

A Reduced-Order Strategy for Solving Inverse Bayesian Shape Identification Problems in Physiological Flows

Andrea Manzoni, Toni Lassila, Alfio Quarteroni, and Gianluigi Rozza

Abstract A reduced-order strategy based on the reduced basis (RB) method is developed for the efficient numerical solution of statistical inverse problems governed by PDEs in domains of varying shape. Usual discretization techniques are infeasible in this context, due to the prohibitive cost entailed by the repeated evaluation of PDEs and related output quantities of interest. A suitable reduced-order model is introduced to reduce computational costs and complexity. Furthermore, when dealing with inverse identification of shape features, a reduced shape representation allows to tackle the geometrical complexity. We address both challenges by considering a reduced framework built upon the RB method for parametrized PDEs and a parametric radial basis functions approach for shape representation. We present some results dealing with blood flows modelled by Navier-Stokes equations.

1 Introduction

In a parametrized context, given a mathematical model of a system the forward problem consists in evaluating some outputs of interest (depending on the PDE solution) for specified parameter inputs. Whenever some parameters are uncertain,

A. Manzoni • G. Rozza (✉)

MATHICSE-CMCS Modelling and Scientific Computing, École Polytechnique Fédérale de Lausanne, Switzerland

Now at SISSA, MathLab, International school for Advanced Studies, Trieste, Italy

e-mail: andrea.manzoni@epfl.ch; gianluigi.rozza@epfl.ch; grozza@sissa.it

T. Lassila

MATHICSE-CMCS Modelling and Scientific Computing, École Polytechnique Fédérale de Lausanne, Switzerland

e-mail: toni.lassila@epfl.ch

A. Quarteroni

MATHICSE-CMCS Modelling and Scientific Computing, École Polytechnique Fédérale de Lausanne, Switzerland

Dipartimento di Matematica, MOX Modeling and Scientific Computing, Politecnico di Milano, Italy

e-mail: alfio.quarteroni@epfl.ch

we aim at inferring their values (and/or distributions) from indirect observations by solving an inverse problem: given an observed output, can we deduce the value of the parameters that resulted in this output? Parameter identification can be performed in two ways, either in a deterministic or in a statistical framework. In the former case, we solve an optimization problem by minimizing (in the least-square sense) the *discrepancy* between the output quantities predicted by the PDE model and observations: this leads to a single-point estimate in the parameter space, provided the optimization problem is feasible. In the latter case, we quantify the relative likelihood of the parameters, which are consistent with the observed output. Following a Bayesian approach, this results in the posterior probability density function, which includes information both on prior knowledge on parameters distribution and on the model used to compute the PDE-based outputs. Inverse problems governed by PDEs entail several computational challenges for current discretization techniques, such as the finite element method. When the parameters to be identified are related with the shape of the domain, the problem is even more complicated. In this framework, computational costs arise from three distinct sources: (i) numerical approximation of the state system (usually a nonlinear system of PDEs); (ii) handling domains of arbitrary shapes; (iii) sampling high-dimensional parameter spaces or performing numerical optimization procedures.

In this paper, we address these challenges by developing a reduced framework based on both *state* and *parameter* reduction, in order to devise a low-dimensional, computationally inexpensive but accurate model that predicts outputs of a high-fidelity, computationally expensive model.

The reduction in state is obtained through a reduced basis (RB) approximation [7]: thanks to a suitable offline/online stratagem, *online* PDE evaluations for any value of input parameters are completely independent of the expensive *offline* computation and storage of the basis functions. On the other hand, when input parameters are related to geometrical features, we rely on low-dimensional but flexible shape parametrizations, able to represent wide families of complex shapes by means of a handful of input parameters.

2 Inverse Problems Governed by PDEs

We introduce a compact description of general inverse problems governed by parametrized PDEs. We denote by $\boldsymbol{\mu} \in \mathcal{D} \subset \mathbb{R}^P$ the finite-dimensional vector of parameters to be identified, and consider an input-output map $\boldsymbol{\mu} \mapsto \mathbf{y}(\boldsymbol{\mu})$ from parameters to observations that is given by two discretized PDEs (taken here linear for notational simplicity):

$$\begin{aligned} \text{State equation:} \quad & A_N(\boldsymbol{\mu})\mathbf{u}_N(\boldsymbol{\mu}) = \mathbf{f}_N(\boldsymbol{\mu}) \\ \text{Observation equation:} \quad & \mathbf{y}_N(\boldsymbol{\mu}) = C_N(\boldsymbol{\mu})\mathbf{u}_N(\boldsymbol{\mu}) \end{aligned} \tag{1}$$

State variables and observed outputs are denoted by $\mathbf{u}_N \in \mathbb{R}^N$ and $\mathbf{y}_N \in \mathbb{R}^M$, respectively. By subscript N we signify the dimension of the state space, $\dim(\mathbf{u}_N) = N$, which in the case of Finite Element (FE) discretizations is typically very large, whereas the dimension of the parameter space, $\dim(\boldsymbol{\mu}) = P$, and of the observation space, $\dim(\mathbf{y}_N) = M$ can be different and typically $P, M \ll N$. In our case, $\boldsymbol{\mu}$ is related to the shape of the domain $\Omega = \Omega(\boldsymbol{\mu})$ where the state problem is posed.

Whereas the forward problem is to evaluate $\mathbf{y}(\boldsymbol{\mu})$ given $\boldsymbol{\mu}$, the inverse problem can be formulated as follows [3,5]: given an observation $\mathbf{y}^* = \mathbf{y} + \boldsymbol{\varepsilon}$ with (additive) noise $\boldsymbol{\varepsilon}$, find the parameter $\boldsymbol{\mu}^*$ that satisfies $\mathbf{y}^* = C_N(\boldsymbol{\mu}^*)\mathbf{u}_N(\boldsymbol{\mu}^*)$. This problem is often ill-posed in one of three basic ways: (i) the solution $\boldsymbol{\mu}^*$ does not exist (e.g. due to $M > P$ and the presence of noise); (ii) the solution $\boldsymbol{\mu}^*$ is not unique (e.g. due to $M < P$ or data degeneracy); or (iii) the solution $\boldsymbol{\mu}^*$ does not depend continuously on \mathbf{y}^* . An example of an inverse problem that is ill-posed in the third sense is the *Calderón problem* of determining the conductivity field inside an object based on the observation of a Dirichlet-to-Neumann or Neumann-to-Dirichlet map on a subsection of the boundary.

2.1 A Deterministic Approach

In order to treat ill-posed inverse problems, the classical approaches [1] are largely based on solving regularized least-squares (RLS) problems of the type:

$$\boldsymbol{\mu}_{\text{RLS}}^* = \arg \min_{\boldsymbol{\mu} \in \mathbb{R}^P} \frac{1}{2} \|\mathbf{y}^* - \mathbf{y}_N(\boldsymbol{\mu})\|_2^2 + \frac{\alpha}{2} \|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{prior}}\|_r^2. \tag{2}$$

The first term minimizes the discrepancy between the observation \mathbf{y}^* and the model prediction $\mathbf{y}_N(\boldsymbol{\mu})$ given by (1). The second term convexifies the problem and assures a unique estimator $\boldsymbol{\mu}_{\text{RLS}}^*$ is recovered. This approach is also sometimes called *variational data assimilation*. The choice of the norm $\|\boldsymbol{\mu}\|_r := \sqrt{\boldsymbol{\mu}^T R \boldsymbol{\mu}}$, the regularization parameter $\alpha > 0$, and the prior value $\boldsymbol{\mu}_{\text{prior}}$ play an important role in the quality of the estimator $\boldsymbol{\mu}_{\text{RLS}}^*$.

2.2 A Bayesian Approach

Under the assumption of independent and identically distributed (i.i.d.) noise, $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma^2 I)$, and Gaussian parameter distribution, $\boldsymbol{\mu} \sim \mathcal{N}(\bar{\boldsymbol{\mu}}, \Sigma)$, it is easy to show that the maximum a posteriori (MAP) estimator

$$\boldsymbol{\mu}_{\text{MAP}}^* := \underset{\boldsymbol{\mu} \in \mathbb{R}^P}{\operatorname{argmax}} \quad \pi_{\boldsymbol{\mu} | \mathbf{y}^*}(\boldsymbol{\mu} | \mathbf{y}^*) \tag{3}$$

obtained by maximizing the conditional probability density function

$$\pi_{\mu|y^*}(\boldsymbol{\mu}|y^*) \sim \exp\left(-\frac{1}{2}\|y^* - y_N(\boldsymbol{\mu})\|_2^2 - \frac{\alpha^2}{2}(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}})^T \Sigma^{-1}(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}})\right)$$

coincides with the Tikhonov-regularized least-squares estimator, that is to say $\boldsymbol{\mu}_{\text{RLS}}^* = \boldsymbol{\mu}_{\text{MAP}}^*$, as long as we choose $\boldsymbol{\mu}_{\text{prior}} = \bar{\boldsymbol{\mu}}$, $\alpha = \sigma^2$, and $R = \Sigma^{-1}$ in (2). In fact, the estimator given by (3) is an example of a wider class of statistical estimators called *Bayesian estimators*. The benefit of using statistical methods for solving inverse problems is that one is able to characterize the variance of the prediction $\boldsymbol{\mu}^*$ due to measurement and model errors more precisely than from the single-point estimates obtained by solving (2).

Bayesian estimators are a subset of statistical estimators that are widely used to solve ill-posed inverse problems. The basic principle of Bayesian inference is that the conditional distribution of the unknown parameters $\boldsymbol{\mu}$ given an observation y^* can be approximated by

$$\pi_{\mu|y^*}(\boldsymbol{\mu}|y^*) = \frac{\pi_{y^*|\mu}(y^*|\boldsymbol{\mu})\pi_{\mu}(\boldsymbol{\mu})}{\pi_{y^*}(y^*)} \approx \frac{\pi_{y^*|\mu}(y^*|\boldsymbol{\mu})\pi_{\mu,\text{prior}}(\boldsymbol{\mu})}{\pi_{y^*}(y^*)} =: \pi_{\mu,\text{post}}(\boldsymbol{\mu}|y^*)$$

obtained by using the Bayes' formula on a prior distribution $\pi_{\mu,\text{prior}}(\boldsymbol{\mu})$ for the unknown parameters. The prior $\pi_{\mu,\text{prior}}$ encapsulates our prior knowledge (structure, regularity, locality, etc.) about the distribution of the uncertain parameters and should be carefully selected based on problem-specific considerations – we do not treat this point in this work since selecting an informative prior is a challenging problem all by itself. The conditional distribution $\pi_{y^*|\mu}$, which in the case of additive noise can be expressed as

$$\pi_{y^*|\mu}(y^*|\boldsymbol{\mu}) = \pi_{\text{noise}}(y^* - y_N(\boldsymbol{\mu})),$$

is called the *likelihood function*. The posterior *distribution* $\pi_{\mu,\text{post}}$ can then be used to compute various estimators for $\boldsymbol{\mu}^*$ and to provide conditional statistics such as covariances for these estimators. The advantage of the Bayesian approach compared to more classical methods is that a prior that carries sufficient information about the true underlying structure of the parameters often provides more meaningful estimates and regularizes the inverse problem in a more natural way than relying on abstract regularization terms, as in (2), that might not have any interpretation.

Statistical methods used to solve an inverse problem can be computationally much more expensive than the deterministic approach due to the necessity of performing sampling in high-dimensional spaces in order to compute sample statistics [3, 5]. This cost is exacerbated by the fact that each evaluation requires the solution of the forward problem in the form of a (potentially large-scale) discrete PDE. To this end we introduce a reduced order model to speed up the computations entailed by statistical inversion.

3 Computational and Geometrical Reduction

We now present a brief description of the two main blocks on which the reduced order model relies: reduced basis method for parametrized PDEs and radial basis functions for low-dimensional shape parametrization. Further methodological aspects and details can be found e.g. in [6].

The reduced basis method provides an efficient way to compute an approximation $\mathbf{u}_n(\boldsymbol{\mu})$ of the solution $\mathbf{u}_N(\boldsymbol{\mu})$ (as well as an approximation $y_n(\boldsymbol{\mu})$ of the output $y_N(\boldsymbol{\mu})$) through a Galerkin projection onto a reduced subspace made up of well-chosen full-order solutions (also called *snapshots*), corresponding to a set of parameter values $S_n = \{\boldsymbol{\mu}^1, \dots, \boldsymbol{\mu}^n\}$ selected by means of a *greedy* algorithm [7]. Let us denote by $Z_n \in \mathbb{R}^{N \times n}$ the matrix

$$Z_n = [\mathbf{u}_1(\boldsymbol{\mu}) \mid \dots \mid \mathbf{u}_n(\boldsymbol{\mu})] \quad (4)$$

obtained by aligning the snapshot vectors (a Gram-Schmidt orthonormalization procedure has to be considered after each basis is added to the reduced space, but for the sake of simplicity we consider the same notation). We denote by $n \ll N$ the dimension of the *reduced* state space. Then, the reduced-order solution is given by a linear combination $Z_n \mathbf{u}_n(\boldsymbol{\mu})$ of the snapshots, being $\mathbf{u}_n \in \mathbb{R}^n$ the solution of the following problem:

$$\begin{aligned} \text{State equation:} \quad & A_n(\boldsymbol{\mu}) \mathbf{u}_n(\boldsymbol{\mu}) = \mathbf{f}_n(\boldsymbol{\mu}) \\ \text{Observation equation:} \quad & y_n(\boldsymbol{\mu}) = C_n(\boldsymbol{\mu}) \mathbf{u}_n(\boldsymbol{\mu}), \end{aligned} \quad (5)$$

where

$$A_n(\boldsymbol{\mu}) = Z_n^T A_N(\boldsymbol{\mu}) Z_n, \quad \mathbf{f}_n = Z_n^T \mathbf{f}_N, \quad C_n = C_N Z_n.$$

To get very fast input/output evaluations, RB methods rely on the assumption of affine parametric dependence in $A_N(\boldsymbol{\mu})$ and $\mathbf{f}_N(\boldsymbol{\mu})$, i.e. on the possibility to express $A_N(\boldsymbol{\mu}) = \sum_{q=1}^{Q_A} \Theta_A(\boldsymbol{\mu}) A_N^q$ and $\mathbf{f}_N(\boldsymbol{\mu}) = \sum_{q=1}^{Q_f} \Theta_f(\boldsymbol{\mu}) \mathbf{f}_N^q$, so that the expensive $\boldsymbol{\mu}$ -independent quantities can be evaluated and stored just once. This is a property inherited by the PDE model, which can be eventually recovered at the discretization stage [7].

Once the reduced model is built in the offline stage, it can be exploited at the online stage to speed up the solution of the optimization problem (2) in the deterministic case or (3) in the Bayesian case. The corresponding reduced-order version of the former reads as follows:

$$\min_{\boldsymbol{\mu} \in \mathbb{R}^P} \frac{1}{2} \|\mathbf{y}^* - y_n(\boldsymbol{\mu})\|_2^2 + \frac{\alpha}{2} \|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{prior}}\|_r^2, \quad (6)$$

whereas in the case of a statistical inverse problem we obtain:

$$\boldsymbol{\mu}_{\text{MAP}}^* := \underset{\boldsymbol{\mu} \in \mathbb{R}^P}{\text{argmax}} \quad \pi_{\boldsymbol{\mu}, \text{post}}(\boldsymbol{\mu} \mid \mathbf{y}^*) \quad (7)$$

being

$$\pi_{\boldsymbol{\mu}, \text{post}}(\boldsymbol{\mu} \mid \mathbf{y}^*) = \frac{\pi_{\text{noise}}(\mathbf{y}^* - \mathbf{y}_n(\boldsymbol{\mu}))\pi_{\boldsymbol{\mu}, \text{prior}}(\boldsymbol{\mu})}{\pi_{\mathbf{y}}(\mathbf{y}^*)}.$$

In this way, state reduction allows to speed up both numerical optimization schemes or sampling algorithms required e.g. to compute statistical estimates based on the posterior distribution.

Concerning parameter space reduction, here we consider a low-dimensional parametrization based on Radial Basis Functions (RBF), an interpolatory technique which allows to define shape deformations through a set of control points (which can be freely chosen, according to the family of deformations to be described), i.e. a linear combination of affine and radial, nonaffine terms; see e.g. [6] for more insights. In this way, parameter space reduction is afforded by selecting only a small set of $P \approx \mathcal{O}(10)$ control points at a preceding stage – state reduction through the RB method is built for a problem where shape parametrization has already been performed. A RB paradigm for simultaneous state and parameter reduction has been introduced in [5] in order to tackle the case of distributed parametric fields (instead of parameter vectors), and represents a possible extension of our current framework.

4 Application and Results

We now apply the reduced framework of the previous section to the solution of an inverse problem arising in modeling of blood flows. Since a strong mutual interaction exists between haemodynamic factors and vessels geometry, improving the understanding of the interplay between flows and geometries may be useful not only for the sake of design of better prosthetic devices [4], but also to characterize pathological risks, such as in the case of narrowing or thickening of an arterial vessel [3]. Typical portions of cardiovascular network where lesions and pathologies may develop are made up by curved vessels and bifurcations; an important segment where vessel diseases are often clinically observed is the human carotid artery [2,6], which supplies blood to the head.¹

¹The common carotid artery (CCA) bifurcates in the lower neck into two branches, the internal and the external carotid arteries (ICA and ECA, respectively). Stenoses, that is the narrowing of the inner portion of an artery, manifest quite often in the ICA.




Fig. 1 *Left*: shape representation of a stenosed carotid artery bifurcation through RBF parametrization. *Right*: velocity profiles (cm/s) in four different carotid bifurcations parametrized with respect to the diameters $d_c = d_c(\mu_1, \mu_2)$ of the CCA at the bifurcation and $d_b = d_b(\mu_3, \mu_4)$ of the mid-sinus level of the ICA

Let us consider a steady, incompressible Navier-Stokes model to describe blood flows in a two-dimensional carotid bifurcation (see Fig. 1):

$$\begin{cases} -\nu \Delta \mathbf{v} + (\mathbf{v} \cdot \nabla) \mathbf{v} + \nabla p = \mathbf{f} & \text{in } \Omega(\mu) \\ \nabla \cdot \mathbf{v} = 0 & \text{in } \Omega(\mu) \\ \mathbf{v} = \mathbf{v}_{in} & \text{on } \Gamma_{in} \\ \mathbf{v} = \mathbf{0} & \text{on } \Gamma_w \\ -p \cdot \mathbf{n} + \nu \frac{\partial \mathbf{v}}{\partial \mathbf{n}} = \mathbf{0} & \text{on } \Gamma_{out} \end{cases} \quad (8)$$

being (\mathbf{v}, p) the velocity and the pressure of the fluid, respectively, and $\nu > 0$ its kinematic viscosity. In view of studying computationally expensive inverse problems, which entail the repeated simulation of these flow equations, we cannot afford at the moment the solution of PDE models involving more complex features,

such as flow unsteadiness and arterial wall deformability – computational costs would be too prohibitive.

In this context, a typical forward problem is the evaluation of flow indices related with geometry variation that assess/measure the occlusion risk. Typical examples are given by vorticity, shear rates, wall shear stresses. On the other hand, we might be interested in recovering some geometrical features by observing some physical index related to flow variables. In particular, the inverse problem we want to solve is the following: is it possible to identify the entity of the occlusions (i.e. the diameters d_c of the CCA at the bifurcation and d_b of the mid-sinus level of the ICA, respectively) from the observation of the mean pressure drop

$$y(\boldsymbol{\mu}) = \int_{\Gamma_{in}} p(\boldsymbol{\mu}) d\Gamma - \int_{\Gamma_{out}} p(\boldsymbol{\mu}) d\Gamma$$

between the internal carotid outflow Γ_{out} and the inflow Γ_{in} ?

To exploit the reduced framework presented in Sect. 3, we represent local shape deformations through a RBF parametrization built over $p = 4$ control points (represented as the bullets in Fig. 1), located in one of the branches and close to the bifurcation. In this case, Gaussian RBFs have been used in order to describe local but moderate deformations representing possible stenoses, being $\boldsymbol{\mu} \in \mathcal{D} = [-0.25, 0.25]^4$ the vector of the displacements of the control point in the horizontal direction; see [6] for more details.

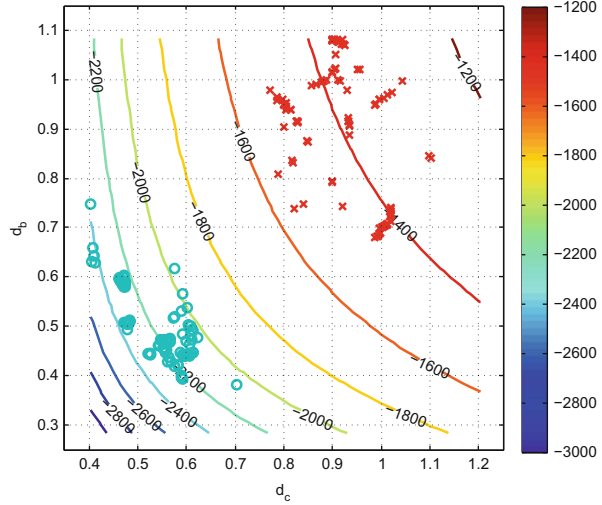
By applying the RB method to the parametrized Navier-Stokes problem (8) we reduce the dimension of the state space from $N \approx 26,000$ ($\mathbb{P}_2/\mathbb{P}_1$ FE discretization) to $n = 45$. Four examples of computed RB solutions are reported in Fig. 1. We remark the strong sensitivity of the flow with respect to varying diameters $d_c = d_c(\mu_1, \mu_2)$ of the CCA at the bifurcation and $d_b = d_b(\mu_3, \mu_4)$ of the mid-sinus level of the ICA, respectively. See e.g. [3, 6] for more insights on RB methodology for nonlinear Navier-Stokes equations.

Thus, we can take advantage of both the deterministic and the Bayesian framework to solve this inverse identification problem, by considering surrogate measurements of the mean pressure drop.

In the first case, we demonstrate the solution of the deterministic inverse problem for two different observed values of the pressure drop, $s^* = -1,400$ and $s^* = -2,200$, by assuming 5 % relative additive noise in the measurements. The results of the inverse identification problem are given in Fig. 2 for 100 realization of random noise in both cases: each point in the graph corresponds to the recovered diameters (d_c, d_b) given a noisy observation. We observe that in the case $s^* = -1,400$ recovered values of the diameters are more smeared out, since locally the pressure drop surface is almost flat, but result is close in values to the considered observation.

Thus, in the former case $s^* = -1,400$ the inverse problem is worse conditioned than in the latter $s^* = -2,200$, where the recovered values (d_c, d_b) lie in a smaller region of the space. However, the solution of a single optimization problem is more feasible in the former case compared to the latter: solving 100 optimization

Fig. 2 Results of the deterministic inverse problems for $s^* = -1,400$ (in red) and $s^* = -2,200$ (in green). Isocontours of the pressure drop (RB Navier-Stokes problem)



problems took about 14h in the former and about 25.6h in the latter case, respectively. We remark that solving 100 inverse problems of this type through a full-order discretization technique would have been infeasible on a standard workstation. Thus, even in presence of small noises, the result of a deterministic inverse problem may be very sensitive – just when one diameter is known, the second one can be recovered. This is due to the fact that several geometrical configurations – in terms of diameters (d_c, d_b) – may correspond to the same output observation.

Following instead the Bayesian approach, we are able to characterize a *set* of configurations, rather than a single configuration: this is done by providing the *joint* probability distribution function for the (uncertain) diameters (d_c, d_b) encapsulating the noise related to measurements, as discussed in Sect. 2.2. Let us denote by $\mathbf{d} = (d_c, d_b)^T \in \mathbb{R}^2$ the vector of the two diameters and assume that the prior distribution is $\pi_{\mathbf{d},\text{prior}}(\boldsymbol{\mu}) \sim \mathcal{N}(\mathbf{d}_M, \Sigma_M)$, being $\mathbf{d}_M \in \mathbb{R}^2$ the (prior) mean and $\Sigma_M \in \mathbb{R}^{2 \times 2}$ the (prior) covariance matrix, encapsulating a possible prior knowledge on the diameters distribution (e.g. from observations of previous shape configurations). By supposing that also the measurements of the pressure drop are expressed by i.i.d. Gaussian variables, such that $\pi_{\text{noise}}(\mathbf{y}^* - \mathbf{y}_n(\mathbf{d})) \sim \mathcal{N}(0, \sigma^2)$, we can compute the explicit form of the posterior probability density $\pi_{\mathbf{d},\text{post}}(\mathbf{d}|\mathbf{y}^*)$. Thus, provided some preliminary information on plausible values of the diameters, the observation of a (large) sample of outputs allows to characterize a set of plausible configurations as the ones maximizing the posterior probability density $\pi_{\mathbf{d},\text{post}}(\mathbf{d}|\mathbf{y}^*)$. In particular, we consider two different realizations of prior normal distributions, obtained by choosing the mean $\mathbf{d}_M = (0.803, 0.684)^T$ as given by the diameters corresponding to the reference carotid configuration, and




Fig. 3 *Top*: two different choices of the prior distribution on diameters $\mathbf{d} = (d_c, d_b)^T$; *left*: $\pi_{\mathbf{d},\text{prior}} \sim \mathcal{N}(\mathbf{d}_M, \Sigma_{M_1})$, *right*: $\pi_{\mathbf{d},\text{prior}} \sim \mathcal{N}(\mathbf{d}_M, \Sigma_{M_2})$. *Center and bottom*: results of the Bayesian inverse problems (*left*: $\pi_{\mathbf{d},\text{prior}} \sim \mathcal{N}(\mathbf{d}_M, \Sigma_{M_1})$, *right*: $\pi_{\mathbf{d},\text{prior}} \sim \mathcal{N}(\mathbf{d}_M, \Sigma_{M_2})$) and observed pressure drop $s^* = -1,400$ (*second row*) and $s^* = -2,200$ (*third row*)

$$\Sigma_{M_1} = \begin{bmatrix} 0.025 & 0 \\ 0 & 0.0125 \end{bmatrix}, \quad \Sigma_{M_2} = \begin{bmatrix} 0.025 & -0.0125 \\ -0.0125 & 0.0125 \end{bmatrix},$$

i.e., we assume that the two diameters are a priori independent (Σ_{M_1} case) or correlated (Σ_{M_2} case), respectively. The two prior distributions, as well as the resulting posterior distribution obtained for two different observed values $s^* = -1,400$ and $s^* = -2,200$ of the pressure drop are reported in Fig. 3. In the case at hand, we do not rely on the Metropolis-Hastings algorithm for the evaluation of the posterior distribution, since its expression can be computed explicitly. Thus, by computing a sample of 1,600 values of pressure drops on a uniform 40×40 grid on the (d_c, d_b) space, we obtain the posterior densities $\pi_{\mathbf{d},\text{post}}(\mathbf{d}|\mathbf{y}^*)$ represented in Fig. 3 in about 0.1 h, since any online evaluation of the reduced Navier-Stokes problem takes about 2.5 s.

Acknowledgements This work was partially funded by the European Research Council Advanced Grant “Mathcard, Mathematical Modelling and Simulation of the Cardiovascular System” (Project ERC-2008-AdG 227058), and by the Swiss National Science Foundation (Projects 122136 and 135444).

References

1. Kaipio, J., Somersalo, E.: Statistical and Computational Inverse Problems. Applied Mathematical Sciences, vol. 160. Springer, New York (2005)
2. Kolachalama, V., Bressloff, N., Nair, P.: Mining data from hemodynamic simulations via Bayesian emulation. *Biomed. Eng. Online* **6**(1), 47 (2007)
3. Lassila, T., Manzoni, A., Quarteroni, A., Rozza, G.: A reduced computational and geometrical framework for inverse problems in haemodynamics. Technical report MATHICSE 12.2011. **29**(7), 741–776 (2013). <http://mathicse.epfl.ch/>
4. Lassila, T., Manzoni, A., Quarteroni, A., Rozza, G.: Boundary control and shape optimization for the robust design of bypass anastomoses under uncertainty. *ESAIM: Math. Mod. Numer. Anal.* **47**(4), 1107–1131 (2013). <http://dx.doi.org/10.1051/m2an/2012059>
5. Lieberman, C., Willcox, K., Ghattas, O.: Parameter and state model reduction for large-scale statistical inverse problems. *SIAM J. Sci. Comput.* **32**(5), 2523–2542 (2010)
6. Manzoni, A., Quarteroni, A., Rozza, G.: Model reduction techniques for fast blood flow simulation in parametrized geometries. *Int. J. Numer. Methods Biomed. Eng.* **28**(6–7), 604–625 (2012)
7. Quarteroni, A., Rozza, G., Manzoni, A.: Certified reduced basis approximation for parametrized partial differential equations in industrial applications. *J. Math. Ind.* **1**, 3 (2011)