

Using 3D Topological Connectivity for Ghost Particle Reduction in Flow Reconstruction

Christina Tsalicoglou Thomas Rösgen
Institute of Fluid Dynamics, ETH Zurich
Zurich, Switzerland

{ctsalico, roesgen}@ethz.ch

Abstract

Volumetric flow velocimetry for experimental fluid dynamics relies primarily on the 3D reconstruction of point objects, which are the detected positions of tracer particles identified in images obtained by a multi-camera setup. By assuming that the particles accurately follow the observed flow, their displacement over a known time interval is a measure of the local flow velocity. The number of particles imaged in a 1 Megapixel image is typically in the order of 10^3 - 10^4 , resulting in a large number of consistent but incorrect reconstructions (no real particle in 3D), that must be eliminated through tracking or intensity constraints. In an alternative method, 3D Particle Streak Velocimetry (3D-PSV), the exposure time is increased, and the particles' pathlines are imaged as "streaks". We treat these streaks (a) as connected endpoints and (b) as conic section segments and develop a theoretical model that describes the mechanisms of 3D ambiguity generation and shows that streaks can drastically reduce reconstruction ambiguities. Moreover, we propose a method for simultaneously estimating these short, low-curvature conic section segments and their 3D position from multiple camera views. Our results validate the theory, and the streak and conic section reconstruction method produces far fewer ambiguities than simple particle reconstruction, outperforming current state-of-the-art particle tracking software on the evaluated cases.

1. Introduction

Volumetric velocimetry techniques for fluid flows enable the analysis of complex flow fields and their complete spatio-temporal evolution, by providing all three velocity components of the flow in discrete positions in a 3D measurement volume. Most volumetric methods require a minimum of two synchronized cameras and record the positions of tracers over time as they follow the flow. The most common approaches are 3D Particle Tracking Velocimetry (3D-PTV) [14] and Tomographic Particle Image Velocime-




Figure 1. Epipolar and epipolar tangency constraints for conic section segments reduce the reconstruction ambiguities, “ghosts”, in 3D flow reconstruction from streaks.

try (Tomo-PIV) [4]. In 3D-PTV, the tracer particles are reconstructed in 3D space by triangulation, while in Tomo-PIV a light intensity field is reconstructed tomographically. Both methods result in “ghost particles”, which are consistently reconstructed particles for which it remains undecidable whether they are real or artifacts of the reconstruction process. The tracer density in these applications is in the order of 10^3 - 10^4 particles in a 1 Mpx image, resulting in a large number of ghost particles if left untreated.

There are several ways to reduce the number of ghost particles and improve the reconstruction quality: increasing the number of cameras, reducing the seeding density of tracer particles, or rejecting particles using criteria such as their intensity, estimated velocity, or trajectory smoothness (e.g. [19]). Still, triangulating particles independently in time, frame-by-frame, results in a loss of information which needs to be recovered using assumptions about the flow field. Handling regions with large velocity differences remains a challenge, as the particle displacement assumptions often fail to eliminate ghost tracks [20].

A less common variant of 3D-PTV is 3D Particle Streak Velocimetry (3D-PSV), where the cameras' exposure time is increased, so that the tracer particles' pathlines over the exposure time are recorded as “streaks” [2, 15, 16, 24]. Therefore, the tracking information is retained in a single image and can be used in the triangulation step to eliminate some of the reconstruction ambiguities. Furthermore, the displacement of both fast and slow-moving particles is en-

coded in the streaks, rendering the method inherently suited to reconstructing flows with a high dynamic velocity range. Therefore, we ask the question, *by how much can the number of ambiguous, “ghost”, reconstructions be reduced by using the topological information inherently contained in long-exposure images?*

We initially treat the streaks as line segments defined by their endpoints and develop a theoretical model that describes how and by how much the number of reconstruction ambiguities is reduced when reconstructing linear streaks instead of points (Sections 3.2.2, 3.2.3). In a second stage, we model the streaks as projections of conic section segments and show that the probability of ghost streak generation is further reduced by imposing correspondence constraints for conic sections over multiple views (Section 3.2.4). Finally, we propose a new method for establishing multi-view matches among conic section segments while satisfying an endpoint correspondence constraint (Section 3.3). We validate our models and reconstruction method on synthetic data (Sections 4.1, 4.2) and test our reconstruction method on experimental data (Section 4.3).

Our theory and validation results fully agree, and ghost reconstructions are virtually eliminated when reconstructing a realistic flow field. Experimental results confirm the practical value of our reconstruction method. Therefore, our work challenges the established preference for point reconstruction and provides theoretical argumentation for why using topological connectivity is beneficial in reducing the number of ghost particles in flow reconstruction.

2. Related work

We build on previous work in experimental fluid dynamics and photogrammetric computer vision. The problem of ghost particles in the reconstruction of densely seeded flows has been identified and described analytically by Maas et al. [13, 14] as a function of the number of particles per pixel, volume depth, and relative position of the cameras, using arguments from epipolar geometry. Elsinga et al. [5] used similar arguments to describe the number of expected ghost particles in Tomo-PIV. Modern 3D-PTV methods manage to significantly reduce ghost particles through Iterative Particle Reconstruction [25], tracking particles over multiple time-steps [19], imposing physical constraints on the reconstructed flow field [11] or using multi-exposed recordings [17]. Other approaches use colored light or focus/defocus to encode the depth information [12, 26, 27].

Here, we harness the information of the endpoint connectivity and shape of streaks that represent particles’ pathlines to reduce the number of ambiguous reconstructions. We employ well-established arguments about point, line, and conic section reconstruction to determine by how much the number of reconstruction ambiguities can be reduced when using long-exposure imaging in 3D-PSV. Conic sec-

tions are chosen to represent the shape of streaks as they are projectively invariant and offer more correspondence conditions than points and lines, without requiring point-wise matching as more complex curves do (e.g., [6]). Fundamental work on conic section matching has been performed by [3, 7, 10, 18, 21]. To estimate and match conic section segments, we use a variant of the method proposed by [21] for our RANSAC-type baseline simulation (supp. mat.), while our optimization problem is similar to that posed by [22].

3. Our approach

Ghost particle generation is one of the fundamental challenges in 3D-PTV, and is caused by ambiguities inherent in projective geometry. To drastically reduce the ambiguities generated by reconstructing particles, we propose reconstructing streaks: short lines or curve segments that result from long-exposure recordings of the pathlines of moving particles. We perform an analysis of ghost streak generation probabilities for linear and curved streaks and propose a method for matching curved streaks across multiple views.

3.1. Preliminaries

The generation of ghost particles is a by-product of projecting 3D points onto two dimensions, as any number of points along a camera’s line of sight are projected to a single position in the image plane. Though the depth can already be reconstructed by using a second camera, constellations can exist for which more consistent 3D positions can be reconstructed than the actual number of corresponding particles imaged by the two cameras.

Maas [13] described the probability of generating such reconstruction ambiguities from two images with randomly distributed particles as follows: a particle in one image, I , defines an epipolar line l' in the second image, I' , and ambiguities occur when more than one particle coincides with l' . Given two images with co-planar image planes and n particles in each, the probability of more than one particle being on l' is then the probability that at least one of the remaining $n - 1$ particles coincides with the epipolar line l' .

In reality, the region where a particle must be located to generate a valid reconstruction is, rather than an epipolar line, an epipolar window f' of width $2d$ ($\pm d$) and length L around l' , where d is the allowed distance from the epipolar line and L is the length of l' . The area of this epipolar window is then $f = 2dL$, and the probability of any one particle lying in this window is f/F , where F is the total area of the image. Therefore, the average number of ambiguous reconstructions is,

$$n_p = (n - 1) \frac{f}{F}. \quad (1)$$

This process must be considered for each of the n particles

in image I . Thus, the total number of ambiguities is

$$N_p = (n^2 - n) \frac{f}{F}. \quad (2)$$

It must be noted that Maas limits the search length of the epipolar line based on the depth of the observed volume. In the following analysis we neglect this constraint.

3.2. Probability models for ghost streak generation

3.2.1 Setup

A linear streak in image I defines, by its endpoints, two epipolar lines l'_a and l'_b in I' (Fig. 2). A 3D streak can be reconstructed when there exists at least one streak connecting the two epipolar lines l'_a, l'_b , and ambiguities occur when more than one streak connects these two epipolar lines. For our analysis we restrict ourselves to a setup with two cameras with co-planar image planes, which helps to simplify the analysis as the epipolar lines are parallel (Fig. 2). Additionally, it is assumed that all particles move to a position within the field of view during the exposure time.

In point reconstruction, and in a uniformly seeded volume, the positions of the frozen particles in one timestep reveal no information about the flow. Conversely, streak images inherently contain information about the flow field. Therefore, the efficacy of streak imaging in reducing reconstruction ambiguities depends on the observed flow field, and an estimate of the expected number of reconstruction ambiguities will be bounded by best-case and worst-case displacement scenarios.

We identify two limiting cases between which the number of streak ambiguities is bound to lie, with the worst-case scenario occurring when the particles experience no displacement and the best-case scenario occurring when they experience random displacements within the image. Between the two extremes, we consider bounded particle displacements to determine the expected number of ambiguities.

3.2.2 Random displacements

The following case presents the scenario where each particle can move to a random position in the image during the exposure time. This case, while physically unrealistic, is useful in establishing a lower bound and introducing the main ghost streak generation mechanism.

The endpoints of a streak in image I define the epipolar lines l'_a and l'_b and their corresponding epipolar windows f'_a and f'_b in I' (Fig. 3a). A reconstruction ambiguity occurs if there exists more than one streak connecting these epipolar windows. For this constellation to occur, the following two conditions must hold. *Condition 1*: at least one more of the remaining points in the image must coincide with f'_a ,




Figure 2. The imaged streaks (black) are projections of the blue streaks in 3D space. Without prior knowledge of the world (blue) streaks, the imaged streaks' volumetric reconstruction can generate four different streaks (blue and red), of which two, three or all four might be real. Here, the two ghost streaks are shown in red.

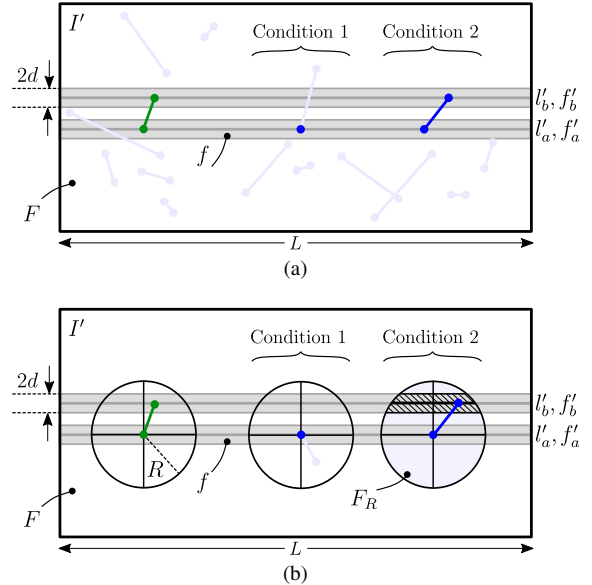


Figure 3. Necessary conditions for ambiguous reconstructions, for streaks generated from (a) random particle displacements and (b) bounded particle displacements. The green streak defines the epipolar constraints, while the blue streak represents every other streak that fulfills the constraints.

and *Condition 2*: the additional particle in f'_a must be the endpoint of a streak with its other endpoint in f'_b .

The expected total number of particles n_p that will coincide with the epipolar window f'_a is known from the work of Maas [13] and Eq. 1. Here, however, twice the number of points, $2n$, must be considered, as the signature of each of the n particles is a line segment with two endpoints in the same image. As both endpoints are randomly distributed in I' , the expected number of particles in f'_a is $n_p = (2n - 2)f/F$.

For *Condition 2*, assuming that the particle displacements are random, the probability that an endpoint lies in

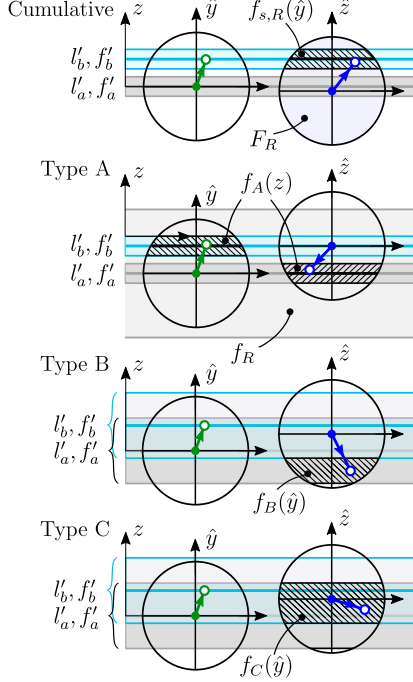


Figure 4. Schematic depiction of the ghost streak generating constellations for the bounded displacement case. The green streak defines the epipolar constraints.

f'_b is f/F , as out of all the possible positions the particle can jump to within the image, it must land in the area f'_b . The average number of particles for which this happens is

$$n_s = 2(n-1) \left(\frac{f}{F} \right)^2 \quad (3)$$

and as this process is repeated for all n streaks in I , the resulting total expected number of streak ambiguities for random displacements is

$$N_s = 2(n^2 - n) \left(\frac{f}{F} \right)^2. \quad (4)$$

Comparing this expression with the one for point correspondences (Eq. 2), one can already see a drastic reduction since $f/F \ll 1$ can be assumed.

3.2.3 Bounded displacements

In a more realistic case, particles will not move to a random position in the image during the exposure time but will instead move to a position within a radius R , thus generating streaks with endpoints at a bounded distance from each other (Fig. 3b). Ambiguities will be generated when two epipolar windows in I' , defined by these endpoints, are connected by more than one streak, as in Section 3.2.2.




Figure 5. Contribution of the different mechanisms of ghost streak generation to the total number of ghost streaks (blue), for increasing radius R in a 1024×1024 px² image.

The conditions for this case are the following. *Condition 1*: at least one more of the remaining points in the image must coincide with f'_a , and *Condition 2*: the additional particle in f'_a must be the endpoint of a streak with its other endpoint in the area of intersection between f'_b and the area F_R of the circle with radius R centered around the first endpoint (Fig. 3b).

We consider the “first” endpoint as randomly distributed and the “second” as conditionally distributed, as it cannot be further away from the first than a distance R . Then, the probability of the first endpoint of a streak falling in f'_a is f/F . The probability of the second endpoint falling in f'_b is the probability that out of all the possible positions in the circle with area F_R , the second endpoint will fall in $f_{s,R}(\hat{y})$, the area of intersection of the circle of radius R around the first endpoint and the epipolar window f'_b (Fig. 3b).

The size of this area depends on the position of the first streak’s second endpoint, which defines l'_b , and follows Wigner’s semi-circle distribution along the \hat{y} axis. Therefore, the expected value $\bar{f}_{s,R}$ of the area $f_{s,R}(\hat{y})$ is described by the integral over the circle’s diameter, of the probability density function of the semi-circle distribution times the area $f_{s,R}(\hat{y})$, integrated over all positions of f'_b (Fig. 4). Relevant formulae are provided in the supp. mat.

By further analyzing the conditions leading to the above cumulative number of ambiguities, we identify three types of streak endpoint constellations that contribute to ambiguity generation (Fig. 4):

Type A: the first endpoint is only in f'_a and the second endpoint is only in f'_b .

Type B: the first and the second endpoint are in f'_a and only the first endpoint is in f'_b .

Type C: the first and the second endpoint are both in f'_a and in f'_b .

It can be seen that the relative importance of these three mechanisms varies depending on the ratio of the epipolar constraint tolerance, d , and the maximum expected particle displacement, R (Fig. 5). The contributions n_A , n_B , n_C

for *Type A*, *B* and *C* constellations to the total number of ambiguities can be described as follows:

$$n_A = 2(n-1) \frac{f_R}{F} \frac{\bar{f}_A^2}{F_R^2}, \quad (5)$$

$$n_B = 2(n-1) \frac{f}{F} \frac{\bar{f}_B}{F_R}, \quad (6)$$

$$n_C = (n-1) \frac{f}{F} \frac{\bar{f}_C}{F_R}. \quad (7)$$

f_R is the area with extent RL on either side of f'_a (Fig. 4) and the formulae for the mean areas of intersection \bar{f}_A^2 , \bar{f}_B , \bar{f}_C are provided in the supplementary material.

Type C ambiguities occur when both endpoints of a streak in one image can be matched to both endpoints of a streak in another image. This contribution increases as the ratio d/R increases. In such cases we assume that one of the matches can be eliminated, so that only one reconstruction is possible from the same streak pair. Therefore, the expected number of ambiguities, $n_{s,R}$, per epipolar line, and total expected number, $N_{s,R}$, of ambiguities reconstructed from images with n randomly distributed streaks, bounded by a maximum displacement radius R is

$$n_{s,R} = 2(n-1) \frac{f}{F} \frac{\bar{f}_{s,R}}{F_R} - n_C, \quad (8)$$

$$= n_A + n_B + n_C \quad (9)$$

$$N_{s,R} = n \cdot n_{s,R}.$$

Note that this expression contains the two limiting cases discussed before. For a non-moving particle $f_R \rightarrow 0$, $\bar{f}_B \rightarrow 0$ and $\bar{f}_C \rightarrow F_R$ so the single point probability is recovered.

3.2.4 Conic section segments

Considering the shapes of the streaks in the matching process further reduces the number of reconstruction ambiguities. Here, we assume that a curved streak can be modeled as a segment of a conic section projected to the cameras from a conic section segment on a plane in the 3D world.

The additional correspondence condition offered by conic sections is that of the epipolar tangency [3, 18, 21]: two matching conics are tangent to corresponding epipolar lines. It follows that, given a conic in image I , tangent to $l_{t,a}$ and $l_{t,b}$, a conic can be reconstructed in 3D if there is a conic tangent to $l'_{t,a}$ and $l'_{t,b}$ in I' (Fig. 6). If multiple conics are tangent to $l'_{t,a}$ and $l'_{t,b}$, we obtain reconstruction ambiguities and, therefore, ghost conics. Although we treat the curved streaks as general conic section segments during reconstruction, we focus the following discussion on ellipses, as they allow for a more geometrically intuitive handling.

We calculate the probability of an ellipse in an image being tangent to two specific epipolar lines, provided that the




Figure 6. Conditions for ghost ellipse generation: endpoint correspondence conditions must hold and, additionally, the ellipses to which the segments belong must be tangent to corresponding epipolar lines within some tolerance d_e (gray area).

lines are horizontal and allowed to be within $y \in [-R, H + R]$, where H is the height of the image and R the maximum expected displacement of a particle. The coordinate y , that defines an ellipse's tangent, also follows a uniform distribution outside of the region between the two endpoints. One endpoint of a streak in the image is uniformly distributed, but the y coordinate of its other endpoint follows the semi-circle distribution due to the circular bounded displacement. Finally, a tolerance, d_e , is defined for the epipolar tangency constraint so that the ellipse's epipolar tangents must be within a distance d_e of the desired epipolar line for the constraint to be fulfilled. The probability \bar{p}_e used to calculate the expected value of the number of ghost ellipses can be found in the supplementary material. The number of ghost ellipse matches per streak, n_e , and the total number of ghost ellipses, N_e , in an image with n streaks are calculated as

$$n_e = n_{s,R} \cdot \bar{p}_e, \quad (10)$$

$$N_e = n \cdot n_e, \quad (11)$$

where $n_{s,R}$ is the number of linear ghost streaks (Section 3.2.3). As $\bar{p}_e \leq 1$, the number of curved ghost streaks will be at most equal to that of linear ghost streaks, with n_e approaching $n_{s,R}$ as d_e increases. The presented models are validated in Section 4.

3.3. Conic section reconstruction

In experiments, the conics that fit through the curved streaks' projected 2D coordinates are unknown and must be estimated. The epipolar tangency constraint, as presented in the previous section, is, however, not practical for matching short segments of conic sections with low curvature such as those obtained from realistic flow images. Therefore, a method is presented here to simultaneously fit the conics and estimate a multi-view correspondence criterion. The method is validated and compared to a RANSAC-based method described in detail in the supp. mat. (Section 4).

3.3.1 Optimization problem

In our matching approach, we optimize the conic fits to the curved streak datapoints over multiple cameras while enforcing a correspondence criterion between the conic segments imaged on the different views. The correspondence is enforced by ensuring that the conic segments on all views are images of the same world conic. The method is described here for three cameras, but it has been implemented for two-camera and three-camera setups, and it can easily be extended to more camera views.

Therefore, a conic is sought that lies on a plane in the 3D world and passes through the two triangulated endpoints of the streak that, when projected to the images, is a good fit to the point clouds that define the curved streak segments in the images (Fig. 7).

A plane $\pi = [\alpha, \beta, \gamma, \delta]$, a conic section C_π with parameters θ_π and three cameras with projection matrices P_1, P_2, P_3 are given. The homography that transforms points \mathbf{x}_π on the plane π to points \mathbf{x}_i on the plane of camera i , so that $\mathbf{x}_i = H_{\pi,i}\mathbf{x}_\pi$ is defined as:

$$H_{\pi,i} = \begin{bmatrix} | & | & | \\ H^1 & H^2 & H^3 \\ | & | & | \end{bmatrix}, \quad (12)$$

$$H^1 = \gamma P_i[:, 1] - \alpha P_i[:, 3], \quad (13)$$

$$H^2 = \gamma P_i[:, 2] - \beta P_i[:, 3], \quad (14)$$

$$H^3 = \gamma P_i[:, 4] - \delta P_i[:, 3], \quad (15)$$

for each of the $i = 1, 2, 3$ cameras.

Then, an ellipse on π can be transformed to the camera plane as $C_i = H_{\pi,i}^{-\top} C_\pi H_{\pi,i}^{-1}$. The goal is then to minimize the Sampson error

$$\mathcal{L}(\pi, \theta_\pi) = \sum_{i,j} \frac{(\mathbf{x}_{i,j}^\top C_i \mathbf{x}_{i,j})^2}{4((C_i \mathbf{x}_{i,j})_1^2 + (C_i \mathbf{x}_{i,j})_2^2)} \quad (16)$$

for each of the i cameras and j datapoints, where $\mathbf{x}_{i,j}$ is datapoint j on camera i .

As only streak combinations that fulfill the epipolar constraints for linear streaks are considered, the streaks' endpoint correspondence is guaranteed. The plane π is therefore a pencil of planes through the line that connects the two endpoints of the streaks, $\mathbf{X}_s, \mathbf{X}_e$, while the conic must necessarily pass through the triangulated endpoints on plane π , $\mathbf{x}_{\pi,s}, \mathbf{x}_{\pi,e}$. This defines a set of equality constraints, resulting in the following optimization problem

$$\begin{aligned} \min_{\pi, \theta_\pi} \quad & \mathcal{L}(\pi, \theta_\pi) \\ \text{s.t.} \quad & \mathbf{x}_{\pi,s}^\top C_\pi \mathbf{x}_{\pi,s} = 0, \\ & \mathbf{x}_{\pi,e}^\top C_\pi \mathbf{x}_{\pi,e} = 0, \\ & \pi^\top \mathbf{X}_s = 0, \\ & \pi^\top \mathbf{X}_e = 0, \end{aligned} \quad (17)$$




Figure 7. Setup of the optimization problem for 3 cameras.

which we solve using a non-linear solver (IPOPT) [23] with CasADi [1]. The tolerance for accepting a good fit is set by estimating the standard deviation of the distance of the point cloud to an unconstrained best-fit conic section. The initial guess for the parameters of π and θ_π is found as a best-fit through five points sampled along the streaks on the two images using the epipolar constraint. To improve the fits for noisy data we use a regularizer (supp. mat.).

4. Results

4.1. Synthetic data - model validation

Synthetic images of particles and their streak signatures are used to validate the estimates of the probability of ghost streak and ghost ellipse generation as described in the previous sections. The particles' coordinates are uniformly distributed in two images and are then displaced either to a random location in each image or to a position within a circle of radius R to simulate bounded particle displacements. We use a co-planar, two-camera setup (Fig. 2) with image sizes of 1024×1024 px².

4.1.1 Linear segments

The points and streaks in the two images are matched using the epipolar constraint and two different strategies: (a) the streaks' endpoints are matched independently of each other and (b) the endpoints are matched jointly, i.e. a match is accepted only if both endpoints of a streak in image I can be matched to the endpoints of a streak in I' .

For each case, the particle generation and matching is repeated with 20 image sets, containing from 1000 to 7000 particles, and the number of reconstructed particles/streaks is averaged. The epipolar constraint is set to $d=1$ px.

The theoretical model and the reconstruction results match well, as shown in Fig. 8 and in Table 1. The number of reconstruction ambiguities lies consistently below that for individual particles by almost an order of magni-

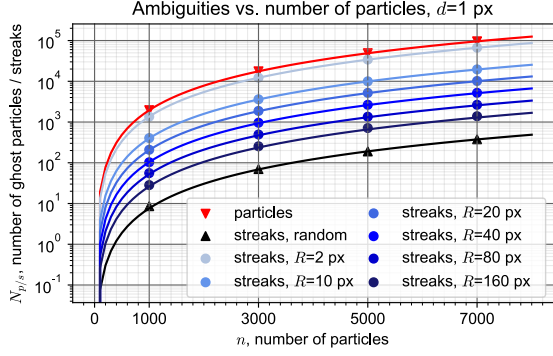


Figure 8. Number of ghost streak and ghost particle reconstructions from synthetic data (markers) and theoretical estimates (lines), for increasing number of particles, n , with $d=1$ px, $F=1024 \times 1024$ px². The bounded cases lie between the two edge cases of individual particle reconstruction and streak reconstruction for streaks with random displacements.

	N_p		$N_{s,R}, R=10$ px		$N_{s,R}, R=40$ px	
	pred.	sim.	pred.	sim.	pred.	sim.
$n=1000$	1951	1968	397	404	104	103
$n=3000$	17572	17739	3573	3614	937	962
$n=5000$	48818	49285	9927	9978	2603	2650
$n=7000$	95689	96543	19458	19579	5102	5206

Table 1. Number of reconstructed ghost particles, N_p and ghost streaks $N_{s,R}$, for two different maximum displacement radii and different numbers of imaged particles, n , with $d=1$ px, $F=1024 \times 1024$ px². The predictions are derived from the theoretical analysis (Section 3.2.3), while the simulation results are obtained as the mean of 20 runs with new particle positions.

tude, even for small maximum displacements. It is evident that ghost streak generation is reduced with increasing maximum allowable displacement radius, as the position of each streak’s second endpoint is allowed more variability. Increasing the initial number of particles n produces the same behavior for ghost streaks as for individual ghost particles.

4.1.2 Ellipse segments

For the ellipse segment matching, ellipses with varying parameters are fitted through the streak endpoints so that the fitted ellipses are tangent to horizontal lines randomly distributed in $y \in [-R, H + R]$. The results were obtained using an epipolar constraint of $d=1$ px, epipolar tangency constraints d_e of 100 and 200 px and $n=1000$ streaks.

A good match between the theoretical estimate (Section 3.2.4) and the reconstructed number of ghost ellipses can be seen in Fig. 9. The number of ghost ellipses lies below that for linear ghost streaks, and it depends on the epipolar constraint tolerance, d_e , which in turn depends on the uncertainty of the estimated ellipse parameters. For example, for $R=40$ px and $d_e=200$ px the number of ambiguities is reduced from 104 ghost streaks to 39 ghost ellipses.




Figure 9. Number of ghost streak and ghost ellipse segment reconstructions from synthetic data (dots) and theoretical predictions (lines), for increasing number of maximum allowable displacement radius, R and two different epipolar tangency tolerances d_e . $d=1$ px, $F=1024 \times 1024$ px².

	N_p		N_s		N_c , optimization		N_c , RANSAC	
	ghosts	ghosts	ghosts	ghosts	ghosts	lost	ghosts	lost
0 px								
$n=1000$	2001.5	189.1	11.6	0.2	7.7	0.1		
$n=3000$	18095.5	1692	115.9	1.4	76	1.3		
0.25 px								
$n=1000$	2000	189.3	52	14.2	31	63		
$n=3000$	18098.4	1696.1	463.8	87.25	275	454		
0.5 px								
$n=1000$	2001	186.5	61.3	28.3	58.8	26.1		
$n=3000$	18115.6	1688.9	553.6	197	543.6	180		

Table 2. Number of reconstructed ghost particles, N_p , ghost streaks, N_s , and curved ghost streaks, N_c , for different initial number of particles n and noise levels. “Ghosts” are the ambiguous reconstructions, while “lost” are the false negatives. The two conic reconstruction methods, optimization-based and RANSAC-based, are compared. $d=1$ px, $F=1024 \times 1024$ px².

4.2. Synthetic data - reconstruction validation

4.2.1 Conic section data

The synthetic data used for the validation of our conic fit and reconstruction method consist of images of conic section segments, obtained by projecting onto the cameras conic section segments generated in 3D, with parameters limited to a range that produces similar results to experimental images. Uniformly distributed noise up to ± 0.5 px is added to the curved streaks’ 2D coordinates.

These conic section segments are then reconstructed using (a) the optimization-based method described in Section 3.3 and (b) our RANSAC-type method (supp. mat). The mean results of 20 3D scenes are presented in Table 2.

In the absence of noise, effectively all segments are reconstructed correctly, and the number of curved ghost streaks, N_c , is significantly lower than the number of linear ghost streaks, N_s . Noise reduces both methods’ performance. Increasing the number of particles increases N_c , which, however, remains at about a third of N_s .

	N_p	N_s	N_c , optimization		N_{CS}	
	ghosts	ghosts	ghosts	lost	ghosts	lost
	0 px					
$n=1000$	542.5	10	4	0	20	71
$n=3000$	4925	111	57	0	149	562
	0.25 px					
$n=1000$	546	11	11	1	-	-
$n=3000$	4951	111	87	6	-	-
	0.5 px					
$n=1000$	551	12	12	1	-	-
$n=3000$	4952.5	107	99	5	-	-

Table 3. Number of reconstructed ghost particles, N_p , ghost streaks, N_s , and curved ghost streaks, N_c , for different initial number of particles n and noise levels, for Hill’s spherical vortex data. “Ghosts” are the ambiguous reconstructions, while “lost” are the false negatives. Our method is compared to a commercial 3D-PTV software (N_{CS}). $d=1$ px, $F=1024 \times 1024$ px².

4.2.2 Hill’s spherical vortex

We reconstruct synthetic streak images of Hill’s spherical vortex [9] to validate our method on realistic flow field data. The mean displacement is 33 px, and the same setup of two co-planar views is used. The results of the conic reconstruction method are compared with results obtained using state-of-the-art commercial 3D-PTV software [8] (Table 3).

Only data without noise are processed with the commercial software, as the purpose of noise is to assess our conic reconstruction method. The software requires range and velocity constraints, and we impose the same constraints on our reconstruction pipeline. Direct comparison with the commercial software is challenging, as in the software, the number of ghost reconstructions can be tuned at the expense of losing more true particles. Therefore, the results of Table 3 are what we consider a fair compromise between recognized true particles and reconstructed ghosts. The detailed evaluation settings, unconstrained results that can be compared to the theoretical estimates, and results obtained with different parameters are provided in the supp. mat.

Our method drastically reduces ghost reconstructions in a flow field with high dynamic velocity range, typically challenging for 3D-PTV. From a single frame, we obtain lower ghost streak numbers than those obtained by an equivalent multi-frame analysis using state-of-the-art commercial software for the examined flow field and setup.

4.3. Experimental data

The flow field we use in our experiments is a vortex ring in air, visualized using helium-filled soap bubbles as tracers. Two synchronized 1 Mpx high-speed cameras record particle images at 1000 fps, with an exposure time of 1 ms. The streaks are obtained by summing the individual frames and are then detected using a CNN-based instance segmentation method. The high frame rate acquisition allows us to perform both 3D-PSV and 3D-PTV analyses.




Figure 10. Projections of the reconstructed experimental data superimposed on one of the camera views (a) with our method and (b) with 3D-PTV. For (a) the velocity, u , is obtained from the 3D arc length. Velocities normalized by maximum velocity, u_{max} .

n Cam. 1 / Cam. 2	$N_{t,p}$	$N_{t,s}$	$N_{t,c}$, optimization	$N_{t,CS}$
2529 / 2265	20366	4928	3477	2759

Table 4. Number of reconstructed particles, $N_{t,p}$, linear streaks, $N_{t,s}$, and curved streaks, $N_{t,c}$, for the vortex ring experiment with $d=1$ px. Our method is compared to 3D-PTV software ($N_{t,CS}$).

2529 and 2265 streaks are detected in the two views, and the number of *total* particle and streak reconstructions, including correct and ghost reconstructions, are shown in Table 4, along with results from the 3D-PTV software [8]. The reconstructed streaks agree well with the recorded images, with visibly more complete trajectories recovered compared to 3D-PTV (Fig. 10). Camera synchronization and calibration are essential to keep the epipolar tolerance, d , and the number of ambiguities, low. Here, $d=1$ px is used. Details on the experimental setup are provided in the supp. mat.

Limitations. A notable limitation of the proposed method is that we require the curved streaks to be planar curves, thus effectively limiting the maximum allowable exposure time. Additionally, streak imaging requires sufficiently accurate streak segmentation and endpoint detection, which can be more challenging than particle detection.

5. Conclusion

We have shown that 3D-PSV is a method inherently suited for reducing the number of reconstruction ambiguities in volumetric flow velocimetry. We have provided models to estimate the expected number of linear and curved ghost streaks and presented a method for simultaneously fitting conic sections through curved streaks and testing their correspondence over multiple views. Simulations validate our theoretical analysis, and streak reconstruction produces significantly fewer ambiguities than point reconstruction. Our endpoint and subsequent conic section matching approach is validated using synthetic flow data, with results outperforming state-of-the-art commercial software on the evaluated cases. Finally, the reconstruction of experimental data shows that our conic matching and reconstruction method can be used to successfully evaluate real data.

References

- [1] Joel A.E. Andersson, Joris Gillis, Greg Horn, James B. Rawlings, and Moritz Diehl. CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019. [6](#)
- [2] Pascal Henry Biwole, Wei Yan, Yanhui Zhang, and Jean Jacques Roux. A complete 3D particle tracking algorithm and its applications to the indoor airflow study. *Measurement Science and Technology*, 20(11), 2009. [1](#)
- [3] Geoffrey Cross and Andrew Zisserman. Quadric reconstruction from dual-space geometry. *Proceedings of the IEEE International Conference on Computer Vision*, pages 25–31, 1998. [2](#), [5](#)
- [4] G. E. Elsinga, F. Scarano, B. Wieneke, and B. W. Van Oudheusden. Tomographic particle image velocimetry. *Experiments in Fluids*, 41(6):933–947, 2006. [1](#)
- [5] G. E. Elsinga, J. Westerweel, F. Scarano, and M. Novara. On the velocity of ghost particles and the bias errors in Tomographic-PIV. *Experiments in Fluids*, 50(4):825–838, 2011. [2](#)
- [6] Ricardo Fabbri and Benjamin Kimia. 3D curve sketch: Flexible curve-based stereo reconstruction and calibration. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1538–1545, 2010. [2](#)
- [7] Olivier Faugeras and Luc Robert. What can two images tell us about a third one? *International Journal of Computer Vision*, 18(1):5–19, 1996. [2](#)
- [8] LaVision GmbH. *DaVis 10.1.2*, 2020. [8](#)
- [9] Micaiah John Muller Hill. On a spherical vortex. *Philosophical Transactions of the Royal Society of London.(A.)*, (185):213–245, 1894. [8](#)
- [10] Kenichi Kanatani and Wu Liu. 3D Interpretation of Conics and Orthogonality. *CVGIP: Image Understanding*, 58(3), 1993. [2](#)
- [11] Katrin Lasinger, Christoph Vogel, and Konrad Schindler. Volumetric Flow Estimation for Incompressible Fluids Using the Stationary Stokes Equations. *Proceedings of the IEEE International Conference on Computer Vision*, 2017-Octob:2584–2592, 2017. [2](#)
- [12] Zhong Li, Yu Ji, Jingyi Yu, and Jinwei Ye. 3D Fluid Flow Reconstruction Using. 4:120–136. [2](#)
- [13] Hans-Gerd Maas. Complexity analysis for the establishment of image correspondences of dense spatial target fields. *International Archives of Photogrammetry and Remote Sensing*, 29:102–107, 1992. [2](#), [3](#)
- [14] H. G. Maas, A. Gruen, and D. Papantoniou. Particle tracking velocimetry in three-dimensional flows - Part 1. Photogrammetric determination of particle coordinates. *Experiments in Fluids*, 15(2):133–146, jul 1993. [1](#), [2](#)
- [15] Matthias Machacek and Thomas Rosgen. Development of a quantitative flow visualization tool for applications in industrial wind tunnels. In *ICIASF 2001 Record, 19th International Congress on Instrumentation in Aerospace Simulation Facilities (Cat. No. 01CH37215)*, pages 125–134. IEEE, 2001. [1](#)
- [16] D. Müller, B. Müller, and U. Renz. Three-dimensional particle-streak tracking (PST) velocity measurements of a heat exchanger inlet flow. *Experiments in Fluids*, 30(6):645–656, 2001. [1](#)
- [17] Matteo Novara, Daniel Schanz, Reinhard Geisler, Sebastian Gesemann, Christina Voss, and Andreas Schröder. Multi-exposed recordings for 3D Lagrangian particle tracking with Multi-Pulse Shake-The-Box. *Experiments in Fluids*, 60(3):1–19, 2019. [2](#)
- [18] Long Quan. Conic Reconstruction and Correspondence from Two Views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(2):1–13, 1996. [2](#), [5](#)
- [19] Daniel Schanz, Sebastian Gesemann, and Andreas Schröder. Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Experiments in Fluids*, 57(5):1–27, 2016. [1](#), [2](#)
- [20] Daniel Schanz, Matteo Novara, and Andreas Schröder. Shake-The-Box particle tracking with variable time-steps in flows with high velocity range (VT-STB). *3rd Workshop and 1st Challenge on Data Assimilation & CFD Processing for PIV and Lagrangian Particle Tracking*, pages 19–20, 2020. [1](#)
- [21] Cordelia Schmid and Andrew Zisserman. Geometry and matching of lines and curves over multiple views. *International Journal of Computer Vision*, 40(3):199–233, 2000. [2](#), [5](#)
- [22] Amnon Shashua, Shai Avidan, and Michael Werman. Trajectory triangulation over conic sections. *Proceedings of the IEEE International Conference on Computer Vision*, 1:330–336, 1999. [2](#)
- [23] Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006. [6](#)
- [24] Huan Wang, Guijin Wang, and Xianting Li. High-performance color sequence particle streak velocimetry for 3D airflow measurement. *Applied Optics*, 57(6):1518, 2018. [1](#)
- [25] Bernhard Wieneke. Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology*, 24(2), 2013. [2](#)
- [26] Jinhui Xiong, Qiang Fu, Ramzi Idoughi, and Wolfgang Heidrich. Reconfigurable rainbow PIV for 3D flow measurement. *IEEE International Conference on Computational Photography, ICCP 2018*, pages 1–9, 2018. [2](#)
- [27] Jinhui Xiong, Ramzi Idoughi, Andres A. Aguirre-Pablo, Abdulrahman B. Aljedaani, Xiong Dun, Qiang Fu, Sigurdur T. Thoroddsen, and Wolfgang Heidrich. Rainbow particle imaging velocimetry for dense 3D fluid velocity imaging. *ACM Transactions on Graphics*, 36(4), 2017. [2](#)