

Physics-informed compressed sensing for PC-MRI: an inverse Navier–Stokes problem

Alexandros Kontogiannis, Matthew P. Juniper

Abstract—We formulate a physics-informed compressed sensing (PICS) method for the reconstruction of velocity fields from noisy and sparse phase-contrast magnetic resonance signals. The method solves an inverse Navier–Stokes boundary value problem, which permits us to jointly reconstruct and segment the velocity field, and at the same time infer hidden quantities such as the hydrodynamic pressure and the wall shear stress. Using a Bayesian framework, we regularize the problem by introducing *a priori* information about the unknown parameters in the form of Gaussian random fields. This prior information is updated using the Navier–Stokes problem, an energy-based segmentation functional, and by requiring that the reconstruction is consistent with the k -space signals. We create an algorithm that solves this inverse problem, and test it for noisy and sparse k -space signals of the flow through a converging nozzle. We find that the method is capable of reconstructing and segmenting the velocity fields from sparsely-sampled (15% k -space coverage), low (~ 10) signal-to-noise ratio (SNR) signals, and that the reconstructed velocity field compares well with that derived from fully-sampled (100% k -space coverage) high (> 40) SNR signals of the same flow.

Index Terms—phase-contrast magnetic resonance imaging (PC-MRI), physics-informed compressed sensing, velocity reconstruction and segmentation

I. INTRODUCTION

IN phase-contrast magnetic resonance imaging one seeks to reconstruct complex images (w) whose magnitude is proportional to the nuclear spin density (ρ) and whose phase difference ($\Delta\varphi$) is proportional to a flow velocity component (u_i). For fully-sampled k -space signals (s^*) there is a one-to-one correspondence between the k -space and the physical (complex) space, which is given by a Fourier transform (\mathcal{F}) such that $s^* = \mathcal{F}w$. Then, for sufficiently high (> 3) signal-to-noise ratios (SNR), the noise in the magnitude and the phase of the complex image can be assumed to be white and Gaussian [1]. In this case, it is reasonable to directly reconstruct (denoise) the complex image in physical space, using either general purpose image denoising algorithms [2], [3], or physics-informed algorithms [4]–[9]. For sparsely-sampled k -space signals (s^*) there is no longer a one-to-one correspondence between k -space and physical space. A naive solution to this problem is to zero-fill the subsampled k -space signal and then perform an inverse Fourier transform. This approach, known as the ‘zero-filling solution’, reconstructs a complex image that is corrupted by artefacts and interference noise. Incoherent subsampling can eliminate artefacts, but interference noise remains a problem. In this case, it is better to reconstruct the complex image using compressed sensing (CS) [10]–[12].

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

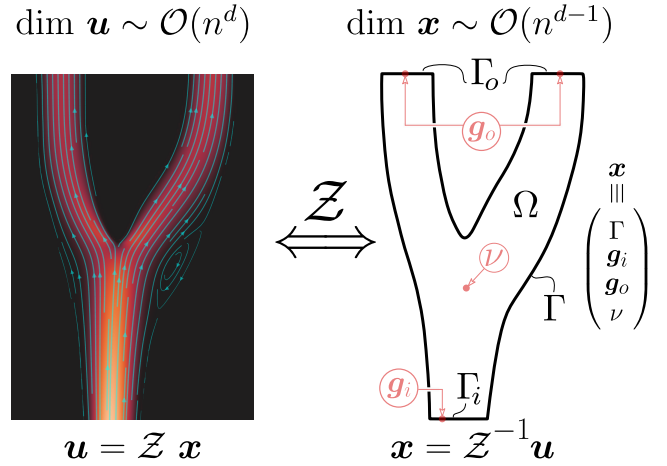


Fig. 1. A d -dimensional velocity field (u) that can be described by a Navier–Stokes problem (Z) has an underlying $(d-1)$ -dimensional structure ($x = Z^{-1}u$), which is the parameter vector containing the shape of the object (Γ), the boundary conditions (g_i, g_o), and the kinematic viscosity (ν). Images of n^d voxels depicting the d velocity components can be compressed/decompressed by solving an inverse/forward N–S problem.

Compressed sensing relies on the assumption that the complex image, or its constituent parts (e.g. the magnitude and the phase), have a sparse representation in a transform domain, and that the transform domain is considered to be known *a priori*. The reconstructed complex image can then be recovered by solving a (generally) nonlinear optimization problem of the form [13]

$$\min_w \|Tw\|_{\ell_1} \quad \text{subject to} \quad \|s^* - \mathcal{P}Fw\|_{\ell_2} < \sqrt{N_s}\varepsilon \quad ,$$

where T is a sparsity-inducing transformation, \mathcal{P} is the sparse sampling pattern, N_s is the total number of sampled k -space points, and $\varepsilon \propto \sigma$ is a user-selected value proportional to the standard deviation of noise (σ). In the above example, the functional to be minimized acts as a regularizer while the constraint enforces consistency between the reconstruction and the measurement. It is also possible to decompose the complex image into real and imaginary [14], [15], or into magnitude and phase components [16] [17, Chapter 9.3], and reconstruct them separately. With this decomposition, different regularizers can be used for each component. This is particularly important for PC-MRI because suitable regularizers for the reconstructed magnitude, e.g. total variation (TV), may produce large errors or artefacts (e.g. staircasing) in the reconstructed phase. This is because the magnitude image often has high-contrast features resembling a piecewise-constant function, which is sparse in the TV-transform domain [12], [13], while the phase usually

has smooth features since it encodes velocity information, and thus is no longer sparse in the same domain. A suitable transform for the phases, and for piecewise-smooth functions in general, is second-order total generalized variation (TGV²) [13], [18]. Complications arise due to the fact that the phase of the complex image alone does not represent velocity, but the phase difference does. Therefore, if there is available *a priori* information on the velocity, the regularization should be imposed on the phase difference rather than on the individual phases [17, Chapter 9.3] [19], [20].

Even though generic CS methods perform very well in magnitude reconstruction, accurate velocity reconstruction from subsampled k -space signals remains a challenge. It has been suggested [17, Chapter 11.1] that even sparser PC-MRI signals could be reconstructed if a regularization method based on the Navier–Stokes (N–S) equations is used. However, the effect of the boundary conditions should be equally important in order to capture the velocity profiles in both the lumen and the near-wall region, and to estimate the wall-shear stress¹ with greater confidence. Based on these observations, we believe that a general and accurate way of injecting *a priori* knowledge for velocity (phase difference) regularization is in the form of a N–S *boundary value problem*, and not in the form of the N–S *equations* alone.

Several physics-informed velocity regularization methods have been proposed in the past [4]–[8], [15], [21], [22], but none of them exploits the full structure of a Bayesian inverse Navier–Stokes problem in which the domain boundary ($\partial\Omega$), the boundary conditions ($\mathbf{g}_i, \mathbf{g}_o$), and the kinematic viscosity (ν) are all considered unknown. A similar approach is discussed in [9], but the method applies only to fully-sampled PC-MRI signals and unwrapped velocity images.

Solving an inverse N–S problem (\mathcal{Z}) amounts to finding its unknown parameters (\mathbf{x}), which produce a modelled velocity ($\mathbf{u} \equiv \mathcal{Z}\mathbf{x}$) that approximates the measured velocity ($\mathbf{u}^* \propto \Delta\varphi$) in an appropriate norm. In this way, not only do we obtain a regularized, noiseless (SNR = ∞) velocity field, but we also infer hidden flow-related quantities such as the hydrodynamic pressure (p) and the wall shear stresses [23], [24], which cannot be measured using conventional MRI or PC-MRI methods, but which naturally arise from the N–S problem. It is also important to note that, because we jointly reconstruct and segment the velocity field, geometric errors are minimized [25], [26].

Our approach to physics-informed CS extends the standard notion of sparsity used in conventional CS methods to a more general notion of a structure [27], which is dictated by the N–S problem. Instead of enforcing sparsity during the nonlinear optimization process, we recover a sparse (hidden) structure of the velocity field by enforcing the Navier–Stokes problem as a constraint. This is better explained in figure 1, which shows that the velocity exhibits an underlying low-dimensional structure that is encoded in the unknown parameters \mathbf{x} . Based on this, a velocity field can be compressed/decompressed by solving an inverse/forward Navier–Stokes problem. In contrast to our approach, a conventional sparsity-promoting and physics-

informed velocity regularization method is described in [28], but it relies on a pre-existing database (library) containing the dominant eigenmodes of a Navier–Stokes problem that is defined in a pre-set geometry, with pre-set inlet boundary condition.

Unlike pure machine learning and library-based algorithms [29], [30], which learn to recognize features in training data (often generated using computational fluid dynamics), our method encapsulates fluid mechanics knowledge in the form of the N–S problem. It can therefore reconstruct flows that it has not yet seen and extrapolate to new flow conditions, enabling patient-specific cardiovascular modelling. While neural networks (NNs) have revolutionized the field of computer vision, some fundamental problems still need to be addressed (e.g. AI-generated hallucinations, and existence of computational algorithms) [31], [32]. Our method, on the other hand, is formulated in a variational framework and is therefore amenable to mathematical analysis. In the future this could provide a reference for the formulation of a rigorous NN-based algorithm that approximates the N–S solution and other partial differential equations using learned operators [33], [34]. To the best of our knowledge, there is currently neither a NN that can solve the N–S problem more efficiently than computational fluid dynamics methods, nor a NN that is proven to approximate the N–S boundary value problem operator.

In this paper, we build upon the work of [9] in order to formulate a physics-informed compressed sensing (PICS) method for the joint reconstruction and segmentation of velocity fields from sparse and noisy k -space signals. We provide an algorithm that solves the reconstruction problem and demonstrate it on k -space signals of a steady axisymmetric flow. Because the acquired signals were originally fully-sampled, we sparsify them using two different sparse sampling patterns and study the effect of sampling density on the velocity reconstruction error. In section II we formulate the physics-informed compressed sensing method, and an algorithm that implements it. In section III we test the method on sparse and noisy k -space signals.

II. PHYSICS-INFORMED COMPRESSED SENSING AS AN INVERSE NAVIER–STOKES PROBLEM

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 $H^k(\Omega)$ denotes the space of square-integrable functions with k square-integrable derivatives. For a given covariance operator, \mathcal{C} , we also define the covariance-weighted $L^2(\Omega)$ spaces, endowed with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{C}(\Omega)} := \langle \mathcal{C}^{-1/2} \cdot, \mathcal{C}^{-1/2} \cdot \rangle_{\Omega}$, which generates the norm $\|\cdot\|_{\mathcal{C}(\Omega)}$. The Euclidean norm in the space of real numbers \mathbb{R}^n is denoted by $|\cdot|_{\mathbb{R}^n}$, and the measure (volume) of the domain Ω by $|\Omega|$. The first variation of a functional $\mathcal{J} : L^2 \rightarrow \mathbb{R}$ with respect to an unknown $z \in L^2$ is defined by

$$\delta_z \mathcal{J} \equiv \left. \frac{d}{d\tau} \mathcal{J}(z + \tau z', \dots) \right|_{\tau=0} = \langle D_z \mathcal{J}, z' \rangle \quad , \quad (1)$$

where $\tau \in \mathbb{R}$, $z' \in L^2$ is an allowed perturbation of z , and $D_z \mathcal{J}$ is the generalized gradient. If z is defined on

¹Wall shear stress is often sought-after in fluid mechanics applications.

a covariance-weighted L^2 space, we furthermore define the steepest ascent direction $\widehat{D}_z \mathcal{J}$ such that

$$\langle \widehat{D}_z \mathcal{J}, z' \rangle_{\mathcal{C}} = \langle \mathcal{C}^{-1} \widehat{D}_z \mathcal{J}, z' \rangle = \langle D_z \mathcal{J}, z' \rangle \quad , \quad (2)$$

therefore $\widehat{D}_z \mathcal{J} = \mathcal{C} D_z \mathcal{J}$. When the covariance takes the form $\mathcal{C} = \sigma^2 \mathbf{I}$, where $\sigma^2 \in \mathbb{R}$ is the variance and \mathbf{I} is the identity operator, we write $\sigma^{-2} \|\cdot\|_{L^2}$ instead of $\|\cdot\|_{\mathcal{C}}$ for simplicity. We use the superscript $(\cdot)^\dagger$ to denote the adjoint of an operator, $(\cdot)^*$ to denote a measurement, $(\cdot)^\circ$ to denote a reconstruction, and $(\cdot)^\bullet$ to denote the ground truth. Note that, for the velocity, \mathbf{u}^* denotes the velocity obtained from the phase differences, and \mathbf{u} denotes the velocity obtained from the Navier–Stokes problem.

A. Phase-contrast magnetic resonance imaging

Phase-contrast magnetic resonance imaging can measure d -dimensional velocity fields \mathbf{u} inside or around an object Ω , by requiring a minimum of d sets of k -space signals $\{s_j\}_{j=1}^d$, one for each velocity component. The velocity component u_i can be recovered from the signal set $s_i = \{s_j\}_{j=1}^4$ by computing the phase difference

$$u_i = c_i(\varphi_1 - \varphi_2 - \varphi_3 + \varphi_4) \equiv c_i(\Delta\varphi)_i \quad , \quad (3)$$

where c_i is a known constant that depends on the gyromagnetic ratio of hydrogen and the gradient pulse properties, and $\varphi_j \equiv \arg(\mathcal{F}^{-1}s_j)$ when s_j is a fully-sampled k -space signal. Note that the last two phases in the above phase difference correspond to zero-flow experiments in order to remove any phase shift contributions that are not caused by the flow.

For sparsely-sampled k -space signals we define the sparse sampling operator $\mathcal{P} : I \rightarrow I_s$, which projects from the full image space $I \subset \mathbb{C}^d$ to the sparse image space $I_s \subset I$. Then, if s_j is sparse, the corresponding complex image is given by

$$w_j \equiv \rho_j e^{i\varphi_j} = \mathcal{F}^{-1} \mathcal{P}^{-1} s_j \quad . \quad (4)$$

Since the backprojection \mathcal{P}^{-1} is ill-posed, we define $\mathcal{P}^{-1} s_j$ as the zero-filled k -space signal in I . We also obtain the zero-filling solution for the velocity by using equation (3) with $\varphi_j = \arg(w_j)$. As was mentioned in section I, the zero-filling velocity solution is corrupted by interference noise and artefacts that strongly depend on the sparse sampling pattern \mathcal{P} . We therefore need to reconstruct the complex images w_j using a different approach.

B. Physics-informed compressed sensing (PICS) formulation

1) *Phase regularization*: A priori knowledge of the velocity, i.e. the phase difference, comes in the form of a Navier–Stokes boundary value problem. Assuming steady, incompressible flow and a Newtonian fluid, the fluid dynamics are governed by the following problem (see figure 1)

$$\left\{ \begin{array}{ll} \mathbf{u} \cdot \nabla \mathbf{u} - \nu \Delta \mathbf{u} + \nabla p = \mathbf{0} & \text{in } \Omega \\ \nabla \cdot \mathbf{u} = 0 & \text{in } \Omega \\ \mathbf{u} = \mathbf{0} & \text{on } \Gamma \\ \mathbf{u} = \mathbf{g}_i & \text{on } \Gamma_i \\ -\nu \partial_\nu \mathbf{u} + p \boldsymbol{\nu} = \mathbf{g}_o & \text{on } \Gamma_o \end{array} \right. , \quad (5)$$

where \mathbf{u} is the velocity, $p \leftarrow p/\rho$ is the reduced hydrodynamic pressure², 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 , $\boldsymbol{\nu}$ is the unit normal vector on $\partial\Omega$, and $\partial_\nu \equiv \boldsymbol{\nu} \cdot \nabla$ is the normal derivative. A zero-velocity (no-slip) boundary condition is imposed on the walls Γ of the object Ω . We furthermore define the projection operator $\mathcal{S} : M \rightarrow I$ that projects from the model space M to the image space I . The model space M is a fixed space that always contains the deformable domain Ω ($\Omega \subset M$), such that $\Gamma_i, \Gamma_o \subset \partial M$, and

$$\partial M \cap \text{Int } \Gamma = \emptyset, \quad \partial M \cap \Gamma_i = \partial M_i, \quad \partial M \cap \Gamma_o = \partial M_o, \quad (6)$$

where Int denotes the interior of a set, i.e. $\text{Int } \Gamma \equiv (\Gamma - \partial\Gamma)$. The domain Ω does not have to be path-connected. In general, it can be defined as the union of disjoint, path-connected sets Ω_j , i.e. $\Omega = \Omega_1 \cup \Omega_2 \cup \dots$, provided that every Ω_j includes at least one inlet and one outlet.

The discrepancy between the measured (\mathbf{u}^*) and the modelled (\mathbf{u}) velocity fields is given by

$$\mathcal{U}(\boldsymbol{\varphi}, \mathbf{u}) \equiv \frac{1}{2} \sum_{i=1}^d \|u_i^* - \mathcal{S}u_i\|_{\mathcal{C}_{u_i}(I_w)}^2 \quad , \quad (7)$$

where $\boldsymbol{\varphi}$ contains the phases needed to compute $\mathbf{u}^* = \mathbf{u}^*(\boldsymbol{\varphi})$ using (3), $I_w \subseteq I$ is a user-selected area (window) of interest, and \mathcal{C}_{u_i} is the covariance operator for the velocity discrepancy of the i -th component. In [9] we had assumed a diagonal covariance operator for the velocity discrepancy because the noise in the phase images can be assumed to be white and Gaussian for fully-sampled k -space signals with $\text{SNR} > 3$ [1]. Here, we model the interference (correlated) noise using an exponential covariance operator such that

$$\mathcal{C}_{u_i} g \equiv \sigma_{u_i}^2 (C * g) \quad \text{for any } g \in L^2(I) \quad , \quad (8)$$

where σ_{u_i} is the standard deviation of noise in the i -th velocity component, $*$ denotes convolution, and C is the exponential kernel

$$C(r) = \frac{1}{d! V_d(\ell)} e^{-|r|/\ell} \quad \text{for any } r \in \mathbb{R}^d \quad , \quad (9)$$

where d is the dimension of the velocity field \mathbf{u} , ℓ is the characteristic length, $V_d(\ell) = \pi^{d/2} \ell^d / \Gamma(d/2+1)$ is the volume of the d -dimensional Euclidean ball of radius ℓ , and Γ is the gamma function. Adopting a Bayesian inference setting similar to [9], we introduce a 2π -periodic prior for the phase [16]. The combined phase regularization functional is then given by

$$\mathcal{R}_\varphi(\boldsymbol{\varphi}, \mathbf{u}) \equiv \mathcal{U} + \frac{1}{2} \sum_{j=1}^{4d} \|e^{i\varphi_j} - \bar{e}^{i\bar{\varphi}_j}\|_{\mathcal{C}_{\varphi_j}(I)}^2 \quad , \quad (10)$$

where $\bar{\varphi}_j$ is the prior assumption of φ_j , $\mathcal{C}_{\varphi_j} = \tilde{\sigma}_{\varphi_j}^2 (C * \cdot)$, $\tilde{\sigma}_{\varphi_j} = \xi_\varphi \sigma_{\varphi_j}$, and $\xi_\varphi \in \mathbb{R}$ is a user-selected parameter that determines the confidence in the prior assumption.

²From now on we refer to p simply as the pressure.

2) *Navier–Stokes parameters regularization*: Additional regularization is required to penalize improbable solutions of the Navier–Stokes parameters \mathbf{x} . The parameters \mathbf{x} consist of the shape of the domain $\partial\Omega$, the Dirichlet boundary condition at the inlet(s) \mathbf{g}_i , the natural boundary condition at the outlet(s) \mathbf{g}_o , and the kinematic viscosity ν . Note that we identify the object Ω and its boundary with a signed distance function $\phi_{\pm} \in L^2(M)$ such that [9, Section 2.4]

$$\Omega = \{y \in \Omega : \phi_{\pm}(y) < 0\}, \quad \partial\Omega = \{y \in \Omega : \phi_{\pm}(y) = 0\}.$$

The regularization functional for \mathbf{x} is given by [9, Section 2.2]

$$\begin{aligned} \mathcal{R}_x(\mathbf{x}) = & \frac{1}{2\sigma_{\phi_{\pm}}^2} \|\bar{\phi}_{\pm} - \phi_{\pm}\|_{L^2(M)}^2 + \frac{1}{2} \|\mathbf{g}_i - \bar{\mathbf{g}}_i\|_{C_{\mathbf{g}_i}(\partial M_i)}^2 \\ & + \frac{1}{2} \|\mathbf{g}_o - \bar{\mathbf{g}}_o\|_{C_{\mathbf{g}_o}(\partial M_o)}^2 + \frac{1}{2\sigma_{\nu}^2} |\nu - \bar{\nu}|_{\mathbb{R}}^2, \quad (11) \end{aligned}$$

where $(\bar{\cdot})$ is a prior assumption, $\sigma_{(\cdot)}$ is the standard deviation of the prior, and $C_{\mathbf{g}_i}, C_{\mathbf{g}_o}$ are covariance operators that are based on the exponential covariance kernel (9), but with additional boundary conditions, which are described in [9].

3) *Magnitude segmentation*: To exploit information about $\partial\Omega$ from the nuclear spin density (magnitude) images we use an energy-based segmentation functional \mathcal{S} , which assumes that the image consists of two regions with approximately uniform magnitudes (e.g. stars in a night sky) [35], [36]. The functional is given by

$$\begin{aligned} \mathcal{S}(\boldsymbol{\rho}, \phi_{\pm}, \alpha, \beta) \equiv & \frac{1}{8d} \sum_{j=1}^{4d} \frac{1}{\sigma_{\rho_j}^2} \left(\|(\rho_j - \alpha) \mathcal{H}(\phi_{\pm})\|_{L^2(I_w)}^2 \right. \\ & \left. + \|(\rho_j - \beta) (\mathcal{H}(\phi_{\pm}) - 1)\|_{L^2(I_w)}^2 \right), \quad (12) \end{aligned}$$

where the mean value of the average magnitude inside Ω is $\alpha \in \mathbb{R}$ and outside Ω is $\beta \in \mathbb{R}$. In addition, σ_{ρ_j} is the standard deviation of noise in the magnitude image, and \mathcal{H} denotes the (modified) Heaviside function such that

$$\mathcal{H}(t) = 1 \quad \text{if } t < 0 \quad \text{else } 0. \quad (13)$$

Taking into account the prior contribution, the combined functional is given by

$$\mathcal{R}_{\rho}(\boldsymbol{\rho}, \phi_{\pm}, \alpha, \beta) \equiv \mathcal{S} + \frac{1}{2} \sum_{j=1}^{4d} \|\rho_j - \bar{\rho}_j\|_{C_{\rho_j}(I)}^2, \quad (14)$$

where $C_{\rho_j} = \tilde{\sigma}_{\rho_j}^2 (C * \cdot)$, $\tilde{\sigma}_{\rho_j} = \xi_{\rho} \sigma_{\rho_j}$, and $\xi_{\rho} \in \mathbb{R}$ is a user-selected parameter that determines the confidence in the prior assumption.

4) *Consistency with measured k -space signal*: We consider the noise in the sparse k -space signals s_j^* to be white and Gaussian, with zero mean and standard deviation σ_j . The discrepancy between the measured sparse signals s_j^* and a signal $s_j \equiv \mathcal{F}(\rho_j e^{i\varphi_j})$ is thus given by

$$\mathcal{E}(\boldsymbol{\varphi}, \boldsymbol{\rho}) \equiv \sum_{j=1}^{4d} \frac{1}{2\sigma_j^2} \|s_j^* - \mathcal{P}\mathcal{F}(\rho_j e^{i\varphi_j})\|_{L^2(I_s)}^2, \quad (15)$$

with $\rho_j e^{i\varphi_j} \equiv w_j$ being the complex image.

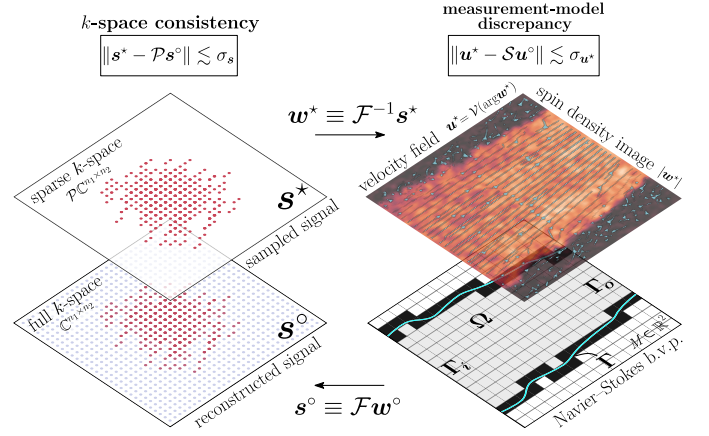


Fig. 2. A set of sparse k -space signals (s^*) can be zero-filled and Fourier-transformed to the physical space in order to produce a set of complex images (w^*), which provide a first estimate for the measured velocity field (w^*) and the ^1H spin density inside an object. A Navier–Stokes problem is then used to reconstruct the measured velocity field and the boundary of the object. The reconstructed complex image (w^o) is then assembled and transformed back to k -space (s^o), where it can be compared with the original sparse signal.

5) *The inverse problem*: We now collect the above results to formulate an inverse problem where the reconstructed phases φ^o , N–S parameters \mathbf{x}^o , magnitudes $\boldsymbol{\rho}^o$, and segmentation constants $\mathbf{c}^o = (\alpha^o, \beta^o)$, are the minimizers of the following nonlinearly-constrained optimization problem

$$(\varphi^o, \mathbf{x}^o, \boldsymbol{\rho}^o, \mathbf{c}^o) = \underset{\boldsymbol{\varphi}, \mathbf{x}, \boldsymbol{\rho}, \mathbf{c}}{\operatorname{argmin}} (\mathcal{R}_{\varphi} + \mathcal{R}_x + \mathcal{R}_{\rho}) \quad (16)$$

$$\begin{aligned} \text{subject to } & \text{i) } \|s_j^* - \mathcal{P}\mathcal{F}(\rho_j e^{i\varphi_j})\|_{L^2(I_s)} \leq \sqrt{|I_s|} \varepsilon_j \text{ for all } j \\ & \text{ii) } \mathbf{u} = \mathcal{Z} \mathbf{x}, \end{aligned}$$

where $\varepsilon_j \propto \sigma_j$ is a user-selected value proportional to the standard deviation of noise. We observe that the N–S problem (5), encoded in the operator \mathcal{Z} , acts as a nonlinear equality constraint in order to ensure that \mathbf{u} is always a N–S solution, which is uniquely defined by the parameters \mathbf{x} ($\mathbf{u}^o = \mathcal{Z}\mathbf{x}^o$).

6) *Relaxing the constraints*: We recast problem (16) into a variational form by defining the Lagrangian

$$\mathcal{J} \equiv \mathcal{R}_{\varphi} + \mathcal{R}_x + \mathcal{R}_{\rho} + \mathcal{E} + \mathcal{M}, \quad (17)$$

where \mathcal{M} is the weak formulation of the N–S problem (5), for test functions $(v, q) \in \mathbf{H}^1(\Omega) \times L^2(\Omega)$, given by

$$\begin{aligned} \mathcal{M}(\mathbf{u}, p, v, q; \mathbf{x}) \equiv & \int_{\Omega} (\mathbf{v} \cdot (\mathbf{u} \cdot \nabla \mathbf{u}) + \nu \nabla v : \nabla \mathbf{u} - (\nabla \cdot \mathbf{v})p \\ & - q(\nabla \cdot \mathbf{u})) + \int_{\Gamma_o} \mathbf{v} \cdot \mathbf{g}_o + \int_{\Gamma \cup \Gamma_i} \mathbf{v} \cdot (-\nu \partial_{\nu} \mathbf{u} + p \boldsymbol{\nu}) \\ & + \mathcal{N}_{\Gamma_i}(\mathbf{v}, q, \mathbf{u}; \mathbf{g}_i) + \mathcal{N}_{\Gamma}(\mathbf{v}, q, \mathbf{u}; \mathbf{0}), \quad (18) \end{aligned}$$

and \mathcal{N} is the Nitsche penalty term [37]

$$\mathcal{N}_{\Gamma}(\mathbf{v}, q, \mathbf{u}; \mathbf{g}) \equiv \int_{\Gamma} (-\nu \partial_{\nu} \mathbf{v} + q \boldsymbol{\nu} + \eta \mathbf{v}) \cdot (\mathbf{u} - \mathbf{g}), \quad (19)$$

which we use to weakly impose a Dirichlet boundary condition $\mathbf{g} \in \mathbf{L}^2(\Gamma)$ on a boundary Γ , for fixed penalty

$\eta \in \mathbb{R}$ [9, Section 2.7].

7) *Euler–Lagrange equations*: Instead of solving (16) as is, it is preferable to search for critical points of (17) by solving the nonlinear Euler–Lagrange (E–L) system for the unknowns

$$v, q, \mathbf{u}, p, \varphi, \mathbf{x}, \boldsymbol{\rho}, \mathbf{c},$$

where v is the adjoint velocity and q is the adjoint pressure, acting as Lagrange multipliers in (17) via \mathcal{M} . We find that the E–L equations

$$\delta_v \mathcal{J} = \delta_v \mathcal{M} = \mathbf{0} \quad , \quad \delta_q \mathcal{J} = \delta_q \mathcal{M} = 0 \quad , \quad (20)$$

are in fact the N–S equations and the weakly-enforced boundary conditions, i.e. problem (5). The next set of equations, namely

$$\delta_{\mathbf{u}} \mathcal{J} = \delta_{\mathbf{u}} (\mathcal{U} + \mathcal{M}) = \mathbf{0} \quad , \quad \delta_p \mathcal{J} = \delta_p \mathcal{M} = 0 \quad , \quad (21)$$

gives rise to the adjoint N–S problem [9, Section 2.3.1]

$$\begin{cases} -\mathbf{u} \cdot (\nabla \mathbf{v} + (\nabla \mathbf{v})^\dagger) - \nu \Delta \mathbf{v} + \nabla q = -D_{\mathbf{u}} \mathcal{J} & \text{in } \Omega \\ \nabla \cdot \mathbf{v} = 0 & \text{in } \Omega \\ \mathbf{v} = \mathbf{0} & \text{on } \Gamma \cup \Gamma_i \\ (\mathbf{u} \cdot \boldsymbol{\nu}) \mathbf{v} + (\mathbf{u} \cdot \mathbf{v}) \boldsymbol{\nu} + \nu \partial_\nu \mathbf{v} - q \boldsymbol{\nu} = 0 & \text{on } \Gamma_o \end{cases} \quad (22)$$

Note that

$$D_{\mathbf{u}} \mathcal{J} = D_{\mathbf{u}} \mathcal{U} = -S^\dagger C_{\mathbf{u}}^{-1} (\mathbf{u}^* - S \mathbf{u}) \quad , \quad (23)$$

and that both v and q vanish when $D_{\mathbf{u}} \mathcal{U} \equiv \mathbf{0}$ in I_w , i.e. when the measured velocity is identical to the modelled velocity in the user-selected area of interest I_w . The rest of the equations comprising the E–L system are

$$\delta_\varphi \mathcal{J} = \delta_\varphi (\mathcal{R}_\varphi + \mathcal{E}) = \mathbf{0} \quad (24)$$

$$\delta_{\mathbf{x}} \mathcal{J} = \delta_{\mathbf{x}} (\mathcal{R}_{\mathbf{x}} + \mathcal{S} + \mathcal{M}) = \mathbf{0} \quad (25)$$

$$\delta_{\boldsymbol{\rho}} \mathcal{J} = \delta_{\boldsymbol{\rho}} (\mathcal{R}_{\boldsymbol{\rho}} + \mathcal{E}) = \mathbf{0} \quad (26)$$

$$\delta_{\mathbf{c}} \mathcal{J} = \delta_{\mathbf{c}} \mathcal{S} = \mathbf{0} \quad , \quad (27)$$

which we expand in Appendix A to obtain an explicit formula for the steepest ascent direction of each unknown. We then use these directions in order to find a critical point of \mathcal{J} using an optimization algorithm.

8) *Solving the E–L system*: For large scale problems such as (16) it is prohibitively expensive to solve the full, implicit E–L system. Instead, we devise a segregated method that solves this nonlinear and nonconvex system, which is briefly explained in algorithm 1. Each iteration in the main loop of algorithm 1 consists of four main stages:

- 1) First, since $\mathcal{J} = \mathcal{J}(\varphi, \mathbf{x}, \boldsymbol{\rho}, \mathbf{c})$, we fix the phases φ , the magnitudes $\boldsymbol{\rho}$, and the segmentation constants \mathbf{c} , and update the unknown N–S parameters \mathbf{x} by solving one iteration of the inverse N–S problem. This stage is actually identical to the algorithm presented in [9], which was devised for the reconstruction of noisy velocity images (fully-sampled PC-MRI signals), where $\varphi, \boldsymbol{\rho}$ are always constant, and the segmentation functional \mathcal{S} is absent, i.e. \mathbf{c} is not involved. The updated N–S

parameters \mathbf{x}_{k+1} are then used to update the modelled velocity field to $\mathbf{u}_{k+1} = \mathcal{Z}^{\mathbf{u}_k} \mathbf{x}_{k+1}$, where $\mathcal{Z}^{\mathbf{u}_k}$ denotes the Oseen linearization of the N–S problem (5) around \mathbf{u}_k . The Oseen linearization around \mathbf{u}_k replaces the nonlinear convective term $\mathbf{u} \cdot \nabla \mathbf{u}$ in (5) with the linear convective term $\mathbf{u}_k \cdot \nabla \mathbf{u}$.

- 2) Next, we fix $\mathbf{x}, \boldsymbol{\rho}$ and \mathbf{c} and reconstruct the phases, drawing information from the modelled velocity field \mathbf{u}_{k+1} , *a priori* phase information, and the k -space signals.
- 3) We then fix φ, \mathbf{x} and $\boldsymbol{\rho}$ in order to compute the segmentation constants α, β , given by the explicit formulas (47) and (48).
- 4) Finally, we fix φ, \mathbf{x} and \mathbf{c} in order to reconstruct the magnitudes, drawing information from the energy-based segmentation functional, *a priori* magnitude information, and the k -space signals.

During the first stage of each iteration, we further update an inverse Hessian approximation of the unknown parameters \mathbf{x} using the damped BFGS quasi-Newton method [38], [39], and this allows us to estimate the uncertainty in the predicted shape, the boundary conditions, and the kinematic viscosity [9, Section 2.6]. Algorithm 1 terminates if either the reconstructed velocity, phases, and magnitudes are consistent with the data and the update for the unknowns $\varphi, \mathbf{x}, \boldsymbol{\rho}, \mathbf{c}$ is below the user-specified tolerance, or the line search of every individual stage fails to reduce \mathcal{J} further. The reconstructed velocity components are consistent with the data when the discrepancy between the velocity component and the respective phase difference, \mathcal{U}_j , i.e. the summands of equation (7), satisfy $\sqrt{2\mathcal{U}_j/|I_s|} < \varepsilon_j$ for every j , where $\varepsilon_j \lesssim 1$. The reconstructed phases and magnitudes are consistent with the data when the k -space consistency norms of the scans, \mathcal{E}_j , i.e. the summands of equation (15), satisfy $\sqrt{2\mathcal{E}_j/|I_s|} < \varepsilon_j$ for every j . Every line search starts with a step size of $\tau = 1$, and the step size is halved until $\mathcal{J}_{k+1} < \mathcal{J}_k$. It is often the case that the search of the N–S parameters (first stage of algorithm 1) converges faster than the reconstruction of the phases and the magnitudes. In this case, the inverse N–S problem within the loop does not need to be solved further, and the reconstructed modelled velocity field (given by the converged N–S parameters) is used to reconstruct the phases and the magnitudes until all stages of the algorithm have converged.

The reconstruction problem (16) that algorithm 1 solves is nonlinear and nonconvex because it involves a moving domain in which a N–S problem is solved, a k -space signal that is decomposed into phase and magnitude components, and a 2π -periodic prior for the phases. Ill-posedness is largely mitigated using a Bayesian regularization framework. The well-posedness of Bayesian inverse N–S problems is addressed in [40]. Problem (16), however, is more complicated than an inverse N–S problem alone. Nevertheless, for the test case that we have studied here, we observe that algorithm 1 successfully recovers the true solution, as it is shown in table II.

9) *Numerics*: To solve boundary value problems, problem (5) and (22), for example, we use an immersed boundary finite element method, the details of which can be found in [9, Section 2.7]. We implement algorithm 1, and all the necessary numerical methods, in Python modules, using Python’s

Algorithm 1: PICS for sparse PC-MRI signals.

Input: sparse PC-MRI signals \mathbf{s}^* , noise variances, initial guesses (priors) for the unknowns $\bar{\varphi}$, $\bar{\mathbf{x}}$, $\bar{\rho}$

begin

$k \leftarrow 0$

Initialization

$\mathbf{u}_0^* \leftarrow$ zero-filled velocity from $\bar{\varphi}$ - eq. (3)

$\bar{\phi}_\pm \leftarrow$ signed distance field from $\sum_j \bar{\rho}_j$ [9]

$(\mathbf{u}, p)_0 \leftarrow \mathcal{Z} \bar{\mathbf{x}}$ - N-S problem (5)

while convergence_criterion_is_not_met

do

1) **inverse N-S - min** $\mathcal{J}(\varphi_k, \mathbf{x}, \rho_k, c_k)$:

$(\mathbf{v}, q)_k \leftarrow$ adjoint N-S problem (22)

$\hat{D}_x \mathcal{J} \leftarrow$ s. a. directions (49)-(55)

$\tilde{H}_x^k, \tau \leftarrow$ approx. inv. Hessian and l. s. [9]

$\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k - \tau \tilde{H}_x^k \hat{D}_x \mathcal{J}$

$(\mathbf{u}, p)_{k+1} \leftarrow \mathcal{Z}^{\mathbf{u}_k} \mathbf{x}_{k+1}$ - lin. N-S problem (5)

2) **phase - min** $\mathcal{J}(\varphi, \mathbf{x}_{k+1}, \rho_k, c_k)$:

$\hat{D}_\varphi \mathcal{J}, \tau \leftarrow$ s. a. directions (43) and l. s.

$\varphi_{k+1} \leftarrow \varphi_k - \tau \hat{D}_\varphi \mathcal{J}$

$\mathbf{u}_{k+1}^* \leftarrow$ measured vel. from φ_{k+1} - eq. (3)

3) **segmentation - min** $\mathcal{J}(\varphi_{k+1}, \mathbf{x}_{k+1}, \rho_k, c)$:

$c_{k+1} \leftarrow$ compute $(\alpha, \beta)_{k+1}$ - eq. (47),(48)

4) **magnitude - min** $\mathcal{J}(\varphi_{k+1}, \mathbf{x}_{k+1}, \rho, c_{k+1})$:

$\hat{D}_\rho \mathcal{J}, \tau \leftarrow$ s. a. directions (44) and l. s.

$\rho_{k+1} \leftarrow \rho_k - \tau \hat{D}_\rho \mathcal{J}$

$k \leftarrow k + 1$

Output:

$\mathbf{x}^\circ \leftarrow \mathbf{x}_k$ (inferred N-S parameters)

$(\mathbf{u}^\circ, p^\circ) \leftarrow (\mathbf{u}, p)_k$ (N-S velocity and pressure)

$(\varphi^\circ, \rho^\circ) \leftarrow (\varphi, \rho)_k$ (reconstructed phases/mag.)

Optional output:

$\gamma_w \leftarrow$ wall shear rate from $\mathbf{u}^\circ, \mathbf{x}^\circ$

s. a.: steepest ascent, l. s.: line search, lin.: linearized

standard libraries for scientific computing: NumPy [41] and SciPy [42].

III. SPARSE PC-MRI RECONSTRUCTION USING PICS

We now apply algorithm 1 to sparse and noisy (low SNR) PC-MRI signals of an axisymmetric flow of water/glycerol through a converging nozzle [9]. The description of the PC-MRI experiment can be found in Appendix B. Since the k -space was fully-sampled, we sparsify the k -space signals using two different sparse sampling patterns \mathcal{P} : i) a sampling pattern \mathcal{P}_\odot , created from a two-dimensional normal distribution, and ii) a sampling pattern \mathcal{P}_\parallel , created from a one-dimensional normal distribution of lines. The sampling pattern \mathcal{P}_\parallel is commonly used in MRI because conventional acquisition protocols sample k -space lines for each system excitation. For these sparse sampling patterns, we further investigate the effect of subsampling on the velocity reconstruction error. The




Fig. 3. Ground truth for the axial velocity component u_z^* ($\text{SNR}_z \simeq 52$), and the radial velocity component u_r^* ($\text{SNR}_r \simeq 43$), for the left-to-right flow through the converging nozzle (colormap scale shown in figure 5). The images were obtained using phase-contrast MRI (see Appendix B).

reconstruction error of a modelled velocity component u_i° is measured by

$$\mathcal{E}_{u_i^\circ} = (\sqrt{|M|} \sigma_{u_i^\bullet})^{-1} \|u_i^\bullet - u_i^\circ\|_{L^2(M)}, \quad (28)$$

where u_i^\bullet is the corresponding ground truth velocity image, which, in this paper, is a high signal-to-noise ratio image of the same flow (figure 3), and $\sigma_{u_i^\bullet}^2$ is the estimated variance of Gaussian white noise in the ground truth image.

A. Noise in the signal, the phase, and the magnitude

The noise in each k -space signal s_j^* is assumed to be Gaussian, white, and additive, with zero mean and variance σ_j^2 . It is common practice in MRI to estimate the noise from regions of the magnitude images with no signal. For fully-sampled signals, the magnitude of the signal s_j^* , given by $\rho_j \equiv |\mathcal{F}^{-1} s_j^*|$, is distributed according to a Rayleigh distribution [1]. The noise variance in s_j^* is thus given by

$$\sigma_j^2 = 2(4 - \pi)^{-1} \sigma_{\rho_j}^2, \quad (29)$$

where $\sigma_{\rho_j}^2$ is the noise variance of the magnitude image ρ_j , estimated from a region with no signal. The SNR of the j -th signal is then computed from

$$\text{SNR}_j = \frac{\mu_{\rho_j}}{\sigma_{\rho_j}}, \quad (30)$$

where μ_{ρ_j} is the mean of ρ_j , taking into account only the active regions (non-zero signal). When $\text{SNR}_j > 3$, the noise in the phase images φ_j approximates a zero-mean Gaussian distribution with variance [1]

$$\sigma_{\varphi_j}^2 = \text{SNR}_j^{-2}. \quad (31)$$

Considering that the measured velocity, given by equation (3), is computed from the phase differences, the noise variance of the k -th velocity component is estimated by

$$\sigma_{u_k}^2 = c_k^2 \sum_j \sigma_{\varphi_j}^2, \quad (32)$$

for $j = 1, \dots, 4d$ and $k \equiv (j - 1) \text{div} 4 + 1$, where div denotes integer division.

Based on the above, for fully-sampled k -space signals the noise ε_{u_k} in the velocity image of the k -th component is distributed according to $\mathcal{N}(0, \sigma_{u_k}^2 \mathbf{I})$, where $\mathcal{N}(m, \mathcal{C})$ denotes the normal distribution with mean m and covariance operator \mathcal{C} . For sparsely-sampled k -space signals we observe that ε_{u_k} is

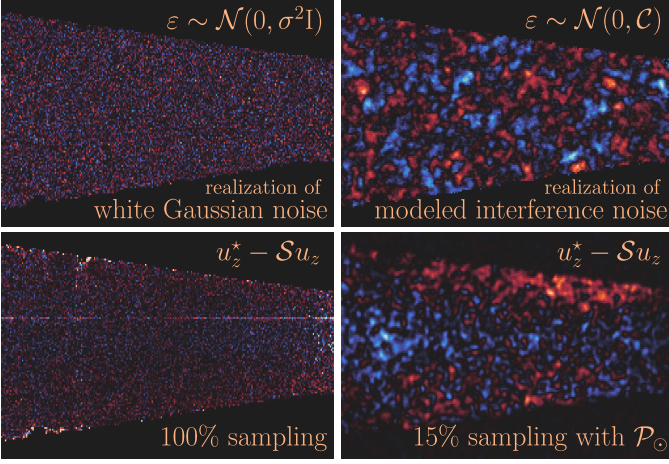


Fig. 4. Left frames: velocity discrepancy for fully-sampled (100% k -space coverage) signals. The noise ε in the measured velocity can be considered to be white and Gaussian when the signals are fully-sampled, i.e. $u_k^* - Su_k \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$. Right frames: velocity discrepancy for sparsely-sampled signals (15% k -space coverage with \mathcal{P}_\odot). For sparsely-sampled signals we observe correlations that depend on the sampling pattern \mathcal{P} . To model these correlations, we introduce a covariance operator \mathcal{C} that is constructed using an exponential kernel, i.e. $u_k^* - Su_k \sim \mathcal{N}(0, \mathcal{C})$.

TABLE I
NOISE LEVEL FOR THE ACQUIRED NOISY SIGNALS AND THE RESULTING VELOCITY IMAGES.

| | | k-space signals s_j^* | | | | | | |
|----------------------------------|------------|-------------------------|-------------------------------|------------|------------|------------|------------|--|
| σ_1 | σ_2 | σ_3 | σ_4 | σ_5 | σ_6 | σ_7 | σ_8 | |
| 38.4 | 38.5 | 14.7 | 14.6 | 43.1 | 42.9 | 19.2 | 19.2 | |
| vel. component | | c_k [cm/s] | $\sigma_{u_{(\cdot)}}$ [cm/s] | SNR $_j$ | | | | |
| axial z ($j = 1, \dots, 4$) | | 1.988 | 0.348 | 11.4 | | | | |
| radial r ($j = 5, \dots, 8$) | | 0.529 | 0.113 | 9.3 | | | | |

correlated. Here, we model this correlation using an exponential covariance function (see section II-B1). Figure 4 shows the velocity discrepancy ($u_z^* - Su_z$) between the measured (u_z^*) and the modelled (u_z) axial velocity components for fully-sampled signals, and 15% sparsely-sampled signals using \mathcal{P}_\odot (see section III-B). For fully-sampled k -spaces, we observe that

$$u_k^* - Su_k \sim \mathcal{N}(0, \sigma_{u_k}^2 \mathbf{I}) \quad , \quad (33)$$

and this justifies the use of the norm $\sigma_{u_k}^{-2} \|u_k^* - Su_k\|_{L^2}^2$ in [9], where fully-sampled signals were considered. For sparsely-sampled k -spaces, we observe that the noise is correlated and we assume that

$$u_k^* - Su_k \sim \mathcal{N}(0, \mathcal{C}_{u_k}) \quad , \quad (34)$$

where \mathcal{C}_{u_k} is given by equation (8), with characteristic length ℓ equal to the smallest resolved length scale in the image I .

B. Generating the sparse sampling patterns \mathcal{P}_\odot and \mathcal{P}_\parallel

To generate the sampling pattern \mathcal{P}_\odot for a discretized two-dimensional k -space, we draw random samples \mathbf{y}_k^s from a normal distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, whose mean and covariance are given by

$$\boldsymbol{\mu} = \left(\frac{n_1}{2}, \frac{n_2}{2} \right) \quad , \quad \boldsymbol{\Sigma} = \text{diag} \left(\frac{\omega_1 n_1}{4}, \frac{\omega_2 n_2}{4} \right) \quad , \quad (35)$$

where n_i denotes the number of points (or pixels), and ω_i denotes the coverage, in the sense of a 2σ interval, along the i -th direction. We further round the samples $\mathbf{y}_k^s = (y_{k1}^s, y_{k2}^s) \mapsto \lfloor \mathbf{y}_k^s \rfloor = (\lfloor y_{k1}^s \rfloor, \lfloor y_{k2}^s \rfloor)$ to the closest integer, so that they correspond to the pixel index. The degree of subsampling is controlled by the sampling density, which is defined by N_s/N , where N_s is the number of admissible sampled k -space points (we reject samples that either repeat on the same pixel or lie outside the k -space domain) and $N = n_1 n_2$ is the total number of k -space points. The sampling pattern $\mathcal{P}_\odot : \mathbb{Z}^{n_1 \times n_2} \rightarrow \mathbb{R}$ is then given by

$$\mathcal{P}_\odot(\mathbf{y}) = 1 \quad \text{if} \quad \mathbf{y} = \lfloor \mathbf{y}_k^s \rfloor \quad \text{else} \quad 0 \quad , \quad (36)$$

for $k = 1, \dots, N_s$. The sampling pattern $\mathcal{P}_\parallel : \mathbb{Z}^{n_1 \times n_2} \rightarrow \mathbb{R}$ consists of random samples of vertical lines, and is given by

$$\mathcal{P}_\parallel(y_1, \cdot) = 1 \quad \text{if} \quad y_1 = \lfloor y_{k1}^s \rfloor \quad \text{else} \quad 0 \quad , \quad (37)$$

for $k = 1, \dots, N_s$.

C. Reconstructing sparse PC-MRI signals

Using the sparse sampling patterns \mathcal{P}_\odot and \mathcal{P}_\parallel , with $\omega_1 = \omega_2 = 0.35$, and for various sampling densities N_s/N , we sparsify the originally full k -space signals that we have acquired for the flow through the nozzle. We use algorithm 1 to reconstruct the sparse, noisy (low SNR) signals s^* (see table I for the noise level), and compare the results with the ground truth (high SNR) images that we have acquired for the same flow (figure 3). The low SNR images required a total scanning time of 2.6 minutes per velocity image (axial and radial components), and the high SNR images required a total scanning time of 68 minutes per velocity image. These scanning times correspond to 100% k -space coverage. Thus, all other things being equal, the 15%-sampled low SNR signals would have been acquired in ~ 23 seconds, assuming that acquisition time linearly scales with sampling density³.

1) *Input and initialization*: The input set of algorithm 1 consists of the k -space signals s^* , the estimated noise variances (table I), and the prior assumption for $\boldsymbol{\varphi}$, $\boldsymbol{\rho}$ and \mathbf{x} . The prior assumptions are in fact Gaussian random fields, and we therefore need to define both their mean and their variance (confidence level of *a priori* knowledge). We choose the prior mean of the phase ($\bar{\boldsymbol{\varphi}}$) and the magnitude ($\bar{\boldsymbol{\rho}}$) images to be their respective zero-filling solution. For the N-S unknowns, we choose the prior mean of the signed distance function (SDF) ($\bar{\boldsymbol{\phi}}_\pm$) to be the SDF that corresponds to the Chan-Vese segmentation [35], [36] of the averaged zero-filling magnitude image. Shape regularization is particularly important at this point because, for very sparse signals with strong artefacts (e.g. 10% \mathcal{P}_\parallel -sampling), the Chan-Vese magnitude segmentation will provide a poor approximation of the true geometry (we show how we regularize the shape and compute the SDF from a segmentation in [9, Section 2.4]). We set the prior mean for the inlet velocity boundary condition to $\bar{\mathbf{g}}_i = (\bar{g}_{iz}, 0)$, where \bar{g}_{iz} is a parabolic velocity profile that satisfies the zero-velocity boundary condition on the boundary $\partial\Gamma_i$ of the inlet

³The validity of this assumption depends on both the sparse sampling pattern and the pulse sequence of the magnetic resonance experiment.

Γ_i , and has peak velocity equal to 2.5 cm/s. The prior mean of the outlet boundary condition is $\bar{g}_o \equiv \mathbf{0}$ (pseudotraction-free boundary condition), and of the kinematic viscosity is $\bar{\nu} \simeq 2.54 \times 10^{-5}$ m²/s [43], [44], which corresponds to a 70% glycerol in water mixture (see Appendix B). We explain how we choose the prior variances and the regularization parameters in Appendix C and in [9, Section 3.6]. Lastly, the initial guess for the measured velocity is obtained using equation (3) with $\bar{\varphi}$, and the initial guess for the modelled velocity by solving the Navier–Stokes problem $\mathcal{Z} \bar{x}$.

TABLE II
RECONSTRUCTION ERROR $\mathcal{E}_z^\bullet / \mathcal{E}_r^\bullet$ FOR EACH SAMPLING PATTERN \mathcal{P} .

| sampling | k-space coverage | | | |
|-------------------------|------------------|-------------|-------------|-------------|
| | 5% | 10% | 15% | 25% |
| \mathcal{P}_\odot | 1.13 / 0.47 | 0.66 / 0.42 | 0.55 / 0.30 | . |
| \mathcal{P}_\parallel | . | 2.41 / 1.78 | 0.63 / 0.39 | 0.61 / 0.29 |
| full sampling | 0.56 / 0.31 | | | |

2) *Reconstruction results:* We first test the algorithm on sparse k -space signals that we generate using \mathcal{P}_\odot for 5%, 10% and 15% sampling. The generated sampling patterns, the reconstructed measured velocity \mathbf{u}^* (obtained from the phase difference), and the reconstructed modelled velocity \mathbf{u} (N–S solution), are shown in figure 5, alongside their corresponding zero-filling solution. The fourth column of figure 5 depicts the result of the image reconstruction algorithm [9], which we use when the signals are fully-sampled. For fully-sampled k -space signals, there is a one-to-one correspondence between the k -space and the physical space, and, consequently, we can directly reconstruct the velocity field in physical space (\mathbf{u}^* is known, but noisy, and fixed during the reconstruction).

In the experiments we selected velocity encoding parameters that cause aliasing due to phase wrapping, in order to increase the SNR of the velocity images. For the 100% sampling case with this flow, we use basic knowledge of fluid mechanics to unwrap the image by adding or subtracting multiples of 2π to aliased pixels of the phase difference image. We can do this because we already know that the velocity components have well-defined regions with fixed signs. We use this simple method of unwrapping only for the 100% sampling case, in order to be able to reconstruct it using the algorithm in [9]. For 5%, 10% and 15% sampling, the present algorithm (PICS, algorithm 1) unwraps and reconstructs the measured velocity simultaneously and autonomously (figure 7). This is possible with algorithm 1 because it treats \mathbf{u}^* as a function of the unknown phases φ , i.e. with \mathbf{u}^* no longer being fixed.

The velocity reconstruction errors, defined by equation (28), are presented in table II for \mathcal{P}_\odot -subsamped and \mathcal{P}_\parallel -subsamped signals. We observe that for 15% \mathcal{P}_\odot -sampling, the reconstruction errors are $\mathcal{E}_z^\bullet = 0.55$ and $\mathcal{E}_r^\bullet = 0.30$, and compare this result to the 100%-sampling reconstruction errors of $\mathcal{E}_z^\bullet = 0.56$ and $\mathcal{E}_r^\bullet = 0.31$. The fact that the \mathcal{P}_\odot -subsamped signal has a lower reconstruction error than the 100%-sampled signal may seem counterintuitive, but in reality the difference in the errors is negligible, and it can

be attributed to either the technical implementation details of this large-scale problem, or in the outliers, the artefacts and the noise in the remaining 85% of the k -space. For the sampling pattern \mathcal{P}_\odot , both the reconstruction error values (table II) and a visual inspection of the velocity images (figures 5,3) show that 15% sampling is sufficient to accurately reconstruct the flow through the converging nozzle. Next, we test the algorithm on sparse k -space signals that we generate using \mathcal{P}_\parallel for 10%, 15% and 25% sampling. This sampling pattern is more coherent than \mathcal{P}_\odot , and, therefore, it produces strong artefacts in the phase and magnitude images for low sampling densities (figure 6), making the reconstruction problem harder to solve. For the sampling pattern \mathcal{P}_\parallel , both the reconstruction error values (table II) and a visual inspection of the velocity images (figures 6,3) show that 25% sampling is sufficient in order to accurately reconstruct the flow through the converging nozzle.

Although we treat the kinematic viscosity ν as an unknown parameter, for all reconstructions we find that the posterior distribution of ν remains effectively unchanged, i.e. $\nu^\circ \simeq \bar{\nu}$ and $\sigma_\nu^\circ \simeq \sigma_\nu$, where σ_ν° is the posterior standard deviation. This is because the prior information is accurate enough (see Appendix C, where the prior variance σ_ν^2 is small) and the model \mathcal{M} cannot further improve this prediction because the velocity reconstruction functional \mathcal{U} is insensitive to such small changes in ν (the prior term in equation (51) dominates) [9, Section 3.5].

D. Wall shear stress estimation

Finally we compute the reconstructed wall shear rate γ_w° (figure 8) using the reconstructed (noiseless, SNR = ∞) velocity field $\mathbf{u}^\circ (= \mathcal{Z}\mathbf{x}^\circ)$ and the inferred shape $\partial\Omega^\circ$, and compare it with the ground truth wall shear rate γ_w^\bullet , computed for the ground truth (high SNR) velocity measurement \mathbf{u}^\bullet (figure 3) on the ground truth (high SNR) shape $\partial\Omega^\bullet$. We further compute the uncertainty in γ_w° by propagating the uncertainty of the reconstructed N–S unknowns \mathbf{x}° through the N–S problem. We observe that the ground truth wall shear rate is particularly noisy because PC-MRI suffers from low resolution and partial volume effects [45], [46] near the boundaries $\partial\Omega$. Using standard image processing algorithms, it is possible to either denoise γ_w^\bullet directly, or compute γ_w^\bullet for a denoised velocity \mathbf{u}^\bullet and a smoothed boundary $\partial\Omega^\bullet$. This denoised velocity field will not, however, be consistent with the new smoothed boundary (the no-slip wall boundary condition will not be satisfied), and the steep, near-wall velocity gradients that dictate the wall shear rate magnitude will not be accurately recovered. On the other hand, our inverse Navier–Stokes boundary value problem (algorithm 1), jointly reconstructs the velocity field and infers the boundary. This accurately reconstructs the near-wall velocity distribution, and, if needed, can be made more accurate by increasing the N–S model resolution. This method infers the boundary by drawing information from both the velocity field and the spin density (magnitude) images so that geometric errors that affect the estimated wall shear stress distribution can be minimized.




Fig. 5. We reconstruct sparse k -space signals with 5%, 10% and 15% \mathcal{P}_{\odot} -sampling using algorithm 1 (first three columns), and compare the results with the reconstruction of the 100%-sampled k -space signal using the algorithm in [9] (last column). For the reconstructed sparse signals (first three columns), we show the zero-filled velocity (rows 2,5), the reconstructed measured velocity (rows 3,6), and the reconstructed modelled velocity (rows 4,7). For the 100%-sampled k -space signal (last column), we show the resulting velocity using Fourier inversion (rows 2,5), the postprocessed (unwrapped and masked) Fourier inversion (rows 3,5) and the reconstructed velocity (rows 4,7), obtained using the algorithm in [9] (flow is from left to right).




Fig. 6. As for figure 5, but for sparse k -space signals with 10%, 15% and 25% \mathcal{P}_{\parallel} -sampling.




Fig. 7. Steps, k , in the iterative process that simultaneously reconstructs and unwraps the velocity. This is shown for the measured axial velocity component u_z^* (left), and the measured radial velocity component u_r^* (right), starting from their respective zero-filled solution (colormap scale shown in figure 6).

E. Further reducing sampling density

1) *Targeted sampling for the N - S unknowns \mathbf{x}* : At this point, it is important to mention that the sampling density can be further decreased (for fixed reconstruction error) if the *a priori* information on \mathbf{x} becomes more accurate. The present method indicates that when the Navier–Stokes problem is an accurate and descriptive model, only the boundary $\partial\Omega$, the boundary conditions $\mathbf{g}_i, \mathbf{g}_o$, and the kinematic viscosity are needed in order to find the velocity field (see figure 1). Therefore, to reduce signal acquisition time further, we can obtain more accurate spin density (magnitude) images separately (without encoding velocity) and use them as priors for $\partial\Omega$ ($\hat{\phi}_\pm$). Similarly, we can obtain $d-1$ dimensional scans for the inlet velocity boundary condition, and use them as priors for $\hat{\mathbf{g}}_i$. Conventional MRI or PC-MRI cannot measure \mathbf{g}_o and

ν , but ν is usually known with high certainty compared to the other parameters. Targeted sampling using a Navier–Stokes problem is an interesting generalization [27] of sparsity in CS, and closely follows the concept of compressed sampling.

2) *Design of optimal sampling patterns \mathcal{P}* : Another way to further decrease the sampling density is to design optimal sampling patterns [47], also taking into account the constraints of the PC-MRI signal acquisition protocols. For example, in section III-C2 we found that the generated 15% \mathcal{P}_\odot -sampling pattern produces slightly better results than the generated 25% \mathcal{P}_\parallel -sampling pattern, but note that the former adequately covers the center of k -space (figure 5), while the latter leaves gaps (figure 6). This can be corrected by fully-sampling the central region of k -space and then sparsely-sampling regions that are not as important. To identify the important regions, we can use our algorithm to backpropagate the errors from the velocity images to the k -space. The physics-informed design of optimal sparse sampling patterns is left for future work.

F. Extension to multi-coil parallel MRI

Our algorithm extends to parallel MRI [48] using a SENSE-type approach [49], in which the individual coil sensitivities are inferred. In brief, the k -space reconstruction error, which in non-parallel MRI is given by (15), is now given by

$$\mathcal{E}(\varphi, \rho, \lambda) \equiv \sum_{j=1}^{4d} \frac{1}{2\sigma_j^2} \sum_{k=1}^{\Lambda} \|s_{jk}^* - \mathcal{P}\mathcal{F}\lambda_k(\rho_j e^{i\varphi_j})\|_{L^2(I_s)}^2, \quad (38)$$

where Λ is the total number of coils, and $\lambda_k \in L^2(I)$ is the (unknown) sensitivity of the k -th coil.

To infer λ we introduce an additional regularization term

$$\mathcal{R}_\lambda \equiv \frac{1}{2} \sum_{k=1}^{\Lambda} \|\lambda_k - \bar{\lambda}_k\|_{\mathcal{C}_{\lambda_k}(I)}^2, \quad (39)$$

where $\bar{\lambda}_k$ is the prior mean (or initial guess), and \mathcal{C}_{λ_k} is a user-selected covariance operator. Since coil sensitivity fields are often assumed to be smooth, a reasonable choice is an operator that is similar to the Bessel potential [50, Chapter V.3]

$$\mathcal{C}_{\lambda_k} = \sigma_{\lambda_k}^2 \left(\mathbf{I} - \ell^2 \Delta \right)^{-s}, \quad (40)$$

where $\sigma_{\lambda_k}^2 \in \mathbb{R}$ is proportional to the variance of the k -th coil sensitivity (level of confidence in the prior), $\ell \in \mathbb{R}$ is a characteristic length scale (length scales smaller than ℓ are suppressed), and $\mathbb{R} \ni s > 0$ is a user-selected parameter based on smoothness assumptions ($s \gtrsim 1$). It is worth noting that the coil sensitivity priors $\bar{\lambda}_k$ can be obtained with a method such as ESPIRiT [51], as in [52].

Replacing functional (15) with (38), and adding (39) to the Lagrangian \mathcal{J} , which is given by (17), the corresponding gradients $D_{\lambda_k} \mathcal{J}$ and the steepest ascent directions $\hat{D}_{\lambda_k} \mathcal{J}$ can be derived. Algorithm 1 must be augmented to include one more reconstruction step, during which the steepest ascent directions $\hat{D}_\lambda \mathcal{J}$ and the step size τ are first computed, and the coil sensitivities are then updated by

$$\lambda_{k+1} \leftarrow \lambda_k - \tau \hat{D}_\lambda \mathcal{J}, \quad (41)$$

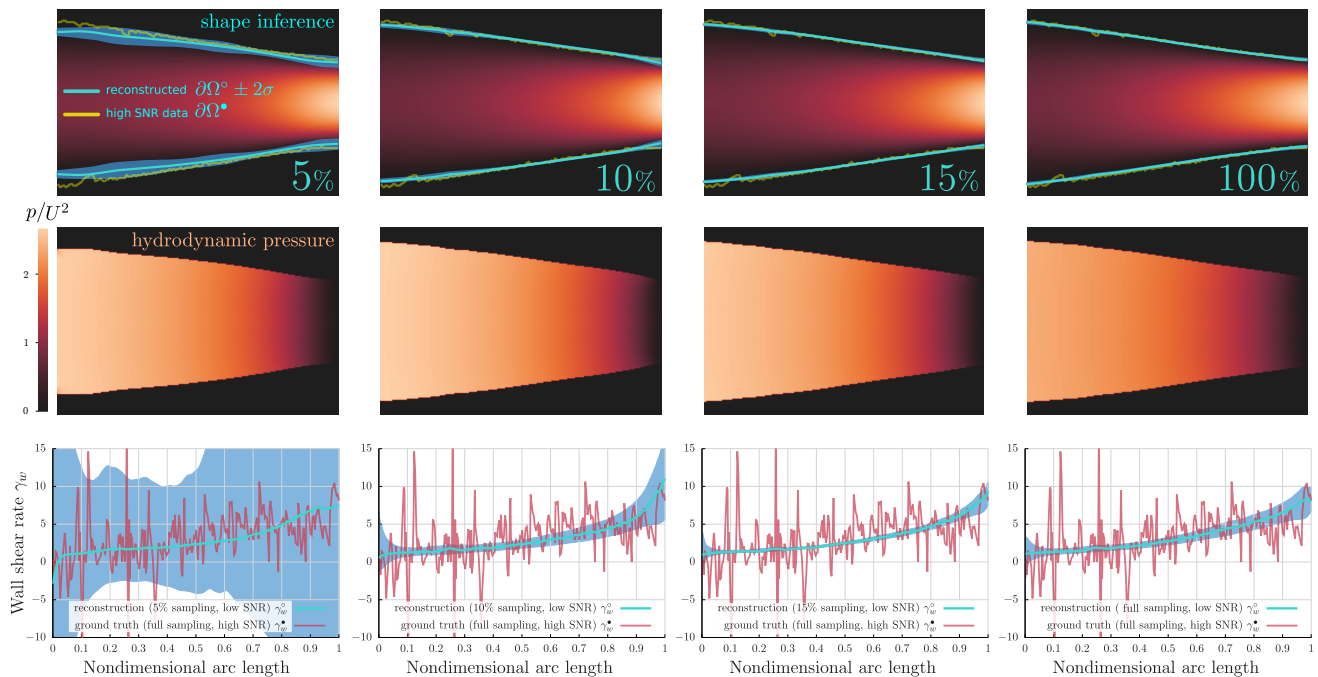


Fig. 8. Inferred shape, $\partial\Omega^\circ$, and velocity magnitude (first row) for 5%, 10% and 15% \mathcal{P}_\odot -sampling using algorithm 1, and for 100%-sampling using the algorithm in [9]. The inferred shape is shown in cyan color and the 2σ confidence interval in blue color. The yellow line denotes the ground truth of $\partial\Omega$, computed from the segmentation of the high SNR magnitude images. The inferred (nondimensional) reduced hydrodynamic pressure (second row) p°/U^2 , where $U = 4.39$ cm/s, is immediately obtained from the Navier–Stokes problem without any additional computation (see algorithm 1). The predicted wall shear rate $\gamma_w^\circ \equiv \boldsymbol{\tau} \cdot \partial_\nu \mathbf{u}^\circ$ (third row), where $\boldsymbol{\tau}$ is the unit tangent vector of $\partial\Omega^\circ$, is shown in cyan, and the 2σ confidence interval, estimated by propagating the uncertainties of \mathbf{x}° through the Navier–Stokes problem, is shown in blue. The red line denotes the wall shear rate computed from the high SNR velocity data (γ_w^\bullet). The wall shear stress τ_w is found by multiplying the wall shear rate by the dynamic viscosity μ ($= \rho\nu$, where $\rho = 1183.6$ kg/m³), i.e. $\tau_w = \mu\gamma_w$.

where k is the iteration index of algorithm 1. In this way we minimize $\mathcal{J}(\varphi_{k+1}, \mathbf{x}_{k+1}, \rho_{k+1}, \mathbf{c}_{k+1}, \boldsymbol{\lambda})$ with respect to the coil sensitivities $\boldsymbol{\lambda}$, with all other unknowns being fixed.

IV. CONCLUSION

We propose a physics-informed compressed sensing (PICS) reconstruction method that optimally combines computational fluid dynamics with phase-contrast magnetic resonance imaging. By adopting a Bayesian framework, we use *a priori* knowledge to regularize the inverse problem and find the most likely measured velocity field, segmentation, and Navier–Stokes twin (maximum *a posteriori* estimates) for the flow through a converging nozzle. We show that the physics-informed reconstruction of the sparse (15% \mathcal{P}_\odot -sampling), low SNR signals, compares well to fully-sampled, high SNR images, of the same flow.

PICS extends the capabilities of conventional compressed sensing methods for phase-contrast MRI. In particular, it integrates an optimal filter (the regularity of the velocity field is controlled by the Navier–Stokes problem), it infers the hydrodynamic pressure without any additional computation, and allows us to estimate the wall shear stresses at negligible additional numerical cost. At the same time it provides us with a digital twin, which can subsequently be used to simulate different flow conditions, and paves the way for the development of new applications in patient-specific cardiovascular modeling. Our formulation naturally extends to three-dimensional and periodic flows in complicated geometries, and the Navier–Stokes problem can be updated to model complex

(e.g. non-Newtonian) fluid flows. We are currently developing efficient numerical algorithms to implement our formulation in three dimensions.

ACKNOWLEDGMENTS

The authors would like to thank Scott V. Elgersma and Andrew J. Sederman for generating the PC-MRI dataset, which is the same as that of [9], and for their useful input and comments during the preparation of this manuscript. Author A.K. is financially supported by the W.D. Armstrong scholarship from the Cambridge Trusts.

APPENDIX A

E–L EQUATIONS FOR φ , ρ , \mathbf{c} AND \mathbf{x}

Here we expand (24)–(27) in order to obtain explicit relations for the corresponding generalized gradients $D_{(\cdot)}\mathcal{J}$, which are defined by (1). Note that the steepest ascent directions, which are defined by (2), are given by

$$\widehat{D}_{(\cdot)}\mathcal{J} = \mathcal{C}_{(\cdot)}D_{(\cdot)}\mathcal{J} \quad . \quad (42)$$

For the phases and magnitudes we find

$$D_{\varphi_j} \mathcal{J} = \chi_{I_w} (-1)^{j \operatorname{div} 2} c_k \mathcal{C}_{u_k}^{-1} (u_k^* - \mathcal{S} u_k) + \mathcal{C}_{\varphi_j}^{-1} \Re \left(i e^{i\varphi_j} (e^{i\varphi_j} - e^{i\bar{\varphi}_j})^\dagger \right) + \sigma_j^{-2} \rho_j \Im \left(e^{i\varphi_j} (\mathcal{F}^{-1} \mathcal{P}^\dagger (s_j^* - \mathcal{P} \mathcal{F} \rho_j e^{i\varphi_j}))^\dagger \right), \quad (43)$$

$$D_{\rho_j} \mathcal{J} = \chi_{I_w} (\sqrt{4d} \sigma_{\rho_j})^{-2} \left((\rho_j - \alpha) \mathcal{S} \mathcal{H}(\phi_\pm) + (\rho_j - \beta) (\mathcal{S} \mathcal{H}(\phi_\pm) - 1) \right) + \mathcal{C}_{\rho_j}^{-1} (\rho_j - \bar{\rho}_j) - \sigma_j^{-2} \Re \left(e^{i\varphi_j} (\mathcal{F}^{-1} \mathcal{P}^\dagger (s_j^* - \mathcal{P} \mathcal{F} \rho_j e^{i\varphi_j}))^\dagger \right), \quad (44)$$

for $j = 1, \dots, 4d$, where $k \equiv (j - 1) \operatorname{div} 4 + 1$, χ_{I_w} is the characteristic function of $I_w \subseteq I$. Note that $D_\varphi \mathcal{J}(\varphi, \mathbf{u}, \boldsymbol{\rho}) \in L^2(I)$, and $D_\rho \mathcal{J}(\boldsymbol{\rho}, \phi_\pm, \varphi, \alpha, \beta) \in L^2(I)$. For the segmentation constants we find

$$D_\alpha \mathcal{J} = - \sum_{j=1}^{4d} (\sqrt{4d} \sigma_{\rho_j})^{-2} \int_\Omega \rho_j - \alpha, \quad (45)$$

$$D_\beta \mathcal{J} = - \sum_{j=1}^{4d} (\sqrt{4d} \sigma_{\rho_j})^{-2} \int_{I_w - \Omega} \rho_j - \beta, \quad (46)$$

where $I_w - \Omega$ denotes the complement of Ω in I_w . Note that $D_\alpha \mathcal{J}(\boldsymbol{\rho}, \phi_\pm, \alpha) \in \mathbb{R}$, and $D_\beta \mathcal{J}(\boldsymbol{\rho}, \phi_\pm, \beta) \in \mathbb{R}$. Fixing $\boldsymbol{\rho}$ and Ω (or, equivalently, ϕ_\pm), and setting (45) and (46) to zero, we find the following explicit relations for the constants α, β that minimize \mathcal{J}

$$\alpha = \sum_{j=1}^{4d} (\sqrt{4d} |\Omega| \sigma_{\rho_j})^{-2} \int_\Omega \rho_j, \quad (47)$$

$$\beta = \sum_{j=1}^{4d} (\sqrt{4d} |I_w - \Omega| \sigma_{\rho_j})^{-2} \int_{I_w - \Omega} \rho_j. \quad (48)$$

The gradients for the N–S unknowns \mathbf{x} are drawn directly from [9, Sections 2.3,2.4]

$$D_{\mathbf{g}_i} \mathcal{J} = \nu \partial_\nu \mathbf{v} - q\mathbf{v} - \eta \mathbf{v} + \mathcal{C}_{\mathbf{g}_i}^{-1} (\mathbf{g}_i - \bar{\mathbf{g}}_i), \quad (49)$$

$$D_{\mathbf{g}_o} \mathcal{J} = \mathbf{v} + \mathcal{C}_{\mathbf{g}_o}^{-1} (\mathbf{g}_o - \bar{\mathbf{g}}_o), \quad (50)$$

$$D_\nu \mathcal{J} = \nabla \mathbf{v} : \nabla \mathbf{u} + \sigma_\nu^{-2} (\nu - \bar{\nu}), \quad (51)$$

$$D_{\phi_\pm} (\mathcal{J} - \mathcal{S}) = \dot{\zeta} + \mathcal{C}_{\phi_\pm}^{-1} (\bar{\phi}_\pm - \phi_\pm), \quad (52)$$

where $\dot{\zeta} \in L^2(M)$ is an extension of the shape gradient $\zeta \in L^2(\Gamma)$ that is generated by an advection-diffusion problem [9, Section 2.4], and the shape gradient is given by

$$\zeta = \partial_\nu \mathbf{u} \cdot (-\nu \partial_\nu \mathbf{v} + q\mathbf{v}). \quad (53)$$

The segmentation functional \mathcal{S} , given by (12), was not included in [9] because the boundary $\partial\Omega$ was inferred only from the velocity images. Here, \mathcal{S} allows us to draw information from the spin density (magnitude) images in order to better infer the boundary. For example, near regions of low velocity gradients (e.g. nearly separated flow) the velocity field alone is not informative. Consequently, taking into account \mathcal{S} , we add the following term in (52)

$$D_{\phi_\pm} \mathcal{S} = \dot{\zeta}_{\mathcal{S}} \in L^2(M), \quad (54)$$

where $\dot{\zeta}_{\mathcal{S}}$ is an advection-diffusion extension of $\zeta_{\mathcal{S}} \in L^2(\Gamma)$, similar to $\dot{\zeta}$, and

$$\zeta_{\mathcal{S}} = \delta_\Gamma \sum_{j=1}^{4d} (\sqrt{8d} \sigma_j)^{-2} \mathcal{S}^\dagger \left((\rho_j - \alpha)^2 - (\rho_j - \beta)^2 \right), \quad (55)$$

where δ_Γ is the Dirac measure on Γ , or, equivalently, the characteristic function of Γ ($\delta_\Gamma = \chi_\Gamma$).

APPENDIX B

PHASE-CONTRAST MRI EXPERIMENT

The PC-MRI experiment for the flow through the converging nozzle is described in detail in [9, Section 3.4]. In this paper, in order to demonstrate the de-aliasing capability of algorithm 1, we use k -space signals (of the same dataset) that produced severely aliased velocity images. These signals correspond to slightly different flow conditions than those described in [9]. They were acquired for a 70 wt%, instead of a 40 wt%, glycerol in water solution, and the Reynolds number was 22.4. We found the T_1 relaxation time of this glycerol solution to be 313ms. For the high SNR images, the repetition time was 0.5s, which resulted in a total acquisition time of 68 minutes per velocity image. For the low SNR images, the repetition time was 150ms, which resulted in a total acquisition time of 2.6 minutes per velocity image. The total acquisition times correspond to fully-sampled k -spaces.

APPENDIX C

PRIOR INFORMATION FOR THE N–S UNKNOWNNS

The prior variances of the Navier–Stokes unknowns are given in table III. Note that, for 5% $\mathcal{P}_\circ/\mathcal{P}_\parallel$ -sampling, instead of $\sigma_{\phi_\pm} = 2h_z$ we set $\sigma_{\phi_\pm} = 10h_z$ as the prior information on the geometry becomes less accurate due to subsampling artefacts. For a more detailed discussion of the regularization parameters, see [9, Section 3.6].

TABLE III
CONFIDENCE LEVEL OF PRIOR INFORMATION.

| \mathcal{P} | shape $\partial\Omega$ ($\sigma_{\phi_\pm}, \operatorname{Re}_{\phi_\pm}, \operatorname{Re}_\zeta$) | inlet \mathbf{g}_i ($\sigma_{\mathbf{g}_i} [\text{cm/s}], \ell$) | viscosity ν $\sigma_\nu [\text{m}^2/\text{s}]$ | resolution $h_z \times h_r [\mu\text{m}]$ |
|-------------------------|--|---|---|--|
| \mathcal{P}_\circ | ($2h_z, 0.025, 0.025$) | (1.0, $3h_r$) | 2.54×10^6 | 165×223 |
| \mathcal{P}_\parallel | ($2h_z, 0.010, 0.010$) | (1.0, $3h_r$) | 2.54×10^6 | 165×223 |

REFERENCES

- [1] H. Gudbjartsson and S. Patz, “The Rician Distribution of Noisy MRI Data,” *Magnetic Resonance in Medicine*, vol. 34, no. 6, pp. 910–914, 1995.
- [2] S. G. Chang, B. Yu, and M. Vetterli, “Adaptive wavelet thresholding for image denoising and compression,” *IEEE Transactions on Image Processing*, vol. 9, no. 9, pp. 1532–1546, 2000.
- [3] G. Pascal, “Rudin–Osher–Fatemi Total Variation Denoising using Split Bregman,” *Image Processing On Line*, vol. 2, no. 1, pp. 74–95, 2012.
- [4] N. Fatourae and A. A. Amini, “Regularization of flow streamlines in multislice phase-contrast MR imaging,” *IEEE Transactions on Medical Imaging*, vol. 22, no. 6, pp. 699–709, 2003.
- [5] F. Ong, M. Uecker, U. Tariq, A. Hsiao, M. T. Alley, S. S. Vasanawala, and M. Lustig, “Robust 4D flow denoising using divergence-free wavelet transform,” *Magnetic Resonance in Medicine*, vol. 73, no. 2, pp. 828–842, 2015.

- [6] J. Mura, A. M. Pino, J. Sotelo, I. Valverde, C. Tejos, M. E. Andia, P. Irrarázaval, and S. Uribe, "Enhancing the Velocity Data From 4D Flow MR Images by Reducing its Divergence," *IEEE Transactions on Medical Imaging*, vol. 35, no. 10, pp. 2353–2364, 2016.
- [7] T. S. Koltukluoğlu and P. J. Blanco, "Boundary control in computational haemodynamics," *Journal of Fluid Mechanics*, vol. 847, pp. 329–364, 2018.
- [8] S. W. Funke, M. Nordaas, Ø. Evju, M. S. Alnæs, and K. A. Mardal, "Variational data assimilation for transient blood flow simulations: Cerebral aneurysms as an illustrative example," *International Journal for Numerical Methods in Biomedical Engineering*, vol. 35, no. 1, pp. 1–27, 2019.
- [9] A. Kontogiannis, S. V. Elgersma, A. J. Sederman, and M. P. Juniper, "Joint reconstruction and segmentation of noisy velocity images as an inverse Navier–Stokes problem," *Journal of Fluid Mechanics*, vol. 944, p. A40, 2022.
- [10] E. J. Candès, J. Romberg, and T. Tao, "Robust Uncertainty Principles : Exact Signal Frequency Information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [11] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [12] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [13] M. Benning, L. Gladden, D. Holland, C. B. Schönlieb, and T. Valkonen, "Phase reconstruction from velocity-encoded MRI measurements - A survey of sparsity-promoting variational approaches," *Journal of Magnetic Resonance*, vol. 238, pp. 26–43, 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.jmr.2013.10.003>
- [14] D. J. Holland, D. M. Malioutov, A. Blake, A. J. Sederman, and L. F. Gladden, "Reducing data acquisition times in phase-encoded velocity imaging using compressed sensing," *Journal of Magnetic Resonance*, vol. 203, no. 2, pp. 236–246, 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.jmr.2010.01.001>
- [15] T. Roberts, N. Kingsbury, and D. J. Holland, "Sparse recovery of complex phase-encoded velocity images using iterative thresholding," *2013 IEEE International Conference on Image Processing, ICIP 2013 - Proceedings*, pp. 350–354, 2013.
- [16] F. Zhao, D. C. Noll, J. F. Nielsen, and J. A. Fessler, "Separate magnitude and phase regularization via compressed sensing," *IEEE Transactions on Medical Imaging*, vol. 31, no. 9, pp. 1713–1723, 2012.
- [17] A. Reci, "Signal sampling and processing in magnetic resonance applications," Ph.D. dissertation, University of Cambridge, 2019.
- [18] K. Bredies, K. Kunisch, and T. Pock, "Total generalized variation," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 492–526, 2010.
- [19] V. Corona, M. Benning, L. F. Gladden, A. Reci, A. J. Sederman, and C.-B. Schönlieb, "Joint Phase Reconstruction and Magnitude Segmentation from Velocity-Encoded MRI Data," *Time-dependent Problems in Imaging and Parameter Identification*, pp. 1–24, 2021. [Online]. Available: <http://arxiv.org/abs/1908.05285>
- [20] J. M. Kollmeier, O. Kalentev, J. Klosowski, D. Voit, and J. Frahm, "Velocity vector reconstruction for real-time phase-contrast MRI with radial Maxwell correction," *Magnetic Resonance in Medicine*, vol. 87, no. 4, pp. 1863–1875, 2022. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29108>
- [21] A. Bakhshinejad, A. Baghaie, A. Vali, D. Saloner, V. L. Rayz, and R. M. D'Souza, "Merging computational fluid dynamics and 4D Flow MRI using proper orthogonal decomposition and ridge regression," *Journal of Biomechanics*, vol. 58, pp. 162–173, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0021929017302531>
- [22] J. Töger, M. J. Zahr, N. Aristokleous, K. Markenroth Bloch, M. Carlsson, and P.-O. Persson, "Blood flow imaging by optimal matching of computational fluid dynamics to 4D-flow data," *Magnetic Resonance in Medicine*, vol. 84, no. 4, pp. 2231–2245, 2020. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28269>
- [23] J. Sotelo, J. Urbina, I. Valverde, C. Tejos, P. Irrarázaval, M. E. Andia, S. Uribe, and D. E. Hurtado, "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 Transactions on Medical Imaging*, vol. 35, no. 6, pp. 1475–1487, 2016.
- [24] J. Zhang, M. C. Brindise, S. Rothenberger, S. Schnell, M. Markl, D. Saloner, V. L. Rayz, and P. P. Vlachos, "4D Flow MRI Pressure Estimation Using Velocity Measurement-Error-Based Weighted Least-Squares," *IEEE Transactions on Medical Imaging*, vol. 39, no. 5, pp. 1668–1680, 2020.
- [25] P. Berg, D. Stucht, G. Janiga, O. Beuing, O. Speck, and D. Thévenin, "Cerebral blood flow in a healthy circle of willis and two intracranial aneurysms: Computational fluid dynamics versus four-dimensional phase-contrast magnetic resonance imaging," *Journal of Biomechanical Engineering*, vol. 136, no. 4, pp. 1–9, 2014.
- [26] J. Id and P. Berg, "2018 (MATCH)— Phase Ib : Effect of morphology on hemodynamics," vol. 2018, 2019.
- [27] M. F. Duarte and Y. C. Eldar, "Structured compressed sensing: From theory to applications," *IEEE Transactions on Signal Processing*, vol. 59, no. 9, pp. 4053–4085, 2011.
- [28] I. Bright, G. Lin, and J. N. Kutz, "Compressive sensing based machine learning strategy for characterizing the flow around a cylinder with limited pressure measurements," *Physics of Fluids*, vol. 25, no. 12, 2013.
- [29] E. Ferdian, A. Suinesiaputra, D. J. Dubowitz, D. Zhao, A. Wang, B. Cowan, and A. A. Young, "4DFlowNet: Super-Resolution 4D Flow MRI Using Deep Learning and Computational Fluid Dynamics," *Frontiers in Physics*, vol. 8, 2020. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/fphy.2020.00138>
- [30] D. R. Rutkowski, A. Roldán-Alzate, and K. M. Johnson, "Enhancement of cerebrovascular 4D flow MRI velocity fields using machine learning and computational fluid dynamics simulation data," *Scientific Reports*, vol. 11, no. 10240, 2021.
- [31] M. J. Muckley, B. Riemenschneider, A. Radmanesh, S. Kim, G. Jeong, J. Ko, Y. Jun, H. Shin, D. Hwang, M. Mostapha, S. Arberet, D. Nickel, Z. Ramzi, P. Ciuciu, J.-L. Starck, J. Teuwen, D. Karkalousos, C. Zhang, A. Sriram, Z. Huang, N. Yakubova, Y. W. Lui, and F. Knoll, "Results of the 2020 fastMRI Challenge for Machine Learning MR Image Reconstruction," *IEEE Transactions on Medical Imaging*, vol. 40, no. 9, pp. 2306–2317, 2021.
- [32] M. J. Colbrook, V. Antun, and A. C. Hansen, "The difficulty of computing stable and accurate neural networks: On the barriers of deep learning and Smale's 18th problem," *Proceedings of the National Academy of Sciences*, vol. 119, no. 12, p. e2107151119, 2022. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.2107151119>
- [33] L. Bungert, M. Burger, Y. Korolev, and C.-B. Schönlieb, "Variational regularisation for inverse problems with imperfect forward operators and general noise models," *Inverse Problems*, vol. 36, no. 12, p. 125014, dec 2020. [Online]. Available: <https://doi.org/10.1088/1361-6420/abc531>
- [34] S. Lunz, A. Hauptmann, T. Tarvainen, C.-B. Schönlieb, and S. Arridge, "On Learned Operator Correction in Inverse Problems," *SIAM Journal on Imaging Sciences*, vol. 14, no. 1, pp. 92–127, 2021. [Online]. Available: <https://doi.org/10.1137/20M1338460>
- [35] T. F. Chan and L. A. Vese, "Active contours without edges," *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 266–277, 2001.
- [36] P. Getreuer, "Chan–Vese Segmentation," *Image Processing On Line*, vol. 2, pp. 214–224, 2012.
- [37] J. Nitsche, "Über ein Variationsprinzip zur Lösung von Dirichlet-Überrn bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind," *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, vol. 36, no. 1, pp. 9–15, 1971.
- [38] R. Fletcher, *Practical Methods of Optimization*. John Wiley & Sons, 2000.
- [39] J. Nocedal and S. J. Wright, *Numerical optimization*, 2006.
- [40] S. L. Cotter, M. Dashti, J. C. Robinson, and A. M. Stuart, "Bayesian inverse problems for functions and applications to fluid mechanics," *Inverse Problems*, vol. 25, no. 11, 2009.
- [41] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, "Array programming with NumPy," *Nature*, vol. 585, no. 7825, pp. 357–362, Sep. 2020. [Online]. Available: <https://doi.org/10.1038/s41586-020-2649-2>
- [42] P. Virtanen *et al.*, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [43] N.-S. Cheng, "Formula for the Viscosity of a Glycerol–Water Mixture," *Industrial & Engineering Chemistry Research*, vol. 47, no. 9, pp. 3285–3288, may 2008. [Online]. Available: <https://doi.org/10.1021/ie071349z>
- [44] A. Volk and C. J. Kähler, "Density model for aqueous glycerol solutions," *Experiments in Fluids*, vol. 59, no. 5, p. 75, 2018. [Online]. Available: <https://doi.org/10.1007/s00348-018-2527-y>
- [45] P. Bouillot, B. M. Delattre, O. Brina, R. Ouared, M. Farhat, C. Chnafa, D. A. Steinman, K. O. Lovblad, V. M. Pereira, and M. I. Vargas, "3D phase contrast MRI: Partial volume correction for robust blood flow quantification in small intracranial vessels," *Magnetic Resonance in Medicine*, vol. 79, no. 1, pp. 129–140, 2018.
- [46] K. Saito, S. Abe, M. Kumamoto, Y. Uchiyama, A. Tanaka, K. Sugie, M. Ihara, M. Koga, and H. Yamagami, "Blood Flow Visualization and

- Wall Shear Stress Measurement of Carotid Arteries Using Vascular Vector Flow Mapping,” *Ultrasound in Medicine and Biology*, vol. 46, no. 10, pp. 2692–2699, 2020.
- [47] L. F. Gladden and A. J. Sederman, “Magnetic Resonance Imaging and Velocity Mapping in Chemical Engineering Applications,” *Annual Review of Chemical and Biomolecular Engineering*, vol. 8, no. 1, pp. 227–247, 2017.
- [48] A. Deshmane, V. Gulani, M. A. Griswold, and N. Seiberlich, “Parallel MR imaging,” *Journal of Magnetic Resonance Imaging*, vol. 36, no. 1, pp. 55–72, 2012. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.23639>
- [49] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, “SENSE: Sensitivity encoding for fast MRI,” *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291522-2594%28199911%2942%3A5%3C952%3A%3AAID-MRM16%3E3.0.CO%3B2-S>
- [50] E. M. Stein, *Singular Integrals and Differentiability Properties of Functions (PMS-30)*. Princeton University Press, 1970. [Online]. Available: <http://www.jstor.org/stable/j.ctt1bpmb07>
- [51] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig, “Espirit—an eigenvalue approach to autocalibrating parallel mri: Where sense meets grappa,” *Magnetic Resonance in Medicine*, vol. 71, no. 3, pp. 990–1001, 2014. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24751>
- [52] A. Sun, B. Zhao, K. Ma, Z. Zhou, L. He, R. Li, and C. Yuan, “Accelerated phase contrast flow imaging with direct complex difference reconstruction,” *Magnetic Resonance in Medicine*, vol. 77, no. 3, pp. 1036–1048, 2017. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26184>