



Extending parametric model embedding with physical information for design-space dimensionality reduction in shape optimization

Andrea Serani¹ · Giorgio Palma¹ · Jeroen Wackers² · Domenico Quagliarella³ · Stefano Gaggero⁴ · Matteo Diez¹

Received: 11 April 2025 / Accepted: 7 September 2025
© The Author(s) 2025

Abstract

Design-space dimensionality reduction is essential to mitigate the cost of high-fidelity simulation-based optimization, especially when dealing with high-dimensional geometric parameterizations. Traditional linear techniques, such as principal component analysis, are widely used but often neglect the physical response of the system and lack invertibility to the design space, i.e., the ability to reconstruct the original design parameters from a reduced representation. This work introduces two physics-aware extensions of the parametric model embedding (PME) framework, aimed at generating reduced representations that incorporate physical information while maintaining analytical backmapping. The first, physics-informed PME (PI-PME), combines geometric and physical variability; the second, physics-driven PME (PD-PME), relies solely on physical responses. The proposed methods enable the construction of interpretable and physically relevant reduced spaces that can be used for design-space exploration, surrogate modeling, and optimization. The approach is demonstrated on multiple engineering configurations, including airfoils, propellers, gliders, and hulls, showing its ability to capture performance-relevant directions and preserve parametric consistency. The methodology is offline and non-intrusive, compatible with low-fidelity simulations, and requires only a modest number of samples to ensure variance convergence.

Keywords Dimensionality reduction · Representation learning · Parametric model embedding · Shape optimization · Vehicle design

1 Introduction

Shape optimization of functional surfaces presents a multifaceted challenge, characterized by numerous geometric, functional, and performance constraints. This is particularly evident in vehicle design, where modern processes must address a broad spectrum of requirements-ranging from energy efficiency and cost effectiveness to safety and environmental sustainability-while simultaneously accounting for aerodynamics, hydrodynamics, structural integrity,

noise emissions, and regulatory standards. As the complexity of each vehicle concept increases, so does the dimensionality of the design space, often leading to an exponential growth in the number of parameters that define the shape and operating conditions. This phenomenon, commonly referred to as the *curse of dimensionality* [1], complicates the exploration, analysis, and optimization tasks, since an extremely large solution space quickly becomes prohibitive to sample thoroughly.

Design-space dimensionality reduction methods for shape optimization [2] have been extensively explored to address these challenges, simplifying high-dimensional design spaces without significantly compromising predictive accuracy and can be broadly categorized into *offline* and *online* approaches. Offline or upfront methods (e.g., principal component analysis and its variants [3, 4]) aim to construct a reduced representation before the actual optimization or design exploration process, relying on a fixed dataset of samples. Conversely, online or adaptive techniques (e.g., active subspace methods [5, 6]) seek to refine

✉ Andrea Serani
andrea.serani@cnr.it

¹ National Research Council-Institute of Marine Engineering, Rome, Italy

² Ecole Centrale de Nantes/CNRS, Nantes, France

³ Italian Aerospace Research Centre, Capua, Italy

⁴ University of Genoa, Genoa, Italy

the reduced space on-the-fly, as new data becomes available. While online approaches may offer enhanced flexibility, they often require substantial computational overhead and complex integration with the optimization workflow. In contrast, offline methods are appealing for their simplicity, interpretability, and suitability for use with high-fidelity simulations, provided that the initial sampling ensures sufficient coverage of the design space. Among them, linear methods, such as principal component analysis (PCA) [3, 7–9] and singular value decomposition (SVD) [10–12] derived from proper orthogonal decomposition (POD) [13–18] and Karhunen-Loève expansion (KLE) [4, 19, 20], have been widely adopted for reducing the dimensionality of design space by capturing the dominant modes of geometric variance. However, traditional PCA-based approaches primarily utilize geometric data, potentially neglecting critical physical phenomena influencing design performance. To enhance predictive capabilities, physics-informed PCA methods have recently emerged [21–25], augmenting geometric data matrices with physical simulation outputs, such as lumped parameters (e.g., efficiency or drag coefficients) or distributed physical quantities (e.g., pressure or velocity fields). Furthermore, geometric moments have been introduced in [26–30] as cost-effective, physics-related descriptors that substitute direct simulation data, further enriching the PCA model without the computational cost of numerical simulations. These enriched representations allow dimensionality reduction frameworks to capture not only geometric but also essential physical variability, significantly improving design relevance and effectiveness.

Despite their widespread use, conventional offline dimensionality reduction techniques such as PCA and SVD present critical limitations in design applications. First, SVD-based approaches, which are often employed in POD and KLE formulations, do not support weighting of the inputs due to their inherent algebraic structure. This prevents users from tuning the relative importance of geometric features or physical observables during the reduction. In contrast, PCA-formulated as an eigenvalue problem of a covariance matrix-can accommodate weighted formulations, enabling generalized PCA variants and physics-informed extensions that incorporate additional information beyond shape geometry. However, these approaches still lack a crucial capability: the ability to reconstruct a feasible set of design parameters from a point in the reduced space. This issue, commonly referred to as the *pre-image problem* [31], severely limits their applicability in optimization workflows or CAD-integrated environments where direct parametric interpretation and reversibility are required.

To overcome the latter limitation, the parametric model embedding (PME) technique [32] was introduced. PME embeds the original parametric variables explicitly into

the reduction process and re-parameterizes [33] the design space by constructing an orthonormal basis that supports both weighted formulations and analytical backmapping, thereby bridging the gap between statistical reduction and design interpretability.

Nevertheless, the original PME was designed to address the pre-image problem only, without accounting for physical integral quantities or distributed fields, which are essential to characterize performance in fluid-dynamic applications. The present work extends PME to incorporate such physical information, exploring two complementary linear offline strategies: physics-informed PME (PI-PME), which embeds physical outputs alongside geometry and design variables, and physics-driven PME (PD-PME), which relies solely on physical features to define the embedding.

PI-PME enriches the data matrix with physical information, including lumped parameters (e.g., force coefficients or efficiencies) or distributed quantities (e.g., surface pressure fields), derived from low-fidelity simulations. Although this enrichment entails an upfront computational cost, the resulting latent space captures both geometric and physical variability, enabling a more informed and reliable assessment of performance trade-offs. By integrating physical insights at the earliest stages of design exploration, PI-PME provides a foundation that can guide subsequent high-fidelity simulations and optimizations.

Building on the same principle, a PD-PME is also proposed, which goes one step further by focusing solely on physical parameters and effectively excluding geometric variance. Although this approach removes explicit shape information from the dimensionality reduction, it can be advantageous in scenarios where the primary objective is to characterize or optimize physical behaviors, and where the link between geometry and physics is well captured through simplified models. This exclusive emphasis on physics can also serve as a natural starting point for multi-fidelity optimization or for data-driven reduced-order modeling, where lumped or distributed physical data are used to predict key performance indicators at reduced computational expense.

Table 1 summarizes the distinguishing features of the considered dimensionality reduction methods with respect to four key aspects: (i) the type of data handled (geometry, physics, or both), (ii) the inclusion of design variables in the embedding, (iii) the availability of weight control, and (iv) the capability to perform backmapping to the original parametric space. Classical methods such as SVD and PCA rely solely on geometric information and, while PCA can incorporate weights to prioritize specific features, neither provides a mechanism for recovering design parameters from reduced coordinates. Physics-informed variants of PCA extend this formulation by integrating physical observables-either lumped or distributed-but still lack the ability to

Table 1 Comparison of dimensionality reduction methods in terms of data types handled, use of design variables, availability of weight control, and possibility of backmapping to the original parametric space

Method	Geometry	Physics	Variables	Weights	Backmapping	Notes
SVD	✓	–	–	✗	✗	Standard decomposition of geometry-only snapshots
PCA	✓	–	–	✓	✗	Can weight geometric features
PI-PCA	✓	✓	–	✓	✗	Can also incorporate physics thanks to weights
PME	✓	–	✓	✓	✓	Embeds geometry and variables
PI-PME	✓	✓	✓	✓	✓	Combines geometry and physics with variables embedding
PD-PME	–	✓	✓	✓	✓	Fully physics-driven embedding

reconstruct input designs. In contrast, the PME framework embeds design variables directly into the reduction process, enabling analytical backmapping and supporting weighted formulations. Its extensions, PI-PME and PD-PME, further enrich the embedding by incorporating physical quantities. While PI-PME retains geometric and parametric information, PD-PME constructs the latent space purely from physical descriptors, which can be advantageous when physical trends are more relevant than shape variance. These distinctions highlight the versatility of PME-based approaches in design-space reduction tasks where physical interpretability and parameter reconstruction are essential.

The effectiveness of PI-PME and PD-PME becomes particularly apparent in real-world applications such as bio-inspired underwater gliders, naval propellers, or classical airfoils (e.g., the RAE-2822). In these cases, fluid-dynamic effects strongly impact performance, and omitting them from the dimensionality reduction step may lead to suboptimal designs and increased resource expenditure. By embedding physical information—whether partially, as in PI-PME, or exclusively, as in PD-PME—into the reduced space, it becomes possible to filter the vast design domain more intelligently, ensuring that only the most relevant configurations are explored. This synergy between geometric and physical factors not only improves the quality of early design-stage decisions but also streamlines the path to final optimization, cutting down both computational time and overall development costs. It is important to emphasize that the test cases discussed throughout the paper are not simplified academic benchmarks, but rather representative of real-world early-stage design workflows in aerospace and marine engineering. Each case reflects a high-dimensional design problem where physical performance plays a central role: the RAE-2822 airfoil is a canonical transonic shape studied for aerodynamic efficiency; the seakeeping hull and propeller designs involve complex hydrodynamic effects under realistic operating conditions; and the glider case addresses the viscous performance of bioinspired geometries. In all scenarios, dimensionality reduction is not an optional step, but a necessary enabler to manage complexity, improve interpretability, and make downstream optimization tractable. PME's flexible structure also makes it directly applicable

to structural optimization. Instead of relying on aerodynamic or hydrodynamic simulations, one may incorporate structural properties (e.g., mass moments of inertia, stiffness distributions, or stress fields) into the PME framework, enabling integrated assessments of mechanical performance [22, 34].

The following sections illustrate the mathematical underpinnings of PME, PI-PME, and PD-PME, including detailed discussions of how generalized PCA is adapted to incorporate physical data, and how these frameworks can be applied to a range of high-dimensional vehicle design problems. Through examples, the benefits of early-stage integration of physical insights will be demonstrated, showing improved design robustness, a decrease in exhaustive computational campaigns, and the potential for more agile innovation cycles in modern vehicle development. It may finally be noted that, while this work focuses on vehicle applications, the proposed method is broadly applicable to general shape optimization problems involving functional surfaces influenced by physical constraints.

2 Design-space dimensionality reduction methods

In shape optimization problems aimed at minimizing an objective function $f(\mathbf{u})$, the optimum shape is uncertain before and during the optimization process. The design variable vector $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^M$, with $\mathcal{U}$ the design domain, can therefore be treated as a random variable, with an associated probability density $p(\mathbf{u})$ which reflects the likelihood of a certain design being optimal, given the existing knowledge. If nothing is known about the design, p can be uniform. This uncertainty propagates through the modeling process, generating uncertainties in both geometric configurations and the associated physical responses.

Dimensionality reduction techniques such as PCA, SVD, POD, and KLE provide a systematic way to extract the dominant directions of variability from such data, i.e., the combinations of design variables which most influence the results. These methods enable the construction of compact, low-dimensional representations that preserve

the most informative features for design analysis and decision-making.

As will be discussed in the following sections, integrating physical insights into these dimensionality reduction frameworks, for design-space assessment and subsequent shape optimization, provides a more robust path toward identifying optimal solutions under uncertainty.

2.1 Parametric model embedding

PME [32] is a design-space dimensionality reduction method that extends the standard PCA approach by incorporating both shape deformations and design variables into a generalized feature space. Specifically, PME solves a generalized eigenvalue problem, applying a generalized PCA to an augmented matrix that includes the discretized shape deformation vector $\mathbf{d}$ and the original design variables $\mathbf{u}$. This extension allows PME to directly map the reduced design space back to the original design variables without the need for reparameterization, which is typically required in standard PCA approaches. As a result, PME offers a more robust and practical method for maintaining the integrity of the original design features while facilitating effective shape optimization.

Consider a manifold $\mathcal{G}$, which identifies the original shape/geometry. This manifold defines the geometric space in which the shape is parameterized by curvilinear coordinates $\xi \in \mathcal{G}$. The coordinates of the original shape are represented by $\mathbf{g}(\xi) \in \mathbb{R}^n$ with $n = 1, 2$, or 3 . Assume that, for the purpose of shape optimization, $\mathbf{g}$ can be transformed to a deformed shape/geometry $\mathbf{g}'(\xi, \mathbf{u})$ by

$$\mathbf{g}'(\xi, \mathbf{u}) = \mathbf{g}(\xi) + \delta(\xi, \mathbf{u}) \quad \forall \xi \in \mathcal{G} \quad (1)$$

where $\delta(\xi, \mathbf{u}) \in \mathbb{R}^n$ is the resulting shape modification vector, defined by arbitrary shape parameterization or

modification method (e.g., CAD parameterization, Bezier surfaces, free-form deformation, NURBS, etc.).

Discretizing $\mathcal{G}$ by L elements (see, e.g., Fig. 1) of measure $\Delta\mathcal{G}_j$ (with $j = 1, \dots, L$), and sampling $\mathcal{U}$ by a statistically convergent number of realizations S , so that $\{\mathbf{u}_k\}_{k=1}^S \sim p(\mathbf{u})$, $\mathbf{d}(\xi, \mathbf{u})$ can be obtained as the discretization of $\delta(\xi, \mathbf{u})$. Organizing $\hat{\mathbf{d}} = \mathbf{d} - \langle \mathbf{d} \rangle$ (with $\langle \cdot \rangle$ the mean value) in a data matrix $\mathbf{D}$ of dimensionality $[nL \times S]$, one obtains

$$\mathbf{D} = \begin{bmatrix} \hat{\mathbf{d}}_1^{(1)} & \dots & \hat{\mathbf{d}}_1^{(S)} \\ \vdots & \vdots & \vdots \\ \hat{\mathbf{d}}_n^{(1)} & \dots & \hat{\mathbf{d}}_n^{(S)} \end{bmatrix} = \begin{bmatrix} \hat{d}_{1,\xi_1}(\mathbf{u}_1) & \dots & \hat{d}_{1,\xi_1}(\mathbf{u}_S) \\ \vdots & \vdots & \vdots \\ \hat{d}_{L,\xi_1}(\mathbf{u}_1) & \dots & \hat{d}_{L,\xi_1}(\mathbf{u}_S) \\ \vdots & \dots & \vdots \\ \hat{d}_{1,\xi_n}(\mathbf{u}_1) & \dots & \hat{d}_{1,\xi_n}(\mathbf{u}_S) \\ \vdots & \vdots & \vdots \\ \hat{d}_{L,\xi_n}(\mathbf{u}_1) & \dots & \hat{d}_{L,\xi_n}(\mathbf{u}_S) \end{bmatrix} \quad (2)$$

where $\hat{d}_{j,\xi_k}$ is the k -th component of the shape modification vector associated to the j -th element. The embedding is achieved by defining the matrix $\mathbf{P}$ of dimensionality $[(nL + M) \times S]$ as follows

$$\mathbf{P} = \begin{bmatrix} \mathbf{D} \\ \mathbf{U} \end{bmatrix} \quad (3)$$

with

$$\mathbf{U} = [\hat{\mathbf{u}}^{(1)} \quad \dots \quad \hat{\mathbf{u}}^{(S)}] = \begin{bmatrix} \hat{u}_{1,1} & \dots & \hat{u}_{1,S} \\ \vdots & \dots & \vdots \\ \hat{u}_{M,1} & \dots & \hat{u}_{M,S} \end{bmatrix} \quad (4)$$

where $\hat{\mathbf{u}} = \mathbf{u} - \langle \mathbf{u} \rangle$. The matrix $\mathbf{U}$ is appended to the data matrix $\mathbf{D}$ and associated to a null weight $\mathbf{W}_u$ such that

$$\mathbf{W}_u = \mathbf{0} \quad \text{and} \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_d & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_u \end{bmatrix} \quad (5)$$

and so leading to a generalized PCA problem in the form

$$\mathbf{A}\mathbf{G}\mathbf{W}\mathbf{Z} = \mathbf{Z}\mathbf{A} \quad \text{with} \quad \mathbf{A} = \frac{1}{S}\mathbf{P}\mathbf{P}^T \quad (6)$$

where

$$\mathbf{G} = \begin{bmatrix} \mathbf{G}_d & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \quad (7)$$




Fig. 1 Example of shape modification and discretization with notation, where $n = 2$ and $L = 4$

and

$$\mathbf{Z} = [\mathbf{z}_1 \ \dots \ \mathbf{z}_S] \quad \text{with} \quad \mathbf{z}_k = \begin{bmatrix} \mathbf{q}_k \\ \mathbf{v}_k \end{bmatrix} \quad (8)$$

is the eigenvectors matrix, whereas $\mathbf{\Lambda}$ is the associated eigenvalues matrix. Here, $\mathbf{q}_k$ and $\mathbf{v}_k$ represent the eigenvector components associated to the shape modification $\mathbf{d}$ and design variable $\mathbf{u}$ vectors, respectively. The matrix $\mathbf{G}_d = \text{diag}(\mathbf{G}_1, \dots, \mathbf{G}_n)$ is block diagonal and has dimensionality $[nL \times nL]$, with each $[L \times L]$ k -th block being a diagonal matrix itself

$$\mathbf{G}_k = \text{diag}(\Delta\mathcal{G}_1, \dots, \Delta\mathcal{G}_L) \quad (9)$$

containing the measure $\Delta\mathcal{G}_j$ of the j -th element. Similarly, $\mathbf{W}_d = \text{diag}(\mathbf{W}_1, \dots, \mathbf{W}_n)$ is a block diagonal matrix of dimensionality $[nL \times nL]$, where each $[L \times L]$ k -th block $\mathbf{W}_k$ ($k = 1, \dots, n$) is itself a diagonal matrix defined as

$$\mathbf{W}_k = \text{diag}(\rho_1, \dots, \rho_L) \quad (10)$$

where ρ_j (for $j = 1, \dots, L$) represents the arbitrary weight given to each element. The columns of $\mathbf{Z}$ are normalized to unit norm with respect to the $\mathbf{GW}$ scalar product. Specifically, each column $\mathbf{z}_k$ is scaled by a scalar γ_k such that:

$$\gamma_k = \sqrt{\mathbf{z}_k^T \mathbf{GW} \mathbf{z}_k}, \quad \mathbf{z}_k^* = \frac{\mathbf{z}_k}{\gamma_k}, \quad (11)$$

leading to the normalized matrix:

$$\mathbf{Z}^* = \mathbf{Z} \mathbf{\Gamma}^{-1}, \quad (12)$$

where

$$\mathbf{\Gamma} = \text{diag} \left(\left[\sqrt{\mathbf{z}_1^T \mathbf{GW} \mathbf{z}_1}, \dots, \sqrt{\mathbf{z}_{nL}^T \mathbf{GW} \mathbf{z}_{nL}} \right] \right). \quad (13)$$

This normalization ensures that each eigenvector contributes equally, avoiding numerical ill-conditioning due to significant variations in vector norms.

The solutions λ_k and the corresponding normalized eigenvectors $\mathbf{v}_k^*$ (columns component of $\mathbf{Z}^*$ associated to the original design variables, see Eq. 7) are used to construct the reduced dimensionality representation of the original parameterization by means of the N reduced design variables $\mathbf{x} = [x_1, \dots, x_N]^T$. Defining the desired confidence level l , with $0 < l \leq 1$, the number of reduced design variables N is chosen such that

$$\sum_{k=1}^N \lambda_k \geq l \sum_{k=1}^{nL} \lambda_k = l \sigma^2 \quad \text{with} \quad \lambda_k \geq \lambda_{k+1}, \quad (14)$$

and the PME of the original design variables is achieved by using these normalized eigenvectors as follows

$$\mathbf{u} \approx \check{\mathbf{u}} = \langle \mathbf{u} \rangle + \sum_{k=1}^N x_k \mathbf{v}_k^*. \quad (15)$$

Thus design parameters $\mathbf{u}$ are assigned zero weight in the inner product definition, but they remain structurally embedded in the generalized eigenvalue problem. This allows PME to retain a direct mapping from the reduced space back to the original parametric domain, thus enabling analytical backmapping (see Eq. 15). As a result, the dimension reduction does not compromise the ability to reconstruct feasible and interpretable design configurations.

Finally, to ensure that all the samples in $\mathbf{P}$ can be reconstructed through the reduced-order representation of the original design space, the reduced design variables $\mathbf{x}$ can be bounded by identifying the minimum and maximum values attained by each component Θ_{jk} of the projection coefficients θ_j , as follows:

$$\min_j \Theta_{jk} \leq x_k \leq \max_j \Theta_{jk} \quad \text{for } k = 1, \dots, N, \quad (16)$$

with $\mathbf{\Theta} = [\theta_1, \dots, \theta_S]^T$ evaluated by projecting the matrix $\mathbf{P}$ onto the normalized basis $\mathbf{Z}^*$, i.e.,

$$\mathbf{\Theta} = \mathbf{P}^T \mathbf{GW} \mathbf{Z}^*. \quad (17)$$

Moreover, it can be shown that the sum of the squared projection coefficients across the dataset equals the sum of the corresponding eigenvalues [4]:

$$\frac{1}{S} \sum_{j=1}^S \Theta_{jk}^2 = \lambda_k \quad \text{for } k = 1, \dots, N. \quad (18)$$

This implies that the variance captured along each principal direction is preserved in the projection.

It may be noted that the overall methodology is independent of the specific shape modification method, which is seen as a black box by PME.

2.2 Physics-informed parametric model embedding

Introducing physical information within PME formulations translates into the definition of, similarly to Eq. 3, a new matrix $\mathbf{P}_I$ as

$$\mathbf{P}_I = \begin{bmatrix} \mathbf{D} \\ \mathbf{U} \\ \mathbf{F} \\ \mathbf{C} \end{bmatrix} \quad (19)$$

with

$$\mathbf{F} = \begin{bmatrix} \hat{\mathbf{f}}_j^{(1)} & \cdots & \hat{\mathbf{f}}_j^{(S)} \\ | & & | \\ | & & | \end{bmatrix} \quad (20)$$

and

$$\mathbf{C} = \begin{bmatrix} \hat{c}_j^{(1)} & \cdots & \hat{c}_j^{(S)} \\ | & & | \\ | & & | \end{bmatrix} \quad (21)$$

where $\mathbf{F}$ is a matrix containing distributed physical information (with $\hat{\mathbf{f}}_j = \mathbf{f}_j - \langle \mathbf{f}_j \rangle$), such as, e.g., pressure distribution and/or velocity components, that don't necessarily have to be defined on the geometry surface but can also belong to the field surrounding the object, like, e.g., wake and vortices, whereas $\mathbf{C}$ contains lumped or scalar quantities of interest (with $\hat{c}_j = c_j - \langle c_j \rangle$), such as, e.g., lift and drag forces. It may be noted that for a given geometry, the physical information can be collected for more than one operating condition.

Analogously, a corresponding block-diagonal weight matrix $\mathbf{W}_I$ is introduced:

$$\mathbf{W}_I = \begin{bmatrix} \mathbf{W}_d & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_u & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{W}_f & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{W}_c \end{bmatrix} \quad (22)$$

$$= \text{diag}(\mathbf{W}_d, \underbrace{\mathbf{0}}_u, \mathbf{W}_f, \mathbf{W}_c),$$

where $\mathbf{W}_d$ is the block-diagonal matrix weighting the geometric entries (as in standard PME), $\mathbf{W}_u$ the one for design variables, $\mathbf{W}_f$ for distributed physical data $\mathbf{F}$, and $\mathbf{W}_c$ for lumped scalars $\mathbf{C}$.

Each block has to be set up to normalize its respective data by the inverse of the estimated variance. In practice, for a block associated with a data vector $y \in \{\mathbf{d}, \mathbf{f}, \mathbf{c}\}$, one calculates

$$\text{Var}(y) = \frac{1}{S} \sum_{j=1}^S (y_j - \langle y \rangle)^2, \quad (23)$$

and then defines

$$\rho_y = 1/\text{Var}(y). \quad (24)$$

If the block contains multiple components (e.g., multi-dimensional fields or multiple operating conditions), each row or column of that block can be assigned its own weight ρ_i . This normalization strategy ensures that both geometric modifications and physical observables contribute on a comparable scale to the PCA procedure—preventing any data block from dominating the embedding solely due to differences in units, magnitude, or raw variance.

Once the augmented data matrix $\mathbf{P}_I$ and the block-diagonal weighting matrix $\mathbf{W}_I$ are assembled, the eigenvalue problem remains analogous to the original PME formulation:

$$\mathbf{A}_I \mathbf{G}_I \mathbf{W}_I \mathbf{Z}_I = \mathbf{Z}_I \mathbf{\Lambda}_I \quad (25)$$

where

$$\mathbf{A}_I = \frac{1}{S} \mathbf{P}_I \mathbf{P}_I^T, \quad \mathbf{G}_I = \begin{bmatrix} \mathbf{G}_d & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{G}_f & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad (26)$$

and

$$\mathbf{Z}_I = [\mathbf{z}_{I,1} \dots \mathbf{z}_{I,S}], \quad \mathbf{z}_{I,k} = \begin{bmatrix} \mathbf{q}_k \\ \mathbf{v}_k \\ \phi_k \\ \pi_k \end{bmatrix}_I \quad (27)$$

Here, $\mathbf{G}_f$ accounts for the element size of the distributed physical vector, which does not necessarily correspond to $\mathbf{G}_d$, while the identity block is also applied to the lumped scalars. The eigenvector solution $\mathbf{Z}_I$ then provides the reduced representation for all data (where ϕ_k and π_k are the eigenvector components that embed the distributed and lumped physical parameters), ensuring that the new PI-PME basis incorporates both shape and physics with consistent normalization. This yields a lower-dimensional yet physics-enriched space suitable for subsequent design-space exploration, optimization, or multi-fidelity modeling.

2.3 Physics-driven parametric model embedding

In some design scenarios, large geometric changes may not necessarily translate into significant variations in the associated physical phenomena. Consequently, the standard geometry-centric approach in dimensionality reduction can yield subspaces that capture a high geometric variance but do not correlate with improved physical performance. PD-PME addresses this issue by removing geometric deformations ($\mathbf{D}$) altogether, focusing only on the physical data.

Such a formulation aims to isolate and amplify the directions in the design space that have the strongest impact on physically relevant quantities, rather than those with merely large geometric variability.

By excluding $\mathbf{D}$, the augmented data matrix reduces to

$$\mathbf{P}_D = \begin{bmatrix} \mathbf{U} \\ \mathbf{F} \\ \mathbf{C} \end{bmatrix}.$$

Because the objective is to let the physical variability guide the principal components, one assembles a block-diagonal matrix

$$\mathbf{W}_D = \text{diag}\left(\underbrace{\mathbf{0}}_{\mathbf{u}}, \mathbf{W}_f, \mathbf{W}_c\right),$$

which assigns, as per standard PME and PI-PME, zero weight to the rows corresponding to $\mathbf{U}$, while normalizing the distributed physical fields $\mathbf{F}$ and the lumped scalars $\mathbf{C}$ by the inverse of their respective variances, as per PI-PME (see, eqs. 23 and 24).

The solution of the generalized PCA problem

$$\mathbf{A}_D \mathbf{G}_D \mathbf{W}_D \mathbf{Z}_D = \mathbf{Z}_D \mathbf{\Lambda}_D, \quad (28)$$

where

$$\mathbf{A}_D = \frac{1}{S} \mathbf{P}_D \mathbf{P}_D^T, \quad (29)$$

yields principal components strictly driven by physics, since the geometry has been removed and design variables are weighted to zero.

A relevant consideration arises when using only lumped (scalar) physical data. If the number of such scalars is smaller than the number of design variables, the rank of $\mathbf{P}_D$ will be limited by the fewer of the two, resulting in a forced dimensionality reduction dictated by the size of the scalar data block rather than by a variance-based choice. Consequently, a minimum requirement for PD-PME to fully capture variability in the design vector is that the dimensionality of the physical information (e.g., the number of scalar data points) be at least equal to or larger than the number of design variables. Otherwise, the principal components will be constrained by the lower-dimensional block and may fail to represent all relevant modes of variation.

3 Test cases

To highlight the effectiveness and flexibility of the proposed approaches (PME, PI-PME, and PD-PME), four different test cases are considered. Each case targets a distinct geometry and set of operating conditions, thus allowing for a comprehensive assessment of how geometric and physical information can be combined (or, in the case of the physics-driven variant, used exclusively) to enable efficient design-space dimensionality reduction. The following subsections detail each test case in turn, describing both the parameterization strategies and the sources of physical data employed for PI-PME and PD-PME.

3.1 Test case 1: RAE-2822 airfoil

The first test case concerns the design optimization of the classical RAE-2822 airfoil (see Fig. 2, bottom left). The design space, defined within the activities of the NATO-AVT-331 Research Task Group [35], includes $M = 20$ design variables; each variable is associated with a different shape function (3 polynomials, 6 Hicks–Henne bumps, and 1 Wagner function [36], see Fig. 2, right), acting either on the upper or lower surface of the airfoil [37]. An in-house code (WG2AER, developed at CIRA) parameterizes the airfoil as a linear combination of the parent geometry $\mathbf{g}(\xi)$ and the modification functions δ .

Discretizing the airfoil with $L = 129$ grid points, the physical information required to train PI-PME and PD-PME is gathered from the XFOIL solver [38], under operating conditions of Mach = 0.4, Reynolds = 6.5×10^6 , and a zero-degree angle of attack. Collected physical data include the pressure coefficient (as distributed information, see Fig. 2 top left) and the lift, drag, and pitching moment coefficients (as lumped parameters).

3.2 Test case 2: autonomous underwater glider

The second test case addresses the design optimization of a bio-inspired autonomous underwater glider (AUG) with a manta-like shape, that has been selected as a test case for the NATO-AVT-404 Research Task Group on “Enhanced Design Processes of Military Vehicles through Machine Learning Methods”. The geometry is constructed as a continuous wing using a section-wise scheme and is by design spanwise symmetric; only half of the body is parameterized and then mirrored at the root section. As illustrated in Fig. 3, the model comprises three zones: the center body, the transition region, and the outer wing. Four sections define the half-span: (i) the root (section 1), (ii) the end of the main body (section 2), (iii) the end of the transition (section 3), and (iv) the tip (section 4). Each section is fully determined




Fig. 2 Test case 1: RAE-2822 (left) original geometry and pressure coefficient C_p with discretization and (right) shape modification functions/ original parameterization




Fig. 3 Test case 2: AUG (left) shape parameterization and (right) pressure coefficient and wake as solver outputs

by ten variables: four parameters for the section's NACA 4-digit airfoil (maximum camber m , maximum camber position p , thickness ratio t , chord c); three parameters for positioning the section's leading edge (x_0, z_0, s_0); three parameters defining the rotation angles (pitch ϑ , roll ϕ , yaw ψ), applied sequentially about the leading edge. Section 1 retains only 2 degrees of freedom (t and c), while all other parameters are fixed. Consequently, the complete parametric model includes $M = 32$ design variables. The manta-like geometry is constructed via the OpenCASCADE CAD kernel [39] and meshed with Gmsh v4.12.1 [40].

The glider surface mesh consists of $L = 784$ elements, and the physical data (pressure coefficients, lift, and drag) are acquired via the PUFFIN solver [41], developed by ENSTA Bretagne, which combines an incompressible potential approach with viscous corrections. The operating conditions correspond to a uniform inflow velocity of 0.25 m/s in seawater at a depth of 1500 m, and an 8-degree angle of attack that yields near-optimal efficiency for the baseline configuration.

3.3 Test case 3: ship propeller

Test case 3 considers the design optimization of a six-blade, right-handed marine propeller of a cruise ship. The reference propeller, shown in Fig. 4, was designed for a nominal advance coefficient ($J = V/nD$, where V is the advance speed, D the propeller diameter, and n the rate of revolution in rps) of about 0.87 and a cavitation index $\sigma_n = 2(p - p_{\text{vapor}})/(\rho n^2 D^2)$ of 2.25. In this functioning condition, the delivered thrust corresponds to a thrust coefficient ($K_T = T/(\rho n^2 D^4)$) of 0.186. The parametric description consists of B-Spline curves that describe the geometrical features, in radial and chordwise directions, of the propeller blade. This approach, extensively validated in several design-by-optimization cases [42, 43], provides a robust and easily controllable representation of the blade geometry in terms of its fairness. Moreover, it allows for the independent choice of quantities, like maximum sectional thickness, using criteria other than hydrodynamic shape optimization (in this case, structural strength of the blade)





Fig. 4 Test case 3: geometry of the reference ship propeller; “Key-blade” with the surface panel discretization used for BEM calculations of performance

which instead are hardly controllable when using free-form deformations or B-Surfaces describing directly the suction and pressure side of the blade.


The B-Spline control polygons define radial distributions of non-dimensional chord (c/D , parameter name c), pitch

(P/D , parameter name pd), rake ($rake/D$, parameter name r), and maximum sectional camber (f_{max}/c , parameter name f) of Fig. 5a. The shape of sectional hydrofoils is included as well in the parametric model by describing, again by using appropriate control polygons (see Fig. 5b), the non-dimensional chordwise sectional thickness ($t/2t_{max}$, parameter name th) and camber (f/f_{max} , parameter name fh) distributions of the blade profile which, radial section per radial section, are finally scaled with the corresponding maximum values. Given some constraints (the radial position of control points at root and tip is fixed, leading- and trailing-edge points are always given, and the chord at the tip never changes to comply with some limits of the flow solver adopted for the hydrodynamic characterization of performance), control points, sequentially numbered from root to tip (or from leading- to trailing-edge), are free to move within assigned ranges. Subscript “ x ” indicates the radial, or chordwise, modification of the control point position, and subscript “ y ” refers instead to changes in the quantity the control point describes. This parametrization leads to a total of $M = 38$ design variables (5 for the chord, 8 for the maximum sectional camber and pitch each, 6 for the rake, 6 for the sectional camber line, and 5 for the sectional thickness non-dimensional distribution) which describe

Fig. 5 Test case 3: distributions of the geometrical features of the propeller



(a) Radial distributions (chord, pitch, rake, and maximum sectional camber, from left to right)



(b) Sectional distributions (chordwise thickness and camber line, left and right)

entirely the blade shape with the exception of the skew that is maintained unchanged and identical to that of the reference propeller. For the hydrodynamic analyses, a low-fidelity boundary element method (BEM) is employed. The BEM code [44] was developed at the University of Genoa since early 2,000 for analysis and design-by-optimization purposes. It is a Morino, Dirichlet-type boundary condition implementation of panel methods for incompressible, potential flow solution, which makes use of the “key blade” approach to deal with stationary or unsteady problems [45]. It includes a cavitation model (the sheet cavity model at leading edge and midchord, both on suction and pressure side proposed in [46]), wake alignment capabilities and the iterative Kutta condition. For current analyses, the “key blade” is discretized with 1,250 hyperboloidal panels ($L = 1,326$ nodes). The collection of physical information was carried out by considering several equivalent steady-state operating conditions representative of the most critical conditions encountered by the blade during a revolution, given the design of the reference propeller, which was developed to operate in the spatially non-uniform wake of the ship shown in Fig. 6.

Despite the computational efficiency of the BEM, addressing the truly unsteady performance of the propeller is excessively demanding for the characterization of the thousands of different configurations needed to feed the physically-informed parametric model embedding methods. Conversely, steady analyses under equivalent flow conditions have consistently yielded the performance indicators of conventional design tools [47, 48] and have been effectively utilized in numerous designs by optimization processes [49] as a surrogate for unsteady functioning predictions. In addition to the equivalent nominal functioning condition (i.e., the propeller operating the circumferentially averaged axial inflow wake), two additional calculations were included in

the analyses to provide a broader view of the propeller performance. One corresponds to the loaded case of the blade passing through the 90° position of the non-uniform inflow wake, where the action of tangential velocities increases the angle of attack. The other, on the contrary, addresses the unloaded blade when at 270° into the wake, where instead the tangential velocities act to reduce the angle of attack. Together, they allow collecting distributed pressure data as well as lumped performance indicators such as thrust coefficient K_T , efficiency η , and tip-vortex intensity Γ related to the risk of different types of cavitation (suction or pressure side, respectively loaded and unloaded equivalent conditions) and to expected unsteady behaviour of the propeller.

3.4 Test case 4: 5415 destroyer-type vessel

The fourth test case targets the international benchmark 5415 hull [50], which is a geosymmetric replica of the DDG-51 (a US Navy destroyer-type vessel). Shape parameterization is based on a recursive set of $M = 27$ global modification functions applied within a hyper-rectangular region enclosing the demi-hull [21, 51]. Figure 7 (left) shows an example of the shape modification magnitude for one design variant.

The surface mesh of the hull comprises $L = 2250$ grid nodes. Physical data include calm-water and seakeeping performance, extracted from the in-house linear potential-flow solver WARP [52] (developed by CNR-INM) with viscous corrections, alongside a linearized strip-theory code SMP [53]. Calm-water analyses provide the distribution of surface pressure (see Fig. 7, right) and the wave-resistance coefficient; seakeeping analyses deliver estimates of the pitch motion (θ) RMS and the vertical acceleration (a_B) RMS at the bridge. These integrated and distributed quantities collectively serve as physical inputs for PI-PME and PD-PME, allowing for a thorough exploration of how variations in hull geometry correlate with performance in real sea conditions.

4 Results

To train the standard PME and its physics-informed extensions (PI-PME and PD-PME), an initial set of $S = 16385$ samples was generated for each test case. To ensure a uniform and representative coverage of the design space, the initial dataset is generated using a Sobol low-discrepancy sequence. This approach provides better space-filling properties than purely random sampling and avoids clustering, which is critical to reliably constructing the reduced manifold. In addition to covering the geometric design space, physical simulations were performed at these samples whenever possible, providing distributed or lumped





Fig. 6 Non-uniform inflow (ship wake) to the propeller, seen from aft

Fig. 7 Test case 4: example of (left) shape modification magnitude with discretization and (right) calm-water numerical solution for the pressure distribution and the wave elevation pattern η generated by the 5415 hull



physical data. Any geometries that produced unfeasible outcomes (e.g., degenerate shapes, solver divergence, or NaN outputs) were discarded. This filtering process alone already serves as a preliminary assessment of the chosen parameterization, as it highlights whether certain geometric variations are physically inadmissible or numerically unstable.

An additional layer of screening was then applied to exclude extreme outliers. Specifically, all geometric–physical samples were retained only if the corresponding physical outputs fell within the ranges $[Q1 - 3IQR, Q3 + 3IQR]$, where $Q1$ and $Q3$ are the first and third quartiles, respectively, and $IQR = Q3 - Q1$. This distribution-agnostic rule was selected to balance permissiveness and robustness: it allows broad variability across designs while excluding pathological outliers. As a result, only physically credible shapes and simulations populate the final data set used for dimensionality reduction. The filtering strategy is intentionally cautious, aiming to eliminate anomalous or ill-posed configurations that could contaminate the reduced basis with non-representative trends. It is finally worth noting that, while the reduced space may still contain directions that could lead to poorly behaved geometries, standard constraint-handling techniques in the subsequent optimization process can be employed to reject such configurations if encountered.

Figure 8 presents how the fraction of variance (σ^2) evolves as a function of the number of reduced components N , contrasting PME (purely geometric variance), PI-PME (geometry + physical data), and PD-PME (physical data only). These results are numerically summarized in Table 2 (e.g., how many modes are needed to reach 95% of the total variance). Meanwhile, Figs. 9, 10, 11, 12 provide deeper insights into how that variance is organized for each mode: on the left, each figure shows the normalized eigenvector components $|v_k|/\max_k(|v_{ik}|)$ along the vertical axis, plotted across the original design variables on the horizontal axis. Large peaks indicate which design variables dominate that particular mode; on the right is shown the participation of

each data source (e.g., geometry d , distributed pressure c_p , or lumped coefficient like c_D , c_L , etc.) to variance retained by each mode. The latter plots reveal whether a mode primarily represents geometric variation, physical variation, or a mixture of both.

Taken together, these figures clarify whether geometry and physics are strongly correlated (leading to fewer modes or shared modes) or largely uncorrelated (leading to more modes, each distinctly owned by geometry or physics). Below is a case-by-case discussion.

PME, PI-PME, and PD-PME curves in the RAE airfoil problem lie close together (see Fig. 8a); each method needs a comparable number of modes ($N = 6$ or 7) to reach 95% or 99% variance. This suggests that geometric changes coincide strongly with variations in the pressure distribution and integrated aerodynamic coefficients. Figure 9 (left) shows a pretty similar embedding with eigenvectors almost coincident for PME, PI-PME, and PD-PME (see first and third mode), in particular, it can be seen how the first mode is mainly participated by u_2 , u_4 , and u_6 , that correspond to the polynomial shape modification function applied to the lower side of the airfoil, highlighting also how the lower side modification is the major region affecting the aerodynamic performance. This is further confirmed by Fig. 9 (right) where is shown that geometry (d) and physical parameters (c_p , c_L , c_D , and c_M) share most of the first modes: apart from mode 2, that is mainly participated by the drag coefficient, no distinct pure physics or pure geometry mode dominates the decomposition. Because airfoil shape changes strongly affect the aerodynamic response under these conditions, the geometry and physics are highly correlated. Hence, adding physical data (PI-PME) does not inflate the necessary dimensionality, and removing geometry altogether (PD-PME) does not drastically reduce the dimension beyond PME, meaning that both geometry and physics point in similar directions in the design space.

For the glider test case, the PME rapidly achieves a high fraction of geometric variance ($N = 8$ modes for 95%),




Fig. 8 Variance resolved as a function of the reduced design variables

Table 2 Design-space dimensionality reduction summary for retaining 95% of the problem variance

TC	Geometry	S	M	N (dimensionality reduction%)		
#	model	Samples	Original	PME	PI-PME	PD-PME
1	RAE-2822	6,022	20	6 (70%)	7 (65%)	6 (70%)
2	AUG	5,330	32	8 (75%)	11 (66%)	5 (84%)
3	Propeller	13,615	38	8 (79%)	12 (68%)	8 (79%)
4	5415	5,842	27	14 (48%)	12 (66%)	7 (74%)

whereas the PI-PME requires more modes ($N = 11$). In contrast, PD-PME dips below the PME line, reaching 95% in fewer modes ($N = 5$), as shown in Fig. 8b. Figure 10, shows, on the left, that while for PME the highest variance (retained by the first mode) is mainly due to the span ($s_{0,4}$) of the tip section, the variance of PI-PME and PD-PME are mainly affected by the maximum camber and the pitch angle of the end of the transition from the center body to the outer wing (m_3 and ϑ_3). As shown in Fig. 10 (right) the first modes appear “decoupled” (i.e., geometry heavy or physics heavy): for both PI-PME and PD-PME the first mode is mainly participated by global forces (L and D), the second mainly by the pressure distribution c_p , and the third mainly by the geometry (for PI-PME only). There is a partial uncorrelation between geometry and physics: large shape modifications do not necessarily produce proportional changes in lift or drag. Hence, PI-PME must retain more overall modes to capture both geometry-dominated and physics-dominated

directions. Meanwhile, if the goal is purely performance-driven, PD-PME focuses on those physical variations alone, achieving a stronger reduction with fewer modes.

Concerning the propeller test case, PME captures geometric variability well in about $N = 8$ modes, but PI-PME demands significantly more ($N = 12$) to cover combined variance from geometry plus thrust, efficiency, vortex data, and pressure distributions. PD-PME, by contrast, can exceed 95% in only $N = 8$ modes if focusing purely on physical data. Figure 11 shows similar results as for the glider. The first mode of PME is participated by the variation of the intermediate control point of the chord ($c_{y,2}$), which obviously causes the most relevant modifications of the propeller shape, interpretable as variations of the expanded area of the blades. The second and third modes of the PME identify the rake (through $r_{y,2}$ and $r_{y,4}$) as the second most influential parameter on the geometry variance. This behavior is also plausible, as the “shift” of the blade sections along




Fig. 9 Design-space dimensionality reduction (left) eigenvectors $\mathbf{v}_k$ that embed the original design variables and (right) participation to the variance retained by each eigenvector of geometrical and physical information for RAE-2822 (TC#1)




Fig. 10 Design-space dimensionality reduction (left) eigenvectors $\mathbf{v}_k$ that embed the original design variables and (right) participation to the variance retained by each eigenvector of geometrical and physical information for AUG (TC#2)

longitudinal plane, which is the modification induced by the rake, contributes more to the differences in the shapes of the blade than the localized modifications induced by changes in camber and thickness (by radial or sectional perturbations of control points). The first two modes of PI-PME and PD-PME are mainly participated by the physical information, while the third of PI-PME is mainly participated by the geometry, even if in this case a partial participation of the

physical quantities is visible. In this respect, it is interesting to observe the very nice correlations between geometrical features and physical information revealed by both PI-PME and PD-PME. The first mode for both approaches, which is mainly participated by efficiency, identifies the pitch, and in particular its values at the tip, as the most important factor responsible for the mode, which is a trend exactly in line with traditional design method outputs and guidelines. The




Fig. 11 Design-space dimensionality reduction (left) eigenvectors $\mathbf{v}_k$ that embed the original design variables and (right) participation to the variance retained by each eigenvector of geometrical and physical information for propeller (TC#3)




Fig. 12 Design-space dimensionality reduction (left) eigenvectors $\mathbf{v}_k$ that embed the original design variables and (right) participation to the variance retained by each eigenvector of geometrical and physical information for 5415 (TC#4)

second mode, for which again PI-PME and PD-PME are perfectly overlapped, establishes a correlation between the delivered thrust and the strength of the tip vortex circulation with the pitch and maximum values of camber again, which are, by analogy with 2D hydrofoils, responsible for sectional forces (lift) through the angle of attack (pitch) and lift coefficient at zero angle of attack (camber). When, instead, the focus is on physical information like the pressure coefficient,

physics-informed models highlight the influence of pitch (responsible for hydrofoil suction peaks at the leading edge) and camber (responsible for pressure distributions on the central part of the hydrofoil) through modes 4 to 8. When purely geometrical analyses were employed, quantities like maximum sectional camber and hydrofoil sectional shape (camber and thickness) were completely discarded, being, from a geometrical point of view, responsible for minimal

and localized shape variations, in favor of the rake that is instead responsible for large geometrical modification but has almost no influence on (any) propeller performance indicators. Only embedding the physical quantities activates these features, which, indeed, regardless of the small geometrical variations they induce, have a critical influence on pressure distributions and the overall performance of the propeller. Geometry and physical performance are weakly correlated. This necessitates many extra modes in PI-PME to accommodate both sets of unaligned directions. PD-PME proves highly compressible because the pure physics space (i.e., lumps + distributions) is lower-dimensional in meaningful directions for performance.

Finally, looking at the 5415 results (see Fig. 12, PME needs $N = 14$ modes to exceed 95% variance, PI-PME, interestingly, needs fewer ($N = 12$) modes to reach the same coverage, while PD-PME reaches 95% variance in only $N = 7$ modes. The first mode in PI-PME is driven almost exclusively by lumped seakeeping parameters. The second mode largely captures the calm-water wave-resistance coefficient. Starting from the third mode onward, geometry and distributed pressure tend to move “in lock-step,” indicating a tight correlation between shape changes and local pressure distribution. At first glance, one might expect adding physics increases the total variance, as per the other three test cases. However, the crucial factor is that many purely geometric variations do not correlate with changes in both calm-water and seakeeping physical quantities. When geometry is combined with physics in PI-PME, the modes that represent “empty” geometric variance (irrelevant to performance) no longer contribute significantly to the principal components. As a result, fewer total modes can end up describing all relevant variations.

5 Discussion

One of the primary motivations behind the proposed PI-/PD-PME extensions is to incorporate physical insight into the dimensionality reduction process without imposing prohibitive computational costs. While enriching the embedding with physical quantities (e.g., lift, drag, or pressure fields) might suggest an increased overhead, these quantities are obtained from low-fidelity solvers (e.g., XFOIL and PUF-FIN), which are computationally inexpensive and widely used in early-stage design. Importantly, in many practical workflows, these evaluations are already available from previous parametric studies, and thus no new high-fidelity simulations are required for training. This enables a cost-effective embedding that reflects relevant physical trends.

This rationale is further supported by an empirical analysis of variance convergence. For all test cases considered, the evolution of both geometric and physical variance was monitored as a function of the number of Sobol samples used in the training dataset. Despite generating up to 16,384 samples, it was observed that variance saturation typically occurs around 1,000 samples. This indicates that the main informative directions in the design space can be captured with a limited number of low-fidelity simulations. Such sampling budgets are compatible with standard preliminary design campaigns, even for problems involving more than 20 design variables. The convergence behavior is illustrated in Fig. 13, where both geometric and physical variance show rapid stabilization across all configurations. Once this saturation is reached, the structure of the reduced manifold becomes largely insensitive to additional samples, supporting the robustness of the method.

Moreover, investing in a physics-informed latent space leads to significant downstream savings. By embedding performance-related directions early, the method avoids wasting computational effort on uninformative or non-promising shape perturbations. This is confirmed, for instance, in the propeller case study, where the dominant




Fig. 13 Geometric and physical variance sensitivity to the number of samples. The curves show the relative error compared to the maximum sample size (16,384), used as reference

reduced basis direction consistently aligns with pitch at tip, a shape feature known to have strong influence on hydrodynamic performance. Similar trends are observed across other configurations, indicating that the reduced space prioritizes design variations with high physical relevance. This is particularly valuable when transitioning to high-fidelity or multi-objective optimization, where the cost of full evaluations is orders of magnitude higher. Unlike purely geometric PCA or POD, which may generate reduced spaces uncorrelated with performance, or black-box regressors that do not provide a true reduction of the design space, PI-/PD-PME yield interpretable, physics-aware coordinates that guide the design exploration more effectively.

It is also important to note that the proposed dimensionality reduction does not require highly accurate physical outputs: it relies on capturing global variance trends rather than absolute values. As long as the low-fidelity solver provides coherent gradients and relative sensitivities across the design space, the method effectively identifies the most informative directions. This assumption is supported by common practice in design optimization, where low-fidelity models are often benchmarked against high-fidelity responses along selected design directions to validate their qualitative behavior. Moreover, the dimensionality of the reduced space is deliberately kept higher than strictly necessary to mitigate potential inaccuracies in the low-fidelity-based embedding and to retain sufficient flexibility for high-fidelity refinement. This allows the embedding to serve not only as a compact representation of the space, but also as a candidate low-fidelity model in a multi-fidelity optimization framework. Furthermore, early sensitivity analyses are made easier by looking at, e.g., the 2D contour plots (see Fig. 14), which reveal how varying just one or two reduced coordinates shifts integral performance (e.g., drag, thrust, seakeeping motion) away from the baseline configuration. This “lumped-level” lens is invaluable for quick screening, where high-fidelity CFD simulations are too expensive to run exhaustively. By building on these low-fidelity models, designers can prune unpromising shape variations early on. Furthermore, this low-fidelity model can be already used to illustrate the trade-offs between multiple objectives (e.g., lowering drag vs. improving lift, increasing thrust vs. reducing vortex intensity, or minimizing calm-water resistance while controlling seakeeping motions), as well as providing a first Pareto front of optimal solutions (see Fig. 15).

This approach allows for quickly visualizing constraints, identifying high-potential regions, and serving as a stepping stone, guiding more elaborate, high-dimensional or high-fidelity optimization, focusing computational resources on the most promising design clusters.


Alongside lumped metrics, the embedding often includes distributed fields—pressure over the surface for the airfoil and

glider, or near-blade flows for the propeller. These data prove critical if the final goal is to train machine-learning reduced-order models (ROMs) capable of predicting integral and/or local quantities of interest. For instance, the pressure distribution, for the RAE-2822, can be aggregated into a small set of modes that reconstruct aerodynamic loads, while, for the propeller, it provides richer mode shapes and ensures variations that affect vortex shedding, cavitation risk, or tip flow details. Since the embeddings reduce the dimension of these distributed fields, surrogate ROMs can be trained more efficiently, predicting crucial integral outputs without re-running the solver for every shape perturbation.


Finally, a distinctive feature of the PI-/PD-PME framework is its ability to accommodate a large number of physical quantities without requiring prior manual selection. When the most informative observables are not known in advance, the method identifies dominant contributors based on their empirical variance and statistical correlations. It should be noted, however, that physical metrics that are redundant or irrelevant for performance are not automatically downweighted: if repeated or exhibiting high variance, they may still dominate the embedding. As such, proper normalization and careful feature selection remain important to avoid unintentional overemphasis. Nonetheless, this unsupervised, data-driven formulation enables the identification of performance-relevant directions and supports the exploration of complex design spaces where physical intuition alone may be insufficient. The relative impact of each physical feature on the reduced modes is visualized in Figs. 9, 10, 11, 12.

6 Conclusions


This work has demonstrated how the parametric model embedding (PME) framework can be extended to include physical information—either partially (PI-PME) or exclusively (PD-PME)—in a manner that seamlessly combines geometry and performance data within a reduced design space for shape optimization. Across a range of test cases, including the RAE-2822 airfoil, a bio-inspired underwater glider, a ship propeller, and the DTMB-5415 destroyer-type hull, the results showed that adding physics to PME does not always raise the number of modes needed for a target variance level, especially when physical behavior strongly correlates with geometric modifications. In cases where geometry and physics are poorly aligned, however, PI-PME incorporates both purely geometric and purely physical directions, inevitably increasing the dimensionality. This difference reveals how important it is to identify whether the chosen geometric parameterization meaningfully covers the performance-critical aspects of the design.




(a) TC#1: RAE-2822



(b) TC#2: AUG



(c) TC#3: Propeller



(d) TC#4: 5415

Fig. 14 Contour maps of lumped physical information as a function of the first two reduced design variables $\mathbf{x}$




Fig. 15 Scatter plot of lumped physical parameters usable as objective functions in a multi-objective problem

PI-PME preserves the comprehensive interplay between shape and physics and is particularly useful when geometric feasibility, manufacturing requirements, and interlinked performance metrics must remain in the same design space. PD-PME, by removing geometry, directly focuses on physics-driven changes and can yield a reduced representation when many shape variations do not meaningfully affect integral objectives such as resistance, lift, or thrust. In this role, PD-PME also serves as a diagnostic to highlight situations where geometric degrees of freedom provide little to no performance advantage. Lumped physical quantities—such as lift, drag, thrust, or calm-water resistance—can be incorporated at lower fidelity for early screening or coarse optimization, while distributed fields like surface pressure or vortex data enrich the understanding of local flow

phenomena and facilitate training machine-learning-based reduced-order models for predicting global performance measures.

A clear benefit of these methods lies in their interpretability. By revealing which original design variables contribute most strongly to each embedding mode, the extended PME approaches guide engineers to focus on shape changes that matter for performance, rather than for mere geometric variety. Early-stage screening becomes more efficient, as shown by low-dimensional contour plots and Pareto fronts generated from PI-PME or PD-PME analyses, allowing for fast identification of promising design regions without incurring the costs of exhaustive high-fidelity simulation. Ultimately, whether one relies on PME alone (when geometry largely dictates performance), integrates physical data

using PI-PME (when shape–physics coupling is important), or centers on PD-PME (when physics alone can drive meaningful decisions), these variations collectively form a flexible toolkit for modeling, analyzing, and optimizing complex design spaces.

Finally, while the proposed methodology effectively incorporates physical information into the reduced space, a limitation lies in its inherently linear structure. In practical applications, both geometric deformations and physical responses often exhibit nonlinear behavior that may not be fully captured by linear embeddings. PME does not impose explicit assumptions regarding the smoothness or continuity of the geometric or physical fields. Its basis construction is entirely data-driven and reflects the empirical variance across samples. As a linear technique, PME can still capture sharp gradients or localized features-provided they manifest as dominant directions of variability. However, in design spaces characterized by strong discontinuities or highly nonlinear relationships, its effectiveness may be limited, and nonlinear extensions could provide more expressive representations. Therefore, future research should explore nonlinear extensions of PME by leveraging advanced dimensionality reduction techniques, such as local PCA [54], kernel PCA [31, 55], autoencoders [56–61], or principal geodesic analysis [62, 63], which allow for the discovery of curved low-dimensional manifolds. These techniques may further enhance the ability of the framework to capture complex dependencies between design variables and physical observables, improving both the fidelity and generalizability of the reduced representations. Nonetheless, to the best of the authors' knowledge, none of these nonlinear approaches currently support the core capabilities offered by PME-namely, the ability to assign custom weights across heterogeneous features and to perform analytical backmapping to the original parametric space. Developing nonlinear extensions of PI-/PD-PME that retain such properties represents a compelling and largely unexplored direction for future research.

Acknowledgements This work has been conducted within the NATO-AVT-404 Research Task Group on “Enhanced Design Processes of Military Vehicles through Machine Learning Methods”. CNR-INM and the University of Genoa acknowledge the support of the Italian Ministry of University and Research (MUR) through the National Recovery and Resilience Plan (PNRR), Sustainable Mobility Center (CNMS), Spoke 3 Waterways, CN00000023 - CUP B43C22000440001. CNR-INM is also grateful to MUR through PRIN 2022 program, project BIODRONES, 20227JNM52 - CUP B53D23005560006. CNR-INM finally acknowledges the CINECA award under the ISCRA initiative, for the availability of high-performance computing resources and support. CIRA work was supported by the CIRA internal project OPTI-WING (OPTImization for WINg Generation).

Author Contributions A.S. conceived and developed the methodology, implemented the models, performed all simulations and analyses, wrote the manuscript, and prepared the figures. M.D. jointly contrib-

uted to the methodological framework, supervised the research, and secured funding. D.Q. defined the RAE-2822 test case, providing solvers and parameterization tools. G.P. co-defined the AUG test case with A.S. and provided the geometry parameterization. S.G. defined the propeller test case, supplied the geometry and dataset, and performed the related simulations. J.W. reviewed and revised the manuscript. All authors reviewed and approved the final version of the manuscript.

Funding Open access funding provided by Consiglio Nazionale Delle Ricerche (CNR) within the CRUI-CARE Agreement.

Data availability The data that support the findings of this study are available from the corresponding author, A. S., upon reasonable request.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Bellman R (1957) Dynamic programming. Princeton University Press. Princeton, NJ, pp. 4–9
2. Serani A, Diez M (2025) A survey on design-space dimensionality reduction methods for shape optimization. Arch Computat Methods Eng. <https://doi.org/10.1007/s11831-025-10349-x>
3. Yonekura K, Watanabe O (2014) A shape parameterization method using principal component analysis in applications to parametric shape optimization. J Mech Des 136(12):121401. <http://doi.org/10.1115/1.4028273>
4. Diez M, Campana EF, Stern F (2015) Design-space dimensionality reduction in shape optimization by karhunen-loève expansion. Comput Methods Appl Mech Eng 283:1525–1544. <https://doi.org/10.1016/j.cma.2014.10.042>
5. Constantine PG, Dow E, Wang Q (2014) Active subspace methods in theory and practice: applications to kriging surfaces. SIAM J Sci Comput 36(4):1500–1524. <https://doi.org/10.1137/130916138>
6. Berguin SH, Mavris DN (2015) Dimensionality reduction using principal component analysis applied to the gradient. AIAA J 53(4):1078–1090. <https://doi.org/10.2514/1.J053372>
7. Harries S, Uharek S (2021) Application of radial basis functions for partially-parametric modeling and principal component analysis for faster hydrodynamic optimization of a catamaran. J Marine Sci Eng 9(10):1069. <https://doi.org/10.3390/jmse9101069>
8. Jun T, Gang S, Liqiang G, Xinyu W (2020) Application of a pca-dbn-based surrogate model to robust aerodynamic design optimization. Chin J Aeronaut 33(6):1573–1588. <https://doi.org/10.1016/j.cja.2020.01.015>
9. Zhang Y, Ma N, Gu X, Shi Q (2024) Geometric space construction method combined of a spline-skinning based geometric variation method and pca dimensionality reduction for ship hull form

- optimization. *Ocean Eng* 302:117604. <https://doi.org/10.1016/j.oceaneng.2024.117604>
10. Poole D, Allen C, Rendall T (2017) High-fidelity aerodynamic shape optimization using efficient orthogonal modal design variables with a constrained global optimizer. *Comput Fluids* 143:1–15. <https://doi.org/10.1016/j.compfluid.2016.11.002>
 11. Allen CB, Poole DJ, Rendall TC (2018) Wing aerodynamic optimization using efficient mathematically-extracted modal design variables. *Optim Eng* 19:453–477. <https://doi.org/10.1007/s11081-018-9376-7>
 12. Poole DJ, Allen CB, Rendall TC (2022) Efficient aeroelastic wing optimization through a compact aerofoil decomposition approach. *Struct Multidiscip Optim* 65(3):81. <https://doi.org/10.1007/s00158-022-03174-4>
 13. Zhang L, Mi D, Yan C, Tang F (2018) Multidisciplinary design optimization for a centrifugal compressor based on proper orthogonal decomposition and an adaptive sampling method. *Appl Sci* 8(12):2608. <https://doi.org/10.3390/app8122608>
 14. Ballarin F, D'Amario A, Perotto S, Rozza G (2019) A pod-selective inverse distance weighting method for fast parametrized shape morphing. *Int J Numer Meth Eng* 117(8):860–884. <https://doi.org/10.1002/nme.5982>
 15. Yanhui D, Wenhua W, Zhang P, Fulin T, Zhaolin F, Guiyu Z, Jiaqi L (2019) Performance improvement of optimization solutions by pod-based data mining. *Chin J Aeronaut* 32(4):826–838. <https://doi.org/10.1016/j.cja.2019.01.014>
 16. Yamazaki W (2020) Efficient multi-objective shape optimization using proper orthogonal decomposition with variable fidelity concept. *J Adv Mech Design Syst Manufact* 14(1):0019–0019. <https://doi.org/10.1299/jamdsm.2020jamdsm0019>
 17. Yang Y, Xue Y, Zhao W, Yang H, Wu C (2024) Aerodynamic shape optimization based on proper orthogonal decomposition reparameterization under small training sets. *Aerosp Sci Technol* 147:109072. <https://doi.org/10.1016/j.ast.2024.109072>
 18. Li J-L, Zhang Y, Zhu B-B, Pang B, Chen G (2024) The aerodynamic optimization of hypersonic vehicles with the proper-orthogonal-decomposition-based CST method. *Aerosp Sci Technol*:109295. <https://doi.org/10.1016/j.ast.2024.109295>
 19. D'Agostino D, Serani A, Diez M (2020) Design-space assessment and dimensionality reduction: an off-line method for shape reparameterization in simulation-based optimization. *Ocean Eng* 197:106852. <https://doi.org/10.1016/j.oceaneng.2019.106852>
 20. Chang H, Wang C, Liu Z, Feng B, Zhan C, Cheng X (2023) Research on the karhunen-loève transform method and its application to hull form optimization. *J Marine Sci Eng* 11(1):230. <https://doi.org/10.3390/jmse11010230>
 21. Serani A, Stern F, Campana EF, Diez M (2022) Hull-form stochastic optimization via computational-cost reduction methods. *Eng Comput* 38(Suppl 3):2245–2269. <https://doi.org/10.1007/s00366-021-01375-x>
 22. Diez M, Pellegrini R, Serani A, Stern F (2023) Design-space dimensionality reduction in structural optimization via parametric model embedding. In: *Proceedings of the 10th International Conference on Computational Methods in Marine Engineering (MARINE 2023)*, Madrid, Spain
 23. Li M, Yu X, Meng D, An G, Liu B (2023) A new approach for deviation modeling in compressors: sensitivity-correlated principal component analysis. *Aerospace* 10(5):491. <https://doi.org/10.3390/aerospace10050491>
 24. Zhang C, Duan Y, Chen H, Lin J, Xu X, Wang G, Liu S (2023) Efficient aerodynamic shape optimization with the metric-based pod parameterization method. *Struct Multidiscip Optim* 66(6):140. <https://doi.org/10.1007/s00158-023-03596-8>
 25. Mingzhi L, Xianjun Y, Dejun M, Guangfeng A, Baojie L (2024) Breaking the geometry-performance tradeoff in compressor deviation modeling: nested principal component analysis. *Chin J Aeronaut* 37(9):131–149. <https://doi.org/10.1016/j.cja.2024.06.012>
 26. Khan S, Kaklis P, Serani A, Diez M, Kostas K (2022) Shape-supervised dimension reduction: extracting geometry and physics associated features with geometric moments. *Comput Aided Des* 150:103327. <https://doi.org/10.1016/j.cad.2022.103327>
 27. Khan S, Kaklis P, Serani A, Diez M (2022) Geometric moment-dependent global sensitivity analysis without simulation data: application to ship hull form optimisation. *Comput Aided Des* 151:103339. <https://doi.org/10.1016/j.cad.2022.103339>
 28. Masood Z, Kostas KV, Khan S, Kaklis PD (2023) Shape-informed dimensional reduction in airfoil/hydrofoil modeling. *J Marine Sci Eng* 11(10):1851. <https://doi.org/10.3390/jmse11101851>
 29. Kostas KV, Manousaridou M (2023) Machine-learning-enabled foil design assistant. *J Marine Sci Eng* 11(7):1470. <https://doi.org/10.3390/jmse11071470>
 30. Masood Z, Usama M, Khan S, Kostas K, Kaklis PD (2024) Generative vs. non-generative models in engineering shape optimization. *J Marine Sci Eng* 12(4): 566. <https://doi.org/10.3390/jmse12040566>
 31. Gaudrie D, Le Riche R, Picheny V, Enaux B, Herbert V (2020) Modeling and optimization with gaussian processes in reduced eigenbases. *Struct Multidiscip Optim* 61:2343–2361. <https://doi.org/10.1007/s00158-019-02458-6>
 32. Serani A, Diez M (2023) Parametric model embedding. *Comput Methods Appl Mech Eng* 404:115776. <https://doi.org/10.1016/j.cma.2022.115776>
 33. Serani A, Diez M, Quagliarella D (2024) Aerodynamic shape optimization in transonic conditions through parametric model embedding. *Aerosp Sci Technol* 155:109611. <https://doi.org/10.1016/j.ast.2024.109611>
 34. Pellegrini R, Diez M, Kubina ER, Stern F (2025) On the use of parametric model embedding for structural optimization with application to high-speed small craft. In: *Research Specialists Meeting AVT-411 on Machine Learning and Artificial Intelligence for Military Vehicle Design*
 35. Beran PS, Bryson D, Thelen AS, Diez M, Serani A (2020) Comparison of multi-fidelity approaches for military vehicle design. In: *AIAA AVIATION 2020 FORUM*, pp. 3158
 36. Hicks R, Henne PA (1978) Wing design by numerical optimization. *J Aircr* 15(7):407–412
 37. Quagliarella D, Clark D, Bryson D, Beran PS, Thelen A, Mainini L, Yildiz S, Nikbay M, Minisci E, Leyland P, Serani A, Diez M (2022) Reproducible industrial multifidelity optimization benchmark problems for air, space, and sea vehicles. In: *Research Workshop AVT-354 on Multifidelity Methods for Military Vehicle Design*
 38. Drela M (1989) Xfoil: an analysis and design system for low reynolds number airfoils. *Low Reynolds Number Aerodynamics* 54
 39. Open Cascade (2024) Open cascade technology: the open source platform for 3D CAD, CAM, CAE. <https://www.opencascade.com/>. Accessed: 2024-07-17
 40. Geuzaine C, Remacle J-F (2009) Gmsh: a 3-d finite element mesh generator with built-in pre-and post-processing facilities. *Int J Numer Meth Eng* 79(11):1309–1331. <https://doi.org/10.1002/nme.2579>
 41. Perali P, Sacher M, Leroux J-B, Wackers J, Augier B, Hauville F, Bot P (2024) Performance prediction of a hydrofoil near the free surface using low (bem) and high (rans) fidelity methods. *Appl Ocean Res* 151:104157. <https://doi.org/10.1016/j.apor.2024.104157>
 42. Bertetta D, Brizzolaro S, Gaggero S, Viviani M, Savio L (2012) Cpp propeller cavitation and noise optimization at different pitches with panel code and validation by cavitation tunnel

- measurements. *Ocean Eng* 53:177–195. <https://doi.org/10.1016/j.oceaneng.2012.06.026>
43. Gaggero S (2020) Numerical design of a rim-driven thruster using a rans-based optimization approach. *Appl Ocean Res* 94:101941. <https://doi.org/10.1016/j.apor.2019.101941>
 44. Gaggero S, Villa D, Brizzolara S (2010) Rans and panel method for unsteady flow propeller analysis. *J Hydrodyn* 22(5 SUPPL. 1):564–569. [https://doi.org/10.1016/S1001-6058\(09\)60253-5](https://doi.org/10.1016/S1001-6058(09)60253-5)
 45. Kinnas SA, Hsin C-Y (1992) Boundary element method for the analysis of the unsteady flow around extreme propeller geometries. *AIAA J* 30(3):688–696. <https://doi.org/10.2514/3.10973>
 46. Fine NE (1992) Nonlinear analysis of cavitating propellers in nonuniform flow. In: PhD thesis, Massachusetts institute of Technology
 47. Lerbs H (1952) Moderately loaded propellers with a finite number of blades and an arbitrary distribution of circulation. *SNAME Trans* 60
 48. Epps BP, Kimball RW (2013) Unified rotor lifting line theory. *J Ship Res* 57(04):181–201. <https://doi.org/10.5957/jsr.2013.57.4.181>
 49. Gaggero S (2025) Robust simulation-based design optimization of marine propellers. *Ocean Eng* 321:120397. <https://doi.org/10.1016/j.oceaneng.2025.120397>
 50. Serani A, Diez M (2024) Hydrodynamic shape optimization of a naval destroyer by machine learning methods. *J Marine Sci Eng* 12(11):1979. <https://doi.org/10.3390/jmse12111979>
 51. Serani A, Fasano G, Liuzzi G, Lucidi S, Iemma U, Campana EF, Stern F, Diez M (2016) Ship hydrodynamic optimization by local hybridization of deterministic derivative-free global algorithms. *Appl Ocean Res* 59:115–128. <https://doi.org/10.1016/j.apor.2016.04.006>
 52. Bassanini P, Bulgarelli U, Campana EF, Lalli F (1994) The wave resistance problem in a boundary integral formulation. *Surv Math Ind* 4:151–194
 53. Meyers WG, Baitis AE (1985) SMP84: improvements to capability and prediction accuracy of the standard ship motion program SMP81. In: Technical Report SPD-0936-04, David Taylor Naval Ship Research and Development Center
 54. D'Agostino D, Serani A, Campana EF, Diez M (2018) Nonlinear methods for design-space dimensionality reduction in shape optimization. In: Machine Learning, Optimization, and Big Data: Third International Conference, MOD 2017, Volterra, Italy, September 14–17, 2017, Revised Selected Papers 3, pp. 121–132. https://doi.org/10.1007/978-3-319-72926-8_11
 55. Zhao H, Gong Z, Gan K, Gan Y, Xing H, Wang S (2024) Supervised kernel principal component analysis-polynomial chaos-kriging for high-dimensional surrogate modelling and optimization. *Knowl Based Syst* 305:112617. <https://doi.org/10.1016/j.knosys.2024.112617>
 56. Seo J, Kim D, Lee I (2024) A study on ship hull form transformation using convolutional autoencoder. *J Comput Design Eng* 11(1):34–48. <https://doi.org/10.1093/jcde/qwad111>
 57. Abbas A, Rafiee A, Haase M (2023) Deepmorpher: deep learning-based design space dimensionality reduction for shape optimisation. *J Eng Des* 34(3):254–270. <https://doi.org/10.1080/09544828.2023.2192606>
 58. Yamazaki W, Buyanbaatar N (2023) Comparative study of dimension reduction methods for efficient design optimization. *J Adv Mech Design Syst Manufact* 17(3):0036–0036. <https://doi.org/10.1299/jamdsm.2023jamdsm0036>
 59. Karafi, Y., Moussaoui, Z., Abou El Majd, B.: Simultaneous surrogate modeling and dimension reduction using unsupervised learning application to parametric wing shape optimization. *Math Model Comput* 11(1): 154–165 (2024). <https://doi.org/10.23939/mmc2024.01.154>
 60. Boncoraglio G, Farhat C (2021) Active manifold and model-order reduction to accelerate multidisciplinary analysis and optimization. *AIAA J* 59(11):4739–4753. <https://doi.org/10.2514/1.J.060581>
 61. Kou J, Botero-Bolívar L, Ballano R, Marino O, De Santana L, Valero E, Ferrer E (2023) Aeroacoustic airfoil shape optimization enhanced by autoencoders. *Expert Syst Appl* 217:119513. <https://doi.org/10.1016/j.eswa.2023.119513>
 62. Bo P, Zhang Y, Junlin L, Xudong W, Chang M, Junqiang B (2024) Data-driven surrogate model for aerodynamic design using separable shape tensor method. *Chin J Aeronaut* 37(9):41–58. <https://doi.org/10.1016/j.cja.2024.03.014>
 63. Doronina OA, Lee B, Grey ZJ, Glaws A (2025) Aerodynamic sensitivities over separable shape tensors. *AIAA J*: 1–14. <https://doi.org/10.2514/1.J064749>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.