A sparse optical flow inspired method for 3D velocimetry

George Lu[1*], Adam Steinberg[2] and Masayuki Yano[1]

[1*]Institute for Aerospace Studies, University of Toronto, 4925
Dufferin St, Toronto, M3H 5T6, Ontario, Canada.
[2]Daniel Guggenheim School of Aerospace Engineering, Georgia
Institute of Technology, 265 North Ave., NW, Atlanta, 30313,
Georgia, United States.

*Corresponding author(s). E-mail(s): g.lu@mail.utoronto.ca;
Contributing authors: adam.steinberg@gatech.edu;
masa.yano@utoronto.ca;

Abstract

We introduce a three-dimensional three-component (3D3C) particle-based velocimetry method that expands the methodology of optical flow to three dimensions. The proposed scheme, sparse particle flow velocimetry (SPFV), uses a sparse representation of intensity fields with kernel functions to facilitate efficient computation in 3D. In addition, to provide robust performance for the large particle displacements seen in images, the sparse representation is combined with a multi-resolution optimization scheme based on an energy functional derived from the displaced frame difference (DFD) equation; however, this formulation is not reliant on linearized coarse-to-fine warping schemes to enable estimations of large displacements at the cost of potentially freezing large scale velocity features. Performance of SPFV is evaluated in terms of accuracy and spatial resolution, using synthetic particle images from a direct numerical simulation of isotropic turbulence. SPFV yields lower errors than tomographic PIV (T-PIV) and is capable of resolving finer scale features, even for large particle displacements and in the presence of artificial tomographic reconstruction artifacts. The method is also validated on experimental images of reacting flows and shows good agreement with T-PIV results.

1 Introduction

Accurate and high-resolution measurements of 3D velocity fields are desirable to better understand and characterize numerous fluid dynamics processes. The past decade has seen the maturation of several relevant techniques – most notably tomographic particle image velocimetry (T-PIV) [1] and shake-the-box (STB) [2] – that can provide high-quality 3D fluid velocity fields. Of course, as with any measurement technique, these methods have some inherent constraints. For T-PIV, tomographic reconstruction of 3D particle fields from a limited number of views is inherently ill-posed, resulting in artefacts called "ghost particles" that reduce the correlation between true particle images in successive interrogation boxes (IB) [3]. Furthermore, the need for IBs that contain many particles is inherently limiting to the spatial resolution and accuracy of the technique; many particles are required to obtain one vector, and velocity gradients at a smaller scale than the IB cannot be resolved. In contrast, STB, an advancement on the commonly employed particle tracking velocimetry (PTV) method, can nearly eliminate ghost particles and provide a particle-by-particle velocity field by tracking individual particles over numerous successive frames. However, this requires relatively expensive measurement hardware and that particles remain in the measurement domain for sufficiently long, which may not be possible for high-speed or highly three-dimensional flows. This requirement for a relatively lengthy tracking window applies to other advancements on PTV, such as the application of deep-learning methods [4] to particle tracking.

The application of optical flow motion estimation (OFME) techniques to analyze fluid flows has garnered increasing attention in recent years as an alternative or supplement to correlation based methods, *viz.* PIV [5–13]. Such techniques are capable of densely estimating the velocity field from two successive particle images, potentially yielding fields with finer details than PIV using only two snapshots [5]. In particular, the resolution of optical flow techniques is not bounded by the minimum size of an IB required to obtain strong correlation.

However, OFME-based fluid velocimetry has previously been constrained to two-dimensional measurements; there are numerous issues that must be addressed if an OFME-inspired formulation is extended to three dimensions. Firstly, optical flow methods generally have a substantially greater computational cost than their statistical counterparts (equations are solved at every pixel of the image) [14], which would increase by more than an order of magnitude with the addition of a third spatial dimension. A formulation intended for three dimensions must therefore be prepared to handle a number of voxels on the order of 10^8 or higher, which results in a large memory footprint for many OFME implementations [15]. Secondly, the aforementioned ghost particles present a new source of noise compared to 2D optical flow.

Furthermore, optical flow formulations commonly assume small displacements in the images analyzed. As particle images displacements for fluid velocimetry typically are rather large (several pixels), multi-scaling algorithms – e.g. employing successively smaller scales of image filtering with down-sampling – are employed in their analysis. However, a common issue in multi-resolution frameworks is the degree of image filtering/downsampling required, with noise being more pronounced at the coarsest scales [16]. In certain strategies, the final estimate is also dependent on the choice of prior resolution scales used in the processing because successive velocity estimates are linearized about estimates at coarser scales [6].

In this paper, we propose a computationally efficient two-frame 3D velocimetry method inspired by OFME that is robust in the presence of noise and large displacements. The proposed method, which is specialized for particle images, features a sparse representation of image intensity, so as to avoid using limited computational resources on regions with no particles. In tandem with the sparse representation of intensity field, the proposed method, sparse particle flow velocimetry (SPFV), incorporates an optimization procedure using an adjoint method that facilitates robust estimation performance in the presence of noise and large displacements.

The remainder of the paper is organized as follows. In section 2, we review concepts in optical flow motion estimation relevant to the proposed method. In section 3, we present the formulation of SPFV. In sections 4 and 5, we assess the accuracy, resolution, and robustness of SPFV, using particle tomogram pairs taken from direct numerical simulation (DNS) [17, 18] and previous T-PIV experiments [19].

2 Optical flow motion estimation

Optical flow is a framework for estimating the motion of objects in a two-dimensional visual scene by estimating the displacements of brightness patterns in the images. It is assumed that the image brightness (discretized as pixel intensities) is a conserved quantity across images of the scene, replete with its own transport equation [20],

$$I(\mathbf{x},t) = I(\mathbf{x}+\mathbf{d}(\mathbf{x}), t+\Delta t), \tag{1}$$

where $I(\mathbf{x},t)$ is image intensity at time t as a function of spatial location $\mathbf{x} = (x,y)^T$, and $\mathbf{d}(\mathbf{x}) = (d_x(\mathbf{x}), d_y(\mathbf{x}))^T$ is the displacement field in the image plane. The objective of OFME is to estimate the optical flow field $\mathbf{u} = \mathbf{d}/\Delta t$, corresponding to the velocity of objects in the scene, which can be done using a variety of formulations. In this work, we characterize differing OFME formulations by (a) the choice of objective function and (b) the approach to regularize the solution.

2.1 Solution by energy functional

The most common way of approaching the problem is with a variational approach [20–24]. It is often assumed that the displacement $\mathbf{d}$ is small, yielding the linearized formulation for divergence-free flows

$$\nabla I(\mathbf{x},t)\cdot\mathbf{u}+\partial_t I(\mathbf{x},t)=0, \tag{2}$$

which is also known as the gradient conservation equation. The prototypical variational scheme [25] seeks to find the velocity $\mathbf{u}$ that is the minimizer of some cost functional J. This cost functional includes a data term that quantifies the consistency of the optical flow field estimate $\mathbf{u}$ with the photometric information [10]. Assuming the gradient conservation assumption holds, the associated data term has the form:

$$J_D(\mathbf{u})=\int_\Omega\left[\partial_t I(\mathbf{x},t)+\mathbf{u}\cdot\nabla I(\mathbf{x},t)\right]^2\,\mathrm{d}\mathbf{x}, \tag{3}$$

where the intensity gradient with respect to time $\partial_t I$ is approximated by forward difference between image pairs, the intensity gradient with respect to position ∇I is approximated by finite difference over the pixel data of the undisplaced image, and the integral over Ω is approximated using the midpoint rule over the pixel data. For a discussion of a more general variational form of (2) that does not assume a divergence-free flow, we refer to Corpetti *et al* [26], who instead assume that $\partial_t I+\nabla\cdot(I\mathbf{u})=0$.

Alternatively, many methodologies for optical flow intended to deal with larger displacements forgo linearization (as in (2)), due to violations of the small displacements and gradient constancy assumptions. Instead, the data term is constructed based on the matching of two consecutive images [27] ($I_0(\mathbf{x}):=I(\mathbf{x},t)$ and $I_1(\mathbf{x}):=I(\mathbf{x},t+1)$), based on the grey value constancy assumption,

$$I_0(\mathbf{x})-I_1(\mathbf{x}+\mathbf{d}(\mathbf{x}))=0. \tag{4}$$

The associated associated energy functional then is

$$J_D(\mathbf{u})=\frac{1}{2}\int_\Omega\left[I_0(\mathbf{x})-I_1(\mathbf{x}+\mathbf{d}(\mathbf{x}))\right]^2\,\mathrm{d}\mathbf{x}. \tag{5}$$

In general, however, a single intensity constraint may not be sufficient to recover $\mathbf{u}$, which consists of multiple unknowns (the different velocity components of $\mathbf{u}$ for flow fields with a dimension greater than one), as the gradient constraint equation only constrains components of $\mathbf{u}$ that are normal to level sets within the image [21]. This is commonly known as the aperture problem. There are a number of schemes that have been developed to address this, and the ones typically employed in fluid velocimetry are regularization by a global smoothing assumption. A classical example of such a regularization term is

$$J_R(\mathbf{u})=\int_\Omega\lambda\left\|\nabla\mathbf{u}\right\|^2\,d\mathbf{x}, \tag{6}$$

proposed by Horn and Schunck [25], which penalizes gradients in order to enforce smoothness in the dense pixel-wise velocity estimate. The cost functional J is then the sum of the terms J_D and J_R, such that $\mathbf{u}$ is the minimizer of $J(\mathbf{u}) = J_D(\mathbf{u}) + J_R(\mathbf{u})$. Such regularizers are of great utility in applying optical flow to images intended for PIV, due to the presence of large regions with nearly uniform low intensity. However, problems can arise from the setting of user-defined regularization parameters, which is done *ad-hoc* and may over-smooth turbulent fields [28].

An alternative to appending explicit regularization terms such as (6) to the cost functional is implicit regularization. Implicit regularization schemes reduce the number of unknowns to solve by adopting a parametric representation of the optical flow field $\mathbf{u}(\mathbf{x}) = \mathbf{\Phi}(\mathbf{x}, \mathbf{c})$, typically some form of basis representation $\mathbf{u}(\mathbf{x}) = \sum_i \phi_i(\mathbf{x})c_i$. The estimation problem is then posed in terms of the parameters $\mathbf{c} = (c_1, ..., c_n)^T$. Commonly employed are piecewise polynomials and wavelet representations of velocity [6].

2.2 Multi-scale estimation

Many implementations of OFME incorporate a coarse-to-fine or multi-scale estimation structure [7, 29, 30]. This helps address temporal aliasing and the small displacement assumption, which does not often hold. Image warping techniques allow for large scale flow features to be first resolved before resolving smaller scale features in a hierarchical image "pyramid" [29]; the original image is blurred using a Gaussian filter and then down-sampled such that only lower spatial frequencies are resolved and pixel-wise displacements are artificially reduced. The motion problem is then solved for this coarse image, the next finest image is then warped about this solution, and the process is repeated (applying successive linearizations about the previous estimate) until an estimate for the motion at the finest scale (the original image) is obtained [30, 31].

However, this strategy has two notable weaknesses, arising from the approximation of the original functional by a series of coarse data terms and by the linearization of the model about a previous estimate. First, estimates at a given scale often freeze features resolved at coarser scales [6, 32], sometimes causing the estimate to depend on the choice of scales. Secondly, noise is more pronounced at the lowest scales [16], which is especially problematic when large displacements demand a high degree of downsampling.

Nonetheless, optical flow methods have shown good performance in comparison to PIV when a multi-scale estimation structure is used to deal with the large displacements in PIV images. Wavelet formulations in particular have seen success in estimations of fluid flows in recent years when combined with the above warping strategy using either implicit [6] or explicit regularization [7, 8]. Such formulations provide higher resolution than PIV when applied to particle images, while remaining accurate for large displacements up to 5 pixels [6, 7], or 10 pixels in hybrid schemes [7]. Recent advances in leveraging

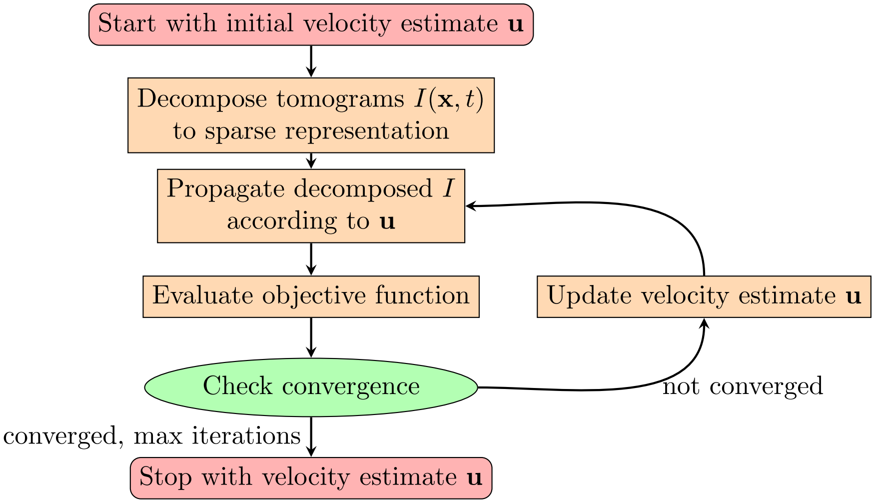

Fig. 1: Overview of procedure for velocity estimation

neural network-enhanced optical flow methods also continue to make use of multi-scale procedures [11–13].

3 3D formulation

In this section, we present the overall structure and operations of SPFV in solving for the displacement field associated with a tomogram pair; the structure is shown graphically in Fig. 1. Note that Fig. 1 describes the velocity estimation at a given scale; implementation in the context of hierarchical multi-scale estimation is described in section 3.4. The tomograms are first decomposed to a sparse representation of the 3D particle intensity distribution. The distribution is then propagated from its original state by an initial guess for $\mathbf{u}$. This initial guess may either be uniformly zero (or a slight perturbation about zero) or generated by the multi-scale estimation process described in section 3.4. An objective function associated with the particle intensity distribution is evaluated. If a convergence criteria is not met, the velocity estimate $\mathbf{u}$ is then updated based on the objective function gradient computed using an adjoint method. These steps are repeated until either the convergence criteria is met or until a maximum number of iterations has been reached.

3.1 Intensity field representation

A typical 3D intensity field reconstruction in a T-PIV experiment will be highly sparse. Evaluating intensity displacements at all N_v voxels in the reconstruction, where N_v can be upwards of 10^8, has extreme requirements in computational cost and memory, while not yielding any new information. We outline here a sparse representation of the intensities using kernel functions, so that tomograms can be processed efficiently and the distribution can be efficiently propagated forwards and backwards in time.

We reconstruct the intensity measurement at a time t using a linear combination of N_p kernel functions centered at voxel centers X^v. That is, kernel centers $X^{p,t} := \{\mathbf{x}_1^{p,t}, ..., \mathbf{x}_{N_p}^{p,t}\}$ are voxel centers $X^v := \{\mathbf{x}_1^v, ..., \mathbf{x}_{N_v}^v\}$ with an intensity measurement above some user-specified threshold, where $\mathbf{x}_1^{p,t}, ..., \mathbf{x}_{N_p}^{p,t}$ are the coordinates of the N_p kernel centers and $\mathbf{x}_1^v, ..., \mathbf{x}_{N_v}^v$ are the coordinates of the N_v voxel centers. Then,

$$I^{\mathrm{obs},t} \approx \sum_{n=1}^{N_p} g(X^v;\ \mathbf{x}_n^{p,t}, \sigma) s_n^t, \tag{7}$$

where $g(X^v;\ \mathbf{x}_n^{p,t}, \sigma)$ is a single non-scaled kernel function centered at $\mathbf{x}_n^{p,t}$ evaluated over X^v at time t, given some width parameter σ if necessary for the chosen function. The coefficients $\mathbf{s}^t = \{s_1^t, ..., s_{N_p}^t\}$ assign a scale factor corresponding to image intensity to each kernel function, and are found to minimize the difference between the reconstruction and original image $I^{\mathrm{obs},t}$ in a least squares sense.

Two choices of kernel are considered in our work, selected for approximation ability and ease of implementation. The first is Gaussians, which provide a generally good approximation of particle intensity profiles [33] and have the expression

$$g(\mathbf{x}^v;\ \mathbf{x}_n^{p,t}, \sigma) = \exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}^v - \mathbf{x}_n^{p,t}) \cdot (\mathbf{x}^v - \mathbf{x}_n^{p,t})\right), \tag{8}$$

where σ is the Gaussian width. The second is a $\mathrm{d}x \times \mathrm{d}y \times \mathrm{d}z$ uniform window (or box) kernel that is analogous to direct propagation of voxel intensities of existing optical flow methods, but adapted to our sparse formulation. The initial reconstruction is then simply the original image but with voxels below the intensity threshold blacked out, as no kernels are used to represent them. When the kernels are propagated by a voxel fraction, its intensity contribution to surrounding voxels of the new intensity field are determined by the fractional overlap of the kernel with the surrounding voxels, preserving the total intensity. More formally, the kernel is evaluated at each voxel position using the expression

$$g(\mathbf{x}^v;\ \mathbf{x}_n^{p,t}, \sigma) = \prod_{i=1}^{N_d} \left(\min\left\{(\sigma - |x_{n,i}^{p,t} - x_i^v|)/\sigma,\ 0\right\}\right), \tag{9}$$

where $x_{n,i}^{p,t}$ and x_i^v are the ith components of $\mathbf{x}_n^{p,t}$ and $\mathbf{x}^v$ respectively, and width parameter $\sigma = \mathrm{d}x$. The number of particles can be as low as 10^{-4} times less than the number of voxels, although typically multiple kernels will be used

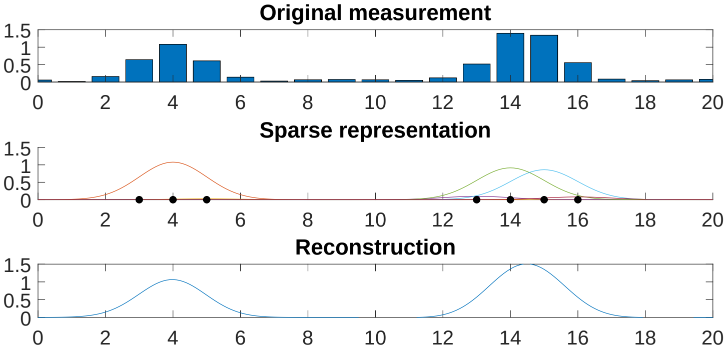

Fig. 2: Example reconstruction of 1-D intensity field using Gaussian kernels centered at voxel centers

to represent a single particle. An illustration of such a reconstruction using Gaussian kernels in 1D is shown in Fig. 2.

3.2 Velocity field representation

We consider a C^1-continuoue piecewise tricubic representation of the velocity field. Given a nodal basis $\{\phi_j\}_{j=1}^N$, the velocity representation is given by

$$\mathbf{u}(\mathbf{x};\ \mathbf{c}_{\{N\}}) = \sum_{j=1}^{N} \phi^j(\mathbf{x}) c_j, \tag{10}$$

where $\mathbf{c}_{\{N\}} = \{c_1, \ldots, c_N\}$ are the nodal basis coefficients. We implicitly regularize the solution by choosing an N that is smaller than the number of particles (and much smaller than the number of voxels N_v). In practice, we have found N in the range of $N = N_{\text{particle}}/5$ to $N = N_{\text{particle}}$ works well; $N = N_{\text{particle}}$ provides maximum resolution for clean tomograms, and $N = N_{\text{particle}}/5$ is a conservative choice that rarely develops instability even for noisy tomograms. An implicit scheme is chosen in order to avoid the *ad-hoc* tuning of explicit regularization parameters necessary for a global smoothing scheme. The velocity field is evaluated using tricubic interpolation within cuboid elements that tessellate the domain, delineated by surrounding nodes. This representation is the minimum order that allows for global C^1 continuity [34], which is necessary for the direct calculations of gradient-dependent fields relevant in fluid mechanics, such as vorticity. We efficiently evaluate the cubic function within each element using the methodology outlined in [34].

3.3 Objective function minimization

In anticipation of large displacements within analyzed tomograms and in order to avoid the evaluation of noise-sensitive finite difference terms, we assume the integrated form of the optical flow constraint equation: $I(\mathbf{x} + \mathbf{d}(\mathbf{x}, t_1), t_1) =$

$I(\mathbf{x}, t_0)$. We combine the intensity field representation in section 3.1 and the velocity representation in section 3.2 for the expression of the predicted intensity field for a given velocity field:

$$I^{\text{pred}}(\mathbf{x}, t_1;\ \mathbf{u}) = \sum_{n=1}^{N_p} g(\mathbf{x};\ \mathbf{x}_n^p(t_1;\ \mathbf{u});\ \sigma) s_n^t, \tag{11}$$

where $\mathbf{x}_n^p(t_1;\ \mathbf{u}) = \mathbf{x}_n^p(t_0) + \mathbf{u}(\mathbf{x}_n^p(t_0);\ \mathbf{c}_{\{N\}})\Delta t$, $n = 1, ..., N_p$. Then, given a pair of voxel-wise intensity measurements $I^{\text{obs},0}, I^{\text{obs},1} \in \mathbb{R}^{N_v}$, we seek the velocity field $\mathbf{u}$ that minimizes the difference between a predicted intensity (based on the initial intensity) and the measured subsequent intensity distribution:

$$\mathbf{u} = \underset{\mathbf{w}}{\text{argmin}}\, \frac{1}{2} \int_\Omega (I^{\text{obs},1} - I^{\text{pred}}(\cdot, t_1;\ \mathbf{w}))^2 \, dx. \tag{12}$$

Specifically, given the velocity representation (10), the velocity is found as the minimizer of the following cost functional, discretized over all voxels:

$$J(\cdot, \mathbf{c}_{\{N\}}) = \frac{1}{2} \sum_{v=1}^{N_v} (I^{\text{obs},1}(X_v) - I^{\text{pred}}(X_v, t_1;\ \mathbf{u}(\cdot, \mathbf{c}_{\{N\}})))^2. \tag{13}$$

We also consider an additional strategy: simultaneously propagating a reconstruction of $I^{\text{obs},1}$ backwards in time by the same estimated velocity field and evenly weighing (13) by the mismatch between the back-propagated reconstruction and $I^{\text{obs},0}$. The incorporation of both original images into the objective allows for increased robustness, although at double the original computational cost. The expression for this modified cost function is

$$J_{\text{double}}(\cdot, \mathbf{c}_{\{N\}}) = J_{\text{fwd}}(\cdot, \mathbf{c}_{\{N\}}) + J_{\text{back}}(\cdot, \mathbf{c}_{\{N\}}), \tag{14}$$

where the expression for $J_{\text{fwd}}(\cdot, c_{\{N\}})$ is identical to the expression in (13), and

$$J_{\text{back}}(\cdot, \mathbf{c}_{\{N\}}) = \frac{1}{2} \sum_{v=1}^{N_v} (I^{\text{obs},0}(X_v) - I^{\text{back}}(X_v, t_0;\ \mathbf{u}(\cdot, \mathbf{c}_{\{N\}})))^2, \tag{15}$$

where $I^{\text{back}}(X_v, t_0;\ \mathbf{u}) = \sum_{n=1}^{N_p} g(X_v;\ \mathbf{x}_n^p(t_0;\ \mathbf{u});\ \sigma) s_n^t$ for $\mathbf{x}_n^p(t_0;\ \mathbf{u}) = \mathbf{x}_n^p(t_1) - \mathbf{u}(\mathbf{x}_n^p(t_1);\ \mathbf{c}_{\{N\}})\Delta t$.

We minimize (13) (or (14) for double propagation) using the unconstrained L-BFGS algorithm [35]. The gradient of (13) with respect to the velocity coefficients $\{c_i\}_{i=1}^N$ is

$$\frac{\partial J_{\mathrm{fwd}}}{\partial \mathbf{c}_{\{N\}}} = -\underbrace{\left(I^{\mathrm{obs},1} - I^{\mathrm{pred}}(X_v, t_1;\, \mathbf{c}_{\{N\}})\right)}_{a_1} \underbrace{\frac{\partial I^{\mathrm{pred}}(X_v, t_1;\, \mathbf{c}_{\{N\}})}{\partial \mathbf{d}^p}}_{a_2} \underbrace{\frac{\partial \mathbf{d}^p}{\partial \mathbf{u}^p}}_{a_3} \underbrace{\frac{\partial \mathbf{u}^p}{\partial \mathbf{c}_{\{N\}}}}_{a_4},$$

where a_1 is the difference between observed and predicted intensities, a_2 is the voxel-wise derivative of intensity with respect to a component of a kernel's displacement, a_3 is a multiplication by Δt, and a_4 is the derivative of kernel velocity with respect to the nodal basis coefficients. We make brief remarks on the implementation of this method. First, the terms a_1 and a_2 can be evaluated in conjunction as an $\mathbb{R}^{n_d N_p}$ gradient row vector by evaluating the respective gradients for each component of kernel displacement as a convolution of the two terms within a compact neighbourhood of each respective kernel. Second, the matrix computed in a_4 is highly sparse, as the velocity of each kernel is computed using $64 \ll N$ local node values in a tricubic interpolation scheme [34]. For double propagation, the gradient for (15) is computed in the same manner.

3.4 Multi-scale estimation

While the proposed formulation lacks, in principle, many of the shortcomings in resolving flow features in the presence of large displacements common to linearized optical flow formulations, large displacements are nonetheless a challenge to any motion estimation framework. Additionally, when noise is introduced into the particle images, greater consideration must be given to the possibility that the nonlinear optimizer converges to a solution with spurious velocity features. In order to address these concerns, we employ a two-step heirarchical estimation procedure to provide a good initial estimate for the algorithm, increasing its robustness. The first step involves blurring and downsampling the intensity fields in order to resolve large scale motions, and in the second step we refine the fidelity of the velocity representation.

In the first step, the intensity fields $I^{\mathrm{obs},0}$ and $I^{\mathrm{obs},1}$ are blurred by applying a Gaussian filter with width l, in order to ensure a sufficient overlap between corresponding particle images. We then generate downsampled fields $I_l^{\mathrm{obs},0}$ and $I_l^{\mathrm{obs},1}$ by sampling the blurred images at every lth voxel in each dimension, which we reconstruct using kernels with inflated width $l\sigma$, where σ is the width parameter to be used to reconstruct the original images. Next, we solve the minimization problem (12) for the downsampled intensity fields $I_l^{\mathrm{obs},0}$ and $I_l^{\mathrm{obs},1}$ with an N-degree of freedom (dof) velocity representation (10). The resulting coefficient vector $\mathbf{c}_{\{N\}}$ for the flow field of the downsampled intensity field is then used as an initial estimate, as the minimization problem is repeated using the original, non-filtered images to refine the solution for $\mathbf{c}_{\{N\}}$.

In the second step, the velocity solution space is refined for progressively larger values of N. Given $N' > N$, we solve the minimization problem (12) to obtain a velocity field estimate with higher resolution, associated with basis coefficients $\mathbf{c}_{\{N'\}}$. In order to ensure that this estimate does not develop spurious features, an initial estimate for the nodal basis coefficients is computed by setting them to represent the lower-fidelity velocity field associated with $\mathbf{c}_{\{N\}}$.

Although this procedure is quite similar to the image pyramid type multi-resolution schemes commonly employed in OFME, discussed in section 2.2, there are a couple important distinctions. The first is that, as the formulation is inherently suited to large displacements, the degree of downsampling necessary is lower than typically used in optical flow type schemes. Furthermore, image warping schemes are not necessary for the initialization operation. In particular, unlike linearized estimates that add finer scales to coarser estimates, the purpose of this procedure is purely to generate a good initial estimate, and the final estimate is otherwise independent of the choice of scales.

4 Validation with Synthetic Data

4.1 Description of reference data

The SPFV framework is designed with reconstructed three-dimensional particle image intensity distributions in mind, such as those reconstructed by the tomographic reconstruction process of T-PIV. In this section, we evaluate its performance using synthetic tomogram pairs intended to replicate the settings of a typical T-PIV experiment. The datasets used consist of $256 \times 256 \times 256$ voxel tomograms, where the domain is randomly seeded with synthetic particles.

Reference velocity fields are retrieved from the DNS turbulent flow database hosted by John Hopkins University (JHU) [17, 18]. Studies are conducted over 30 snapshots of the homogenous isotropic turbulence (HIT) solution, with a Taylor-scale Reynolds number of $\mathrm{Re}_\lambda = 418$. A Runge-Kutta method is used to propagate synthetic particles between each pair of intensity fields. The flow is interpolated so that two viscous length scales are resolved for every 32 voxels in each dimension. Due to the homogeneity and isotropy of the data, this provides sufficient statistically independent spatial information in the data. Further details regarding the selection and interpolation of snapshots are provided in appendix B.

The synthetic particles used for the study in section 4.2 have a scattering intensity profile approximated by a Gaussian with a 4 vx particle diameter, and particles are seeded with a density of 3.7×10^{-4} particles per voxel; the same particle characteristics are used for all DNS studies below unless otherwise noted. The particles are propagated for a normalized time step of $\frac{\Delta t \|u\|_{\max}}{\mathrm{d}x} = 6.9$ vx, where $\|u\|_{\max}$ is the maximum velocity within the domain, calculated as an average over the maximum velocities of all snapshots, and $\mathrm{d}x$ is the spacing between voxels.

Two different T-PIV velocity estimates are used for comparing accuracy and resolution of the proposed technique. The first "T-PIV" estimate is an idealized case generated by averaging the reference velocity over each IB. The

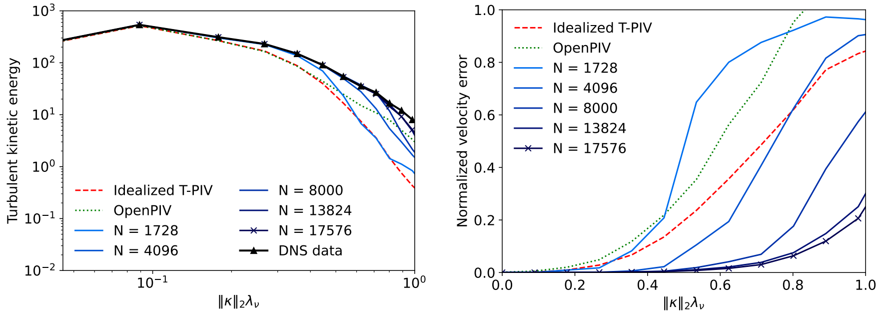

(a) **(b)**

Fig. 3: **(a)** Energy spectra and **(b)** error spectra of velocity estimates computed by SPFV and T-PIV

second is actual calculated T-PIV fields determined using an open source package (OpenPIV [36]). IBs for the T-PIV are chosen to have a width of 30 voxels with a 50% overlap in interrogation boxes. The default seeding conditions then correspond to a particle density of 10 particles per IB, a common setting in T-PIV experiments [33].

We have also conduct similar studies using the transitional boundary layer flow in the JHU's turbulent flow database to assess the method for a flow featuring strong mean shear. The results are included in appendix A.

4.2 Resolution of SPFV

In order to assess the resolution of the proposed method, we analyze the average turbulent kinetic energy spectra of the reference and estimated velocity fields, which we average across all snapshots. These spectra are approximated using discrete cosine transformation (DCT). Also of interest is the spectrum of the normalized estimation error, defined as

$$\widehat{E}_{\text{tke}} := \frac{|\widehat{f}(u_{\text{est}} - u_{\text{ref}})|^2 + |\widehat{f}(v_{\text{est}} - v_{\text{ref}})|^2 + |\widehat{f}(w_{\text{est}} - w_{\text{ref}})|^2}{|\widehat{f}(u_{\text{ref}})|^2 + |\widehat{f}(v_{\text{ref}})|^2 + |\widehat{f}(w_{\text{ref}})|^2}, \tag{16}$$

where Fourier transforms $\widehat{f}(\cdot)$ are approximated by DCT. T-PIV estimates are interpolated using cubic splines onto the $256 \times 256 \times 256$ voxel grid for the discretized tomogram domain prior to evaluation of these quantities.

Figure 3 shows the TKE spectra and error spectra (16) computed using the box-kernel intensity representation and double propagation, plotted against normalized wave number. The wave number $\|\kappa\|_2$ is non-dimensionalized by multiplying it by the Taylor microscale λ_ν. Solutions with more degrees of freedom N preserve the energy of the system at scales approaching the Taylor microscale $\|\kappa\|_2 \lambda_\nu = 1$, while idealized T-PIV diverges from the TKE spectrum of the original velocity at lower wavenumbers than even the coarsest

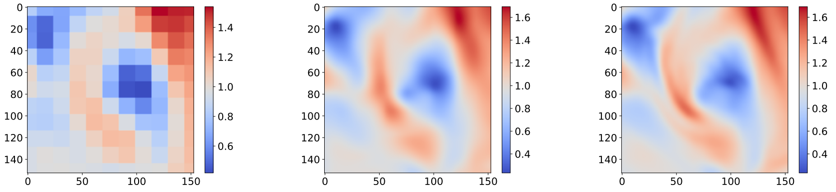

(a) **(b)** **(c)**

Fig. 4: Velocity magnitude of an initial DNS snapshot over a truncated slice at the center of the domain. Plots depict **(a)** an idealized T-PIV estimate, **(b)** an SPFV estimate ($N = 17576$), and **(c)** the reference velocity.

estimates (smallest N) computed by SPFV. However, it is possible for spurious features in the velocity solution to artificially increase the captured kinetic energy for certain wavenumbers. This is apparent in some of the coarser SPFV solutions at high wavenumbers, and is especially so for the OpenPIV solution. Therefore, it is necessary to look at the spectral error with respect to wave number (16) in order to draw further conclusions on the resolution of SPFV.

Low values of the error spectrum in Fig. 3(b) indicate more accurate solutions at a given wavenumber. Although coarse SPFV solutions can have higher error than idealized T-PIV for larger wavenumbers, the SPFV error curves for $N \geq 5832$ lie entirely under the error curve of idealized T-PIV, with monotonic improvements in accuracy as solution dof N is increased. Furthermore, SPFV solutions with $N > 1728$ have lower error than the T-PIV computed using actual cross-correlated particle images from OpenPIV.

Figure 4 shows the velocity magnitude for a truncated slice of one of the DNS snapshots used to evaluate the results in Fig. 3, along with an idealized T-PIV estimate, and the SPFV estimate for $N = 17576$. Finer features present in the DNS solution are depicted more clearly in the estimate by SPFV.

4.3 Robustness of SPFV

We now study the robustness of SPFV with respect to the following four parameters: the presence of ghost particles (section 4.3.1); particle image size (section 4.3.2); particle image density (section 4.3.3); and particle displacement magnitude (section 4.3.4).

4.3.1 Impact of ghost particles

In order to test the robustness of SPFV to reconstruction artefacts, we corrupt the synthetic tomograms by introducing a number of ghost particles equal to 20% of the number of real particles; the ghost particles are distributed randomly within the domain. Figure 5 shows the kinetic energy spectra and associated error spectra for these noisy images. Both the energy and error spectra show similar trends to the noiseless data. However, the error in the SPFV data with the highest N-values are larger for the noisy data, particularly

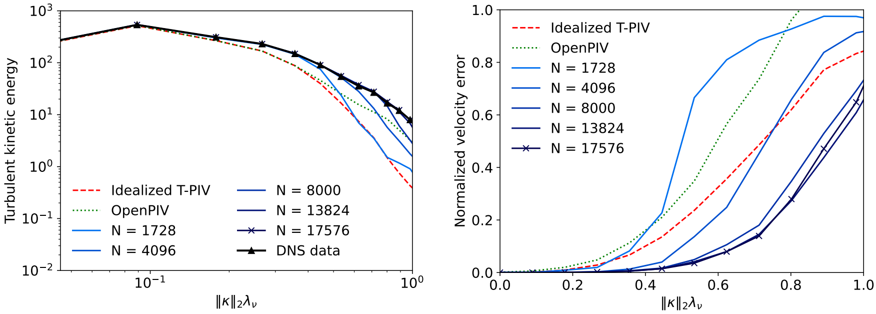

(a) **(b)**

Fig. 5: **(a)** Energy spectra and **(b)** error spectra of velocity estimates computed by SPFV and T-PIV. Analyzed tomograms are corrupted by synthetic ghost particles

at the highest wavenumbers. Furthermore, there is greater error in the $N = 17596$ spectrum (the highest N) compared to the next two highest N-values, in contrast to the trend observed in Fig. 3, though this is a relatively minor difference. Indeed, the most accurate N tested with the noisy data is $N = 13824$. Despite this, SPFV is more accurate than T-PIV for all $N \gtrsim 1728$ and more accurate than idealized T-PIV for $N \gtrsim 8000$.

4.3.2 Impact of particle image size

We now assess the impact of the size of the imaged particles on the accuracy of the velocity estimates. In these tests, the tomogram pairs are generated with different particle sizes, with other parameters otherwise held constant.

Figure 6 shows the dependence of error on the particle diameter, where error is assessed in a normalized L^2 norm

$$E_{L^2} := \frac{\|\mathbf{u}_{\text{est}} - \mathbf{u}_{\text{ref}}\|_{L^2(\Omega)}}{\|\mathbf{u}_{\text{ref}}\|_{L^2(\Omega)}}. \tag{17}$$

Compared to T-PIV, which begins to see lower quality results for large particle images, performance for the proposed method stays roughly constant after reaching a minimum error level. In particular, T-PIV sees optimal performance in the range of 2–3 voxel diameter particles, where particles are large enough so that their centers can be determined reliably, but not so large that displacements become ambiguous [37]. The use of Gaussian kernels mimics this trend, but maintains a greater degree of accuracy for larger particles. In contrast, the use of box kernels sees a nearly monotonic improvement in accuracy with increasing particle size in the range studied, consistent with larger image features being favorable to optical flow type formulations.

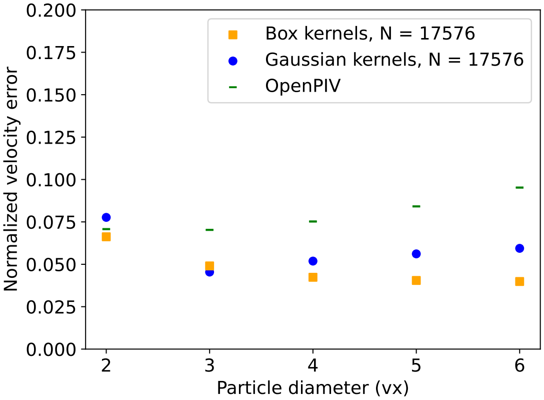

Fig. 6: Normalized L^2 error of estimation results for different particle diameters

4.3.3 Impact of particle image density

We now assess the impact of the seeding density on the accuracy of the velocity estimates. In this set of tests, the seeding densities of the tomograms are varied by creating them with different numbers of synthetic particles.

Figure 7 shows the dependence of error on the particle seeding density. The SPFV solutions observe monotonic improvements in accuracy with increasing particle density even though the solution dof remains unchanged, due to increased support for the basis projection and thus greater stability. While accuracy improves for the T-PIV result as well, it is seemingly not able to leverage the additional information as efficiently, with the limiting factor in accuracy being more in its representational power than in stability [10]. Despite the fact that SPFV sees improving performance with higher seeding densities, in practical scenarios, obtaining accurate estimates for higher particle densities is challenged by the increase in ghost particles during the tomographic reconstruction process.

4.3.4 Impact of particle displacement magnitude

One of the primary challenges optical flow type methods face is in resolving the velocity fields present in images featuring large displacements. This is addressed by multi-resolution schemes, such as that used by SPFV, which is evaluated in this section. A number of tomogram pairs are generated sequentially over a range of particle displacements, and the performance of both SPFV and T-PIV are evaluated according to the normalized L^2 error (17).

Figure 8 shows the average normalized L^2 errors of the snapshot motion estimates over a range of maximum particle displacements, for clean synthetic tomograms. SPFV solutions are refined up to $N = 17576$, though, in cases where instabilities develop in the solution, the error recorded is that of the most accurate intermediate solution.

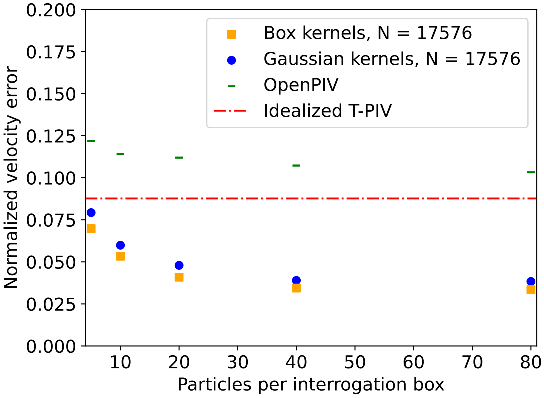

Fig. 7: Normalized L^2 error of estimation results for different seeding densities

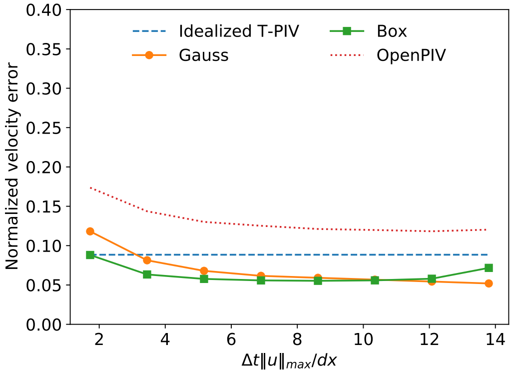

Fig. 8: Normalized L^2 error of estimation results over a range of displacements

SPFV solutions have a lower error than idealized T-PIV for displacements greater than approximately three voxels for the selected dataset. Furthermore, the trend in accuracy with respect to displacement of SPFV solutions computed using Gaussian kernels is similar to the trend of OpenPIV, with near monotonic improvements in accuracy over the range of displacements studied. The accuracy of SPFV with box kernels is reduced for larger displacements but is increased for small displacements, and estimates remain more accurate than idealized T-PIV over the entire range studied.

We also verify that the observed trends in accuracy statistically apply to smaller scale fluctuations, since the normalized L^2 error metric considers all scales in tandem. Figure 9 records the smallest normalized wave-number $\|\kappa\|\,\lambda$ at which the spectral error for the estimate meets or exceeds an error threshold of 0.25. The trends in resolution are largely reflective of the observed trends in accuracy (as recorded by the normalized L^2 error), with SPFV demonstrating

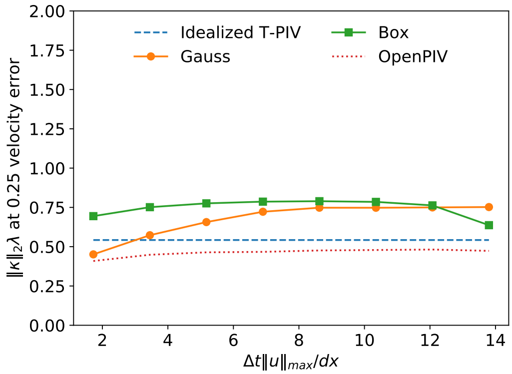

Fig. 9: Resolution of estimation results over a range of displacements. Resolution is defined as the normalized wave number at which the normalized spectral error is equal to 0.25

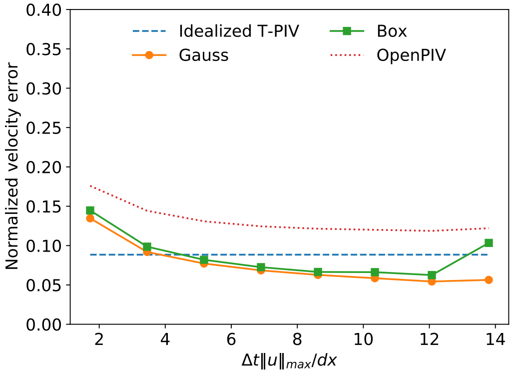

Fig. 10: Normalized L^2 error of estimation results over a range of displacements. Analyzed tomograms are corrupted by synthetic ghost particles

consistent improvement over T-PIV over much of the range of observation. In fact, through this metric, SPFV with a box-kernel representation of intensity demonstrates tangible improvements in resolution for smaller displacements.

As the effects of noise are pronounced at coarser scales, an evaluation of the robustness of the multiscale estimation procedure necessitates an evaluation of SPFV's performance with respect to particle displacement magnitude when noise is present in the data. Figure 10 shows the normalized L^2 error across a range of displacements for the previous synthetic tomograms after being corrupted by introducing a number of ghost particles equal to 20% of the number of real particles; the ghost particles are distributed randomly within the domain. The introduction of tomographic noise minimally impacts the

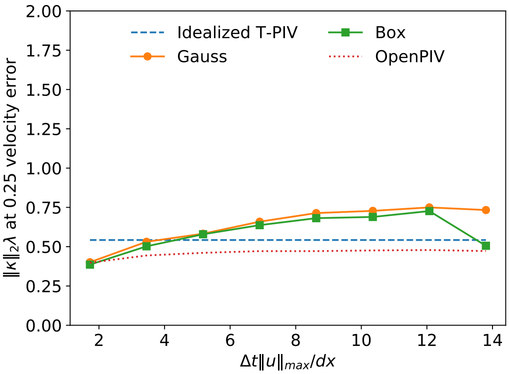

Fig. 11: Resolution of estimation results over a range of displacements. Resolution is defined as the normalized wave number at which the normalized spectral error is equal to 0.25. Analyzed tomograms are corrupted by synthetic ghost particles

accuracy trend of the OpenPIV estimates, due to its nature as a statistical method. SPFV loses greater amounts of accuracy for both small and large displacements, particularly in the box kernel representation of intensity. However, SPFV estimates remain more accurate than idealized T-PIV for displacements of greater than four voxels and, in the case of a Gaussian kernel intensity representation, improve nearly monotonically in accuracy over the range of displacements studied.

Figure 11 shows the resolution according to the proposed metric for the corrupted tomograms. Compared to the overall accuracy of the method, the introduction of tomographic noise has a larger effect on the resolution of SPFV relative to T-PIV for small displacements, suggesting a smoothing effect on field estimates in this range. The trends in resolution for SPFV otherwise reflect those of accuracy for medium and large range displacements.

5 Application to real images

In order to assess the applicability of SPFV in the presence of real-world imperfections in the input data, the method is applied to T-PIV tomogram pairs captured in experiments conducted by Kazbekov et al using LaVision DaVis 8.4. [19, 38]. The flow is a swirling turbulent air jet emitted from a nozzle into the ambient environment with a Reynolds number of $\mathrm{Re}_j = 27{,}000$ and swirl number of $S = 0.55$. The experimental configuration, which has been extensively studied in the literature, is typical of a gas turbine swirl combustor. In all cases provided, ghost particle levels were estimated at 16–25%, and particles were displaced an average of 8–11 voxels. For each case, results of SPFV and T-PIV were compared for a 450 × 450 × 285 voxel slice centered

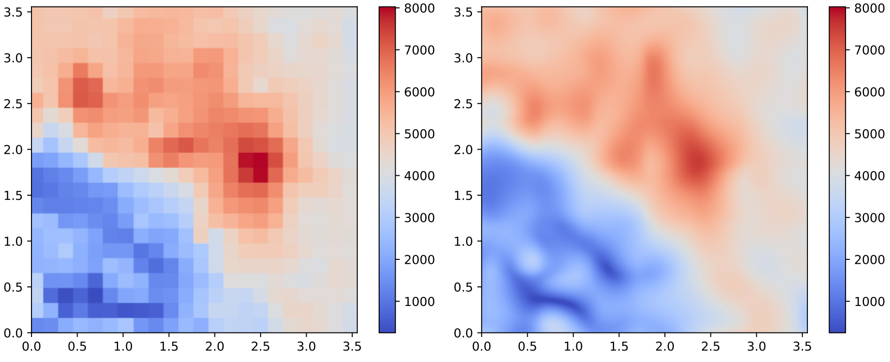

(a) **(b)**

Fig. 12: Closeup cross sections of the magnitudes of **(a)** the T-PIV velocity estimate (LaVision DaVis 8.4) and **(b)** the SPFV velocity estimate of provided experimental data

at the edge of the non-reacting jet, solved with $N = 3888$ and the box kernel representation of intensity, and without image blurring or downsampling.

No ground truth is available for this data. In lieu of this, we ensure the consistency of the velocity estimates through qualitative comparisons. Figure 12 shows the magnitude of a T-PIV velocity estimate within a cross section of the analyzed region for a select case, as well as that of a velocity estimate by SPFV. The SPFV solution corresponds well to the solution provided by T-PIV, but with a greater degree of regularity and lack of pixelation associated with the piece-wise constant T-PIV estimates within each interrogation box. Finer scales are visible in the SPFV. Importantly, the consistency with T-PIV implies that SPFV remains robust in the face of sources of noise present in real world scenarios, unaccounted for in synthetic studies.

Although the ground truth velocity is not actually known, T-PIV is generally considered to be a reliable and robust method. Therefore, we augment the visual comparison of Fig. 12 by demonstrating the statistical agreement of the proposed method with T-PIV. The T-PIV vectors are super-sampled to the original voxel grid using nearest-neighbor interpolation, and we evaluate the L^2-normalized error of our estimated vectors with respect to that of the super-sampled T-PIV estimate on a voxel-by-voxel basis. A histogram of the errors at each voxel is shown in figure 13 to display the statistics of the agreement, demonstrating the velocity field estimated by SPFV generally is consistent with TPIV in direction and magnitude.

6 Conclusions

We introduced a velocimetry method based on the principles of optical flow as extended to three dimensions with specific adaptions for particle images. The optical flow type formulation allows for the resolution of finer scale flow

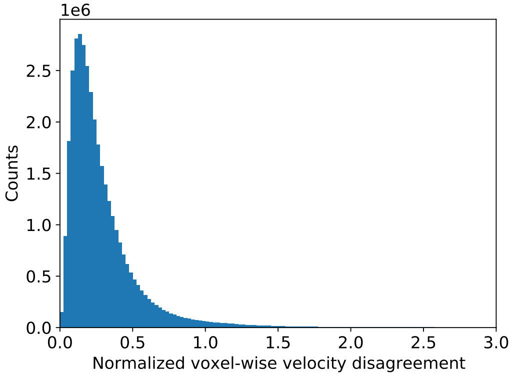

Fig. 13: Histogram of the normalized L^2 “errors” of SPFV vectors taken with respect to T-PIV (LaVision DaVis 8.4) at each voxel of the real tomograms

features than correlation-based velocimetry methods, with the sparse representation of particle images allowing for the efficient extension of the method to 3D. An objective function based on intensity displacements allows for a robust BFGS optimization process, regularized by an implicit representation of the velocity field by tricubic polynomial elements. The process does not inherently require a multi-scale resolution process, and avoids freezing coarser features during the optimization process. SPFV is found to be capable of resolving finer scale flow features relative to T-PIV over a large range of particle displacements in studies on synthetic particle images. Good agreement with T-PIV is also observed for real reacting flow experimental data. We remark here, that recent developments in the tomographic reconstruction process [39] may further improve agreement with T-PIV for real flows. Usage of other kernels, or combination of different kernels, may also yield better results. However, while SPFV allows for the principles of optical flow to be extended to 3D, existing implementations are more computationally demanding than T-PIV– details on this are included in appendix C. One issue stems from the nonlinear optimization problem formulation. Future improvements may also include leveraging developments in parallel computing with GPUs in order to provide low-level optimizations. An additional practical concern that must be addressed is the determination of an appropriate stopping condition for the refinement process. At present, N is selected conservatively if an appropriate scale cannot be determined in post-processing. This may additionally lead to the development of methods for adaptive refinement in the velocity space as well.

Acknowledgements

This work was supported by the US Air Force Office of Scientific Research under Grant FA9550-17-1-0011 (Project Monitor Dr. Chiping Li). Computations were performed on the Niagara supercomputer at the SciNet HPC

Consortium. SciNet is funded by the Canada Foundation for Innovation; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto. Experimental Tomographic PIV data was provided by Askar Kazbekov of the Georgia Institute of Technology.

A Transitional boundary layer flow

In this section, we assess the performance of SPFV using a transitional boundary layer flow. The boundary layer flow represents a real world turbulent flow similar to the T3A benchmark experiment by Roach and Brierley [40] and provides an assessment of SPFV for a flow featuring strong mean shear. As many of the observations for the HIT case also apply to this transitional boundary layer case, we present only key results.

A.1 Description of reference data

The transitional boundary layer dataset is produced from DNS of incompressible flow over a flat plate with an elliptical leading edge. We consider synthetic $256 \times 256 \times 256$ voxel tomograms constructed as in section 4.1. We select 20 snapshots at evenly spaced intervals from the dataset; the number of snapshots used is smaller than the HIT case, as we have observed smaller variation in performance compared to the HIT dataset. Appendix B provides further details regarding the selection of snapshots.

A.2 Resolution of SPFV

Figure 14 shows the magnitude of the adjusted velocity field, $\mathbf{u} - \mathbf{u}_{\text{inflow}}$, associated with the idealized T-PIV estimate, the SPFV estimate with a box kernel intensity representation, and the original reference velocity. These estimates are computed for a representative snapshot at the end of the domain chosen for demonstrational purposes, where turbulence is more developed. For this test, we propagate particles for a normalized time step $\frac{\Delta t \|u\|_{\max}}{dx} = 15$. Finer features present in the DNS solution are captured more clearly in the SPFV estimate.

A.3 Robustness of SPFV

We now apply the error analysis in section 4.3.4, which provides the most complete assessment out of the studies in section 4, to the trasitional boundary layer case. Figure 15 shows the average normalized L^2 errors of the velocity estimates over a range of maximum particle displacements, for clean synthetic tomograms without ghost particles. For each case, we select the value of N that yields the most accurate solution during the refinement process; the selected value of N varies from 1728 to X. SPFV solutions have a lower error than idealized T-PIV for medium and large displacements. Performance remains notably robust for large displacements for which OpenPIV has difficulty resolving finer scale features. We do not make the claim of superior robustness to T-PIV, as

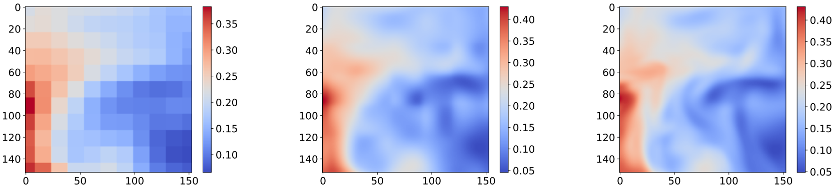

(a) **(b)** **(c)**

Fig. 14: Magnitude of the adjustded velocity field, $\mathbf{u} - \mathbf{u}_{\text{inflow}}$, for an initial DNS snapshot over a truncated slice at the center of the domain of the transitional boundary layer flow. Plots depict **(a)** an idealized T-PIV estimate, **(b)** an SPFV estimate ($N = 17576$), and **(c)** the reference velocity.

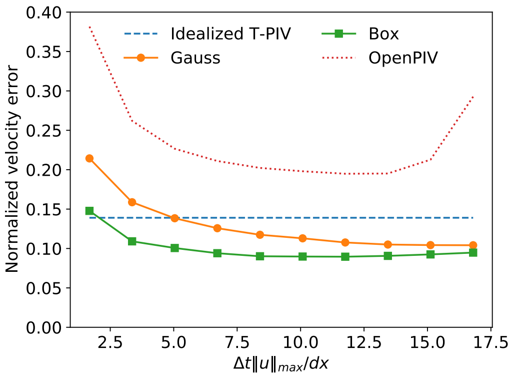

Fig. 15: Normalized L^2 error of estimation results over a range of displacements for the transitional boundary layer case.

the exploration of coarse-to-fine T-PIV strategies is beyond the scope of this paper, but nonetheless note this result.

We now apply the methods to noisy tomograms corrupted by a number of synthetic ghost particles equal to 20% the number of real particles, as in section 4.3.4. Figure 16 shows the normalized L^2 error across a range of displacements. We observe that SPFV with the box representation of intensity fields is vulnerable to tomographic reconstruction artifacts, with worse performance than idealized T-PIV for most small displacements. In contrast, the effect of noise on SPFV estimates using a Gaussian representation are minimal. SPFV estimates are accurate at high displacements despite the presence of noise. These trends are consistent with the robustness study conducted with the HIT reference velocity.

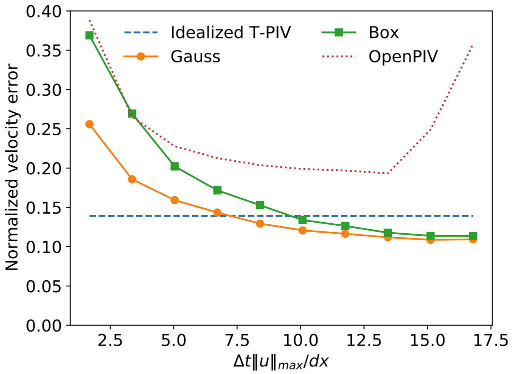

Fig. 16: Normalized L^2 error of estimation results over a range of displacements for the transitional boundary layer case. Analyzed tomograms are corrupted by synthetic ghost particles

B Synthetic image generation

For the purposes of benchmarking and reproducability, we detail the procedure used to generate reference velocity fields used to assess methods in this this work.

The HIT velocity fields in the JHU database are computed over the physical domain $[0, 2\pi]\times[0, 2\pi]\times[0, 2\pi]$ discretized into $1024\times1024\times1024$ nodes [17, 18]. A total of 5028 snapshots are available for this solution from the database. To generate reference velocity fields $\mathbf{u}_{\mathrm{HIT}}$, we extract a $47 \times 47 \times 47$ node section from the center of the original JHU domain and super-sample it to a $256\times256\times256$ voxels using cubic spline interpolation. We chose the first, 101st, and 251st snapshots, as well as subsequent snapshots at an interval of 500 snapshots (i.e. the 501st, 601st, 751st, and so on), for a total of 30 snapshots.

The transitional boundary layer dataset in the JHU database is for flow over a flat plate with an elliptical leading edge [17, 18]. The data in the database is associated with the domain over the flat region of the upper plate: physical domain $[30.2185, 1000.065] \times [0.0036, 26.4880] \times [0, 240]$ discretized over $N_x \times N_y \times N_z = 3320 \times 224 \times 2048$ nodes. To obtain reference velocity fields $\mathbf{u}_{\mathrm{BL}}$, we extract a $256 \times 256 \times 256$ voxel section encompassing physical domain $[350, 356]\times[0, 6]\times[0, 6]$, which is centered at the region where the flow transitions from a laminar to turbulent regime. We chose the 38th and 88th snapshots, as well as subsequent snapshots at an interval of 50 snapshots, for a total of 20 snapshots.

C Remark on computation cost

We now discuss the computational time observed for the particular implementation of SPFV used in this study. We caution that our implementation is not highly optimized; careful code optimization, as well as more efficient sparsification and kernel selection, may significantly reduce the computational time. Nevertheless, we note that the observed computational time of the current implementation of SPFV is substantially higher than that of T-PIV. We recall that SPFV requires the solution of the nonlinear optimization problem (12) using BFGS. For a single $256 \times 256 \times 256$ voxel tomogram pair at a typical seeding density, a single BFGS iteration on a 40-core workstation takes the order of 5 seconds, and achieving the desired optimization convergence requires $\mathcal{O}(100)$ BFGS iterations. Hence the total computational time for SPFV is, at very minimum, of order of 500 seconds or greater depending on settings of scales. In contrast, OpenPIV estimation and filtering with $30 \times 30 \times 30$ voxel interrogation boxes and a 50% overlap requires a total of approximately 40 seconds on the same machine for the same tomogram pair.

The high computational cost of SPFV can be attributed to (i) the sheer cost of evaluating the difference between two 3d intensity fields to evaluate the objective function and (ii) the large number of BFGS iterations required to solve the nonlinear optimization problem. As regard (i), we reduce the the cost of evaluating the difference between two intensity fields using the sparse representation; as discussed in section 3.1, the number of particles can be as low as 10^{-4} the number of voxels, and hence the sparsification results in a significant reduction in computational time (and memory footprint). We also note that the construction of the initial sparse representation itself using (7) is typically less than 5% of the overall cost. As regard (ii), while computationally expensive, this nonlinear formulation is arguably essential to provide the resolution and accuracy in the presence of larger displacement. This work demonstrates promising accuracy and robustness of SPFV, which motivate future work to further reduce the computational cost.

References

[1] Scarano, F.: Tomographic PIV: principles and practice. Measurement Science and Technology **24**(1), 012001 (2012). https://doi.org/10.1088/0957-0233/24/1/012001

[2] Schanz, D., Gesemann, S., Schröder, A.: Shake-The-Box: Lagrangian particle tracking at high particle image densities. Experiments in Fluids **57**(5), 70 (2016). https://doi.org/10.1007/s00348-016-2157-1

[3] Elsinga, G.E., Westerweel, J., Scarano, F., Novara, M.: On the velocity of ghost particles and the bias errors in Tomographic-PIV. Experiments in Fluids **50**(4), 825–838 (2011). https://doi.org/10.1007/s00348-010-0930-0

[4] Mallery, K., Shao, S., Hong, J.: Dense particle tracking using a learned predictive model. Experiments in Fluids **61**(10), 223 (2020). https://doi.org/10.1007/s00348-020-03061-y

[5] Chen, X., Zillé, P., Shao, L., Corpetti, T.: Optical flow for incompressible turbulence motion estimation. Experiments in Fluids **56**(1), 8 (2015). https://doi.org/10.1007/s00348-014-1874-6

[6] Dérian, P., Héas, P., Herzet, C., Mémin, E.: Wavelets and optical flow motion estimation. Numerical Mathematics: a Journal of Chinese Universities **6**, 116–137 (2013). https://doi.org/10.4208/nmtma.2013.mssvm07

[7] Schmidt, B., Sutton, J.: High-resolution velocimetry from tracer particle fields using a wavelet-based optical flow method. Experiments in Fluids **60**, 37 (2019). https://doi.org/10.1007/s00348-019-2685-6

[8] Schmidt, B.E., Sutton, J.A.: Improvements in the accuracy of wavelet-based optical flow velocimetry (wOFV) using an efficient and physically based implementation of velocity regularization. Experiments in Fluids **61**(2), 32 (2020). https://doi.org/10.1007/s00348-019-2869-0

[9] Liu, T.: OpenOpticalFlow: An open source program for extraction of velocity fields from flow visualization images. Journal of Open Research Software **5**, 29 (2017). https://doi.org/10.5334/jors.168

[10] Kumashiro, K., Steinberg, A.M., Yano, M.: A functional error analysis of differential optical flow methods. Experiments in Fluids **62**(8), 159 (2021). https://doi.org/10.1007/s00348-021-03244-1

[11] Cai, S., Liang, J., Gao, Q., Xu, C., Wei, R.: Particle image velocimetry based on a deep learning motion estimator. IEEE Transactions on Instrumentation and Measurement **69**(6), 3538–3554 (2020). https://doi.org/10.1109/TIM.2019.2932649

[12] Lagemann, C., Lagemann, K., Mukherjee, S., Schröder, W.: Deep recurrent optical flow learning for particle image velocimetry data. Nature Machine Intelligence **3**, 641–651 (2021). https://doi.org/10.1038/s42256-021-00369-0

[13] Lagemann, C., Lagemann, K., Mukherjee, S., Schröder, W.: Generalization of deep recurrent optical flow estimation for particle-image velocimetry data. Measurement Science and Technology **33** (2022). https://doi.org/10.1088/1361-6501/ac73db

[14] Vig, D.K., Hamby, A.E., Wolgemuth, C.W.: On the quantification of cellular velocity fields. Biophysical Journal **110**(7), 1469–1475 (2016). https://doi.org/10.1016/j.bpj.2016.02.032

[15] Kumashiro, K.: A physics-constrained three-dimensional three-component particle-based velocimetry method for constant-density flows. Master's thesis, University of Toronto (2019)

[16] Heitz, D., Héas, P., Mémin, E., Carlier, J.: Dynamic consistent correlation-variational approach for robust optical flow estimation. Experiments in Fluids **45**(4), 595–608 (2008). https://doi.org/10.1007/s00348-008-0567-4

[17] Perlman, E., Burns, R., Li, Y., Meneveau, C.: Data exploration of turbulence simulations using a database cluster. In: SC '07: Proceedings of the 2007 ACM/IEEE Conference on Supercomputing, pp. 1–11 (2007). https://doi.org/10.1145/1362622.1362654

[18] Li, Y., Perlman, E., Wan, M., Yang, Y., Meneveau, C., Burns, R., Chen, S., Szalay, A., Eyink, G.: A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence. Journal of Turbulence **9**, 31 (2008) https://doi.org/10.1080/14685240802376389. https://doi.org/10.1080/14685240802376389

[19] Kazbekov, A.: Enstrophy transport in premixed swirl flames by askar kazbekov a thesis submitted in conformity with the requirements. (2018)

[20] Fleet, D., Weiss, Y.: In: Paragios, N., Chen, Y., Faugeras, O. (eds.) Optical Flow Estimation, pp. 237–257. Springer, Boston, MA (2006). https://doi.org/10.1007/0-387-28831-7_15

[21] Heitz, D., Mémin, E., Schnörr, C.: Variational fluid flow measurements from image sequences: synopsis and perspectives. Experiments in Fluids **48**(3), 369–393 (2010). https://doi.org/10.1007/s00348-009-0778-3

[22] Aubert, G., Deriche, R., Kornprobst, P.: Computing optical flow via variational techniques. SIAM Journal on Applied Mathematics **60** (1999). https://doi.org/10.1137/S0036139998340170

[23] Chen, J., Duan, H., Song, Y., Cai, Z., Yang, G., Liu, T.: Motion estimation for complex fluid flows using helmholtz decomposition. IEEE Transactions on Circuits and Systems for Video Technology, 1–1 (2022). https://doi.org/10.1109/TCSVT.2022.3219438

[24] Wang, B., Cai, Z., Shen, L., Liu, T.: An analysis of physics-based optical flow. Journal of Computational and Applied Mathematics **276**, 62–80 (2015). https://doi.org/10.1016/j.cam.2014.08.020

[25] Horn, B., Schunck, B.: Determining optical flow. Artificial Intelligence **17**, 185–203 (1981). https://doi.org/10.1016/0004-3702(81)90024-2

[26] Corpetti, T., Memin, E., Perez, P.: Dense estimation of fluid flows. IEEE Transactions on Pattern Analysis and Machine Intelligence **24**(3), 365–380 (2002). https://doi.org/10.1109/34.990137

[27] Barron, J.L., Fleet, D.J., Beauchemin, S.S.: Performance of optical flow techniques. International Journal of Computer Vision **12**(1), 43–77 (1994). https://doi.org/10.1007/BF01420984

[28] Cai, S., Mémin, E., Dérian, P., Xu, C.: Motion estimation under location uncertainty for turbulent fluid flows. Experiments in Fluids **59**(1), 8 (2017). https://doi.org/10.1007/s00348-017-2458-z

[29] Anandan, P.: A computational framework and an algorithm for the measurement of visual motion. International Journal of Computer Vision **2**(3), 283–310 (1989). https://doi.org/10.1007/BF00158167

[30] Ruhnau, P., Yuan, J., Schnörr, C.: On variational methods for fluid flow estimation. In: Jähne, B., Mester, R., Barth, E., Scharr, H. (eds.) Complex Motion, pp. 124–145. Springer, Berlin, Heidelberg (2007)

[31] Brox, T., Bruhn, A., Papenberg, N., Weickert, J.: High accuracy optical flow estimation based on a theory for warping. In: Pajdla, T., Matas, J. (eds.) Computer Vision - ECCV 2004, pp. 25–36. Springer, Berlin, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24673-2_3

[32] Fortun, D., Bouthemy, P., Kervrann, C.: Optical flow modeling and computation: a survey. Computer Vision and Image Understanding **134**, 21 (2015). https://doi.org/10.1016/j.cviu.2015.02.008

[33] Raffel, M., Willert, C.E., Scarano, F., Kähler, C.J., Wereley, S.T., Kompenhans, J.: Particle Image Velocimetry: a Practical Guide. Springer, Cham ZG (2018). https://doi.org/10.1007/978-3-319-68852-7

[34] Lekien, F., E., M.: Tricubic interpolation in three dimensions. International Journal for Numerical Methods in Engineering **63**, 455–471 (2005). https://doi.org/10.1002/nme.1296

[35] Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. Mathematical Programming **45**(1), 503–528 (1989). https://doi.org/10.1007/BF01589116

[36] Liberzon, A., Lasagna, D., Aubert, M., Bachant, P., Käufer, T., Jakirkham, Bauer, A., Vodenicharski, B., Dallas, C., Borg, J., Tomerast, Ranleu: OpenPIV/openpiv-python: OpenPIV - Python (v0.22.2) with a new extended search PIV grid option. Zenodo (2020). https://doi.org/10.5281/ZENODO.3930343. https://zenodo.org/record/3930343

[37] Scharnowski, S., Kähler, C.J.: Particle image velocimetry - classical operating rules from today's perspective. Optics and Lasers in Engineering **135**, 106185 (2020). https://doi.org/10.1016/j.optlaseng.2020.106185

[38] Kazbekov, A., Kumashiro, K., Steinberg, A.M.: Enstrophy transport in swirl combustion. Journal of Fluid Mechanics **876**, 715–732 (2019). https://doi.org/10.1017/jfm.2019.551

[39] Liang, J., Cai, S., Xu, C., Chu, J.: Filtering enhanced tomographic piv reconstruction based on deep neural networks. IET Cyber-Systems and Robotics **2**(1), 43–52 (2020) https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/iet-csr.2019.0040. https://doi.org/10.1049/iet-csr.2019.0040

[40] Savill, A.M.: Evaluating turbulence model predictions of transition. Applied Scientific Research **51**(1), 555–562 (1993). https://doi.org/10.1007/BF01082590

[41] Lu, G.: A high-resolution optical flow-inspired 3d3c velocimetry method based on sparse representation of 3d particle images. Master's thesis, University of Toronto (2021)