{\|f_s\|^2}\right) \tag{2.6}$$

where $\|f_s\|^2 = f_x^2 + f_y^2$. So the MLE can be thought of as a "generalised normal" flow, where the residual $y - H\hat{v}$ is measured with respect to the norm $\|\cdot\|_{R^{-1}}$ instead of $\|\cdot\|_2$. Also note that (2.5) is a linear transformation of $y \sim N(Hv, R)$:

$$\hat{v} \sim N(v, (H'R^{-1}H)^{-1}),$$

and $\hat{v}$ is a consistent estimator of v.

The observation y and the optical flow v are both considered random variables in the Bayesian paradigm. This means that prior knowledge of the uncertainty in v is summarised in a prior probability density $P(v)$, which is taken here to follow a Gaussian distribution $v \sim N(\mu_0, P_0)$. The posterior density $P(v|y)$ follows a Gaussian distribution $N(\mu, P)$:

$$\begin{aligned} \mu &= (P_0^{-1} + H'R^{-1}H)^{-1}[(H'R^{-1}H)\hat{v} + P_0^{-1}\mu_0] \\ P &= [P_0^{-1} + H'R^{-1}H]^{-1}. \end{aligned} \tag{2.7}$$

The formulation in (2.7) confirms that the posterior mean is a linear combination of the prior mean and the generalised normal flow. The "Kalman update" form in (2.8), where the Kalman gain is

$K = P_0H'(HP_0H' + R)^{-1}$, is less computationally expensive because inversion of the $2N \times 2N$ matrix P_0 is avoided:

$$\begin{aligned} \mu &= \mu_0 + K(y - H\mu_0) \\ P &= (I_{2N} - KH)P_0. \end{aligned} \tag{2.8}$$

2.1 Specification of prior and observation error covariance matrices

Specification of prior and observation error covariance matrices is a key aspect of the Bayesian paradigm. They could be chosen to promote smoothness of the solution and ameliorate the problem of rank deficiency in matrix inversion.

The gradient of v could be minimised by choosing $P_0 = L$, where L is a numerical approximation of the spatial differentiation operator. Note that L must be symmetric and positive semi-definite for it to have a statistical interpretation. The matrix P_0 could alternatively be based on a function allowing specification of different correlation length scales in the x- and y-directions, for example the fifth-order piecewise rational function by Gaspari and Cohn (1999, Appendix A) [28].

The choice of R depends on our knowledge of how errors from several sources enter H. Assume that they enter (2.2) as follows [18]:

$$y = H(v + \epsilon_1) + \epsilon_2.$$

The random vectors $\epsilon_1 \in \mathbb{R}^{2N}$ and $\epsilon_2 \in \mathbb{R}^N$ are assumed to be independent and follow Gaussian distributions $N(0, \Gamma_1)$ and $N(0, \Gamma_2)$. Although ϵ_1 represents error in H, it does not do this solely: error in y is quantified by ϵ_2, which can be induced by the terms neglected in (2.1) besides instrumentation. These error terms are combined as follows:

$$y = Hv + \epsilon,$$

where $\epsilon \in \mathbb{R}^N$ follows a Gaussian distribution $N(0, R)$ with $R = H\Gamma_1H' + \Gamma_2$. The structure of Γ_1 and Γ_2 is usually unknown, so take $\Gamma_1 = \sigma_1^2 I_{2N}$ and $\Gamma_2 = \sigma_2^2 I_N$ for simplicity:

$$R = \sigma_1^2 HH' + \sigma_2^2 I_N. \tag{2.9}$$

Note that R is diagonal in (2.9), with diagonal entries equal to $\sigma_1^2\|f_s\|^2 + \sigma_2^2$: the generalised normal flow solution in (2.5) is the same as the standard normal flow solution in (2.6). This holds for any choice of non-singular, diagonal R due to the block diagonal structure of H.

Example 1: Let $\mu_0 = 0$ and $P_0 = \alpha I_{2N}$ ($\alpha > 0$). The optical flow at a pixel s is:

$$v_x = -\frac{\alpha f_x f_t}{r}$$
$$v_y = -\frac{\alpha f_y f_t}{r},$$

where $r = ((\alpha + \sigma_1^2)\|f_s\|^2 + \sigma_2^2)$, using (2.8). It has the same orientation as the normal flow, whilst the σ_2^2 term operates as a "regularisation" to prevent a blown-up solution in a zero-contrast region. Note that using too large a value of σ_2^2 can make posterior variance reduction (PVR) impotent. The posterior variances and covariances at a pixel s are given by:

$$Var(v_x(s)) = \alpha\left(1 - \frac{\alpha f_x^2}{r}\right)$$
$$Var(v_y(s)) = \alpha\left(1 - \frac{\alpha f_y^2}{r}\right)$$
$$Cov(v_x(s), v_y(s)) = -\frac{\alpha^2 f_x f_y}{r}$$
$$Cov(v_x(s), v_x(s')) = Cov(v_y(s), v_y(s')) = Cov(v_x(s), v_y(s')) = 0 \qquad s \neq s'.$$

The posterior variance is always less than the prior variance, except for at pixels with zero contrast where the variance is unchanged. For a high contrast pixel where the first term in r dominates the second, the minimum variance is approximately $\frac{\alpha^2}{\alpha+\sigma_1^2}$.

The choice of P_0 and R can be viewed more generally as a use of $\|\cdot\|_{P_0^{-1}}$ and $\|\cdot\|_{R^{-1}}$ for penalising the data-fidelity and constraint (or regularisation) terms, respectively. So P_0^{-1} and R^{-1} are weight matrices for the above weighted norms. However, using the class of symmetric and positive semi-definite matrices for both P_0 and R is the only way to guarantee the scale-invariant property of the solution.

2.2 Choosing σ_1^2 and σ_2^2

The variances σ_1^2 and σ_2^2 are chosen by maximising the marginal likelihood of the observation y:

$$P(y) = \int P(y|v)P(v)dv = N(H\mu_0, Q), \tag{2.10}$$

where $Q = R + HP_0H'$. The dependence on σ_1^2 and σ_2^2 is implicitly presented in R. Note that $P(v|y)$ may be used instead of $P(v)$ in the above calculation, but this would favour values of σ_1^2 and σ_2^2 that overfit the data.

Maximising (2.10) is equivalent to minimising the log likelihood:

$$E := -2 \log P(y) = \log |Q| + \delta z' Q^{-1} \delta z + N \log(2\pi), \tag{2.11}$$

where $\delta z = y - H\mu_0$ is the innovation. A direct evaluation of (2.11) is numerically intensive in high dimensions, which is always the case for images. In particular, direct computation of the determinant term can lead to an indeterminate result. The following approximation is used instead [29]:

$$\log |Q| = \log |R| + \log |I + A|,$$

where $A = R^{-1} H P_0 H'$.

3 Test results

3.1 Synthetic data experiments

Synthetic images depicting spatial tracer distribution are used to reconstruct the known underlying flow field, subsequently referred to as the truth. This is simulated on a domain Ω, which represents a slice of marine sediment whose depth is assumed to be sufficiently "small" that flow is constrained to the 0.2 m $\times$ 0.2 m cross-sectional area.

Figure 1: Boundary conditions imposed on Ω.

The pressure boundary conditions indicate that: there is no advective flux across the left and bottom boundaries because they represent solid walls; there is no advective flux across the right-hand boundary, except for at an inlet Λ midway down, because it represents a symmetry axis; and fluid can flow across the open top boundary, where pressure is set to a fixed value P_{hyd} reflecting hydrostatic pressure. The pressure at the inlet P_{inj} is set to 205 Pa above P_{hyd}. This yields a fluid injection rate $\mathcal{Q} = 1.8$ ml min^{-1} in this numerical implementation, which is representative of lugworm pumping behaviour [30].

The pressure field satisfies $\nabla^2 p = 0$, representing incompressibility. The truth is then computed using Darcy's Law in (3.1) with sediment permeability $k = 1 \times 10^{-10}$ m^2 and dynamic viscosity $\hat{\mu} = 1 \times 10^{-3}$ Pa s:

$$v = -\frac{k}{\hat{\mu}\phi} \nabla p. \tag{3.1}$$

Fluid with tracer concentration $C_{inj} = 1$ is injected at the inlet. Zero gradient conditions are applied on all boundaries (including the top), because the images cover a time period before a "meaningful" amount of tracer reaches the upper boundary. Initially at zero, the concentration field is evolved forward in time subject to advection and diffusion (with $D = 7.0 \times 10^{-11}$ m^2 s^{-1}) using a finite volume toolbox for MATLAB [31]. So the images represent snapshots of the concentration field, under conditions in which the flow field is known.

Figure 2: A synthetic image depicting spatial tracer distribution at time $t_0 = 100$ s. The magenta contours denote the concentration field at time $t_1 = 150$ s, and the turquoise arrows denote the true flow field.

The prior mean in Example 1 is adopted here, and the fifth-order piecewise rational function is used to construct P_0 [28]. The correlation length scales are $\ell_x = 1.4\delta x$ m in the x-direction and $\ell_y = 1.4\delta y$ m in the y-direction (where δx and δy are pixel width and height); these are chosen to be "small" enough that spurious correlations are avoided, whilst being "large" enough to ensure that the optical flow is spatially smooth. The main diagonal of P_0 contains ones [28]: these variances are unrealistically large, because the truth is (on average) $O(10^{-5})$ m s^{-1}. We pre-multiply P_0 by a scalar $\alpha = 1 \times 10^{-8}$: the non-informative prior mean is accounted for with a standard deviation that is one order of magnitude greater than the average speed of the truth.

The approach in Section 2.2 is used to select observation error variances, with respect to the domain $(\sigma_1^2, \sigma_2^2) \in [1 \times 10^{-8}, 1 \times 10^{-4}) \times [1 \times 10^{-6}, 1 \times 10^{-3})$. The truth's speed and the observations are (on average) $O(10^{-5})$ m s^{-1} and $O(10^{-4})$ s^{-1}; the associated standard deviations σ_1 and σ_2 being at least one order of magnitude greater account for lack of prior knowledge. We choose $\sigma_1^2 = 1 \times 10^{-8}$ m^2 s^{-2} and $\sigma_2^2 = 1 \times 10^{-6}$ s^{-2} (see Figure 3). Figure 4 illustrates the importance of this choice: normalising posterior variances by α generates a number on the interval $(0, 1]$, where those close to one indicate a low level of PVR.

Figure 3: Log likelihood versus $\log \sigma_1^2$ and $\log \sigma_2^2$.

Figure 4: Variances (normalised by α) of x- (left panel) and y- (right panel) components of the optical flow arising from the Bayesian approach when $\sigma_1^2 = 9 \times 10^{-5}$ m^2 s^{-2} and $\sigma_2^2 = 9 \times 10^{-4}$ s^{-2}.

For comparative purposes, an optical flow is computed using the deterministic multi-resolution Horn-Schunck method [32] besides the Bayesian approach. The level of smoothing by the constraint term is controlled by the parameter $\lambda \in \mathbb{R}$ [8]; given the spatial intensity derivatives in Figure 6, choosing $\lambda = 60$ means that neither the data-fidelity nor the constraint term dominate the other.

Figure 5: Optical flows arising from the multi-resolution Horn-Schunck method (left panel) [32] and the Bayesian approach when $\sigma_1^2 = 1 \times 10^{-8}$ m^2 s^{-2} and $\sigma_2^2 = 1 \times 10^{-6}$ s^{-2} (right panel), overlaid on top of their angular errors.

Figure 6: Spatial intensity derivatives in the x- (left panel) and y- (middle panel) directions, and the temporal intensity derivative (right panel), arising from Figure 2.

The innovation is $-f_t$ because the prior mean is a zero flow. Figure 6 illustrates how the absolute values of f_t are greatest in a "half-ring" surrounding the inlet; this is where the prior mean is updated most. So the speed and divergence of both optical flows in Figure 5 increase as we initially depart from the inlet. This is unrealistic because fluid is injected at the inlet, so flow should be fastest there.

When the left- and right-hand panels in Figure 7 are overlaid, the region of maximum PVR roughly coincides with that where the angular errors in the right-hand panel of Figure 5 are smallest. This suggests that, for this application, posterior variances can act as a proxy for residual errors in optical flows.

Figure 7: Variances (normalised by α) of x- (left panel) and y- (right panel) components of the optical flow in Figure 5 arising from the Bayesian approach.

3.2 Real data experiment

Residual errors cannot be computed when dealing with real images because the truth is unknown. The Bayesian algorithm is applied, and posterior variances are used to try quantifying uncertainty. The images in Figure 8 capture a slice of marine sediment with a 0.22 m $\times$ 0.445 m cross-sectional area, as observed in a thin aquarium. Unlike the setup in Section 3.1, there is no inlet here: fluid movement instead arises from the bioirrigation activity (pumping) of the lugworm. These images are converted to greyscale and downsampled from a resolution of 2848 $\times$ 4256 pixels to a more tractable resolution of 89 $\times$ 133. Because (2.1) only applies to saturated porous media, the downsampled images are cropped to remove the overlying water as well as aquaria walls (see the red boundaries in Figure 9). Flow is determined using 41 images, captured at 10 s intervals, over a period where pumping activity is evident.

Figure 8: Real images capturing spatial tracer distribution at times $t_0 = 0$ s (left panel), $t_1 = 190$ s (middle panel) and $t_2 = 390$ s (right panel).

Figure 9: The images in Figure 8 converted to greyscale. The red lines denote cropping boundaries.

The prior distribution at time $t_0 = 0$ s and the observation error variances are the same as those in Section 3.1 because the physical setups of the two experiments are so similar. The correlation length scales are rescaled to $\ell_x = 5.0\delta x$ m and $\ell_y = 2.2\delta y$ m though, since the size of the aquarium is different to that in Section 3.1.

Flow speeds range from $O(10^{-5})$ m s^{-1} to $O(10^{-4})$ m s^{-1} (see Figure 10). The Péclet number is greater than one at these speeds for length scales greater than the grain scale. So diffusion is negligible, which justifies one of the assumptions made in deriving (2.2). Figure 10 also shows a decrease in flow speed during the initial 100 s. One possible explanation is that the lugworm pumps less vigorously over time. This is consistent with the computed divergence fields in Figures 11-13. The right-hand panel in Figure 11 depicts "large", positive divergences close to the likely location of injection. Such a pattern suggests "source-like" dynamics, where fluid is forced towards the image plane and spreads out when it reaches the aquarium boundary.

Figure 10: Histograms of flow speeds at each time point.

Figure 11: The optical flow at time $t_0 = 0$ s overlaid on top of the greyscale image at that time which has been downsampled and cropped (left panel), the average of the x- and y- component variances (middle panel), and flow divergence (right panel).

Figure 12: The optical flow at time $t_1 = 190$ s overlaid on top of the greyscale image at that time which has been downsampled and cropped (left panel), the average of the x- and y- component variances (middle panel), and flow divergence (right panel).

Figure 13: The optical flow at time $t_2 = 390$ s overlaid on top of the greyscale image at that time which has been downsampled and cropped (left panel), the average of the x- and y- component variances (middle panel), and flow divergence (right panel).

Injection of fluid into the sediment takes place near where there is most spatial contrast initially. So flow should be fastest there. Over time, tracer spreads out (see the left-hand panels in Figures 11-13). The magnitude of f_t temporally decreases near the likely location of injection, suggesting that the prior distribution is updated less there.

The prior distribution can be better updated elsewhere on the domain as the tracer, and the region of high spatial contrast, spreads. This means that the region of maximum PVR, i.e. where we can be most certain of the optical flow, grows. And if posterior variances can indeed be used as a proxy for residual errors in optical flows, as Section 3.1 suggests, then spatial mean residual errors should temporally decrease.

4 Conclusion

A Bayesian algorithm to estimate unobserved velocity fields in porous media from images, whose intensities reflect the concentration of a fluorescent tracer, has been developed. The presence of a variance means we know where we can be most confident that the optical flow mimics reality. We have noticed, in this application, how posterior variances and angular errors appear to be smallest in the same spatial regions. The computation of variances has enabled uncertainty quantification when dealing with the real experimental images, i.e. when the truth is unknown.

This Bayesian algorithm requires some modifications to ensure that optical flows are incompressible away from the likely location of injection. One option is to estimate the stream function rather than the velocity field [9, 33], because the divergence of the curl operator is zero. This would lead to a nonlinear problem, which would require a further redesign of this Bayesian algorithm. This would be advantageous because the observation operator could then be easily modified to account for reactive tracers, which are measured routinely using planar optodes [34].

Appendix A Fifth-order piecewise rational function

The fifth-order piecewise rational function proposed by Gaspari and Cohn (1999) is given by [28]:

$$\rho = \begin{cases} -\frac{1}{4}\left(\frac{d}{\ell}\right)^5 + \frac{1}{2}\left(\frac{d}{\ell}\right)^4 + \frac{5}{8}\left(\frac{d}{\ell}\right)^3 - \frac{5}{3}\left(\frac{d}{\ell}\right)^2 + 1; & 0 \leq d \leq \ell \\ \frac{1}{12}\left(\frac{d}{\ell}\right)^5 - \frac{1}{2}\left(\frac{d}{\ell}\right)^4 + \frac{5}{8}\left(\frac{d}{\ell}\right)^3 + \frac{5}{3}\left(\frac{d}{\ell}\right)^2 - 5\left(\frac{d}{\ell}\right) + 4 - \frac{2}{3}\left(\frac{\ell}{d}\right); & \ell < d \leq 2\ell \\ 0; & d > 2\ell \end{cases}$$

where $d \in \mathbb{R}$ is the physical distance between two pixels and $\ell \in \mathbb{R}$ is the correlation length scale.

References

[1] Heitz D, Mémin E, Schnörr C. Variational fluid flow measurements from image sequences: synopsis and perspectives. Experiments in Fluids. 2010;48(3):369–393.

[2] Kaipio J, Somersalo E. Statistical and Computational Inverse Problems. vol. 160. Antman S, Marsden J, Sirovich L, editors. Springer Science & Business Media; 2006.

[3] Sun J, Bollt E. Statistical Inverse Formulation of Optical Flow with Uncertainty Quantification. arXiv:161101230. 2016;.

[4] Cohen I, Herlin I. Optical flow and phase portrait methods for environmental satellite image sequences. In: Buxton B, Cipolla R, editors. European Conference on Computer Vision. vol. 1065. Springer; 1996. p. 141–150.

[5] Santitissadeekorn N, Bollt E. The infinitesimal operator for the semigroup of the Frobenius-Perron operator from image sequence data: Vector fields and transport barriers from movies. Chaos: An Interdisciplinary Journal of Nonlinear Science. 2007;17(2):023126.

[6] Basnayake R, Bollt E. A Multi-time Step Method to Compute Optical Flow with Scientific Priors for Observations of a Fluidic System. In: Bahsoun W, Bose C, Froyland G, editors. Ergodic Theory, Open Dynamics, and Coherent Structures. Springer; 2014. p. 59–79.

[7] Corpetti T, Mémin E, Pérez P. Adaptation Of Standard Optic Flow Methods To Fluid Motion. In: 9th International Symposium on Flow Visualisation; 2000. p. 1–10.

[8] Horn B, Schunck B. Determining optical flow. Artificial Intelligence. 1981;17(1-3):185–203.

[9] Luttman A, Bollt E, Basnayake R, Kramer S, Tufillaro N. A framework for estimating potential fluid flow from digital imagery. Chaos: An Interdisciplinary Journal of Nonlinear Science. 2013;23(3):033134.

[10] Souhila K, Karim A. Optical Flow Based Robot Obstacle Avoidance. International Journal of Advanced Robotic Systems. 2007;4(1):2.

[11] Beauchemin S, Barron J. The computation of optical flow. ACM Computing Surveys. 1995;27(3):433–466.

[12] Ullman S. The interpretation of visual motion. Massachusetts Institute of Technology; 1979.

[13] Aubert G, Deriche R, Kornprobst P. Computing Optical Flow via Variational Techniques. SIAM Journal on Applied Mathematics. 1999;60(1):156–182.

[14] Fortun D, Bouthemy P, Kervrann C. Optical flow modeling and computation: A survey. Computer Vision and Image Understanding. 2015;134:1–21.

[15] Chantas G, Gkamas T, Nikou C. Variational-Bayes Optical Flow. Journal of Mathematical Imaging and Vision. 2014;50(3):199–213.

[16] Vogel C. Computational Methods for Inverse Problems. SIAM; 2002.

[17] Corpetti T, Heitz D, Arroyo G, Mémin E, Santa-Cruz A. Fluid experimental flow estimation based on an optical-flow scheme. Experiments in Fluids. 2006;40(1):80–97.

[18] Simoncelli E, Adelson E. Computing Optical Flow Distributions Using Spatio-temporal Filters. Vision and Modeling Group, Media Laboratory, Massachusetts Institute of Technology; 1991.

[19] Simoncelli E, Adelson E, Heeger D. Probability distributions of optical flow. In: Computer Vision and Pattern Recognition. IEEE; 1991. p. 310–315.

[20] McCane B, Novins K, Crannitch D, Galvin B. On Benchmarking Optical Flow. Computer Vision and Image Understanding. 2001;84(1):126–143.

[21] Santos I, Eyre B, Huettel M. The driving forces of porewater and groundwater flow in permeable coastal sediments: A review. Estuarine, Coastal and Shelf Science. 2012;98:1–15.

[22] Meile C, Cappellen P. Global estimates of enhanced solute transport in marine sediments. Limnology and Oceanography. 2003;48(2):777–786.

[23] Roskosch A, Morad M, Khalili A, Lewandowski J. Bioirrigation by Chironomus plumosus: advective flow investigated by particle image velocimetry. Journal of the North American Benthological Society. 2010;29(3):789–802.

[24] Roskosch A, Hupfer M, Nützmann G, Lewandowski J. Measurement techniques for quantification of pumping activity of invertebrates in small burrows. Fundamental and Applied Limnology/Archiv für Hydrobiologie. 2011;178(2):89–110.

[25] Delefosse M, Kristensen E, Crunelle D, Braad P, Dam J, Thisgaard H, et al. Seeing the Unseen-Bioturbation in 4D: Tracing Bioirrigation in Marine Sediment Using Positron Emission Tomography and Computed Tomography. PloS ONE. 2015;10(4):e0122201.

[26] Meile C, Koretsky C, Cappellen P. Quantifying bioirrigation in aquatic sediments: An inverse modeling approach. Limnology and Oceanography. 2001;46(1):164–177.

[27] Riisgård H, Banta G. Irrigation and deposit feeding by the lugworm Arenicola marina, characteristics and secondary effects on the environment. A review of current knowledge. Vie Milieu. 1998;48(4):243–257.

[28] Gaspari G, Cohn S. Construction of correlation functions in two and three dimensions. Quarterly Journal of the Royal Meteorological Society. 1999;125(554):723–757.

[29] Ipsen I, Lee D. Determinant Approximations. arXiv preprint arXiv:11050437. 2011;.

[30] Riisgård H, Berntsen I, Tarp B. The lugworm (Arenicola marina) pump: characteristics, modelling and energy cost. Marine Ecology Progress Series. 1996;138(1-3):149–156.

[31] Eftekhari A. FVTool: a finite volume toolbox for Matlab (Version v0.11); 2015. Accessed: 20/09/2017. Zenodo. http://doi.org/10.5281/zenodo.18156.

[32] Black M, Anandan P. The Robust Estimation of Multiple Motions: Parametric and Piecewise-Smooth Flow Fields. Computer Vision and Image Understanding. 1996;63(1):75–104.

[33] Azijli I, Dwight R, Bijl H. A Bayesian Approach to Physics-Based Reconstruction of Incompressible Flows. In: Abdulle A, Deparis S, Kressner D, Nobile F, Picasso M, editors. Numerical Mathematics and Advanced Applications-ENUMATH 2013. Springer; 2015. .

[34] Glud R, Ramsing N, Gundersen J, Klimant I. Planar optrodes: a new tool for fine scale measurements of two-dimensional O_2 distribution in benthic communities. Marine Ecology Progress Series. 1996;140(1-3):217–226.

[35] Brox T, Bruhn A, Papenberg N, Weickert J. High Accuracy Optical Flow Estimation Based on a Theory for Warping. Computer Vision-ECCV 2004. 2004;3024:25–36.

[36] Volkenborn N, Hedtkamp S, van Beusekom J, Reise K. Effects of bioturbation and bioirrigation by lugworms (Arenicola marina) on physical and chemical sediment properties and implications for intertidal habitat succession. Estuarine, Coastal and Shelf Science. 2007;74(1-2):331–343.

[37] Weickert J, Bruhn A, Papenberg N, Brox T. Variational Optic Flow Computation: From Continuous Models to Algorithms. IWCVIA. 2003;3:1–6.