Inverse Problems 34 (2018) 105008 (26pp)

https://doi.org/10.1088/1361-6420/aad7cc

# Bayesian optical flow with uncertainty quantification

# Jie Sun<sup>1,2,3,4</sup>, Fernando J Quevedo<sup>4,5</sup> and Erik Bollt<sup>1,4,6</sup>

<sup>1</sup> Department of Mathematics, Clarkson University, Potsdam, NY 13699, United States of America

<sup>2</sup> Department of Physics, Clarkson University, Potsdam, NY 13699,

United States of America

<sup>3</sup> Department of Computer Science, Clarkson University, Potsdam, NY 13699, United States of America

<sup>4</sup> Clarkson Center for Complex Systems Science ( $C^3S^2$ ), Clarkson University, Potsdam, NY 13699, United States of America

<sup>5</sup> Department of Mechanical and Aerospace Engineering, Clarkson University, Potsdam, NY 13699, United States of America

<sup>6</sup> Department of Electrical and Computer Engineering, Clarkson University, Potsdam, NY 13699, United States of America

# E-mail: sunj@clarkson.edu

Received 12 October 2017, revised 6 July 2018 Accepted for publication 3 August 2018 Published 20 August 2018



#### Abstract

Optical flow refers to the visual motion observed between two consecutive images. Since the degree of freedom is typically much larger than the constraints imposed by the image observations, the straightforward formulation of optical flow as an inverse problem is ill-posed. Standard approaches to determine optical flow rely on formulating and solving an optimization problem that contains both a data fidelity term and a regularization term, the latter effectively resolves the otherwise ill-posedness of the inverse problem. In this work, we depart from the deterministic formalism, and instead treat optical flow as a statistical inverse problem. We discuss how a classical optical flow solution can be interpreted as a point estimate in this more general framework. The statistical approach, whose 'solution' is a distribution of flow fields, which we refer to as Bayesian optical flow, allows not only 'point' estimates (e.g. the computation of average flow field), but also statistical estimates (e.g. quantification of uncertainty) that are beyond any standard method for optical flow. As application, we benchmark Bayesian optical flow together with uncertainty quantification using several types of prescribed ground-truth flow fields and images.

Keywords: optical flow, uncertainty quantification, statistical inversion, inverse problems

(Some figures may appear in colour only in the online journal)

# 1. Introduction

Optical flow reflects the visual motion between consecutive images. Determination of optical flow is important for applications ranging from machine learning and computer vision [28], artificial intelligence and robotics [4, 25], to scientific applications from oceanography to weather forecasting [5, 6, 10, 19, 20], to name a few. In classical approaches optical flow is determined by solving a variational optimization problem which requires inclusion of a regularization term in additional to data fitting for the problem to be well-posed [1, 16, 35]. The choice of regularization parameter turns out critical for the satisfactory inference of optical flow. Despite the many (competing) methods of selecting the regularization parameter, none seems to be most 'natural' comparing to the others [15, 30, 34]. Another important feature of classical optical flow approaches is that they produce a single solution (by design) as a 'point estimate'. In practice, the magnitude of the flow as well as measurement error and noise can vary significantly from one part of the images to another, and thus at a given pixel the inferred flow vector can be associated with either large or small uncertainty. Although not captured in classical optical flow, such 'uncertainty' would provide valuable information if attainable as part of the solution.

In this work we adopt a statistical inversion framework for the estimation of optical flow. In this framework the computation of optical flow is reformulated from the Bayesian perspective as a statistical inversion problem. From this new, statistical perspective, different types of information and model assumptions are collectively fused to give rise to a posterior distribution of the unknowns of interest. The posterior distribution, which is typically sampled using some appropriately designed Markov chain Monte Carlo scheme, can be further used to derive various estimates. Importantly, unlike the point estimate of optical flow obtained by classical variational approaches, the proposed statistical inversion approach is a methodological way of uncertainty propagation to produce a distribution of candidate optical flow fields from which various statistical properties can be extracted, including ensemble average estimation and uncertainty quantification.

The rest of the paper is organized as follows. In section 2 we review the standard optical flow setup and discuss how it can be treated as an inverse problem defined in finite dimensions. In section 3 we we review the basics of inverse problems, including the standard least squares approach and Tikhonov regularization. Then, in section 4 we present a statistical inversion approach for optical flow, and discuss the choice of priors, models, sampling procedure and computational algorithms. In section 5 we showcase the utility statistical inversion approach for optical flow using several benchmark examples of flow fields and noisy image data. Finally, section 6 contains a conclusion and discussion of several unsolved and unattempted issues that might be of interest for future research.

# 2. Optical flow as an inverse problem

Given a sequence of consecutively captured images, the visual relative motion between them is commonly referred to as optical flow, which can often provide insight about the actual physical motion. The inference of optical flow is an outstanding scientific question, which requires making assumptions about the underlying motion as well as the measurement process. Consider two single-channeled (typically grayscale) digital images ('pictures') taken from the same scene at two nearby time instances. The image data are represented by two matrices  $F = [F_{ij}]_{n_x \times n_y}$  and  $G = [G_{ij}]_{n_x \times n_y}$ . Thus each image contains  $n_x \times n_y$  pixels, defined on a common two-dimensional subspace  $\Omega$ . The goal of optical flow estimation is to infer a flow field, defined by matrices  $U = [U_{ij}]_{n_x \times n_y}$  and  $V = [V_{ij}]_{n_x \times n_y}$  where  $\langle U_{ij}, V_{ij} \rangle$  represents the optical 'velocity' occurring at the (i, j)th pixel inferred from the two images. The image data *F* and *G* are often regarded as sampled data from smooth functions F(x, y) and G(x, y), with  $F_{ij} = F(x_i, y_j)$  and  $G_{ij} = G(x_i, y_j)$  where  $\{(x_i, y_j)\}_{(i=1,...,n_x;j=1,...,n_y)}$  are grid points from a spatial domain  $\Omega$ . Thus *U* and *V* can be viewed as discrete spatial samples of a smooth 2D flow field  $\vec{W}(x, y) = \langle U(x, y), V(x, y) \rangle$ , which is defined on  $\Omega$  that captures the visual optical motion occurring between the two observed images.

#### 2.2. Variational approach of inferring optical flow

The classical variational approach of optical flow starts by defining an 'energy' functional whose minimization yields an estimation of the optical flow field [1]. One of the most widely used functional was proposed by Horn and Schunck in 1981 [16], given by

$$E(U,V) = \iint_{\Omega} (F_x U + F_y V + F_t)^2 dx dy + \alpha \iint_{\Omega} (\|\nabla U\|^2 + \|\nabla V\|^2) dx dy,$$
(1)

where U(x, y) and V(x, y) are smooth functions defined over  $\Omega$  which represent a candidate flow field. In the Horn–Schunck functional, the first term is often referred to as *data fidelity* as it measures the deviation of the total image intensity from being conservative, that is, how much does the model fit deviates from the condition

$$dF/dt = F_x U + F_y V + F_t = 0.$$
 (2)

The second term measures the *solution regularity* by penalizing solutions that have large spatial gradients and is called the regularization term. The relative emphasis of smoothness as compared to 'fitting' the total image intensity conservation equation (2) is controlled by the positive scalar  $\alpha$  which is called a *regularization parameter*. The main role of the regularization term is to ensure that the minimization of the functional is a well posed problem. Without the regularization term, the problem is ill-posed.

Given  $\alpha$ , the functions U and V that minimize the Horn–Schunck functional (1) satisfy the Euler–Lagrange equations

$$\begin{cases} F_x(F_xU + F_yV + F_t) = \alpha(U_{xx} + U_{yy}), \\ F_y(F_xU + F_yV + F_t) = \alpha(V_{xx} + V_{yy}), \end{cases}$$
(3)

which are typically solved by some iterative scheme over a finite set of spatial grid points [16, 29] to produce an estimation of the optical flow. Alternatively, one could also discretize the functional (1) itself to yield a finite-dimensional inverse problem as discussed in section 2.3 below with solution strategy reviewed in section 3.

#### 2.3. Finite-dimensional representation of the variational optical flow functional

As discussed in the previous section, the classical variational approach of optical flow works by first formulating and minimizing a functional over smooth vector fields, and then evaluating the obtained vector field at the grid points on which the original image data are given. Here we approach the problem from a different route, by first discretizing the functional (1) to convert the functional minimization (an infinite-dimensional problem) into a finite-dimensional linear inverse problem, and then solving the inverse problem to yield solutions which give the values of a vector field defined over the same grid points as the image data.

The remainder of this section will be focused on the conversion of the functional (1) into a finite-dimensional function defined on a set of uniformly distributed grid points

$$\{(x_i, y_j)\}_{i=1,2,\dots,n_x; j=1,2,\dots,n_y},\tag{4}$$

where  $x_{i+1} - x_i = \Delta x$  and  $y_{j+1} - y_j = \Delta y$  are the spacing in the *x* direction and the *y* direction, respectively. The conversion will be achieved by approximating the integrals in equation (1) with appropriately derived summations over the grid points. For notational convenience, we use a bold lowercase variable to denote the vectorization of a matrix. For example, the bold-face vector  $\boldsymbol{q}$  denotes the column vector obtained by 'vertically stacking' the columns of a matrix  $Q = [\vec{Q}_1, \vec{Q}_2, \dots, \vec{Q}_n]$  in order [12], where  $\vec{Q}_i$  denotes the *i*th column of Q. That is,

$$\boldsymbol{q} = \operatorname{vec}(\boldsymbol{Q}) = \begin{bmatrix} \vec{Q}_1 \\ \vec{Q}_2 \\ \vdots \\ \vec{Q}_n \end{bmatrix}.$$
 (5)

First let us consider the data fidelity term:  $\iint_{\Omega} (F_x U + F_y V + F_t)^2 dx dy$ . The spatial derivatives  $F_x(x,y)$  and  $F_y(x,y)$  can be approximated by a finite difference scheme. For example, the simple *forward difference* yields the approximations:

$$\begin{cases} F_x(x,y) \approx \frac{1}{\Delta x} \left[ F(x + \Delta x, y) - F(x, y) \right], \\ F_y(x,y) \approx \frac{1}{\Delta y} \left[ F(x, y + \Delta y) - F(x, y) \right]. \end{cases}$$
(6)

We next express these derivatives as operations on the column vector f. To do this, we define matrix  $S_k = [S_k^{(ij)}]_{k \times k}$  as

$$S_{k}^{(ij)} = \begin{cases} -\delta_{ij} + \delta_{i+1,j}, & \text{if } i < k; \\ -\delta_{i,j-1} + \delta_{i,j}, & \text{if } i = k. \end{cases}$$
(7)

The forward difference applied to f can be represented as

$$\begin{cases} f_x \approx Q_x f, \\ f_y \approx Q_y f, \end{cases}$$
(8)

where

$$\begin{cases} Q_x \equiv \frac{1}{\Delta x} [I_n \otimes S_m], \\ Q_y \equiv \frac{1}{\Delta y} [S_n \otimes I_m]. \end{cases}$$
(9)

The temporal derivative can be estimated from the data by the direct pixel-wise difference between the two images, to yield

$$\boldsymbol{f}_t \approx \boldsymbol{\tilde{g}} - \boldsymbol{f}. \tag{10}$$

With these definitions and approximations, we obtain a discretized version of the conservation equation (2) expressed as a finite-dimensional linear system:

$$A\boldsymbol{x} = \boldsymbol{b},\tag{11}$$

where

$$\begin{cases}
A = [\operatorname{diag}(f_x), \operatorname{diag}(f_y)], \\
\mathbf{x} = [\mathbf{u}^\top, \mathbf{v}^\top]^\top, \\
\mathbf{b} = -f_t.
\end{cases}$$
(12)

Here diag( $f_x$ ) and diag( $f_y$ ) represent the diagonal matrices whose diagonal elements are given by the entries of the vector  $f_x$  and  $f_y$ , respectively. From this connection we approximate the first integral in the functional (1) as  $||A\mathbf{x} - \mathbf{b}||^2$  where  $|| \cdot ||$  denotes the standard Euclidean norm.

Next we develop a finite-dimensional approximation of the regularization term in the functional (1). This requires discretization of  $\nabla U$  and  $\nabla V$ . Using a similar forward difference to approximate the partial derivatives, we obtain

$$\begin{cases} \nabla U \approx \frac{1}{\Delta x} \left[ U(x + \Delta x, y) - U(x, y) \right] + \frac{1}{\Delta y} \left[ U(x, y + \Delta y) - U(x, y) \right], \\ \nabla V \approx \frac{1}{\Delta x} \left[ V(x + \Delta x, y) - V(x, y) \right] + \frac{1}{\Delta y} \left[ V(x, y + \Delta y) - V(x, y) \right]. \end{cases}$$
(13)

For the vectorized variables  $\boldsymbol{u}$  and  $\boldsymbol{v}$ , we have

$$\begin{cases} \nabla \boldsymbol{u} = \boldsymbol{u}_x + \boldsymbol{u}_y \approx (Q_x + Q_y)\boldsymbol{u}, \\ \nabla \boldsymbol{v} = \boldsymbol{v}_x + \boldsymbol{v}_y \approx (Q_x + Q_y)\boldsymbol{v}, \end{cases}$$
(14)

where  $Q_x$  and  $Q_y$  are defined in equation (9). Consequently, we obtain the approximation of the second integral in the Horn–Schunck functional (1) as

$$\begin{cases} \iint_{\Omega} \|\nabla U\|^2 dx dy \approx \boldsymbol{u}_x^\top \boldsymbol{u}_x + \boldsymbol{u}_y^\top \boldsymbol{u}_y \approx \boldsymbol{u}^\top [\boldsymbol{Q}_x^\top \boldsymbol{Q}_x + \boldsymbol{Q}_y^\top \boldsymbol{Q}_y] \boldsymbol{u}, \\ \iint_{\Omega} \|\nabla V\|^2 dx dy \approx \boldsymbol{v}_x^\top \boldsymbol{u}_y + \boldsymbol{v}_y^\top \boldsymbol{v}_y \approx \boldsymbol{v}^\top [\boldsymbol{Q}_x^\top \boldsymbol{Q}_x + \boldsymbol{Q}_y^\top \boldsymbol{Q}_y] \boldsymbol{v}, \end{cases}$$
(15)

which then gives

$$\iint_{\Omega} (\|\nabla u\|^2 + \|\nabla v\|^2) dx dy \approx \mathbf{x}^{\top} Q \mathbf{x},$$
(16)

where the matrix

$$Q = I_2 \otimes [Q_x^\top Q_x + Q_y^\top Q_y].$$
<sup>(17)</sup>

Therefore, the variational optical flow formulation (1) can be reformulated at a finite spatial resolution as an inverse problem, where, for a given parameter value  $\alpha$ , the corresponding solution is given by solving the following (regularized) optimization:

$$\min\left(\|A\boldsymbol{x} - \boldsymbol{b}\|^2 + \alpha \boldsymbol{x}^\top \boldsymbol{Q} \boldsymbol{x}\right). \tag{18}$$

This is a standard approach in inverse problems, formulated as a least squares problem with Tikhonov regularization, with more details to be presented in the next section.

# 3. Inverse problem in finite dimensions

Although there has been a great deal of progress on the mathematical characterization of inverse problems in the field of functional analysis, a practical problem often concerns finding a solution in a finite-dimensional space. At a fundamental level, the most common inverse problem stems from a linear model [15, 17, 34]

(19)

$$b = Ax + \eta$$
,

where  $\boldsymbol{b} \in \mathbb{R}^m$  is a column vector of *observed data*,  $A = [a_{ij}]_{m \times n} \in \mathbb{R}^{m \times n}$  is a (known) matrix representing the underlying model, the column vector  $\boldsymbol{\eta} \in \mathbb{R}^n$  denotes (additive) noise, and  $\boldsymbol{x} \in \mathbb{R}^n$  is the vector of *unknowns* to be inferred.

Given A and **b**, the problem of inferring **x** in equation (19) is called an *inverse problem* because rather than direct 'forward' computation from the model, it requires a set of indirect, 'backward', or 'inverse' operations to determine the unknowns [34]. Depending on the rank and conditioning of the matrix A, the problem may be ill-posed or ill-conditioned. In classical approaches, these issues are dealt with by adjusting the original problem to a (slightly) modified optimization problem as discussed in section 2.1 whose solution is meant to represent the original, as discussed in section 3.2.

We note that in the classical setting a solution to the inverse problem is a vector  $\mathbf{x}$  as a result of solving an optimization problem. Such a solution is referred to as an *point estimate* because it gives one solution vector without providing any information about how reliable (or uncertain) the solution is [15, 34]. On the other hand, the statistical inversion approach to inverse problems provides an *ensemble* of solutions—defined by the *posteriori* distribution which not only point estimates can be made but also their uncertainty quantification [11, 17].

#### 3.1. Least squares solution

The classical least squares solution to the inverse problem is given by [12]

$$\boldsymbol{x}_{\ell_2} = A^{\dagger} \boldsymbol{b}, \tag{20}$$

where  $A^{\dagger}$  denotes the pseudo-inverse of A which can be obtained from the singular value decomposition of A [12]. Depending on the rank of A, the least squares solution  $\mathbf{x}_{\ell_2}$  is associated with one of the minimization problems:

$$\begin{cases} \min_{A\boldsymbol{x}=\boldsymbol{b}} \|\boldsymbol{x}\|_{2}, & \text{if rank}(A) < n; \\ \min_{\boldsymbol{x}} \|A\boldsymbol{x}-\boldsymbol{b}\|_{2}, & \text{if rank}(A) = n. \end{cases}$$
(21)

Here  $\|\cdot\|_2$  denotes the  $\ell_2$  (Euclidean) norm. Let the true solution to equation (19) be  $x^*$ , that is,  $b = Ax^* + \eta$ . It follows that

$$\mathbf{x}_{\ell_2} - \mathbf{x}^* = A^{\dagger} \boldsymbol{\eta}. \tag{22}$$

In practice, even when the matrix A has full column rank  $(\operatorname{rank}(A) = n)$ , the discrepancy between the true and least squares solutions is typically dominated by noise when some singular values of A are close to zero, rendering A an ill-conditioned matrix and the solution  $x^*$  unstable and sensitive to small changes in data [34].

# 3.2. Tikhonov regularization

A powerful approach to resolve the instabilities due to noise and the near-singularity of A is to *regularize* the problem. In the classical Tikhonov regularization one adds a quadratic regularization term  $\alpha x^T L x$  for some prescribed matrix L to penalize non-smoothness, giving rise to a regularized optimization problem [31–34]:

$$\min_{\mathbf{x}} \left( \|A\mathbf{x} - \mathbf{b}\|_2^2 + \alpha \mathbf{x}^T L \mathbf{x} \right).$$
(23)

In the regularized problem the positive scalar parameter  $\alpha$  controls the weight of regularization and *L* is typically a symmetric positive definite matrix, both of which need to be chosen appropriately for the problem to be uniquely defined [31–34]. We refer to the two terms  $||A\mathbf{x} - \mathbf{b}||_2^2$  and  $\alpha \mathbf{x}^\top L \mathbf{x}$  in (23) as *data fidelity* and *solution regularity*, respectively. In a simplistic description, they can be described as 'selecting' a solution  $\mathbf{x}_{\alpha}$  that balances the desire to 'solve'  $A\mathbf{x} = \mathbf{b}$  and to be 'regular' as measured by  $\mathbf{x}^\top L \mathbf{x}$ . The regularization parameter  $\alpha$ therefore dictates the extent to which the compromise is made between the two.

For a given parameter  $\alpha$ , we denote the corresponding regularized solution by

$$\boldsymbol{x}_{\alpha} = \operatorname{argmin}_{\boldsymbol{x}} \{ \|A\boldsymbol{x} - \boldsymbol{b}\|^2 + \alpha \boldsymbol{x}^\top L \boldsymbol{x} \}.$$
(24)

By standard vector calculus, it can be shown that  $x_{\alpha}$  is in fact a solution to the modified linear system

$$(A^{\top}A + \alpha L)\boldsymbol{x}_{\alpha} = A^{\top}\boldsymbol{b}, \tag{25}$$

which is typically well-posed for appropriate choices of L and  $\alpha$ . When the matrices are large and sparse, equation (25) is often solved by iterative methods rather than a direct matrix inversion since the latter tends to be numerically costly and unstable [12].

In Tikhonov regularization, a key issue is how to choose the regularization parameter  $\alpha$  appropriately. In theory, a 'good' regularizer has the property that in the absence of noise, the solution to the regularized problem converges to the true solution when the regularization parameter  $\alpha \rightarrow 0$ . However, in practice, under the presence of noise, it is always a challenge to try to determine a good value for  $\alpha$ . If  $\alpha$  is too small, the instability and sensitivity of the original problem would still persist; whereas for too large of an  $\alpha$  the solution will be overregularized and not fit the data well. A good balance is thus paramount. Despite the existence and ongoing development of many competing methods for selecting  $\alpha$  most of which focus on asymptotical optimality as the number of data points approach infinity, none of them stands out as a best 'natural' choice unless specific *priori* information about the noise in the data are available (see chapter 7 of [34]). In the following section we discuss how the problem of selecting an exact regularization parameter is no longer required from a Bayesian perspective; instead, it suffices to start with some loose range of values represented by a probability distribution, unless more specific knowledge is available about the problem in which case the distribution can be chosen to reflect such information.

# 4. Statistical inversion approach

The statistical inversion approach to an inverse problem starts by treating all variables as random variables (e.g. x, b and  $\eta$  in equation (19)), and representing our knowledge (or absence of knowledge) of the unknowns as prior distributions. With observational data and a forward model (e.g. matrix A in equation (19)), the 'inversion' leads to an ensemble of points together with a posterior distribution from which various statistical estimates can be made [8, 11, 17, 21]. The key in the inversion is to use the Bayes rule to express the *posterior* distribution p(x|b), which is the conditional distribution of the 'solution vector' x given the observed data b, as [11, 17]

$$p(\mathbf{x}|\mathbf{b}) = \frac{1}{p(\mathbf{b})} p(\mathbf{b}|\mathbf{x}) \cdot p(\mathbf{x}).$$
(26)

Here the *likelihood* function  $p(\boldsymbol{b}|\boldsymbol{x})$  is the probability density function (pdf) of the random variable  $\boldsymbol{b}$  given  $\boldsymbol{x}$  which is determined by the underlying model;  $p(\boldsymbol{x})$  is the *priori* distribution

of x; and p(b) > 0 acts as a normalization constant which does not affect the solution procedure or the final solution itself.

Thus, in the statistical inversion formulation, each candidate solution  $\mathbf{x}$  is associated with the probability  $p(\mathbf{x}|\mathbf{b})$  that is determined (up to a normalization constant  $1/p(\mathbf{b})$ ) once the likelihood function and the prior distribution are given. For a given inverse problem, the likelihood function can be obtained using the underlying model such as equation (19) including the noise distribution. On the other hand, the prior distribution  $p(\mathbf{x})$  is typically constructed according to some prior knowledge of the solution. Unlike standard approaches which produce 'point estimates', in statistical inversion it is the posterior distribution that is considered as the 'solution' to the inverse problem. Based on the posterior distribution, one can further extract useful information such as point estimates and uncertainty quantification, by sampling from the distribution. Efficient sampling methods will be reviewed toward the end of this section.

The unique feature of enabling information fusion and uncertainty quantification has made the statistical inversion approach to inverse problems an attractive venue for the development of new theory and applications. In image processing applications, it has been utilized for many problems such as image denoising and deblurring [2, 18], sparse signal reconstruction [23], and more recently attempted for optical flow computation [9]. In particular, we note that our approach, although different in many of the technical aspects, shares a similar statistical inversion perspective as [9].

# 4.1. Statistical interpretation of optical flow obtained by Tikhonov regularization

Under the statistical inversion framework, solution to an inverse problem is the posterior distribution, from which point estimates can be further obtained. Among these point estimates, a particularly common one is the *maximum a posteriori* (MAP) estimator, which is defined as

$$\boldsymbol{x}_{\text{MAP}} = \operatorname{argmax}_{\boldsymbol{x}} p(\boldsymbol{x}|\boldsymbol{p}) = \operatorname{argmin}_{\boldsymbol{x}} \{-\ln p(\boldsymbol{x}|\boldsymbol{b})\}$$
(27)

$$= \operatorname{argmin}_{\mathbf{x}} \{-\ln p(\boldsymbol{b}|\boldsymbol{x}) - \ln p(\boldsymbol{x})\}.$$
(28)

As noted in [2, 17], the MAP estimator given by equation (28) produces a vector  $\mathbf{x}$  that is identical to the solution of a Tikhonov regularization specified in equation (24) upon appropriate choice of the model and prior pdf. In particular, consider the model given by equation (19) with independent and identically distributed (iid) Gaussian noise of variance  $\lambda^{-1}$ . It follows that

$$p(\boldsymbol{b}|\boldsymbol{x}) = p(\boldsymbol{\eta}) \propto \exp\left(-\frac{\lambda}{2} \|A\boldsymbol{x} - \boldsymbol{b}\|^2\right),$$
(29)

where the symbol ' $\propto$ ' means 'proportional to'. Under a Gaussian prior,

$$p(\mathbf{x}) \propto \exp\left(-\frac{\delta}{2}\mathbf{x}^{\top}L\mathbf{x}\right),$$
(30)

the term  $-\ln p(\boldsymbol{x}|\boldsymbol{b})$  in the MAP estimator becomes

$$-\ln p(\boldsymbol{x}|\boldsymbol{b}) \propto \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|^2 + (\delta/\lambda)\boldsymbol{x}^\top \boldsymbol{L}\boldsymbol{x}.$$
(31)

The choice of  $\alpha = \delta/\lambda$  in the Tikhonov regularization then yields a solution  $\mathbf{x}_{\alpha}$  that equals the vector  $\mathbf{x}_{MAP}$  given by the same MAP estimator. Thus, with these assumptions of the form of the noise, the distribution of the prior, there is a logical bridge between two different philosophies for inverse problems, in that a solution from Tikhonov regularization can be interpreted

as an MAP estimate of the *posteriori* distribution under appropriate choices of prior distributions and likelihood functions.

Furthermore, and crucially, as shown in the next section, more important information exists in the statistical inversion framework. Specifically, by sampling from the posterior distribution the statistical inversion approach allows for not only a point estimate but also other statistical properties associated with the solution, in particular spread around a point estimate which enables uncertainty quantification.

#### 4.2. Choice of priors and hyperpriors

In general, the statistical inversion framework allows flexibility in choosing prior distributions about unknowns and noise. When these prior distributions themselves contain unknown parameters, these parameters can themselves be thought of as random variables which follow hyperprior distributions. In this paper we will consider some elementary choice of the priors and hyperpriors, with the main purpose of showing how to setup the procedure for sampling the posterior distribution.

We focus on (multivariate) Gaussian priors for both the unknown x and noise  $\eta$ . In the absence of knowledge of how 'spread-out' the prior distributions are, we use additional parameters for these priors, where these parameters are taken to be random variables drawn from hyperprior distributions. This type of approach has been previously adopted and implemented in image denoising applications [2]. In particular, we consider a specific class of the noise and prior distributions:

$$\begin{cases} \text{likelihood function}: p(\boldsymbol{b}|\boldsymbol{x},\lambda) \propto \lambda^{m/2} \exp\left(-\frac{\lambda}{2} \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|^2\right),\\ \text{prior distribution}: p(\boldsymbol{x}|\delta) \propto \delta^{n/2} \exp\left(-\frac{\delta}{2} \boldsymbol{x}^\top \boldsymbol{L} \boldsymbol{x}\right). \end{cases}$$
(32)

Here the noise is assumed to be additive, Gaussian, and independent of the measured data, with variance  $\lambda^{-1}$ , giving rise to the form of the likelihood function. On the other hand, the prior distribution is considered to be Gaussian with covariance matrix  $(\delta L)^{-1}$  (matrix  $\delta L$  is referred to as the *precision* matrix). We choose L = Q where Q is given by equation (17) which corresponds to a spatial regularization measure. As it turns out, this choice of L is closely related to the selection of prior according to a spatial Gaussian Markov random field which is common in tackling spatial inverse problems [3, 14].

As we pointed out, to completely specify the posterior distribution, we also need to choose prior distributions for the parameters  $\lambda$  and  $\delta$ . These are often called *hyperpriors*. Following [2], we choose the priors for  $p(\lambda)$  and  $p(\delta)$  to be Gamma distributions, as

$$\begin{cases} p(\lambda) \propto \lambda^{\alpha_{\lambda}-1} \exp(-\beta_{\lambda}\lambda), \\ p(\delta) \propto \delta^{\alpha_{\delta}-1} \exp(-\beta_{\delta}\delta). \end{cases}$$
(33)

Such choice ensures that  $p(\lambda)$  and  $p(\delta)$  are conjugate hyperpriors, and makes the design of sampling of posterior distributions more convenient. In the absence of knowledge of the values of  $\lambda$  and  $\delta$ , one needs to choose the values of  $\alpha_{\lambda}$ ,  $\beta_{\lambda}$ ,  $\alpha_{\delta}$ , and  $\beta_{\delta}$  to ensure the distributions  $p(\lambda)$  and  $p(\delta)$  to be 'wide', allowing the Markov chain to explore a potentially larger region of the parameter space. Following [2, 3, 14], we set  $\alpha_{\lambda} = \alpha_{\delta} = 1$  and  $\beta_{\lambda} = \beta_{\delta} = 10^{-4}$ unless otherwise noted. We tested other choice of parameters as well and they mainly affect the length of the transient in the MCMC sampling process and do not seem to have a strong influence on the asymptotic outcome. Details of these will be discussed in section 5 along with numerical examples.

#### 4.3. Sampling from the posterior distribution

In the statistical inversion formalism, once the form of the posterior distribution is derived, the remaining part of the work is devoted to efficient sampling from the posterior distribution. Typically a Markov chain Monte Carlo (MCMC) sampling approach is adopted. The main idea is to generate a sequence of samples according to a prescribed Markov chain whose unique stationary distribution is the desired posterior distribution.

Note that under the Gaussian priors and gamma hyperpriors as we discussed above, the full conditional distributions that relate to the posterior distribution are given by

$$\begin{cases} p(\boldsymbol{x}|\lambda, \delta, \boldsymbol{b}) \propto \exp\left(-\frac{\lambda}{2} \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|^2 - \frac{\delta}{2} \boldsymbol{x}^\top \boldsymbol{L} \boldsymbol{x}\right), \\ p(\lambda|\boldsymbol{x}, \delta, \boldsymbol{b}) \propto \lambda^{m/2 + \alpha_{\lambda} - 1} \exp\left(\lambda \left[-\frac{1}{2} \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|^2 + \beta_{\lambda}\right]\right), \\ p(\delta|\boldsymbol{x}, \lambda, \boldsymbol{b}) \propto \delta^{n/2 + \alpha_{\delta} - 1} \exp\left(\delta \left[-\frac{1}{2} \boldsymbol{x}^\top \boldsymbol{L} \boldsymbol{x} - \beta_{\delta}\right]\right). \end{cases}$$
(34)

In other words,

$$\begin{cases} \boldsymbol{x}|\lambda, \delta, \boldsymbol{b} \sim \mathcal{N}\left((\lambda A^{\top}A + \delta L)^{-1}\lambda A^{\top}\boldsymbol{b}, (\lambda A^{\top}A + \delta L)^{-1}\right), \\ \lambda|\boldsymbol{x}, \delta, \boldsymbol{b} \sim \Gamma\left(m/2 + \alpha_{\lambda}, \frac{1}{2} \|A\boldsymbol{x} - \boldsymbol{b}\|^{2} + \beta_{\lambda}\right), \\ \delta|\boldsymbol{x}, \lambda, \boldsymbol{b} \sim \Gamma\left(n/2 + \alpha_{\delta}, \frac{1}{2}\boldsymbol{x}^{\top}L\boldsymbol{x} + \beta_{\delta}\right). \end{cases}$$
(35)

We adopt the block Gibbs sampler developed in [2, 17] as a specific MCMC procedure to sample the posterior distribution. In theory the sample distribution asymptotically converges to the true posterior distribution. The approach contains the following steps.

(i) Initialize  $\delta_0$  and  $\lambda_0$ , and set k = 0. (ii) Sample  $\mathbf{x}^k \sim \mathcal{N} \left( (\lambda A^\top A + \delta L)^{-1} \lambda A^\top \mathbf{b}, (\lambda A^\top A + \delta L)^{-1} \right)$ . (iii) Sample  $\lambda_{k+1} \sim \Gamma \left( m/2 + \alpha_{\lambda}, \frac{1}{2} \| A \mathbf{x}^k - \mathbf{b} \|^2 + \beta_{\lambda} \right)$ . (iv) Sample  $\delta_{k+1} \sim \Gamma \left( n/2 + \alpha_{\delta}, \frac{1}{2} (\mathbf{x}^k)^\top L (\mathbf{x}^k) + \beta_{\delta} \right)$ . (v) Set  $k \leftarrow k + 1$  and return to step (ii).

Here the computational burden is mainly due to step 1, which requires drawing samples from a multivariate Gaussian variable, which is equivalent to solving the following linear system at each iteration for  $x^k$ :

$$(\lambda_k A^{\top} A + \delta_k L) \mathbf{x}^k = \lambda_k A^{\top} \mathbf{b} + \mathbf{w}, \text{ where } \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \lambda_k A^{\top} A + \delta_k L).$$
 (36)

For large matrices, instead of a direct solve using Gauss elimination, an iterative method is usually preferred. Among the various notable iterative methods such as Jacobi, Gauss–Seidel (G–S), and conjugate gradient (CG) [22], we adopted the CG for all the numerical experiments as reported in this paper, with a starting vector of all zeros, maximum of 500 iterations, and error tolerance of  $10^{-6}$ . Note that the computational bottleneck in solving equation (36) can in principle be tackled using more efficient methods, for example, by exploring history of solutions with similar parameters to provide 'good' initial guess or by utilizing the Karhunen–Loève expansion to reduce the dimensionality of the problem.

# 5. Results: Bayesian optical flow from statistical inversion

#### 5.1. Benchmark flow fields and noisy image pairs

To benchmark the proposed statistical inversion approach for Bayesian optical flow, we consider 5 qualitatively different flow fields that span a broad diversity of possibilities. The flow fields, which are all defined on the normalized spatial domain  $[-1, 1] \times [-1, 1]$ , are generally not divergence-free. In particular, the flow fields are defined as follows.

Flow field 1:  $\langle U(x,y), V(x,y) \rangle = \langle x, y \rangle$ . Flow field 2:  $\langle U(x,y), V(x,y) \rangle = \langle -y, x \rangle$ . Flow field 3:  $\langle U(x,y), V(x,y) \rangle = \langle y, \sin(x) \rangle$ . Flow field 4:  $\langle U(x,y), V(x,y) \rangle = \langle -\pi \sin(0.5\pi x) \cos(0.5\pi y), \pi \cos(0.5\pi x) \sin(0.5\pi y) \rangle$ . Flow field 5:  $\langle U(x,y), V(x,y) \rangle = \langle -\pi \sin(\pi x) \cos(\pi y), \pi \cos(\pi x) \sin(\pi y) \rangle$ .

We first consider synthetic image pairs, where the first image is constructed by the equation

$$F(x,y) = \frac{1}{2} \left[ \cos(\pi x) \cos(\pi y) + 1 \right].$$
(37)

Then, for each optical flow field, we generate the second image G using the equation

$$g = f - f_x u - f_y v + \eta, \tag{38}$$

where f and g represent the vectorization of F and G, respectively,  $\eta$  denotes (multivariate) noise whose individual components are independently drawn from a Gaussian distribution with zero mean and fixed standard deviation  $\sigma$ , and spatial derivatives are numerically implemented using the forward difference scheme. All synthetic images are constructed at the fixed resolution of 30-by-30 pixels, with uniform spacing in both directions, represented by a set of 30-by-30 matrices.

#### 5.2. Bayesian optical flow with uncertainty quantification

For each image pair (F, G), we adopt the MCMC-Gibbs sampling procedure and corresponding choice of prior pdf and hyerpriors presented in section 4 to obtain an empirical posterior distribution p(U, V) as the solution of the statistical inversion problem, which we refer to as Bayesian optical flow. Unlike classical optical flow which provides a point estimate, the Bayesian optical flow can be thought of as an ensemble of flow fields each associated with some probability. From such an ensemble and associated probability distribution (the posterior distribution), we can further extract useful information. For example, the mean flow field can be computed using the sampling mean of the posterior distribution, which is shown in figures 1 through 6 to compare with the true underlying flow field. In particular, in each figure, the top row (a1)–(a3) shows the image data of the first image F (a1), the second image G generated from equation (38) with no noise (a2) and with noise under standard deviation  $\sigma = 0.02$  (a3), respectively. The middle rows (b1)–(b3) show the true optical flow field (b1) compared with the inferred 'mean' optical flow fields together with uncertainty quantification from the MCMC samples (b2) and (b3). The uncertainty regions are computed as follows. At each point z = (x, y), we construct a 2d normal pdf  $\mathcal{N}(\mu, \Sigma)$  by using the sample mean  $\mu$  and sample covariance  $\Sigma$  estimated from the MCMC samples after discarding the initial transients. This allows us to obtain a 'mean' optical flow at point z defined as  $\langle u(z), v(z) \rangle = \langle \mu_1, \mu_2 \rangle$ . Uncertainty is quantified by computing a confidence region that contains q probability mass of the fitted multivariate normal distribution given by [24].

$$(z-\mu)^{+}\Sigma^{-1}(z-\mu) \leqslant \chi_{2}^{2}(q).$$
(39)

Here  $\chi_2^2(q)$  denotes the *q*th quantile of the Chi-squared distribution with two degrees of freedom, that is,  $\chi_2^2(q) = K^{-1}(q)$  where *K* is the cdf of  $\chi_2^2$ . These confidence regions (shaded ellipses) are shown for the zoomed-in plots for the inferred optical flow fields. Finally, the last row (c1)–(c3) in each figure shows how the MCMC procedure produces a distribution of the effective regularization parameter  $\delta/\lambda$ . Panel (c1) shows the change of  $\delta/\lambda$  over time in the



**Figure 1.** Bayesian optical flow computed using statistical inversion, based on image pairs (F, G) under example flow field 1, where  $\langle U(x, y), V(x, y) \rangle = \langle x, y \rangle$ . Top row (a1)–(a3): image data generated according to equation (38) for *F* (a1) and using equation (37) for *G* either without noise (a2) or with noise standard deviation  $\sigma = 0.02$  (a3). Middle rows (b1)–(b3): true flow field (b1) and the mean flow field from the posterior distribution estimated using the MCMC samples. In the panels below (b2) and (b3), we also show the uncertainty regions (as shaded ellipses), also from MCMC samples as given by equation (39). Bottom row (c1)–(c3): time evolution (c1) as well as the distribution (c2)–(c3) of the effective regularization parameter  $\sigma/\lambda$ , where the distributions are obtained after discarding the initial transient in the MCMC sampling process.



**Figure 2.** Bayesian optical flow computed using statistical inversion, based on image pairs (F, G) under example flow field 2, where  $\langle U(x, y), V(x, y) \rangle = \langle -y, x \rangle$ . Top row (a1)–(a3): image data generated according to equation (38) for *F* (a1) and using equation (37) for *G* either without noise (a2) or with noise standard deviation  $\sigma = 0.02$  (a3). Middle rows (b1)–(b3): true flow field (b1) and the mean flow field from the posterior distribution estimated using the MCMC samples. In the panels below (b2) and (b3), we also show the uncertainty regions (as shaded ellipses), also from MCMC samples as given by equation (39). Bottom row (c1)–(c3): time evolution (c1) as well as the distribution (c2)–(c3) of the effective regularization parameter  $\sigma/\lambda$ , where the distributions are obtained after discarding the initial transient in the MCMC sampling process.



**Figure 3.** Bayesian optical flow computed using statistical inversion, based on image pairs (F, G) under example flow field 3, where  $\langle U(x, y), V(x, y) \rangle = \langle y, \sin(x) \rangle$ . Top row (a1)–(a3): image data generated according to equation (38) for *F* (a1) and using equation (37) for *G* either without noise (a2) or with noise standard deviation  $\sigma = 0.02$  (a3). Middle rows (b1)–(b3): true flow field (b1) and the mean flow field from the posterior distribution estimated using the MCMC samples. In the panels below (b2) and (b3), we also show the uncertainty regions (as shaded ellipses), also from MCMC samples as given by equation (39). Bottom row (c1)–(c3): time evolution (c1) as well as the distribution (c2)–(c3) of the effective regularization parameter  $\sigma/\lambda$ , where the distributions are obtained after discarding the initial transient in the MCMC sampling process.



**Figure 4.** Bayesian optical flow computed using statistical inversion, based on image pairs (F, G) under example flow field 4, where  $\langle U(x, y), V(x, y) \rangle = \langle -\pi \sin(0.5\pi x) \cos(0.5\pi y), \pi \cos(0.5\pi x) \sin(0.5\pi y) \rangle$ . Top row (a1)–(a3): image data generated according to equation (38) for *F* (a1) and using equation (37) for *G* either without noise (a2) or with noise standard deviation  $\sigma = 0.02$  (a3). Middle rows (b1)–(b3): true flow field (b1) and the mean flow field from the posterior distribution estimated using the MCMC samples. In the panels below (b2) and (b3), we also show the uncertainty regions (as shaded ellipses), also from MCMC samples as given by equation (39). Bottom row (c1)–(c3): time evolution (c1) as well as the distribution (c2) and (c3) of the effective regularization parameter  $\sigma/\lambda$ , where the distributions are obtained after discarding the initial transient in the MCMC sampling process.



**Figure 5.** Bayesian optical flow computed using statistical inversion, based on image pairs (F, G) under example flow field 5, where  $\langle U(x, y), V(x, y) \rangle =$  $\langle -\pi \sin(\pi x) \cos(\pi y), \pi \cos(\pi x) \sin(\pi y) \rangle$ . Top row (a1)-(a3): image data generated according to equation (38) for *F* (a1) and using equation (37) for *G* either without noise (a2) or with noise standard deviation  $\sigma = 0.02$  (a3). Middle rows (b1)-(b3): true flow field (b1) and the mean flow field from the posterior distribution estimated using the MCMC samples. In the panels below (b2) and (b3), we also show the uncertainty regions (as shaded ellipses), also from MCMC samples as given by equation (39). Bottom row (c1)-(c3): time evolution (c1) as well as the distribution (c2) and (c3) of the effective regularization parameter  $\sigma/\lambda$ , where the distributions are obtained after discarding the initial transient in the MCMC sampling process.

MCMC sampling procedure, indicating convergence to a stationary distribution typically after a quick initial transient. The remaining panels (c2) and (c3) show the distribution of  $\delta/\lambda$  after discarding the initial transient, for both the case of no noise (c2) and the case with noise (c3).

We point out a few observations from the numerical experiments. First, the estimated mean optical flow compare reasonably well with the true flow in all examples of the qualitatively different optical flow fields, supporting the utility of the proposed statistical inversion approach (see panels (b1)–(b3) in all figures). Secondly, we again point out that the MCMC procedure used in our statistical inversion approach to optical flow does not require an active prior choice of the regularization parameter. The MCMC samples seems to quickly converge to a stationary distribution for the effective parameter (panel (c1) in all figures), from which the distributions of parameters and solutions can be determined. Finally, comparing to the noise free images, the estimation of optical flow becomes less accurate when noise is added. It is worth mentioning that the statistical inversion approach in fact allows us to 'predict' this difference without knowing the ground-truth optical flow, by quantifying and comparing the uncertainty of solutions (panels (b2) and (b3) in all figures). Note that although here we only show the Bayesian optical flow results based on Gaussian noise, we have also performed simulations using uniform noise and Laplace-distributed noise, and found that the results are quite similar to what is shown in figures 1–6.

#### 5.3. Simulations on real images

Finally, we conduct numerical experiments for the computation of Bayesian optical flow from using real images together with the benchmark flow fields. In each example, we consider matrix F defined by a real image from the Middlebury dataset (http://vision.middlebury.edu/stereo/, also see [26]), resized to a fixed size of 60-by-60 pixels in grayscale, with intensity normalized so that each pixel intensity is in the range [0, 1]. We consider a total of six such images, as shown in figure 6.

For each real image *F* and a given flow field  $\langle U, V \rangle$ , we generate a resulting second image *G* according to the flow equation (38), where the noise is taken to be a multivariate Gaussian distribution with zero mean and covariance matrix  $\sigma I$  with  $\sigma = 0.05$ . This choice of standard deviation ensures a non-negligible effect on the pixel intensities of the image pairs, since each  $F_{ij} \in [0, 1]$ . These 'noisy second images' are shown in each of the first column of figures 7–12. For comparison, we also consider what we call a 'true second image', denoted as  $\overline{G}$ , which is defined by the same flow equation (38) but in the absence of noise. These true second images are shown in the second column in figures 7 through 12, for each one of the image example and flow field. Thus, each baseline image *F* and flow field gives rise to a noisy image pair (*F*, *G*), and there is a total of 30 such pairs given the 6 real images and 5 flow fields.

For each noisy image pair (F, G), we use the same methodology with the same choice of priors and hyperpriors as in section 5.2 to obtain a Bayesian optical flow from sampling the posterior distribution—in practice because of the randomness of the MCMC process, we observe that sometimes the sample distribution does not 'converge' to a stationary distribution even after thousands of iterations; when this occurs we simply restart the MCMC from a different initial seed. Recall that different from a standard regularization approach which require careful choice of the regularization parameter, here such parameter is itself treated as a random variable that has its own prior which can be taken to be a 'wide' distribution in the absence of additional knowledge. From the posterior distribution, we take the mean as an estimate of the average flow field. To test the usefulness of the reconstructed flow field, we use it together with the first image F to obtain an estimated second image  $\hat{G}$  from equation (38), setting the noise term to be zero. The estimated second image  $\hat{G}$  seems to not only resemble the given noisy image G, through 12. Interestingly, the estimated  $\hat{G}$  seems to not only resemble the given noisy image G.



Figure 6. Real images used for the computations of Bayesian optical flow in section 5.3.

but even more similar to the noiseless second image  $\bar{G}$ . This observation is confirmed quantitatively, as shown in the last column of figures 7 through 12, where the values of  $\hat{G}_{ij}$  are plotted against both those of  $G_{ij}$  and of  $\bar{G}_{ij}$ . These results confirm the validity of Bayesian optical flow obtained by a statistical inversion approach, and, additional suggests that accurate reconstruction of optical flow can be potentially useful for image smoothing and denoising applications.

# 6. Discussion and conclusions

In this paper we take a statistical inversion perspective to the optical flow inference problem. From this perspective all relevant variables in an otherwise standard inverse problem are treated as random variables, and the key is to form an efficient process to construct and sample the posterior distribution by utilizing knowledge about the form of model, noise, and other prior information. From a Bayesian perspective, the various priors and the forward model combine to produce a posterior distribution describing the propagation of prior information in context of the problem. We have shown that optical flow estimates given by traditional variational approaches such as the seminal work developed by Horn and Schunck [16] can in fact be interpreted in the statistical inversion framework under particular choices of models, noise, and priors. Thus we recap that there are major advantages over the classical variational calculus approach to inverse problems where by necessity the ill-posedness is dealt with by adding an ad hoc regularity term that hopefully agrees with expected physical interpretation. From a Bayesian perspective, the ill-posedness is dealt with naturally under the statistical inversion framework by restating as a well-posed extended problem in a larger space of probability distributions [11, 17]. This therefore naturally removes a key difficulty of having to choose exactly an appropriate regularity parameter encountered in classical methods. Instead, in contrast to classical optical flow methods which only yield single solutions as 'point estimates',



**Figure 7.** Optical flow computed from noisy image pairs (F, G), where *F* is given by the image shown in figure 6(a). Each row corresponds to the choice of a different flow field, each of which is defined in section 5.1. From left to right in each row: noisy second image *G* obtained from the flow equation (38) using Gaussian noise with  $\sigma = 0.05$ , true second image  $\overline{G}$  from the same flow equation in the absence of noise, estimated  $\hat{G}$  based on using the mean flow field, and plots of  $\hat{G}$  against *G* versus  $\hat{G}$  against  $\overline{G}$ .



**Figure 8.** Optical flow computed from noisy image pairs (F, G), where *F* is given by the image shown in figure 6(b). Each row corresponds to the choice of a different flow field, each of which is defined in section 5.1. From left to right in each row: noisy second image *G* obtained from the flow equation (38) using Gaussian noise with  $\sigma = 0.05$ , true second image  $\overline{G}$  from the same flow equation in the absence of noise, estimated  $\hat{G}$  based on using the mean flow field, and plots of  $\hat{G}$  against *G* versus  $\hat{G}$  against  $\overline{G}$ .



**Figure 9.** Optical flow computed from noisy image pairs (F, G), where *F* is given by the image shown in figure 6(c). Each row corresponds to the choice of a different flow field, each of which is defined in section 5.1. From left to right in each row: noisy second image *G* obtained from the flow equation (38) using Gaussian noise with  $\sigma = 0.05$ , true second image  $\overline{G}$  from the same flow equation in the absence of noise, estimated  $\hat{G}$  based on using the mean flow field, and plots of  $\hat{G}$  against *G* versus  $\hat{G}$  against  $\overline{G}$ .



**Figure 10.** Optical flow computed from noisy image pairs (F, G), where *F* is given by the image shown in figure 6(d). Each row corresponds to the choice of a different flow field, each of which is defined in section 5.1. From left to right in each row: noisy second image *G* obtained from the flow equation (38) using Gaussian noise with  $\sigma = 0.05$ , true second image  $\overline{G}$  from the same flow equation in the absence of noise, estimated  $\hat{G}$  based on using the mean flow field, and plots of  $\hat{G}$  against *G* versus  $\hat{G}$  against  $\overline{G}$ .



**Figure 11.** Optical flow computed from noisy image pairs (F, G), where *F* is given by the image shown in figure 6(e). Each row corresponds to the choice of a different flow field, each of which is defined in section 5.1. From left to right in each row: noisy second image *G* obtained from the flow equation (38) using Gaussian noise with  $\sigma = 0.05$ , true second image  $\overline{G}$  from the same flow equation in the absence of noise, estimated  $\hat{G}$  based on using the mean flow field, and plots of  $\hat{G}$  against *G* versus  $\hat{G}$  against  $\overline{G}$ .



**Figure 12.** Optical flow computed from noisy image pairs (*F*, *G*), where *F* is given by the image shown in figure 6(f). Each row corresponds to the choice of a different flow field, each of which is defined in section 5.1. From left to right in each row: noisy second image *G* obtained from the flow equation (38) using Gaussian noise with  $\sigma = 0.05$ , true second image  $\bar{G}$  from the same flow equation in the absence of noise, estimated  $\hat{G}$  based on using the mean flow field, and plots of  $\hat{G}$  against *G* versus  $\hat{G}$  against  $\bar{G}$ .

the statistical inversion approach produces a distribution of points which can be sampled in terms of most appropriate estimators and also for uncertainty quantification. Specifically in the context of an optic flow problem, we expect a distribution of regularity parameters, and correspondingly a distribution of optical flow vectors at each pixel.

In this paper we focused on a statistical inversion formulation with the choice of priors inspired from the classical Horn-Schunck functional. Although such choices appear to be reasonable for the synthetic flow fields considered herein, real images and optical flow fields are much more complicated and, consequently, require additional efforts in forming the appropriate priors. Such priors can come from previous knowledge of images and flow fields from similar systems, taken under similar scenes, or from other physical measurements. In addition, different noise distributions should be considered, and these will require modifying the standard Horn-Schunck framework, which assumed rigid body motion and conservation of brightness. In particular, with regard to the regularization term, one example is to consider different *p*-norms other than just the standard 2-norm, and possibly with a kernel to weigh in additional information about the physical embedding of objects [27]. Other data fidelity terms [5, 6, 19]can be formulated to describe the physics of the underlying application, for example for fluid and oceanographic problems where a stream function, or even a quasi-static approximation to assume such physics as corriolis can be used; or in scenarios where divergence-free flows are estimated by utilizing vorticity-velocity formalism or more general data assimilation tools [7, 13, 36]. Likewise, regularity in time and multiple time steps may be appropriate [5], as these correspondingly more complex formulations nonetheless come back to a linear inverse problem tenable in the framework of this paper. Specifically within the statistical inversion framework developed in the current paper, all of these could be recast so that the data fidelity term may be incorporated and should not require significant modification of the general framework as introduced here, which we plan in future work. Likewise, other numerical differentiation and integration schemes can be used as well in place of the simple forward difference used here.

# Acknowledgments

This work was funded in part by the Simons Foundation Grant No. 318812, the Army Research Office Grant No. W911NF-16-1-0081, and the Office of Naval Research Grant No. N00014-15-1-2093.

# **ORCID** iDs

Jie Sun lo https://orcid.org/0000-0003-2553-1804

#### References

- Aubert G, Deriche R and Kornprobst P 1999 Computing optical flow via variational techniques SIAM J. Appl. Math. 60 156–82
- Bardsley J M 2012 MCMC-based image reconstruction with uncertainty quantification SIAM J. Sci. Comput. 34 A1316–32
- [3] Bardsley J M 2013 Gaussian Markov random field priors for inverse problems Inverse Problems Imaging 7 397–416
- [4] Barron J L, Fleet D J and Beauchemin S S 1994 Performance of optical flow techniques Int. J. Comput. Vis. 12 43–77
- [5] Basnayake R and Bollt E M 2014 A multi-time step method to compute optical flow with scientific priors for observations of a fluidic system *Ergodic Theory, Open Dynamics, and Coherent Structures* ed G Froyland and W Bahsoun (New York: Springer)

- [6] Basnayake R, Luttman A and Bollt E M 2013 A lagged diffusivity method for computing total variation regularized fluid flow *Contemp. Math* 586 57–64
- [7] Béréziat D and Herlin I 2011 Solving ill-posed image processing problems using data assimilation *Numer. Algorithms* 56 219–52
- [8] Cavalier L 2008 Nonparametric statistical inverse problems Inverse Problems 24 034004
- [9] Chantas G, Gkamas T and Nikou C 2014 Variational-Bayes optical flow J. Math. Imaging Vis. 50 199–213
- [10] Cohen I and Herlin I 1996 Optical flow and phase portrait methods for environmental satellite image sequences Proc. European Conf. Computer Vision (Cambridge: Springer) pp 141–50
- [11] Gelman A, Carlin J B, Stern H S, Dunson D B, Vehtari A and Rubin D B 2014 Bayesian Data Analysis 3rd edn (London: Chapman and Hall)
- [12] Golub G H and Van Loan C F 2012 Matrix Computations 4th edn (Baltimore, MD: Johns Hopkins University Press)
- [13] Herlin I and Béréziat D 2012 Divergence-free motion estimation European Conf. on Computer Vision pp 15–27
- [14] Higdon D 2002 Space and space-time modeling using process convolutions *Quantitative Methods for Current Environmental Issues* ed C W Anderson, V Barnett, P C Chatwin and A H El-Shaarawi (London: Springer)
- [15] Hansen P C 1998 Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion (Philadelphia, PA: SIAM)
- [16] Horn B K P and Schunck B G 1981 Determining optical flow Artif. Intell. 17 185-203
- [17] Kaipio J and Somersalo E 2005 Statistical and Computational Inverse Problems (New York: Springer)
- [18] Lebrun M, Buades A and Morel J M 2013 A nonlocal Bayesian image denoising algorithm SIAM J. Imaging Sci. 6 1665–88
- [19] Luttman A, Bollt E, Basnayake R and Kramer S 2011 A stream function approach to optical flow with applications to fluid transport dynamics *Proc. Appl. Math. Mech.* 11 855–6
- [20] Luttman A, Bollt E M, Basnayake R, Kramer S and Tufillaro N B 2013 A framework for estimating potential fluid flow from digital imagery *Chaos* 23 033134
- [21] O'Sullivan F A 1986 A statistical perspective on ill-posed inverse problems Stat. Sci. 1 502–18
- [22] Saad Y 2003 Iterative Methods for Sparse Linear Systems 2nd edn (Philadelphia, PA: SIAM)
- [23] Seeger M W and Nickisch H 2011 Large scale Bayesian inference and experimental design for sparse linear models SIAM J. Imaging Sci. 4 166–99
- [24] Siotani M 1964 Tolerance regions for a multivariate normal population Ann. Inst. Stat. Math. 16 135–53
- [25] Souhila K and Karim A 2007 Optical flow based robot obstacle avoidance Int. J. Adv. Robot. Syst. (https://doi.org/10.5772/5715)
- [26] Sun D, Roth S and Black M J 2010 Secrets of optical flow estimation and their principles Proc. IEEE Conf. Computer Vision and Pattern Recognition
- [27] Sun J, Quevedo F J and Bollt E 2017 Data fusion reconstruction of spatially embedded complex networks (arXiv:1707.00731)
- [28] Szeliski R 2011 Computer Vision (New York: Springer)
- [29] Le Tarnec L, Destrempes F, Cloutier G and Garcia D 2014 A proof of convergence of the Horn– Schunck optical flow algorithm in arbitrary dimension SIAM J. Imaging Sci. 7 277–93
- [30] Thompson A M, Brown J C, Kay J W and Titterington D M 1991 A study of methods of choosing the smoothing parameter in image restoration by regularization *IEEE Trans. Pattern Anal. Mach. Intell.* 13 326–39
- [31] Tikhonov A N 1963 Regularization of incorrectly posed problems Sov. Math. Dokl. 4 1624-7
- [32] Tikhonov A N and Arsenin V 1977 Solutions of Ill-Posed Problems (New York: Wiley)
- [33] Tikhonov A N, Goncharsky A V, Stepanov V V and Yagola A G 1990 Numerical Methods for the Solution of Ill-Posed Problems (Dordrecht: Kluwer)
- [34] Vogel C R 2002 Computational Methods for Inverse Problems (Philadelphia, PA: SIAM)
- [35] Weickert J and Schnörr C 2001 Variational optic flow computation with a spatio-temporal smoothness constraint J. Math. Imaging Vis. 14 245–55
- [36] Zhuk S, Tchrakian T, Akhriev A, Lu S and Hamann H 2017 Where computer vision can aid physics: dynamic cloud motion forecasting from satellite images Proc. of 2017 IEEE 56th Conf. on Decision and Control (CDC)