AERODYNAMIC SHAPE OPTIMIZATION OF RAE-2822 AIRFOIL BY PHYSICS-INFORMED PARAMETRIC MODEL EMBEDDING

AeroBest 2025
Lisboa, 22–24 April 2025

AERODYNAMIC SHAPE OPTIMIZATION OF RAE-2822 AIRFOIL BY PHYSICS-INFORMED PARAMETRIC MODEL EMBEDDING

Damiano Squillace[1], Umberto Iemma[1], Domenico Quagliarella[2], Matteo Diez[3], and Andrea Serani[3]

[1]Department of Civil, Computer Science and Aeronautical Technologies Engineering
Università Roma tre
{damiano.squillace,umberto.iemma}@uniroma3.it, https://www.uniroma3.it/en/

[2]CIRA — Italian Aerospace Research Centre
Via Maiorise, 1, 81043 Capua CE, Italy
d.quagliarella@cira.it, https://www.cira.it/en/

[3]National Research Council – Institute of Marine Engineering
Via di Vallerano, 139, 00128 Roma RM, Italy
{matteo.diez,andrea.serani}@cnr.it, https://www.inm.cnr.it/

Abstract. *The efficiency of optimization algorithms in engineering design is often hindered by the curse of dimensionality, where increasing problem complexity degrades algorithmic performance due to the high dimensionality of the design space. To address this challenge, dimensionality reduction techniques that minimize the number of variables or constrain their variability are essential. This study introduces physics-informed parametric model embedding (PI-PME), a novel approach specifically designed for aerodynamic shape optimization of transonic airfoils. PI-PME extends the parametric model embedding (PME) framework by integrating physics-based insights, enhancing its capability to guide the optimization process toward physically relevant variations in the design space. By leveraging principal component analysis, PI-PME efficiently reduces dimensionality while maintaining the ability to reconstruct variables in the original space, avoiding the limitations of purely geometry-based formulations. To demonstrate its potential, PI-PME is applied to the optimization of the RAE-2822 airfoil at Mach 0.3, a benchmark case in aerodynamic design. Results are compared against PME and full-domain optimization, showing that PI-PME not only reduces computational complexity but also identifies superior configurations in the reduced design space. These findings underscore the practicality and robustness of PI-PME for industrial applications requiring efficient and reliable optimization strategies.*

Keywords: Curse of Dimensionality, Dimensionality Reduction, Simulation-based Optimization, Aerodynamic Shape Optimization, Principal Component Analysis, Parametric Model Embedding

1 INTRODUCTION

In recent years, a new approach has emerged in the design and development cycle of industrial products, ranging from simple mechanical components to highly innovative aircraft. Whereas the construction of a physical or digital model was once considered the final step in the production process, it is now an integral part of the preliminary study, allowing significant reductions in both economic and time inefficiencies.

Simulation-Based Design Optimization (SBDO) involves defining an algorithm that, once the variables and their domain are set, enables the creation of a simulated version of the final product. By analysing the behaviour of this model under design conditions, key values can be extracted, leading to the optimization process, which focuses on identifying the design that provides the best outcome.

However, this process requires considerable computational time, especially when high-fidelity solvers are used to achieve the most accurate possible representation of physical phenomena. Fortunately, recent technological advancements have significantly reduced these times, enabling the study of problems characterized by an increasing number of variables. The expansion in the number of variables is particularly relevant for two reasons: the need to explore more complex domains in the search for previously unidentified optimal configurations and the possibility of analyzing highly innovative structures for which the necessary number of variables is not known a priori. However, this leads to the well-known Curse of Dimensionality (CoD) [1], where the performance of the optimization algorithm deteriorates as the design space domain increases. For this reason, it is crucial to study and develop techniques capable of reducing the dimensionality of the problem (Dimensionality Reduction, DR) before the optimization process, particularly in industrial contexts where available resources are limited. This challenge can be addressed in two ways: by reducing the total number of variables or by restricting the domain of individual variables.

The main methods currently in use are Sensitivity Analysis (SA) and Principal Component Analysis (PCA) [2]. In this work, the focus is on further developing the technique known as Parametric Model Embedding (PME) [3], which is based on PCA. PME offers the advantage of immediately identifying a correlation between the variables in the reduced space and those in the original space. This makes it a particularly effective tool in optimization processes, as it allows for easily constraining the problem to prevent invalid configurations. This method is also particularly useful for applications such as CAD, which rely on a predefined structure and cannot operate with variables in the reduced space. The objective of this work is to integrate the relevant physical information into PME to obtain a reduced design space that is potentially more effective in identifying the optimal configuration [4].

The paper is structured as follows: the second section introduces and explains the functioning of PME and its physics-informed version (PI-PME). Sections three and four present the case study, the tools used, and the various results obtained. Finally, the fifth section provides the conclusions.

2 DESIGN-SPACE DIMENSIONALITY REDUCTION

Consider a manifold $\mathcal{G}$, which represents the original geometry, whose coordinates in three-dimensional space are given by $\mathbf{g}(\boldsymbol{\xi}) \in \mathbb{R}^3$, where $\boldsymbol{\xi} \in \mathcal{G}$ are the curvilinear coordinates defined on $\mathcal{G}$. Assume that, for the purpose of shape optimization, $\mathbf{g}$ can be transformed into a deformed shape/geometry $\mathbf{g}'(\boldsymbol{\xi}, \mathbf{u})$ through:

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

where $\boldsymbol{\delta}(\boldsymbol{\xi}, \mathbf{u}) \in \mathbb{R}^3$ represents the resulting shape modification vector, which is determined by a chosen shape parameterization or modification technique, and $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^M$ is the design variable vector. Consider the identification of the optimal design through the following SBDO optimization problem.

$$\underset{\mathbf{u} \in \mathcal{U}}{\text{minimize}} f(\mathbf{u}) \tag{2}$$

where $\mathbf{u}$ is treated as an uncertain or random parameter due to epistemic uncertainty. The optimal design is unknown before optimization, but exists. The vector of the design variable $\mathbf{u}$ is assigned a probability density function (PDF) $p(\mathbf{u})$, reflecting the belief in finding the optimal solution within a specific region of the design space. This PDF can be based on prior knowledge or experience; if unavailable, a uniform distribution may be used. Once $p(\mathbf{u})$ is defined, the shape modification vector $\boldsymbol{\delta}$ becomes stochastic and can be analyzed as a random field, for example, using the Karhunen-Loève Expansion (KLE), equivalent to Proper Orthogonal Decomposition (POD).

2.1 From Karhunen-Loève Expansion to Parametric Model Embedding

The KLE provides an optimal basis for the linear representation of a random process in a Hilbert space. The goal is to decompose the vector $\hat{\boldsymbol{\delta}}(\boldsymbol{\xi}, \mathbf{u}) = \boldsymbol{\delta} - \langle \boldsymbol{\delta} \rangle$ into orthonormal basis functions $\boldsymbol{\phi}_k(\boldsymbol{\xi})$, such that the geometric variance is retained. The expansion is expressed as:

$$\hat{\boldsymbol{\delta}}(\boldsymbol{\xi}, \mathbf{u}) \approx \sum_{k=1}^{N} x_k(\mathbf{u}) \boldsymbol{\phi}_k(\boldsymbol{\xi}), \tag{3}$$

where $x_k(\mathbf{u})$ are the projection coefficients obtained by:

$$x_k(\mathbf{u}) = (\hat{\boldsymbol{\delta}}, \boldsymbol{\phi}_k)_\rho = \int_{\mathcal{G}} \rho(\boldsymbol{\xi}) \hat{\boldsymbol{\delta}}(\boldsymbol{\xi}, \mathbf{u}) \cdot \boldsymbol{\phi}_k(\boldsymbol{\xi}) \, \mathrm{d}\boldsymbol{\xi}. \tag{4}$$

The KLE aims to retain the maximum variance of the random process through the basis functions $\boldsymbol{\phi}_k(\boldsymbol{\xi})$. The total variance σ^2 is given by:

$$\sigma^2 = \sum_{k=1}^{\infty} \langle (\hat{\boldsymbol{\delta}}, \boldsymbol{\phi}_k)_\rho^2 \rangle. \tag{5}$$

The optimal basis functions $\boldsymbol{\phi}_k$ are solutions to the variational problem, maximizing the functional:

$$\begin{aligned} \underset{\boldsymbol{\phi} \in \mathcal{L}_\rho^2 \mathcal{G}}{\text{maximize}} \quad & \mathcal{J}(\boldsymbol{\phi}) = \langle (\hat{\boldsymbol{\delta}}, \boldsymbol{\phi}_k)_\rho^2 \rangle \\ \text{subject to} \quad & (\boldsymbol{\phi}, \boldsymbol{\phi})_\rho^2 = 1 \end{aligned} \tag{6}$$

which leads to the eigenvalue equation:

$$\mathcal{L} \boldsymbol{\phi}(\boldsymbol{\xi}) = \lambda \boldsymbol{\phi}(\boldsymbol{\xi}), \tag{7}$$

where λ are the eigenvalues representing the variance retained for each mode $\boldsymbol{\phi}_k(\xi)$.

Upon discretizing the domain and applying Monte Carlo sampling to the random variables, it is possible to move from the continuous KLE formulation to PCA. In PCA, having $\mathbf{d}(\boldsymbol{\xi}, \mathbf{u})$ as the discretization of $\boldsymbol{\delta}(\boldsymbol{\xi}, \mathbf{u})$ and sampling $\mathcal{U}$ by a statistically convergent number of realizations S, it is possible to define the data matrix $\mathbf{D}$:

$$\mathbf{D} = \begin{bmatrix} \hat{\mathbf{d}}_1^{(1)} & \dots & \hat{\mathbf{d}}_1^{(S)} \\ \vdots & \vdots & \vdots \\ \hat{\mathbf{d}}_n^{(1)} & \dots & \hat{\mathbf{d}}_n^{(S)} \end{bmatrix} = \begin{bmatrix} \hat{d}_{1,\xi_1}(\mathbf{u}_1) & & \hat{d}_{1,\xi_1}(\mathbf{u}_S) \\ \vdots & & \vdots \\ \hat{d}_{L,\xi_1}(\mathbf{u}_1) & & \hat{d}_{L,\xi_1}(\mathbf{u}_S) \\ \vdots & \dots & \vdots \\ \hat{d}_{1,\xi_n}(\mathbf{u}_1) & & \hat{d}_{1,\xi_n}(\mathbf{u}_S) \\ \vdots & & \vdots \\ \hat{d}_{L,\xi_n}(\mathbf{u}_1) & & \hat{d}_{L,\xi_n}(\mathbf{u}_S) \end{bmatrix} \tag{8}$$

where $\hat{\mathbf{d}} = \mathbf{d} - \langle \mathbf{d} \rangle$ (with $\langle \cdot \rangle$ the mean value) and $\hat{d}_{i,\xi_j}$ representing the j-th component of the shape modification vector associated with the element i.

To further improve the model's usability, especially in optimization contexts, the PCA methodology is extended into PME. The challenge in KLE/PCA-based dimensionality reduction lies in transforming back to the original design variables, which is addressed by PME. This is achieved by embedding the original design variables along with the shape modification vector in the data matrix $\mathbf{P}$:

$$\mathbf{P} = \begin{bmatrix} \mathbf{D} \\ \mathbf{U} \end{bmatrix} \quad \text{with} \quad \mathbf{U} = \begin{bmatrix} \hat{\mathbf{u}}^{(1)} & \dots & \hat{\mathbf{u}}^{(S)} \end{bmatrix} = \begin{bmatrix} \hat{u}_{1,1} & & \hat{u}_{1,S} \\ \vdots & \dots & \vdots \\ \hat{u}_{M,1} & & \hat{u}_{M,S} \end{bmatrix} \tag{9}$$

where $\mathbf{U}$ is the matrix of the original design variables to witch is associated a null weight such that

$$\mathbf{W}_{\mathrm{u}} = 0 \quad \text{and} \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_{\mathrm{d}} & 0 \\ 0 & \mathbf{W}_{\mathrm{u}} \end{bmatrix} \tag{10}$$

The eigenvalue problem for PME is then formulated as follows:

$$\mathbf{AGWZ} = \mathbf{Z}\boldsymbol{\Lambda} \quad \text{with} \quad \mathbf{A} = \frac{1}{S}\mathbf{P}\mathbf{P}^\top. \tag{11}$$

where

$$\mathbf{G} = \begin{bmatrix} \mathbf{G}_{\mathbf{d}} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \quad \text{and} \quad \mathbf{Z} = \begin{bmatrix} \mathbf{z}_1 & \dots & \mathbf{z}_S \end{bmatrix} \quad \text{with} \quad \mathbf{z}_k = \begin{bmatrix} \mathbf{e}_k \\ \mathbf{v}_k \end{bmatrix} \tag{12}$$

Here, $\mathbf{e}_k$ and $\mathbf{v}_k$ represent the eigenvector components associated to the shape modification $\mathbf{d}$ and design variable $\mathbf{u}$ vectors, respectively. Both $\mathbf{G}_{\mathbf{d}} = \mathrm{diag}(\mathbf{G}_1, ..., \mathbf{G}_n)$ and $\mathbf{W}_{\mathbf{d}} = \mathrm{diag}(\mathbf{W}_1, ..., \mathbf{W}_n)$ are block diagonal matrices with dimensionality $[nL \times nL]$, each block $[L \times L]$ being a diagonal matrix itself.

$$\mathbf{G}_k = \mathrm{diag}\left(\Delta\mathcal{G}_1, \dots, \Delta\mathcal{G}_L\right) \quad \text{and} \quad \mathbf{W}_k = \mathrm{diag}\left(\rho_1, \dots, \rho_L\right) \tag{13}$$

The solutions λ_k and $\mathbf{z}_k$ of Equation (11) are finally used to construct the reduced dimensionality representation of the original parameterization, and the reconstruction of the original variables is achieved as follows

$$\mathbf{u} \approx \hat{\mathbf{u}} = \langle \mathbf{u} \rangle + \sum_{k=1}^{N} x_k \mathbf{v}_k. \tag{14}$$

2.2 Physics Informed Parametric Model Embedding

The incorporation of physics-related information into the optimization problem represents the final step of the study proposed in this paper. Similarly to what was described in the previous section, physics is implemented in PME by defining a new matrix $\mathbf{P}_I$ as:

$$\mathbf{P}_I = \begin{bmatrix} \mathbf{D} \\ \mathbf{U} \\ \mathbf{F} \\ \mathbf{C} \end{bmatrix} \quad \text{with} \quad \mathbf{F} = \begin{bmatrix} | & \vdots & | \\ \hat{\mathbf{f}}_j^{(1)} & \dots & \hat{\mathbf{f}}_j^{(S)} \\ | & \vdots & | \end{bmatrix} \quad \text{and} \quad \mathbf{C} = \begin{bmatrix} | & \vdots & | \\ \hat{c}_j^{(1)} & \dots & \hat{c}_j^{(S)} \\ | & \vdots & | \end{bmatrix} \tag{15}$$

where $\mathbf{F}$ is a matrix containing distributed physical information (with $\hat{\mathbf{f}}_j = \mathbf{f}_j - \langle \mathbf{f}_j \rangle$) and $\mathbf{C}$ contains lumped or scalar quantities of interest $\hat{c}_j = c_j - \langle c_j \rangle$). As seen before, a block-diagonal weight matrix $\mathbf{W}_I$ is introduced:

$$\mathbf{W}_I = \begin{bmatrix} \mathbf{W}_\mathrm{d} & 0 & 0 & 0 \\ 0 & \mathbf{W}_\mathrm{u} & 0 & 0 \\ 0 & 0 & \mathbf{W}_\mathbf{f} & 0 \\ 0 & 0 & 0 & \mathbf{W}_c \end{bmatrix} = \mathrm{diag}(\mathbf{W}_\mathbf{d}, \underbrace{0}_{\mathbf{u}}, \mathbf{W}_\mathbf{f}, \mathbf{W}_c), \tag{16}$$

where the new physical information, distributed $\mathbf{F}$ and scalar $\mathbf{C}$, are weighted respectively through the new matrices $\mathbf{W}_\mathbf{f}$ and $\mathbf{W}_c$. Each block has to be set up to normalize its respective data by the inverse of the estimated variance, ensuring that geometric modifications and physical data all contribute comparably to the subsequent PCA procedure. A further formulation of Eq. 11 is than obtained:

$$\mathbf{A}_I \mathbf{G}_I \mathbf{W}_I \mathbf{Z}_I = \mathbf{Z}_I \mathbf{\Lambda}_I \tag{17}$$

where

$$\mathbf{A}_I = \frac{1}{S} \mathbf{P}_I \mathbf{P}_I^\mathsf{T}, \quad \mathbf{G}_I = \begin{bmatrix} \mathbf{G}_\mathbf{d} & 0 & 0 & 0 \\ 0 & \mathbf{I} & 0 & 0 \\ 0 & 0 & \mathbf{G}_\mathbf{f} & 0 \\ 0 & 0 & 0 & \mathbf{I} \end{bmatrix}, \quad \mathbf{Z}_I = [\mathbf{z}_{I,1} \dots \mathbf{z}_{I,S}], \quad \mathbf{z}_{I,k} = \begin{bmatrix} \mathbf{e}_k \\ \mathbf{v}_k \\ \boldsymbol{\phi}_k \\ \boldsymbol{\pi}_k \end{bmatrix}_I \tag{18}$$

In this context,$\mathbf{G}_\mathbf{f}$ represents the element size of the distributed physical vector, which does not necessarily coincide with $\mathbf{G}_\mathbf{d}$. Meanwhile, the identity block is also applied to the lumped scalars. The eigenvector solution $\mathbf{Z}_I$ subsequently provides a reduced representation of all data, where $\boldsymbol{\phi}_k$ and $\boldsymbol{\pi}_k$ serve as the eigenvector components embedding the distributed and lumped physical parameters. As a result, a lower-dimensional yet physics-informed space is obtained, facilitating further design-space exploration and optimization.

3 TEST CASE: OPTIMIZATION OF THE RAE-2822

The chosen problem aimed to optimize the shape of the RAE-2822 airfoil, which represents a classical test case for evaluating the performance of optimization algorithms in the aeronautical field. The operating conditions under which the experiments were conducted are Mach equal to 0.3 and Reynolds equal to 6.5×10^6, while the optimization problem was defined as follows:

$$\begin{aligned} \text{minimize} \quad & C_D(\mathbf{u}) \\ \text{subject to} \quad & \alpha = 0 \\ & C_L(\mathbf{u}) \geq 0.212 \\ & t/c = 0.1211 \\ & t_{85}/c \geq 0.02 \\ & \mathbf{u}_l \leq \mathbf{u} \leq \mathbf{u}_u \end{aligned} \tag{19}$$

where C_D represents the drag coefficient, while c denotes the airfoil chord. The decision was made to impose the constraint on the angle of attack α rather than on C_L due to time constraints related to the high-fidelity solver used (SU2). To still obtain meaningful results, a constraint on C_L was imposed to identify only configurations with aerodynamic performance better than the original airfoil. Additionally, constraints were imposed on the maximum thickness-to-chord ratio (t/c) and the minimum thickness at 85% of the chord.

The design space was defined by 20 variables, each associated with a different shape function applied to the upper and lower surfaces of the airfoil. In this case, the selected functions are: 3 polynomial functions, 6 Hicks-Henne bump functions, and one Wagner function, shown in Figure 1 and described as follows:

$$\begin{cases} x = (0.5x^3 - 1.5x^2 + x) \cdot 0.52 \\ x = -0.4x^2 + 0.x \\ x = (-0.5x^3 + 0.5x) \cdot 0.52 \end{cases}$$

$$\begin{cases} x = 0.888(\sqrt{x}(1-x)/e^{13.28x}) \\ x = 0.57(\sqrt{x}(1-x)/e^{5x}) \\ x = 0.1[sin(\pi x^{0.431})]^3 \\ x = 0.1[sin(\pi x^{0.757})]^3 \\ x = 0.1[sin(\pi x^{1.357})]^3 \\ x = 0.1[sin(\pi x^{3.106})]^3 \end{cases}$$

$$x = 0.87((\theta + sin(\theta))/\pi - sin(\theta/2)^2) \quad where \quad \theta = 2sin\sqrt{x}$$

The modification of the airfoil geometry is performed using a code developed by the Italian Aerospace Research Center (CIRA) called WG2AER, which parameterizes the airfoil as a linear combination of the baseline geometry and the modification functions.

To conduct this study, a well-defined methodological approach was followed. Initially, the necessary database was built to perform PME and PI-PME. To create it, a code was written that, by uniformly sampling each variable within its domain, was able to generate an S number of modified configurations of the original airfoil. Subsequently, these configurations were analyzed using the XFOIL software [6] to obtain crucial information

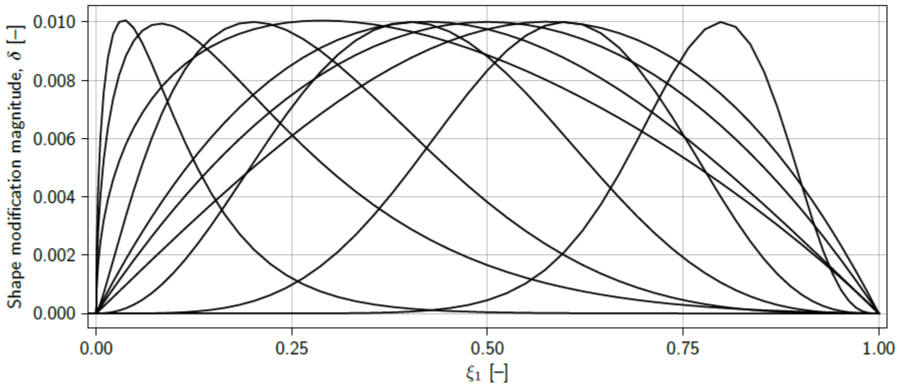

Figure 1: Shape modification functions used [5].

about the physics involved. Once the quality of the obtained dataset was verified, a MATLAB code was developed to perform PME and its physics-informed version. Finally, to optimize the aerodynamic profiles, a code developed by CNR based on Particle Swarm Optimization (PSO) [7] was used. During the optimization process, each configuration identified by individual particles was carefully studied and analyzed through simulations performed using the SU2 software [8]. The number of particles used for every optimization is equal to $4 *$ (**Number of variables**), and the convergence criterion is set as the maximum number of function evaluations, equal to 1500.

3.1 Analysis of the Dimensionality Reduction Processes

Before optimizing the profile by leveraging reduced-dimensionality design spaces, the effect of individual physical information (in this case, C_P and C_D) on the new design space was analyzed.

The first noticeable aspect in the graphs in Figure 2 is that the exclusive use of C_D does not alter the solution of the eigenvalue problem in a way that provides a different number of reduced variables compared to the baseline case (PME).

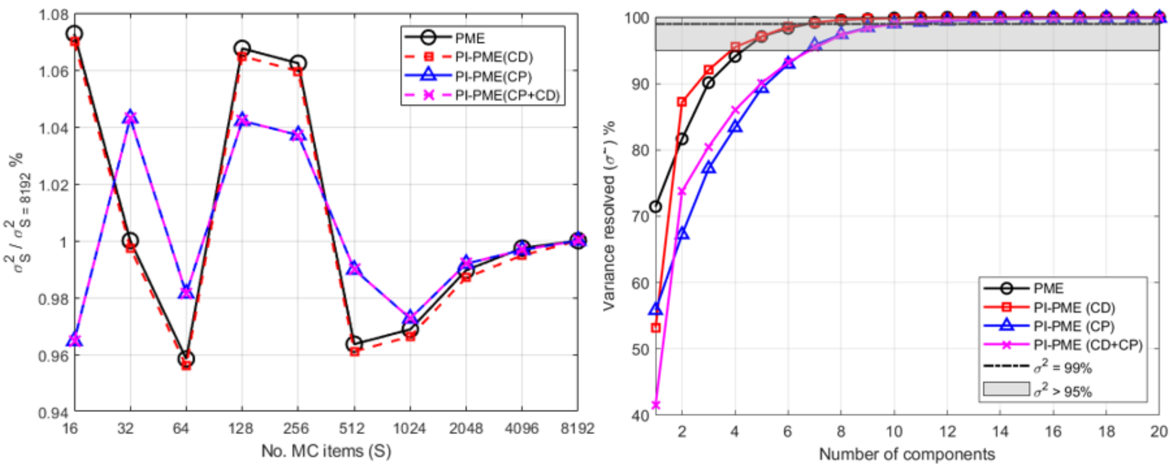

Figure 2: Comparison of the four studies conducted for variance convergence as a function of the MC elements used (left) and the variance retained by the eigenvectors (right).

A similar behavior can also be observed when comparing PI-PME (C_P) and PI-PME

$(C_P + C_D)$. Although the first eigenvector in the C_D-informed studies is capable of representing a significantly lower percentage of the total variance of the problem compared to its uninformed counterparts, from the second eigenvector onward, it becomes evident that the solution increasingly aligns with the "original" one. Regarding the determination of the number of Monte Carlo samples (S) to be used, referring again to Figure 2, it is clear that meaningful solutions can be obtained using 1024 configurations for any study considered. In this work, the maximum available number of MC samples was chosen since the time required to perform the DR varied imperceptibly as their number increased.

Based on the results obtained, three different optimizations were performed. In addition to the study in the original domain and in the domain obtained using PME, the final optimization was carried out in the reduced domain obtained by leveraging both available physical information. This choice proved to be the most appropriate, not only because the presence of C_D alone did not affect the number of reduced variables but also because it allowed for the analysis of the problem with the maximum amount of available information.

4 ANALYSIS OF OPTIMIZATION RESULTS

Once the optimization domains were selected, the aerodynamic performance of the original unmodified profile was initially evaluated. Subsequently, the three optimizations were performed in the various domains and finally a comparative analysis of the results was performed.

To allow a comparison with the results obtained through the optimization processes, an initial aerodynamic analysis of the unmodified RAE-2822 airfoil was conducted. The simulation performed using SU2 provided various useful outputs, summarized in Table 1, which allowed for the definition of the minimum required constraint C_L for the optimized airfoil, ensuring better efficiency compared to the initial configuration.

Table 1: Data related to the unoptimized RAE-2822 airfoil.

RAE 2822 - Original	
C_L	0.211748
C_D	0.0093607
C_M	-0.06556
E	22.62071
t_{85}/c	0.031766

4.1 Optimization in the original domain

The analysis in the full domain initially involved uniformly sampling the 20 variables between -2 and 2. However, this often resulted in invalid geometries, such as a particularly sharp Leading Edge (LE) or intersecting top and bottom surfaces. These profiles, which could not be evaluated using XFOIL, significantly slowed down the process of building the required database for the subsequent PI-PME. Thus, the range of the 20 variables was adjusted, as shown in Table 2, to minimize the occurrence of these invalid geometries. This adjustment represented the best compromise found to ensure the minimum number of intersecting geometries and an adequate variety of representable profiles during the dataset creation.

Table 2: Definition of design variables in the full domain.

Variable	Upper bound	Lower bound
x_1	1.800	-1.500
x_2	1.800	-1.500
x_3	1.800	-1.500
x_4	1.800	-1.000
x_5	1.800	-1.000
x_6	1.800	-1.000
x_7	1.000	-0.500
x_8	1.000	-0.500
x_9	1.000	-0.500
x_{10}	1.800	-1.000
x_{11}	1.800	-1.500
x_{12}	1.800	-1.000
x_{13}	1.800	-1.500
x_{14}	1.800	-1.500
x_{15}	1.800	-1.500
x_{16}	1.800	-1.500
x_{17}	1.800	-1.000
x_{18}	1.800	-1.000
x_{19}	1.800	-1.500
x_{20}	1.800	-1.000

Once the variable domain, the number of particles used (80), and the maximum number of function evaluations were defined, the optimization process was started. After about 250 hours, the process reached the stopping condition.

The aerodynamic characteristics of the new airfoil (shown in Figure 3a) are summarized in Table 3. Notably, the drag coefficient is $C_D = 9.086 \cdot 10^{-3}$ and the lift coefficient is $C_L = 2.273 \cdot 10^{-1}$, representing an improvement of approximately 2.93% in C_D and 10.63% in efficiency.

Among the various results obtained from the simulations, an interesting comparison with the original airfoil is the distribution of the pressure coefficient (C_P) along the profile, as shown in Figure 3b.

Table 3: Data related to the RAE-2822 airfoil optimized using the original domain.

RAE 2822 - Full Domain	
C_L	0.227392
C_D	0.00908617
C_M	-0.0680325
E	25.02621
t_{85}/c	0.022638

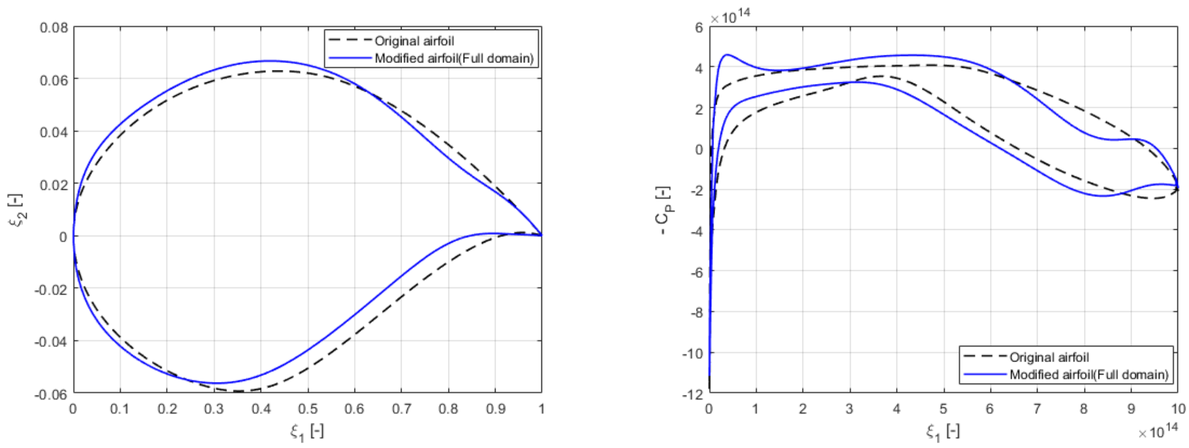

(a) Comparison of geometries (full domain) (b) Comparison of C_P (full domain)

Figure 3: Results of the optimization in the full domain.

4.2 Optimization in the reduced domain via PME

The PME was trained using a set of $S = 8192$ MC elements. The initial geometry is described by $L = 364$ nodes, providing a data matrix P of dimensions $[748 \times 8192]$. The coefficients ρ_i that compose the matrices $\mathbf{W}$ were uniformly set to $1/\sigma_g^2$, where σ_g^2 denotes the variance of the shape modification vector.

The number of reduced design variables required to retain at least 95% of the original geometric variance is $N = 5$, thus achieving a 75% reduction in dimensionality, whereas $N = 8$ reduced design variables are needed to preserve at least 99% of the geometric variance.

The optimization was configured by setting the number of particles to 20 to form the PSO swarm. The five new variables used are those obtained through PME contained in Table 4

Table 4: Definition of design variables in the reduced domain using PME.

Variable	Upper Bound	Lower Bound
x_1	3.49762	-3.73830
x_2	1.04279	-1.90947
x_3	1.24933	-1.16896
x_4	0.81568	-0.67094
x_5	0.69107	-0.92062

The new optimized geometry, shown in Figure 4a compared to the original configuration, is characterized by a $C_D = 9.058 \cdot 10^{-3}$ and a $C_L = 2.444 \cdot 10^{-1}$, indicating an improvement in C_D of 3.22% and a total efficiency increase of 19.3%. Other relevant parameters are listed in Table 5, while Figure 4b illustrates the distribution of C_P for each point of the profile compared to the original.

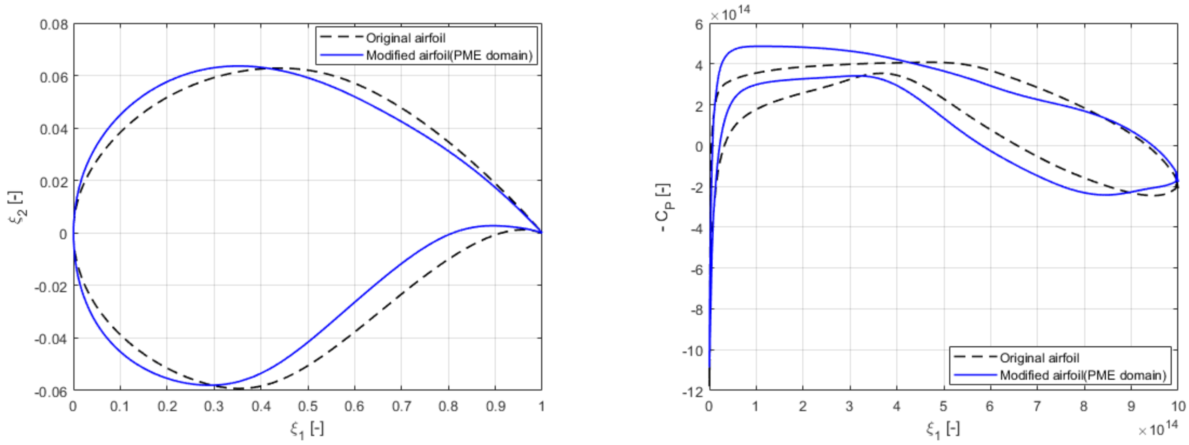

(a) Comparison of geometries (PME) (b) Comparison of C_P (PME)

Figure 4: Results of the optimization in the domain reduced via PME.

Table 5: Data related to the RAE-2822 airfoil optimized using PME.

RAE 2822 - PME	
C_L	0.24447
C_D	0.00905892
C_M	-0.0739756
E	26.9867
t_{85}/c	0.022577

4.3 Optimization in the reduced domain via PI-PME

The physical information used was extracted from simulations performed using XFOIL. As in the previous case, the PI-PME was trained using S = 8192 MC elements.

Since C_P is defined for each point of the profile and C_D is unique for each configuration, a matrix B of size [1113 x 8192] is obtained. Similarly to what was done with the matrix $\mathbf{W}$, the matrices $\mathbf{W_f}$ and $\mathbf{W}_c$, associated to C_P and C_D respectively, were also weighted with respect to their variances $\sigma^2_{C_P}$ and $\sigma^2_{C_D}$.

Table 6: Definition of design variables in the reduced space using PI-PME.

Variable	Upper bound	Lower bound
x_1	4.94878	-4.06976
x_2	2.60132	-5.56868
x_3	1.85518	-1.59955
x_4	1.883284	-1.9525
x_5	1.71518	-1.43197
x_6	1.18975	-1.47970
x_7	1.26270	-1.19997

The number of reduced design variables required to retain 95% of the original variance is N = 7, while to retain 99% of the original variance, N = 10 is required.

In this case, the optimization was initiated by setting the number of particles to 28, and the new seven variables were defined as in Table 6.

The newly obtained geometry, visible in Figure 5a, defined by the parameters present in Table 7, is characterized by a $C_D = 8.941 \cdot 10^{-3}$ and a $C_L = 2.308 \cdot 10^{-1}$, which indicate an improvement of C_D by 4.48% and an overall increase in efficiency of 14. 10%. As with previous optimizations, Figure 5b illustrates the trend of C_P.

Table 7: Data related to the RAE-2822 airfoil optimized using PI-PME.

RAE 2822 - PI-PME	
C_L	0.23076
C_D	0.0089409
C_M	-0.054516
E	25.8093
t_{85}/c	0.022577

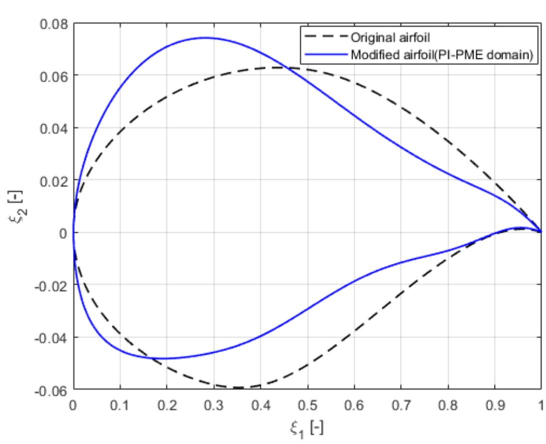

(a) Comparison of geometries (PI-PME)

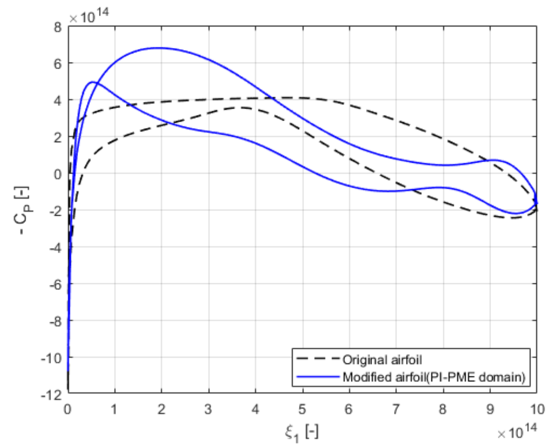

(b) Comparison of C_P (PI-PME)

Figure 5: Results of the optimization in the domain reduced via PI-PME.

4.4 Comparative Analysis of the Results

Figure 6, which shows the convergence of the optimization algorithms, clearly highlights how the implementation of PI-PME, compared to the study in the full domain (Full) and reduced through PME alone, has proven to be particularly advantageous.

Around the 180th iteration, a better configuration was identified compared to those obtained at the end of the other two tests, indicating a greater speed in identifying optimal geometries for the given problem. Furthermore, the configuration found upon reaching the imposed 1500 function evaluations is significantly better than the two previously obtained. A common feature observed across all three optimization cases, shown in Fig. 7, is the thinning of the lower surface. This suggests that the optimization process systematically reduces the thickness of this region and the extent of this modification varies among the different optimization domains.

In the “FULL” case, the geometric modifications appear to be more evenly distributed along the entire profile. The optimized shape closely follows the original contour but

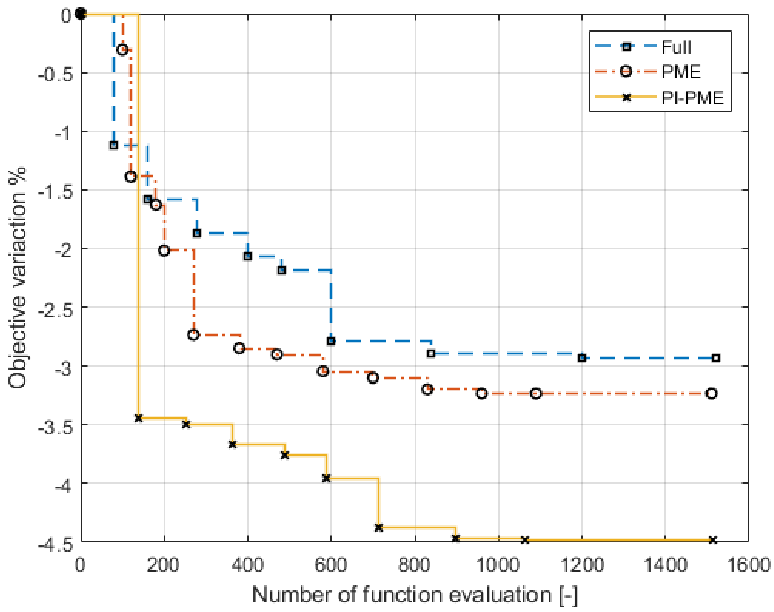

Figure 6: Comparison of convergence towards the optimal configuration identified for the three studies conducted.

with subtle refinements. The "PME" case, on the other hand, exhibits a more localized geometric transformation. While the overall shape remains similar to the original, certain regions show a more pronounced deviation, particularly on the upper surface. The last case displays the most significant alterations in both geometry and pressure distribution. The optimized profile exhibits greater curvature variations, especially along the upper surface, which could imply a more aggressive approach in this optimization domain. The C_P plot further supports this observation, as it presents a noticeable shift in pressure levels, particularly at the leading edge and peak suction regions.

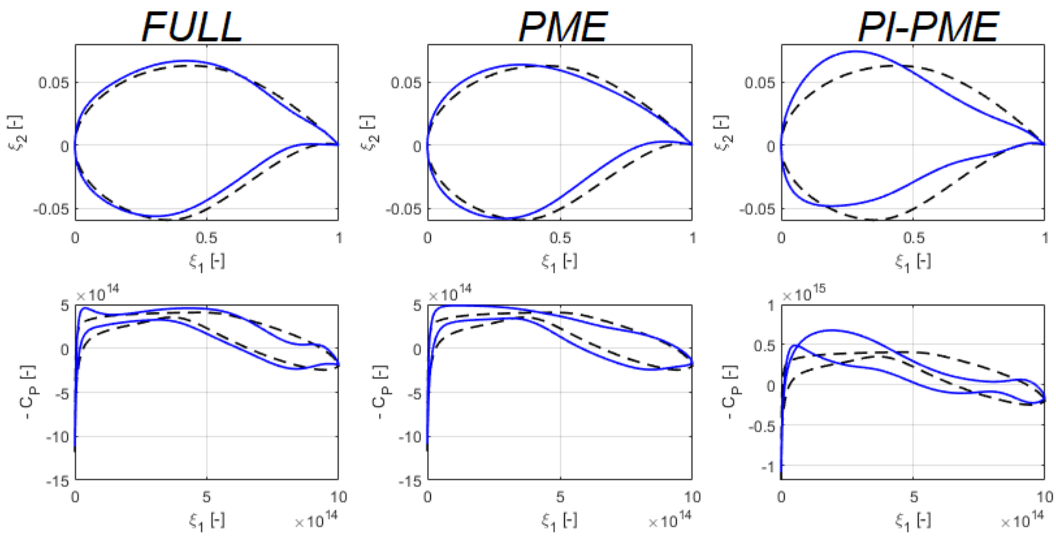

Figure 7: Comparison of the geometry and pressure coefficient for the three studies conducted.

5 CONCLUSIONS

In this work, the problem of shape optimization in engineering has been addressed through the use of DR techniques, taking into account the ”Curse of Dimensionality” and presenting a solution based on PI-PME, an extension of PME. Through critical and experimental analysis, the effectiveness of this technique in tackling high-dimensional optimization problems has been demonstrated.

Not only did PI-PME accelerate the convergence of the optimization algorithm, but the final geometry also proved to be significantly better in terms of the objective function compared to those obtained within the complete and reduced domains using PME.

Regarding future developments, several interesting directions can be considered. First, the application of the PI-PME methodology could be extended to test cases with more challenging boundary conditions to further explore the limitations and potential of this technique. Additionally, conducting a study where the initial dataset is created using XFOIL under boundary conditions in which this software is known to produce accurate results would be of great importance. Once the reduced-dimensional space is built with the obtained data, it would be interesting to perform the airfoil optimization under more complex boundary conditions using specialized software such as SU2. This approach would maximize the strengths of XFOIL for generating a reliable starting point while leveraging SU2 for handling more complex flight conditions.

ACKNOWLEDGEMENTS

This work was carried out under the National Recovery and Resilience Plan (PNRR) scholarship on Methods and Models for Sustainable Engineering at the University of Roma Tre, according to the Ministerial Decree (DM) 117/2023, in collaboration with the Italian Aerospace Research Center (CIRA). It was also conducted within the framework of the NATO-AVT-404 activity on 'Enhanced Design Processes of Military Vehicles through Machine Learning Methods'.

REFERENCES

[1] R. Bellman. Dynamic programming. *Cambridge Studies in Speech Science and Communication. Princeton University Press*, 1957.

[2] A. Serani and M. Diez. A survey on design-space dimensionality reduction methods for shape optimization. *arXiv preprint arXiv:2405.13944*, 2024.

[3] A. Serani and M. Diez. Parametric model embedding. *Computer Methods in Applied Mechanics and Engineering*, 404, 2023.

[4] A. Serani, G. Palma, J. Wackers, D. Quagliarella, S. Gaggero, and M. Diez. Extending parametric model embedding with physical information for design-space dimensionality reduction in shape optimization. *arXiv preprint arXiv:2504.05863*, 2025.

[5] A. Serani, M. Diez, and D. Quagliarella. Aerodynamic shape optimization in transonic conditions through parametric model embedding. *Aerospace Science and Technology*, 155:109611, 2024.

[6] M. Drela. XFOIL: An analysis and design system for low reynolds number airfoils. *Low Reynolds Number Aerodynamics*, 54, 1989.

[7] J. Kennedy and R. Eberhart. Particle swarm optimization. *Proc. Int. Conf. Neural Netw*, 4:1942–1948, 1995.

[8] 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. *AIAAJ*, 54:828–846., 2016.