

Model order reduction by convex displacement interpolation

Simona Cucchiara¹, Angelo Iollo², Tommaso Taddei², Haysam Telib¹

¹ *Optimad Engineering, 10134 Torino, Italy*

simona.cucchiara@optimad.it, haysam.telib@optimad.it

² *Univ. Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, F-33400 Talence, France*

Inria Bordeaux Sud-Ouest, Team MEMPHIS, 33400 Talence, France,

angelo.iollo@inria.fr, tommaso.taddei@inria.fr

Abstract

We present a nonlinear interpolation technique for parametric fields that exploits optimal transportation of coherent structures of the solution to achieve accurate performance. The approach generalizes the nonlinear interpolation procedure introduced in [Iollo, Taddei, J. Comput. Phys., 2022] to multi-dimensional parameter domains and to datasets of several snapshots. Given a library of high-fidelity simulations, we rely on a scalar testing function and on a point set registration method to identify coherent structures of the solution field in the form of sorted point clouds. Given a new parameter value, we exploit a regression method to predict the new point cloud; then, we resort to a boundary-aware registration technique to define bijective mappings that deform the new point cloud into the point clouds of the neighboring elements of the dataset, while preserving the boundary of the domain; finally, we define the estimate as a weighted combination of modes obtained by composing the neighboring snapshots with the previously-built mappings. We present several numerical examples for compressible and incompressible, viscous and inviscid flows to demonstrate the accuracy of the method. Furthermore, we employ the nonlinear interpolation procedure to augment the dataset of simulations for linear-subspace projection-based model reduction: our data augmentation procedure is designed to reduce offline costs — which are dominated by snapshot generation — of model reduction techniques for nonlinear advection-dominated problems.

Keywords: model order reduction; nonlinear approximations.

1 Introduction

1.1 Nonlinear interpolation of parametric fields

Despite the many recent contributions, model order reduction (MOR) of parametric problems with compactly-supported features — such as shocks or shear layers — remains an outstanding task for state-of-the-art techniques due to the fundamental inadequacy of linear approximations. The aim of this work is to devise a general — i.e., independent of the underlying equations — interpolation technique for steady-state parametric problems, with emphasis on fluid dynamics applications.

During the past decade, several authors have proposed *Lagrangian methods* to deal with this class of problems [11, 18, 34, 36, 37, 40, 50, 48]. We denote by μ the vector of model parameters in the region $\mathcal{P} \subset \mathbb{R}^P$ and we denote by $\Omega \subset \mathbb{R}^d$ the open computational domain; then, we introduce the parametric field of interest $u : \Omega \times \mathcal{P} \rightarrow \mathbb{R}^D$ and the solution manifold $\mathcal{M} = \{u_\mu := u(\cdot; \mu) : \mu \in \mathcal{P}\}$. Lagrangian approximations rely on the ansatz

$$\hat{u}_\mu = \tilde{u}_\mu \circ \Phi_\mu^{-1}, \quad (1a)$$

where \tilde{u}_μ is a linear (or affine) approximation of the form

$$\tilde{u}_\mu(x) = \sum_{i=1}^n \hat{\omega}_\mu^i \zeta_i(x), \quad x \in \Omega, \mu \in \mathcal{P}, \quad (1b)$$

for proper choices of the weights $\hat{\omega}_\mu^1, \dots, \hat{\omega}_\mu^n$ and the parameter-independent fields $\zeta_1, \dots, \zeta_n : \Omega \rightarrow \mathbb{R}^D$, and $\Phi : \Omega \times \mathcal{P} \rightarrow \Omega$ is a suitably-chosen bijection that tracks the coherent structures of the solution; here, D denotes the number of state variables, while d is the spatial dimension. Lagrangian approaches are motivated by the observation that in many problems of interest coherent structures that are troublesome for linear approximations vary smoothly with the parameter and they hence can be tracked through a low-rank parameter-dependent mapping Φ .

In [19], we (Iollo, Taddei) proposed a general method dubbed convex displacement interpolation (CDI) that relies on optimal transportation to perform accurate nonlinear interpolations between solution snapshots; the approach was proposed for databases of two snapshots and one-dimensional parameter domains. Similarly to

Lagrangian approaches, CDI relies on the assumption that the location of coherent features of the solution field depends smoothly on the parameter; however, unlike in (1), it does not rely on the definition of a reference configuration where the location of the coherent features is (approximately) frozen. In this work, we discuss an extension of CDI to multi-dimensional parameter domains and to datasets of several snapshots.

1.2 Convex displacement interpolation

CDI exploits the standard offline/online decomposition to predict the state map $\mu \in \mathcal{P} \mapsto u_\mu \in \mathcal{M}$; Algorithm 1 summarizes the general procedure and highlights the key steps of the methodology that will be discussed in the next sections. Given the solution manifold \mathcal{M} , we introduce the training set $\mathcal{P}_{\text{train}} = \{\mu^k\}_{k=1}^{n_{\text{train}}} \subset \mathcal{P}$ and the dataset of solutions $\mathcal{D}_{\text{train}} = \{u_\mu : \mu \in \mathcal{P}_{\text{train}}\}$. During the offline (or learning) stage, for each element of $\mathcal{D}_{\text{train}}$, we determine a set of points $X_\mu^{\text{raw}} = \{x_{i,\mu}^{\text{raw}}\}_{i=1}^{N_\mu} \subset \bar{\Omega}$ that describe local features of the solution field that we wish to track: note that the sets $\{X_\mu^{\text{raw}} : \mu \in \mathcal{P}_{\text{train}}\}$ do not necessarily have the same number of elements and are not matched with each other. Then, we define the reference points $X^{\text{ref}} := \{x_i^{\text{ref}}\}_{i=1}^N$ and the new sensors $X_\mu := \{x_{i,\mu} = x_i^{\text{ref}} + v_{i,\mu}\}_{i=1}^N$ such that the set X_μ approximates — in a sense to be defined — X_μ^{raw} for all $\mu \in \mathcal{P}_{\text{train}}$. During the online stage, given the new parameter value $\mu \notin \mathcal{P}_{\text{train}}$, first, we apply a regression method to predict the new sensor locations $\hat{X}_\mu = \{\hat{x}_{i,\mu} = x_i^{\text{ref}} + \hat{v}_{i,\mu}\}_{i=1}^N$; second, we identify a set of κ nearest neighbors $\mathcal{P}_{\text{nn}}^\mu = \{\nu^i\}_{i=1}^\kappa \subset \mathcal{P}_{\text{train}}$; third, we define the mappings $\{\Phi_\nu : \nu \in \mathcal{P}_{\text{nn}}^\mu\}$ such that each mapping Φ_ν is bijective in Ω and $\Phi_\nu(\hat{x}_{i,\mu}) \approx \hat{x}_{i,\nu}$ for $i = 1, \dots, N$ and $\nu \in \mathcal{P}_{\text{nn}}^\mu$; fourth, we return the (generalized) CDI as

$$\hat{u}_\mu = \sum_{\nu \in \mathcal{P}_{\text{nn}}^\mu} \omega_\mu^\nu \tilde{u}_\nu, \quad \text{where } \tilde{u}_\nu = u_\nu \circ \Phi_\nu, \quad (2)$$

for a proper choice of the weights $\{\omega_\mu^\nu : \nu \in \mathcal{P}_{\text{nn}}^\mu\}$. Figure 1 illustrates the computational procedure for a simple one-dimensional problem.

Algorithm 1 : offline/online decomposition.

Offline stage

performed once

- 1: Generate the dataset $\mathcal{D}_{\text{train}} = \{u_\mu : \mu \in \mathcal{P}_{\text{train}}\}$.
- 2: Identify the point clouds $\{X_\mu^{\text{raw}} : \mu \in \mathcal{P}_{\text{train}}\}$. section 2.1
- 3: Define the template set X^{ref} and the sorted point clouds $\{X_\mu : \mu \in \mathcal{P}_{\text{train}}\}$. section 2.2

Online stage

performed for any $\mu \in \mathcal{P}$

- 1: Estimate the new points $\hat{X}_\mu = \{\hat{x}_{i,\mu}\}_{i=1}^N$. section 2.3
 - 2: Select the neighboring parameters $\mathcal{P}_{\text{nn}}^\mu = \{\nu^i\}_{i=1}^\kappa \subset \mathcal{P}_{\text{train}}$. section 2.4
 - 3: Compute the mappings Φ_ν based on \hat{X}_μ and \hat{X}_ν for all $\nu \in \mathcal{P}_{\text{nn}}^\mu$. section 2.5
 - 4: Compute the weights $\{\omega_\mu^\nu : \nu \in \mathcal{P}_{\text{nn}}^\mu\}$ and return the estimate (2). section 2.6
-




Figure 1: schematic of the CDI procedure. (a) dataset of snapshots $\{u_{\mu^1}, u_{\mu^2}\}$ and corresponding sensors (shock location), \hat{x}_μ denotes the predicted shock location for the parameter μ . (b) mapped snapshots $\{\tilde{u}_{\mu^1}, \tilde{u}_{\mu^2}\}$ in (2).

The practical implementation of CDI requires to carefully choose appropriate methods for each step of the offline and online stages of Algorithm 1.

Sensor selection There is a broad literature in fluid mechanics and scientific computing on the definition of quantitative criteria for the identification of coherent flow features such as shear layers, shocks, or vortices: representative examples include shock sensors for high-order methods in CFD [39, 42], and local criteria for vortex identification [21].

Point cloud matching The problem of determining sorted point clouds based on a dataset of unmatched point clouds and a template set is referred to as point-set registration (PSR, [17, 27, 28, 29, 56]): PSR methods implicitly introduce a global smooth mapping that coherently deforms the points of the template set X^{ref} to match — in a sense to be defined — the target cloud X_μ , for all $\mu \in \mathcal{P}_{\text{train}}$.

Regression The problem of estimating the sensor locations for new parameter values is a standard multi-target regression problem that can be handled through off-the-shelves machine learning routines including radial basis function (RBF, [54]) regression. As in non-intrusive MOR methods (e.g., [16]), we first apply proper orthogonal decomposition (POD, [47, 53]) to determine a low-rank representation of the sorted point clouds and then we apply RBF to the POD coefficients.

Registration The mapping generated by PSR methods does not fulfill the bijectivity condition $\Phi(\Omega) = \Omega$; in order to avoid extrapolation of full-order solutions outside the domain of definition, it is hence necessary to devise an efficient boundary-aware registration technique that ensures bijectivity in Ω . Note that the registration step should be significantly less expensive than full-order solves in order for the CDI method to be beneficial.

Choice of the weights The weights $\{\omega_\mu^\nu : \nu \in \mathcal{P}_{\text{nn}}^\mu\}$ should be chosen to ensure relevant properties of the interpolation method (cf. section 2.7). Since our emphasis is on datasets of very modest size, we rely on the inverse distance weighting (IDW, [46]) method.

1.3 Contributions and relation to previous works

The past decade has witnessed a surge in the development of nonlinear MOR methods to mitigate the Kolmogorov barrier [1, 4, 5, 25, 41]. As discussed in the previous section, our method shares important features with Lagrangian methods based on the ansatz (1); the ansatz (2) is also formally equivalent to front tracking methods proposed in [24, 44]. Note also that our approach relies on a piecewise-nonlinear approximation, which is motivated by the need to reduce the mesh interpolation and registration costs that scale linearly with κ and also by the particular strategy employed to define the weights $\{\omega_\mu^\nu : \nu \in \mathcal{P}_{\text{nn}}^\mu\}$. In this respect, the method is related to previously-proposed localized MOR techniques [2, 13].

The CDI procedure outlined in Algorithm 1 generalizes the method in [19]: in more detail, it reduces to the two-field CDI estimate for one-dimensional parameters and $\kappa = 2$, and a proper choice of the mappings $\{\Phi_\nu : \nu \in \mathcal{P}_{\text{nn}}^\mu\}$ (cf. Remark 2.1). We observe that, similarly to Lagrangian approximations, CDI relies on geometric mappings to track local features of the solution field and ultimately improve state predictions; on the other hand, CDI does not rely on the definition of a reference domain and the estimate is inherently Eulerian. We note that the definition of the template X^{ref} is instrumental to the construction of the sorted point clouds $\{X_\mu : \mu \in \mathcal{P}_{\text{train}}\}$, which enable the application of standard regression algorithms during the online stage for the prediction of \widehat{X}_μ . We also observe that during the online stage each mapping Φ_ν is computed based on the point clouds \widehat{X}_μ and \widehat{X}_ν (in lieu of X_ν): thanks to this choice, we prove that CDI is an interpolation method (cf. section 2.7).

In this work, we rely on well-established techniques for sensor selection, point cloud matching, multi-target regression, and weight selection; on the other hand, we emphasize that the development of boundary-aware registration techniques to determine bijective mappings in the domain Ω of interest is an active research area with many recent contributions. In this work, we propose a new elasticity-based method that borrows ideas from penalization methods for PDE discretization on Cartesian meshes (see, e.g., [10, 3]); we also note that the method shares important features with morphing techniques for mesh adaptation [51, 52] and image processing [7, 8]. We further consider the optimization-based method first proposed in [48] and further developed in several subsequent works.

The ultimate goal of this work is to devise a nonlinear interpolation strategy that can cope with datasets of very modest size. The evaluation of the CDI estimate (2) is not necessarily *online-efficient* — that is, the prediction cost scales with the cost of the underlying high-fidelity (HF) mesh — but it is significantly less expensive than the solution to the HF model. In this respect, CDI might be employed to initialize nonlinear

solvers for nonlinear problems [22], or to enrich the snapshot dataset that can later be used to train a projection-based reduced-order model (PROM) — in effect, *data augmentation* [6]. We discuss the application of CDI to data augmentation through the vehicle of a parametric two-dimensional inviscid flow problem.

The outline of the paper is as follows. In section 2, we present the methods for each step of the offline and online stages of Algorithm 1, and we show relevant properties of the interpolation procedure. In section 3, we present a motivating example that illustrates the many features of the approach in a simplified setting. In section 4 we present extensive numerical results for three different test cases that cover compressible and incompressible, viscous and inviscid flows. In section 5, we illustrate the application to data augmentation. Section 6 completes the paper.

2 Methodology

2.1 Sensor (or feature) selection

We rely on the definition of a problem-dependent *scalar testing function* \mathcal{F} to identify unmatched point clouds that can be exploited to derive out-of-sample predictions (cf. section 2.3). Given the domain Ω , we define a discrete set of points $P_{\text{hf}} \subset \bar{\Omega}$ and we compute $\{\mathcal{F}(x, u_{\mu}^{\text{hf}}) : x \in P_{\text{hf}}\}$ for all $\mu \in \mathcal{P}_{\text{train}}$, where u_{μ}^{hf} denotes the high-fidelity (HF) estimate of the solution field for $\mu \in \mathcal{P}$; then, we set

$$X_{\mu}^{\text{raw}} = \{x_{i,\mu}^{\text{raw}}\} = \{x \in P_{\text{hf}} : \mathcal{F}(x, u_{\mu}^{\text{hf}}) \geq t_{\mu}\} \quad (3)$$

for some threshold t_{μ} that is specified below.

The choice of the scalar testing function is highly problem-dependent and requires a deep understanding of the underlying physical model and of the features we wish to track. For compressible flow problems in the transonic or supersonic regime, we rely on the Ducros sensor [35, 39] to identify the shocks: we denote by $P_{\text{hf}} = \{x_k^{\text{hf}}\}_{k=1}^{N_e}$ the cell centers of the HF mesh and we compute

$$\mathcal{F}(x_k^{\text{hf}}; u_{\mu}^{\text{hf}}) = \sup_{x \in \mathbf{D}_k} |\phi(x; u_{\mu}^{\text{hf}})| \quad (4a)$$

where \mathbf{D}_k denotes the k -th cell of the mesh with center x_k^{hf} , ϕ is the Ducros sensor

$$\phi(x; u) := \frac{(-\nabla \cdot v)^+}{\sqrt{(\nabla \cdot v)^2 + \|\nabla \times v\|_2^2 + a^2}} \frac{\|\nabla p\|_2}{p + \epsilon} \|v\|_2, \quad (4b)$$

v is the velocity field, p is the pressure, a is the sound speed and $\epsilon > 0$ is a user-defined tolerance that is set equal to 0.01. Finally, the threshold t_{μ} in (3) is chosen equal to the γ_{thr} quantile over the training set:

$$t_{\mu} = \text{quantile}(\{\mathcal{F}(x, u_{\mu}^{\text{hf}}) : x \in P_{\text{hf}}\}, \gamma_{\text{thr}}). \quad (4c)$$

For the incompressible flow problem of section 4, we rely on the isolines of the streamfunction to identify the recirculation area, for two-dimensional channel flows. In more detail, we define the streamline function

$$\Psi_{\mu}(x = [x_1, x_2]) = \int_{y_{\text{btm}}(x_1)}^{x_2} (u_{\mu}(x_1, s))_1 ds, \quad (5)$$

where $y_{\text{btm}}(x_1) = \inf\{x_2 : [x_1, x_2] \in \Omega\}$ denotes the bottom boundary of the channel; then, we set $\mathcal{F}(x, u_{\mu}^{\text{hf}}) = -\Psi_{\mu}(x)$ and $t_{\mu} \equiv 0$; that is,

$$X_{\mu}^{\text{raw}} = \{x \in P_{\text{hf}} : \Psi_{\mu}(x) \leq 0\}.$$

We remark that we also tried to exploit the Q-criterion and the triple decomposition of relative motion (TDM) [23]: however, in our experience, these two indicators were not able to distinguish between pure shearing motions and the actual swirling motion of the vortex.

2.2 Point cloud matching

The point clouds in (3) are not necessarily of the same size and are not sorted; therefore, we resort to PSR to find matched point clouds:

$$X_{\mu} = \text{PSR}(X^{\text{ref}}, X_{\mu}^{\text{raw}}), \quad X^{\text{ref}} := X_{\mu^{\text{ref}}}^{\text{raw}}, \quad \mu^{\text{ref}} \in \mathcal{P}_{\text{train}}. \quad (6)$$

Given the point clouds $X = \{x_i\}_{i=1}^N$ and $Y = \{y_j\}_{j=1}^M$, the problem of PSR consists in finding a map $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that (approximately) minimizes the distance

$$\text{dist}(Y, T(X)) = \max_{y \in Y} \left(\min_{x \in X} \|y - T(x)\|_2 \right)$$

in a suitable model class; the output of the algorithm is the deformed set of points $T(X)$. A broad range of PSR methods — including the approach considered in this work — relies on a probabilistic interpretation of the problem: these methods rely on the assumption that $\{x_i\}_{i=1}^N$ and $\{y_j\}_{j=1}^M$ are independent identically distributed (iid) samples from a given probability distribution with unknown parameters. As discussed in the introduction, Algorithm 1 can cope with any PSR technique that can be written in the form (6); we refer to a future work for the development of a specialized PSR procedure for MOR. In the remainder of this section, we review the PSR algorithm employed in the numerical experiments of sections 4 and 5.

We rely on Gaussian-based PSR. Given the point cloud $X = \{x_i\}_{i=1}^N$, we resort to maximum likelihood estimation (MLE, see, e.g., [45, Chapter 8]) to estimate mean and covariance matrix:

$$\mu_X = \frac{1}{N} \sum_{i=1}^N x_i, \quad \Sigma_X = \frac{1}{N} \sum_{i=1}^N (x_i - \mu_X)(x_i - \mu_X)^\top, \quad (7)$$

where $(\cdot)^\top$ is the transpose operator, and we define the Gaussian model $g_X = \mathcal{N}(\mu_X, \Sigma_X)$. Similarly, given the point cloud $Y = \{y_j\}_{j=1}^M$, we define the MLE estimates μ_Y, Σ_Y and the corresponding Gaussian model $g_Y = \mathcal{N}(\mu_Y, \Sigma_Y)$. Finally, we apply the optimal transport map (cf. [31]) from g_X to g_Y

$$T_{X,Y}(x) = \mu_Y + \Sigma_X^{-1/2} \left(\Sigma_X^{1/2} \Sigma_Y \Sigma_X^{1/2} \right)^{1/2} \Sigma_X^{-1/2} (x - \mu_X). \quad (8)$$

to define the deformed point cloud $\tilde{X} = \{\tilde{x}_i = T_{X,Y}(x_i)\}_{i=1}^N$.

2.3 Regression of sorted point clouds

We rely on RBF regression to determine the approximation $\mu \in \mathcal{D} \mapsto \hat{X}_\mu \in \mathbb{R}^{N \times d}$. Given the sorted point clouds $\{X_\mu : \mu \in \mathcal{D}_{\text{train}}\}$, we first apply POD to determine the equivalent representation

$$X_{\mu^k} = \sum_{i=1}^M Z_i \beta_k^i, \quad \text{with } Z_1, \dots, Z_M \in \mathbb{R}^{N \times d}, \quad \beta_1^1, \dots, \beta_{n_{\text{train}}}^M \in \mathbb{R},$$

and $M = \min\{N \cdot d, n_{\text{train}}\}$. Then, we apply RBF regression to each coefficient of the expansion separately,

$$\hat{\beta}^i = \arg \min_{\beta \in \mathfrak{H}_\phi} \lambda \|\beta\|_{\mathfrak{H}_\phi}^2 + \sum_{k=1}^{n_{\text{train}}} (\beta(\mu^k) - \beta_k^i)^2, \quad i = 1, \dots, M. \quad (9)$$

Here, \mathfrak{H}_ϕ denotes the native space associated with the kernel ϕ and $\lambda > 0$ is a regularization coefficient. We rely on the Gaussian kernel, and we exploit cross-validation to select λ and the kernel width; furthermore, we estimate the out-of-sample R-squared indicator and we keep the coefficients of the expansion for which the latter exceeds the threshold 0.5. In conclusion, we obtain the estimate

$$\hat{X}_\mu = \sum_{i=1}^m Z_i \hat{\beta}_\mu^i, \quad (10)$$

with $m \leq M$.

2.4 Choice of the nearest neighbors

Given the parameter value $\mu \in \mathcal{D}$, we select the parameters $\nu^1, \dots, \nu^\kappa \in \mathcal{D}_{\text{train}}$ that minimize the Euclidean distance $\text{dist}(\mu, \nu) = \|\mu - \nu\|_2$. The Euclidean distance is invariant under rototranslations of the parameters: as discussed in section 2.7, this property is key to prove the frame indifference of the CDI procedure.

2.5 Boundary-aware registration

The mapping $T_{X,Y}$ in (8) does not preserve the boundary of the domain: for each $\nu \in \mathcal{D}_{\text{in}}^\mu$, we should hence find a mapping $\Phi_\nu : \Omega \rightarrow \Omega$ that is bijective in Ω and satisfies $\Phi_\nu(x_i^{\text{ref}}) \approx \hat{x}_{i,\nu}$ for $i = 1, \dots, N$. Below, we present two registration techniques to accomplish this task. Note that we here decouple the problem of PSR (*point cloud matching*) from the problem of determining a bijective mapping in Ω : this enables the use of off-the-shelves PSR routines for the matching problem. The development of a general registration procedure for arbitrary domains is the subject of ongoing research and a thorough review of the subject is beyond the scope of the present paper. It is straightforward to verify that both registration methods return the identity map if $x_i^{\text{ref}} = \hat{x}_{i,\mu}$ for $i = 1, \dots, N$: this observation is instrumental to prove the interpolation property of CDI in section 2.7.

2.5.1 Elasticity-based registration

We denote by $\tilde{v} = T_{X,Y} - \text{id}$ the spatial displacement field associated with the PSR algorithm where id is the identity map; we assume that \tilde{v} is defined over \mathbb{R}^d — this assumption holds for several PSR algorithms including the probabilistic method (8) and its generalizations based on Gaussian mixture models (e.g., [38]). Given the reference point cloud $\{x_i\}_{i=1}^N$, we define the region A_η such that

$$A_\eta = \bigcup_{i=1}^N \mathcal{B}_\eta(x_i), \quad (11a)$$

where $\mathcal{B}_\eta(x)$ denotes the ball of radius η centered in x ; we introduce the regularized signed distance function

$$H_{\eta,\delta}(x) = f_\delta(\phi(x, A_\eta)), \quad \text{with} \quad \begin{cases} f_\delta(t) = \frac{1}{2} \left(\tanh\left(\frac{t}{\delta}\right) + 1 \right), \\ \phi(x, A_\eta) = \begin{cases} \text{dist}(x, \partial A_\eta) & x \in A_\eta \\ -\text{dist}(x, \partial A_\eta) & x \in \partial A_\eta. \end{cases} \end{cases} \quad (11b)$$

Finally, we define the modified field

$$\begin{cases} \Delta v + \nabla(\nabla \cdot v) = \frac{1}{\epsilon}(v - \tilde{v})H_{\eta,\delta} & x \in \Omega, \\ \mathbf{t} \cdot (\nabla v \mathbf{n}) = 0, \quad v \cdot \mathbf{n} = 0 & x \in \partial\Omega, \end{cases} \quad (11c)$$

where \mathbf{n} is the normal vector and \mathbf{t} denotes the tangent vector to $\partial\Omega$. As in [3], the right-hand-side of (11)₁ weakly enforces the condition $v = \tilde{v}$ in the region A_η in (11a); on the other hand, the condition $v \cdot \mathbf{n} = 0$ ensures the non-penetration condition, for $\epsilon \ll 1$.

We observe that the field v depends on the hyper-parameters ϵ, η, δ ; we discuss their choice in the numerical experiments of section 4. Finally, we define the mapping Φ as follows:

$$\Phi(\xi) := \mathfrak{X}(\xi, 1), \quad \text{where} \quad \begin{cases} \partial_t \mathfrak{X}(\xi, t) = v(\mathfrak{X}(\xi, t)) & t \in (0, 1), \\ \mathfrak{X}(\xi, 0) = \xi. \end{cases} \quad (12)$$

We observe that the field v corresponds to the Eulerian flow velocity, while $\mathfrak{X}(\xi, \cdot) : (0, 1) \rightarrow \mathbb{R}^d$ describes the temporal evolution of the particle located at ξ at time $t = 0$. The condition $v \cdot \mathbf{n} = 0$ ensures that $\mathfrak{X}(\xi, t) \in \Omega$ for all $t > 0$ at the continuous level: in practice, we discretize (12) using an explicit Euler scheme and we properly choose a time step Δt to ensure that the geometry error is below a given threshold. We remark that our approach does not ensure that $\Phi(x_i) = y_i$ for $i = 1, \dots, N$ nor it ensures the bijectivity of the mapping in Ω ; nevertheless, the numerical experiments of this manuscript show that in practice the method provides proper deformations of the domain Ω that approximately satisfy the interpolation condition $\Phi(x_i) = y_i$ for $i = 1, \dots, N$. The development of more effective elasticity-based registration techniques as well as their rigorous analysis is the subject of ongoing research.

2.5.2 Optimization-based registration

An alternative approach consists in formulating the problem of registration as a minimization problem of the form

$$\min_{\Phi \in \mathcal{W}_\Omega} \frac{1}{N} \sum_{i=1}^N \|\Phi(x_i^{\text{ref}}) - \hat{x}_{i,\mu}\|_2^2 + \mathfrak{P}(\Phi), \quad (13)$$

that can be solved using a gradient-based (quasi-Newton) method. The optimization statement in (13) depends on the choice of the penalty term \mathfrak{P} and of the search space \mathcal{W}_Ω : the former should enforce local bijectivity in Ω and also promote the smoothness (in a Sobolev sense) of the mapping; the latter should ensure that locally-bijectivity — i.e., $\det(\nabla \Phi) > 0$ in $\bar{\Omega}$ — implies global bijectivity in Ω and should be minimal for $\Phi = \text{id}$ to ensure the interpolation property of CDI. We refer to [48, 15] for a detailed presentation of the registration procedure for two-dimensional domains: the generalization of the approach to arbitrary domains is the subject of ongoing research.

2.6 Choice of the weights

We determine the weights $\{\omega_\mu^\nu : \nu \in \mathcal{P}_{\text{nn}}^\mu\}$ using the inverse distance weighting (IDW, [46]) method. Given the nearest-neighbor training parameters $\nu^1, \dots, \nu^\kappa \in \mathcal{P}_{\text{train}}$ and the new parameter $\mu \in \mathcal{P}$, we consider

$$\left\{ \begin{array}{l} \omega_\mu^\nu = \frac{\bar{\omega}_\mu^\nu}{\sum_{\nu' \in \mathcal{P}_{\text{nn}}^\mu} \bar{\omega}_\mu^{\nu'}}, \quad \text{with } \bar{\omega}_\mu^{\nu'} = \frac{1}{\text{dist}^p(\mu, \nu')}, \quad \text{if } \mu \notin \mathcal{P}_{\text{train}}; \\ \omega_\mu^\nu = \begin{cases} 1 & \text{if } \mu = \nu \\ 0 & \text{otherwise} \end{cases} \quad \text{if } \mu \in \mathcal{P}_{\text{train}}. \end{array} \right. \quad (14a)$$

The algorithm can cope with arbitrary choices of the distance $\text{dist} : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}_+$. We here consider the Euclidean distance

$$\text{dist}(\mu, \nu) = \|\mu - \nu\|_2. \quad (14b)$$

The power parameter $p \geq 1$ influences the relative importance of neighboring parameters: large values of p assign greater influence to values closest to the interpolated parameter. Since we consider relatively small values of κ , we expect that the sensitivity with respect to p is limited. In the numerical experiments, we set $p = 1$.

Remark 2.1. Let $\mathcal{P} \subset \mathbb{R}$ and let $\kappa = 2$. Given $\mu \in [\nu^0, \nu^1]$ with $\nu_0, \nu_1 \in \mathcal{P}_{\text{train}}$, we define $s = \frac{\mu - \nu^0}{\nu^1 - \nu^0}$; then, exploiting (14a) with $p = 1$, we obtain

$$\hat{u}_\mu = (1 - s)\tilde{u}_{\nu^0} + s\tilde{u}_{\nu^1},$$

which corresponds to the two-field CDI proposed in [19], for a proper choice of the mappings.

2.7 Properties of the interpolation procedure

The CDI (2) with the choice of the weights (14) has four noteworthy properties that are of interest for interpolation of fluid mechanics fields.

1. *Interpolation.* Let $\mu \in \mathcal{P}_{\text{train}}$. Then, $\hat{u}_\mu = u_\mu$.

Proof. The proof exploits the discussion of the previous sections. Recalling the choice of $\mathcal{P}_{\text{nn}}^\mu$ in section 2.4, we find $\mu \in \mathcal{P}_{\text{nn}}^\mu$. Then, exploiting (14a), we find $\hat{u}_\mu = \tilde{u}_\mu = u_\mu \circ \Phi_\mu$. Recalling the correction in section 2.3, we have $\hat{X}_\mu = X_\mu$ and hence (cf. section 2.5) $\Phi_\mu = \text{id}$. \square

2. *Maximum principle.* Let $b \in \mathbb{R}^D$; then, $\sup_{x \in \Omega} |b \cdot \hat{u}_\mu(x)| \leq \sup_{x \in \Omega, \nu \in \mathcal{P}} |b \cdot u_\nu(x)|$ for all $\mu \in \mathcal{P}$.

Proof. We have

$$\begin{aligned} |b \cdot \hat{u}_\mu(x)| &\stackrel{(i)}{\leq} \sum_{\nu \in \mathcal{P}_{\text{nn}}^\mu} \omega_\mu^\nu |b \cdot \tilde{u}_\nu(x)| \stackrel{(ii)}{\leq} \max_{\nu \in \mathcal{P}_{\text{nn}}^\mu} \left(\sup_{x \in \Omega} |b \cdot \tilde{u}_\nu(x)| \right) \\ &\stackrel{(iii)}{=} \max_{\nu \in \mathcal{P}_{\text{nn}}^\mu} \left(\sup_{x \in \Omega} |b \cdot u_\nu(x)| \right) \leq \sup_{\nu \in \mathcal{P}, x \in \Omega} |b \cdot u_\nu(x)|. \end{aligned}$$

In (i), we applied the definition of CDI, the triangle inequality and the positivity of the weights $\{\omega_\mu^\nu\}_\nu$; in (ii), we used the fact that $\sum_{\nu \in \mathcal{P}_{\text{nn}}^\mu} \omega_\mu^\nu = 1$; in (iii), we used that $\Phi_\nu(\Omega) = \Omega$ for all $\nu \in \mathcal{P}$. \square

3. *Minimum principle.* Let $b \in \mathbb{R}^D$; then, $\inf_{x \in \Omega} b \cdot \hat{u}_\mu(x) \geq \inf_{x \in \Omega, \nu \in \mathcal{P}} b \cdot u_\nu(x)$ for all $\mu \in \mathcal{P}$.

Proof. Given $x \in \Omega, \mu \in \mathcal{P}$, we have indeed that $\omega_\mu^\nu \geq 0$, $\sum_{\nu \in \mathcal{P}_{\text{nn}}^\mu} \omega_\mu^\nu = 1$, and $b \cdot \tilde{u}_\nu(x) \geq c := \inf_{x \in \Omega, \nu \in \mathcal{P}} (b \cdot u_\nu(x))$ for all $\nu \in \mathcal{P}_{\text{nn}}^\mu$. Therefore,

$$b \cdot \hat{u}_\mu(x) = \sum_{\nu \in \mathcal{P}_{\text{nn}}^\mu} \omega_\mu^\nu (b \cdot \tilde{u}_\nu(x)) \geq c \sum_{\nu \in \mathcal{P}_{\text{nn}}^\mu} \omega_\mu^\nu = c.$$

\square

4. *Frame indifference.* Given the parametric field $\mu \mapsto u_\mu$, consider the rototranslation $p(\mu) = R\mu + b$ for some rotation matrix R and vector $b \in \mathbb{R}^P$, and the parametric field $p \mapsto v_p = u_{\mu(p)}$. Then, $\hat{u}_{p(\mu)} = \hat{u}_\mu$ for all $\mu \in \mathcal{P}$.

Proof. It suffices to check that all the steps of the CDI procedure are invariant under roto-translations. First, we observe that RBF interpolation/regression is invariant under roto-translations (cf. [54]). Second, since the choice of the neighboring elements and the choice of the CDI weights (14) are both based on the Euclidean distance, we find that both steps also invariant under roto-translations. Finally, the registration procedure is independent of the parameterization considered and in particular is invariant under roto-translations. \square

The maximum and minimum principles can be used to ensure that the CDI prediction is physically meaningful for out-of-sample parameters. To provide a concrete reference, if we apply CDI to the solution to the Euler or Navier-Stokes equations in primitive variables, we can guarantee that density and pressure are strictly positive and do not exceed the maximum value attained in the training set. Similarly, since the mappings $\{\Phi_\nu\}_\nu$ in (2) are bijective in Ω , if the far-field conditions are parameter-independent and constant on $\partial\Omega$, we can exploit the same argument to prove that the CDI is consistent with the far-field conditions.

3 Motivating example

3.1 Model problem

We consider the parametric problem:

$$\begin{cases} -\frac{\partial^2 u_\mu}{\partial x^2} = f_\mu & \text{in } \Omega = (-1, 1), \\ u_\mu(-1) = u_\mu(1) = 0; \end{cases} \quad f_\mu(x) = \frac{1}{\sigma} \exp\left(\frac{(x-\mu)^2}{\sigma^2}\right), \quad (15)$$

with $\mu \in \mathcal{P} = [-0.9, 0.9]$ and $\sigma > 0$. We introduce the space $\mathcal{U} = H_0^1(\Omega)$ endowed with the inner product

$$(w, v) = \int_{-1}^1 \partial_x w \partial_x v + wv \, dx, \quad (16a)$$

and the induced norm $\|\cdot\| = \sqrt{(\cdot, \cdot)}$. Exploiting integration by part, we find the variational problem associated with (15):

$$\text{find } u_\mu \in \mathcal{U} : a(u_\mu, v) = F_\mu(v) \quad \forall v \in \mathcal{U}, \quad (16b)$$

where

$$a(w, v) = \int_{-1}^1 \partial_x w \partial_x v \, dx, \quad F_\mu(v) = \int_{-1}^1 f_\mu v \, dx. \quad (16c)$$

By tedious but straightforward calculations, we can show that

$$\frac{\pi}{2 + \pi} \|u\|^2 \leq a(u, u) \leq \|u\|^2 \quad \forall u \in \mathcal{U}. \quad (17)$$

3.2 Linear method

We denote by $\mathcal{Z} \subset \mathcal{U}$ an n -dimensional linear subspace of \mathcal{U} ; we further introduce the Galerkin reduced-order model:

$$\text{find } \hat{u}_\mu \in \mathcal{Z} : a(\hat{u}_\mu, v) = F_\mu(v) \quad \forall v \in \mathcal{Z}. \quad (18)$$

Exploiting C ea Lemma for symmetric coercive problems (e.g., [43, Lemma 4.3]) and the estimate (17), we find

$$\|\hat{u}_\mu - u_\mu\| \leq \sqrt{1 + \frac{2}{\pi} \min_{v \in \mathcal{Z}} \|v - u_\mu\|}, \quad \forall \mu \in \mathcal{P}. \quad (19)$$

We conclude that for this model problem Galerkin projection achieves near optimal performance over the entire parameter domain, for any choice of the reduced space \mathcal{Z} .

We generate $n_{\text{train}} = 15$ snapshots for equispaced parameter values $\mathcal{P}_{\text{train}} \subset \mathcal{P}$ and for two values of σ , $\sigma = 10^{-1}$ and $\sigma = 10^{-3}$. For each value of σ , we assess performance of the POD Galerkin ROM (18) for three values of the reduced basis size n , $n = 5, 10, 15$. Figure 2(a) shows the behavior of the relative H^1 error over \mathcal{P} for three choices of n for $\sigma = 10^{-1}$, while Figure 2(b) shows the solution and the Galerkin approximations for the parameter that maximizes the prediction error for $n = 15$. Figure 3 replicates the results for $\sigma = 10^{-3}$.

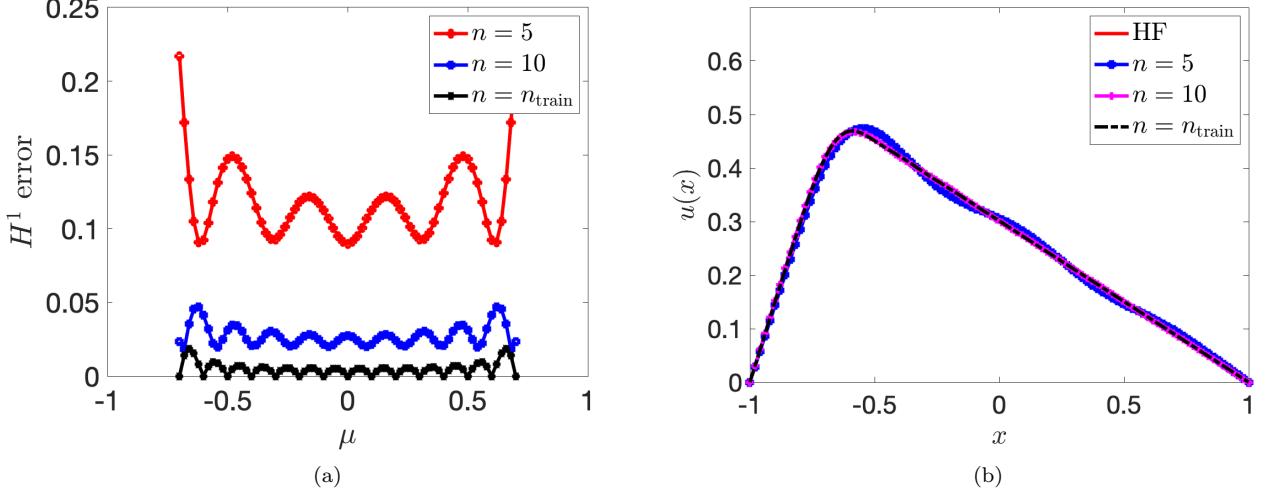


Figure 2: motivating example; linear ROM for $\sigma = 10^{-1}$. (a) H^1 relative error. (b) comparison of HF and ROM solutions for the parameter that maximizes the error for $n = n_{\text{train}}$.

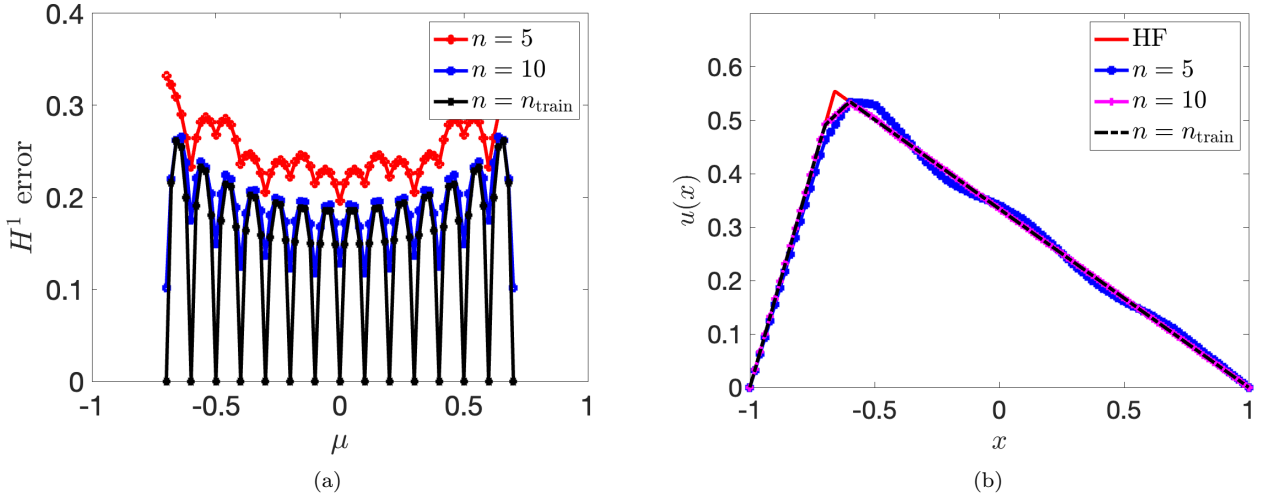


Figure 3: motivating example; linear ROM for $\sigma = 10^{-3}$. (a) H^1 relative error. (b) comparison of HF and ROM solutions for the parameter that maximizes the error for $n = n_{\text{train}}$.

We observe that the ROM for $n = n_{\text{train}}$ provides exact reconstructions for in-sample parameters: this immediately follows from (19). We note, however, that for $\sigma = 10^{-3}$ the error is large — in the order of 20% — and does not decrease as we increase n from $n = 10$ to $n = 15$. We hence conclude that the available snapshot set is not sufficient to achieve an accurate *linear* representation of the elements of the solution manifold.

We draw two important conclusions. First, the inadequacy of linear approximations primarily depends on the geometry of the solution manifold — which depends on the PDE model and on the parameterization — rather than on the class of PDEs of interest (hyperbolic, elliptic, parabolic); on the other hand, the PDE model affects the sub-optimality (cf. (19)) of Galerkin projection. Second, we observe that several nonlinear MOR methods (e.g., [4, 5]) seek state estimates that are nonlinear functions of the generalized coordinates but, nevertheless, belong to the span of the training snapshots: these methods hence do not overcome the issue of the inaccuracy of the snapshot set.

3.3 Convex displacement interpolation

Figures 4 and 5 show the performance of CDI for the model problem (16) for $\sigma = 10^{-1}$ and $\sigma = 10^{-3}$: Figures 4(a) and 5(a) show the relative H^1 error, while Figures 4(b) and 5(b) show the behavior of the solution and its estimate for the parameter that maximizes the prediction error. We here set $X_\mu = \{\mu\}$, we consider $\kappa = 2$ and

we define the bijective maps Φ_ν for $\nu \in \mathcal{D}_{\text{nn}}^\mu$ so that $\Phi_\nu(\Omega) = \Omega$, Φ_ν is piecewise linear in the intervals $(-1, \mu)$ and $(\mu, 1)$ and $\Phi_\nu(\mu) = \nu$. We notice that the error is below 0.5% for both values of σ .




Figure 4: motivating example; CDI for $\sigma = 10^{-1}$. (a) H^1 relative error. (b) comparison of HF and ROM solutions for the parameter that maximizes the prediction error.




Figure 5: motivating example; CDI for $\sigma = 10^{-3}$. (a) H^1 relative error. (b) comparison of HF and ROM solutions for the parameter that maximizes the prediction error.

3.4 Data augmentation

As final test, we perform Galerkin projection based on the parameter-dependent reduced space

$$\mathcal{Z}_\mu = \text{span} \{ \widehat{u}_\mu^{\text{cdi}}, u_{\nu^1}, u_{\nu^2} \}, \quad \text{with } \mathcal{D}_{\text{nn}}^\mu = \{ \nu^1, \nu^2 \}.$$

Figures 6 and 7 show the results: we observe that the error is roughly three times lower than the error associated with CDI. This result suggests the potential of data augmentation for Galerkin ROMs: nonlinear interpolation is employed to rapidly enhance the accuracy of the reduced space, while Galerkin projection is employed to ensure stable and near-optimal estimates. We further investigate the combination of CDI and data augmentation in the next section.




Figure 6: motivating example; Galerkin ROM with data augmentation for $\sigma = 10^{-1}$. (a) H^1 relative error. (b) comparison of HF and ROM solutions for the parameter that maximizes the prediction error.

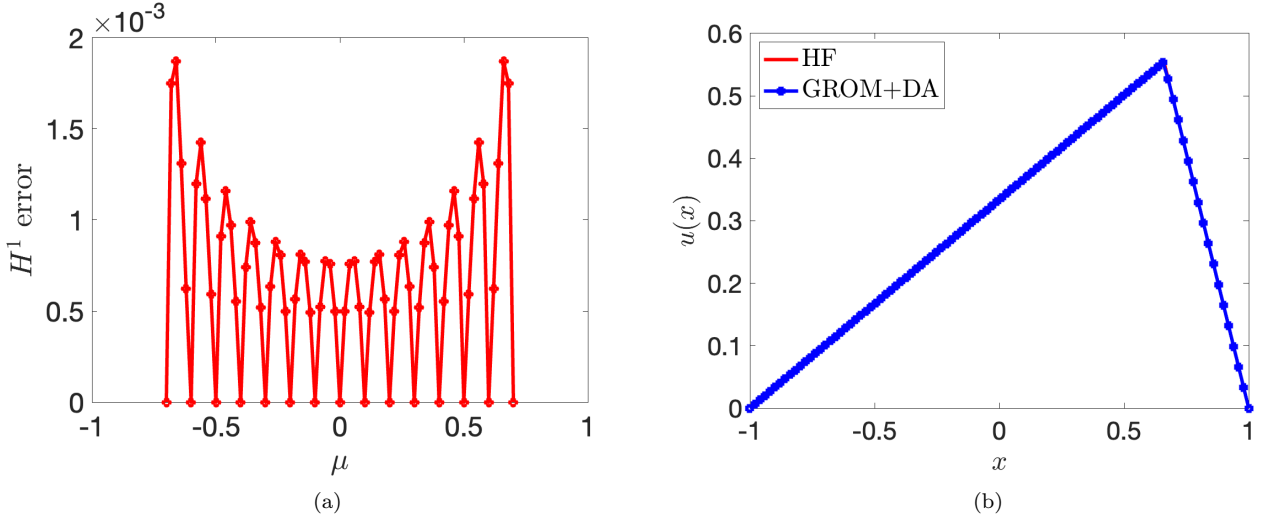


Figure 7: motivating example; Galerkin ROM with data augmentation for $\sigma = 10^{-3}$. (a) H^1 relative error. (b) comparison of HF and ROM solutions for the parameter that maximizes the prediction error.

4 Nonlinear interpolation of parametric flow fields

In this section, we illustrate the performance of our nonlinear interpolation procedure for three model problems. We rely on the elasticity-based approach for registration (cf. section 2.5.1): in all our tests, we set $\epsilon = 10^{-8}$, $\delta = 50$, and $\eta = 10^{-2}$ in (11a) and (11c); we further consider the time step $\Delta t = 5 \cdot 10^{-3}$ for the integration of (12). We study the sensitivity of the interpolation parameter s , for datasets with two parameters μ^0, μ^1 and $P = 1$. Given the point clouds $X_i = \{x_j^i\}_{j=1}^N \subset \Omega$ for $i = 0, 1$, we define $X(s) = \{(1-s)x_j^0 + sx_j^1\}_{j=1}^N$ and the mappings $\Phi_i(s)$ such that $\Phi_i(X(s), s) \approx X_i$ for $i = 0, 1$; finally, we define the two-field CDI

$$\hat{u}(s) = (1-s)u_0 \circ \Phi_0(s) + su_1 \circ \Phi_1(s). \quad (20)$$

and the linear counterpart (convex interpolation, CI)

$$\hat{u}(s) = (1-s)u_0 + su_1. \quad (21)$$

In the absence of prior information, we here set $s = \frac{\mu - \mu_0}{\mu_1 - \mu_0}$ (“linear relation”, cf. Remark 2.1); however, we also investigate the behavior of the optimal choice of s as a function of μ ,

$$s_\mu^{\text{opt}} = \arg \min_{s \in [0,1]} \|\hat{u}(s) - u_\mu\|_*, \quad (22)$$

where the norm $\|\cdot\|_*$ is introduced below for each model problem.

4.1 Compressible flow past a RAE2822 airfoil

4.1.1 Model problem

We study the compressible viscous flow past a RAE2822 airfoil for varying inflow Mach number Ma_∞ and angle of attack α . We model the flow using the compressible Reynolds-averaged Navier-Stokes equations with Menter’s SST closure model [33]; we rely on the open-source solver SU2 [12] for snapshot generation. In the remainder, we apply the nonlinear interpolation procedure to estimate the pressure coefficient

$$C_p = \frac{p - p_\infty}{\frac{1}{2}\rho_\infty \|v_\infty\|_2^2}, \quad (23)$$

where p, ρ, v denote flow pressure, density and velocity, and the subscript $(\bullet)_\infty$ denotes the free-stream value of the flow quantity \bullet .

4.1.2 One-dimensional parameterization

We first study variations of the free-stream Mach number in the range $[0.81, 0.85]$ for $\alpha = 2^\circ$. Figure 8(a) shows the computational domain; Figures 8(b) and (c) illustrate the behavior of the pressure coefficients for $\text{Ma}_\infty = 0.81$ and $\text{Ma}_\infty = 0.85$. We observe that the solution develops two shocks on the lower and upper sides of the airfoil whose position and intensity strongly depend on the parameter value.




Figure 8: Compressible flow past a RAE2822 airfoil. (a) computational domain. (b)-(c) pressure coefficients for $\text{Ma}_\infty = 0.81$ and $\text{Ma}_\infty = 0.85$, $\alpha = 2^\circ$.

We rely on the Ducros sensor (4) to identify the markers, with $\gamma_{\text{thr}} = 99.6\%$ (cf. (4c)). Figure 9 shows the marked cells of the mesh for two parameter values. Since the flow exhibits two distinct coherent structures (shocks), we apply the procedure of section 2.2 to the lower and upper sides of the domain separately (see also [19, section 5.2.2]).




Figure 9: compressible flow past a RAE2822 airfoil. Marked cells for $Ma_\infty = 0.81$ and $Ma_\infty = 0.85$, $\alpha = 2^\circ$.

Figure 10 investigates the performance of CDI (20) and CI (21) over $Ma_\infty \in [0.81, 0.85]$ for linear choice of the interpolation parameter s on the boundary Γ_2 . We observe that CDI preserves the structure of the solution and is also significantly more accurate. Figure 11 compares the behavior of the optimal choice of s (“experimental”) with the linear relationship — optimality is measured with respect to the $L^2(\Gamma_2)$ norm over the airfoil (“integral error”). We observe that for this model problem s_μ^{opt} varies linearly with respect to the parameter $\mu = Ma_\infty$.

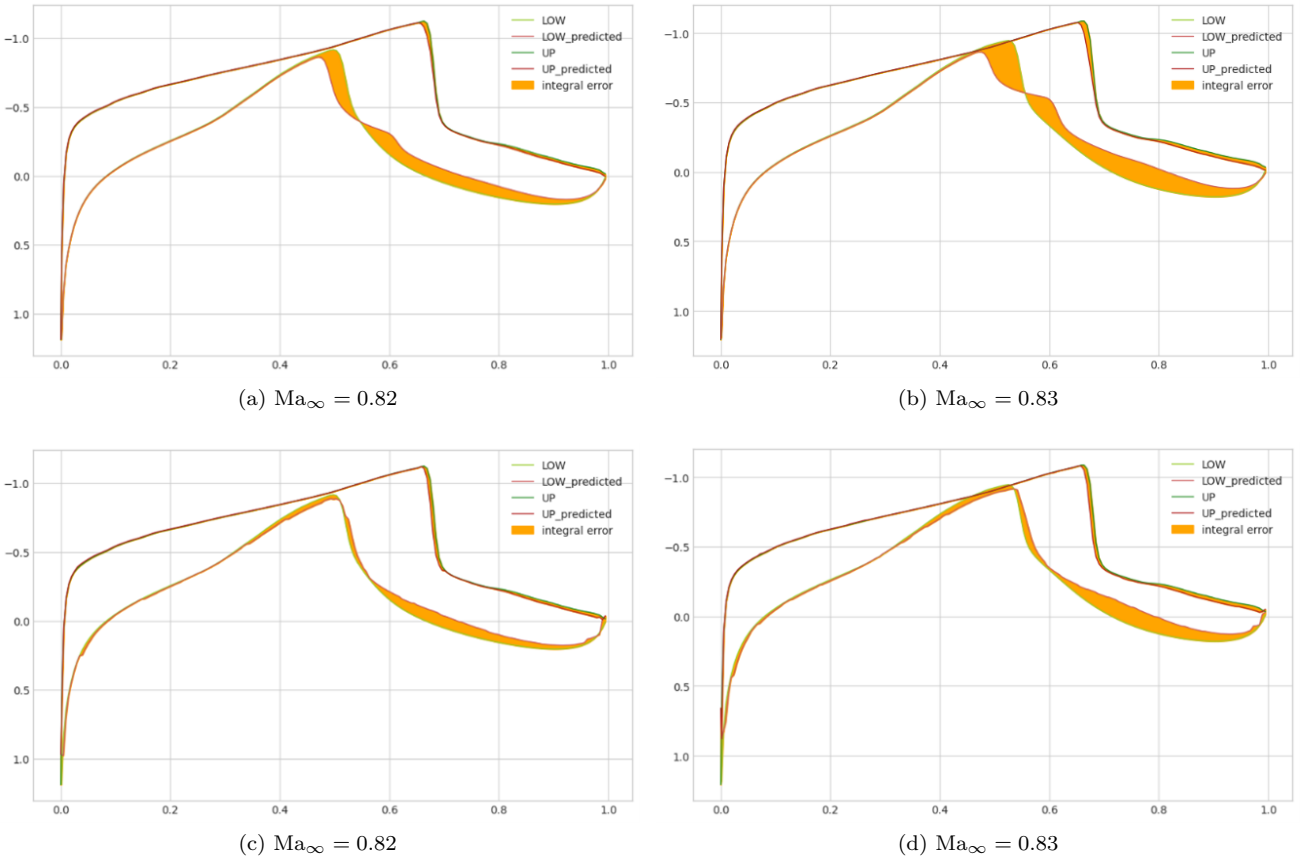


Figure 10: compressible flow past a RAE2822 airfoil, one-dimensional parameterization. Behavior of truth and predicted pressure coefficient on lower (LOW) and upper (UP) sides of the blade, for two values of Ma_∞ and $\alpha = 2^\circ$. (a)-(b) convex interpolation. (c)-(d) convex displacement interpolation.




Figure 11: compressible flow past a RAE2822 airfoil, one-dimensional parameterization. (a) behavior of the optimal value of s in (22) for CDI and CI. (b) behavior of the integral error over the airfoil.

4.1.3 Two-dimensional parameterization

We consider the problem of estimating the solution for $\mu = (\text{Ma}_\infty, \alpha) = [0.83, 2]$, based on the training points $\mu^1 = [0.82, 2]$, $\mu^2 = [0.84, 1.5]$ and $\mu^3 = [0.84, 2.5]$. We compare the performance of convex interpolation and CDI. Given the (very) modest amount of datapoints, we rely on linear interpolation for the regression step (cf. Algorithm 1, Line 1, Online stage). Figure 12 compares the predictions of truth and predicted pressure coefficient for the out-of-sample parameter μ . We observe that nonlinear interpolation is qualitatively more accurate, especially in the neighborhood of the shock. We observe that the $L^2(\Omega)$ error is 5.62% for CI and 4.10% for CDI; the $L^2(\Gamma_2)$ error is 2.69% for CI and 2.14% for CDI. We emphasize that the L^2 metric does not fully capture local improvements in the regions where the shocks develop.




Figure 12: compressible flow past a RAE2822 airfoil, two-dimensional parameterization. Behavior of truth and predicted pressure coefficient for an out-of-sample configuration. (a) convex interpolation. (b) convex displacement interpolation.

4.2 Incompressible flow past a backward facing step


4.2.1 Model problem

We study the incompressible flow past a backflow facing step for varying inflow Reynolds number Re and step's inclination angle α ; we model the flow using the incompressible Navier-Stokes equations and we rely on the open-source software OpenFOAM [20]. Mesh morphing is based on the Optimad proprietary software `mimic`¹. The coherent structures of interest are recirculation bubbles (vortices), whose extensions and positions vary as a function of the parameters. We apply the nonlinear interpolation procedure to the velocity field.


¹See <https://optimad.github.io/mimmo/>.

4.2.2 One-dimensional parameterization

We first study variations of the inflow Reynolds number in the range $[50, 200]$ for $\alpha = 90^\circ$. Figure 13 shows the velocity magnitude for two parameter values, while Figure 14 shows the recirculation area as predicted by the indicator (5). We observe that the recirculation area is very sensitive to the value of the inflow Reynolds number.



(a) $Re = 50$



(b) $Re = 200$

Figure 13: incompressible flow past a backward facing step. Velocity magnitude for two values of Re and $\alpha = 90^\circ$.




(a) $\text{Re} = 50$




(b) $\text{Re} = 200$

Figure 14: incompressible flow past a backward facing step. Marked cells for two values of Re and $\alpha = 90^\circ$.


Figure 15 compares of CDI (20) and CI (21) for $\text{Re} = 100$. We observe that the nonlinear interpolation is qualitatively more accurate to represent the recirculation area. Figure 16 investigates the behavior of the optimal choice of s (“experimental”) and compares it with the linear relationship — optimality is measured with respect to the $L^2(\Omega)$ norm. We observe that s_μ^{opt} depends almost linearly on Re ; we further observe that the improvement due to the nonlinear interpolation is less significant than for the previous case; we again emphasize that the global L^2 norm does not properly capture local-in-space features of the error.



(a) high-fidelity model



(b) nonlinear interpolation



(c) linear interpolation

Figure 15: incompressible flow past a backward facing step, one-dimensional parameterization. Predictions of velocity streamlines for $Re = 100$ based on the high-fidelity model, nonlinear interpolation and linear interpolation.





Figure 16: incompressible flow past a backward facing step, one-dimensional parameterization. (a) behavior of the optimal value of s in (22) for CDI and CI. (b) behavior of the relative $L^2(\Omega)$ error for the velocity field.


4.2.3 Two-dimensional parameterization

We consider the problem of estimating the solution for $\mu = (\text{Re}, \alpha) = [120, 70^\circ]$ based on the training points $\mu^1 = [100, 50^\circ]$, $\mu^2 = [100, 90^\circ]$ and $\mu^3 = [180, 70^\circ]$. We compare performance of convex interpolation and CDI in terms of streamlines and horizontal velocity prediction. As in the previous test case, prediction of the markers' locations for out-of-sample configurations is performed through linear interpolation. Figure 17 shows the results.


We observe that the improvement due to nonlinear interpolation in terms of global $L^2(\Omega)$ relative error in velocity prediction is very modest (1.31% vs 1.29%). However, we notice that the streamlines of the flow field that are predicted through nonlinear interpolation are qualitatively more accurate than the ones associated with linear interpolation; similarly, the prediction of the horizontal velocity based on the nonlinear model over the horizontal slice $\{x : x_1 \in (0.1, 0.7), x_2 = -0.01\}$ are significantly more accurate in the recirculation region, especially in the proximity of the minimum of the horizontal velocity.



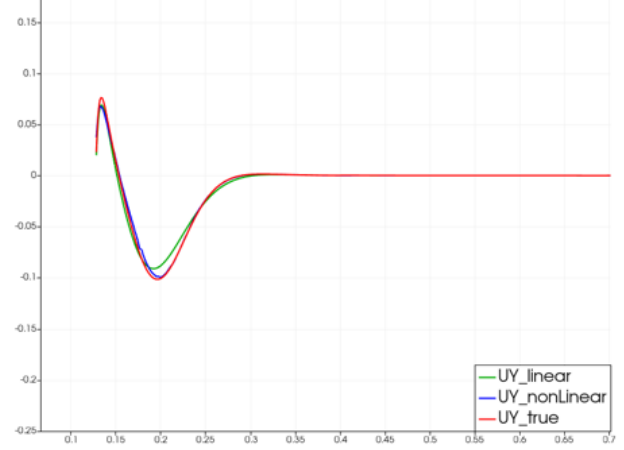
(a) high-fidelity model



(b) nonlinear interpolation



(c) linear interpolation



(d)

Figure 17: incompressible flow past a backward facing step two-dimensional parameterization. Comparison of high-fidelity, nonlinear interpolation, and linear interpolation predictions. (a)-(b)-(c) streamlines of the velocity field. (d) horizontal velocity over the horizontal slice $\{x : x_1 \in (0.1, 0.7), x_2 = -0.01\}$.

4.3 Compressible flow past a sphere

4.3.1 Model problem

We study the compressible supersonic inviscid flow past a sphere for inflow Mach number $Ma_\infty \in [2, 5]$. We model the flow using the compressible Euler equations for ideal gases and we rely on SU2 for snapshot generation. As in section 4.1, we apply the nonlinear interpolation strategy to estimate the pressure coefficient (23) and we rely on the Ducros sensor with $\gamma_{\text{thr}} = 99.6\%$ to identify the coherent structure. Figure 18 shows the behavior of C_p for $Ma_\infty = 2$ and $Ma_\infty = 5$: we observe that the solution exhibits a bow shock ahead of the sphere whose location depends on the parameter. Figure 19 shows the marked cells for the same parameter values.




Figure 18: compressible flow past a sphere. Pressure coefficients for two values of Ma_∞ .




Figure 19: compressible flow past a sphere. Marked cells for $Ma_\infty = 2$ and $Ma_\infty = 5$.

Figure 20 illustrates the performance of CDI (20) and CI (21) for $Ma_\infty \in [2, 5]$: Figure 20(a) investigates the behavior of the optimal value of s in (22) for CDI and CI; Figure 20(b) shows the behavior of the relative $L^2(\Omega)$ error for the pressure coefficient. Interestingly, we observe that the behavior of s_μ^{opt} is strongly nonlinear with respect to the Mach number. As for the previous cases, nonlinear interpolation is superior to linear interpolation, especially when the parameter s in (20) is chosen in an optimal manner.




Figure 20: compressible flow past a sphere. (a) behavior of the optimal value of s in (22) for CDI and CI. (b) behavior of the relative $L^2(\Omega)$ error for the pressure coefficient.

5 Data augmentation

5.1 Model problem: inviscid flow past a bump

We consider the problem of estimating the solution to the two-dimensional Euler equations in the transonic regime in a channel with a Gaussian bump. We denote by ρ the density, by v the velocity field, by E the total energy, and by $p = (\gamma - 1) (E - \frac{1}{2}\rho\|v\|_2^2)$ the pressure, where $\gamma > 0$ is the ratio of specific heats; we further denote by $u = [\rho, \rho v, E]$ the vector of conserved variables. We introduce the computational domain $\Omega_\mu = \{x \in (-1.5, 1.5) \times (0, 0.8) : x_2 > h e^{-25x_1^2}\}$ where $h > 0$ is a given parameter. We consider the conservation law:

$$\nabla \cdot F(u^{\text{true}}) = 0, \text{ where } F(u) = \begin{bmatrix} \rho v^\top \\ \rho v v^\top \\ (E + p)v^\top \end{bmatrix}, \quad (24)$$

completed with wall boundary conditions on top and bottom boundaries, subsonic inlet condition (total temperature, total pressure and flow direction) at the left boundary and subsonic outlet condition (static pressure) at the right boundary. We express the free-stream field u_∞ in terms of the Mach number Ma_∞ ,

$$T_\infty = 1, \quad p_\infty = \frac{1}{\gamma}, \quad \rho_\infty = 1, \quad u_\infty = [\text{Ma}_\infty, 0],$$

where $T = \frac{p}{R\rho}$ is the temperature and $R = \gamma - 1$ is the (non-dimensional) specific gas constant. Finally, we introduce the parameter vector $\mu = [h, \text{Ma}_\infty]$ and the parameter region $\mathcal{P} = [0.05, 0.065] \times [0.68, 0.78]$. Figures 21(a) and (b) show the behavior of the Mach field for two values of the parameters in \mathcal{P} .




Figure 21: inviscid flow over a Gaussian bump. Visualization of the Mach field for two parameter values.

5.2 Geometric parameterization and high-fidelity solver

The computational domain Ω_μ depends on the geometric parameter h ; therefore, we should introduce a geometric mapping to recast the problem over a parameter-independent configuration. We here resort to a Gordon-Hall transformation: we denote by $\Phi_{\text{geo},\mu}$ the geometric mapping from the rectangle $\tilde{\Omega} = (-1.5, 1.5) \times (0, 0.8)$ to the physical domain Ω_μ (see, e.g., [15, section 2]). We further define the domain $\Omega := \Omega_{\bar{\mu}}$, where $\bar{\mu}$ is the centroid of \mathcal{P} and the additional geometric mapping $\Psi_{\text{geo},\mu} := \Phi_{\text{geo},\mu} \circ \Phi_{\text{geo},\bar{\mu}}^{-1} : \Omega \rightarrow \Omega_\mu$.

The reference configuration $\tilde{\Omega}$ is used below to define the sensors and the regression procedure; on the other hand, the HF mesh is defined in the domain Ω . The latter choice is justified by the observation that for small variations of h the mapping Ψ_{geo} is (significantly) closer to the identity map than Φ_{geo} ; therefore, defining the mesh in Ω offers significantly enhanced control over the quality of the deformed mesh.

Simulations are performed in Matlab 2022a [30] based on an in-house code, and executed over a commodity Linux workstation (RAM 32 GB, Intel i7 CPU 3.20 GHz x 12). We consider a P2 mesh with $N_e = 12748$ elements; given a new value of the parameter μ , we define the deformed mesh using the geometric mapping Ψ_{geo} , and then we call the DG solver to compute the solution field. We measure performance and we compute reduced spaces in the reference domain Ω ; we denote by (\cdot, \cdot) the $[L^2(\Omega)]^4$ inner product and by $\|\cdot\| = \sqrt{(\cdot, \cdot)}$ the induced norm.

5.3 Model order reduction

We rely on a projection-based ROM based on a least-square Petrov-Galerkin (LSPG, [9]) formulation and empirical quadrature (EQ, [14, 55]) hyper-reduction; we refer to [15] for a detailed presentation of the method.

5.4 Convex displacement interpolation

We apply Algorithm 1 to construct the CDI estimate.

Feature selection. We rely on the Ducros sensor (4) to identify the elements of the mesh that contain the shock. We set $\gamma_{\text{thr}} = 99\%$ in (4c) to determine the threshold t_μ . In order to take into account geometric variations, we define the ‘‘mapped’’ point clouds $\tilde{X}_\mu^{\text{raw}} = \{\Phi_{\text{geo},\mu}^{-1}(x_{i,\mu}^{\text{raw}})\}_{i=1}^{N_\mu} \subset \tilde{\Omega}$, where $\Phi_{\text{geo},\mu} : \tilde{\Omega} \rightarrow \Omega_\mu$ is the Gordon-Hall map. Then, we set $\tilde{X}^{\text{ref}} = \tilde{X}_{\mu^{\text{ref}}}^{\text{raw}}$ where μ^{ref} is the parameter in $\mathcal{P}_{\text{train}}$ that minimizes the distance from the centroid of \mathcal{P} , and we rely on Gaussian-based PSR to determine the sorted point clouds $\{\tilde{X}_\mu : \mu \in \mathcal{P}_{\text{train}}\}$.

Regression. We apply the regression procedure described in section 2.3 to find the RBF model $\mu \in \mathcal{P} \mapsto \hat{X}_\mu^* = \{\hat{x}_{i,\mu}^*\}_{i=1}^N \subset \tilde{\Omega}$. Then, we define the point cloud in physical domain using the geometric mapping Φ_{geo} : $\hat{X}_\mu = \{\hat{x}_{i,\mu} = \Phi_{\text{geo},\mu}(\hat{x}_{i,\mu}^*)\}_{i=1}^N$. The choice of applying RBF and PSR in the reference (parameter-independent) domain $\tilde{\Omega}$ robustifies the learning process and ultimately improves prediction performance.

Nearest neighbors, registration, and choice of the weights. We build the CDI based on $\kappa = 4$ neighbors; we consider the Euclidean distance to identify the neighbors and also in the IDW strategy for the weights (cf. section 2.6). Finally, we resort to optimization-based registration (cf. (13)) to determine the mapping Φ : we consider the penalty term f_{jac} ([49, section 4.3]) that penalizes small values of the Jacobian determinant, and we rely on a piecewise-polynomial search space with $M = 198$ degrees of freedom.

5.5 Data augmentation procedure

We denote by $\mathcal{P}_{\text{train}} = \{\mu^i\}_{i=1}^{n_{\text{train}}} \subset \mathcal{P}$ the training set of parameters for which HF simulations are available and by $\mathcal{P}_{\text{train,cdi}} = \{\nu^i\}_{i=1}^{n_{\text{train,cdi}}} \subset \mathcal{P}$ an additional set of parameters. We first define the Lagrangian space $\mathcal{Z}_0 = \text{span}\{u_{\mu^i}^{\text{hf}}\}_{i=1}^{n_{\text{train}}}$ and the projection operator $\Pi_{\mathcal{Z}_0} : [L^2(\Omega)]^4 \rightarrow \mathcal{Z}_0$; then, we define the POD space

$$\mathcal{Z}_n = \mathcal{Z}_0 \oplus \mathcal{Z}_{\text{lf}}, \quad \mathcal{Z}_{\text{lf}} = \text{POD}(\mathcal{D}_{\text{train,lf}}, n - n_{\text{train}}, (\cdot, \cdot)), \quad (25)$$

where \oplus indicates the direct sum of orthogonal spaces, $\mathcal{D}_{\text{train,lf}} = \{\hat{u}_\mu - \Pi_{\mathcal{Z}_0} \hat{u}_\mu : \mu \in \mathcal{P}_{\text{train,cdi}}\}$ and \mathcal{Z}_{lf} is the POD space of size $n - n_{\text{train}}$ associated with the snapshots $\mathcal{D}_{\text{train,lf}}$ and the $L^2(\Omega)$ inner product (\cdot, \cdot) . We observe that the space \mathcal{Z}_n contains the HF training set: potential inaccuracies of the CDI estimates do not pollute the performance of the reduced space for $\mu \in \mathcal{P}_{\text{train}}$.

5.6 Numerical results: nonlinear interpolation

We first investigate the performance of CDI for this model problem. Towards this end, we study the behavior of the average relative $L^2(\Omega)$ state prediction error of the projection-based linear ROM (LSPG), CDI, and convex interpolation (CI), for tensorized grids of increasing size n_{train} ; in all our experiments we set $\kappa = 4$ for CDI and CI, while the reduced space of LSPG includes all the training snapshots ($n = n_{\text{train}}$). Figure 22 shows the results.

We observe that CDI outperforms CI for all parameter values and all training set sizes n_{train} . Interestingly, CDI is more accurate than LSPG for $n_{\text{train}} = 4$: this can be explained by observing that, unlike the LSPG ROM, CDI provides estimates that do not belong to the span of the dataset of HF simulations; for small values of n_{train} , CDI can hence compensate the fact that it does not exploit the knowledge of the PDE model for prediction.

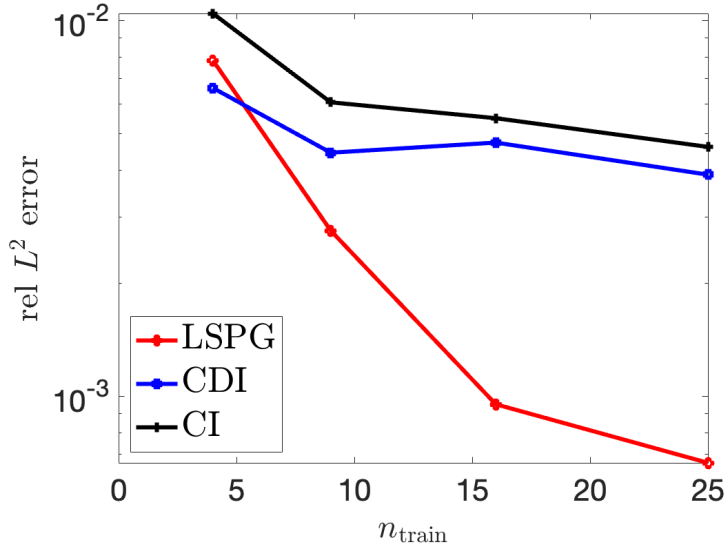


Figure 22: nonlinear interpolation. Behavior of the average relative L^2 error for LSPG, CDI and CI for snapshot sets of increasing size n_{train} .

5.7 Numerical results: data augmentation

Figure 23 illustrates the performance for $n_{\text{train}} = 9$, $n_{\text{train,cdi}} = 20$. We consider a tensorized 3×3 grid for $\mathcal{P}_{\text{train}}$ and we generate $\mathcal{P}_{\text{train,cdi}}$ using latin hyper-cube sampling [32]. We consider a POD space in (25) of dimension $n = 29$; as for the previous experiment, we assess performance based on a dataset of $n_{\text{test}} = 20$ out-of-sample parameters. Figure 23(a) shows boxplots of the relative $L^2(\Omega)$ error for LSPG ROM, CDI and LSPG ROM based on $n = 9$ modes (i.e., without data augmentation) (LSPG-0). We observe that CDI is not as accurate as LSPG-0; however, it can be used to augment the dataset of simulations and ultimately contribute to reduce the prediction error. Figure 23(b) shows the projection error over the test for the POD space built using (25) (“mixed training”), the POD space based on $n_{\text{train}} + n_{\text{train,cdi}}$ HF simulations (“HF training”) and the POD space based n_{train} HF simulations (“HF training ($n_{\text{train}} = 9$)”): we observe that the mixed training strategy based on CDI is not as accurate as HF training, particularly for large values of n ; however, it enables significant reduction of the projection error without resorting to new simulations.




Figure 23: data augmentation $n_{\text{train}} = 9$, $n_{\text{train,cdi}} = 20$. Behavior of LSPG ROM with CDI-based data augmentation. (a) boxplots of relative L^2 error for LSPG ROM, CDI and projection error ($n = 29$). (b) projection error.

Figure 24 replicates the results of Figure 23 for $n_{\text{train}} = 4$, $n_{\text{train,cdi}} = 25$. Note that CDI is more accurate than LSPG-0, which is consistent with the results of Figure 22.




Figure 24: data augmentation $n_{\text{train}} = 4$, $n_{\text{train,cdi}} = 25$. Behavior of LSPG ROM with CDI-based data augmentation. (a) boxplots of relative L^2 error for LSPG ROM, CDI and projection error ($n = 29$). (b) projection error.

Table 1 provides an overview of the computational costs. The HF solver is roughly 14 times slower than CDI and 625 times slower than the LSPG ROM; the construction of the LSPG ROM, which comprises the construction of the trial and test spaces and hyper-reduction, is roughly as expensive as a single full-order solve. We remark that approximately 99% of the total cost of CDI is associated with mesh interpolation: we hence envision that more sophisticated implementations of this step (see, e.g., [26]) might lead to a significant reduction of the overall CDI cost.

| | cost [s] |
|--|----------|
| HF solver (avg) | 481.21 |
| ROM generation (data compression+hyperreduction) | 431.24 |
| Convex displacement interpolation (avg) | 33.49 |
| LSPG ROM ($n = 29$, avg) | 0.77 |

Table 1: data augmentation. Wall-clock costs of HF and MOR procedures ($n_{\text{train}} = 9$, $n_{\text{train,cdi}} = 20$).

6 Conclusions

We presented a general (i.e., independent of the underlying equations) nonlinear interpolation procedure for parametric fields with local compactly-supported coherent structures; the approach is dubbed convex displacement interpolation (CDI) and extends the previous proposal of [19]. The methodology comprises five main building blocks that were all discussed in section 2: sensor selection, point cloud matching, regression of sorted point clouds, boundary-aware registration, and piecewise-multivariate interpolation. CDI assumes that coherent structures of the solution field — such as shocks, shear layers, vortices — vary smoothly with the parameter. The numerical experiments suggest that this assumption holds for a broad and diverse range of parametric problems in continuum mechanics.

We illustrated the application of CDI to data augmentation. Since CDI generates estimates that do not belong to the span of the training snapshots, we can exploit predictions for out-of-sample parameters to augment the dataset of simulations that are later used for the construction of projection-based ROMs. We envision that combination of lower and higher fidelity snapshots for MOR training has the potential to significantly reduce offline costs, and ultimately enable the application of MOR techniques to complex engineering problems for which extensive parameter explorations based on HF solves are not practical.

Acknowledgements

The authors acknowledge the support by European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Actions, grant agreement 872442 (ARIA).

References

- [1] D. Amsallem and C. Farhat. Interpolation method for adapting reduced-order models and application to aeroelasticity. *AIAA journal*, 46(7):1803–1813, 2008.
- [2] D. Amsallem, M. J. Zahr, and C. Farhat. Nonlinear model order reduction based on local reduced-order bases. *International Journal for Numerical Methods in Engineering*, 92(10):891–916, 2012.
- [3] P. Angot, C.-H. Bruneau, and P. Fabrie. A penalization method to take into account obstacles in incompressible viscous flows. *Numerische Mathematik*, 81(4):497–520, 1999.
- [4] J. Barnett and C. Farhat. Quadratic approximation manifold for mitigating the Kolmogorov barrier in nonlinear projection-based model order reduction. *Journal of Computational Physics*, 464:111348, 2022.
- [5] J. L. Barnett, C. Farhat, and Y. Maday. Neural-network-augmented projection-based model order reduction for mitigating the kolmogorov barrier to reducibility of CFD models. *arXiv preprint arXiv:2212.08939*, 2022.
- [6] F. Bernard, A. Iollo, and S. Riffaud. Reduced-order model for the BGK equation based on POD and optimal transport. *Journal of Computational Physics*, 373:545–570, 2018.
- [7] V. Camion and L. Younes. Geodesic interpolating splines. In *International workshop on energy minimization methods in computer vision and pattern recognition*, pages 513–527. Springer, 2001.
- [8] Y. Cao, M. I. Miller, R. L. Winslow, and L. Younes. Large deformation diffeomorphic metric mapping of vector fields. *IEEE transactions on medical imaging*, 24(9):1216–1230, 2005.
- [9] K. Carlberg, C. Farhat, J. Cortial, and D. Amsallem. The GNAT method for nonlinear model reduction: effective implementation and application to computational fluid dynamics and turbulent flows. *Journal of Computational Physics*, 242:623–647, 2013.
- [10] F. Chantalat, C.-H. Bruneau, C. Galusinski, and A. Iollo. Level-set, penalization and cartesian meshes: A paradigm for inverse problems and optimal design. *Journal of Computational Physics*, 228(17):6291–6315, 2009.
- [11] D. S. Ching, P. J. Blonigan, F. Rizzi, and J. A. Fike. Model reduction of hypersonic aerodynamics with residual minimization techniques. In *AIAA SCITECH 2022 Forum*, page 1247, 2022.
- [12] T. D. Economon, F. Palacios, S. R. Copeland, T. W. Lukaczyk, and J. J. Alonso. SU2: An open-source suite for multiphysics simulation and design. *Aiaa Journal*, 54(3):828–846, 2016.
- [13] J. L. Eftang, A. T. Patera, and E. M. Rønquist. An “hp” certified reduced basis method for parametrized elliptic partial differential equations. *SIAM Journal on Scientific Computing*, 32(6):3170–3200, 2010.

- [14] C. Farhat, T. Chapman, and P. Avery. Structure-preserving, stability, and accuracy properties of the energy-conserving sampling and weighting method for the hyper reduction of nonlinear finite element dynamic models. *International journal for numerical methods in engineering*, 102(5):1077–1110, 2015.
- [15] A. Ferrero, T. Taddei, and L. Zhang. Registration-based model reduction of parameterized two-dimensional conservation laws. *Journal of Computational Physics*, 457:111068, 2022.
- [16] M. Guo and J. S. Hesthaven. Reduced order modeling for nonlinear structural analysis using Gaussian process regression. *Computer methods in applied mechanics and engineering*, 341:807–826, 2018.
- [17] R. Horaud, F. Forbes, M. Yguel, G. Dewaele, and J. Zhang. Rigid and articulated point registration with expectation conditional maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(3):587–602, 2010.
- [18] A. Iollo and D. Lombardi. Advection modes by optimal mass transfer. *Physical Review E*, 89(2):022923, 2014.
- [19] A. Iollo and T. Taddei. Mapping of coherent structures in parameterized flows by learning optimal transportation with Gaussian models. *Journal of Computational Physics*, 471:111671, 2022.
- [20] H. Jasak. OpenFOAM: Open source CFD in research and industry. *International Journal of Naval Architecture and Ocean Engineering*, 1(2):89–94, 2009.
- [21] J. Jeong and F. Hussain. On the identification of a vortex. *Journal of fluid mechanics*, 285:69–94, 1995.
- [22] T. Kadeethum, D. O’Malley, F. Ballarin, I. Ang, J. N. Fuhg, N. Bouklas, V. L. Silva, P. Salinas, C. E. Heaney, and C. C. Pain. Enhancing high-fidelity nonlinear solver with reduced order model. *Scientific Reports*, 12(1):20229, 2022.
- [23] V. Kolář. Vortex identification: New requirements and limitations. *International journal of heat and fluid flow*, 28(4):638–652, 2007.
- [24] P. Kraus, S. Büchholz, M. Häring, and J. Reiss. Front transport reduction for complex moving fronts. *Journal of scientific computing*, 96(28), 2023.
- [25] K. Lee and K. T. Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics*, 404:108973, 2020.
- [26] E. Luke, E. Collins, and E. Blades. A fast mesh deformation method using explicit interpolation. *Journal of Computational Physics*, 231(2):586–601, 2012.
- [27] J. Ma, W. Qiu, J. Zhao, Y. Ma, A. L. Yuille, and Z. Tu. Robust L_2 E estimation of transformation for non-rigid registration. *IEEE Transactions on Signal Processing*, 63(5):1115–1129, 2015.
- [28] J. Ma, J. Wu, J. Zhao, J. Jiang, H. Zhou, and Q. Z. Sheng. Nonrigid point set registration with robust transformation learning under manifold regularization. *IEEE transactions on neural networks and learning systems*, 30(12):3584–3597, 2018.
- [29] B. Maiseli, Y. Gu, and H. Gao. Recent developments and trends in point set registration methods. *Journal of Visual Communication and Image Representation*, 46:95–106, 2017.
- [30] MATLAB. *R2022a*. The MathWorks Inc., Natick, Massachusetts, 2022.
- [31] R. J. McCann. A convexity principle for interacting gases. *Advances in mathematics*, 128(1):153–179, 1997.
- [32] M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 2000.
- [33] F. R. Menter. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA journal*, 32(8):1598–1605, 1994.
- [34] M. A. Mirhoseini and M. J. Zahr. Model reduction of convection-dominated partial differential equations via optimization-based implicit feature tracking. *Journal of Computational Physics*, 473:111739, 2023.
- [35] D. Modesti and S. Pirozzoli. A low-dissipative solver for turbulent compressible flows on unstructured meshes, with Open-FOAM implementation. *Computers & Fluids*, 152:14–23, 2017.

- [36] R. Mojjani and M. Balajewicz. Arbitrary Lagrangian Eulerian framework for efficient projection-based reduction of convection dominated nonlinear flows. In *APS Division of Fluid Dynamics Meeting Abstracts*, pages M1–008, 2017.
- [37] R. Mojjani and M. Balajewicz. Low-rank registration based manifolds for convection-dominated PDEs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 399–407, 2021.
- [38] A. Myronenko and X. Song. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2262–2275, 2010.
- [39] F. Nicoud and F. Ducros. Subgrid-scale stress modelling based on the square of the velocity gradient tensor. *Flow, turbulence and Combustion*, 62(3):183–200, 1999.
- [40] M. Ohlberger and S. Rave. Nonlinear reduced basis approximation of parameterized evolution equations via the method of freezing. *Comptes Rendus Mathematique*, 351(23-24):901–906, 2013.
- [41] B. Peherstorfer. Model reduction for transport-dominated problems via online adaptive bases and adaptive sampling. *SIAM Journal on Scientific Computing*, 42(5):A2803–A2836, 2020.
- [42] P.-O. Persson and J. Peraire. Sub-cell shock capturing for discontinuous Galerkin methods. In *44th AIAA aerospace sciences meeting and exhibit*, page 112, 2006.
- [43] A. Quarteroni. *Numerical models for differential problems*, volume 2. Springer, 2009.
- [44] J. Reiss, P. Schulze, J. Sesterhenn, and V. Mehrmann. The shifted proper orthogonal decomposition: A mode decomposition for multiple transport phenomena. *SIAM Journal on Scientific Computing*, 40(3):A1322–A1344, 2018.
- [45] J. A. Rice. *Mathematical statistics and data analysis*. Cengage Learning, 2006.
- [46] D. Shepard. A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM national conference*, pages 517–524, 1968.
- [47] L. Sirovich. Turbulence and the dynamics of coherent structures. I. Coherent structures. *Quarterly of applied mathematics*, 45(3):561–571, 1987.
- [48] T. Taddei. A registration method for model order reduction: data compression and geometry reduction. *SIAM Journal on Scientific Computing*, 42(2):A997–A1027, 2020.
- [49] T. Taddei. Compositional maps for registration in complex geometries. *arXiv preprint arXiv:2308.15307*, 2023.
- [50] T. Taddei, S. Perotto, and A. Quarteroni. Reduced basis techniques for nonlinear conservation laws. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(3):787–814, 2015.
- [51] T. E. Tezduyar, M. Behr, S. Mittal, and J. Liou. A new strategy for finite element computations involving moving boundaries and interfaces—the deforming-spatial-domain/space-time procedure: II. computation of free-surface flows, two-liquid flows, and flows with drifting cylinders. *Computer methods in applied mechanics and engineering*, 94(3):353–371, 1992.
- [52] P. Tonon, R. A. K. Sanches, K. Takizawa, and T. E. Tezduyar. A linear-elasticity-based mesh moving method with no cycle-to-cycle accumulated distortion. *Computational Mechanics*, 67:413–434, 2021.
- [53] S. Volkwein. Model reduction using proper orthogonal decomposition. *Lecture Notes, Institute of Mathematics and Scientific Computing, University of Graz*. see <http://www.uni-graz.at/imawww/volkwein/POD.pdf>, 1025, 2011.
- [54] H. Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- [55] M. Yano and A. T. Patera. An LP empirical quadrature procedure for reduced basis treatment of parametrized nonlinear PDEs. *Computer Methods in Applied Mechanics and Engineering*, 344:1104–1123, 2019.
- [56] B. Zitova and J. Flusser. Image registration methods: a survey. *Image and vision computing*, 21(11):977–1000, 2003.