

3D reconstruction of a thin flexible disc in a vortical flow

Ibarra Eric · Verhille Gautier · Bartoli Adrien

Received: 05 Apr 2023 / Accepted: 19 Sep 2023

Abstract To improve our understanding of how deformable objects are transported in flows, it is necessary to develop new experimental tools capable of accurately measuring the evolution of their deformations. We present a reconstruction process for sheet-like objects utilizing Thin-Plate Splines (TPS), providing access to the object's 3D position and deformation over time. We tested the technique on a simple configuration: a thin-heavy-flexible disc within a vortical flow driven by impellers in a cubic tank. The vortical flow field is characterized using Particle Image Velocimetry (PIV) and is seen to be well approximated as a Lamb-Oseen vortex within the volume of reconstruction. The disc is imaged using three cameras, which are calibrated using the pinhole model. The reconstruction process uses shape-from-silhouette to define an initial 3D reconstruction, which is subsequently refined by minimizing a cost function based on physical and visual criteria. This process is shown to be generalizable to other thin geometries, offering a starting point towards studying the dynamics of more complex sheet-like objects, such as plastic pollutants and vegetation.

Keywords particle tracking · anisotropic particles · vortex

1 Introduction

Particles in fluid flows are ubiquitous in nature, industry, and biology: from the transport of sediment, vegetation, organisms, snowflakes, and pollutants in geophysi-

cal flows, down to the scale of cells, proteins, and bacteria. Recently, the fate of plastic pollutants in the ocean has motivated several investigations to understand their drifting due to waves and currents (DiBenedetto et al., 2018) and their fragmentation rate due to hydrodynamical constraints, as seen in Allende et al. (2020) and Brouzet et al. (2021). These last two studies demonstrated a strong relationship between the deformation of fibers and the resulting fragment size distribution. Before it is possible to explore these findings in relation to 2D objects, such as plastic bags and films, a method to reconstruct and characterize the dynamics of complex 2D objects in 3D flows is needed.

We present a process to measure the deformations of 2D objects and test its efficacy on a simple case where a smooth, thin, flexible silicone disc is placed in a vortical flow. Our process represents a deformed 2D object by a Thin-Plate Spline (TPS) (Duchon, 1976), composed of three 2D to 1D functions sharing their source control points. The TPS maps points from a reference 2D frame, representing the object laid flat, to their 3D position. TPS fitting techniques have been effectively used in medical applications in the representation of organ tissues and the detection of anomalies (Bookstein, 1989). A computationally efficient and readily available TPS creation and fitting technique is provided in the toolset created by Bartoli et al. (2010), which also provides the built-in ability to calculate a range of desired values needed to characterize the deformation of the disc. Our process would help one to characterize the Lagrangian transport of elastic particles in flows. Furthermore, the presented work can be generalized to measure the deformations of other flexible objects, such as the dynamic ecosystem of kelp forests, the migration and locomotion of jellyfish swarms, or aggregations of plas-

Eric Ibarra
E-mail: eric.IBARRA@univ-amu.fr

Gautier Verhille
E-mail: gautier.verhille@irphe.univ-mrs.fr


tic debris (fabrics, bags, bottles) in large-scale oceanic gyres. Building a collection of resources for the investigation of these flexible sheet-like particles would allow for a new range of problems to be explored in fundamental flows that will, in turn, better enable the study of richer, complex large-scale questions.

This paper presents our work to reconstruct a thin-heavy-flexible disc in a vortical flow. The experimental set-up is described first, along with details about the flow and disc characteristics. Then, the general process is presented, followed by a discussion of the resulting 3D reconstruction. The paper ends with concluding remarks and our parting thoughts.


2 Experimental set-up

The experiments are performed in a 60 cm cubic tank filled with water. The flow can be illuminated or viewed through the acrylic windows on each face of the cube. The generation of the flow field in the tank is provided by 8 rotors fitted with impeller disks with diameter of 17 cm equipped with six straight blades 5 mm in height. Each impeller is located at a vertex and points towards the center of the cube. The rotation rate Ω_F of each impeller can be set independently between 5 and 19 Hz. In the present study, a single pair of diagonally opposing impellers are actuated at the same frequency with opposite directions. Video acquisition is carried out by three synchronized high-speed cameras (VEO 710L) each equipped with a ZEISS Milvus 2/100M ZF.2. The images used for the reconstruction were recorded at a sampling rate of 3000 Hz with an exposure of 330 μ s. The cameras are positioned to provide 3 approximately orthogonal views of the volume of interest. An illustration of the relative positions of the tank, motors, and cameras is provided in figure 1(a). Through the windows on the opposite side of the cameras, the backlighting for the captures is provided by three LED panels. The LED panels are energized by DC power supplies to ensure the illumination is continuous. The reconstruction measurements of the flexible disc have been performed in a cubic volume of nearly 10 cm \times 10 cm \times 10 cm at the center of the tank.

The flow field of the vortex was characterized for a range of forcing frequencies, Ω_F , using multi-planar PIV measurement. For these measurements, the tank is seeded with polyamide seeding particles of diameter 50 μ m (Dantec Dynamics). The PIV images were captured using a Dantec Flowsense EO camera and illumination was provided by a Nd:YAG pulsed laser 532 nm. The investigation was carried out using DynamicStudio, for



(a) Experimental set-up



(b) PIV planes

Fig. 1 Figure 1(a) provides an illustration of the tank, motors, and position of the cameras. Figure 1(b) shows the location of the planes where PIV measurements were taken with respect to the physical center of the tank in the laboratory frame of reference. The blue dotted line [\dots] is provided to indicate the expected position of the vortex centerline along the diagonal of the tank

both the capture and PIV data analysis, with post-processing carried out in MATLAB. Figure 1(b) provides an illustration of the positions of the laser sheet used for the characterization. As the axis of the vortex is not aligned with any of the viewing planes of the tank, the planes of the PIV measurements are not orthogonal to the axis of the vortex. However, the azimuthal velocity of the vortex can be inferred by adjust-

ing for difference in alignment between the measured plane's normal and the axis of the vortex. Due to the meandering of the vortex core across image pairs, an averaged value for the velocity field was achieved by a spatial shift in the velocity fields to align the two local velocity magnitude maxima across each respective run (Heyes et al., 2004). Figure 2 provides a sample of a comparison between the time-averaging approaches of the velocity magnitude for motor forcing frequency of 9 Hz. Figure 2(a) is the result of time-averaging the velocity magnitude in the laboratory frame of reference. The signatures of the core of the vortex are more readily identifiable when averaging in the vortex frame of reference, shown in figure 2(b). The flow in the center of the tank is modelled (to first order) as a Lamb-Oseen vortex (Lamb, 1945). The values of circulation, Γ , and scale of the vortex core, σ^2 , plotted in figure 2(c) are averaged across the measurement planes with bar half-length representing the standard deviation. Both parameters are found to evolve monotonically with the forcing frequency. The standard deviations are provided to give a sense of the similarity between the flows across the measurement planes at each frequency. The velocity fields, averaged in the vortex frame of reference, were consistent across the multiple planes of measurement, shown in figure 1(b), at each forcing frequency. The vortical flow does not appear to vary along the axis of the vortex within our volume of interest and the measured flow field agrees reasonably well with that of a Lamb-Oseen vortex.

The smooth flexible disc is made of EC 13 silicone, from Esprit Composite, with a specific gravity of 1.08 and Young's modulus of $E = 130 \pm 20$ kPa. To make the silicone semi-transparent, a small amount of fine pigment was well mixed into the silicone prior to curing. To form the silicone sheet, the silicone is poured and pulled onto a glass substrate (within a known gap height determining the thickness of the sheet) and allowed to cure. The roughness of the the disc is expected to be on the order of that of the glass substrate ($\lesssim 10 \mu\text{m}$) used for the molding. The disc is cut using a 16 mm inner-diameter hollow punch (BOEHM) from a $h = 312 \pm 7 \mu\text{m}$ thick silicone sheet.

3 General process

The process required to measure the position and shape of the flexible disc is separated into four main steps, covered in the following sections. Step 1 calibrates the cameras. Step 2 reconstructs an initial estimate of the disc's 3D perimeter. Step 3 creates an initial TPS to model the surface of the disc. Step 4 optimizes the fit




Fig. 2 Time averaged velocity magnitudes from 2D PIV measurements from center plane ($x = 0$). Figure 2(a) is the result of time averaging in the lab frame of reference. Figure 2(b) is the result of time averaging after accounting for the spatial shift of the vortex core. Figure 2(c) provides the estimated Lamb-Oseen vortex parameters obtained by averaging across the 4 measurement planes, with bar half-length representing the standard deviation

based on desired physical (bending and stretching energies) and visual (reprojection quality) criteria.

3.1 Camera calibration

To quantify the 3D position of the disc from images, a calibration for the three cameras must be created.


The cameras are calibrated by approximating their optics with the pinhole model. The 11 model parameters represent the position and orientation of the camera in the laboratory frame, two scaling factors, the skew parameter and the principal point. They are determined through a specific process based on multiple correspondences of 3D and 2D points (Faugeras and Luong (2001); Hartley and Zisserman (2003); Verhille and Bartoli (2016)).

The calibration is performed within the fluid to account for the variation of the refractive index through the multiple flat interfaces that light traverses during a capture (water/acrylic/air) (Agrawal et al., 2012). In our implementation, the point correspondences are obtained by tracking a target sphere through a sweep of the volume of interest. The center of the sphere is positioned by a set of 3 linear translation stages in an assembly. Each linear stage is actuated by a corresponding stepper motor and the direction of each linear stage is fixed to move orthogonal to one another. As the translation apparatus moves to each known position, the 3 cameras simultaneously capture an image of the sphere. The projection of the sphere onto each camera’s image plane is well approximated¹ as a circle, therefore the MATLAB function `imfindcircles` is used to estimate the coordinates of the center of the sphere in each camera’s image plane. The 11 parameters are then determined from the correspondences between the known 3D coordinates of the sphere’s center in laboratory space and the measured 2D coordinates of its projection in the image planes. This calibration step is only required once per adjustment of the optics.


3.2 Initial perimeter reconstruction

Working with the calibrated cameras, a set of raw images is taken and used to isolate the pixel locations related to the disc’s perimeter. These pixel locations are then used to estimate the 3D location of the disc’s perimeter. Figures 3(a-c) present cropped views of the disc as seen from the three camera views. The backlit, semi-transparent disc reduces the amount of light a camera receives from the LED panel it is facing. The pixel-wise intensity value has an inverse connection with the integral of the disc’s volume along the line of sight between the light source and camera sensor. Sharp spatial changes of the intensity value in the image can indicate an edge of the disc, which is a useful cue for


¹ Due to the quality of the lenses used and the sphere’s size in combination with its distance from each of the cameras, image distortions and the effect of local perspective are negligible.



(a) View 1



(b) View 2



(c) View 3

Fig. 3 Cropped raw image captures from the three orthogonal views of the disc. The cropped region is 155 px \times 155 px in size

determining the disc’s perimeter. Changes in the intensity value can also indicate local variations of the surface orientation. However, in this case, they tend to be smooth variations.

Figure 4(a) shows the result of view 1 processed by a Canny edge detection; this is identically applied to the other views. The process provides a binary im-

age with values of 1 where a strong intensity gradient was found. Defocused particles outside the volume of interest's focal depth are passively removed due to their weaker edges. The position of points corresponding to the perimeter of the disc are estimated using a volume based shape-from-silhouette method (Cheung et al. (2005); De La Rosa Zambrano et al. (2018)). The volume of reconstruction is divided into cubic voxels of $750 \mu\text{m}$ in length. Each voxel is projected onto each image plane. The voxels with projections that fall on locations of strong image edges in all three image planes are preserved and the non-matching points are discarded. This predominately preserves the points that represent the disc's perimeter, while suppressing strong image edges incurred by bends along the surface of the disc that are not consistently seen in all views. The retained set of perimeter points are reconstructed in the experimental 3D frame of reference, visualized in figure 4(b) by red dots. Image masks are also constructed using the points of the silhouette's edge, \underline{s} . The masks are used to classify the two regions for each view: inside the disc's image and outside.

3.3 Initial surface reconstruction

We define a TPS to represent the initial surface reconstruction. The TPS is obtained from 2D to 3D point correspondences, which we create on the perimeter. To accomplish this, the coordinates of the reconstructed perimeter points are resampled at equispaced intervals along its closed curve. These resampled 3D points are mapped to equispaced points along the perimeter of the disc's 2D reference model: these points are referred to as boundary control points. This correspondence of the control points defines an initial TPS estimating the disc's surface.

The TPS, ϕ , is a parametric mapping taking points from a 2D space, $\underline{X} = (X, Y)$, to any desired dimension. We instantiate it to create a surface, $\underline{x} = (x, y, z)$. Fitting the TPS requires knowledge of corresponding points from the 2D model reference frame to their position in 3D space. These are referred to as control points. The TPS function will satisfy the constraints imposed by the control points while minimizing the *integral bending norm* (Duchon, 1976):

$$I[\phi] = \iint_{\mathbf{R}^2} \left((\phi_{XX})^2 + 2(\phi_{XY})^2 + (\phi_{YY})^2 \right) dS \quad (1)$$

The subscripts of ϕ denote the partial derivative with respect to the noted reference coordinate; dS represents a surface element in the domain \mathbf{R}^2 . The 2D model of




Fig. 4 Figure 4(a) provides the result after a Canny edge detection process is applied on view 1. In figure 4(b), red dots [•] show the original reconstructed perimeter points using shape-from-silhouette, the magenta line [⊕] is the connecting line between the initial reconstructed control points, and the goldenrod mesh [■] provides an illustration of the disc's initial reconstruction in space. The axis is with respect to the experimental coordinate system, as measured from camera calibration. Figure 4(c) illustrates the disc reference model, the mesh points are represented by black dots [•], the reference boundary control points are shown as magenta circles [⊙], internal control points are represented by magenta squares [■], and the model's perimeter is outlined by blue dots [•].

the disc is a circle with a radius equal to the physical radius of the disc in the experiment. Figure 4(c) provides an illustration of the reference model. The area of the disc model is sampled as a regular mesh with grid spacing of $\Delta x = \Delta y = \frac{R}{N}$. Results using a value of $N = 33$ are presented here, corresponding to a total number of mesh nodes of $M = 793$, represented by black dots in figure 4(c).

For this work, the control points in the 2D reference frame are sampled at equispaced locations along the perimeter of the model disc. Since we do not have access to an unambiguous one-to-one mapping of the material points of the physical disc, we cannot measure the spinning about the principal axis of the disc. However, this ambiguity in the mapping requirements leads to a beneficial invariance in the selection of the boundary control points, which are only required to be in order and equispaced along the length of the perimeter. The set of reconstructed 3D points are parameterized and fitted using the spherical coordinates. The 3D boundary control points are subsequently sampled at equispaced locations along the length of the 3D perimeter fit to match those from the 2D reference model. Because the control points are related between the reference and the experimental frame, the initial TPS mapping can be solved for using the toolset provided in Bartoli et al. (2010). The number and location of the control points and mesh points are set per investigation. The kernels used for the mapping and evaluation of derivatives on the surface are calculated once. These kernels are matrices that formulate a system of equations to efficiently calculate their respective outputs once given the 3D locations of the control points.

Figure 4(b) provides an illustration of the initial mapped model mesh in goldenrod. The magenta line segments link the locations of the control points, shown as magenta circles. This preliminary fit is sampled to initialize the internal control points. The location of the internal control points on the reference model are marked as magenta squares in figure 4(c). The reprojection of the initial fit, as shown in figure 5, is deficient in visually reproducing the image of the disc in all 3 views, most notably in figure 5(a). To produce a more faithful reconstruction, the 3D position of the boundary and internal control points are numerically optimized to iterate towards desired physical and visual criteria.

3.4 Reconstruction refinement

The optimization process consists of verifying quantifiable criteria to improve the quality of the mapping. These criteria come from both physical characteristics




Fig. 5 Reprojection of the initial boundary control points [\ominus] and mesh [\boxplus].

of the disc and visually derived metrics based on reprojections of the mapped mesh onto the source images.

The physical disc in its equilibrium position is flat. As such, the reference configuration for the disc model is an unbent and unstretched surface. Once in the flow, the physical disc adopts a shape that balances the hydrodynamical forces and the internal restorative forces,

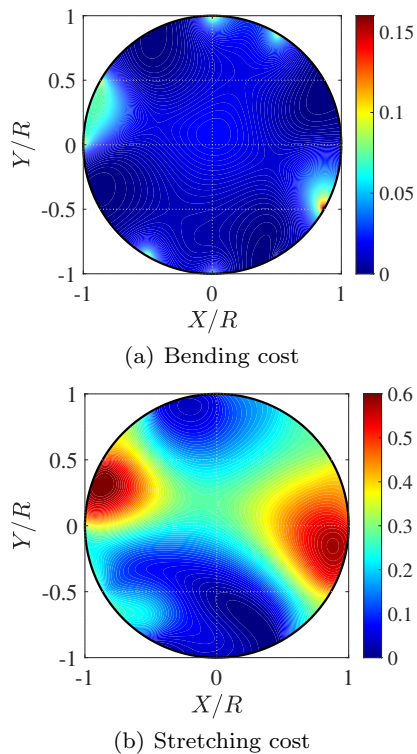


Fig. 6 Local bending cost figure 6(a) and stretching cost 6(b) of the initial TPS fit

causing the disc to tend back to its equilibrium shape which minimizes its elastic energy.

Bending cost. The disc's surface is observed as smooth, with no cusps, while advected by the flow. One cost involved in the optimization aims to minimize the bending energy of the particle. Mapped control points, along with derivative kernels, are used to compute the principle curvatures at all points of the mapped mesh. Subscript $[i]$ is used to denote a local mesh point value. A nodal bending cost value is defined as:

$$c_{\text{bend}}[\phi, \underline{x}_{[i]}] = [\kappa_1^2 + \kappa_2^2 + 2p\kappa_1\kappa_2]_{[i]} \quad (2)$$

where p is the Poisson's ratio of the disc and $\kappa_{1,2}$ represent the principal curvatures. This cost is proportionally equivalent to the bending energy per unit area as shown in Verbeek and Van Vliet (1993). For silicone, a working value of $p = 0.5$ is used. The bending cost is calculated as the mean cost over all the nodes in the mesh:

$$c_{\text{bend}}[\phi] = \frac{1}{M} \sum_{i=1}^M (c_{\text{bend}}[\phi, \underline{x}_{[i]}]) \quad (3)$$

Stretching cost. By minimizing the stretching cost of the model, the optimization tends towards an isometric

mapping. This isometry would preserve the local distance between two points on the surface, taken from the model to the mapped frame of reference. This is done through the stretching cost:

$$c_{\text{stch}}[\phi, \underline{x}_{[i]}] = \left[\left((\phi_X)^2 - 1 \right)^2 + \left((\phi_Y)^2 - 1 \right)^2 \right]_{[i]} \quad (4)$$

As with the bending cost, the mean cost over all the nodes in the mesh is calculated:

$$c_{\text{stch}}[\phi] = \frac{1}{M} \sum_{i=1}^M c_{\text{stch}}[\phi, \underline{x}_{[i]}] \quad (5)$$

This value's trend has parity with the summed strain over the surface. The local strain, ϵ_j , is calculated from the Euclidean distances of neighboring points in the mesh grid between the model, ΔR_j and mapped configurations, Δr_j .

$$\epsilon_{[j]} = \left[\frac{|\Delta r - \Delta R|}{\Delta R} \right]_{[j]} \quad (6)$$

The subscript $[j]$ denotes a unique neighborly connection. Each internal point of the model mesh has 4 Cartesian neighbors considered. With mirrored connections excluded, the unique number of connections is less than $4M$.

Reprojection cost. The model mesh is mapped to the estimated physical 3D position and conformation in the laboratory frame. This mapped mesh is then reprojected onto the image plane of each camera. The image masks of the disc are used to verify the TPS fit's reprojection quality in each view. This is done by determining if the mapped mesh points are reprojected onto the image coinciding with a location outside or inside the area of the disc. For a particular image space there are T silhouette points $\underline{s}_{[t]}$ delineating the region outside and inside the disc, P mesh points reprojecting outside $\underline{h}_{[p]}$, and Q number of mesh points reprojecting inside $\underline{g}_{[q]}$, where $P+Q = M$, the total number of mesh points. The values of $\underline{h}_{[p]}$ and $\underline{g}_{[q]}$ and their quantities depend on the TPS ϕ .

Every outer point $\underline{h}_{[p]}$ is penalized by its Euclidean distance to the nearest point $\underline{s}_{[t]}$. The mean of these distances for all the outer points is taken and normalized by the disc scale term δ of the image:

$$c_{\text{out},[k]}[\phi] = \left[\frac{1}{\delta P} \sum_{p=1}^P \min \left(\text{dist}(\underline{h}_{[p]}, \underline{s}_{[1:T]}) \right) \right]_{[k]} \quad (7)$$

The subscript $[k]$ is used to denote a particular image space, where $k \in \{1, 2, 3\}$. Per image, the disc scale

term is calculated by taking the mean of the maximum distances of all points on the silhouette to another point on the silhouette:

$$\delta_{[k]} = \left[\frac{1}{T} \sum_{t=1}^T \max \left(\text{dist}(\underline{s}_{[t]}, \underline{s}_{[1:T]}) \right) \right]_{[k]} \quad (8)$$

Each cost associated with points reprojected outside the disc's image space per view is squared then summed over all views, which results in a scalar cost denoted as $\mathcal{C}_{\text{out}}[\phi]$:

$$\mathcal{C}_{\text{out}}[\phi] = \sum_{k=1}^3 (c_{\text{out},[k]}[\phi])^2 \quad (9)$$

The under-utilization of the disc's inner region is penalized by using the distance of each inner $\underline{s}_{[t]}$ point to its nearest $\underline{g}_{[q]}$ as:

$$c_{\text{in},[k]}[\phi] = \left[\frac{1}{\delta T} \sum_{t=1}^T \min \left(\text{dist}(\underline{s}_{[t]}, \underline{g}_{[1:Q]}) \right) \right]_{[k]} \quad (10)$$

$$\mathcal{C}_{\text{in}}[\phi] = \sum_{k=1}^3 (c_{\text{in},[k]}[\phi])^2 \quad (11)$$

The cost $\mathcal{C}_{\text{in}}[\phi]$ aids in approaching a fit that reprojects more effectively within the desired area in each image.

Total cost. Each of the four costs is multiplied by its respective weight and then summed together into the total cost to be minimized as:

$$\mathcal{C}_{\text{tot}}[\phi] = \mathcal{W}_{\text{out}}\mathcal{C}_{\text{out}}[\phi] + \mathcal{W}_{\text{in}}\mathcal{C}_{\text{in}}[\phi] + \dots \quad (12)$$

$$\mathcal{W}_{\text{stch}}\mathcal{C}_{\text{stch}}[\phi] + \mathcal{W}_{\text{bend}}\mathcal{C}_{\text{bend}}[\phi]$$

with each weight tuned to have the mapping converge to a desired result. Specifically, the weights were selected to prioritize the enforcement of their respective attributes as they pertain to their relative contributions to the total cost.

More specifically for our application, the elastic energy can be deconstructed into two parts: bending and stretching. The bending and stretching energy are proportional to the bending and stretching costs defined in equation (3) and (5), respectively. From elasticity (Landau et al., 1986, Powers, 2010), we know that for a plate bent by a length δ the force scales as:

$$F_{\text{bend}} \sim EhR \cdot \kappa h \sim EhR \cdot (\delta/R^2)h \quad (13)$$

Similarly, for a plate stretched by a length δ the force scales as:

$$F_{\text{stch}} \sim EhR \cdot \delta/R \quad (14)$$

In both cases, the energy scales as $F \cdot \delta$; therefore, the ratio of the bending over the stretching energy scales as h/R . It follows that the ratio of the bending and stretching weights should then scale as:

$$\mathcal{W}_{\text{bend}}/\mathcal{W}_{\text{stch}} \sim F_{\text{bend}}/F_{\text{stch}} \sim h/R \quad (15)$$

In practice, the bending and stretching contribution, $\mathcal{W}_{[i]}\mathcal{C}_{[i]}$, to the total cost, \mathcal{C}_{tot} , should scale in relation to their expected contribution to the total elastic energy of the disc. For our application, the the bending weight is lower than the stretching weight by a scalar factor of this ratio. This relative weighting allows the refinement step to move away from the initial reconstruction, towards one that agrees the physical scaling of the costs' respective energy analogues. Because we wish to enforce a faithful reprojected of the disc within the silhouette of the disc in the 3 views, the reprojected weights are set to be even higher than that of stretching.

4 Discussion

Bending. Figures 6(a) and 7(a) allow for comparison of the local bending cost between the initial and optimized fits. The initial fit is the minimization of the integral bending norm (Duchon, 1976); as would be expected, figure 6(a) shows its internal local bending costs are low and smoothly evolving. The local bending cost for the optimized fit, in figure 7(a), shows a band of higher bending cost along a diameter ($\sim 10^\circ$ clockwise from vertical axis). The highly localized peaks in the curvature at the control points are characteristic of fitting schemes that are enforced with the second derivative.

Stretching. Figures 6(b) and 7(b) present a comparison of the local stretching cost between the initial and optimized mapping. The initial fit has no obligation to produce an isometric mapping, which can be seen as regions of high local stretching cost values in figure 6(b). The mean local strain value is $\mu_\epsilon = 0.173$ with a standard deviation of $\sigma_\epsilon = 0.100$ for this initial mapping. Figure 7(b) illustrates that the optimized mapping drastically lowers the local stretching cost over the entirety of the disc. The mean local strain value is $\mu_\epsilon = 0.0107$ with a standard deviation of $\sigma_\epsilon = 0.0092$ for this optimized mapping. A change of color mapping used between figures 6(b) and 7(b) draws attention to the order of magnitude difference while preserving the visibility of the cost function's spatial variations.




Fig. 7 Local bending cost figure 7(a) and stretching cost 7(b) of the optimized TPS fit. Different color-maps between figures 6(b) and 7(b) are used to facilitate the inspection of the variations while emphasizing the difference of scales between the two results. In figure 7(c), red dots [•] show the original reconstructed perimeter points using shape-from-silhouette, the blue line [—] is the perimeter of the optimized reconstruction, and the cyan mesh [■] provides an illustration of the reconstructed disc in space. The axis is with respect to the experimental coordinate system as measured from the camera calibration

Reprojection. Figure 5(a-c) presents the reprojection of the initial boundary control points and a mesh of the fit on the raw capture of each view. The reprojection of the initial fit is in poor agreement with the area of the disc in the raw captures. This can be most notably seen in 5(a), as the reprojection under-occupies the disc area. Figure 8 provides a visual illustration associated




Fig. 8 Color bar denotes pixel distance from silhouette points [—] to their respective nearest mesh point reprojected inside. This distance is used for the evaluation of the optimization cost, $\mathcal{C}_{in}[\phi]$

with \mathcal{C}_{in} to compare the initial and final distances of silhouette points to the points reprojected inside the area of the disc in the image of view 1. The silhouette points are linked to their nearest internally reprojected point and colored according to the magnitude of their distance. Figure 8(a) illustrates an under-occupation of the internal area by the reprojected initial points, with distances reaching 32 px. Figure 8(b) shows a significant improvement with distances falling below 5.5 px. A visual illustration associated with \mathcal{C}_{out} is not provided because the initial reconstruction of this sample falls reasonably well within the image area of the disc in the 3 views and therefore does not show an appreciable improvement once optimized. A correlation can be observed between the density of the mesh points and the relatively lowered local intensity value of the image. The mesh points serve as a proxy for the physical location of the disc's volume; the higher density of points is analogous to light having to travel a longer distance

through the semi-transparent material of the disc in a given view. Figure 9(a-c) presents the reprojection of the optimized fit on the raw capture of each view. When compared to the initial fit, figure 5(a-c), the optimized fit occupies a fuller internal range of the disc in the image. In contrast to the other costs, the bending cost globally increases after the optimization. This was expected, as the TPS is the solution that minimizes the integral bending norm based on the location of the initial boundary control points. After optimization, one can see that the disc is predominately bent along one diameter aligned with the likely local axis of the vortex. The high curvatures localized around the control points are not physical in this case, as the curvature is expected to evolve at the scale of the radius of the disc. These peaks are artifacts stemming from the finite number of control points used for the model, similar to using a discrete number of pointed poles in trying to pitch a tent. To prevent these non-physical local peaks, two options are possible. The first one is to spatially filter the curvature after the optimization step; the second is to increase the stiffness of the TPS used to fit the data. Both cases require some *a priori* on the expected deformation. Finally, by comparing the bending cost and the stretching cost, one can extrapolate that the disc is an effectively inextensible area.

Figure 7(c) provides a 3D illustration of the finalized mesh mapping. To begin investigating the evolution of the flexible disc, measurements of the disc’s position, orientation, and deformation are required. Using the proposed reconstruction process, the disc is tracked through a short video sequence. Figure 10(a) uses black dots to illustrate the path traced by the center of the mapped disc model in time alongside the reconstruction of the model at the start and end of the trajectory in this sequence. Figure 10(b) presents the time evolution of the center of the disc with each coordinate of trajectory normalized by the disc’s radius, R , and C_{bend} . Reconstructions with cost values two standard deviations away from the mean in time are considered sub-optimal, treated as outliers, and naively excluded. These exclusions can be noted as irregular spacing in the time evolution. Of the 512 frames reconstructed in the truncated series, 58 frames were excluded with this simplistic metric.

To demonstrate the reconstruction’s sensitivity to the blurring produced by a particle being out of the focal depth of a given view, we applied an artificial Gaussian blur to our images with a range of blurring strengths and then applied the reconstruction process to the blurred frames. This artificial blur was applied using the MATLAB function `imgaussfilt`. We verified the sensitiv-




Fig. 9 Reprojection of the optimized model’s perimeter [—] and mesh [■■] of resultant mapping

ity of the center of gravity of the disc, bending cost, and stretching cost as σ was varied from 0.3 to 3. Per frame reconstructed, the center of gravity of the reconstructed particle progressively shifted as the intensity of the blur strengthened. Whereas the center of gravity of the reconstructed disc was measured as the mean of all the mesh points’ positions, the bending cost is C_{bend} , as presented in equation (3), and the stretching




Fig. 10 A 3D visualization of the disc’s trajectory is presented in figure 10(a). The measured position of the center of the disc is plotted with black dots [•]. The coordinate system \underline{x}_0 is with respect to the likely axis of the vortex, denoted with a black dashed line [—]. The blue [■] and green [■] mapped model meshes, with center points, are provided for the first and last reconstructions of this truncated series. The evolution in time of the measured position of the center of the disc normalized by its radius, R , and C_{bend} is provided in figure 10(b).

cost is C_{stch} , as presented in equation (5). The mean and standard deviation values are evaluated over all frames and blur strengths. Per frame, the reconstruction process is applied to the raw images, without blurring, to produce a reference reconstruction. The shift distance between the center of gravity of the blurred reconstruction and that of the reference reconstruction is measured for each respective frame, for each blur strength. This difference is normalized by the disc’s radius. The mean value of this spatial shift of the center of gravity is 8.22% of the disc’s radius. The mean bending cost magnitude multiplied by R^2 , for normalization, is 2.44 with a standard deviation of 1.98×10^{-1} . The mean stretching cost is 2.36×10^{-3} with a standard deviation of 5.81×10^{-4} . This sensitivity characterization to uniform blurring does not provide insight to the process’s sensitivity to motion blur resulting from large exposure times, but it does provide some insight into the reconstruction process’s sensitivity to blur in general, such as when the particle in the view of all three cameras falls outside of their focal depth.

The accuracy of the reconstruction was verified using a reference disc. This reference disc was fabricated using a stereolithography 3D printer with $100 \mu\text{m}$ resolution. The resin used was semi-transparent with a Young’s modulus of $E = \mathcal{O}(1)$ GPa. The 3D printed disc was fabricated to have a radius of 8 mm, a constant radius of curvature of 8 mm for the concave side, and thickness of $750 \mu\text{m}$. A laser profilometer was used to measure the convex side of the reference disc after the curing process, which assessed the upper bound of the disc’s radius of curvature as 8.86 mm with a mean absolute error of $27 \mu\text{m}$. The relatively large thickness, compared to the thickness of the flexible disc, was selected to ensure the rigidity and dimensions of the disc. The reference disc was measured using the same experimental set-up and process as presented above. The disc’s radius, based on the the reconstructed perimeter, was 5.89% above the expected value according to the disc fabrication specifications. Similarly, the radius of the disc based on the measured surface area of the reconstructed disc had an error of 6.32%. Using the reconstruction process, the measured radius of curvature of the reference disc was 8.77 mm. These errors are on the order of magnitude of the disc’s aspect ratio, between its thickness and radius, $h/R \simeq 10\%$. This suggests that the finite thickness of the disc might need to be taken into consideration to improve the reconstruction. This source of error is elaborated upon in the closing remarks.

5 Closing remarks

The presented framework provides a robust starting point for the investigation of smooth flexible sheets in various flows. This preliminary work focuses on a circular disc model, yet the principles behind the process are not limited to this singular geometry. This work can be extended to various thin particles and models, given an adequate matching of control points between the model and the target. Figure 11 demonstrates the results from the same process applied to square flexible sheets. Tracking the discernible features of the particle’s geometry (i.e., the corners of the flexible square sheet) provides a basis to measure rotation of the particle about its principal axis. This extension to different geometries opens our investigation to explore the effect that shape might have on the observed equilibrium position or preferential shape alignment with the vortex; it additionally presents the possibility of studying the dynamics of similar anisotropic particles in other forced flows.

The current study focuses only on thin flexible particles fabricated to be uniform in thickness and smooth, with no surface singularities. Smooth, for the purpose of these measurements, indicates that no roughness or imperfections are discernible at the resolution of the recording. In general, the resolution of the recording depends on the desired application, requiring a compromise between spatial resolution and the field of view. It is possible that our 3D reconstruction process could be used to reconstruct more complex particles (inhomogeneous thickness/material, roughness), but this has currently not been tested. Other possible approaches to handle reconstructing complex geometries include utilizing digital image correlation (DIC) on surface patterns or textures (Pan, 2018) or using Fourier Transformation Profilometry (FTP) with structured illumination (Su and Zhang, 2010).

The TPS mapping has a parameter that allows the user to tune the stiffness of the mapping, as mentioned briefly in the previous section. For the work presented here, the parameter was preserved at the default value of unity. This parameter should be definable based on the mechanical properties of the disc: the modulus of the material and its aspect ratio.

The model approximates the physical disc as a surface of zero thickness; this simplification can make it difficult for the model’s reprojection to properly occupy the image of a physical volume. Using the process described in this paper, lamination can simply be placed normal to the surface of the fit at a distance commensurate with its thickness, h , to be used towards the





Fig. 11 Reprojection of the optimized mapped model perimeter [—] and mapped model mesh [] for a square sheet particle. The cropped region is 343 px \times 343 px in size

visual reprojection costs. This strategy was unnecessary in our investigation due to the aspect ratio of the disc, $\approx \mathcal{O}(10^2)$, and our image resolution. Within the volume of reconstruction, the resolution of the disc provided sufficient points to well characterize its silhouette while maintaining a thickness on the order of $\mathcal{O}(1)$ px, which allowed the thin model approximation to perform effectively.

The efficiency of the optimization step and quality of the resultant fit is dependent on the accuracy of the target information. A poor estimation of the disc’s silhouette, and the subsequent mask, results in a poor mapping. Particular care is taken to tune the opacity of the silicone during fabrication to obtain a desired signal-to-noise ratio in the images: one that allows for discernible detail when the disc is occluded while still providing a strong contrast between the disc and background. If we were to imprint the disc with a known texture or designed pattern, DIC techniques would be able to circumvent this difficulty while also facilitating a unique fitting of control points to the features. An unambiguous one-to-one mapping would allow the ability to track the spinning of circular discs. DIC approaches are promising for applications with a limited number of particles, where unique patterning can be used to reconstruct complex surface features. However, for applications dealing with a high number of flexible particles, this method would be limited due to the substantial volume of fabrication required and the challenges of measuring the numerous unique, unstretched reference patterns.

When working with a large number of particles, different complications arise in the event of multiple discs in the region of interest. When sufficiently far apart, multiple discs can be treated independently from one another. Discs overlapping in images create a segmentation challenge. Both DIC and FTP methods would face considerable challenges in experiments where particle-particle occlusions or self occlusion (resulting from strong deformations of the particle) are prevalent. In our experiments, the semi-transparent property of the smooth thin flexible disc allows an image from a single view to provide information about the location of the perimeter of the particle, even when self-occlusion is present. This feature is desired for tracking multiple discs, where occlusions are more prevalent. Yet, the ambiguity of a composite region may cause the optimization step to create a mapping that could prioritize the reprojection of the mesh points onto a region not associated with the corresponding disc. Accurately segmenting regions, such as silhouette points pertaining to separate discs, is an issue. To be able to process multi-disc flows, future work will require the incorporation of a robust segmentation routine. Machine learning, as applied to shape segmentation for autonomous driving and bio-medical applications (Ronneberger et al., 2015), provides a framework for constructing an efficient tool for the segmentation of overlapping discs. The range of simple geometries used for our investigation allows a tractable creation of synthetic data based

on our particle models, as described in Rozantsev et al. (2015) and in Nikolenko (2021), to provide an adequate dataset needed to train a deep neural network. Advanced tools such as domain adaptation will then be needed to make sure that the neural network copes with real images.

6 Acknowledgements

The authors would like thank our colleagues Eric Bertrand and Marie-Julie Dalbe for aiding us with the 3D-printing and laser profilometer measurement of our reference disc. This work received support from the French government under the France 2030 investment plan, as part of the Initiative d’Excellence d’Aix-Marseille Université - A*MIDEX - AMX-19-IET-010. This work was carried out in the framework of NetFlex Project (ANR-21-CE30-0040) funded by the French National Research Agency (ANR).

7 Declarations

7.1 Ethical Approval

[not applicable]

7.2 Competing interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

7.3 Authors' contributions

E. Ibarra carried out the experiments and the data analysis, generated figures for the manuscript, and wrote the original manuscript. G. Verhille and A. Bartoli contributed by providing coding toolsets, guidance, and discussions needed for the work presented. All authors contributed by reviewing the manuscript.

7.4 Funding

This work received support from the French government under the France 2030 investment plan, as part of the Initiative d'Excellence d'Aix-Marseille Université - A*MIDEX - AMX-19-IET-010. This work was carried out in the framework of NetFlex Project (ANR-21-CE30-0040) funded by the French National Research Agency (ANR).

7.5 Availability of data and materials

Data used in the production of this paper can be obtained by contacting E. Ibarra or G. Verhille. The MATLAB toolset used to generate Thin-Plate Splines, as written by A. Bartoli, is available under **Deformable Image Registration** - [GTPSW: Generalized Thin-Plate Spline Warps] at:

http://encov.ip.uca.fr/ab/code_and_datasets/

References

- A. Agrawal, S. Ramalingam, Y. Taguchi, and V. Chari. A theory of multi-layer flat refractive geometry. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3346–3353. IEEE, 2012. doi: 10.1109/CVPR.2012.6248073.
- S. Allende, C. Henry, and J. Bec. Dynamics and fragmentation of small inextensible fibres in turbulence. *Philosophical Transactions of the Royal Society A*, 378(2175):20190398, 2020. doi: 10.1098/rsta.2019.0398.
- A. Bartoli, M. Perriollat, and S. Chambon. Generalized thin-plate spline warps. *International Journal of Computer Vision*, 88:85–110, 2010. doi: <https://doi.org/10.1007/s11263-009-0303-4>.
- F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on pattern analysis and machine intelligence*, 11(6):567–585, 1989. doi: 10.1109/34.24792.
- C. Brouzet, R. Guiné, M.-J. Dalbe, B. Favier, N. Vandenberghe, E. Villermaux, and G. Verhille. Laboratory model for plastic fragmentation in the turbulent ocean. *Physical Review Fluids*, 6(2):024601, 2021. doi: 10.1103/PhysRevFluids.6.024601.
- K.-M. Cheung, S. Baker, and T. Kanade. Shape-from-silhouette across time part i: Theory and algorithms. *International Journal of Computer Vision*, 62:221–247, 2005. doi: <https://doi.org/10.1007/s11263-005-4881-5>.
- H. M. De La Rosa Zambrano, G. Verhille, and P. Le Gal. Fragmentation of magnetic particle aggregates in turbulence. *Phys. Rev. Fluids*, 3:084605, 2018. doi: 10.1103/PhysRevFluids.3.084605.
- M. H. DiBenedetto, N. T. Ouellette, and J. R. Kosoff. Transport of anisotropic particles under waves. *Journal of Fluid Mechanics*, 837:320–340, 2018. doi: 10.1017/jfm.2017.853.
- J. Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *Revue française d'automatique, informatique, recherche opérationnelle. Analyse numérique*, 10(R3):5–12, 1976.
- O. Faugeras and Q.-T. Luong. *The geometry of multiple images: the laws that govern the formation of multiple images of a scene and some of their applications*. MIT press, 2001.
- R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- A. Heyes, R. Jones, and D. Smith. Wandering of wing-tip vortices. In *Proceedings of the 12th international symposium on applications of laser techniques to fluid mechanics*, pages 35–3, 2004.

- H. Lamb. Hydrodynamics. *Victoria University of Manchester, Cambridge University Press, Dover Publications Inc., New York, Printed in the USA, Card Number 46-1891, ISBN: 0 486 60256 7*, 1945.
- L. D. Landau, E. M. Lifshitz, A. M. Kosevich, and L. P. Pitaevskii. *Theory of elasticity: volume 7*, volume 7. Elsevier, 1986.
- S. I. Nikolenko. *Synthetic data for deep learning*, volume 174. Springer, 2021.
- B. Pan. Digital image correlation for surface deformation measurement: historical developments, recent advances and future goals. *Measurement Science and Technology*, 29(8):082001, 2018.
- T. R. Powers. Dynamics of filaments and membranes in a viscous fluid. *Reviews of Modern Physics*, 82(2):1607, 2010. doi: 10.1103/RevModPhys.82.1607.
- O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. doi: 10.1007/978-3-319-24574-4_28.
- A. Rozantsev, V. Lepetit, and P. Fua. On rendering synthetic images for training an object detector. *Computer Vision and Image Understanding*, 137:24–37, 2015. URL 10.1016/j.cviu.2014.12.006.
- X. Su and Q. Zhang. Dynamic 3-d shape measurement method: a review. *Optics and Lasers in Engineering*, 48(2):191–204, 2010.
- P. Verbeek and L. Van Vliet. Curvature and bending energy in digitized 2D and 3D images. In *8th Scandinavian Conference on Image Analysis, Tromsø, Norway*, 1993.
- G. Verhille and A. Bartoli. 3D conformation of a flexible fiber in a turbulent flow. *Experiments in Fluids*, 57(7):117, 2016. doi: 10.1007/s00348-016-2201-1.