

Bayesian inference of mean velocity fields and turbulence models from flow MRI

By A. Kontogiannis[†], P. Nair, M. Loecher, D. B. Ennis, A. Marsden
AND M. P. Juniper[†]

We solve a Bayesian inverse Reynolds-averaged Navier–Stokes (RANS) problem that assimilates mean flow data by jointly reconstructing the mean flow field and learning its unknown RANS parameters. We devise an algorithm that learns the most likely parameters of an algebraic effective viscosity model, and estimates their uncertainties, from mean flow data of a turbulent flow. We conduct a flow MRI experiment to obtain mean flow data of a confined turbulent jet in an idealized medical device known as the FDA (Food and Drug Administration) nozzle. The algorithm successfully reconstructs the mean flow field and learns the most likely turbulence model parameters without overfitting. The methodology accepts any turbulence model, be it algebraic (explicit) or multi-equation (implicit), as long as the model is differentiable, and naturally extends to unsteady turbulent flows.

1. Introduction

Probability theory (Cox 1946; Jaynes 2003) provides a fundamental and rigorous framework for extracting information from data when we know, to some degree of certainty, what the data contain. In this approach, some parameters are fixed while others are expressed as probability distributions (MacKay 2003; Stuart 2010; Ghattas & Willcox 2021). These distributions are initially asserted by the user and are then updated deterministically with Bayesian inference as data arrive. If more than one model could explain the data, the likelihood of each model is calculated and the most likely model is selected (MacKay 2003). This framework provides a principled way to tune and choose models, given the available data. This is also known as inverse problem theory, described in seminal texts such as Kaipio & Somersalo (2005); Idier (2002); Sullivan (2015).

The challenge with performing Bayesian inference on partial differential equations (PDEs) in general and fluid dynamics in particular is that function evaluations are usually expensive. This means that sampling methods such as Markov chain Monte Carlo (MCMC) or Hamiltonian Monte Carlo (HMC) quickly become intractable (Kaipio & Somersalo 2005; Sullivan 2015). For problems in which the outputs vary smoothly with their parameters, however, we can exploit the fact that Gaussian prior distributions lead to nearly Gaussian posterior distributions, meaning that Laplace’s approximation (a second-order Taylor expansion of the model output with respect to its parameters) becomes sufficiently accurate to be useful (MacKay 2003; Sullivan 2015). Laplace’s approximation can then be combined with adjoint methods such that many parameters of large PDE problems can be inferred simultaneously and cheaply.

In previous studies, we have demonstrated adjoint-accelerated Bayesian inference on experimental flow MRI data from a laminar flow through a nozzle (Kontogiannis *et al.*

[†] Department of Engineering, University of Cambridge, United Kingdom

2022) and a 3D-printed aortic arch (Kontogiannis *et al.* 2024a). Flow MRI data provides three-component velocity fields in a large 3D volume. The spatial resolution is around 1 mm³ per voxel, which is coarse compared with many optical diagnostics, and the signal-to-noise ratio (SNR) is relatively low. Nevertheless, when combined with prior knowledge that this velocity field must obey Navier–Stokes and the no-slip boundary condition, this coarse noisy flow MRI data is sufficient to infer the flow with surprising accuracy. For example, Figure 1 shows a reconstruction of 3D three-component flow MRI velocity data of steady flow through a 3D-printed aortic arch at Reynolds number 550. The raw data [Figure 1(a)] at typical SNR is combined with prior information expressed through a Finite Element Method (FEM), which enforces conservation of mass and momentum [Figure 1(b)]. The velocity field is inferred rapidly with adjoint-accelerated Bayesian inference [Figure 1(c)]. The inferred velocity is almost identical to the high SNR velocity data [Figure 1(d)], which the algorithm did not see. Furthermore, the reconstruction in Figure 1(c) has higher spatial resolution than the high SNR data and also calculates the position of the aorta’s boundary [thin grey line in Figure 1(c)]. The other velocity components are presented in Kontogiannis *et al.* (2024a).

The above studies assumed uniform viscosity. In Kontogiannis *et al.* (2024b), we measured 3D steady flow MRI data of a shear-thinning jet and assimilated the flow and viscosity model simultaneously. The inferred viscosities were nearly identical to those measured in 18 separate rheometry experiments over four orders of magnitude of shear rate. This showed that it is possible to assimilate velocity data into the Navier–Stokes equations with a non-uniform viscosity model.

This study extends Kontogiannis *et al.* (2024b) to a turbulent flow from a 10-mm diameter jet at Reynolds number 6500. The spatial resolution of flow MRI is insufficient to resolve turbulent structures in this flow, meaning that turbulence models are required to assimilate this flow MRI data. We use adjoint-accelerated Bayesian inference to combine flow MRI velocity data with prior knowledge that the data are from a turbulent jet that obeys the no-slip condition, and the RANS equations with a prescribed eddy viscosity model. We note that the local turbulent kinetic energy (TKE) can be calculated from the decay rate of the MR signal, although we do not use this information in the current paper.

2. Bayesian inversion of RANS problems

Given noisy flow MRI data of a mean flow, \mathbf{u}^* , we assume that there is a RANS solution, $\mathbf{u}^\circ = \mathbf{u}(\mathbf{x}^\circ)$, that can fit the data for suitable parameters, \mathbf{x}° , i.e.,

$$\mathbf{u}^* = \mathcal{Z}\mathbf{x}^\circ + \boldsymbol{\varepsilon} \quad , \quad (2.1)$$

where $\mathcal{Z} := \mathcal{S}\mathcal{Q}$ is the operator that maps model parameters to model solutions projected into the data space, \mathcal{S} is the model space to data space projection operator, \mathcal{Q} is the operator that encapsulates the physical model, and $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathcal{C}_{\mathbf{u}^*})$ is Gaussian noise with zero mean and covariance operator $\mathcal{C}_{\mathbf{u}^*}$. We further assume that we have prior knowledge of the probability distribution of the model parameters, \mathbf{x} . We call this the prior parameter distribution, $\mathcal{N}(\bar{\mathbf{x}}, \mathcal{C}_{\bar{\mathbf{x}}})$, which we assume to be Gaussian with prior mean $\bar{\mathbf{x}}$, and prior covariance operator $\mathcal{C}_{\bar{\mathbf{x}}}$. We then use Bayes’ theorem, which states that the posterior probability density function (PDF) of \mathbf{x} , given the data \mathbf{u}^* , $\pi(\mathbf{x}|\mathbf{u}^*)$,




FIGURE 1: (a) A 2D slice of 3D three-component velocity data from flow MRI through a 3D-printed aorta. This data has typical signal-to-noise-ratio (SNR) and is colored by the velocity into the page, u_y . (b) A diagram of our prior knowledge that this velocity data comes from a fluid that obeys the Navier–Stokes equations within a domain Ω bounded by a no-slip surface Γ , a Dirichlet velocity condition on Γ_i , and a stress-free boundary on Γ_o . This prior knowledge is expressed through a finite element CFD solver in which the free parameters define the shape of Γ , the velocity on Γ_i , and the viscosity μ . The optimal parameter values and their uncertainties are found with adjoint-accelerated Bayesian inference giving (c) the inferred velocity field and boundary position, which can be compared with (d) high SNR data gathered over several hours, which was not seen by the assimilation algorithm. The reconstructed field in panel (c) is a higher resolution twin of the high SNR data in panel (d) despite assimilating only the low SNR data in panel (a) and the prior knowledge in panel (b). This figure is adapted from Kontogiannis *et al.* (2024a).

is proportional to the data likelihood, $\pi(\mathbf{u}^*|\mathbf{x})$, times the prior PDF of \mathbf{x} , $\pi(\mathbf{x})$, i.e.,

$$\begin{aligned} \pi(\mathbf{x}|\mathbf{u}^*) &\propto \pi(\mathbf{u}^*|\mathbf{x}) \pi(\mathbf{x}) \\ &= \exp\left(-\frac{1}{2}\|\mathbf{u}^* - \mathcal{Z}\mathbf{x}\|_{\mathcal{C}_{\mathbf{u}^*}}^2 - \frac{1}{2}\|\mathbf{x} - \bar{\mathbf{x}}\|_{\mathcal{C}_{\bar{\mathbf{x}}}}^2\right) \quad , \end{aligned} \quad (2.2)$$

where $\pi(\cdot)$ is the Gaussian PDF, and $\|\cdot, \cdot\|_{\mathcal{C}}^2 := \langle \cdot, \mathcal{C}^{-1} \cdot \rangle$ is the covariance-weighted L^2 -norm. The most likely parameters \mathbf{x}° (maximum *a-posteriori* probability, or MAP estimator), are given implicitly as the solution of the nonlinear optimization problem

$$\mathbf{x}^\circ \equiv \underset{\mathbf{x}}{\operatorname{argmin}} \mathcal{J}(\mathbf{x}) \quad , \quad \text{where} \quad \mathcal{J} := -\log\left(\pi(\mathbf{u}^*|\mathbf{x})\pi(\mathbf{x})\right) \quad . \quad (2.3)$$

For \mathbf{x} in a neighborhood of the MAP point, \mathbf{x}° , we can Laplace-approximate (MacKay 2003, Chapter 27) the posterior distribution such that

$$\pi(\mathbf{x}|\mathbf{y}^*) \simeq \tilde{\pi}(\mathbf{x}|\mathbf{y}^*) := \exp\left(-\frac{1}{2}\|\mathbf{x} - \mathbf{x}^\circ\|_{\mathcal{C}_{\mathbf{x}^\circ}}^2 - \text{const.}\right) \quad , \quad (2.4)$$

where $\mathcal{C}_{\mathbf{x}^\circ} := (\mathcal{G}_{\mathbf{x}^\circ}^* \mathcal{C}_{\mathbf{u}^*}^{-1} \mathcal{G}_{\mathbf{x}^\circ} + \mathcal{C}_{\bar{\mathbf{x}}}^{-1})^{-1}$ is the posterior covariance operator, $\mathcal{G}_{\mathbf{x}^\circ}$ is a linearization of operator \mathcal{Z} around \mathbf{x}° , and $\mathcal{G}_{\mathbf{x}^\circ}^*$ is the adjoint operator of $\mathcal{G}_{\mathbf{x}^\circ}$ (see Table 1). The Bayesian inversion of Navier–Stokes problems is treated in more detail in Kontogiannis *et al.* (2022, 2024a).

\mathcal{S}	Model-to-data space projection	\mathcal{Q}	RANS problem: maps \mathbf{x} to \mathbf{u}
\mathcal{Z}	($:= \mathcal{S}\mathcal{Q}$) maps \mathbf{x} to $\mathcal{S}\mathbf{u}$	$\mathcal{G}_{\mathbf{x}_k}$	Linearization of \mathcal{Z} around \mathbf{x}_k
$\mathcal{C}_{\mathbf{u}^*}$	Data covariance operator	$\mathcal{C}_{\mathbf{x}_k}$	Model parameter covariance operator around \mathbf{x}_k


TABLE 1: Operators used in the Bayesian RANS inversion (Kontogiannis *et al.* 2024a)

FIGURE 2: Bayesian inversion of RANS problems: we assimilate mean flow data, \mathbf{u}^* , in order to jointly reconstruct the mean flow field, \mathbf{u} , and learn the unknown RANS parameters, \mathbf{x} . The most likely RANS parameters, \mathbf{x}° , are found by solving problem (2.3), and the inferred mean velocity field is given by $\mathbf{u}^\circ = \mathcal{Q}\mathbf{x}^\circ$ (Kontogiannis *et al.* 2024a).

2.1. RANS problem formulation

The RANS boundary value problem (BVP) in $\Omega \subset I_m \subset \mathbb{R}^3$ (see Figure 2) is given by

$$\left\{ \begin{array}{l} \mathbf{u} \cdot \nabla \mathbf{u} - \nabla \cdot (2\nu_e \nabla^s \mathbf{u}) + \nabla p = \mathbf{0} \quad \text{in } \Omega \\ \nabla \cdot \mathbf{u} = 0 \quad \text{in } \Omega \\ \mathbf{u} = \mathbf{0} \quad \text{on } \Gamma \\ \mathbf{u} = \mathbf{g}_i \quad \text{on } \Gamma_i \\ -2\nu_e \nabla^s \mathbf{u} \cdot \boldsymbol{\nu} + p\boldsymbol{\nu} = \mathbf{g}_o \quad \text{on } \Gamma_o \end{array} \right. , \quad \underbrace{\mathbf{x} = (\Omega, \mathbf{g}_i, \mathbf{g}_o, \nu_e)}_{\text{Unknown RANS parameters}} , \quad (2.5)$$

where \mathbf{u} is the mean velocity, $p \leftarrow p/\rho$ is the reduced pressure, ρ is the density, ν_e is the effective kinematic viscosity, $\nabla^s \mathbf{u} \equiv (\nabla^s \mathbf{u})_{ij} := \frac{1}{2}(\partial_j u_i + \partial_i u_j)$ is the strain-rate tensor, \mathbf{g}_i is the Dirichlet boundary condition (BC) at the inlet Γ_i , \mathbf{g}_o is the natural BC at the outlet Γ_o , and $\boldsymbol{\nu}$ is the unit normal vector on the boundary $\partial\Omega = \Gamma \cup \Gamma_i \cup \Gamma_o$, where Γ is the part of the boundary where a no-slip BC is imposed. Problem (2.5) can model many different wall-bounded incompressible laminar flows for Newtonian or generalized Newtonian fluids (i.e., power-law non-Newtonian fluids), as well as turbulent mean flows. Here, since we are interested in modeling turbulent mean flows, we define

$$\mu_e := \mu_\ell + \mu_t \quad , \quad \mu_t := \ell_c^2 \dot{\gamma} \quad , \quad \dot{\gamma} := \sqrt{2\nabla^s \mathbf{u} : \nabla^s \mathbf{u}} \quad , \quad (2.6)$$

where $\mu_e = \rho\nu_e$ is the effective dynamic viscosity, μ_ℓ is the laminar viscosity, μ_t is the eddy viscosity, ℓ_c is a compound mixing length, and $\dot{\gamma}$ is the shear rate magnitude. In this paper, in order to fit the mean flow of a confined turbulent jet, we introduce the compound mixing length

$$\ell = \ell(d_s, d_w; \alpha, \beta, x_c, c, d_{s_0}) := \alpha \mathcal{H}_\eta (d_s - d_{s_0}) + \beta(1 - \mathcal{H}_\eta) d_w \quad , \quad (2.7)$$

where d_s is the distance measured along the flow (streamline) main direction, d_w is the distance from the wall, $\alpha, \beta \in \mathbb{R}$ are the unknown proportionality constants, d_{s_0} is an offset,

$$\mathcal{H}_\eta = \mathcal{H}_\eta(x_c, c) := \left(1 + \exp \left(\frac{(d_s - x_c)^2 - c^2}{\eta c^2} \right) \right)^{-1} \quad (2.8)$$

is a smooth activation (Heaviside-like) function, centered around x_c , with width $2c$, and $\eta = 0.2$. Note that, the compound mixing length model that we have introduced is designed for confined turbulent jets in order to deal with the two different flow regions: (i) the jet breakdown region, where the mixing length is a function of the streamwise distance, and (ii) the pipe flow region, where the mixing length is a function of the wall distance. For fixed density ρ , μ_e (or, equivalently, ν_e) is then described by an algebraic six-parameter model such that $\mu_e = \mu_e(\ell_c, \dot{\gamma}; \mathbf{p})$, where $\mathbf{p} := (\mu_\ell, \alpha, \beta, x_c, c, d_{s_0})$. The unknown RANS parameters thus become $\mathbf{x} = (\Omega, \mathbf{g}_i, \mathbf{g}_o, \mathbf{p})$.

3. The flow MRI experiment

3.1. Phantom design

The phantom was designed in SolidWorks based on the FDA nozzle geometry definition with modifications to fully develop the inlet flow and accommodate direct pressure measurements. The nozzle design, which is shown in Figure 3, includes the following main features:

- (1) *Conical inlet baffle*. A conical inlet baffle upstream of the inlet expansion zone was added to reduce the opening angle and thereby prevent formation of a jet. The 1:1 ratio of the areas of the two paths at the inlet was maintained at the outlet to the baffle.
- (2) *Flow straightener*. A grid-like flow straightener was added prior to the converging section to avoid the introduction of jets in the region of interest. The structured grid has square channels 0.2×0.2 cm with 0.1-cm walls between them.
- (3) *Conical convergence*. A convergence zone of 20° brings the diameter of the nozzle down to the diameter of the throat.
- (4) *Throat*. A narrow segment with a diameter of 1 cm and a length of 10 cm (ten times the diameter).
- (5) *Sudden expansion*. The diameter of the nozzle expands to 3 cm (three times the diameter of the throat) resulting in creation of a jet at sufficient Reynolds numbers.
- (6) *Pressure ports*. Two Luer lock ports were added. One port is proximal to the sudden expansion and one port is downstream (after the jet has dissipated). Pressure catheters were inserted through these ports to acquire estimates of pressure drop between the two points.

3.2. Flow bench setup

Given restrictions on the printer test bed size, the phantom design was split into two parts right after the throat (see Figure 4). The two pieces were printed individually using






FIGURE 3: Phantom design schematic (length units in centimeters).



(a)



(b)

FIGURE 4: Nozzle phantom assembled on the benchtop. (a) Assembled phantom embedded into a solid ballistics gel. (b) Bench setup schematic.

Clear Loctite 3D IND405 in a Carbon L1 3D printer to have rigid walls with a wall thickness of 0.3 cm. The two parts were screwed together with a rubber gasket in between to ensure the interface was watertight. The assembled phantom was embedded into a solid ballistics gel (ClearBallistics) for spatial stabilization and to provide static MRI signal regions for velocity (phase) offset corrections. An MRI-compatible programmable flow pump (CardioFlow 5000 MR, Shelley Medical Imaging Technologies) was used to pump flow through the phantom at a steady flow rate of 200 mL/s, corresponding to a Reynolds number of 6500. A glycerol-water mixture (40/60 ratio) was used as a blood-mimicking fluid. Ferumoxytol, a contrast medium, was added for MRI signal enhancement. The density of the fluid was 1.1 g/cm^3 and the dynamic viscosity was 4.2 mPa.s. Estimates of pressure drop were acquired using two pressure transducers (SPR-350S, Millar), one inserted proximal to the sudden expansion and another further downstream after the

jet had dissipated. MRI data was acquired on a 3T scanner (Siemens Vida Fit) in the phantom with a conventional 4D-flow measurement sequence using an 18-channel chest coil. The scan was acquired with a resolution of $1 \times 1 \times 1 \text{ mm}^3$ (isotropic), a matrix size of $256 \times 256 \times 44$, TE/TR = 3.00/7.14 ms, a flip angle of 7° , GRAPPA R = 2, and an acquisition time of 4 minutes. Four different data sets were acquired with velocity encoding strength (VENC) of [500, 200, 100, 50] cm/s. The final imaging data set was generated as a magnitude weighted average of the different VENC acquisitions, unwrapped against the VENC=500 cm/s data set.

4. Reconstruction of flow MRI data

We apply the RANS inversion methodology [see section 2 and Kontogiannis *et al.* (2024a)[†]] to reconstruct 3D mean flow data of a confined turbulent jet measured using flow MRI (see section 3). We assume that the noise in the measured velocity field is white Gaussian (Gudbjartsson & Patz 1995), with zero mean and covariance operator $\mathcal{C}_{\mathbf{u}^*} = \sigma \mathbf{I}$, where $\sigma = 5 \text{ cm/s}$. We fix the geometry, Ω , which is found by segmenting the magnitude image of the flow MRI scan, and set the outlet BC to $\mathbf{g}_o \equiv \mathbf{0}$. We thus only need infer the unknowns $\mathbf{x} = (\mathbf{g}_i, \mathbf{p})$ in order to reconstruct the mean flow data, \mathbf{u}^* . The dimensionality of the data and model spaces is given in Table 2.

We define the component-wise data-model discrepancy by

$$\mathcal{E}(u_i) := |\Omega|^{-1} \|u_i^* - \mathcal{S}u_i\|_{L^2(\Omega)} \quad . \quad (4.1)$$

We find that the prior-data discrepancies are $(\mathcal{E}(\bar{u}_x), \mathcal{E}(\bar{u}_y), \mathcal{E}(\bar{u}_y)) / \sigma = (1.45, 1.44, 5.30)$, and the MAP-data discrepancies are $(\mathcal{E}(u_x^o), \mathcal{E}(\bar{u}_y^o), \mathcal{E}(\bar{u}_y^o)) / \sigma = (1.30, 1.31, 2.08)$. The Bayesian RANS inversion algorithm takes 48 iterations to converge; 95% of the error reduction, however, is achieved in the first 20 iterations (see Figure 5). The RANS model reconstruction (i.e., the inferred mean flow) and the discrepancies are shown in Figure 6(f). The inversion algorithm infers the velocity inlet BC, \mathbf{g}_i , and the effective viscosity field parameters, \mathbf{p} , such that the discrepancies between the RANS model and the flow MRI data are minimized. The inferred algebraic viscosity model parameters, as well as their uncertainties, are shown in Table 3. The prior and the learned (MAP) effective viscosity fields, which are functions of the parameters \mathbf{p} , are shown in Figure 7.

4.1. Discussion

In this paper we use a five-parameter algebraic turbulence model for the effective viscosity field (Boussinesq approximation), which is based on a compound mixing length, such that $\mu_t = \ell_c^2 \dot{\gamma}$. Figure 6(f) shows that the algorithm has managed to find model parameters that cause the inferred velocity field to closely match the measured velocity field. To do so, it has significantly altered the effective viscosity and compound mixing length fields from their prior values, as shown in Figure 7. In particular, the RANS inversion algorithm learns the compound mixing length parameters that generate a modeled velocity field that fits both the jet breakdown region, and the pipe flow region. This shows that this RANS model is sufficiently descriptive for this confined jet flow. These results prove the concept of assimilating RANS model parameters from flow MRI data.

Future investigations will consider different turbulence models. A different, dimensionally consistent algebraic model based on the turbulence kinetic energy, k , is given by

[†] For the Bayesian RANS inversion algorithm, we use the same numerical stabilization parameters as in Kontogiannis *et al.* (2024a).

Data voxels	Data resolution	Model voxels	Model resolution
$40 \times 45 \times 169$	$1 \times 1 \times 1$ mm	$53 \times 60 \times 225$	$0.75 \times 0.75 \times 0.75$ mm

TABLE 2: Data and model spaces for the Bayesian inverse RANS problem.

Inferred (MAP estimate)					
μ_ℓ [mPa.s]	α	β	x_c [cm]	c [cm]	d_{s_o} [cm]
4.2 ± 0.15	0.722 ± 0.193	19.2 ± 12.6	3.25 ± 2.87	7.32 ± 2.44	0.159 ± 0.0881
Prior					
μ_ℓ [mPa.s]	α	β	x_c [cm]	c [cm]	d_{s_o} [cm]
4.2 ± 0.15	2 ± 6	10 ± 15	0 ± 30	8.5 ± 30	0 ± 15

TABLE 3: Compound mixing length parameters.

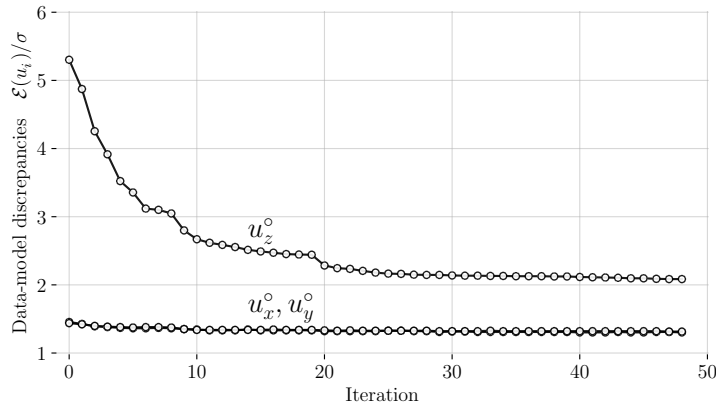


FIGURE 5: Optimization log of the data-model discrepancy of each velocity component.

$\mu_t = \ell_c k^{1/2}$. This model requires TKE data, and has been used in Casel *et al.* (2022) yielding good results. This will be examined next; however, it is important to note that algebraic models based on a learned compound mixing length and TKE data can be made to fit the measured mean flow but are not expected to extrapolate well at different flow conditions. A further improvement will be obtained from one- or two-equation turbulence models, such as the k - ϵ turbulence model, because it should extrapolate well outside the training regime.

5. Conclusions

We have solved a Bayesian inverse Reynolds-averaged Navier–Stokes (RANS) problem that assimilates 3D velocimetry data of a mean turbulent flow in order to jointly reconstruct the flow field and learn the unknown RANS parameters. By incorporating an algebraic effective viscosity model into the RANS problem, we devise an algorithm that




FIGURE 6: Noisy 3D flow MRI data, \mathbf{u}^* , and RANS model reconstruction, \mathbf{u}° . (a) Flow MRI data, \mathbf{u}^* . (b) Inferred (MAP), u_z° . Images (a,b) show velocity streamlines in cyan color and velocity magnitude in red color. (c) Flow MRI data, u_z^* . (d) Inferred (MAP), u_z° . (c,d) The mid-plane axial velocity (u_z) slice is shown and the dashed cyan line depicts the walls of the confined turbulent jet. (e) Data-prior discrepancy, $\mathcal{S}^*(u_z^* - \mathcal{S}\bar{u}_z)$. (f) Data-MAP discrepancy, $\mathcal{S}^*(u_z^* - \mathcal{S}u_z^\circ)$. (e,f) The data-model discrepancy is shown before and after assimilation of the flow MRI data. Panels (c-f) share the same colormap scale, which is $[0, 100\text{cm/s}]$.

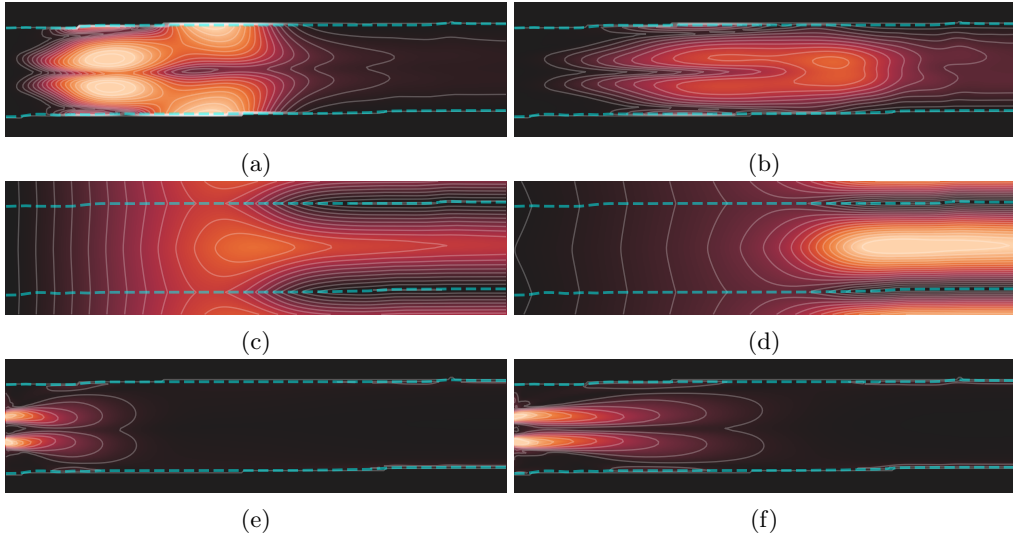


FIGURE 7: (a,c,e) Prior assumption and (b,d,f) inferred MAP estimate for (a,b) the effective viscosity field, $\mu_e := \mu_\ell + \ell_c^2 \dot{\gamma}$, (c,d) the compound mixing length, ℓ_c , and (e,f) the shear rate magnitude, $\dot{\gamma}(\mathbf{u})$ (the dashed cyan line depicts the walls of the confined turbulent jet). Panels (a,b), (c,d), and (e,f) share the same colormap scale.

learns the most likely viscosity model parameters, and estimates their uncertainties, from velocimetry data. Then we conduct a flow MRI experiment to obtain 3D velocimetry data of a confined turbulent jet through an idealized medical device, known as the Food and Drug Administration (FDA) nozzle. We show that the algorithm successfully reconstructs the noisy flow field, and, at the same time, learns the viscosity model parameters and their uncertainties. In this paper we have considered a five-parameter algebraic viscosity model. The present methodology, however, accepts any type of turbulence model, be it algebraic, one-equation, or multi-equation, as long as the model is differentiable. More complicated turbulence models, such as the k - ϵ model, can be used in order to reconstruct the flow and learn a turbulence model that can better extrapolate at different flow conditions. The turbulence model parameters can be learned by assimilating velocimetry data into the RANS problem and turbulent kinetic energy (TKE) data into the turbulence model of the kinetic energy equation. This pilot study has proven the concept. Our next step will be to assimilate more complicated turbulence models.

REFERENCES

- CASEL, M., OBERLEITHNER, K., ZHANG, F., ZIRWES, T., BOCKHORN, H., TRIMIS, D. & KAISER, T. L. 2022 Resolvent-based modelling of coherent structures in a turbulent jet flame using a passive flame approach. *Combust. and Flame* **236**, 111695.
- COX, R. T. 1946 Probability, frequency and reasonable expectation. *Am. J. Phys.* **14**, 1–13.
- GHATTAS, O. & WILLCOX, K. 2021 Learning physics-based models from data: perspectives from inverse problems and model reduction. *Acta Numer.* **30**, 445–554.
- GUDBJARTSSON, H. & PATZ, S. 1995 The Rician distribution of noisy MRI data. *Magn. Reson. Med.* **34**, 910–914.
- IDIER, J. 2002 *Approche bayésienne pour les problèmes inverses*. Hermes Science/Lavoisier.
- JAYNES, E. T. 2003 *Probability Theory: The Logic of Science*. Cambridge University Press.
- KAPIO, J. & SOMERSALO, E. 2005 *Statistical and Computational Inverse Problems*. Springer.
- KONTOGIANNIS, A., ELGERSMA, S. V., SEDERMAN, A. J. & JUNIPER, M. 2024a Bayesian inverse navier-stokes problems: joint flow field reconstruction and parameter learning. *Inverse Problems* .
- KONTOGIANNIS, A., ELGERSMA, S. V., SEDERMAN, A. J. & JUNIPER, M. P. 2022 Joint reconstruction and segmentation of noisy velocity images as an inverse Navier–Stokes problem. *J. Fluid Mech.* **944**, A40.
- KONTOGIANNIS, A., HODGKINSON, R. & MANCHESTER, E. L. 2024b *Learning rheological parameters of non-Newtonian fluids from velocimetry data*. arXiv:2408.02604 [physics.flu-dyn].
- MACKEY, D. J. C. 2003 *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.
- STUART, A. M. 2010 Inverse problems: A Bayesian perspective. *Acta Numer.* **19**, 451–459.
- SULLIVAN, T. 2015 *Introduction to Uncertainty Quantification*. Springer.