

Joint reconstruction and segmentation of noisy velocity images as an inverse Navier–Stokes problem

Alexandros Kontogiannis¹, Scott V. Elgersma², Andrew J. Sederman² and Matthew P. Juniper^{1,†}

¹Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK

²Department of Chemical Engineering & Biotechnology, University of Cambridge, Philippa Fawcett Drive, Cambridge CB3 0AS, UK

(Received 6 July 2021; revised 14 April 2022; accepted 31 May 2022)

We formulate and solve a generalized inverse Navier–Stokes problem for the joint velocity field reconstruction and boundary segmentation of noisy flow velocity images. To regularize the problem, we use a Bayesian framework with Gaussian random fields. This allows us to estimate the uncertainties of the unknowns by approximating their posterior covariance with a quasi-Newton method. We first test the method for synthetic noisy images of two-dimensional (2-D) flows and observe that the method successfully reconstructs and segments the noisy synthetic images with a signal-to-noise ratio (SNR) of three. Then we conduct a magnetic resonance velocimetry (MRV) experiment to acquire images of an axisymmetric flow for low ($\simeq 6$) and high (> 30) SNRs. We show that the method is capable of reconstructing and segmenting the low SNR images, producing noiseless velocity fields and a smooth segmentation, with negligible errors compared with the high SNR images. This amounts to a reduction of the total scanning time by a factor of 27. At the same time, the method provides additional knowledge about the physics of the flow (e.g. pressure) and addresses the shortcomings of MRV (i.e. low spatial resolution and partial volume effects) that otherwise hinder the accurate estimation of wall shear stresses. Although the implementation of the method is restricted to 2-D steady planar and axisymmetric flows, the formulation applies immediately to three-dimensional (3-D) steady flows and naturally extends to 3-D periodic and unsteady flows.

Key words: computational methods, variational methods

† Email address for correspondence: mpj1001@cam.ac.uk

1. Introduction

Experimental measurements of fluid flows inside or around an object often produce velocity images that contain noise. These images may be post-processed to either reveal obscured flow patterns or to extract a quantity of interest (e.g. pressure or wall shear stress). For example, magnetic resonance velocimetry (MRV) (Fukushima 1999; Mantle & Sederman 2003; Elkins & Alley 2007; Markl *et al.* 2012; Demirkiran *et al.* 2021) can measure all three components of a time varying velocity field but the measurements become increasingly noisy as the spatial resolution is increased. To achieve an image of acceptable signal-to-noise ratio (SNR), repeated scans are often averaged, leading to long signal acquisition times. To address that problem, fast acquisition protocols (pulse sequences) can be used, but these may be difficult to implement and can lead to artefacts depending on the magnetic relaxation properties and the magnetic field homogeneity of the system studied. Another way to accelerate signal acquisition is by using sparse sampling techniques in conjunction with a reconstruction algorithm. The latter approach is an active field of research, commonly referred to as compressed sensing (Donoho 2006; Lustig, Donoho & Pauly 2007; Benning *et al.* 2014; Corona *et al.* 2021; Peper *et al.* 2020). Compressed sensing (CS) algorithms exploit *a priori* knowledge about the structure of the data, which is encoded in a regularization norm (e.g. total variation, wavelet bases), but without considering the physics of the problem. Even though the present study concerns the reconstruction of fully-sampled, noisy MRV images, the method that we present here can be applied to sparsely sampled MRV data.

For images depicting fluid flow, *a priori* knowledge can come in the form of a Navier–Stokes (N–S) problem. The problem of reconstructing and segmenting a flow image then can be expressed as a generalized inverse Navier–Stokes problem whose flow domain, boundary conditions and model parameters have to be inferred for the modelled velocity to approximate the measured velocity in an appropriate metric space. This approach not only produces a reconstruction that is an accurate fluid flow inside or around the object (a solution to a Navier–Stokes problem), but also provides additional physical knowledge (e.g. pressure), which is otherwise difficult to measure. Inverse Navier–Stokes problems have been intensively studied during the last decade, mainly enabled by the increase of available computing power. Recent applications in fluid mechanics range from the forcing inference problem (Hoang, Law & Stuart 2014) to the reconstruction of scalar image velocimetry (SIV) (Gillissen *et al.* 2018; Sharma *et al.* 2019) and particle image velocimetry (PIV) (Gillissen, Bouffanais & Yue 2019) signals, as well as the identification of optimal sensor arrangements (Mons, Chassaing & Sagaut 2017; Verma *et al.* 2019). Regularization methods that can be used for model parameters are reviewed by Stuart (2010) from a Bayesian perspective and by Benning & Burger (2018) from a variational perspective. The well-posedness of Bayesian inverse Navier–Stokes problems is addressed by Cotter *et al.* (2009).

Recently, Koltukluoğlu & Blanco (2018) treat the reduced inverse Navier–Stokes problem of finding only the Dirichlet boundary condition for the inlet velocity that matches the modelled velocity field to MRV data for a steady three-dimensional (3-D) flow in a glass replica of the human aorta. They measure the model-data discrepancy using the L^2 -norm and introduce additional variational regularization terms for the Dirichlet boundary condition. The same formulation is extended to periodic flows by Koltukluoğlu (2019), Koltukluoğlu, Cvijetić & Hiptmair (2019), using the harmonic balance method for the temporal discretization of the Navier–Stokes problem. Funke *et al.* (2019) address the problem of inferring both the inlet velocity (Dirichlet) boundary condition and the initial condition, for unsteady blood flows and four-dimensional (4-D) MRV data,

with applications to cerebral aneurysms. We note that the above studies consider rigid boundaries and require *a priori* an accurate, and time-averaged, geometric representation of the blood vessel.

To find the shape of the flow domain, e.g. the blood vessel boundaries, computed tomography (CT) or magnetic resonance angiography (MRA) is often used. The acquired image is then reconstructed, segmented and smoothed. This process not only requires substantial effort and the design of an additional experiment (e.g. CT, MRA), but it also introduces geometric uncertainties (Morris *et al.* 2016; Sankaran *et al.* 2016), which, in turn, affect the predictive confidence of arterial wall shear stress distributions and their mappings (Katritsis *et al.* 2007; Sotelo *et al.* 2016). For example, Funke *et al.* (2019) report discrepancies between the modelled and the measured velocity fields near the flow boundaries, and they suspect they are caused by geometric errors that were introduced during the segmentation process. In general, the assumption of rigid boundaries either implies that a time-averaged geometry has to be used or that an additional experiment (e.g. CT, MRA) has to be conducted to register the moving boundaries to the flow measurements.

A more consistent approach to this problem is to treat the blood vessel geometry as an unknown when solving the generalized inverse Navier–Stokes problem. In this way, the inverse Navier–Stokes problem simultaneously reconstructs and segments the velocity fields and can better adapt to the MRV experiment by correcting the geometric errors and improving the reconstruction.

In this study, we address the problem of simultaneous velocity field reconstruction and boundary segmentation by formulating a generalized inverse Navier–Stokes problem, whose flow domain, boundary conditions and model parameters are all considered as unknown. To regularize the problem, we use a Bayesian framework and Gaussian measures in Hilbert spaces. This further allows us to estimate the posterior Gaussian distributions of the unknowns using a quasi-Newton method, which has not yet been addressed for this type of problem. We provide an algorithm for the solution of this generalized inverse Navier–Stokes problem, and demonstrate it on synthetic images of two-dimensional (2-D) steady flows and real MRV images of a steady axisymmetric flow.

This paper consists of two parts. In § 2, we formulate the generalized inverse Navier–Stokes problem and an algorithm that solves it. In § 3, we test the method using both synthetic and real MRV velocity images and describe the set-up of the MRV experiment.

2. An inverse Navier–Stokes problem for noisy flow images

In this section, we formulate the generalized inverse Navier–Stokes problem and provide an algorithm for its solution. In what follows, $L^2(\Omega)$ denotes the space of square-integrable functions in Ω , with inner product $\langle \cdot, \cdot \rangle$ and norm $\|\cdot\|_{L^2(\Omega)}$, and $H^k(\Omega)$ the space of square-integrable functions with k square-integrable derivatives in Ω . For a given covariance operator, \mathcal{C} , we also define the covariance-weighted L^2 spaces, endowed with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{C}} := \langle \cdot, \mathcal{C}^{-1} \cdot \rangle$, which generates the norm $\|\cdot\|_{\mathcal{C}}$. The Euclidean norm in the space of real numbers \mathbb{R}^n is denoted by $|\cdot|_{\mathbb{R}^n}$. We use the superscript $(\cdot)^*$ to denote a measurement, $(\cdot)^\circ$ to denote a reconstruction and $(\cdot)^\bullet$ to denote the ground truth.

2.1. The inverse Navier–Stokes problem

An n -dimensional velocimetry experiment usually provides noisy flow velocity images on a domain $I \subset \mathbb{R}^n$, depicting the measured flow velocity \mathbf{u}^* inside an object $\Omega \subset I$




Figure 1. Given the images of a measured velocity field \mathbf{u}^* , we solve an inverse Navier–Stokes problem to infer the boundary Γ (or $\partial\Omega$), the kinematic viscosity and the inlet velocity profile on Γ_i . The solution to this inverse problem is a reconstructed velocity field \mathbf{u}° , from which the noise and the artefacts ($\mathbf{u}^* - S\mathbf{u}^\circ$) have been filtered out.

with boundary $\partial\Omega = \Gamma \cup \Gamma_i \cup \Gamma_o$ (figure 1). An appropriate model is the Navier–Stokes problem

$$\left. \begin{aligned} \mathbf{u} \cdot \nabla \mathbf{u} - \nu \Delta \mathbf{u} + \nabla p &= \mathbf{0} && \text{in } \Omega, \\ \nabla \cdot \mathbf{u} &= 0 && \text{in } \Omega, \\ \mathbf{u} &= \mathbf{0} && \text{on } \Gamma, \\ \mathbf{u} &= \mathbf{g}_i && \text{on } \Gamma_i, \\ -\nu \partial_\nu \mathbf{u} + p\mathbf{v} &= \mathbf{g}_o && \text{on } \Gamma_o \end{aligned} \right\} \quad (2.1)$$

where \mathbf{u} is the velocity, $p \leftarrow p/\rho$ is the reduced pressure, ρ is the density, ν is the kinematic viscosity, \mathbf{g}_i is the Dirichlet boundary condition at the inlet Γ_i , \mathbf{g}_o is the natural boundary condition at the outlet Γ_o , \mathbf{v} is the unit normal vector on $\partial\Omega$ and $\partial_\nu \equiv \mathbf{v} \cdot \nabla$ is the normal derivative.

We denote the data space by \mathbf{D} and the model space by \mathbf{M} , and assume that both spaces are subspaces of L^2 . In the 2-D case, $\mathbf{u}^* = (u_x^*, u_y^*)$ and we introduce the covariance operator

$$\mathcal{C}_u = \text{diag}(\sigma_{u_x}^2 \mathbf{I}, \sigma_{u_y}^2 \mathbf{I}), \quad (2.2)$$

where $\sigma_{u_x}^2, \sigma_{u_y}^2$ are the Gaussian noise variances of u_x^*, u_y^* , respectively, and \mathbf{I} is the identity operator. The discrepancy between the measured velocity field $\mathbf{u}^* \in \mathbf{D}$ and the modelled velocity field $\mathbf{u} \in \mathbf{M}$ is measured on the data space \mathbf{D} using the reconstruction error functional

$$\mathcal{E}(\mathbf{u}) \equiv \frac{1}{2} \|\mathbf{u}^* - S\mathbf{u}\|_{\mathcal{C}_u}^2 := \frac{1}{2} \int_I (\mathbf{u}^* - S\mathbf{u}) \mathcal{C}_u^{-1} (\mathbf{u}^* - S\mathbf{u}), \quad (2.3)$$

where $S : \mathbf{M} \rightarrow \mathbf{D}$ is the L^2 -projection from the model space \mathbf{M} to the data space \mathbf{D} . (Since the discretized space consists of bilinear quadrilateral finite elements (see § 2.7), this projection is a linear interpolation.)

Our goal is to infer the unknown parameters of the Navier–Stokes problem (2.1) such that the model velocity \mathbf{u} approximates the noisy measured velocity \mathbf{u}^* in the

covariance-weighted L^2 -metric defined by \mathcal{E} . In the general case, the unknown model parameters of (2.1) are the shape of Ω , the kinematic viscosity ν and the boundary conditions $\mathbf{g}_i, \mathbf{g}_o$. This inverse Navier–Stokes problem leads to the nonlinearly constrained optimization problem

$$\text{find } \mathbf{u}^\circ \equiv \operatorname{argmin}_{\Omega, \mathbf{x}} \mathcal{E}(\mathbf{u}(\Omega; \mathbf{x})) \quad \text{such that } \mathbf{u} \text{ satisfies (2.1),} \tag{2.4}$$

where \mathbf{u}° is the reconstructed velocity field and $\mathbf{x} = (\mathbf{g}_i, \mathbf{g}_o, \nu)$. Like most inverse problems, (2.4) is ill-posed and hard to solve. To alleviate the ill-posedness of the problem, we need to restrict our search of the unknowns (Ω, \mathbf{x}) to function spaces of sufficient regularity.

2.2. Regularization

If $x(t) \in L^2(\mathbb{R})$ is an unknown parameter, one way to regularize the inverse problem (2.4) is to search for minimizers of the augmented functional $\mathcal{J} \equiv \mathcal{E} + \mathcal{R}$, where

$$\mathcal{R}(x) = \sum_{j=0}^k \int_{\mathbb{R}} \alpha_j \left| \partial_x^j (x - \bar{x}) \right|^2 \tag{2.5}$$

is a regularization norm for a given (and fixed) prior assumption $\bar{x}(t) \in H^k(\mathbb{R})$, weights $\alpha_j \in \mathbb{R}$ and positive integer k . This simple idea can be quite effective because by minimizing \mathcal{R} , we force x to lie in a subspace of L^2 having higher regularity, namely H^k , and as close to the prior value \bar{x} as α_j allow. (The regularization term, given by (2.5), can be further extended to fractional Hilbert spaces by defining the norm $\|x\|_{H^s(\mathbb{R})} := \|(1 + |t|^s)\mathcal{F}x\|_{L^2(\mathbb{R})}$ for non-integer s , with $0 < s < \infty$, and where \mathcal{F} denotes the Fourier transform. Interestingly, under certain conditions, which are dictated by Sobolev’s embedding theorem (Evans 2010, Chapter 5), these Hilbert spaces can be embedded in the more familiar spaces of continuous functions.) However, as Stuart (2010) points out, in this setting, the choice of α_j , and even the form of \mathcal{R} , is arbitrary.

There is a more intuitive approach that recovers the form of the regularization norm \mathcal{R} from a probabilistic viewpoint. In the setting of the Hilbert space L^2 , the Gaussian measure $\gamma \sim \mathcal{N}(m, \mathcal{C})$ has the property that its finite-dimensional projections are multivariate Gaussian distributions, and it is uniquely defined by its mean $m \in L^2$, and its covariance operator $\mathcal{C} : L^2 \rightarrow L^2$ (Appendix A). It can be shown that there is a natural Hilbert space H_γ that corresponds to γ , and that (Bogachev 1998; Hairer 2009)

$$H_\gamma = \sqrt{\mathcal{C}}(L^2). \tag{2.6}$$

In other words, if x is a random function distributed according to γ , any realization of x lies in H_γ , which is the image of $\sqrt{\mathcal{C}}$. Furthermore, the corresponding inner product

$$\langle x, x' \rangle_{\mathcal{C}} = \langle \mathcal{C}^{-1/2}x, \mathcal{C}^{-1/2}x' \rangle \tag{2.7}$$

is the covariance between x and x' , and the norm $\|x\|_{\mathcal{C}}^2 = \langle x, x \rangle_{\mathcal{C}}$ is the variance of x . Therefore, if x is an unknown parameter for which *a priori* statistical information is available, and if the Gaussian assumption can be justified, we can choose

$$\mathcal{R}(x) = \frac{1}{2} \|x - \bar{x}\|_{\mathcal{C}}^2. \tag{2.8}$$

In this way, $\mathcal{J} \equiv \mathcal{E} + \mathcal{R}$ increases as the variance of x increases. Consequently, minimizing \mathcal{J} penalizes improbable realizations.

As mentioned in § 2.1, the unknown model parameters of the Navier–Stokes problem (2.1) are the kinematic viscosity ν , the boundary conditions $\mathbf{g}_i, \mathbf{g}_o$ and the shape of Ω . Since we consider the kinematic viscosity ν to be constant, the regularizing norm is simply

$$\frac{1}{2} |\nu - \bar{\nu}|_{\Sigma_\nu}^2 = \frac{1}{2\sigma_\nu^2} |\nu - \bar{\nu}|_{\mathbb{R}}^2, \tag{2.9}$$

where $\bar{\nu} \in \mathbb{R}$ is a prior guess for ν and $\sigma_\nu^2 \in \mathbb{R}$ is the variance. For the Dirichlet boundary condition, $\mathbf{g}_i \in L^2(\Gamma_i)$, we choose the exponential covariance function

$$C(x, x') = \frac{\sigma_{\mathbf{g}_i}^2}{2\ell} \exp\left(-\frac{|x - x'|}{\ell}\right) \tag{2.10}$$

with variance $\sigma_{\mathbf{g}_i}^2 \in \mathbb{R}$ and characteristic length $\ell \in \mathbb{R}$. For zero-Dirichlet (no-slip) or zero-Neumann boundary conditions on $\partial\Gamma_i$, (2.10) leads to the norm (Tarantola 2005, Chapter 7.21)

$$\|\mathbf{g}_i\|_{C_{\mathbf{g}_i}}^2 \simeq \frac{1}{\sigma_{\mathbf{g}_i}^2} \int_{\Gamma_i} \mathbf{g}_i^2 + \ell^2 (\nabla \mathbf{g}_i)^2. \tag{2.11}$$

Using integration by parts, we find that the covariance operator is

$$C_{\mathbf{g}_i} = \sigma_{\mathbf{g}_i}^2 (\mathbf{I} - \ell^2 \tilde{\Delta})^{-1}, \tag{2.12}$$

where $\tilde{\Delta}$ is the L^2 -extension of the Laplacian Δ that incorporates the boundary condition $\mathbf{g}_i = \mathbf{0}$ on $\partial\Gamma_i$. For the natural boundary condition, $\mathbf{g}_o \in L^2(\Gamma_o)$, we can use the same covariance operator, but equip $\tilde{\Delta}$ with zero-Neumann boundary conditions, i.e. $\partial_\nu \mathbf{g}_o = 0$ on $\partial\Gamma_o$. Lastly, for the shape of Ω , which we implicitly represent with a signed distance function ϕ_\pm (defined in § 2.4), we choose the norm

$$\frac{1}{2} \|\bar{\phi}_\pm - \phi_\pm\|_{C_{\phi_\pm}}^2 = \frac{1}{2\sigma_{\phi_\pm}^2} \|\bar{\phi}_\pm - \phi_\pm\|_{L^2(I)}^2, \tag{2.13}$$

where $\sigma_{\phi_\pm} \in \mathbb{R}$ and $\bar{\phi}_\pm \in L^2(I)$. Additional regularization for the boundary of Ω (i.e. the zero level-set of ϕ_\pm) is needed and it is described in § 2.4. Based on the above results, the regularization norm for the unknown model parameters is

$$\begin{aligned} \mathcal{R}(\mathbf{x}, \phi_\pm) &= \frac{1}{2} |\nu - \bar{\nu}|_{\Sigma_\nu}^2 + \frac{1}{2} \|\mathbf{g}_i - \bar{\mathbf{g}}_i\|_{C_{\mathbf{g}_i}}^2 \\ &\quad + \frac{1}{2} \|\mathbf{g}_o - \bar{\mathbf{g}}_o\|_{C_{\mathbf{g}_o}}^2 + \frac{1}{2} \|\bar{\phi}_\pm - \phi_\pm\|_{C_{\phi_\pm}}^2. \end{aligned} \tag{2.14}$$

2.3. Euler–Lagrange equations for the inverse Navier–Stokes problem

Testing the Navier–Stokes problem (2.1) with functions $(\mathbf{v}, q) \in \mathbf{H}^1(\Omega) \times L^2(\Omega)$ and after integrating by parts, we obtain the weak form

$$\begin{aligned} \mathcal{M}(\Omega)(\mathbf{u}, p, \mathbf{v}, q; \mathbf{x}) &\equiv \int_{\Omega} (\mathbf{v} \cdot (\mathbf{u} \cdot \nabla \mathbf{u}) + \nu \nabla \mathbf{v} : \nabla \mathbf{u} - (\nabla \cdot \mathbf{v})p - q(\nabla \cdot \mathbf{u})) + \int_{\Gamma_o} \mathbf{v} \cdot \mathbf{g}_o \\ &\quad + \int_{\Gamma \cup \Gamma_i} \mathbf{v} \cdot (-\nu \partial_\nu \mathbf{u} + p\mathbf{v}) + \mathcal{N}_{\Gamma_i}(\mathbf{v}, q, \mathbf{u}; \mathbf{g}_i) + \mathcal{N}_{\Gamma}(\mathbf{v}, q, \mathbf{u}; \mathbf{0}) = 0, \end{aligned} \tag{2.15}$$

where \mathcal{N} is the Nitsche (1971) penalty term

$$\mathcal{N}_T(\mathbf{v}, q, \mathbf{u}; \mathbf{z}) \equiv \int_T (-\nu \partial_\nu \mathbf{v} + q\mathbf{v} + \eta\mathbf{v}) \cdot (\mathbf{u} - \mathbf{z}), \tag{2.16}$$

which weakly imposes the Dirichlet boundary condition $\mathbf{z} \in L^2(T)$ on a boundary T , given a penalization constant η . (The penalization η is a numerical parameter with no physical significance (see § 2.7).) We define the augmented reconstruction error functional

$$\mathcal{J}(\Omega)(\mathbf{u}, p, \mathbf{v}, q; \mathbf{x}) \equiv \mathcal{E}(\mathbf{u}) + \mathcal{R}(\mathbf{x}, \phi_\pm) + \mathcal{M}(\Omega)(\mathbf{u}, p, \mathbf{v}, q; \mathbf{x}), \tag{2.17}$$

which contains the regularization terms \mathcal{R} and the model constraint \mathcal{M} , such that \mathbf{u} weakly satisfies (2.1). To reconstruct the measured velocity field \mathbf{u}^* and find the unknowns (Ω, \mathbf{x}) , we minimize \mathcal{J} by solving its associated Euler–Lagrange system.

2.3.1. Adjoint Navier–Stokes problem

To derive the Euler–Lagrange equations for \mathcal{J} , we first define

$$\mathcal{U}' = \left\{ \mathbf{u}' \in H^1(\Omega) : \mathbf{u}'|_{\Gamma \cup \Gamma_i} \equiv \mathbf{0} \right\} \tag{2.18}$$

to be the space of admissible velocity perturbations \mathbf{u}' and $\mathcal{P}' \subset L^2(\Omega)$ to be the space of admissible pressure perturbations p' , such that $(-\partial_\nu \mathbf{u}' + p'\mathbf{v})|_{\Gamma_o} \equiv \mathbf{0}$. We start with

$$\begin{aligned} \delta_{\mathbf{u}} \mathcal{E} &\equiv \frac{d}{d\tau} \mathcal{E}(\mathbf{u} + \tau \mathbf{u}') \Big|_{\tau=0} = \int_{\Omega} -\mathcal{C}_u^{-1}(\mathbf{u}^* - S\mathbf{u}) \cdot S\mathbf{u}' \\ &= \int_{\Omega} -S^\dagger \mathcal{C}_u^{-1}(\mathbf{u}^* - S\mathbf{u}) \cdot \mathbf{u}' \equiv \left\langle D_{\mathbf{u}} \mathcal{E}, \mathbf{u}' \right\rangle_{\Omega}. \end{aligned} \tag{2.19}$$

Adding together the first variations of \mathcal{M} with respect to (\mathbf{u}, p) ,

$$\delta_{\mathbf{u}} \mathcal{M} \equiv \frac{d}{d\tau} \mathcal{M}(\cdot)(\mathbf{u} + \tau \mathbf{u}', \dots) \Big|_{\tau=0}, \quad \delta_p \mathcal{M} \equiv \frac{d}{d\tau} \mathcal{M}(\cdot)(\dots, p + \tau p', \dots) \Big|_{\tau=0}, \tag{2.20a,b}$$

and after integrating by parts, we find

$$\begin{aligned} \delta_{\mathbf{u}} \mathcal{M} + \delta_p \mathcal{M} &= \int_{\Omega} \left(-\mathbf{u} \cdot (\nabla \mathbf{v} + (\nabla \mathbf{v})^\dagger) - \nu \Delta \mathbf{v} + \nabla q \right) \cdot \mathbf{u}' + \int_{\Omega} (\nabla \cdot \mathbf{v}) p' \\ &\quad + \int_{\partial \Omega} ((\mathbf{u} \cdot \mathbf{v})\mathbf{v} + (\mathbf{u} \cdot \mathbf{v})\mathbf{v} + \nu \partial_\nu \mathbf{v} - q\mathbf{v}) \cdot \mathbf{u}' \\ &\quad + \int_{\Gamma \cup \Gamma_i} \mathbf{v} \cdot (-\nu \partial_\nu \mathbf{u}' + p'\mathbf{v}) + \mathcal{N}_{\Gamma_i \cup \Gamma}(\mathbf{v}, q, \mathbf{u}'; \mathbf{0}). \end{aligned} \tag{2.21}$$

Since \mathcal{R} does not depend on (\mathbf{u}, p) , we can use (2.19) and (2.21) to assemble the optimality conditions of \mathcal{J} for (\mathbf{u}, p)

$$\left\langle D_{\mathbf{u}} \mathcal{J}, \mathbf{u}' \right\rangle_{\Omega} = 0, \quad \left\langle D_p \mathcal{J}, p' \right\rangle_{\Omega} = 0. \tag{2.22}$$

For (2.22) to hold true for all perturbations $(\mathbf{u}', p') \in \mathcal{U}' \times \mathcal{P}'$, we deduce that (\mathbf{v}, q) must satisfy the following adjoint Navier–Stokes problem

$$\left. \begin{aligned} -\mathbf{u} \cdot (\nabla \mathbf{v} + (\nabla \mathbf{v})^\dagger) - \nu \Delta \mathbf{v} + \nabla q &= -D_{\mathbf{u}} \mathcal{E} && \text{in } \Omega, \\ \nabla \cdot \mathbf{v} &= 0 && \text{in } \Omega, \\ \mathbf{v} &= \mathbf{0} && \text{on } \Gamma \cup \Gamma_i, \\ (\mathbf{u} \cdot \mathbf{v})\mathbf{v} + (\mathbf{u} \cdot \mathbf{v})\mathbf{v} + \nu \partial_\nu \mathbf{v} - q\mathbf{v} &= \mathbf{0} && \text{on } \Gamma_o. \end{aligned} \right\} \tag{2.23}$$

In this context, \mathbf{v} is the adjoint velocity and q is the adjoint pressure, which both vanish when $\mathbf{u}^* \equiv S\mathbf{u}$. Note also that we choose boundary conditions for the adjoint problem (2.23) that make the boundary terms of (2.21) vanish and that these boundary conditions are subject to the choice of \mathbf{U}' , which, in turn, depends on the boundary conditions of the (primal) Navier–Stokes problem.

2.3.2. Shape derivatives for the Navier–Stokes problem

To find the shape derivative of an integral defined in Ω , when the boundary $\partial\Omega$ deforms with speed \mathcal{V} , we use Reynold’s transport theorem. For the bulk integral of $f : \Omega \rightarrow \mathbb{R}$, we find

$$\frac{d}{d\tau} \left(\int_{\Omega(\tau)} f \right) \Big|_{\tau=0} = \int_{\Omega} f' + \int_{\partial\Omega} f (\mathcal{V} \cdot \mathbf{v}), \tag{2.24}$$

while for the boundary integral of f , we find (Walker 2015, Chapter 5.6)

$$\frac{d}{d\tau} \left(\int_{\partial\Omega(\tau)} f \right) \Big|_{\tau=0} = \int_{\partial\Omega} f' + (\partial_{\mathbf{v}} + \kappa)f(\mathcal{V} \cdot \mathbf{v}), \tag{2.25}$$

where f' is the shape derivative of f (due to \mathcal{V}), κ is the summed curvature of $\partial\Omega$ and $\mathcal{V} \equiv \zeta \mathbf{v}$, with $\zeta \in L^2(\partial\Omega)$, is the Hadamard parametrization of the speed field. Any boundary that is a subset of ∂I , i.e. the edge of the image I , is non-deforming and therefore the second term of the above integrals vanishes. The only boundary that deforms is $\Gamma \subset \partial\Omega$. For brevity, let $\delta_{\mathcal{V}}I$ denote the shape perturbation of an integral I . Using (2.24) on \mathcal{E} , we compute

$$\delta_{\mathcal{V}}\mathcal{E} = \left\langle D_{\mathbf{u}}\mathcal{E}, \mathbf{u}' \right\rangle_{\Omega}, \tag{2.26}$$

where $D_{\mathbf{u}}\mathcal{E}$ is given by (2.19). Using (2.24) and (2.25) on \mathcal{M} , we obtain the shape derivatives problem for (\mathbf{u}', p')

$$\left. \begin{aligned} \mathbf{u}' \cdot \nabla \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u}' - \nu \Delta \mathbf{u}' + \nabla p' &= \mathbf{0} && \text{in } \Omega, \\ \nabla \cdot \mathbf{u}' &= 0 && \text{in } \Omega, \\ \mathbf{u}' &= -\partial_{\mathbf{v}}\mathbf{u}(\mathcal{V} \cdot \mathbf{v}) && \text{on } \Gamma, \\ \mathbf{u}' &= \mathbf{0} && \text{on } \Gamma_i, \\ -\nu \partial_{\mathbf{v}}\mathbf{u}' + p'\mathbf{v} &= \mathbf{0} && \text{on } \Gamma_o, \end{aligned} \right\} \tag{2.27}$$

which can be used directly to compute the velocity and pressure perturbations for a given speed field \mathcal{V} . We observe that $(\mathbf{u}', p') \equiv \mathbf{0}$ when $\zeta \equiv \mathcal{V} \cdot \mathbf{v} \equiv 0$. Testing the shape derivatives problem (2.27) with (\mathbf{v}, q) , and adding the appropriate Nitsche terms for the weakly enforced Dirichlet boundary conditions, we obtain

$$\begin{aligned} \delta_{\mathcal{V}}\mathcal{M} &= \int_{\Omega} (\mathbf{v} \cdot (\mathbf{u}' \cdot \nabla \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u}') + \nu \nabla \mathbf{v} : \nabla \mathbf{u}' - (\nabla \cdot \mathbf{v})p' - q(\nabla \cdot \mathbf{u}')) \\ &+ \int_{\Gamma \cup \Gamma_i} \mathbf{v} \cdot (-\nu \partial_{\mathbf{v}}\mathbf{u}' + p'\mathbf{v}) + \mathcal{N}_{\Gamma_i}(\mathbf{v}, q, \mathbf{u}'; \mathbf{0}) + \mathcal{N}_{\Gamma}(\mathbf{v}, q, \mathbf{u}'; -\zeta \partial_{\mathbf{v}}\mathbf{u}) = 0. \end{aligned} \tag{2.28}$$

If we define I_i to be the first four integrals in (2.21), integrating (2.28) by parts yields

$$\delta_{\mathcal{V}}\mathcal{M} = \sum_{i=1}^4 I_i + \mathcal{N}_{\Gamma_i}(\mathbf{v}, q, \mathbf{u}'; \mathbf{0}) + \mathcal{N}_{\Gamma}(\mathbf{v}, q, \mathbf{u}'; -\zeta \partial_{\mathbf{v}}\mathbf{u}) = 0, \tag{2.29}$$

and, due to the adjoint problem (2.23), we find

$$\delta_{\mathcal{V}} \mathcal{M} = -\delta_{\mathcal{V}} \mathcal{E} + \int_{\Gamma} (-\nu \partial_{\mathbf{v}} \mathbf{v} + q\mathbf{v}) \cdot \zeta \partial_{\mathbf{v}} \mathbf{u} = 0, \quad (2.30)$$

since $\mathbf{u}'|_{\Gamma_i} \equiv \mathbf{0}$ and $\mathbf{u}|_{\Gamma} \equiv \mathbf{0}$. Therefore, the shape perturbation of \mathcal{J} is

$$\delta_{\mathcal{V}} \mathcal{J} \equiv \left\langle D_{\mathcal{V}} \mathcal{J}, \mathcal{V} \cdot \mathbf{v} \right\rangle_{\Gamma} \equiv \left\langle D_{\zeta} \mathcal{J}, \zeta \right\rangle_{\Gamma} = \delta_{\mathcal{V}} \mathcal{E} + \delta_{\mathcal{V}} \mathcal{M} + \delta_{\mathcal{V}} \mathcal{R} = 0, \quad (2.31)$$

which, due to (2.30) and $\delta_{\mathcal{V}} \mathcal{R} \equiv 0$, takes the form

$$\left\langle D_{\zeta} \mathcal{J}, \zeta \right\rangle_{\Gamma} = \left\langle \partial_{\mathbf{v}} \mathbf{u} \cdot (-\nu \partial_{\mathbf{v}} \mathbf{v} + q\mathbf{v}), \zeta \right\rangle_{\Gamma}, \quad (2.32)$$

where $D_{\zeta} \mathcal{J}$ is the shape gradient. Note that the shape gradient depends on the normal gradient of the (primal) velocity field and the pseudotraction, $(-\nu \nabla \mathbf{v} + q\mathbf{I}) \cdot \mathbf{v}$, that the adjoint flow exerts on Γ .

2.3.3. Generalized gradients for the unknown model parameters \mathbf{x}

The unknown model parameters \mathbf{x} have an explicit effect on \mathcal{M} and \mathcal{R} , and can therefore be obtained by taking their first variations. For the Dirichlet-type boundary condition at the inlet, we find

$$\begin{aligned} \left\langle D_{\mathbf{g}_i} \mathcal{J}, \mathbf{g}'_i \right\rangle_{\Gamma_i} &= \left\langle \nu \partial_{\mathbf{v}} \mathbf{v} - q\mathbf{v} - \eta\mathbf{v} + C_{\mathbf{g}_i}^{-1} (\mathbf{g}_i - \bar{\mathbf{g}}_i), \mathbf{g}'_i \right\rangle_{\Gamma_i} \\ &= \left\langle C_{\mathbf{g}_i} (\nu \partial_{\mathbf{v}} \mathbf{v} - q\mathbf{v} - \eta\mathbf{v}) + \mathbf{g}_i - \bar{\mathbf{g}}_i, \mathbf{g}'_i \right\rangle_{C_{\mathbf{g}_i}} = \left\langle \hat{D}_{\mathbf{g}_i} \mathcal{J}, \mathbf{g}'_i \right\rangle_{C_{\mathbf{g}_i}}, \end{aligned} \quad (2.33)$$

where $-\hat{D}_{\mathbf{g}_i} \mathcal{J}$ is the steepest descent direction that corresponds to the covariance-weighted norm. For the natural boundary condition at the outlet, we find

$$\begin{aligned} \left\langle D_{\mathbf{g}_o} \mathcal{J}, \mathbf{g}'_o \right\rangle_{\Gamma_o} &= \left\langle \mathbf{v} + C_{\mathbf{g}_o}^{-1} (\mathbf{g}_o - \bar{\mathbf{g}}_o), \mathbf{g}'_o \right\rangle_{\Gamma_o} \\ &= \left\langle C_{\mathbf{g}_o} \mathbf{v} + \mathbf{g}_o - \bar{\mathbf{g}}_o, \mathbf{g}'_o \right\rangle_{C_{\mathbf{g}_o}} = \left\langle \hat{D}_{\mathbf{g}_o} \mathcal{J}, \mathbf{g}'_o \right\rangle_{C_{\mathbf{g}_o}}. \end{aligned} \quad (2.34)$$

Lastly, since the kinematic viscosity is considered to be constant within Ω , its generalized gradient is

$$\begin{aligned} \left\langle D_{\nu} \mathcal{J}, \nu' \right\rangle_{\mathbb{R}} &= \left\langle \int_{\Omega} \nabla \mathbf{v} : \nabla \mathbf{u} + \Sigma_{\nu}^{-1} (\nu - \bar{\nu}), \nu' \right\rangle_{\mathbb{R}} \\ &= \left\langle \Sigma_{\nu} \int_{\Omega} \nabla \mathbf{v} : \nabla \mathbf{u} + \nu - \bar{\nu}, \nu' \right\rangle_{\Sigma_{\nu}} = \left\langle \hat{D}_{\nu} \mathcal{J}, \nu' \right\rangle_{\Sigma_{\nu}}. \end{aligned} \quad (2.35)$$

For a given step size $\mathbb{R} \ni \tau > 0$, the steepest descent directions (2.33)–(2.35) can be used either to update an unknown parameter x through

$$x_{k+1} = x_k + \tau s_k, \quad (2.36)$$

with $s_k = -\hat{D}_x \mathcal{J}$, or to reconstruct an approximation \tilde{H} of the inverse Hessian matrix, in the context of a quasi-Newton method, and thereby to compute $s_k = -\tilde{H} \hat{D}_x \mathcal{J}$. We adopt the latter approach, which is discussed in § 2.5.

2.4. Geometric flow

To deform the boundary $\partial\Omega$ using the simple update formula (2.36), we need a parametric surface representation. Here we choose to implicitly represent $\partial\Omega$ using signed distance functions ϕ_{\pm} . The object Ω and its boundary $\partial\Omega$ are then identified with a particular function ϕ_{\pm} such that

$$\Omega = \{x \in \Omega : \phi_{\pm}(x) < 0\}, \quad \partial\Omega = \{x \in \Omega : \phi_{\pm}(x) = 0\}. \quad (2.37a,b)$$

2.4.1. Implicit representation of Ω using signed distance functions

A signed distance function ϕ_{\pm} for Ω can be obtained by solving the Eikonal equation

$$|\nabla\phi_{\pm}(x)| = 1 \quad \text{subject to } \phi_{\pm}|_{\partial\Omega} = 0 \quad x \in I. \quad (2.38)$$

One way to solve this problem is with level-set methods (Osher & Sethian 1988; Sethian 1996; Burger 2001, 2003; Burger & Osher 2005; Yu, Juniper & Magri 2019). There is, however, a different approach, which relies on the heat equation (Varadhan 1967b,a; Crane, Weischedel & Wardetzky 2017). The main result that we draw from Varadhan (1967b), to justify the use of the heat equation for the approximation of ϕ_{\pm} , states that

$$d(x, \partial\Omega) = \lim_{\tau_1 \rightarrow 0} \left(-\frac{\sqrt{\tau_1}}{2} \log u(x, \tau_1) \right), \quad x \in I, \quad (2.39)$$

where $d(x, \partial\Omega)$ is the Euclidean distance between any point $x \in I$ and $\partial\Omega$, and u is the solution of heat propagation away from $\partial\Omega$

$$\left. \begin{aligned} (I - \tau_1 \Delta) u &= 0 && \text{in } I, \\ u &= 1 && \text{on } \partial\Omega. \end{aligned} \right\} \quad (2.40)$$

Crane *et al.* (2017) used the above result to implement a smoothed distance function computation method which they called the ‘heat method’. Here, we slightly adapt this method to compute signed distance functions ϕ_{\pm} in truncated domains (figure 2b). To compute ϕ_{\pm} , we therefore solve (2.40) for $\tau_1 \ll 1$ and then obtain ϕ_{\pm} by solving

$$\left. \begin{aligned} \nabla \cdot \nabla\phi_{\pm} &= \nabla \cdot X && \text{in } I, \\ \partial_{\mathbf{v}}\phi_{\pm} &= X \cdot \mathbf{v} && \text{on } \partial I, \\ \phi_{\pm} &= 0 && \text{on } \partial\Omega, \end{aligned} \right\} \quad X = -\text{sgn}(\psi) \frac{\nabla u}{|\nabla u|}, \quad (2.41)$$

with X being the normalized heat flux and ψ being a signed function such that $\psi(x)$ is negative for points x in Ω and positive for points x outside Ω . This intermediate step (the solution of two Poisson problems (2.40)–(2.41) instead of one) is taken to ensure that $|\nabla\phi_{\pm}| = 1$.

2.4.2. Propagating the boundary of Ω

To deform the boundary $\partial\Omega$, we transport ϕ_{\pm} under the speed field $\mathcal{V} \equiv \zeta \mathbf{v}$. The convection-diffusion problem for $\phi_{\pm}(x, t)$ reads

$$\left. \begin{aligned} \partial_t \phi_{\pm} + \mathcal{V} \cdot \nabla \phi_{\pm} - \epsilon_{\phi_{\pm}} \Delta \phi_{\pm} &= 0 && \text{in } I \times (0, \tau], \\ \phi_{\pm} &= \phi_{\pm 0} && \text{in } I \times \{t = 0\}, \end{aligned} \right\} \quad \epsilon_{\phi_{\pm}} = \frac{|\mathcal{V}|_{\infty} \iota}{Re_{\phi_{\pm}}}, \quad (2.42)$$

where $\phi_{\pm 0}$ denotes the signed distance function of the current domain Ω , $\epsilon_{\phi_{\pm}}$ is the diffusion coefficient, ι is a length scale, $Re_{\phi_{\pm}}$ is a Reynolds number and $\mathcal{V} : I \rightarrow \mathbb{R} \times \mathbb{R}$ is

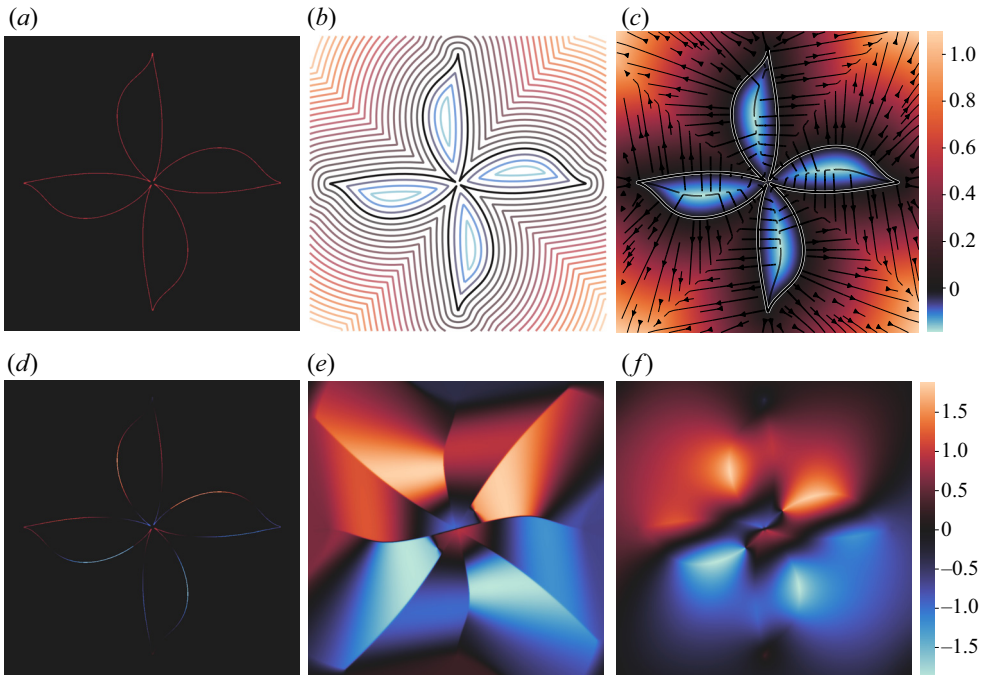


Figure 2. Geometric flow of $\partial\Omega$ (figure 2a) relies on the computation of its signed distance field ϕ_{\pm} (figure 2b) and its normal vector extension $\mathring{\mathbf{v}}$ (figure 2c). The shape gradient ζ (figure 2d), which is initially defined on $\partial\Omega$, is extended to the whole image I (ζ in figure 2e,f). Shape regularization is achieved by increasing the diffusion coefficient ϵ_{ζ} to mitigate small-scale perturbations when assimilating noisy velocity fields \mathbf{u}^* . Figure 2(f) shows results at a lower value of Re_{ζ} than figure 2(e). (a) Shape of $\partial\Omega$. (b) Level-sets of ϕ_{\pm} . (c) Magnitude of ϕ_{\pm} and $\mathring{\mathbf{v}}$. (d) Shape gradient ζ on $\partial\Omega$. (e) ζ in I ($Re_{\zeta} = 1$). (f) ζ in I ($Re_{\zeta} = 0.01$).

an extension of $\mathcal{V} : \partial\Omega \rightarrow \mathbb{R} \times \mathbb{R}$. If we solve (2.42) for $\phi_{\pm}(x, \tau)$, we obtain the implicit representation of the perturbed domain Ω_{τ} at time $t = \tau$ (the step size), but to do so, we first need to extend \mathcal{V} to the whole space of the image I .

To extend \mathcal{V} to I , we extend the normal vector \mathbf{v} and the scalar function ζ , which are both initially defined on $\partial\Omega$. The normal vector extension (figure 2c) is easily obtained by

$$\mathring{\mathbf{v}}(x) = \frac{\nabla\phi_{\pm}}{|\nabla\phi_{\pm}|} = \nabla\phi_{\pm}, \quad x \in I, \tag{2.43}$$

since $|\nabla\phi_{\pm}| = 1$, and an outward-facing extension is given by

$$\mathring{\mathbf{v}}_o = \text{sgn}(\phi_{\pm})\mathring{\mathbf{v}}. \tag{2.44}$$

We then use the extended normal vector $\mathring{\mathbf{v}}_o$ to extend $\zeta \in L^2(\partial\Omega)$ to $\mathring{\zeta} \in L^2(I)$, using the convection-diffusion problem

$$\left. \begin{aligned} \partial_t \mathring{\zeta} + \mathring{\mathbf{v}}_o \cdot \nabla \mathring{\zeta} - \epsilon_{\zeta} \Delta \mathring{\zeta} &= 0 & \text{in } I \times (0, \tau_{\zeta}], \\ \mathring{\zeta} &= \zeta & \text{on } \partial\Omega \times (0, \tau_{\zeta}], \\ \mathring{\zeta} &\equiv 0 & \text{in } I \times \{t = 0\}, \end{aligned} \right\} \epsilon_{\zeta} = \frac{|\mathring{\mathbf{v}}_o|_{\infty} t}{Re_{\zeta}}. \tag{2.45}$$

In other words, we convect ζ along the predefined $\mathring{\mathbf{v}}_o$ -streamlines and add isotropic diffusion for regularization (figure 2e,f). The choice of $\epsilon_{\phi_{\pm}}$ in (2.42) and ϵ_{ζ} in (2.45)

has been made for the shape regularization to depend only on the length scale ι and the Reynolds numbers $Re_{\phi_{\pm}}$, Re_{ζ} . More precisely, the shape regularization depends only on $Re_{\phi_{\pm}}$ and Re_{ζ} because we fix the length scale ι to equal the smallest possible length scale of the modelled flow, which is the numerical grid spacing h for a uniform Cartesian grid. For illustration, if we consider ζ to be the concentration of a dye on $\partial\Omega$ (figure 2d), using a simplified scaling argument similar to the growth of a boundary layer on a flat plate, we observe that the diffusing dye at distance d from $\partial\Omega$ will extend over a width δ such that

$$\delta \sim \sqrt{\frac{\epsilon_{\zeta} d}{|\mathring{v}_o|_{\infty}}} = \sqrt{\frac{\iota d}{Re_{\zeta}}} \quad \text{or} \quad \frac{\delta}{\iota} \sim \sqrt{\frac{\alpha}{Re_{\zeta}}} \quad \text{when } d = \alpha \iota. \quad (2.46)$$

The above scaling approximation describes the dissipation rate of small-scale features such as roughness away from $\partial\Omega$. This is therefore how $Re_{\phi_{\pm}}$ and Re_{ζ} control the regularity of the boundary $\partial\Omega_{\tau}$ at time $t = \tau$, which is given by (2.42). We take τ_{ζ} to be large enough to find a steady-state for (2.45). We recast the linear initial value problems (2.42) and (2.45) into their corresponding boundary value problems using backward-Euler temporal discretization because the time-dependent solution does not interest us here.

The extended shape gradient (2.32), after taking into account the regularizing term for ϕ_{\pm} , is therefore given by

$$\begin{aligned} \left\langle D_{\zeta}^{\circ} \mathcal{J}, \mathring{\zeta}' \right\rangle_I &= \left\langle \mathring{\zeta} + C_{\phi_{\pm}}^{-1} (\bar{\phi}_{\pm} - \phi_{\pm}), \mathring{\zeta}' \right\rangle_I \\ &= \left\langle C_{\phi_{\pm}} \mathring{\zeta} + \bar{\phi}_{\pm} - \phi_{\pm}, \mathring{\zeta}' \right\rangle_{C_{\phi_{\pm}}} = \left\langle \hat{D}_{\zeta}^{\circ} \mathcal{J}, \mathring{\zeta}' \right\rangle_{C_{\phi_{\pm}}}, \end{aligned} \quad (2.47)$$

where $\mathring{\zeta}$ is the extension of the shape gradient $\zeta(x) = \partial_v \mathbf{u} \cdot (-v \partial_v \mathbf{v} + q \mathbf{v})$, for x on Γ .

2.5. Segregated approach for the Euler–Lagrange system

The inverse Navier–Stokes problem for the reconstruction and the segmentation of noisy velocity images \mathbf{u}^* can be written as the saddle point problem (Benzi, Golubt & Liesen 2005)

$$\text{find } \mathbf{u}^{\circ} \equiv \arg \min_{\Omega, \mathbf{x}} \max_{\mathbf{v}, q} \mathcal{J}(\Omega)(\mathbf{u}, p, \mathbf{v}, q; \mathbf{x}), \quad (2.48)$$

where \mathcal{J} is given by (2.17). The above optimization problem leads to an Euler–Lagrange system whose optimality conditions were formulated in §2.3. We briefly describe our segregated approach to solve this Euler–Lagrange system in algorithm 1.

To precondition the steepest descent directions (2.33)–(2.35) and (2.47), we reconstruct the approximated inverse Hessian \tilde{H} of each unknown using the BFGS quasi-Newton method (Fletcher 2000) with damping (Nocedal & Wright 2006). Due to the large scale of the problem, it is only possible to work with the matrix-vector product representation of \tilde{H} . Consequently, the search directions are given by

$$s = - \begin{pmatrix} \tilde{H}_{\zeta}^{\circ} & \cdot & \cdot & \cdot \\ \cdot & \tilde{H}_{g_i} & \cdot & \cdot \\ \cdot & \cdot & \tilde{H}_{g_o} & \cdot \\ \cdot & \cdot & \cdot & \tilde{H}_v \end{pmatrix} \begin{pmatrix} \hat{D}_{\zeta}^{\circ} \mathcal{J} \\ \hat{D}_{g_i} \mathcal{J} \\ \hat{D}_{g_o} \mathcal{J} \\ \hat{D}_v \mathcal{J} \end{pmatrix}, \quad (2.49)$$

and the unknown variables \mathbf{x} are updated according to (2.36). The signed distance function ϕ_{\pm} is perturbed according to (2.42), with $\mathring{\mathcal{V}} \equiv -(\tilde{H}_{\zeta}^{\circ} \hat{D}_{\zeta}^{\circ} \mathcal{J}) \mathring{v}$.

Algorithm 1: Reconstruction and segmentation of noisy flow velocity images.

Input: \mathbf{u}^* , initial guesses for the unknowns (Ω_0, x_0) , regularization parameters.

begin

$k \leftarrow 0$

$(\phi_{\pm})_k \leftarrow$ signed distance field ((2.40) and (2.41))

$(\mathbf{u}, p)_k \leftarrow$ Navier–Stokes problem for $(\phi_{\pm}, \mathbf{x})_k$ (2.1)

while convergence criterion is not met **do**

$(\mathbf{v}, q) \leftarrow$ adjoint Navier–Stokes problem with \mathbf{u}_k (2.23)

$\hat{D}_{(\cdot)} \mathcal{J} \leftarrow$ steepest ascent directions ((2.33)–(2.35) and (2.47))

$s, \tau \leftarrow$ search directions and step-size (2.49)

$(\phi_{\pm}, \mathbf{x})_{k+1} \leftarrow$ perturb ϕ_{\pm} (2.42) and model parameters \mathbf{x} (2.36)

$(\mathbf{u}, p)_{k+1} \leftarrow$ linearized Navier–Stokes problem for $(\phi_{\pm}, \mathbf{x})_{k+1}$ (2.50)

$k \leftarrow k + 1$

$(\mathbf{u}^{\circ}, p^{\circ}) \leftarrow (\mathbf{u}, p)_k$

$(\Omega^{\circ}, \mathbf{x}^{\circ}) \leftarrow (\phi_{\pm}, \mathbf{x})_k$

Output: reconstruction $(\mathbf{u}^{\circ}, p^{\circ})$ and inferred model parameters $(\Omega^{\circ}, \mathbf{x}^{\circ})$.

Optional output: wall shear rate γ_w° from \mathbf{u}° and $\partial\Omega^{\circ}$.

We start every line search with a global step size $\tau = 1$ and halve the step size until $\mathcal{J}((\phi_{\pm}, \mathbf{x})_{k+1}) < \mathcal{J}((\phi_{\pm}, \mathbf{x})_k)$. To update the flowfield \mathbf{u}_k to \mathbf{u}_{k+1} , we solve the Oseen problem for the updated parameters $(\phi_{\pm}, \mathbf{x})_{k+1}$

$$\mathbf{u}_k \cdot \nabla \mathbf{u}_{k+1} - \nu \Delta \mathbf{u}_{k+1} + \nabla p_{k+1} = \mathbf{0}, \quad \nabla \cdot \mathbf{u}_{k+1} = 0, \quad (2.50)$$

with the boundary conditions given by (2.1). Algorithm 1 terminates if either the covariance-weighted norm for the perturbations of the model parameters is below the user-specified tolerance or the line search fails to reduce \mathcal{J} .

2.6. Uncertainty estimation

We now briefly describe how the reconstructed inverse Hessian \tilde{H} can provide estimates for the uncertainties of the model parameters. To simplify the description, let x denote an unknown parameter distributed according to $\mathcal{N}(x_k, C_x)$. The linear approximation to the data \mathbf{u}^* is given by

$$\mathbf{u}^* = \mathcal{Z}x + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, C_u), \quad (2.51a,b)$$

where $\mathbf{u} = \mathcal{Z}x$, where \mathcal{Z} is the operator that encodes the linearized Navier–Stokes problem around the solution \mathbf{u}_k . To solve (2.48), we update x as

$$x_{k+1} = x_k + C \mathcal{Z}^{\dagger} C_u^{-1} (\mathbf{u}^* - \mathcal{Z}x_k), \quad \text{with } C = (\mathcal{Z}^{\dagger} C_u^{-1} \mathcal{Z} + C_x^{-1})^{-1}, \quad (2.52)$$

where \mathcal{Z}^{\dagger} is the operator that encodes the adjoint Navier–Stokes problem and C is the posterior covariance operator. It can be shown that (Tarantola 2005, Chapter 6.22.8)

$$C = (\mathcal{Z}^{\dagger} C_u^{-1} \mathcal{Z} + C_x^{-1})^{-1} = (C_x \mathcal{Z}^{\dagger} C_u^{-1} \mathcal{Z} + \mathbf{I})^{-1} C_x \simeq \tilde{H}_x C_x, \quad (2.53)$$

where \tilde{H}_x is the reconstructed inverse Hessian for x . Note that \tilde{H} by itself approximates $(C_x \mathcal{Z}^{\dagger} C_u^{-1} \mathcal{Z} + \mathbf{I})^{-1}$ and not C , because we use the steepest ascent directions $\hat{D}_{(\cdot)} \mathcal{J}$ (prior-preconditioned gradients), instead of the gradients $D_{(\cdot)} \mathcal{J}$, in the BFGS formula.

Therefore, if $\tilde{C} \equiv \tilde{H}_x C_x$ is the approximated covariance matrix, then samples x_{k+1}^s from the posterior distribution can be drawn using the Karhunen–Loève expansion

$$x_{k+1}^s = x_k + \sum_k \eta_k \sqrt{\lambda_k} \varphi_k, \quad \text{with } \eta_k \sim \mathcal{N}(0, 1), \quad (2.54)$$

where $(\lambda, \varphi)_k$ is the eigenvalue/eigenvector pair of \tilde{C} . The variance of x_{k+1} can then be directly computed from the samples.

2.7. Numerics

To solve the above boundary value problems numerically, we use an immersed boundary finite element method. In particular, we implement the fictitious domain cut-cell finite element method (FEM), which was introduced by Burman (2010) and Burman & Hansbo (2012) for the Poisson problem, and extended by Burman, and Massing *et al.* (2014); Massing, Schott & Wall (2018) to the Stokes and the Oseen problems. We define \mathcal{T}_h to be a tessellation of I produced by square cells (pixels) $K \in \mathcal{T}_h$, having sides of length h . We also define the set of cut-cells $\mathcal{T}_h^\triangleright$ consisting of the cells that are cut by the boundary $\partial\Omega$, and \mathcal{T}_h^\square the set of cells that are found inside Ω and which remain intact (not cut) (see figure 1). We assume that the boundary $\partial\Omega$ is well-resolved, i.e. $\ell_{\partial\Omega}/h \gg 1$ where $\ell_{\partial\Omega}$ is the smallest length scale of $\partial\Omega$. For the detailed assumptions on $\partial\Omega$, we cite Burman & Hansbo (2012). The discretized space is generated by assigning a bilinear quadrilateral finite element \mathcal{Q}_1 to every cell K . To compute the integrals, we use standard Gaussian quadrature for cells $K \in \mathcal{T}_h^\square$, while for cut-cells $K \in \mathcal{T}_h^\triangleright$, where integration must be considered only for the intersection $K \cap \Omega$, we use the approach of Mirtich (1996), which relies on the divergence theorem and simply replaces the integral over $K \cap \Omega$ with an integral over $\partial(K \cap \Omega)$. The boundary integral on $\partial(K \cap \Omega)$ is then easily computed using one-dimensional Gaussian quadrature (Massing, Larson & Logg 2013). Since we use an inf-sup unstable finite element pair $(\mathcal{Q}_1-\mathcal{Q}_1)$ (Brenner & Scott 2008), we use a pressure-stabilizing Petrov–Galerkin formulation (Tezduyar 1991; Codina 2002) and ∇ -div stabilization for preconditioning (Benzi & Olshanskii 2006; Heister & Rapin 2013). Typical values and formulae for numerical parameters, e.g. Nitsche’s penalization η , are given by Massing *et al.* (2014, 2018). Here, we take $\eta = \gamma\nu/h$ (Massing *et al.* 2018), with $\gamma = 100$. To solve the Navier–Stokes problem, we use fixed-point iteration (Oseen linearization) and at each iteration, we solve the coupled system using the Schur complement; with an iterative solver (LGMRES) for the outer loops and a direct sparse solver (UMFPACK) for the inner loops. The immersed FEM solver and all the necessary numerical operations of algorithm 1 are implemented in Python, using its standard libraries for scientific computing, namely SciPy (Virtanen *et al.* 2020) and NumPy (Harris *et al.* 2020). Computationally intensive functions are accelerated using Numba (Lam, Pitrou & Seibert 2015) and CuPy (Okuta *et al.* 2017).

3. Reconstruction and segmentation of flow images

In this section, we reconstruct and segment noisy flow images by solving the inverse Navier–Stokes problem (2.48) using algorithm 1. We then use the reconstructed velocity field to estimate the wall shear rate on the reconstructed boundary. First, we apply this to three test cases with known solutions by generating synthetic 2-D Navier–Stokes data. Next, we perform a magnetic resonance velocimetry experiment to acquire images of a 3-D axisymmetric Navier–Stokes flow, and apply algorithm 1 to these images.

		Image dimension	Model dimension	σ_{u_x}/U	σ_{u_y}/U
Converging channel	(2-D)	192 ²	200 ²	3.97×10^{-1}	8.97×10^{-3}
<i>Regularization</i>		$\sigma_{\phi_{\pm}}/D$	σ_{g_i}/U	σ_v/UD	$Re_{\phi_{\pm}}$
Converging channel	(2-D)	1.0	2.0	.	Re_{ξ}
					ℓ/h
				0.025	0.025
					3

Table 1. Input parameters for the inverse 2-D Navier–Stokes problem.

We define the SNR of the u_x^* image as

$$SNR_x = \frac{\mu_x}{\sigma_{u_x}}, \quad \mu_x \equiv \frac{1}{|\Omega^\bullet|} \int_{\Omega^\bullet} |u_x^\bullet|, \quad (3.1a,b)$$

where σ_{u_x} is the standard deviation, Ω^\bullet is the ground truth domain, $|\Omega^\bullet|$ is the volume of this domain and $|u_x^\bullet|$ is the magnitude of the ground truth x -velocity component in Ω^\bullet . We also define the component-wise averaged, noise relative reconstruction error \mathcal{E}_x^\bullet and the total relative reconstruction error \mathcal{E}^\bullet by

$$\mathcal{E}_x^\bullet \equiv \log \left(\frac{1}{|\Omega|} \int_{\Omega} \frac{|u_x^\bullet - \mathcal{S}u_x^\circ|}{\sigma_{u_x}} \right) \quad \text{and} \quad \mathcal{E}^\bullet \equiv \frac{\|u^\bullet - \mathcal{S}u^\circ\|_{L^1(I)}}{\|u^\bullet\|_{L^1(I)}}, \quad (3.2a,b)$$

respectively. Similar measures also apply for the u_y^* image.

We define the volumetric flow rate Q , the cross-section area at the inlet A and the diameter at the inlet D . The Reynolds number is based on the reference velocity $U \equiv Q/A$ and the reference length D .

3.1. Synthetic data for 2-D flow in a converging channel

We start by testing [algorithm 1](#) on a flow through a symmetric converging channel having a taper ratio of 0.67. To generate synthetic 2-D Navier–Stokes data, we solve the Navier–Stokes problem (2.1) for a parabolic inlet velocity profile (g_i), zero-pseudotraction boundary conditions at the outlet ($g_o \equiv \mathbf{0}$) and $Re \simeq 534$, to obtain the ground truth velocity u^\bullet . We then generate the synthetic data u^* by corrupting the components of u^\bullet with white Gaussian noise such that $SNR_x = SNR_y = 3$. For this test case, we are only trying to infer Ω and g_i . Note that, in our method, the initial guess x_0 of an unknown x equals the mean of its prior distribution \bar{x} , i.e. $x_0 \equiv \bar{x}$. We start the algorithm using bad initial guesses (high uncertainty in priors) for both the unknown parameters (see [table 1](#)). The initial guess for Ω , labelled Ω_0 , is a rectangular domain with height equal to $0.7D$, centred in the image domain. For g_{i0} , we take a parabolic velocity profile with a peak velocity of approximately $2U$ that fits the inlet of Ω_0 . For comparison, g_i^\bullet has a peak velocity of $1.5U$, while it is also defined on a different domain, namely Ω^\bullet .

The algorithm manages to reconstruct and segment the noisy flow images in 39 iterations, with a total reconstruction error $\mathcal{E}^\bullet \simeq 1.44\%$. The results are presented in [figures 3](#) and [4](#). We observe that the inverse Navier–Stokes problem performs very well in filtering the noise ($u^* - \mathcal{S}u^\circ$) ([figure 3c,f](#)), providing noiseless images for each component of the velocity ([figure 3b,e](#)). As we expect, the discrepancies $C_u^{-1}(u^* - \mathcal{S}u^\circ)$ ([figure 3c,f](#)) consist mainly of Gaussian white noise, except at the corners of the outlet ([figure 3f](#)), where there is a weak correlation. For a more detailed presentation of the denoising effect, we plot slices of the reconstructed velocity ([figure 4c](#)) and the ground




Figure 3. Reconstruction (algorithm 1) of synthetic noisy velocity images depicting the flow (from left to right) in a converging channel. Figures 3(a,b) and 3(d,e) show the horizontal, u_x , and vertical, u_y , velocities and share the same colourmap (colourbar not shown). Figure 3(c,f) shows the discrepancy between the noisy velocity images and the reconstruction (colourbars apply only to figure 3c,f). (a) Synthetic image u_x^* . (b) Our reconstruction u_x° . (c) Discrepancy $\sigma_{u_x}^{-1}(u_x^* - Su_x^\circ)$. (d) Synthetic image u_y^* . (e) Our reconstruction u_y° . (f) Discrepancy $\sigma_{u_y}^{-1}(u_y^* - Su_y^\circ)$.

truth velocity (figure 4d). The reconstructed pressure p° , which is consistent with the reconstructed velocity \mathbf{u}° to machine precision accuracy, is, in effect, indistinguishable from the ground truth p^* (figure 5).

Having obtained the reconstructed velocity \mathbf{u}° , we can compute the wall shear rate γ_w° on the reconstructed boundary $\partial\Omega^\circ$, which we compare with the ground truth γ_w^* in figure 6. Using the upper ($\partial\Omega_+^\circ$) and lower ($\partial\Omega_-^\circ$) limits of the 2σ confidence region for $\partial\Omega^\circ$ (figure 4a), we estimate a confidence region for γ_w° ; although this has to be interpreted carefully. Note that, for example, $\partial\Omega_+^\circ$ and $\partial\Omega_-^\circ$ can be smoother than the mean $\partial\Omega^\circ$, and, therefore, $\partial\Omega^\circ$ may be found outside this confidence region. A better estimate of the confidence region could be obtained by sampling the posterior distribution of $\partial\Omega^\circ$ solving a Navier–Stokes problem for each sample $\partial\Omega_k$ and finding the distribution of γ_w° . Since the latter approach would be computationally intensive, we only provide our estimate, which requires the solution of only two Navier–Stokes problems.

3.2. Synthetic data for 2-D flow in a simulated abdominal aortic aneurysm

Next, we test algorithm 1 in a channel that resembles the cross-section of a small abdominal aortic aneurysm, with $D_{max}/D \simeq 1.5$, where D_{max} is the maximum diameter at the midsection. We generate synthetic images for \mathbf{u}^* as in §3.1, again for $\text{SNR}_x = \text{SNR}_y = 3$, but now for $Re = 153$. The ground truth domain Ω^* has horizontal symmetry but the inlet velocity profile deliberately breaks this symmetry. The inverse problem is the same as that in §3.1 but with different input parameters (see table 2). The initial guess Ω_0 is a rectangular domain with height equal to $0.85D$, centred in the image domain. For \mathbf{g}_{i0} , we take a skewed parabolic velocity profile with a peak velocity of approximately $2U$ that fits the inlet of Ω_0 .

The algorithm manages to reconstruct and segment the noisy flow images in 39 iterations, with total reconstruction error $\mathcal{E}^\bullet \simeq 2.87\%$. The results are presented in

Reconstruction and segmentation of noisy velocity images

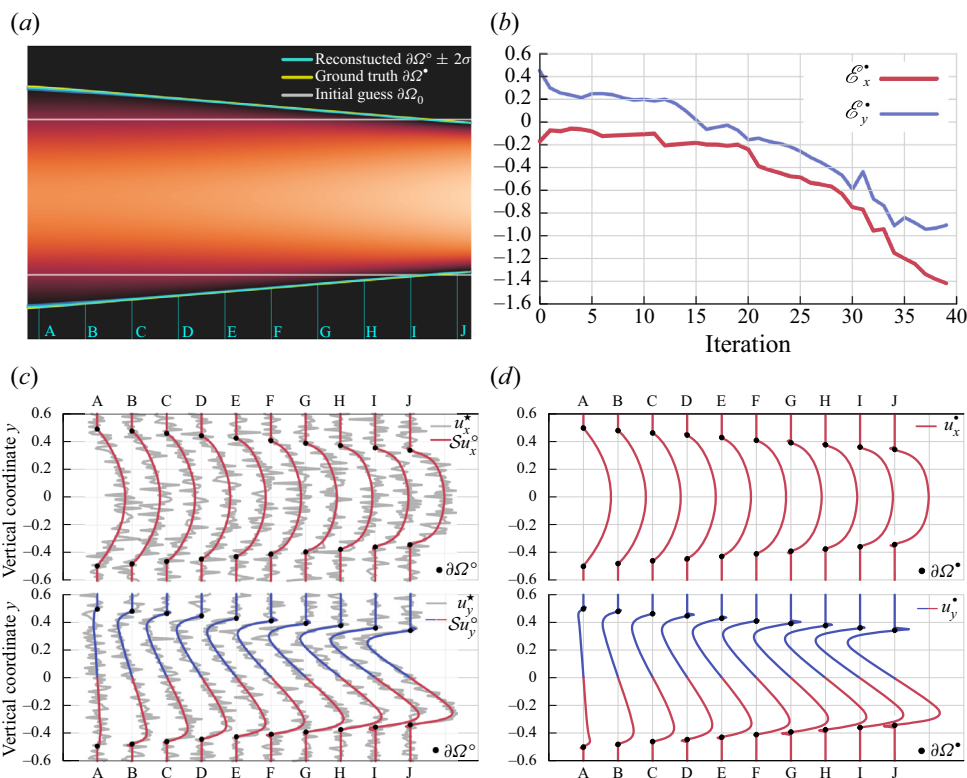


Figure 4. Reconstruction (algorithm 1) of synthetic images depicting the flow (from left to right) in a converging channel. Figure 4(a) depicts the reconstructed boundary $\partial\Omega^\circ$ (cyan line), the 2σ confidence region computed from the approximated posterior covariance $\tilde{\mathcal{C}}_\zeta \equiv \tilde{H}_\zeta \mathcal{C}_{\phi_\pm}$ (blue region), the ground truth boundary $\partial\Omega^*$ (yellow line) and the initial guess $\partial\Omega_0$ (white line). Figure 4(b) shows the reconstruction error as a function of iteration number. Velocity slices are drawn for 10 equidistant cross-sections (labelled with the letters A to J) for both the reconstructed images (figure 4c) and the ground truth (figure 4d), coloured red for positive values and blue for negative. (a) Velocity magnitude $|\mathbf{u}^\circ|$ and shape $\partial\Omega^\circ$. (b) Reconstruction error history. (c) Noisy data (grey) and reconstruction. (d) Ground truth velocity distributions.




Figure 5. (a) Reconstructed and (b) ground truth reduced hydrodynamic pressure (p) for the flow (from left to right) in the converging channel in figure 4. (a) Our reconstruction p° . (b) Ground truth p^* .

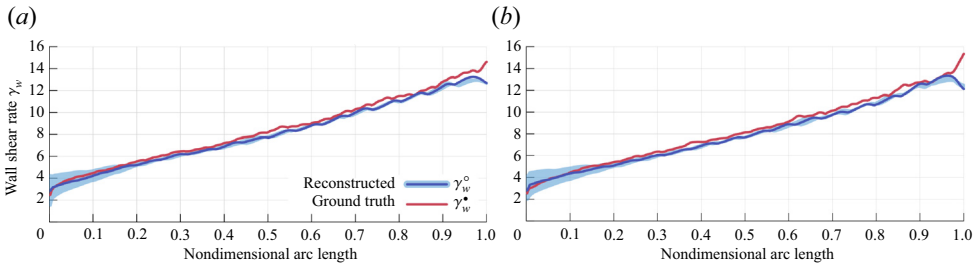


Figure 6. Wall shear rate $\gamma_w \equiv \boldsymbol{\tau} \cdot \partial_\nu \mathbf{u}$, where $\boldsymbol{\tau}$ is the unit tangent vector of $\partial\Omega$, for the converging channel flow in figure 4. The wall shear stress is found by multiplying this by the viscosity. The reconstructed wall shear rate (γ_w°) is calculated on $\partial\Omega^\circ$ and for \mathbf{u}° , while the ground truth (γ_w^\bullet) is calculated on $\partial\Omega^\bullet$ and for \mathbf{u}^\bullet . The blue region is bounded by the two wall shear rate distributions for \mathbf{u}° , calculated on the upper ($\partial\Omega_+^\circ$) and lower ($\partial\Omega_-^\circ$) limits of the 2σ confidence region of $\partial\Omega^\circ$. Note that the reconstructed solution can sometimes be found outside the blue region because the reconstructed shape $\partial\Omega^\circ$ may be less regular than $\partial\Omega_+^\circ$ or $\partial\Omega_-^\circ$. (a) Lower boundary. (b) Upper boundary.

		Image dim.	Model dim.	σ_{u_x}/U	σ_{u_y}/U
Simul. abd. aortic aneurysm	(2-D)	192^2	200^2	2.80×10^{-1}	5.26×10^{-3}
	<i>Regularization</i>	σ_{ϕ_\pm}/D	σ_{g_i}/U	σ_v/UD	Re_{ϕ_\pm}
Simul. abd. aortic aneurysm	(2-D)	1.0	2.0	.	0.1
					Re_ζ
					ℓ/h
					3

Table 2. Input parameters for the inverse 2-D Navier–Stokes problem.

figures 7 and 8. We observe that the discrepancy (figure 7c,f) consists mainly of Gaussian white noise. Again, some correlations are visible in the discrepancy of the y -velocity component at the upper inlet corner and the upper boundary of the simulated abdominal aortic aneurysm. The latter correlations (figure 7f) can be explained by the associated uncertainty in the predicted shape $\partial\Omega^\circ$ (figure 8a), which is well estimated for the upper boundary but slightly underestimated for the upper inlet corner. It is interesting to note that the upward skewed velocity profile at the inlet creates a region of low velocity magnitude on the lower boundary. The velocity profiles in this region produce low wall shear stresses, as seen in figures 8(c,d), and 10(b). These conditions are particularly challenging when one tries to infer the true boundary $\partial\Omega^\bullet$ because the local SNR is low ($\text{SNR} \ll 1$), meaning that there is considerable information loss there. Despite the above difficulties, algorithm 1 manages to approximate the posterior distribution of $\partial\Omega^\circ$ well, and successfully predicts extra uncertainty in this region (figure 8a). Again, the reconstructed pressure p° is indistinguishable from the ground truth p^\bullet (figure 9).

Using the reconstructions \mathbf{u}° and $\partial\Omega^\circ$, we compute the wall shear rate and we compare it with the ground truth in figure 10(b). We observe that the reconstructed solution approximates the ground truth well, even for very low SNRs ($\text{SNR} = 3$). Note that the waviness of the ground truth γ_w^\bullet is due to the relatively poor resolution of the level set function that we intentionally used to implicitly define this domain.

3.3. Synthetic data for 2-D flow in a simulated aortic aneurysm

Next, we test algorithm 1 in a channel that resembles the cross-section of an aorta that has an aneurysm in its ascending part. This test case is designed to demonstrate that

Reconstruction and segmentation of noisy velocity images




Figure 7. As for figure 3 but for the synthetic images depicting the flow (from left to right) in the simulated 2-D model of an abdominal aortic aneurysm. (a) Synthetic image u_x^* . (b) Our reconstruction u_x^* . (c) Discrepancy $\sigma_{u_x}^{-1}(u_x^* - S u_x^{\circ})$. (d) Synthetic image u_y^* . (e) Our reconstruction u_y^* . (f) Discrepancy $\sigma_{u_y}^{-1}(u_y^* - S u_y^{\circ})$.

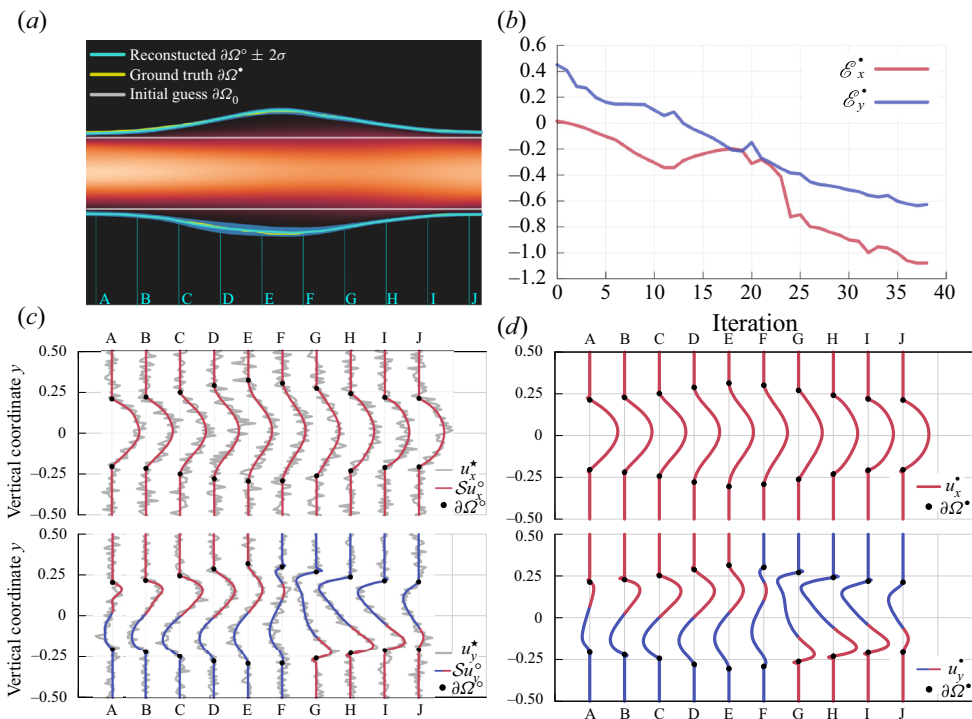


Figure 8. As for figure 4 but for the synthetic images depicting the flow (from left to right) in the simulated 2-D model of an abdominal aortic aneurysm. (a) Velocity magnitude $|u^{\circ}|$ and shape $\partial\Omega^{\circ}$. (b) Reconstruction error history. (c) Noisy data (grey) and reconstruction. (d) Ground truth velocity distributions.




Figure 9. (a) Reconstructed and (b) ground truth reduced hydrodynamic pressure (p) for the flow (from left to right) in the simulated 2-D model of an abdominal aortic aneurysm in figure 8. (a) Our reconstruction p° . (b) Ground truth p^\star .




Figure 10. As for figure 6 but for the synthetic images depicting the flow (from left to right) in the simulated 2-D model of an abdominal aortic aneurysm in figure 8. (a) Lower boundary. (b) Upper boundary.

		Image dim.	Model dim.		σ_{u_x}/U	σ_{u_y}/U		
Simul. aortic aneurysm	(2-D)	300^2	325^2		1.17×10^{-1}	2.62×10^{-1}		
	<i>Regularization</i>	σ_{ϕ_\pm}/D	$\sigma_{g_{ix}}/U$	$\sigma_{g_{iy}}/U$	σ_v/UD	Re_{ϕ_\pm}	Re_ζ	ℓ/h
Simul. aortic aneurysm	(2-D)	0.025	0.5	0.4	.	1	1	5

Table 3. Input parameters for the inverse 2-D Navier–Stokes problem.

the algorithm is applicable to realistic geometries with multiple inlets/outlets and for abnormal flow conditions (e.g. separation and recirculation zones). We generate synthetic images for \mathbf{u}^\star as in § 3.1, but for $SNR_x = SNR_y = 2.5$ and for $Re = 500$. For increased Reynolds numbers ($Re = 1000, 1500$), we observed vortex shedding within the aneurysm and we could not find a steady flow solution to generate synthetic images of steady flow. The inverse problem is the same as that in § 3.1 but with different input parameters (see table 3). The initial guess for the boundary of Ω_0 (figure 11a) is generated by using the Chan–Vese segmentation method (Chan & Vese 2001; Getreuer 2012a; Van Der Walt *et al.* 2014) on the noisy mask of the ground truth domain Ω^\star (figure 11b). The prior standard deviation σ_{ϕ_\pm} corresponds to the length of approximately 7 pixels of the noisy mask. The initial guess for the inlet velocity profile \mathbf{g}_{i0} is also shown in figure 11(a). Using the prior information of the boundary and the inlet velocity profile, algorithm 1 generates an initial guess for the Navier–Stokes velocity field (figure 12a,b) during its zeroth iteration.

Reconstruction and segmentation of noisy velocity images




Figure 11. Initial guesses (input for algorithm 1) for the geometry ($\partial\Omega_0$) and the inlet velocity profile ($g_{i,0}$) versus their corresponding ground truth (figure 11a) for the flow in the simulated 2-D model of an aortic aneurysm. The initial guess $\partial\Omega_0$ (figure 11a) is generated by segmenting the noisy mask (figure 11b) of the ground truth domain Ω^* . (a) Initial guesses (priors) for $\partial\Omega$ and g_i . (b) Noisy mask of Ω^* .




Figure 12. Zeroth iteration (N-S solution for the initial guesses in figure 11) velocity images (figure 12a,b), streamlines (figure 12c) and discrepancies with the data (figure 12d-f) for the flow in the simulated 2-D model of an aortic aneurysm. Streamlines are plotted on top of the velocity/discrepancy magnitude image and streamline thickness increases as the velocity magnitude increases. Figures 12(a-c) (colourbar not shown) and 12(d-f) (colourbar shown on the right) share the same colourmap. (a) $(u_x)_0$. (b) $(u_y)_0$. (c) u_0 . (d) $\sigma_{u_x}^{-1}(u_x^* - \mathcal{S}(u_x)_0)$. (e) $\sigma_{u_y}^{-1}(u_y^* - \mathcal{S}(u_y)_0)$. (f) $\mathcal{C}_u^{-1}(u^* - Su_0)$.

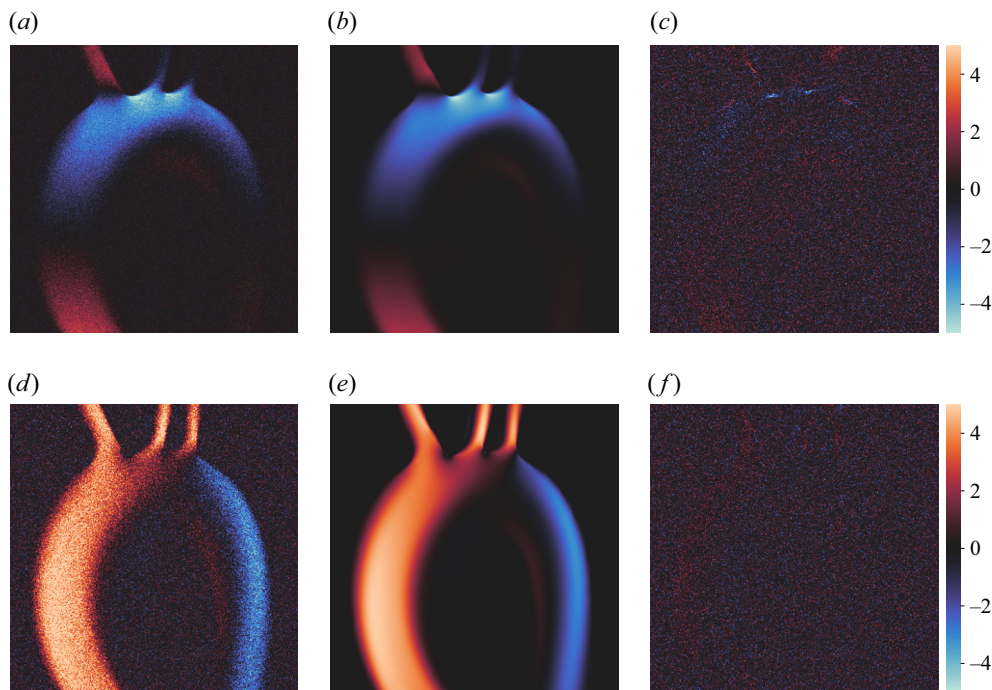


Figure 13. Reconstruction (final iteration of [algorithm 1](#)) of synthetic noisy velocity images depicting the flow in the simulated 2-D model of an aortic aneurysm. [Figures 13\(a,b\)](#) and [13\(d,e\)](#) show the horizontal, u_x , and vertical, u_y , velocities and share the same colourmap (colourbar not shown). [Figure 13\(c,f\)](#) shows the discrepancy between the noisy velocity images and the reconstruction (colourbars apply only to [figure 13c,f](#)). (a) Synthetic image u_x^* . (b) Our reconstruction u_x° . (c) Discrepancy $\sigma_{u_x}^{-1}(u_x^* - \mathcal{S}u_x^\circ)$. (d) Synthetic image u_y^* . (e) Our reconstruction u_y° . (f) Discrepancy $\sigma_{u_y}^{-1}(u_y^* - \mathcal{S}u_y^\circ)$.

The algorithm manages to reconstruct and segment the noisy flow images in 15 iterations, with a total reconstruction error $\mathcal{E}^\bullet \simeq 5.73\%$. The results are presented in [figures 13](#) and [14](#). We observe that the discrepancy of the last iteration ([figure 13c,f](#)) consists mainly of Gaussian white noise. Some correlations are visible in the discrepancy of the x -velocity component near the stagnation points of the upper branches, but these correlations are explained by the extra uncertainty in the predicted shape $\partial\Omega^\circ$ ([figure 14a](#)). By comparing [figures 13\(c,f\)](#) with [figure 12\(d,e\)](#), we confirm that the algorithm has successfully assimilated the remaining information from the noisy velocity measurements. [Figure 15](#) shows the pressure of the zeroth iteration ([figure 15a](#)), and the reconstructed pressure p° ([figure 15b](#)), which compares well to the ground truth pressure p^\bullet ([figure 15c](#)).

We further compare the performance of [algorithm 1](#) with a state-of-the-art image denoising algorithm, namely total variation denoising using Bregman iteration (TV-B) (Getreuer 2012b; Van Der Walt *et al.* 2014) in [figure 16](#). We first observe that [algorithm 1](#) denoises the velocity field without losing contrast near the walls of the aorta and accurately identifies the low-speed vortical structure within the aneurysm, which is obscured by noise. We then test three different values of the TV-B parameter λ/λ_0 , which controls the total variation regularization, and observe that even though TV-B manages to denoise the velocity field and reveal certain large scale vortices, there is considerable loss of contrast near the walls of the aorta and a systematic error (e.g. decreasing peak velocity)

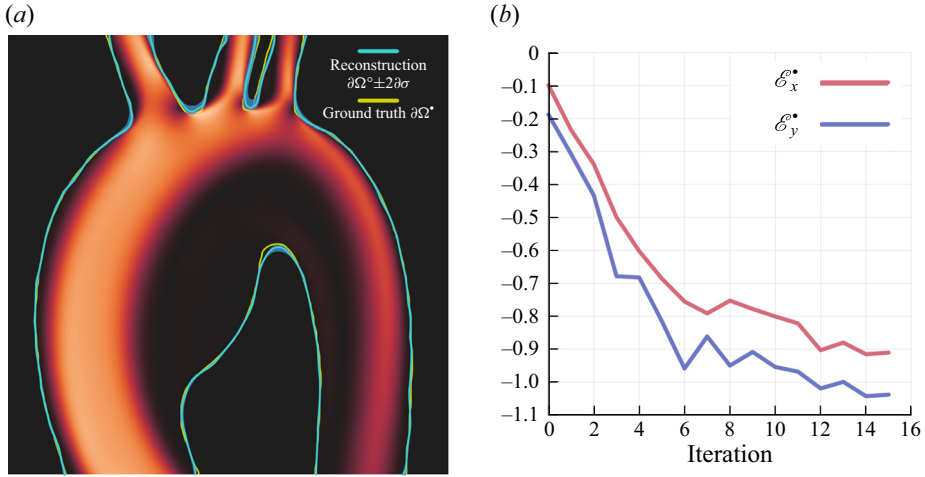


Figure 14. Reconstruction (final iteration of [algorithm 1](#)) of synthetic images depicting the flow in the simulated 2-D model of an aortic aneurysm. [Figure 14\(a\)](#) depicts the reconstructed boundary $\partial\Omega^\circ$ (cyan line), the 2σ confidence region computed from the approximated posterior covariance $\tilde{C}_\xi \equiv \tilde{H}_\xi C_{\phi_\pm}$ (blue region) and the ground truth boundary $\partial\Omega^\bullet$ (yellow line). [Figure 14\(b\)](#) shows the reconstruction error as a function of iteration number. (a) Velocity magnitude $|\mathbf{u}^\circ|$ and shape $\partial\Omega^\circ$. (b) Reconstruction error history.




Figure 15. (a) Zeroth iteration (N–S solution for the initial guesses in [figure 11](#)), (b) reconstructed (final iteration of [algorithm 1](#)) and (c) ground truth reduced hydrodynamic pressure for the simulated 2-D model of an aortic aneurysm in [figure 14](#). All panels share the same colourmap (symmetric logarithmic scale) and the same colourbar. (a) Zeroth iteration p_0 . (b) Our reconstruction p° . (c) Ground truth p^\bullet .

that increases as λ decreases. (The parameter $\lambda_0 = \lambda_0(\sigma)$, where σ is the noise standard deviation in the image, is given by [Getreuer \(2012b\)](#) as an optimal value for λ .)

Using the reconstructions \mathbf{u}° and $\partial\Omega^\circ$, we compute the reconstructed wall shear rate (γ_w°) and compare it with the ground truth (γ_w^\bullet) ([figure 17](#)). We observe that γ_w° approximates γ_w^\bullet well and that discrepancies are well accounted for by the $\gamma_w^\circ \pm 2\sigma$ -bounds.

3.4. Magnetic resonance velocimetry experiment

We measured the flow through a converging nozzle using magnetic resonance velocimetry ([Fukushima 1999](#); [Mantle & Sederman 2003](#); [Elkins & Alley 2007](#)). The nozzle converges from an inner diameter of 25 mm to an inner diameter of 13 mm, over a length of 40 mm ([figure 18b](#)). On either side of the converging section, the entrance-to-exit length equals




Figure 16. Streamlines for the flow in the simulated 2-D model of an aortic aneurysm (figures 13 and 14), and comparison with total variation denoising using Bregman iteration (TV-B) with different weights λ . Streamlines are plotted on top of the velocity magnitude image and streamline thickness increases as the velocity magnitude increases. (a) Synthetic data \mathbf{u}^* . (b) Our reconstruction \mathbf{u}° . (c) Ground truth \mathbf{u}^* . (d) TV-B $\lambda/\lambda_0 = 0.1$. (e) TV-B $\lambda/\lambda_0 = 0.01$. (f) TV-B $\lambda/\lambda_0 = 0.001$.

10 times the local diameter (figure 18b) to ensure the absence of entrance/exit effects. We acquired velocity images for a Reynolds number of 162 (defined at the nozzle outlet). We used a 40 wt% glycerol in water solution (Cheng 2008; Volk & Kähler 2018) as the working fluid to increase the viscosity and minimize the effect of thermal convection in the resulting velocity field due to the temperature difference between the magnet bore and the working fluid. The nozzle is made of polyoxymethylene to minimize magnetic susceptibility differences between the nozzle wall and the working fluid (Wapler *et al.* 2014). Figure 18(a) depicts the schematic of the flow loop of the MRV experiment. To pump the water/glycerol solution, we used a Watson Marlow 505S peristaltic pump (Watson Marlow, Falmouth UK) with a 2 l dampening vessel at its outlet to dampen flow oscillations introduced by the peristaltic pump. To make the flow uniform, we installed porous polyethylene distributor plates (SPC technologies, Fakenham UK) at the entrance and the exit of the nozzle.

We acquired the velocity images on a Bruker Spectrospin DMX200 with a 4.7 T superconducting magnet, which is equipped with a gradient set providing magnetic field gradients of a maximum strength of 13.1 G cm^{-1} in three orthogonal directions, and a birdcage radiofrequency coil tuned to a ^1H frequency of 199.7 MHz with a diameter and a length of 6.3 cm. To acquire 2-D velocity images, we used slice-selective spin-echo imaging (Edelstein *et al.* 1980) combined with pulsed gradient spin-echo (PGSE) (Stejskal & Tanner 1965) for motion encoding (figure 18c). We measured each of the three orthogonal velocity components in a 1 mm-thick transverse slice through the converging

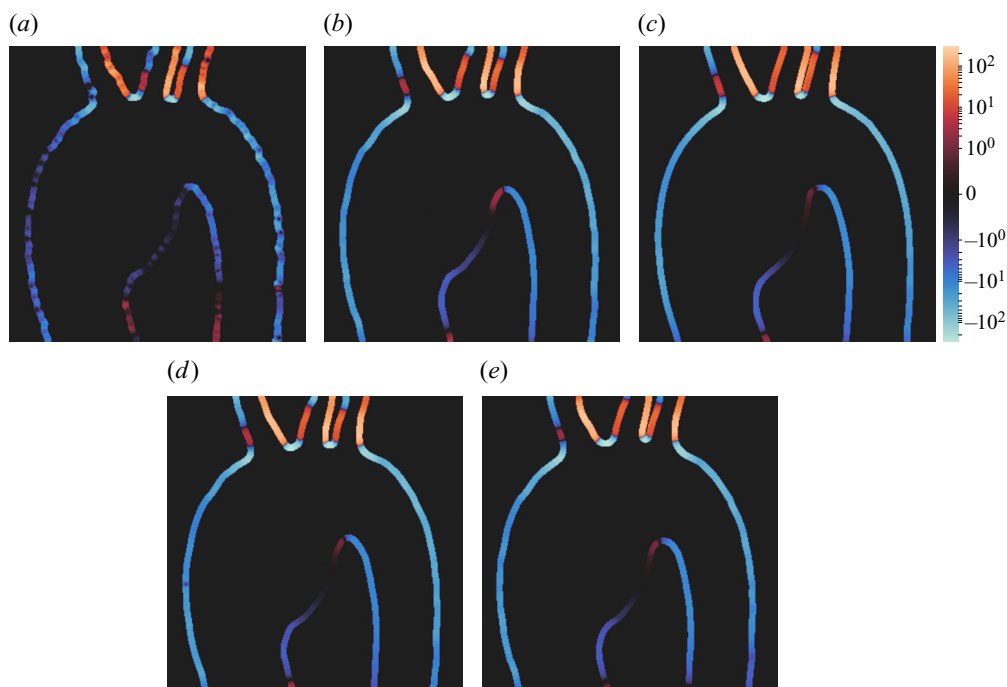


Figure 17. Wall shear rate $\gamma_w \equiv \boldsymbol{\tau} \cdot \partial_v \mathbf{u}$, where $\boldsymbol{\tau}$ is the unit tangent vector of $\partial\Omega$, for the flow in the simulated 2-D model of an aortic aneurysm in figure 14. The wall shear stress is found by multiplying this by the viscosity. The reconstructed wall shear rate (γ_w°) is calculated on $\partial\Omega^\circ$ and for \mathbf{u}° , while the ground truth (γ_w^\bullet) is calculated on $\partial\Omega^\bullet$ and for \mathbf{u}^\bullet . The $\pm 2\sigma$ -bounds are calculated on the upper ($\partial\Omega_+^\circ$) and lower ($\partial\Omega_-^\circ$) limits of the confidence region of $\partial\Omega^\circ$. All panels share the same colourmap (symmetric logarithmic scale) and the same colourbar. (a) Zeroth iteration $(\gamma_w)_0$. (b) Our reconstruction γ_w° . (c) Ground truth γ_w^\bullet . (d) Lower confidence bound $\gamma_w^\circ - 2\sigma$. (e) Upper confidence bound $\gamma_w^\circ + 2\sigma$.

section of the nozzle, which is centred along the nozzle centreline. The flow images we acquired have a field of view of 84.2×28.6 mm at 512×128 pixels, giving an in-plane resolution of 165×223 μm . For velocity measurements in the net flow direction, we used a gradient pulse duration, δ , of 0.3–0.5 ms and flow observation times, Δ , of 9–12 ms. For velocity measurements in the perpendicular to the net flow direction, we used an increased gradient pulse duration, δ , of 1.0 ms and an increased observation time, Δ , of 25–30 ms, due to the lower velocity magnitudes in this direction. We set the amplitude, g , of the flow encoding gradient pulses to $\pm 3 \text{ G cm}^{-1}$ for the direction parallel to the net flow and to $\pm 1.5 \text{ G cm}^{-1}$ for the direction perpendicular to the net flow to maximize phase contrast whilst avoiding velocity aliasing by phase wrapping. To obtain an image for each velocity component, we took the phase difference between two images acquired with flow encoding gradients having equal magnitude g but opposite signs. To remove any phase shift contributions that are not caused by the flow, we corrected the measured phase shift of each voxel by subtracting the phase shift measured under zero-flow conditions. The gradient stabilization time that we used is 1 ms and we acquired the signal with a sweep width of 100 kHz. We used hard 90° excitation pulses with a duration of 85 μs , and a 512 μs Gaussian-shaped soft 180° pulse for slice selection and spin-echo refocusing. We found the T_1 relaxation time of the glycerol solution to be 702 ms, as measured by an inversion recovery pulse sequence. To allow for magnetization recovery between the




Figure 18. Schematic of the rig that we use to conduct the MRV experiment consisting of: (1) 20 l holding tank; (2) peristaltic pump; (3) 2 l vessel; (4) clamp valves; (5) porous polyethylene distributor; (6) radiofrequency probe; (7) converging nozzle; (8) 4.7 T superconducting magnet; (9) volumetric cylinder for flow measurements. Figure 18(b) shows a sketch of the converging nozzle with the active area of the spectrometer shown by a red box. The pulse sequence that we use for 2-D velocity imaging is shown in figure 18(c). (a) Magnet and flow loop. (b) Converging nozzle. (c) Spin-echo pulse sequence with slice selective refocusing and flow encoding.

acquisitions, we used a repetition time of 1.0 s. To eliminate unwanted coherences and common signal artefacts, such as DC offset, we used a four step phase cycle.

To be consistent with the standard definition used in MRI/MRV, we define the SNR of each MRV image using (3.1a,b), but with μ_x replaced by the mean signal intensity (images of the ^1H spin density) over the nozzle domain (μ_I) and σ_{u_x} replaced by the standard deviation of the Rayleigh distributed noise in a region with no signal (σ_I) (Gudbjartsson & Patz 1995). The standard deviation for the phase is therefore $\sigma_\phi = 1/\text{SNR}$. The MRV images are acquired by taking the sum/difference of four phase images and then multiplying by the constant factor $1/2\gamma g\delta\Delta$, where γ is the gyromagnetic ratio of ^1H (linear relation between the image phase and the velocity). The error in the MRV measured velocity is therefore $\sigma_u = \sigma_\phi/\gamma g\delta\Delta$. To acquire high SNR images (figure 19), we averaged 32 scans, resulting in a total acquisition time of 137 minutes per velocity image (~ 4.6 h for both velocity components). To evaluate the denoising capability of the algorithm, we acquired poor SNR images by averaging only four scans (the minimum requirement




Figure 19. High SNR images (average of 32 scans with $\text{SNR}_z \simeq 44$, $\text{SNR}_r \simeq 34$) that we acquired for the flow through the converging nozzle using MRV (units in cm s^{-1}). (a) u_z^* . (b) u_r^* .

for a full phase cycle) and decreasing the repetition time to 300 ms, resulting in a total acquisition time of 5.1 min per velocity image (10.2 min for both velocity components).

To verify the quantitative nature of the MRV experiment, we compared the volumetric flow rates calculated from the MRV images (using 2-D slice-selective velocity imaging in planes normal to the direction of net flow) with the volumetric flow rates measured from the pump outlet. The results agree with an average error of $\pm 1.8\%$.

3.5. Magnetic resonance velocimetry data in a converging nozzle

We now use [algorithm 1](#) to reconstruct and segment the low SNR images (\mathbf{u}^*) that we acquired during the MRV experiment (§ 3.4), and compare them with the high SNR images of the same flow (\mathbf{u}^* in [figure 19](#)). The flow is axisymmetric with zero swirl. The subscript ‘ x ’ is replaced by ‘ z ’, which denotes the axial component of velocity, and the subscript ‘ y ’ is replaced by ‘ r ’, which denotes the radial component of velocity. The low SNR images ($\text{SNR}_z = 6.7$, $\text{SNR}_r = 5.8$) required a total scanning time of 5.1 minutes per velocity image (axial and radial components) and the high SNR images ($\text{SNR}_z = 44.2$, $\text{SNR}_r = 34.4$) required a total scanning time of 137 minutes per velocity image. Since the signal intensity of an MRV experiment corresponds to the ^1H spin density, we segment the spin density image using a thresholding algorithm (Otsu 1979) to obtain a mask ψ , such that $\psi = 1$ inside Ω (the nozzle) and $\psi = 0$ outside Ω . We consider ψ to be the prior information for the geometry of the nozzle, which also serves as an initial guess for Ω (Ω_0). For \mathbf{g}_{i0} , we take a parabolic velocity profile with a peak velocity of $0.6U$, where $U \simeq 5 \text{ cm s}^{-1}$ is the characteristic velocity for this problem. In this case, we treat the kinematic viscosity as an unknown, with a prior distribution $\mathcal{N}(\bar{\nu}, (0.1\bar{\nu})^2)$ and $\bar{\nu} = 4 \times 10^{-6} \text{ m}^2 \text{ s}^{-1}$. Note that the axis of the nozzle is not precisely known beforehand, and since we only solve an axisymmetric Navier–Stokes problem on the $z - r$ half-plane, we also introduce an unknown variable for the vertical position of the axis (see [Appendix C](#)).

Using the input parameters of [table 4](#), the algorithm manages to reconstruct the noisy velocity image and reduce segmentation errors in just six iterations, with a total reconstruction error $\mathcal{E}^* \simeq 5.94\%$. The results are presented in [figures 20](#) and [21](#). We observe that [algorithm 1](#) manages to filter out the noise, the outliers and the acquisition artefacts of the low SNR MRV images depicting the axial u_z^* ([figure 20a](#)) and the radial u_r^* ([figure 20d](#)) components of velocity. A notable difference between these real MRV images and the synthetic MRV images in §§ 3.1 and 3.2 is that the real MRV images

Nozzle	(3-D)	Image dimension	Model dimension	σ_{u_z}/U		σ_{u_r}/U	
		255×128	300×130 (half-plane)	1.4168×10^{-1}	3.0679×10^{-2}		
Regularization		$\sigma_{\phi_{\pm}}/D$	σ_{g_i}/U	σ_v/UD	$Re_{\phi_{\pm}}$	Re_{ξ}	ℓ/h
Nozzle	(3-D)	0.25	0.5	6.2×10^{-4}	0.025	0.025	3

Table 4. Input parameters for the inverse 3-D axisymmetric Navier–Stokes problem.




Figure 20. Reconstruction (algorithm 1) of low SNR MRV velocity images depicting the axisymmetric flow (from left to right) in the converging nozzle (figure 18b). Figures 20(a,b) and 20(d,e) show the horizontal, u_x , and vertical, u_y , velocities and share the same colourmap (colourbar not shown). Figure 20(c,f) show the discrepancy between the noisy velocity images and the reconstruction (colourbars apply only to figure 20c,f). The reconstructed flow \mathbf{u}° is axisymmetric by construction; therefore, u_z° depicts an even reflection and u_r° depicts an odd reflection, so that they can be compared with the MRV images (see Appendix C). (a) Low SNR MRV image u_z^* . (b) Our reconstruction u_z° . (c) Discrepancy $\sigma_{u_z}^{-1}(u_z^* - \mathcal{S}u_z^\circ)$. (d) Low SNR MRV image u_r^* . (e) Our reconstruction u_r° . (f) Discrepancy $\sigma_{u_r}^{-1}(u_r^* - \mathcal{S}u_r^\circ)$.

display artefacts and contain outliers. We have not pre-processed the MRV images, for example, by removing outliers. The estimated posterior uncertainty of $\partial\Omega^\circ$ is depicted in figure 21(a), in which we observe that regions with gaps in the data coincide with regions of higher uncertainty. Although we treat the kinematic viscosity ν as an unknown parameter, the posterior distribution of ν remains effectively unchanged. More precisely, we infer a kinematic viscosity of $\nu^\circ = 3.995 \times 10^{-6} \simeq \bar{\nu}$, with a posterior variance of $(0.1005\bar{\nu})^2$. This is because we use a Bayesian approach to this inverse problem, where the prior information for ν is already rich enough. Technically, the reconstruction functional \mathcal{E} is insensitive to small changes of ν (or $1/Re$), and, as a result, the prior term in the gradient of ν (2.35) dominates; i.e. the model \mathcal{M} is not informative. Physically, it is not possible to infer ν (with reasonable certainty) for this particular flow without additional information on pressure.

As in § 3.3, we compare the denoising performance of algorithm 1 (figure 20) with TV-B (Getreuer 2012b; Van Der Walt *et al.* 2014) (figure 22). We again observe that algorithm 1 has managed to filter out both the noise and the artefacts, while the TV-B-denoised images present artefacts, loss of contrast and a systematic error that depends on the parameter λ .

Reconstruction and segmentation of noisy velocity images

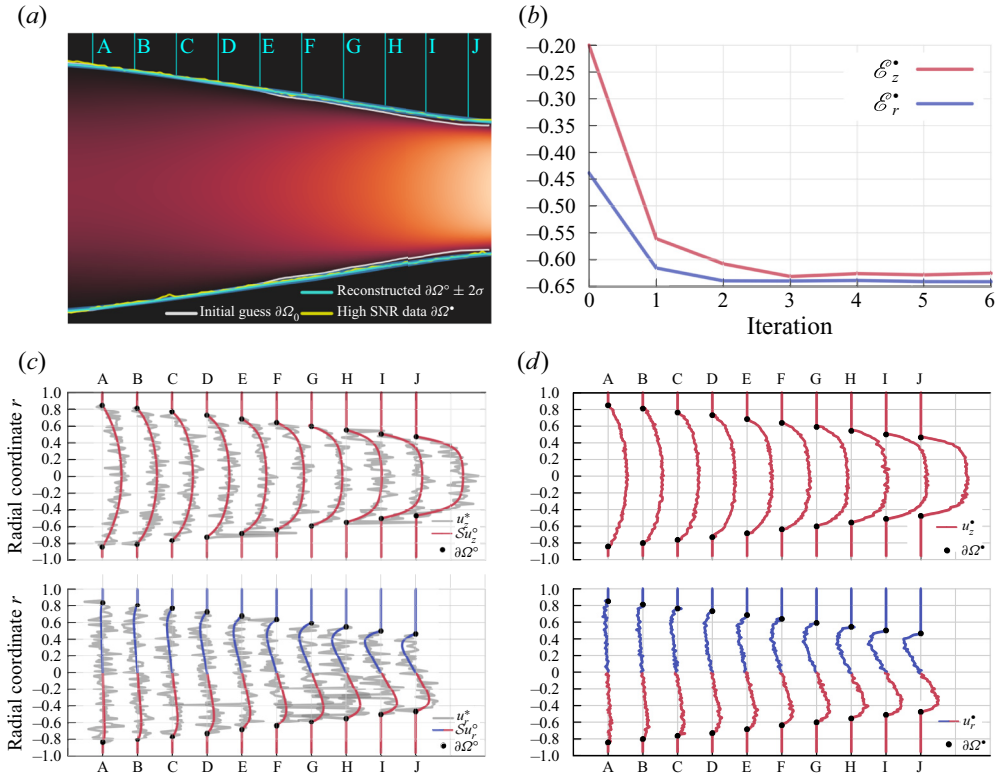


Figure 21. Reconstruction (algorithm 1) of synthetic images depicting the axisymmetric flow (from left to right) in the converging nozzle (figure 18b). Figure 21(a) depicts the reconstructed boundary $\partial\Omega^\circ$ (cyan line), the 2σ confidence region computed from the approximated posterior covariance $\tilde{\mathcal{C}}_\zeta \equiv \tilde{H}_\zeta \mathcal{C}_{\phi_\pm}$ (blue region), the ground truth boundary $\partial\Omega^\bullet$ (yellow line) and the initial guess $\partial\Omega_0$ (white line). Figure 21(b) shows the reconstruction error as a function of iteration number. Velocity slices are drawn for 10 equidistant cross-sections (labelled with the letters A to J) for both the reconstructed images (figure 21c) and the high SNR images (figure 21d), coloured red for positive values and blue for negative. (a) Velocity magnitude $|\mathbf{u}^\circ|$ and shape $\partial\Omega^\circ$. (b) Reconstruction error history. (c) Low SNR data (grey) and reconstruction. (d) High SNR velocity data.

Figure 23(a) shows the reconstructed wall shear rate γ_w° , computed for the reconstructed velocity field \mathbf{u}° on the segmented shape $\partial\Omega^\circ$, and compares it with the ground truth wall shear rate γ_w^\bullet computed for the high SNR velocity field \mathbf{u}^\bullet (figure 19) on the high SNR shape $\partial\Omega^\bullet$ (^1H spin density). We observe that the ground truth wall shear rate is particularly noisy, as MRV suffers from low resolution and partial volume effects (Bouillot *et al.* 2018; Saito *et al.* 2020) near the boundaries $\partial\Omega$. Certainly, it is possible to smooth the boundary $\partial\Omega^\bullet$ (which we obtained using the method of Otsu (1979) for the ^1H spin density) using conventional image processing algorithms. However, the velocity field \mathbf{u}^\bullet will not be consistent with the new smoothed boundary (the no-slip boundary condition will not be satisfied). The method that we propose here for the reconstruction and segmentation of MRV images tackles exactly this problem: it infers the most likely shape of the boundary ($\partial\Omega^\circ$) from the velocity field itself, without requiring an additional experiment (e.g. CT, MRA) or manual segmentation using another software. Furthermore, in this Bayesian setting, we can use the ^1H spin density to introduce *a priori* knowledge of $\partial\Omega$ in the form of a prior, which would prove useful in areas of low velocity




Figure 22. Total variation denoising using Bregman iteration with different weights λ for the low SNR MRV images (figure 20a,d) depicting the axisymmetric flow (from left to right) in the converging nozzle. (a) u_z , TV-B $\lambda/\lambda_0 = 0.1$. (b) u_z , TV-B $\lambda/\lambda_0 = 0.01$. (c) u_z , TV-B $\lambda/\lambda_0 = 0.001$. (d) u_r , TV-B $\lambda/\lambda_0 = 0.1$. (e) u_r , TV-B $\lambda/\lambda_0 = 0.01$. (f) u_r , TV-B $\lambda/\lambda_0 = 0.001$.




Figure 23. (a) Wall shear rates (as for figure 6) and (b) reduced hydrodynamic pressure inferred from the MRV images depicting the axisymmetric flow in the converging nozzle. (a) Wall shear rates γ_w^o and γ_w^* . (b) Our pressure reconstruction p^o .

magnitudes where the velocity field itself does not provide enough information to segment the boundaries, e.g. flow within an aneurysm or a heart ventricle (Demirkiran *et al.* 2021). As a result, algorithm 1 performs very well in estimating the posterior distribution of wall shear rate, a quantity which depends both on the velocity field and the boundary shape, and which is hard to measure otherwise.

3.6. Choosing the regularization parameters

Regularization is crucial to successfully reconstruct the velocity field and segment the geometry of the nozzle in the presence of noise, artefacts and outliers. Regularization comes from the Navier–Stokes problems (primal and adjoint) (\mathcal{M}), and the regularization of the model parameters (\mathcal{R}).

3.6.1. Notes on prior information for the Navier–Stokes unknowns

By adopting a Bayesian inference framework, we assume that the prior information of an unknown x is a Gaussian random field with mean \bar{x} and covariance \mathcal{C}_x , i.e. $x \sim \mathcal{N}(\bar{x}, \mathcal{C}_x)$ (see § 2.2 and Appendix A). We, therefore, need to provide algorithm 1 with a prior mean and a prior covariance for every N–S unknown. For the inlet velocity boundary condition, $\bar{\mathbf{g}}_i$ can be a smooth approximation to the noisy velocity data at the inlet and then $\sigma_{\mathbf{g}_i}$ is the prior standard deviation around this mean. For the outlet natural boundary condition, $\bar{\mathbf{g}}_o$ can be 0 and then $\sigma_{\mathbf{g}_o}$ determines the confidence of the user regarding whether or not the outlet is a pseudotraction-free boundary. For both the inlet and the outlet boundary conditions, the parameter ℓ , which can be different for each boundary condition, controls the regularity of the functions \mathbf{g}_i and \mathbf{g}_o , i.e. length scales smaller than ℓ are suppressed. For the shape, $\bar{\phi}_{\pm}$ can be a rough segmentation of the original geometry and then $\sigma_{\phi_{\pm}}$ is the prior standard deviation around this mean. For example, in § 3.3, we set $\sigma_{\phi_{\pm}}$ approximately equal to a length of seven pixels by visually inspecting the noisy mask (figure 11b). The same methodology applies to the determination of prior information regarding the kinematic viscosity ν .

The advantage of this probabilistic framework is that when prior information is available, it can be readily exploited to regularize the inverse problem and facilitate its numerical solution. However, if there is no prior information available regarding an unknown, we can assume that this unknown is distributed according to a zero-mean Gaussian distribution with a sufficiently large standard deviation σ .

3.6.2. Notes on shape regularization and the choice of $Re_{\phi_{\pm}}, Re_{\zeta}$

For the axisymmetric nozzle (see § 3.5), we avoid overfitting the shape $\partial\Omega$ by choosing the Reynolds numbers for the geometric flow to be $Re_{\phi_{\pm}} = Re_{\zeta} = 0.025$. Increasing these Reynolds numbers to approximately 1.0, we start noticing that the assimilated boundary becomes more susceptible to noise in the image. However, for the simulated aortic aneurysm (see § 3.3), we chose $Re_{\phi_{\pm}} = Re_{\zeta} = 1.0$ to preserve high curvature regions. From numerical experiments, we have observed that typical successful values for the Reynolds numbers $Re_{\phi_{\pm}}, Re_{\zeta}$ lie in the interval (0.01, 0.1) for low SNR images (SNR < 10) with relatively flat boundaries, in (0.1, 1.0) for higher SNR images (SNR \geq 10) with relatively flat boundaries and ≥ 1 for geometries with regions of high curvature. Physical intuition that justifies the use of $Re_{\phi_{\pm}}, Re_{\zeta}$ as the preferred shape regularization parameters is provided in § 2.4.2.

4. Conclusions

We have formulated a generalized inverse Navier–Stokes problem for the joint reconstruction and segmentation of noisy velocity images of steady incompressible flow. To regularize the inverse problem, we adopt a Bayesian framework by assuming Gaussian prior distributions for the unknown model parameters. Although the inverse problem is formulated using variational methods, every iteration of the nonlinear problem is actually equivalent to a Gaussian process in Hilbert spaces. We implicitly define the boundaries of the flow domain in terms of signed distance functions and use Nitsche’s method to weakly enforce the Dirichlet boundary condition on the moving front. The moving of the boundaries is expressed by a convection-diffusion equation for the signed distance function, which allows us to control the regularity of the boundary by tuning an artificial diffusion coefficient. We use the steepest ascent directions of the model parameters in conjunction with a quasi-Newton method (BFGS), and we show how the posterior

Gaussian distribution of a model parameter can be estimated from the reconstructed inverse Hessian.

We devise an algorithm that solves this inverse Navier–Stokes problem and test it for noisy (SNR = 2.5, 3) 2-D synthetic images of Navier–Stokes flows. The algorithm successfully reconstructs the velocity images, infers the most likely boundaries of the flow and estimates their posterior uncertainty. We then design a magnetic resonance velocimetry (MRV) experiment to obtain images of a 3-D axisymmetric Navier–Stokes flow in a converging nozzle. We acquire MRV images of poor quality (SNR \simeq 6), intended for reconstruction/segmentation, and images of higher quality (SNR > 30) that serve as the ground truth. We show that the algorithm performs very well in reconstructing and segmenting the poor MRV images, which were obtained in just 10.2 minutes, and that the reconstruction compares well to the high SNR images, which required a total acquisition time of \sim 4.6 h. Lastly, we use the reconstructed images and the segmented (smoothed) domain to estimate the posterior distribution of the wall shear rate and compare it with the ground truth. Since the wall shear rate depends on both the shape and the velocity field, we note that our algorithm provides a consistent treatment to this problem by jointly reconstructing and segmenting the flow images, avoiding the design of an additional experiment (e.g. CT, MRA) for the measurement of the geometry or the use of external (non-physics-informed) segmentation software.

The present method has several advantages over general image reconstruction and segmentation algorithms, which do not respect the underlying physics and the boundary conditions, and, at the same time, provides additional knowledge of the flow physics (e.g. pressure field and wall shear stress), which is otherwise difficult to measure. It can be used to substantially decrease signal acquisition times and provides additional knowledge of the physical system being imaged. Although our current implementation is restricted to 2-D planar and axisymmetric flows, the method naturally extends to periodic and unsteady Navier–Stokes problems in complicated 3-D geometries.

Declaration of interests. The authors report no conflict of interest.

Author ORCIDs.

✉ Alexandros Kontogiannis <https://orcid.org/0000-0001-6353-3427>;

✉ Andrew J. Sederman <https://orcid.org/0000-0002-7866-5550>;

✉ Matthew P. Juniper <https://orcid.org/0000-0002-8742-9541>.

Appendix A. Gaussian measures in Hilbert spaces

The mean of a Gaussian measure $\gamma \sim \mathcal{N}(m, C)$ in L^2 is given by

$$m \equiv \mathbb{E}h := \int_{L^2} h \gamma(dh). \tag{A1}$$

The covariance operator $C : L^2 \rightarrow L^2$ and the covariance $C : L^2 \times L^2 \rightarrow \mathbb{R}$ are defined by

$$Cx := \int_{L^2} h \langle x, h \rangle \gamma(dh), \quad C(x, x') := \int_{L^2} \langle x, h \rangle \langle x', h \rangle \gamma(dh), \tag{A2a,b}$$

noting that $\langle Cx, x' \rangle = C(x, x')$. The above (Bochner) integrals define integration over the function space L^2 and under the measure γ , and are well defined due to Fernique’s theorem (Hairer 2009). These integrals can be directly computed by sampling the Gaussian measure γ with Karhunen–Loève expansion (as in § 2.6).

Appendix B. Euler–Lagrange system

The integration by parts formulae for the nonlinear term (2.21) are

$$\begin{aligned} \int_{\Omega} (\mathbf{u}' \cdot \nabla \mathbf{u}) \cdot \mathbf{v} &= \int_{\Omega} u'_j \partial_j u_i v_i = - \int_{\Omega} u'_j u_i \partial_j v_i + \int_{\partial\Omega} u'_j v_j u_i v_i \\ &= - \int_{\Omega} (\mathbf{u} \cdot (\nabla \mathbf{v})^\dagger) \cdot \mathbf{u}' + \int_{\partial\Omega} (\mathbf{u} \cdot \mathbf{v})(\mathbf{v} \cdot \mathbf{u}'), \end{aligned} \tag{B1}$$

$$\begin{aligned} \int_{\Omega} (\mathbf{u} \cdot \nabla \mathbf{u}') \cdot \mathbf{v} &= \int_{\Omega} u_j \partial_j u'_i v_i = - \int_{\Omega} \partial_j u_j u'_i v_i + \int_{\partial\Omega} u_j u'_i v_j v_i \\ &= - \int_{\Omega} (\mathbf{u} \cdot \nabla \mathbf{v}) \cdot \mathbf{u}' + \int_{\partial\Omega} (\mathbf{u} \cdot \mathbf{v})(\mathbf{v} \cdot \mathbf{u}'). \end{aligned} \tag{B2}$$

Appendix C. Axisymmetric inverse Navier–Stokes problem

The axisymmetric Navier–Stokes problem is

$$\mathbf{u} \cdot \nabla \mathbf{u} - \nu \Delta \mathbf{u} + \nabla p + \mathbf{f} = \mathbf{0}, \quad \nabla \cdot \mathbf{u} = 0, \tag{C1}$$

where

$$\left. \begin{aligned} \mathbf{u} &= u_z \hat{\mathbf{z}} + u_r \hat{\mathbf{r}}, \quad \nabla \mathbf{u} = (\partial_z \mathbf{u}, \partial_r \mathbf{u}), \quad \Delta \mathbf{u} = \partial_z^2 u_z + \partial_r^2 u_r + \frac{1}{r} \partial_r u_r, \\ \nabla \cdot \mathbf{u} &= \partial_z u_z + \partial_r u_r + \frac{u_r}{r}, \quad \mathbf{f} = \left(0, \frac{\nu u_r}{r^2} \right), \end{aligned} \right\} \tag{C2}$$

and the nonlinear term $\mathbf{u} \cdot \nabla \mathbf{u}$ retains the same form as in the Cartesian frame.

To compare the axisymmetric modelled velocity field with the MRV images, we introduce two new operators: (i) the reflection operator $\mathcal{R} : \mathbb{R}^+ \times \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}$ and (ii) a rigid transformation $\mathcal{T} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$. The reconstruction error is then expressed by

$$\mathcal{E}(\mathbf{u}) \equiv \frac{1}{2} \|\mathbf{u}^\star - ST\mathcal{R}\mathbf{u}\|_{C_u}^2 := \frac{1}{2} \int_I (\mathbf{u}^\star - ST\mathcal{R}\mathbf{u}) C_u^{-1} (\mathbf{u}^\star - ST\mathcal{R}\mathbf{u}) \, dx \, dy. \tag{C3}$$

We introduce an unknown variable for the vertical position of the axisymmetry axis by letting $\mathcal{T}u = u(x, y + y_0)$, for $y_0 = \text{const}$. Then, the generalized gradient for y_0 is

$$\langle D_{y_0} \mathcal{J}, y'_0 \rangle_{\mathbb{R}} = \left\langle - \int_I C_u^{-1} (\mathbf{u}^\star - ST\mathcal{R}\mathbf{u}) (ST\mathcal{R} \partial_y \mathbf{u}), y'_0 \right\rangle_{\mathbb{R}}, \tag{C4}$$

and y_0 is treated in the same way as the inverse Navier–Stokes problem unknowns \mathbf{x} .

REFERENCES

BENNING, M. & BURGER, M. 2018 Modern regularization methods for inverse problems. *Acta Numer.* **27**, 1–111.
 BENNING, M., GLADDEN, L., HOLLAND, D., SCHÖNLIEB, C.B. & VALKONEN, T. 2014 Phase reconstruction from velocity-encoded MRI measurements – A survey of sparsity-promoting variational approaches. *J. Magn. Reson.* **238**, 26–43.
 BENZI, M., GOLUBT, G.H. & LIESEN, J. 2005 Numerical solution of saddle point problems. *Acta Numer.* **14**, 1–137.
 BENZI, M. & OLSHANSKII, M.A. 2006 An augmented Lagrangian-based approach to the Oseen problem. *SIAM J. Sci. Comput.* **28** (6), 2095–2113.
 BOGACHEV, V.I. 1998 *Gaussian Measures*. American Mathematical Society.

- BOUILLOT, P., DELATTRE, B.M.A., BRINA, O., OUARED, R., FARHAT, M., CHNAFA, C., STEINMAN, D.A., LOVLAD, K.O., PEREIRA, V.M. & VARGAS, M.I. 2018 3D phase contrast MRI: partial volume correction for robust blood flow quantification in small intracranial vessels. *Magn. Reson. Med.* **79** (1), 129–140.
- BRENNER, S. & SCOTT, R. 2008 *The Mathematical Theory of Finite Element Methods*, vol. 15, Texts in Applied Mathematics. Springer.
- BURGER, M. 2001 A level set method for inverse problems. *Inverse Probl.* **17** (5), 1327–1355.
- BURGER, M. 2003 A framework for the construction of level set methods for shape optimization and reconstruction. *Interfaces Free Bound.* **5** (3), 301–329.
- BURGER, M. & OSHER, S.J. 2005 A Survey in Mathematics for Industry: A survey on level set methods for inverse problems and optimal design. *Eur. J. Appl. Maths* **16** (2), 263–301.
- BURMAN, E. 2010 La pénalisation fantôme. *C. R. Math.* **348** (21–22), 1217–1220.
- BURMAN, E. & HANSBO, P. 2012 Fictitious domain finite element methods using cut elements: II. A stabilized Nitsche method. *Appl. Numer. Maths* **62** (4), 328–341.
- CHAN, T.F. & VESE, L.A. 2001 Active contours without edges. *IEEE Trans. Image Process.* **10** (2), 266–277.
- CHENG, N.-S. 2008 Formula for the Viscosity of a Glycerol – Water Mixture. *Ind. Engng Chem. Res.* **47** (9), 3285–3288.
- CODINA, R. 2002 Stabilized finite element approximation of transient incompressible flows using orthogonal subscales. *Comput. Meth. Appl. Mech. Engng* **191** (39–40), 4295–4321.
- CORONA, V., BENNING, M., GLADDEN, L.F., RECI, A., SEDERMAN, A.J. & SCHOENLIEB, C.-B. 2021 Joint phase reconstruction and magnitude segmentation from velocity-encoded MRI data, In *Time-dependent Problems in Imaging and Parameter Identification*. Springer International Publishing.
- COTTER, S.L., DASHTI, M., ROBINSON, J.C. & STUART, A.M. 2009 Bayesian inverse problems for functions and applications to fluid mechanics. *Inverse Probl.* **25** (11), 115008.
- CRANE, K., WEISCHEDEL, C. & WARDETZKY, M. 2017 The heat method for distance computation. *Commun. ACM* **60** (11), 90–99.
- DEMIRKIRAN, A., *et al.* 2021 Clinical intra-cardiac 4D flow CMR: acquisition, analysis, and clinical applications. *Eur. Heart J. Cardiovasc. Imag.* **23** (2), 154–165.
- DONOHO, D.L. 2006 Compressed sensing. *IEEE Trans. Inf. Theory* **52** (4), 1289–1306.
- EDELSTEIN, W.A., HUTCHISON, J.M.S., JOHNSON, G. & REDPATH, T. 1980 Spin warp NMR imaging and applications to human whole-body imaging. *Phys. Med. Biol.* **25** (4), 751–756.
- ELKINS, C.J. & ALLEY, M.T. 2007 Magnetic resonance velocimetry: applications of magnetic resonance imaging in the measurement of fluid motion. *Exp. Fluids* **43** (6), 823–858.
- EVANS, L.C. 2010 *Partial Differential Equations*, 2nd edn. American Mathematical Society.
- FLETCHER, R. 2000 *Practical Methods of Optimization*. John Wiley & Sons.
- FUKUSHIMA, E. 1999 Nuclear magnetic resonance as a tool to study flow. *Annu. Rev. Fluid Mech.* **31**, 95–123.
- FUNKE, S.W., NORDAAS, M., EVJU, Ø., ALNÆS, M.S. & MARDAL, K.A. 2019 Variational data assimilation for transient blood flow simulations: cerebral aneurysms as an illustrative example. *Intl J. Numer. Meth. Biomed. Engng* **35** (1), 1–27.
- GETREUER, P. 2012a Chan-veese segmentation. *Image Process. On Line* **2**, 214–224.
- GETREUER, P. 2012b Rudin–Osher–Fatemi total variation denoising using split Bregman. *Image Process. On Line* **2** (1), 74–95.
- GILLISSEN, J.J.J., VILQUIN, A., KELLAY, H., BOUFFANAIS, R. & YUE, D.K.P. 2018 A space-time integral minimisation method for the reconstruction of velocity fields from measured scalar fields. *J. Fluid Mech.* **854**, 348–366.
- GILLISSEN, J.J.J., BOUFFANAIS, R. & YUE, D.K.P. 2019 Data assimilation method to de-noise and de-filter particle image velocimetry data. *J. Fluid Mech.* **877**, 196–213.
- GUDBJARTSSON, H. & PATZ, S. 1995 The Rician distribution of noisy MRI data. *Magn. Reson. Med.* **34** (6), 910–914.
- HAIRER, M. 2009 An Introduction to Stochastic PDEs. Lecture Notes., [arXiv:0907.4178](https://arxiv.org/abs/0907.4178).
- HARRIS, C.R., *et al.* 2020 Array programming with NumPy. *Nature* **585** (7825), 357–362.
- HEISTER, T. & RAPIN, G. 2013 Efficient augmented Lagrangian-type preconditioning for the Oseen problem using Grad-Div stabilization. *Intl J. Numer. Meth. Fluids* **71** (1), 118–134.
- HOANG, V.H., LAW, K.J.H. & STUART, A.M. 2014 Determining white noise forcing from Eulerian observations in the Navier–Stokes equation. *Stoch. Partial Differ. Equ.* **2** (2), 233–261.
- KATRITSIS, D., KAIKTSIS, L., CHANIOTIS, A., PANTOS, J., EFSTATHOPOULOS, E.P. & MARMARELIS, V. 2007 Wall shear stress: theoretical considerations and methods of measurement. *Prog. Cardiovasc. Dis.* **49** (5), 307–329.

- KOLTUKLUOĞLU, T.S. 2019 Fourier Spectral Dynamic Data Assimilation: Interlacing CFD with 4D Flow MRI. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019* (ed. D. Shen, T. Liu, T.M. Peters, L.H. Staib, C. Essert, S. Zhou, P.-T. Yap & A. Khan), pp. 741–749. Springer International Publishing.
- KOLTUKLUOĞLU, T.S. & BLANCO, P.J. 2018 Boundary control in computational haemodynamics. *J. Fluid Mech.* **847**, 329–364.
- KOLTUKLUOĞLU, T.S., CVIJEIĆ, G. & HIPTMAIR, R. 2019 Harmonic balance techniques in cardiovascular fluid mechanics. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019* (ed. D. Shen, T. Liu, T.M. Peters, L.H. Staib, C. Essert, S. Zhou, P.-T. Yap & A. Khan), pp. 486–494. Springer International Publishing.
- LAM, S.K., PITROU, A. & SEIBERT, S. 2015 Numba: a LLVM-based Python JIT compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pp. 1–6.
- LUSTIG, M., DONOHO, D. & PAULY, J.M. 2007 Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.* **58** (6), 1182–1195.
- MANTLE, M.D. & SEDERMAN, A.J. 2003 Dynamic MRI in chemical process and reaction engineering. *Prog. Nucl. Magn. Reson. Spectrosc.* **43** (1), 3–60.
- MARKL, M., FRYDRYCHOWICZ, A., KOZERKE, S., HOPE, M. & WIEBEN, O. 2012 4D flow MRI. *J. Magn. Reson. Imag.* **36** (5), 1015–1036.
- MASSING, A., LARSON, M.G. & LOGG, A. 2013 Efficient implementation of finite element methods on nonmatching and overlapping meshes in three dimensions. *SIAM J. Sci. Comput.* **35** (1), C23–C47.
- MASSING, A., LARSON, M.G., LOGG, A. & ROGNES, M.E. 2014 A stabilized nitsche fictitious domain method for the Stokes problem. *J. Sci. Comput.* **61** (3), 604–628.
- MASSING, A., SCHOTT, B. & WALL, W.A. 2018 A stabilized Nitsche cut finite element method for the Oseen problem. *Comput. Meth. Appl. Mech. Engng* **328**, 262–300.
- MIRTICH, B. 1996 Fast and accurate computation of polyhedral mass properties. *J. Graph. Tools* **1** (2), 31–50.
- MONS, V., CHASSAING, J.C. & SAGAUT, P. 2017 Optimal sensor placement for variational data assimilation of unsteady flows past a rotationally oscillating cylinder. *J. Fluid Mech.* **823**, 230–277.
- MORRIS, P.D., *et al.* 2016 Computational fluid dynamics modelling in cardiovascular medicine. *Heart* **102** (1), 18–28.
- NITSCHKE, J. 1971 Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* **36** (1), 9–15.
- NOCEDAL, J. & WRIGHT, S.J. 2006 *Numerical optimization*. Springer.
- OKUTA, R., UNNO, Y., NISHINO, D., HIDO, S. & LOOMIS, C. 2017 CuPy: a NumPy-Compatible Library for NVIDIA GPU Calculations. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*.
- OSHER, S. & SETHIAN, J.A. 1988 Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. *J. Comput. Phys.* **79** (1), 12–49.
- OTSU, N. 1979 A threshold selection method from gray-level histograms. *IEEE Trans. Syst. Man Cybern.* **9** (1), 62–66.
- PEPER, E.S., GOTTWALD, L.M., ZHANG, Q., COOLEN, B.F., VAN OOIJ, P., NEDERVEEN, A.J. & STRIJKERS, G.J. 2020 Highly accelerated 4D flow cardiovascular magnetic resonance using a pseudo-spiral Cartesian acquisition and compressed sensing reconstruction for carotid flow and wall shear stress. *J. Cardiovasc. Magn. Reson.* **22**, 7.
- SAITO, K., ABE, S., KUMAMOTO, M., UCHIHARA, Y., TANAKA, A., SUGIE, K., IHARA, M., KOGA, M. & YAMAGAMI, H. 2020 Blood flow visualization and wall shear stress measurement of carotid arteries using vascular vector flow mapping. *Ultrasound Med. Biol.* **46** (10), 2692–2699.
- SANKARAN, S., KIM, H.J., CHOI, G. & TAYLOR, C.A. 2016 Uncertainty quantification in coronary blood flow simulations: impact of geometry, boundary conditions and blood viscosity. *J. Biomech.* **49** (12), 2540–2547.
- SETHIAN, J.A. 1996 A fast marching level set method for monotonically advancing fronts. *Proc. Natl Acad. Sci. USA* **93** (4), 1591–1595.
- SHARMA, A., RYPINA, I.I., MUSGRAVE, R. & HALLER, G. 2019 Analytic reconstruction of a two-dimensional velocity field from an observed diffusive scalar. *J. Fluid Mech.* **871**, 755–774.
- SOTELO, J., URBINA, J., VALVERDE, I., TEJOS, C., IRARRAZAVAL, P., ANDIA, M.E., URIBE, S. & HURTADO, D.E. 2016 3D quantification of wall shear stress and oscillatory shear index using a finite-element method in 3D CINE PC-MRI data of the thoracic aorta. *IEEE Trans. Med. Imag.* **35** (6), 1475–1487.

- STEJSKAL, E.O. & TANNER, J.E. 1965 Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *J. Chem. Phys.* **42** (1), 288–292.
- STUART, A.M. 2010 Inverse problems: A Bayesian perspective. *Acta Numer.* **19** (2010), 451–459.
- TARANTOLA, A. 2005 *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM.
- TEZDUYAR, T.E. 1991 Advances in Applied Mechanics. *Stabilized Finite Element Formulations for Incompressible Flow Computations* (ed. J.W. Hutchinson & T.Y. Wu), vol. 28. Elsevier.
- VAN DER WALT, S., SCHÖNBERGER, J.L., NUNEZ-IGLESIAS, J., BOULOGNE, F., WARNER, J.D., YAGER, N., GOUILLART, E. & YU, T. 2014 scikit-image: image processing in Python. *PeerJ*, 2:e453.
- VARADHAN, S.R.S. 1967*a* Diffusion processes in a small time interval. *Commun. Pure Appl. Maths* **20** (4), 659–685.
- VARADHAN, S.R.S. 1967*b* On the behavior of the fundamental solution of the heat equation with variable coefficients. *Commun. Pure Appl. Maths* **20** (2), 431–455.
- VERMA, S., PAPADIMITRIOU, C., LÜTHEN, N., ARAMPATZIS, G. & KOUMOUTSAKOS, P. 2019 Optimal sensor placement for artificial swimmers. *J. Fluid Mech.* **884**, A24.
- VIRTANEN, P., *et al.* 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Meth.* **17**, 261–272.
- VOLK, A. & KÄHLER, C.J. 2018 Density model for aqueous glycerol solutions. *Exp. Fluids* **59** (5), 75.
- WALKER, S.W. 2015 *The Shapes of Things: A Practical Guide to Differential Geometry and the Shape Derivative*. Advances in Design and Control, SIAM.
- WAPLER, M.C., LEUPOLD, J., DRAGONU, I., VON ELVERFELD, D., ZAITSEV, M. & WALLRABE, U. 2014 Magnetic properties of materials for MR engineering, micro-MR and beyond. *J. Magn. Reson.* **242**, 233–242.
- YU, H., JUNIPER, M.P. & MAGRI, L. 2019 Combined state and parameter estimation in level-set methods. *J. Comput. Phys.* **399**, 108950.