



Research paper

Physics-informed dimensionality reduction for propeller shape optimization

Stefano Gaggero ^a , Andrea Serani ^b ,*^a University of Genoa, Via Montallegro, 1, Genoa, 16145, Italy^b National Research Council-Institute of Marine Engineering, Via di Vallerano 139, Rome, 00128, Italy

ARTICLE INFO

Keywords:

Marine propellers
 Propeller optimization
 Simulation-based design optimization
 Dimensionality reduction
 Parametric model embedding
 Physics-informed methods
 Computational fluid dynamics

ABSTRACT

The design of marine propellers is challenged by high-dimensional parameter spaces and the need to balance efficiency with cavitation avoidance. Dimensionality reduction techniques offer a cost-effective way to address the curse of dimensionality, but geometry-based approaches such as parametric model embedding (PME) may neglect local features with strong hydrodynamic relevance. This work introduces the application of physics-informed PME (PI-PME) to propeller shape optimization, where physical observables, including pressure distributions and performance indicators, are embedded into the reduced-order space. A multi-objective optimization framework, based on boundary element method analyses and validated with Reynolds-averaged Navier–Stokes simulations, is applied to a cruise-ship propeller. Comparisons between original, PME-reduced, and PI-PME-reduced design spaces demonstrate that PI-PME preserves critical sectional features and significantly improves optimization results. The results highlight the benefits of integrating physical information into dimensionality reduction, enabling reliable, efficient, and physics-aware design optimization of marine propellers.

1. Introduction

The hydrodynamic design of marine propellers has long represented one of the central challenges in naval architecture. As a key component of propulsion systems, propellers influence vessel speed, fuel consumption, emissions, and onboard comfort through vibration and noise. Traditional design methods based on lifting-line, vortex-lattice, and lifting-surface theories (Kerwin and Leopold, 1964) provided valuable insights for decades, but their simplified assumptions limited the capability to simultaneously address efficiency, cavitation, and structural constraints.

Over the last two decades, the advent of simulation-based design optimization (SBDO, Serani et al. 2024b) has transformed the design process. Medium- and high-fidelity (Bertetta et al., 2012; Gaggero, 2018, 2020, 2025) computational fluid dynamics (CFD) analyses, ranging from lifting surface and boundary element methods (BEM) to Reynolds-averaged Navier–Stokes (RANS) simulations, have enabled the systematic exploration of multi-objective trade-offs. These approaches allow designers to account for a broad spectrum of requirements, including energy efficiency, noise and vibration mitigation, cavitation suppression, and compliance with stricter environmental regulations. As a result, optimization frameworks (Serani et al., 2022; Gaggero et al., 2017) are increasingly applied not only in academia but also in industrial propeller design practice. However, SBDO comes at a cost:

the parametrization of complex blade geometries typically leads to high-dimensional design spaces, and the *curse of dimensionality* (Bellman et al., 1957) severely inflates the computational cost of robust exploration and optimization.

To alleviate this issue, dimensionality reduction (DR, Serani and Diez 2025) techniques have become a focus of attention in marine engineering. Classical approaches based on proper orthogonal decomposition (POD) and principal component analysis (PCA) have been applied to simplify parametric spaces while retaining essential shape variability. In marine applications, POD/PCA has improved hydrodynamic performance by optimizing the designs of mono- (Harries and Abt, 2019; Ballarin et al., 2019; D'Agostino et al., 2020; Chang et al., 2023; Zhang et al., 2024; Villa et al., 2020) and multi-hulls (Diez et al., 2015; Chen et al., 2015; Harries and Uharek, 2021), hull vanes (Çelik et al., 2021), propellers (Gaggero et al., 2020), and turbines (Masood et al., 2021). More recently, parametric model embedding (PME, Serani and Diez 2023, Serani et al. 2024a), which extends PCA by using a generalized feature space that includes shape modification and design variables vector together, have gained prominence. PME offers two key advantages: it reduces the number of effective design variables while preserving the coherence of the underlying parametric model, and it provides an analytical back-mapping to the original design variables.

* Corresponding author.

E-mail address: andrea.serani@cnr.it (A. Serani).

This ensures that reduced-order representations can be directly reinterpreted in terms of meaningful design features such as pitch, chord, or camber distributions.

Nonetheless, PME, and geometry-based DR approaches in general, are limited by their reliance on purely geometric variance. Local shape modifications with minor geometric footprint, but strong hydrodynamic impact, may be filtered out by the reduction process. For marine propellers, this limitation is particularly relevant for sectional camber and thickness distributions, which play a decisive role in cavitation inception and avoidance. Neglecting these features reduces the ability of reduced-order models to capture critical trade-offs in performance.

To overcome this limitation, physics-informed PME (PI-PME, Serani et al. 2025) extends PME by embedding physical observables into the dimensionality reduction process. By enriching the data matrix with distributed quantities (e.g., surface pressure fields) and lumped indicators (e.g., thrust coefficient, efficiency, or vortex circulation), PI-PME creates reduced spaces that are not only geometrically coherent but also physically aware. This enhancement allows reduced-order spaces to preserve subtle features with strong influence on cavitation and efficiency, enabling a more balanced and reliable exploration of design trade-offs.

The objective of this work is to assess the effectiveness of PME and PI-PME in the hydrodynamic optimization of a six-bladed cruise ship propeller. A multi-objective optimization framework is established, combining BEM analyses (Gaggero et al., 2010; Gaggero and Villa, 2017) for computational efficiency with RANS simulations (Siemens Digital Industries Software, 2023) for validation. Results from optimizations carried out in the original and reduced spaces are compared, demonstrating that PI-PME significantly improves the ability to mitigate cavitation while preserving or enhancing efficiency. The study highlights how embedding physics into dimensionality reduction fosters more robust and efficient design processes, bridging the gap between advanced mathematical techniques and practical propeller optimization.

The remainder of the paper is organized as follows. Section 2 introduces the test case and presents the reference propeller performance. Section 3 formulates the optimization problem, while Section 4 describes the dimensionality reduction methodologies, with emphasis on PME and PI-PME. Section 5 reports optimization results in the original and reduced design spaces, followed by Section 6, where the performance of selected optimal propellers is assessed through unsteady BEM and RANS validation. Finally, Section 7 concludes the paper by summarizing the findings and outlining directions for future research.

2. Test case and reference performance

The test case selected for this application is a six-blade, right-handed, marine propeller (see Fig. 1) preliminarily designed by Fincontieri for a cruise ship.

The open water performance (BEM and RANS), and the unsteady functioning calculations with cavitation in the nominal wake of the ship (BEM only) are employed to evidence the possible issues in terms of risk of cavitation to be addressed in the design by optimization. Furthermore, compared to experiments, they provide minimal validation (since only open water measured data are available) of the tools employed through the optimization process itself. Since the focus is on the feasibility of a design by optimization via reduced-order models and not on the absolute accuracy of the analyses, a dedicated verification of the numerical models (both BEM and RANS) is not provided. The best practices developed over the years for accurate but computationally efficient analyses for optimization-type applications (Gaggero and Villa, 2017, 2018; Gaggero, 2022) guide the setup of all the proposed numerical analyses. All BEM calculations (steady, equivalent steady, and unsteady through the “key blade” of Kinna and Hsin 1992) are carried out using the code developed at the University of Genoa since




Fig. 1. Geometry of the reference propeller.

early 2000 (Gaggero et al., 2010) with a surface mesh of 1250 hyperboloidal panels per blade (25 radial sections, 50 chordwise panels, then $L = 1326$ grid nodes). Cavitation prediction is based on the sheet cavity model (Fine, 1992) at the leading edge and midchord (both on suction and pressure side), coupled with wake alignment capabilities and the iterative Kutta condition. In the absence of other data, the high-fidelity RANS serves to validate the medium-fidelity BEM on which the entire design process is based. RANS (StarCCM+, Siemens Digital Industries Software 2023), used in current analyses only for open-water calculations (at equivalent open water advance coefficient to compare with the circumferentially averaged, radially varying, inflow used with BEM during the optimization, Section 3), employs a cell-centered implementation of the finite-volume method on polyhedral grids that count, on average, between 2 and 3 million cells per blade passage. Fully turbulent analyses adopt the $k - \omega$ SST turbulence model with a moving reference frame formulation of the equations and periodic boundary conditions to exploit the axial symmetry of the problem. Only when employed to assess the occurrence (or the reduction) of tip vortex cavitation through detailed characterization of the pressure drop in the core of the vortex, grids are refined up to 30 million cells using refinement zones derived from the trajectory of the vortex computed through the wake alignment procedure of the BEM.

The open water performance of the reference propeller, computed with BEM and RANS, is compared with measurements in Fig. 2. The agreement is very good, as expected for this type of analysis, especially at the design advance coefficient ($J_A = V(1 - w)/nD = 0.87$), where the experimental thrust and torque are 0.1866 and 0.3768, respectively. For BEM, the accuracy is higher in predicting the thrust coefficient, which is less affected by viscous effects, only approximately accounted for by potential flow methods. At the design functioning in open water conditions, BEM ($K_T = 0.1855$, $10K_Q = 0.3891$) under-predicts less than 1% the thrust, while torque is about 4% higher than the measured value. RANS ($K_T = 0.1813$, $10K_Q = 0.3804$) slightly underestimates the thrust (-2%), but torque is less than 1% overestimated. This trend is consistent for a relatively wide range of the advance coefficient, since the differences are never greater than 4%, also at very low ($J = 0.5$) and very high ($J = 1.1$) advance coefficient. Combined, the thrust and torque of BEM and RANS determine very similar predicted efficiencies, which are 2% lower than experiments across the design functioning.

For the sake of the optimization process, better detailed in Section 3, accurate predictions of the pressure distributions over the blade are of fundamental importance to correctly assess the risk of cavitation and guide the selection of the optimal geometry. Since the design process is based on BEM calculations, the comparison of pressure distributions ($C_{p,N} = 2(p - p_{ref})/\rho n^2 D^2$) with those from higher-fidelity codes like RANS provides an acceptable validation of these key performance indicators extensively employed for the optimization (Gaggero et al.,




Fig. 2. Open water propeller performance; comparison of BEM and RANS results with experiments.

2017). At the design advance coefficient, the comparison at several radial positions from root to tip is given in Fig. 3. As seen in terms of global performances, the agreement between BEM and RANS is very good, since the values of pressure, as well as the most important features of its distribution, are qualitatively and quantitatively captured also by the medium-fidelity flow solver. In uniform inflow, there is an evident tendency to flow inversion at the leading edge, with suction acting on the face side (rather than on the back) towards the tip of the blade. This characteristic, observed by RANS, is slightly anticipated by BEM, since the suction on the face side at the leading edge is clearly visible from $r/R = 0.7$. This suggests a risk of unwanted pressure side cavitation that is moreover clear from the unsteady calculations, again using BEM, of the propeller under the non-uniform wake of the ship (for simplicity, the nominal wake from dedicated full-scale RANS calculations, given in Fig. 4).

Unsteady pressure distributions over the blades of the reference propeller (Fig. 5) evidence, when the propeller works at the design cavitation index ($\sigma_N = 2(p_{ref} - p_{vap})/\rho n^2 D^2$ equal to 2.25), the occurrence of pressure side cavitation over a large portion of the blade revolution, especially where the combination of axial and tangential flow is less favorable for pressure side phenomena. At the leading edge of the face side, the suction is significantly lower than the cavitation index, and this turns into a cavity bubble extended over a relevant portion of the pressure side when the cavitation model of the BEM solver is turned on, as in Fig. 6. The risk of cavitation is also high on the suction side, in correspondence to the 90 deg angular position of the blade, where the localized reduction of axial velocity (wake of the upstream shaft bracket), combined with the tangential velocity pointing upward, increases the angle of attack of the blade. The risk of suction side cavitation turns into a fully developed midchord cavity bubble when the functioning conditions are slightly altered to mimic, in addition to the variety of working conditions naturally given by the non-homogeneous inflow, the functioning under uncertainties a propeller may be subjected to. In the specific, an advance coefficient J_A equal to 0.84 and a cavitation index of 1.8 have been added to the analysis to have a broader view of the performance and the resilience of the reference propeller against variable inputs. Namely, the reduced cavitation index (at $J_A = 0.87$) corresponds to the condition of maximum ship speed. The higher angle of attack associated with the reduced value of advance coefficient fosters the development of suction side cavitation, which in any case, always appears as midchord

(bubble) cavitation at intermediate radial positions along the blade, and only for the reduced cavitation index. Pressure side cavitation is still visible around the 270 deg position of the blade over a relatively large portion of the leading edge from root to tip. These features suggest the design criteria adopted for the reference propeller, which has been heavily unloaded at the tip to prevent (or retard) the appearance of a cavitating tip vortex. This unloading, especially under the influence of an uneven (tangential) inflow wake, determines the dominance of leading-edge pressure-side cavitation and the appearance of suction-side phenomena at the inner positions on the blade, where an increase in the angle of attack is necessary to balance the tip unloading. The persistence of face-side cavitation also in correspondence of loaded conditions (i.e., $J_A = 0.84$) suggests that the unloading is probably excessive and represents one of the issues, in contrast to the need to maintain as low as possible the risk of tip vortex cavitation, to be addressed in the optimization problem.

3. Propeller design optimization problem

The multi-objective optimization procedure is formulated conventionally, independently of the specific design space and the associated design parameters x . The constraint is imposed on the delivered thrust, fixed to the reference value (K_{Tref}) of the baseline propeller at the design point, consistently calculated with the same numerical setup. Within this framework, efficiency is maximized while cavitation is simultaneously minimized. Nevertheless, as highlighted in Section 2, the treatment of unsteady performance necessitates a number of simplifications to ensure computational tractability of the shape optimization process. In fact, despite the adoption of a BEM solver, conducting cavitating unsteady simulations over several hundred thousand design candidates remains computationally prohibitive.

As extensively shown in Gaggero et al. (2017) and in several other similar activities when using RANS solvers in place of BEM (Gaggero, 2020), the risk of cavitation can be inferred from inexpensive non-cavitating analyses through the simplest inception criterion based on the cavitation index, i.e., when $-C_{pN}$ is greater than σ_N . To enrich the quality of the results and to distinguish between sheet and bubble cavitation simply based on the zone of occurrence, the cavitation inception through the maximum of $-C_{pN}$ is monitored separately at the leading edge and at the midchord, on both the suction and pressure sides (Fig. 7), since the type of cavitation occurring in these zones is very different (and with different acceptability) despite the same cavitation risk. The phenomena of the leading edge are collected in a zone up to 20% of the local chord of the blade (midchord from $x/c = 20\%$ to the trailing edge). The tip region is identified when the radial position exceeds 70% of the blade radius.

The unsteady functioning is simplified as well, employing kinematic equivalent functioning conditions to represent the most relevant loading conditions and the relative cavitating phenomena (through steady BEM analyses only) which the propeller blades are most likely to encounter during a revolution. These kinematic equivalent functioning conditions are expressed in terms of radially varying and circumferentially averaged axial and tangential velocity distributions. The design functioning condition corresponds to the circumferentially averaged, radial position per radial position, axial velocity of the nominal wake, which, once averaged over the entire propeller disk, leads to the design advance coefficient $J_A = 0.87$ in the uniform inflow mentioned in Section 2. Under this radially non-uniform axial inflow, the predicted thrust coefficient by the BEM is 0.1860, very close to the open water, equivalent uniform inflow case (0.1855 at the same advance, Section 2) given the relatively constant distribution of the circumferentially averaged inflow along the blade radius, and only slightly higher than the averaged thrust (0.1781) predicted by the unsteady BEM, naturally affected by the spatial non-homogeneous inflow distribution and the associated unsteady effects.

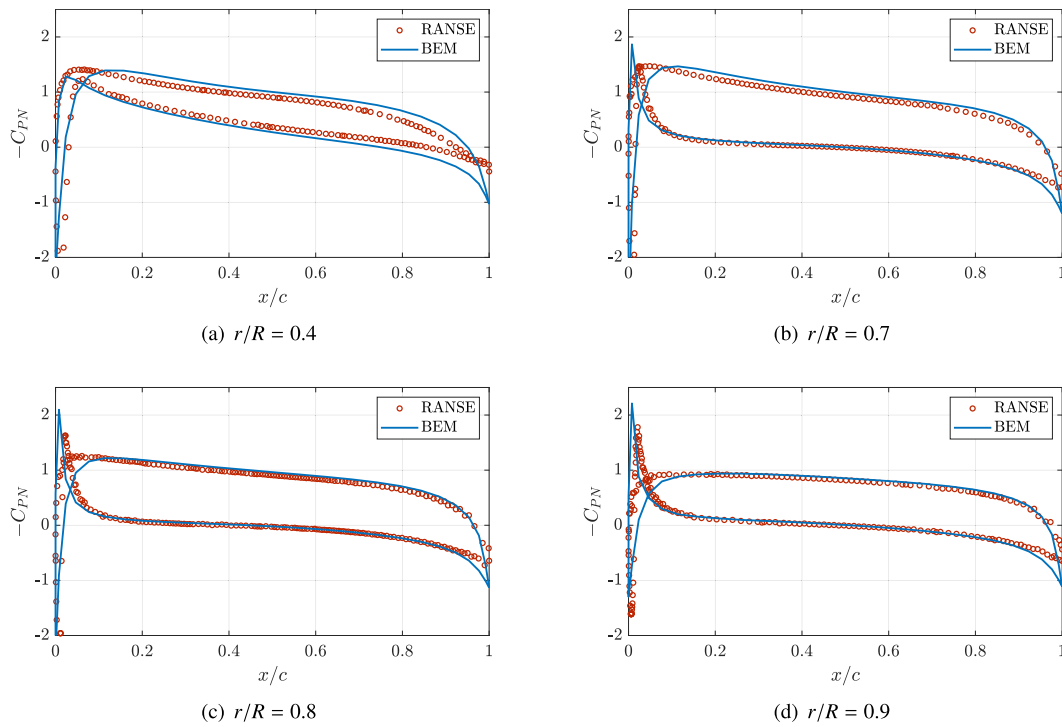


Fig. 3. Sectional pressure distributions computed with BEM and RANS for the reference propeller at the design advance coefficient.

Table 1

Inception of cavitation (maximum values of suction on representative zones) for the reference propeller at the selected functioning conditions and thresholds for filtering the space of solutions of the designs by optimization.

	Design	Opt. Thresh.	Loaded	Opt. Thresh.	Unloaded	Opt. Thresh.
$\max(-C_{PN,back,LE}^{tip})$	1.394	1.3	1.665	1.5	1.145	1.1
$\max(-C_{PN,back,LE}^{root})$	1.568	1.5	1.887	1.8	1.287	1.2
$\max(-C_{PN,back,mid}^{tip})$	1.170	1.2	1.322	1.4	1.026	1.1
$\max(-C_{PN,back,mid}^{root})$	1.246	1.2	1.389	1.4	1.103	1.1
$\max(-C_{PN,face,LE}^{tip})$	2.042	0.8	1.418	0.4	2.975	1.7
$\max(-C_{PN,face,LE}^{root})$	1.798	1.2	1.307	1.0	2.385	1.6




Fig. 4. Nominal wake, seen from aft, for unsteady analyses.

Heavy- and lightly-loaded conditions with the highest risk of suction-side and pressure-side cavitation, respectively, during a revolution, correspond to the distributions of velocity to the propeller observed across the 90 and 270 deg blade positions. While the average thrust of the propeller, as well as the efficiency, is monitored at the design functioning only, data on the risk of cavitation are extracted also for loaded and unloaded conditions to provide a more balanced design in the non-uniform wake, thereby simplifying the formal robust optimization approach (Gaggero, 2025), in principle needed to address the propeller design in a non-uniform inflow without directly employing unsteady calculations. This is even more important given the observed behavior of the reference propeller, which is dangerously unbalanced towards pressure side cavitation.

Under these simplified conditions, from which the key performance indicators of the design are derived, cavitation inception is minimized at the suction-side leading edge, independently of the operating condition. At the mid-chord and on the pressure side, a constraint on the maximum cavitation risk is imposed, defined with respect to the design cavitation index and calibrated against data obtained for the reference propeller, evaluated under identical inflow conditions and using the same numerical setup. The baseline performance of the reference propeller, based on the kinematic equivalent functioning conditions and adopted for the optimization, are listed in Table 1. At the design point, the propeller delivers a thrust coefficient of 0.1860 with an efficiency

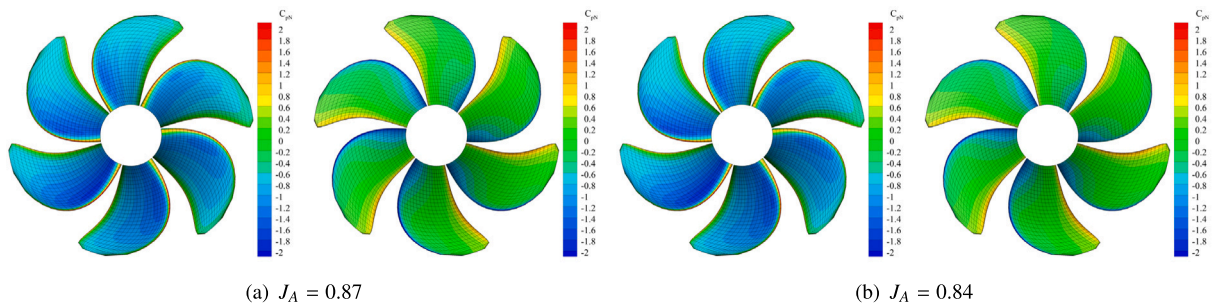


Fig. 5. Unsteady pressure distributions (C_{pN}) on the reference propeller under non-uniform nominal wake at different advance coefficients.




Fig. 6. Predicted unsteady cavity extension on suction and pressure side of the reference propeller at different functioning conditions including maximum ship speed (c) and design functioning (d).

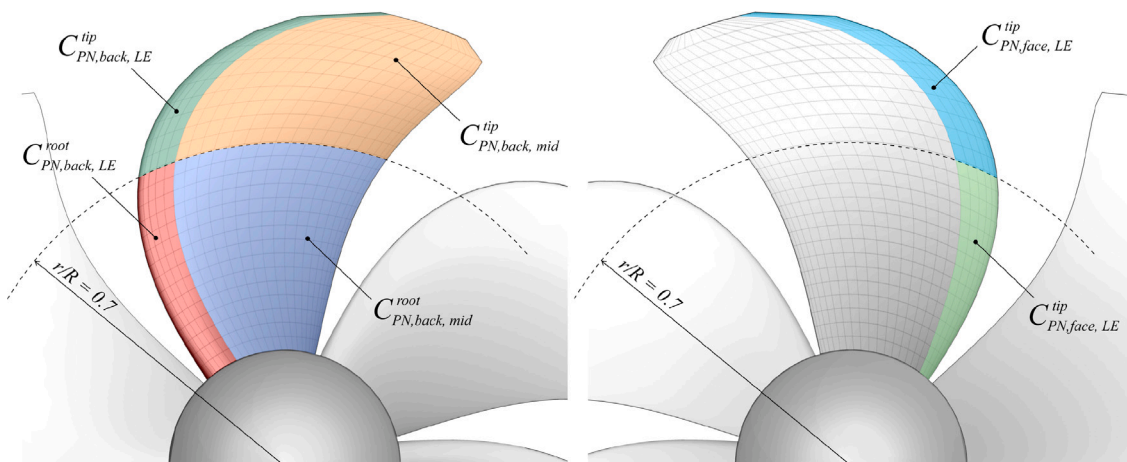


Fig. 7. Blade zones (suction and pressure side) subdivision for cavitation inception monitoring.

of 0.6558. Cavitation (i.e., the maximum values of suction) starts at the leading edge of the pressure side (tip region) while suction side cavitation, coherently with the fully unsteady analyses of Section 2, appears at the root and only for the loaded function condition.

Since the tip unloading observed in the reference propeller reflects the necessity of mitigating the risk of tip-vortex cavitation, any modification of the radial loading distribution intended to delay pressure-side cavitation must account for this trade-off, as pressure-side and tip-

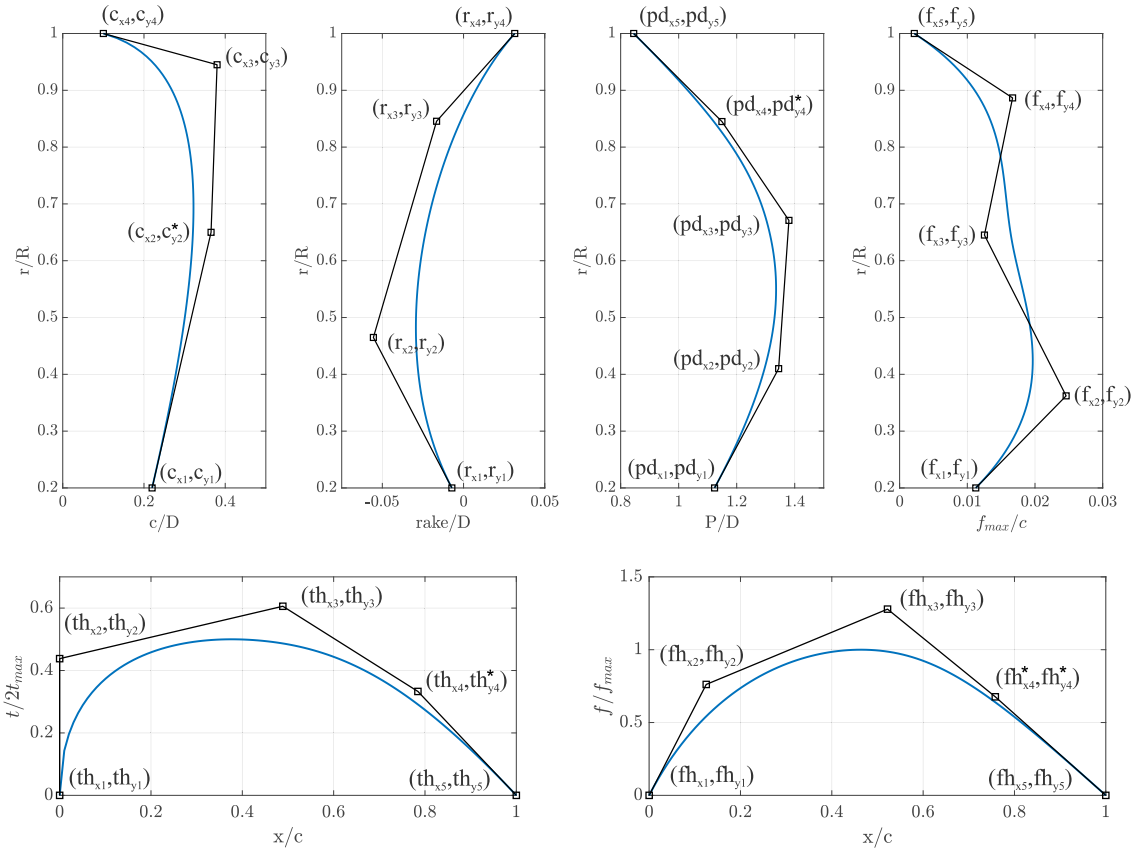


Fig. 8. Parametric description of the propeller blade through radial and sectional B-Spline parametric curves.

vortex cavitation represent conflicting design objectives. In the absence of a dedicated predictive model for tip-vortex cavitation such as those proposed in Lee and Kinna (2004), Bosschers (2018, 2019) and Lidtke et al. (2022), the strength of the tip vortex (Γ_{tip}^{Design}), evaluated from the blade circulation computed by the BEM, is minimized and used as a qualitative indicator of the associated cavitation risk. This formulation ultimately yields the eight-objective optimization problem defined in Eq. (1):

$$\begin{cases} \text{maximize } \eta_o(\mathbf{x}) \\ \text{minimize } -C_{PN_i}(\mathbf{x}) \quad \forall i = 1 : 6 \\ \text{minimize } \Gamma_{des}(\mathbf{x}) \\ \text{subject to } -C_{PN_j}(\mathbf{x}) \leq -C_{PNmax} \quad \forall j = 1 : 12 \\ \text{subject to } K_{Tref} - \delta \leq K_{Tdes}(\mathbf{x}) \leq K_{Tref} + \delta \end{cases} \quad (1)$$

where i is the index of the six combined blade patches and functioning conditions (suction side, leading edge at root and blade tip, design and equivalent 90 and 270 deg blade positions) that require minimization of the risk of cavitation. The index j , instead, identifies the 12 combinations of patches and operating conditions (mid-chord, suction side at root and tip of blade, leading edge at root and tip of blade of pressure side) where the risk of cavitation is restricted to be lower than given thresholds deduced from the performance of the reference propeller under the same operating conditions (Table 1). To accelerate convergence, the constraint on the thrust delivered in the design function is met within an assigned tolerance $\delta = \pm 1.5\%$. In addition to these objectives and constraints, the cavitation area, i.e., the portion of the blade where cavitation inception is predicted, is also monitored as an additional measure of the effectiveness of low pressure in developing large cavitation bubbles.

The parametric description \mathbf{x} of the geometry, through which the optimization algorithm is applied, adopts the conventional B-Spline representation of non-dimensional radial and chordwise distributions

of pitch, chord, camber, and other geometrical quantities, which is a formulation extensively employed in hydrodynamic shape optimization studies (Bertetta et al., 2012; Gaggero et al., 2017). This approach provides a compact representation of the blade geometry through the coordinates of the spline control points, while offering sufficient degrees of freedom to expand the design space without compromising control over the resulting blade shape. Additional features, such as the radial distribution of the maximum sectional thickness, can be constrained independently (for example, according to structural strength requirements), a flexibility not available in alternative methods such as free-form deformation or direct application of B-Spline surfaces to the three-dimensional blade geometry. Furthermore, the parametric model readily accommodates the introduction of nonlinear constraints, ensuring that geometries lying outside the applicability range of medium-fidelity BEM solvers are excluded a priori. This is achieved by hard-coding appropriate bounds and constraints on parameter variability, thereby preventing the occurrence of spurious optima induced by the intrinsic limitations of the adopted medium-fidelity physical model for the hydrodynamic predictions of the optimization process.

For this particular problem, the blade geometry is parametrized through the radial distributions of the 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. 8. The shape of sectional hydrofoils is included as well in the parametric model by describing, again by using appropriate control polygons, 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. As for the original propeller, the parametric description of the blade geometry adopts an identical non-dimensional sectional hydrofoil shape everywhere along the blade span, finally scaled, radial section per radial section, with the corresponding maximum values provided by the relative radial distributions.

Table 2

Range of variation of the design parameters, where * identifies parameters constrained by the values of previous and subsequent control points to avoid inflections or cusps.

Param.	Min	Max	Param.	Min	Max
c_{x2}	0.55	0.8	r_{x2}	0.4	0.6
c_{x3}	0.88	0.97	r_{x3}	0.7	0.9
c_{y1}	0.18	0.24	r_{y1}	-0.025	0
c_{y2}^*	0.27	0.51	r_{y2}	-0.07	-0.04
c_{y3}	0.37	0.45	r_{y3}	-0.03	-0.01
pd_{x2}	0.35	0.6	r_{y4}	0.02	0.06
pd_{x3}	0.65	0.75	fh_{x2}	0.03	0.15
pd_{x4}	0.85	0.95	fh_{x3}	0.4	0.6
pd_{y1}	0.9	1.2	fh_{x4}^*	0.72	0.88
pd_{y2}	1	1.45	fh_{y2}	0.7	0.9
pd_{y3}	1	1.45	fh_{y3}	1	1.35
pd_{y4}^*	0.74	1.31	fh_{y4}^*	0.2	0.87
pd_{y5}	0.7	0.9	th_{x3}	0.3	0.5
f_{x2}	0.3	0.5	th_{x4}	0.7	0.9
f_{x3}	0.55	0.8	th_{y2}	0.35	0.55
f_{x4}	0.85	0.98	th_{y3}	0.5	0.8
f_{y1}	0.0	0.04	th_{y4}^*	0.13	0.61
f_{y2}	0.01	0.05			
f_{y3}	0.01	0.05			
f_{y4}	0.01	0.05			
f_{y5}	0	0.008			

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 BEM solver), control points, sequentially numbered from root to tip (or from leading-to trailing-edge), are free to move within assigned ranges (see Table 2). 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 entirely the blade shape except the skew that is maintained unchanged and identical to that of the reference propeller.

This parametric description, which defines the original design space and natively generates all geometries within the optimization process, is likewise employed for the construction of blade geometries in the reduced spaces. As discussed in Section 4, one of the innovative aspects of the PME approaches proposed in this work, unlike standard PCA, is the capability of mapping back any geometry identified in the reduced space to the corresponding set of parameters in the original design space. Consequently, the parametric description module of the optimization framework remains common to all design spaces, since the reduced-order model reconstructs, from the selected reduced parameters, not only the three-dimensional coordinates of the blade geometry (as in Gaggero et al. 2020, Khan et al. 2023) but also the underlying parametrization in the original space. This eliminates the need for post-processing when comparing geometrical features across spaces and enables more robust surface grid generation for hydrodynamic analyses, as the reconstruction is consistently based on the native B-Spline module of the BEM solver.

4. Design-space dimensionality reduction methodologies

The design of propeller blades often involves a large number of parameters, which makes direct optimization computationally demanding. Dimensionality reduction provides a way to identify a compact set of latent variables that retain the most relevant geometric and physical variability of the design space. In the next subsections, a brief introduction to PME and PI-PME Embedding, which are the two approaches adopted and compared in the current work, is provided.

4.1. Parametric model embedding

Let $\mathbf{u} \in \mathbb{R}^M$ denote the vector of design parameters and $\mathbf{d} \in \mathbb{R}^{nL}$ the discretized shape modifications over L surface points. For a set of S samples, PME builds the augmented data matrix

$$\mathbf{P} = \begin{bmatrix} \mathbf{D} \\ \mathbf{U} \end{bmatrix}, \quad (2)$$

where $\mathbf{D} \in \mathbb{R}^{nL \times S}$ contains geometric deviations from the reference geometry and $\mathbf{U} \in \mathbb{R}^{M \times S}$ contains the corresponding design variables.

A principal component analysis is then performed by solving

$$\mathbf{A} \mathbf{W} \mathbf{Z} = \mathbf{Z} \mathbf{A}, \quad \mathbf{A} = \frac{1}{S} \mathbf{P} \mathbf{P}^T, \quad (3)$$

with $\mathbf{W} = \text{diag}(\mathbf{W}_d, 0)$ a weighting matrix. The eigenvalues $\{\lambda_i\}$ of \mathbf{A} quantify the variance associated with each principal direction, and the total variance of the dataset is given by

$$\sigma^2 = \sum_i \lambda_i. \quad (4)$$

Given a prescribed variance retention level $\eta \in (0, 1)$, the reduced dimension N is selected as the smallest integer such that

$$\frac{\sum_{i=1}^N \lambda_i}{\sum_i \lambda_i} \geq \eta, \quad (5)$$

where the eigenvalues are assumed to be ordered in decreasing order. The corresponding first N eigenvectors define the reduced subspace and the associated latent variables $\mathbf{x} \in \mathbb{R}^N$. Each eigenvector \mathbf{z}_k can be partitioned as

$$\mathbf{z}_k = \begin{bmatrix} \mathbf{q}_k \\ \mathbf{v}_k \end{bmatrix}, \quad (6)$$

where \mathbf{q}_k corresponds to the geometric part and \mathbf{v}_k to the design-variable part. The analytical back-mapping to the original parameters then reads

$$\mathbf{u} \approx \langle \mathbf{u} \rangle + \sum_{k=1}^N x_k \mathbf{v}_k. \quad (7)$$

In short, PME reduces the dimensionality of the problem while preserving a direct link to the original parametric description.

4.2. Physics-informed PME

Standard PME focuses on directions of maximum geometric variance. As a result, local features with small geometric footprint but strong hydrodynamic relevance (e.g., sectional camber or thickness) may be neglected.

To address this limitation, PI-PME augments the data matrix with physical observables:

$$\mathbf{P}_I = \begin{bmatrix} \mathbf{D} \\ \mathbf{U} \\ \mathbf{F} \\ \mathbf{C} \end{bmatrix}, \quad (8)$$

where \mathbf{F} collects distributed fields (e.g., surface pressure distributions), and \mathbf{C} contains lumped performance indicators (e.g., thrust or efficiency).

A block-diagonal weighting matrix is introduced to normalize each block, ensuring that geometry and physics contribute on comparable scales:

$$\mathbf{W}_I = \text{diag}(\mathbf{W}_d, 0, \mathbf{W}_f, \mathbf{W}_c). \quad (9)$$

The dimensionality reduction is then obtained by solving

$$\mathbf{A}_I \mathbf{W}_I \mathbf{Z}_I = \mathbf{Z}_I \mathbf{A}_I, \quad \mathbf{A}_I = \frac{1}{S} \mathbf{P}_I \mathbf{P}_I^T. \quad (10)$$




Fig. 9. Variance resolved as a function of the reduced design variables.

The resulting latent variables capture both geometric deformations and their physical implications, yielding a reduced-order space that is physics-aware.

In summary, PME provides a geometry-based reduction with analytical back-mapping, while PI-PME enhances the method by incorporating physical observables, thus retaining design features that are crucial for accurate propeller performance prediction.

Further details on the mathematical formulation and validation of PME and PI-PME can be found in Serani and Diez (2023), Serani et al. (2025).

5. Dimensionality reduction results

The effectiveness of the proposed design-space dimensionality reduction methods is assessed by comparing the performance of optimal geometries obtained from the standard optimization process in the full (original) design space with those derived from the various reduced spaces (differing in dimensionality and reduction methodology) considered in this study. This approach enables evaluating the influence of different levels of retained variance, as well as the impact of the physical quantities incorporated into the reduction process (PI-PME models), directly in terms of the quality of the resulting optimal geometries rather than relying solely on the, at times less informative, comparison of resolved geometrical and physical variance associated with different numbers of reduced parameters (or “modes”). Both PME and PI-PME strategies are analyzed and the results in terms of variance retained are shown in Fig. 9. For PME, only the configuration with 13 reduced variables is retained, as it accounts for more than 98.5% of the geometric variance; cases with fewer parameters proved incapable of generating feasible geometries within the constraints of the given design problem. In contrast, two PI-PME configurations are considered, employing 12 and 18 reduced parameters, respectively. The 12-parameter case corresponds to retaining 95% of the combined geometry–physics variance, a commonly accepted threshold for adequate reduced-order representations. The 18-parameter case matches the variance (98.5%) of the selected PME configuration, thus enabling a consistent and fair comparison between the two reduction approaches.

Given the complexity of the problem, a multi-objective genetic algorithm (MOGA-II, Esteco 2022) was used to perform the optimization. For each design space, the algorithm started from an initial population of 2500 configurations generated via Sobol sampling and evolved over 50 generations, resulting in a total of 125,000 evaluated geometries per design space. To ensure the statistical reproducibility of the process, multiple independent optimization runs were conducted (five for the PME and ten for the PI-PME). No significant variations were observed in the resulting Pareto fronts, indicating adequate repeatability of the genetic algorithm.

The outcomes of the optimization in the various design spaces, along with a comparison of the corresponding optimal geometries in terms of the key performance indicators of the design problem, are summarized in Sections 5.1 to 5.4. Detailed unsteady BEM analyses and equivalent open-water RANS simulations, aimed at providing a final assessment

of both the optimization strategy (including the necessary design simplifications) and the different design space reduction techniques, are presented in Section 6.

5.1. Original design space

The solution of the optimization problem in the original design space is not different from many similar cases proposed in the literature: all the $M = 38$ coordinates of B-Spline control points of the relevant distributions describing the blade geometry are used as the design parameters of the problem in a *try-and-error* procedure which identifies as feasible the geometries satisfying the required constraints, and advances those providing maximum efficiency and minimum risk of cavitation. A qualitative measure of the geometrical variability embedded in this original design space is given by the envelopes of all the possible radial and chordwise distributions of the blade geometrical features ensured by the native B-Spline parametrization of the shape. With the exception of the rake, which is relevant from the pure geometrical point of view but has limited influence, especially when the simpler kinematically equivalent steady performance that defines the key performance indicators of the design is employed, a view of the shapes allowed in the original design space is given in Fig. 10. Qualitatively, by means of statistical indicators associated with the geometries populating the corresponding Design of Experiments, such as mean distributions, 95% confidence intervals, inter-quartile ranges, and the absolute minimum and maximum values attained by the radial and sectional shapes, it also provides a reference framework for assessing the quality of the dimensionality reduction achieved with the various choices of reduced parameters and models.

Except that at tip, which is fixed to comply with the numerical limitations of the BEM in dealing with very warped panels, control points of the chord permit quite large variations of the expanded area ratio to include from very efficient geometries (potentially with a higher risk of anticipated cavitation inception) to very conservative geometries in terms of chordwise distribution of the load. The non-linear constraint on c_y^2 effectively avoids any inflection point in the chord distribution. In contrast, inflections are allowed for the pitch (only at root) and for the camber (everywhere). The aim was to force, directly by the allowed variability of the design parameters, a monotonic and smooth decrease of the angle of attack at tip to facilitate the avoidance of the tip vortex cavitation by minimizing the vortex strength. This was achieved by the non-linear constraints on pd_y^4 , which, for the nature of the B-Spline curves, has only a local influence on the shape of the pitch distribution. The radial distribution of the camber, which works together with the angle of attack to ensure *shock-free* functioning of the sectional hydrofoil, was instead not constrained in shape. Only the range of variations of the parameters was prescribed to avoid geometries usually out of the reliability of the BEM solver. The range at tip was forced to ensure a decrease of the maximum sectional camber starting from about $r/R = 0.9$. Similar criteria were adopted for the sectional distributions of the thickness and the camber line, mainly to ensure the maximum possible variance without inflection points, not observed for the original geometry. These choices, mostly, turn into modifications of the chordwise position of the maximum of the distributions, a change in the slope at the trailing edge, and a modification of the curvature of the distributions at the leading edge of the hydrofoil.

The results of the optimization in the original design space are given in Fig. 11. Due to the multi-objective nature of the design problem, for visualization purposes, the performances of configurations analyzed by the optimization algorithm are projected on significant planes (couples of key performance indicators) of the space of solutions, where Pareto fronts of contrasting objectives are clearly identifiable. For the sake of compactness of the presentation of the data, the risk of cavitation (maximum of $-C_{pN}$) at the leading edge ($x/c < 0.2$) as well as at midchord ($x/c \geq 0.2$) on both pressure and suction side includes both

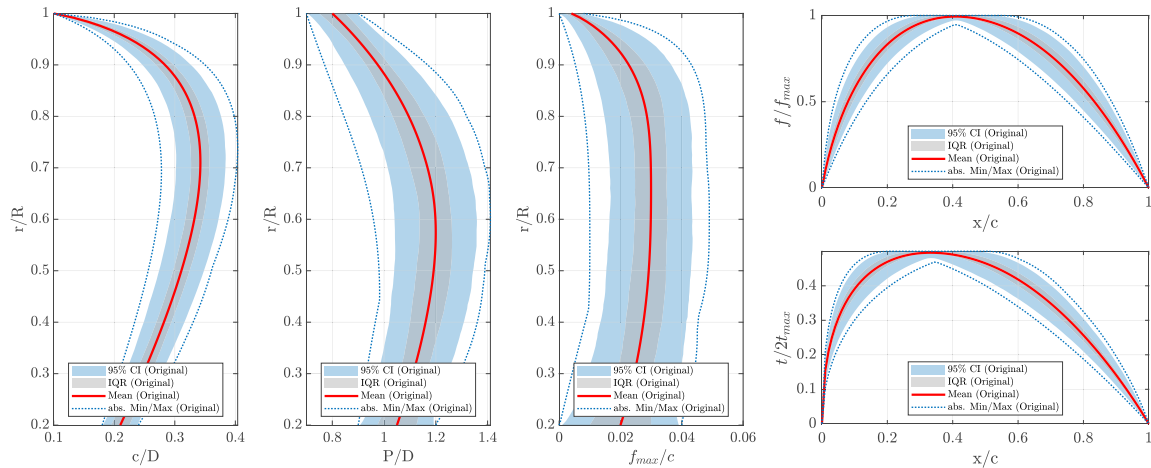


Fig. 10. Envelopes of all the possible geometries within the original design space (mean distributions, minimum and maximum values, 95% intervals of confidence and inter-quartile ranges).




Fig. 11. Results of the optimization: Pareto front on relevant planes in the space of outputs, original design space.

root ($r/R < 0.7$) and tip ($r/R \geq 0.7$, worst value between the two zones). In the optimization process (and for the selection of the optimal geometry), in any case, they are handled separately. Furthermore, a qualitative measure of the Pareto convergence of the process is provided in Fig. 12, where the Pareto front in the efficiency/tip vortex strength projection plane is analyzed as a function of the evolution of the genetic process.

Compared to the reference propeller, which has been already observed to have very poor performance in Fig. 6 and in Table 1, several configurations in the original design space provide optimal performance simultaneously at design, loaded (back side cavitation) and unloaded (face side cavitation) conditions, and for all contrasting objectives included in the design problem.

In a multi-objective setting there is no absolute optimum; candidate designs lie on a Pareto front. Selecting a final geometry therefore requires balancing conflicting objectives (with weights informed by prior experience), accounting for solver credibility in the relevant regimes, and applying qualitative criteria not yet encoded in the KPIs—for example, the operational consequences of distinct cavitation modes and whether a controlled amount of cavitation is tolerable if it suppresses more dangerous effects. Given the performance of the reference propeller, the criteria adopted for the choice of an alternative design were based on the avoidance of pressure side cavitation, and the reduction of the risk of tip vortex cavitation, possibly without worsening the efficiency and the margins against suction side cavitation observed for the reference propeller, which indeed were selected very close (see the thresholds of $C_{PN,back,LE}^{tip}$, $C_{PN,back,LE}^{root}$, $C_{PN,back,mid}^{tip}$ and $C_{PN,back,mid}^{root}$ in Table 1) to those of the original case to filter the most likely optimal




Fig. 12. Convergence of the Pareto front in the efficiency/tip vortex strength projection plane of the space of solutions (original design space).

geometries among the feasible solutions. In this respect, it is worth noting that the original geometry was already not affected by suction side cavitation at the design functioning, and only at maximum speed there was the risk of bubble cavitation at the root. Ensuring only

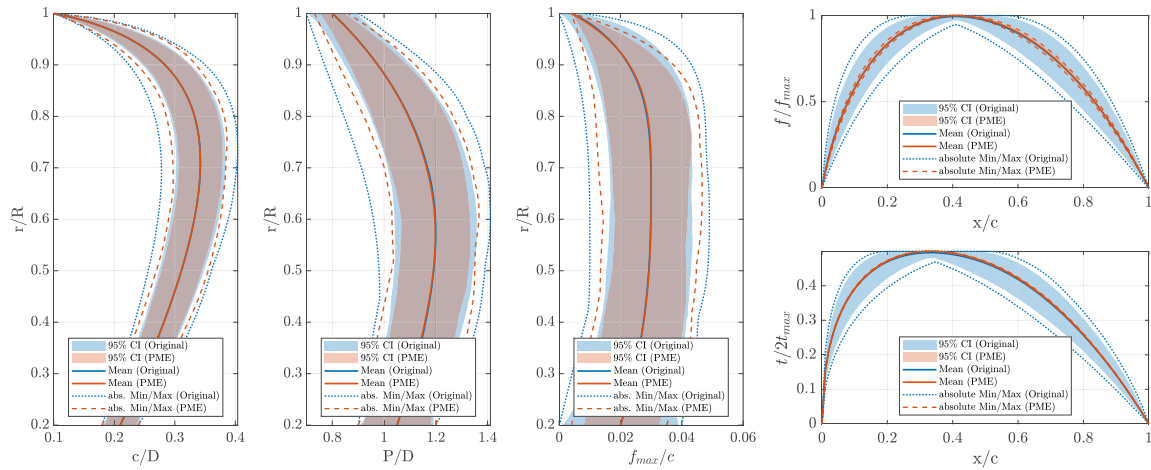


Fig. 13. Envelopes of all the possible geometries within the PME reduced design space (mean distributions, minimum and maximum values, 95% intervals of confidence).

slightly lower values of cavitation inception on the suction side (i.e., not accepting any geometry worse than the original) seemed sufficient to focus the design towards the improvements of the critical aspects of the reference blade shape. After post-processing a set of Pareto candidates, ID-156471 was identified as the preferred geometry among the feasible designs. While maintaining substantially unchanged efficiency on the basis of the equivalent steady BEM calculations of the optimization ($\eta_o = 0.6595$ versus 0.6558 of the reference propeller), it delays pressure-side cavitation and retards suction-side phenomena, even for the loaded condition. Inception on the pressure side is almost halved at the design functioning (from $-C_{PN} = 2$ to $-C_{PN} = 1.13$) and a similar improvement is observed also for the equivalent unloaded condition, naturally more prone to pressure side cavitation (from $-C_{PN} = 2.98$ to $-C_{PN} = 1.41$). This is even more satisfactory when these improvements on the pressure side are discussed considering the achieved performances on the suction side, which are better than those of the reference propeller as well. At the design functioning, in particular, inception is postponed at $-C_{PN} = 1.33$, which is far better than the 1.57 of the original configuration. For the loaded condition, the maximum suction for the selected geometry is 1.68 , making this geometry completely non-cavitating under the equivalent steady state calculations, also at the maximum ship speed functioning condition (cavitation index equal to 1.8).

These performance are the result of some significant, potentially not expected modifications in the blade geometry, some of which affect the shape of the sectional hydrofoil. The most relevant changes in the radial distributions with respect to the original blade shape can be seen in any of the comparisons dedicated to the reduced spaces, for instance Fig. 16. The expanded area ratio of this optimal propeller is almost unchanged; only the maximum value of the chord is slightly moved towards the root of the blade causing, consequently, the reduction of the maximum sectional thickness at the root. The pitch is heavily unloaded at tip, and this contributes to the almost halved value of the vortex strength that characterizes ID-156471 and that reasonably will operate to postpone the inception of the tip vortex (Section 6). Despite the much lower value of pitch at tip (counterintuitive given the risk of pressure side cavitation observed for the reference propeller), pressure side cavitation is avoided thanks to a simultaneous reduction of the camber, not simply the maximum sectional value but especially the curvature at leading edge (the new hydrofoil shape has a “more flat” nose), and the shift of the maximum of the sectional thickness at about $x/c = 0.48$ (for the reference geometry it was at about 0.38). In the central portion of the blade, where margins against cavitation inception were already higher, the optimized geometry is characterized by higher values of pitch that balance the reduction at tip to ensure

the required delivered thrust. Given these performances, ID-156471 appears as a promising alternative to the original geometry and as the ideal candidate for assessing which kind of improvements (and to what extent with respect to a traditional, but expensive, design by optimization) are possible when the various reduced design spaces replace the original parametrization of the geometry.

5.2. Reduced design space: PME approach

For building the reduced spaces, a set of 4000 geometries was generated by sampling with a Sobol algorithm the original design space. The corresponding discretized geometries, through the $L = 1326$ surface nodes of the BEM surface grid (then, 3978 metric values of x -, y - and z - of each grid node) represent the snapshots based on which the dimensionality reduction is achieved, along with the corresponding $M = 38$ design variables of the original parametrization from which any sample is derived.

The selected design space consists of 13 dimensions (“modes”), which represents a 70% reduction in the design space dimensionality and retains more than 98.5% of the variance associated with the blade geometry. This choice, as discussed, was dictated by the need to comply with the multi-objective, constrained nature of the optimization problem, for which it was possible to achieve feasible geometries with no fewer than 13 reduced design variables. Despite the very high value of (geometrical) variance embedded in the 13-dimensions reduced design space, the limitations of a reduction strategy based solely on geometrical features when applied to the complex geometry of a propeller, as seen in Gaggero et al. (2020) and Villa et al. (2020), are evident just from the comparison of the envelopes of all the possible geometries permitted in this PME-reduced space (Fig. 13) with those from the original design space (Fig. 10).

While global variations, such as those in chord and pitch, are indeed retained by the reduced space because they have the highest purely geometrical weight, local modifications, such as those to the sectional hydrofoil, are almost completely sacrificed to dimensionality reduction. The envelopes of the radial distributions reconstructed in the PME-reduced space (defined as the absolute minimum and maximum values attainable across the space) are narrower than those observed in the original space. Nevertheless, the confidence intervals remain largely unchanged. This indicates that most of the information related to the extension of the design space is preserved by the reduction process. With these metrics, no conclusions can be drawn regarding the variability of the reconstructed shapes. The only noticeable difference is a slight loss of information in the tails of the statistical distributions describing the geometry shape. On the contrary the resulting camber

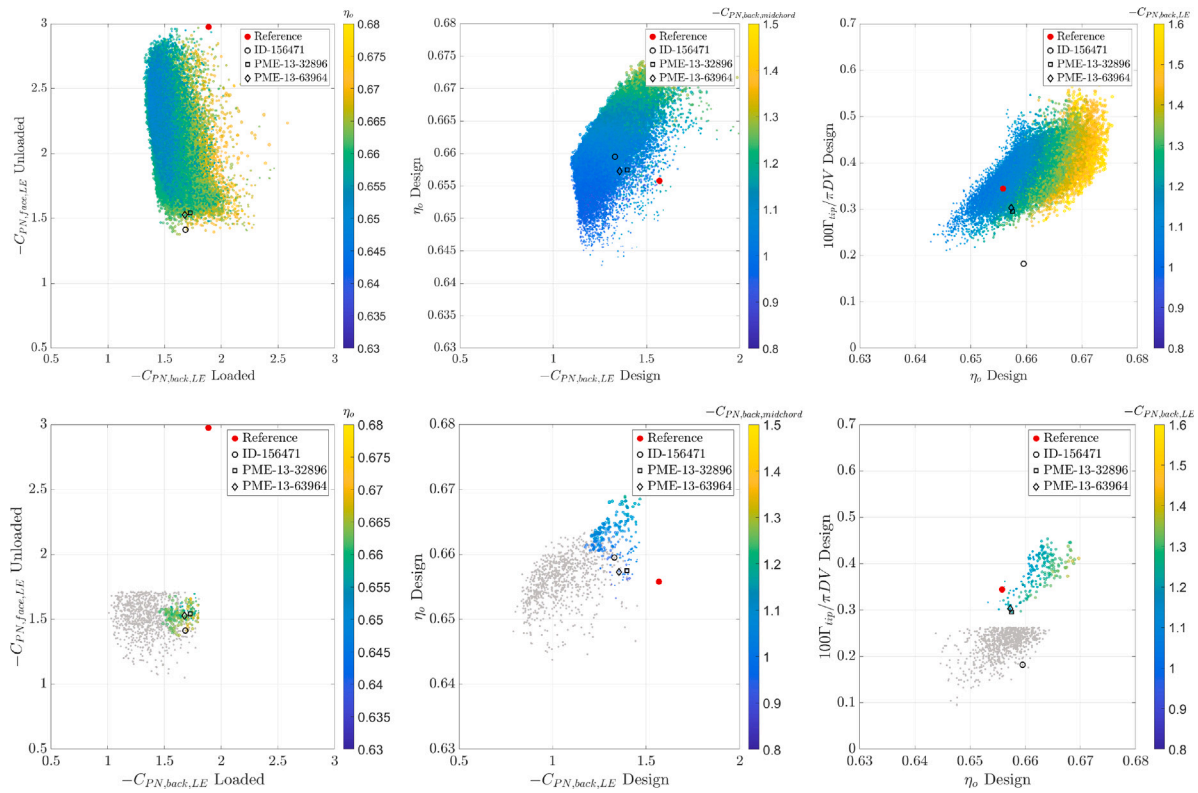


Fig. 14. Results of the optimization using the PME-13 reduced design space (13 “modes”); on the top, Pareto front on relevant planes in the space of outputs. On the bottom, comparison with the feasible geometries (gray dots) from the original design space, filtered with the thresholds of Table 1; the performance of the selected optimal propeller in the original design space (ID-156471) are given as a further measure of the “limits” of the reduced space.




Fig. 15. Convergence of the Pareto front in the efficiency/tip vortex strength projection plane of the space of solutions (within PME-13 reduced space).

and thickness sectional distributions are almost unchanged compared to the reference shape, because from the absolute geometrical point of view, the variation in space of a grid node due to a change in sectional camber (or thickness) is order of magnitude less relevant than the displacement due to a variation in the chord (or the rake distribution). This is evidenced by the envelopes of these distributions being nearly coincident with their mean values, as well as by the confidence intervals, which are narrowly concentrated around the mean. Losing variability in the camber and thickness sectional distributions

narrows the space of solutions compared to the non-reduced domain, and prevents features modifications, such as those due to sectional camber and thickness, which may have a significant influence on performance and cavitation avoidance based on the outcomes of the previous optimization in the non-reduced space. Indeed, although a performance improvement is possible compared to the reference geometry (especially concerning the risk of face-side cavitation, see Fig. 14) and the convergence of the Pareto front is smooth and fast as a consequence of the very small number of design parameters (see Fig. 15), there are significantly low margins of action about suction-side cavitation under loaded conditions, or tip vortex strength, since only a fraction of the reduction foreseen by the optimization in the original space is possible. The limitations of the PME-reduced design space in providing enough (or appropriate) shape variance to approach the optimal performances observed in the original design space are even more evident when a direct comparison with the original ID-156471 geometry (or with all feasible configurations of the original design space having inception values lower than the selection thresholds listed in Table 1) is addressed. A measure of the limits of the PME method is given by the value of vortex strength at the tip. Setting the threshold to the value adopted in the original design space to restrict the feasible geometries to a limited set having overall better performance than the reference prevents any PME reduced geometry from being selected. Acceptable geometries are obtained only if the margin on tip-vortex strength is relaxed. Even then, their overall performance is lower than the corresponding designs from the original design space. None of the PME reduced geometries, even at higher values of tip vortex, permits indeed the simultaneous substantial reduction of inception at the loaded (back side cavitation) and at the unloaded equivalent functioning conditions (face side cavitation) observed instead in the original design space, nor the substantial improvements at the design functioning possible instead in the original space.




Fig. 16. Radial distributions of chord, pitch, maximum thickness and camber for PME-13-32896 and PME-13-63964 compared to the reference and the selected optimal propeller from the non-reduced design space; data normalized with the maximum of distributions of the original geometry.

Given the poor performance of the reference geometry on the pressure side, substantially driven by the pitch, which is a geometrical feature successfully addressed in the PME reduced design space, some sub-optimal configurations can be identified even in this reduced space. The selection criteria, with the exceptions discussed above, are similar to those adopted in the non-reduced space: favoring the reduction of the risk of cavitation on the pressure side rather than postponing suction side phenomena further, as these were already within acceptable limits even in the loaded condition (and in the end not fully possible in this design space). Two configurations (PME-13-32896 and PME-13-63964) comply with these criteria and, among the filtered designs, they are those with the greatest possible reduction in vortex strength without worsening sheet cavitation or efficiency. These results are obtained by a certain increase of the expanded area ratio (Fig. 16) that is, for both, about 10% higher than the reference, and a reduction of the sectional maximum camber at tip. The maximum sectional thickness, which is strictly correlated with the chord distribution, also changes and the selected geometries provide identical structural resistance through thinner but longer root sections. Pitch is increased at root, to compensate for the reduction in camber, and is almost identical to the reference at the tip. This behavior, which is identical for both the selected geometries, suggests a certain interdependence between modifications at root and at tip when the reduced variables are employed, as if a simultaneous increase of pitch at root and its reduction at tip would not be possible by the combination of the “modes” of the reduced design space, explaining the relatively poor performances at tip of both propellers despite the apparent possibility (see Fig. 13) of further reducing the pitch at the outermost positions of the blade.

5.3. Reduced design space: PI-PME approach with 12 parameters

The dimensionality reduction using the PI-PME requires the enrichment of the geometrical variations of PME with physical quantities, which may be lumped performance indicators (such as thrust and efficiency), as well as distributed variables (such as pressure distributions). In light of the current optimization problem, adding pressure fields to the description of the snapshots based on which reduced variables are obtained appears to be the most promising strategy. In the definition of the combined geometrical and physical variance of the reduced space, this includes information (pressure, then inception) relevant for the objectives and the constraints of the design. To this end, the same 4000 samples adopted for the PME underwent an hydrodynamic analysis process similar to that adopted in the optimization process, exactly aimed at collecting the same lumped and distributed parameters and

performance indicators in accordance with the three operating conditions relevant to the design problem of Eq. (1). This results in snapshots composed each by 7769 elements, including the 3-D coordinated of each surface grid points defining the blade geometry (\mathbf{D} in Eq. (8)), the value of the pressure in the center of each surface panels for the three functioning conditions relevant for the design ($1250 \cdot 3$, \mathbf{F} in Eq. (8)), the parameters of the original parametric description associated with that shape (\mathbf{U}), and the values of thrust, efficiency, and tip vortex strength at the design functioning as lumped indicators (\mathbf{C} in the same equation).

Using the approach briefly described in Eq. (10) and detailed in Serani et al. (2025), physics-informed reduced spaces, in which the importance of “modes”, and thus the identification of the most relevant geometric features concurring to the dimensionality reduction depends not only on the displacement of points in space but also on the variation of physical quantities, can be defined. In this respect, physical information gives importance to (local) geometrical features that would otherwise be disregarded by purely geometrical dimensionality reduction-driven methods due to their very low importance in terms of pure spatial displacement, regardless of the crucial role they play in terms of influence on performances. This is the case of the sectional thickness and camber distributions, which, compared to PME, take advantage of the significant influence they have on pressure distributions and, already from a 12-parameters reduced design space, contribute significantly to the variability of the allowed shapes of the PI-PME model. With a variance of 95% provided by the 12 “modes”, the envelopes of all possible geometries (see Fig. 17) are much richer than the PME with 13 parameters (equal to 98.5% of the geometrical variance). Chords, or pitch, which are naturally sustained by the large geometrical displacement they realize, have envelopes using the PI-PME-12 already comparable to those from PME-13 or to those from the original design space, for sure in terms of intervals of 95% confidence (but also the absolute maximum and minimum values are a bit closer to those of the reference space compared to the PME case). At the same time, despite the reduction of about 60% in the original dimensionality, the sectional hydrofoil variability is included and is comparable to that achieved in the original design space. The confidence intervals associated with both thickness and camber obtained using the 12 “modes” PI-PME model remain narrower than those observed in the original design space. Nevertheless, the extrema reached in the reduced design space fully bound these intervals, indicating that the reduced geometries are capable of reproducing the most probable configurations of the original space. Compared to PME, with only 12 parameters the PI-PME-12 favors variability of the sectional distributions of camber (or thickness) rather than its maximum sectional value (especially at

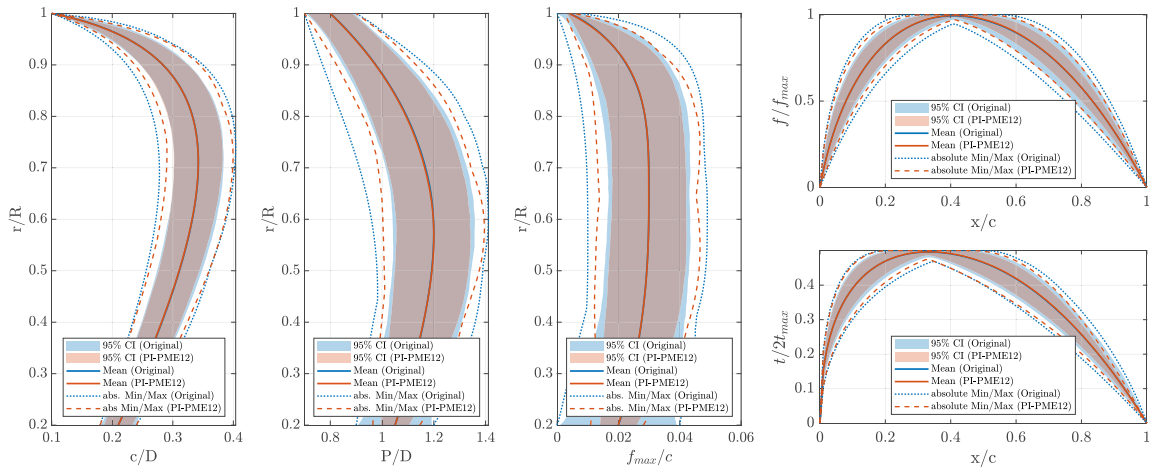


Fig. 17. Envelopes of all the possible geometries within PI-PME-12 reduced design space (mean distributions, minimum and maximum values, 95% intervals of confidence).

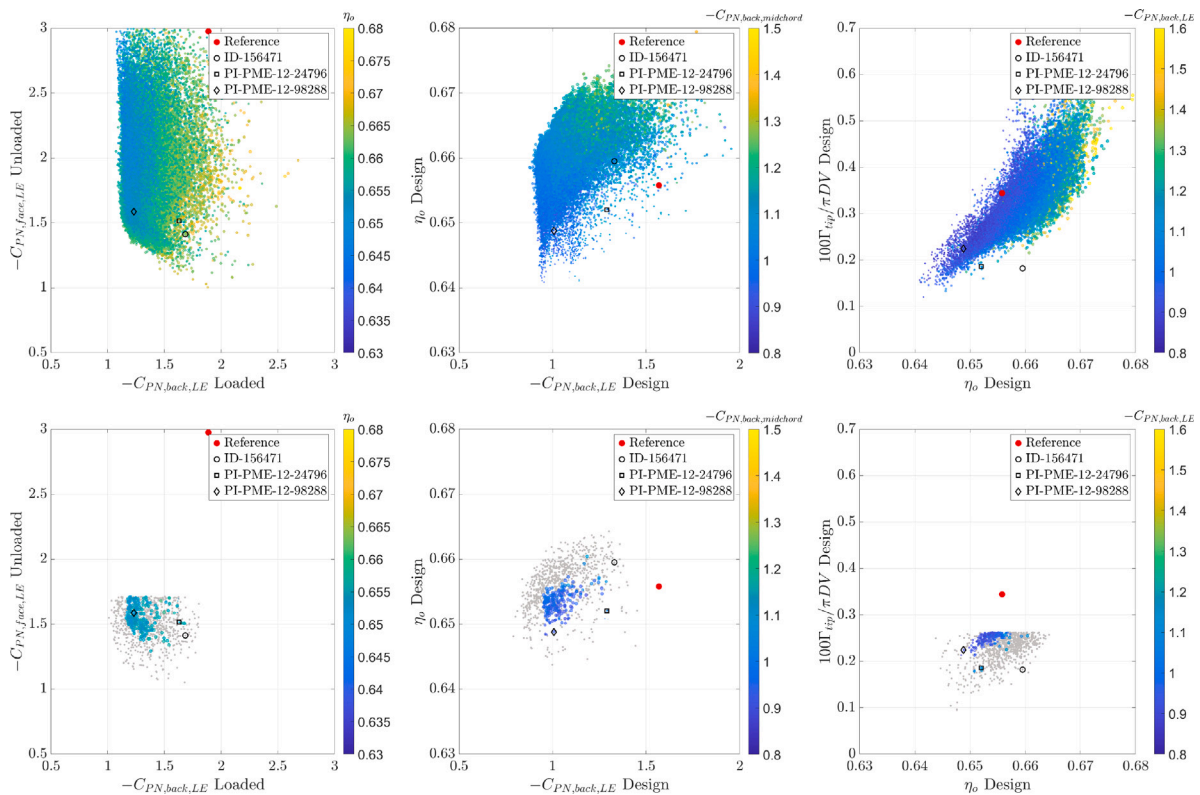


Fig. 18. Results of the optimization using the PI-PME-12 reduced design space (12 “modes”); on the top, Pareto front on relevant planes in the space of outputs. On the bottom, comparison with the feasible geometries (gray dots) from the original design space, filtered with the thresholds of Table 1; the performance of the selected optimal propeller in the original design space (ID-156471) are given as a further measure of the “limits” of the reduced space.

root), contributing to a substantial widening of the space of solutions (see Fig. 18), which is more densely populated and qualitatively more similar to that from the original design space. As for the PME-13 case, the convergence of the solution observed in the Pareto fronts of Fig. 19 is smooth and faster than in the original design space thanks to the smaller number of design parameters, but with the inclusion of physical information in the selection process of the reduced parameters the space of solution is less affected by the reduction process. Even if the number of filtered designs is not comparable to that of the original space, no relaxation of the tip vortex strength selection criterion was necessary to identify potentially optimal geometries, which instead show performance sufficiently close to those of ID-156471. At

comparable levels of tip vortex strength, there are geometries with very interesting margins against suction side cavitation (design, as well as loaded functioning condition) at almost nullified pressure side cavitation, since also in this case the worst cavitating phenomenon of the reference propeller is more than halved in terms of risk of inception. The efficiency of the filtered geometries is, instead, overall lower than the values achieved by the PME configurations (the higher tip vortex strength required to enable the choice of potentially non-cavitating geometries in that reduced design space easily explain this difference), but it seems an acceptable cost for achieving tip-unloaded, completely non-cavitating, geometries potentially at a fraction of the effort required in the reference design space.




Fig. 19. Convergence of the Pareto front in the efficiency/tip vortex strength projection plane of the space of solutions (PI-PME-12 reduced space).

Among these geometries, PI-PME-12-24796 and PI-PME-12-98288 seem rational optimal choices, especially from a comparative point of view, since only their efficiency is marginally lower than that of ID-156471 or that of the reference propeller. Both, instead, have a very reduced tip vortex strength (that of PI-PME-12-24796 is identical to ID-156471) and similar pressure side performances; PI-PME-12-98288 provides significantly better margins against suction side cavitation, which is an interesting feature to be investigated in fully unsteady functioning or through the equivalent open water RANS analyses of Section 6. This difference is a consequence of the different combination of pitch and camber observable in Fig. 20: PI-PME-12-98288, at intermediate radial positions, has a pitch 5% lower than PI-PME-12-24796 and a maximum sectional camber of about 3% of the local chord versus the 2% of PI-PME-12-24796. Chords are overall larger, and the expanded area ratio is the biggest among all the selected geometries (reference, original, and any reduced design space). At tip, pitch is almost identical, and the small difference in tip vortex strengths can be associated with the slightly different radial distributions of load provided by the mutual influence of the geometrical characteristics of the two propellers.

5.4. Reduced design space: PI-PME approach with 18 parameters

Adding “modes” increases the combined variance of the reduced space. It reaches the 98.5% achieved by purely geometrical variations of the 13-dimensions PME approach when the number of reduced parameters of the PI-PME model is equal to 18. The consequences of this increase, which still corresponds to a relevant reduction (about 53%) in the dimensionality of the original design space, are evident both in terms of possible shapes included in the design space (see Fig. 21) and in terms of widening of the space of solutions.

Compared to the previous ones, the envelopes of all the possible radial and sectional distributions of geometrical features of the blade are the largest, almost overlapped (at least in terms of achievable intervals of confidence and, in the case of sectional distributions, also in terms of minimum and maximum values) with those from the original design space. If the inclusion of physical quantities enabled, already with 12 reduced parameters, a certain variability in the sectional distributions of camber and thickness, the 18-dimensional design space permits larger variations for pitch and maximum sectional camber (similar to those of the reference design space) also at the tip and at root. As seen in Fig. 22, this has consequences for the performance when this reduced

space replaces PME-13 or PI-PME-12 in the design process. The space of solution is much larger, and the geometric variability allowed by the 18 parameters permits performance improvements comparable to those enabled in the original design space. This feature is particularly clear through the comparison between the filtered designs (using the thresholds of Table 1) from this PI-PME-18 design space and those from the original design. The only noticeable difference is a slightly worse Pareto front in the tip vortex/propeller efficiency projection of the space of solutions, while all other performances from geometries belonging to the PI-PME-18 design space similarly replicate those from the original space. The convergence trend (see Fig. 23) is fast and smooth as well, showing that for the couple of efficiency and tip vortex strength indicators, almost no improved performance from the tenth generation, at least over a large portion of the Pareto front, can be furthermore achieved. For the sake of comparisons, PI-PME-18-24924 and PI-PME-18-66781 are well-balanced geometries, able to nullify any risk of cavitation and reduce the tip vortex strength to values comparable to that of ID-156471: PI-PME-18-66781 is particularly effective in improving also the margins against suction side cavitation; PI-PME-24924 at identical efficiency of the reference propeller has the lowest circulation at tip without any risk of pressure side cavitation. The modifications of the geometry responsible for these performances are an increase in the expanded area ratio, which remains higher than that of the reference propeller and of ID-156471, and a combined modification of pitch and camber which, as expected, makes these two geometries very similar to the optimal configuration in the original design space, as observable in Fig. 24: tip unloading is compensated by a slightly higher pitch in the central part of the blade, while the maximum sectional camber is reduced everywhere from root to tip. The sectional hydrofoil is modified as well. The thickness distribution is almost identical between PI-PME-18-24924 and PI-PME-18-66781 and very similar to a standard NACA16 shape. Camber line, instead, approximately maintains the distribution of the reference propeller for PI-PME-18-24924 while it is flattened at the leading edge for PI-PME-18-66781, as for ID-156471.

5.5. PME and PI-PME parameters composition

The procedure adopted for constructing the PME and PI-PME reduced design spaces allows for a detailed assessment of the contribution that each of the original design parameter has to each reduced-space parameter, as well as of the relative influence of the underlying information sources (geometrical or physical) in defining the “importance” of the different “modes”—that is, the share of variance each reduced parameter retains within the reduced space. These information are collected in Figs. 25 and 26. The analysis is supported by a *t*-test performed on the data from the original design space (see Fig. 27). This statistical test, although relatively basic, provides an indication of both the strength and the direction (positive or negative) of the relationship between input and output variables. While the results cannot be regarded as conclusive, given the limited number of samples and their non-uniform distribution across the design space, being the outcome of an optimization process, they remain useful for identifying underlying trends. When combined with established expertise in propeller design, these trends serve to reinforce the findings obtained from the reduced parameter space analysis.

Fig. 25 shows the normalized components of the first 12 *k* “modes” of the PME and PI-PME reduced spaces along the vertical axis, plotted against the original design variable on the horizontal axis. Large peaks indicate which design variable dominates that particular mode. The participation of each data source (i.e., geometry, lumped performance indicators, like thrust coefficient and efficiency, and distributed performance indicators like the pressure coefficient) to the variance retained by each of the first 12 “modes” of both the reduced design spaces is shown, instead, in Fig. 26. This latter diagram, in particular, is of utmost importance, since it reveals whether a mode primarily




Fig. 20. Radial distributions of chord, pitch, maximum thickness, and camber for PI-PME-12-24796 and PI-PME-12-98288 compared to the reference and the selected optimal propeller from the non-reduced design space; data normalized with the maximum of distributions of the original geometry.

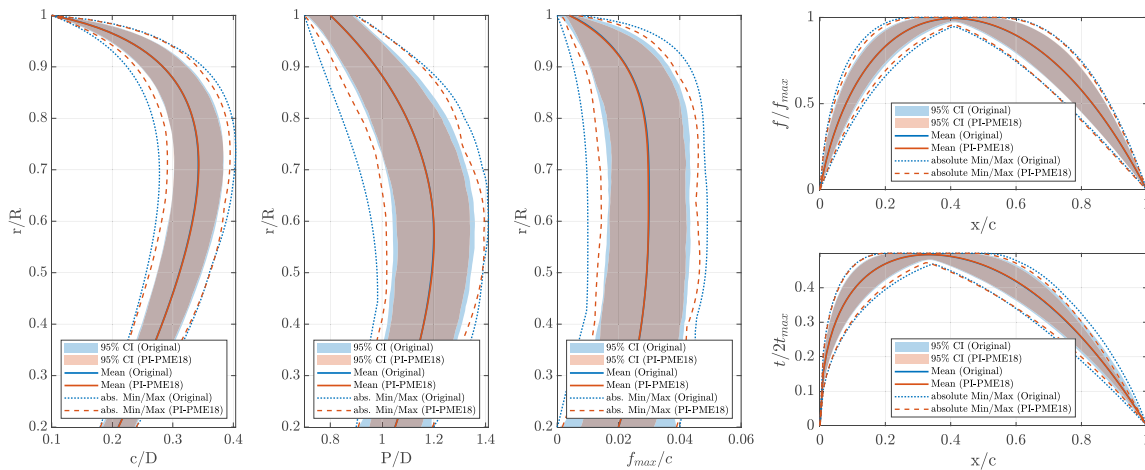


Fig. 21. Envelopes of all the possible geometries within PI-PME-18 reduced design space (mean distributions, minimum and maximum values, 95% intervals of confidence).

represents (or takes importance from) geometric variations, physical quantities or a mix of the two and, together with the information circa the contribution to each “modes” of the parameters of the original design space, it clarifies the level of correlation between geometry and physics, leading to fewer “modes” (or shared ones) when the correlation is strong or to more “modes”, each owned by geometry or physical quantities in case of weak correlation. Then, it is particularly useful for discussing the PI-PME reduced space and the relevance of the pressure distributions for the identification of the most important modes oriented to cavitation avoidance.

The first “mode” (the one carrying the largest part of the variance) of the PME space is participated mainly by the variation of the intermediate control point of the chord (c_{y2}), which causes the most relevant modifications of the propeller shape, interpretable as variations of the expanded area. The second and the third “mode” of the PME identify, instead, the rake (through parameters r_{y2} and r_{y4}), as the second most influential parameter to the geometry variance. This behavior is plausible too, since the “shift” of blade sections in the longitudinal plane, which is the modification induced by the rake, contributes more to differences in blade shapes than the localized modifications induced by changes in the camber and thickness (by radial or sectional perturbations of control points) that, indeed, never appear in the first 12 “modes” under considerations as the most important contributors of

Table 3

Comparison of propeller performances at the design point using different tools: kinematic equivalent BEM (steady, non-uniform radial wake), unsteady BEM in the nominal wake of Fig. 4, open water RANS at equivalent advance coefficient.

	Equiv. BEM		Unst. BEM		Equiv. O.W. RANS	
	K_T	η_o	K_T	η_o	K_T	η_o
Ref.	0.1860	0.6558	0.1781	0.6428	0.1813	0.6588
ID-156471	0.1859	0.6595	0.1746	0.6449	0.1814	0.6551
PME-13-32896	0.1865	0.6575	0.1756	0.6434	0.1815	0.6527
PME-13-63964	0.1879	0.6572	0.1764	0.6426	0.1811	0.6521
PI-PME-12-24796	0.1886	0.6520	0.1770	0.6371	0.1830	0.6488
PI-PME-12-98288	0.1888	0.6488	0.1768	0.6327	0.1842	0.6447
PI-PME-18-24924	0.1872	0.6568	0.1756	0.6415	0.1796	0.6531
PI-PME-18-66781	0.1851	0.6530	0.1737	0.6382	0.1770	0.6487

any reduced parameters. PME, by its nature, is participated only by the geometry (Fig. 26) and it is expected that the reduced parameters of this space take care of the most relevant geometrical modifications, which are, first of all, those promoting the largest displacements of the grid points of the geometry. On the contrary, for the PI-PME reduced model, the geometry participates significantly to the variance retained by the modes only from the third (and even more in “mode” 4 and 5), because

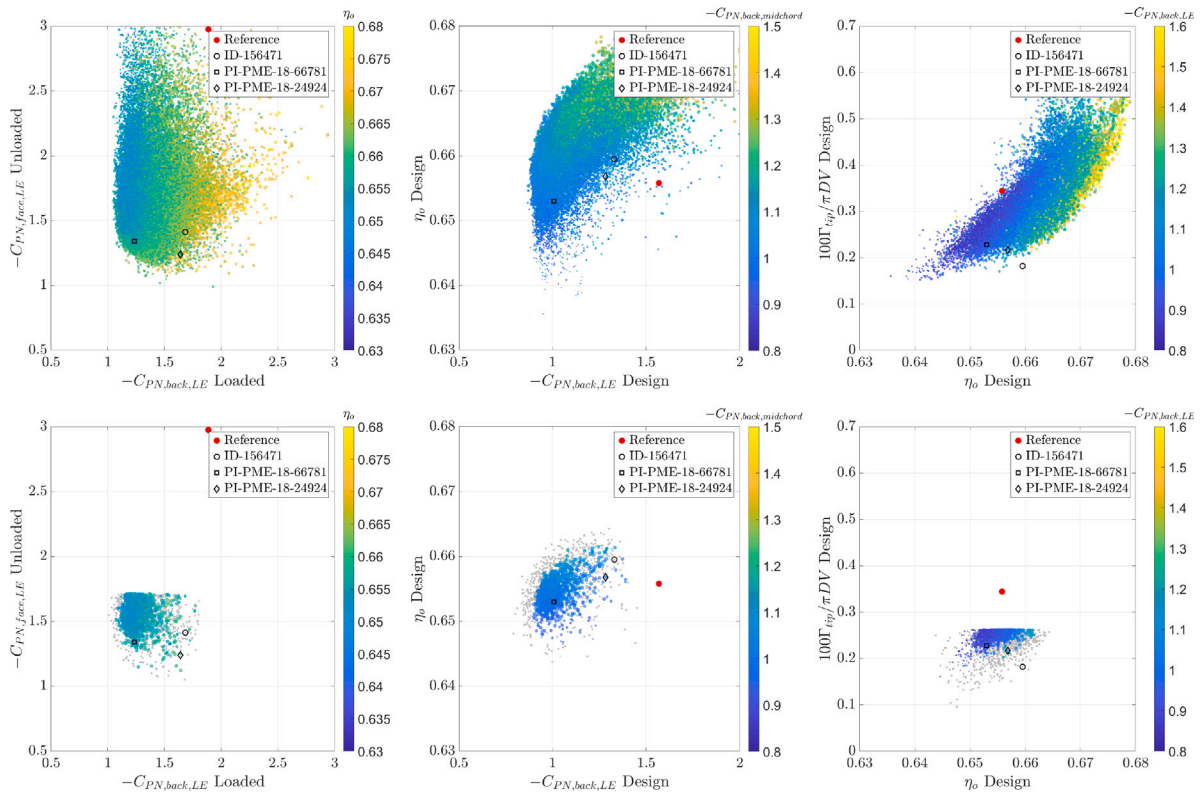


Fig. 22. Results of the optimization using the PI-PME-18 reduced design space (18 “modes”); on the top, Pareto front on relevant planes in the space of outputs. On the bottom, comparison with the feasible geometries (gray dots) from the original design space, filtered with the thresholds of Table 1; the performance of the selected optimal propeller in the original design space (ID-156471) are given as a further measure of the “limits” of the reduced space.




Fig. 23. Convergence of the Pareto front in the efficiency/tip vortex strength projection plane of the space of solutions (PI-PME-18 reduced space).

for the two most important ones, the contribution to the variance comes from physical quantities. In this respect, it is interesting to observe the very nice correlations between the geometrical features and the physical data revealed by the PI-PME method, which quantitatively explain the discussion based on the results of the optimization process. The first mode, which is mainly participated by efficiency, identifies the pitch, and in particular its values at tip through pd_{y3} , pd_{y4} and pd_{y5} , and the camber, through f_{y2} , f_{y3} and f_{y4} , as the most important contributors of

the mode. Regardless the effect, positive or negative, which cannot be guessed by this analysis, this trend is exactly in line with the outputs of traditional design methods and established design guidelines and, more quantitatively, it agrees with the results, extracted from the database of the optimization in the original design space, through the t -test of Fig. 27. Efficiency, also from this analysis, is mainly influenced by the value of the pitch at tip and by the radial distribution of camber through the same set of parameters. Combining the information from Figs. 25 and 26, it is clear that the second mode establishes a correlation between the delivered thrust and the strength of the tip vortex circulation, again with the pitch and maximum values of camber. The comparison with the t -test (where the increase of pitch at tip, pd_{y5} is the most important parameter responsible for the increase of Γ) reinforces the validity of this obvious correlation that substantiates the method’s capability to effectively capture the relationship between geometry and physics and to preserve it in the formulation of the reduced variables.

The third mode of the PI-PME method, which is the first in which geometry plays a significant role, is an example of the coupled contribution of geometric and physical information. The primary source of information in this mode is efficiency, followed by geometry. The mode is participated by pitch and tip camber (as in the first mode, given obviously their influence on efficiency), but also by chord. The chord is particularly relevant, as it not only drives most of the geometrical variance as shown in the PME reduced space due to the large grid-point displacements induced by its modification, but also enhances efficiency, since its reduction (i.e., of the expanded area ratio) decreases the wetted area and consequently the frictional forces.

When physical information such as the pressure coefficient is considered, the physics-informed model links the relevant features to pitch (suction peaks at the leading edge) and camber (pressure distributions over the mid-section of the hydrofoil), primarily through modes 4 and 8–12. Again, this outcome is not particularly different from the results




Fig. 24. Radial distributions of chord, pitch, maximum thickness, and camber for PI-PME-18-24924 and PI-PME-18-66781 compared to the reference and the selected optimal propeller from the non-reduced design space; data normalized with the maximum of distributions of the original geometry.

of the statistical analyses on the database of solutions in the original design space, which identify through the sectional camber and thickness distributions the relationship (at least the most relevant one) between the pressure coefficient and the geometry, but it remarks the capability of physical informed design space reduction methodologies to provide relevance to geometrical features crucial for the objectives (or constraints) of the design problem but uninfluential if addressed from the pure geometrical (i.e., spatial displacement) point of view. These modes, indeed, project into the reduced space variance contributions that correspond to alternative geometrical configurations of pitch and sectional blade shape, induced by the pressure field or, more generally, by performance indicators of the design problem. Specifically, mode 4 involves pitch and sectional camber, while modes 8–12 incorporate sectional camber and thickness distributions, which are features always discarded in any of the PME modes. These are examples of weakly correlated geometry and physical performance, causing the need for some extra modes to accommodate both sets of unaligned directions.

Considerations on rake further highlight the advantages of PI-PME over purely PME-based reduction techniques. The t -test conducted on the original database indicates that rake has no statistically significant influence on the design objectives and, in fact, exerts one of the weakest effects on the analyzed performance indicators. Within the PME framework, however, rake contributes substantially to the most influential modes (the second and third), such that geometries reconstructed from the reduced space are primarily distinguished by variations in rake distribution rather than by more relevant features, such as radial (or sectional) camber, which are critical for optimization. By contrast, under PI-PME the rake is confined to the sixth mode, of comparatively low importance, and is used almost exclusively to represent geometric variation in the reduced space (see Fig. 26). The preceding dominant modes are, through the physical information, allocated to pitch and sectional camber, which are variables that decisively govern the achievable performance.

6. Detailed performances of the optimized propellers

As a final assessment of the optimization results, aimed at verifying the quality of the optimal geometries from the dimensionality reduction strategies and the reliability of the simplified design criteria (i.e., cavity inception and tip vortex), all the selected geometries were analyzed in fully unsteady conditions (BEM, design advance coefficient) and, using RANS, at equivalent open water conditions. For the latter, $J = 0.8$, $J = 0.87$, and $J = 0.95$, respectively, are the advance coefficients in uniform flow selected to enable simpler steady analyses comparable to the

loaded, design, and unloaded conditions included in the optimization process by the kinematic equivalent radial distributions of velocity representative of the various blade positions inside the nominal wake. The comparison of performances is listed in Table 3 while the instantaneous pressure coefficient at the design advance coefficient are shown in Figs. 28 and 29. Pressure distributions are preferred to representations of the cavity extensions at first for a coherent comparison with the data from the optimization. Moreover, since any type of cavitation, regardless of the position of the blade in the wake, is avoided for any of the geometries under investigation, no useful information could have been gathered from cavitating analyses. Calculations for the reduced advance coefficient ($J_A = 0.84$), considered for the reference propeller, are also available. However, since there is no risk of cavitation, which further confirms the success of the optimization process, the discussion can be limited to the unsteady functioning at the design advance coefficient for the sake of conciseness. Performances from the kinematic equivalent BEM calculations are those already discussed in the optimization problem and during the selection of the optimal geometries. The thrust is within the 1.5% threshold of the design problem and the efficiency is predicted to be higher than the reference in the original design space. Also the geometries from the PME reduction, as a consequence of the need to relax the strength of the tip vortex to enable acceptable geometries in terms of cavity inception, have a higher efficiency, and only one geometry from the PI-PME models (PI-PME-18-24924) permits similar results. These trends are overall replicated by the averaged unsteady performances and the data from the equivalent open water RANS.

Despite the difference between the kinematic equivalent BEM steady and averaged BEM unsteady performances already discussed when analyzing the reference geometry, any optimized geometry delivers a thrust within an acceptable tolerance with respect to the corresponding value (same computational setup) of the reference propeller. In fully unsteady conditions, all the geometries provide a thrust within 1.5%, with the exception of ID-156471 and PI-PME-18-66781, which are between 2 and 2.4% lower than the averaged unsteady thrust of the reference. In open water, at the equivalent advance coefficient, RANS results are even closer, within 1.5% except for PI-PME-18-66781, which, as for the averaged unsteady calculations, delivers a thrust in open water conditions 2.2% lower than the corresponding RANS analysis of the reference. The trends of the efficiency are a bit less coherent among the different calculation strategies, but together with the data of thrust, overall, confirm the reliability of the simplified calculation methods employed in the design process. The unsteady functioning using the same BEM of the equivalent steady analyses (then the same calculation




Fig. 25. Eigenvectors (“modes”) embedding the original design variables.

method for frictional forces) confirms the geometry from the original design space and one from the PME as those having values of efficiency higher than the reference propeller. The second geometry from the PME and PI-PME-18-24924 have a comparable efficiency, and also when addressed with the RANS solver in uniform inflow, this set of geometries (original design, those from PME and PI-PME-18-24924), despite having lower efficiency than the reference, confirm the trend of the optimization process, which identified those cases as the most efficient. In any case, the differences in efficiency with respect to the corresponding reference propeller case are lower than 1.2%, which is a value perfectly within the precision of the employed CFD tools.

The analysis of the unsteady pressure fields (and their comparison with the equivalent open water data from RANS), which is more

useful in the case of completely non-cavitating propellers, evidences the improvements achieved by the optimization process. The heavy suction observed at the leading edge on the pressure side (see Fig. 5), which causes the sheet cavitation shown in Fig. 6, is completely avoided for any selected blade and at any blade position, even where the influence of the tangential velocities (270 deg blade position) is stronger; geometries from the PI-PME-18 reduced space have even better inception than the optimized propeller in the original design space, as observed when considering the unloaded condition in the Pareto plots of the design process, which was selected exactly as representative of the blade at 270 deg position in the wake. This result is confirmed by the RANS analyses since in the sectional pressure distributions of the optimized geometries the heavy inversion of the flow (i.e., local negative angle




Fig. 26. Participation of geometrical and physical information in the variance retained by each eigenvector (“mode”).

of attack) observed for the reference propeller (already visible at the design point and further magnified for $J = 0.95$, also at very inner radial positions) is almost avoided close to the hub ($r/R = 0.5$ in Fig. 30) and significantly reduced at the tip of the blades ($r/R = 0.9$, Fig. 31).

As further proof of the quality of the optimized geometries, the comparison of the pressure fields on the suction side evidences the successful balancing of the contrasting objectives of the design, since the avoidance of the pressure side cavitation is achieved at improved inception also on the back. The unsteady pressure fields show local values far from the inception, and only PME-13-32896 and PME-13-63964 have, at 90 deg blade position, values close to the design cavitation index of 2.25. This confirms, once more, the reliability of

the design criteria, which identified these two geometries as those having the worst performances on the suction side at the design condition and at the loaded functioning, and for which the suction at the loading edge on the back was higher when computed at the lowest advance coefficient in open water by RANS. This very nice coherence among results also unveils the influence of the different geometrical variances included in the reduced design spaces. The sectional pressure distributions, indeed, are naturally more influenced by the sectional hydrofoil geometry than by pitch and chord once the correct angle of attack has been identified. The pressure distributions of PME-13-32896 and PME-13-63964 are qualitatively closer to those of the reference geometry, and show a greater tendency towards suction side/pressure side inversion at the leading edge, mitigated but similar with respect to

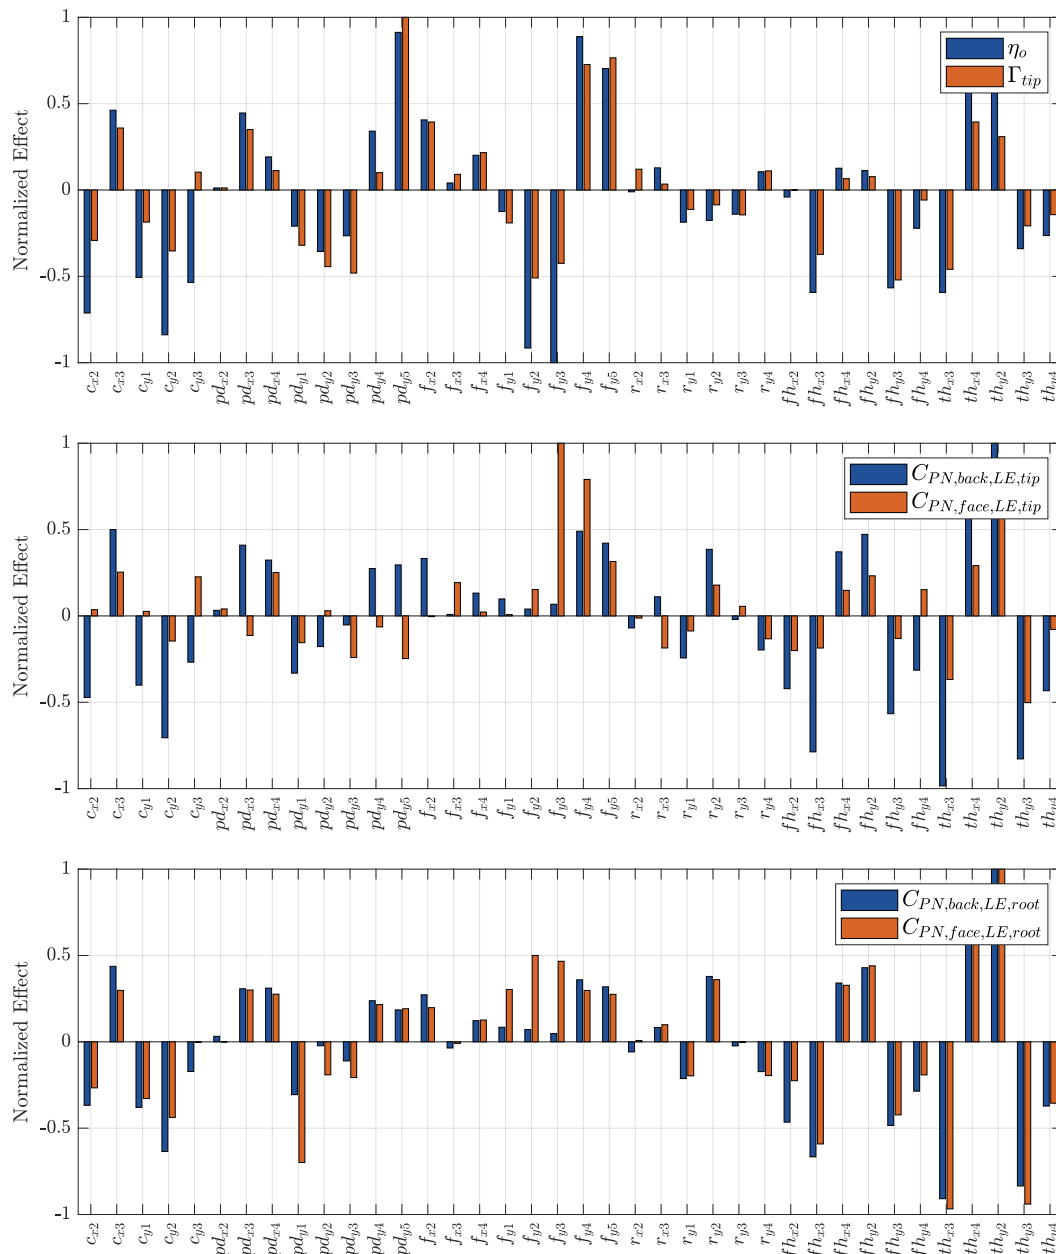


Fig. 27. Relative influence of the design parameters on several propeller key performance indicators (efficiency, tip vortex strength, cavitation inception at the design functioning, design functioning condition) through t -test.

that of the reference propeller, as a consequence of an almost identical sectional shape, not included in the reduced parameter of the PME model. Instead, the pressure distributions of PI-PME-12 and PI-PME-18 mirror the flat plateau at mid-chord observed in the optimal geometry using the non-reduced domain, thanks to the significant change in the sectional thickness distribution permitted by the physically informed reduced models.

Finally, the pressure drops at the vortex core in Fig. 32 uphold the feasibility, at a very early design stage, of a simple criterion on tip vortex cavitation through the use of data (vortex strength through circulation) from the simplest steady, non-cavitating BEM analyses. The inception of tip-vortex cavitation is greatly reduced for ID-156471 and PI-PME-12/PI-PME-18 geometries, which have comparable values of suction in the core of the vortex, coherently with the selection criteria that preferred, at nullified cavitation, those geometries with the lowest strength of the tip vortex. Both geometries from the PME reduced space, for which the avoidance of cavitation was possible by acting on fewer

geometrical parameters (pitch and maximum sectional camber rather than also on the sectional hydrofoil shape), show, instead, a limited reduction of the suction at the vortex core, confirming the results of Fig. 14.

7. Conclusions and future work

This study provides a comprehensive demonstration of the potential of parametric model embedding (PME) and, more importantly, of physics-informed parametric model embedding (PI-PME) as effective strategies for dimensionality reduction in complex marine propeller design optimization problems, typically hindered by high-dimensional design spaces and the curse of dimensionality. Compared to principal component analyses, PME-based methodologies, including those enriched with physical information, preserve the link between the reduced-order model and the original, full parametric description of the design space. This provides clear advantages in terms of interpretability




Fig. 28. Unsteady pressure distributions (C_{pN}) on the suction side under the non-uniform nominal wake at the design advance coefficient $J_A = 0.87$; comparison between the reference geometry (a) and the optimized configurations from the various design spaces.

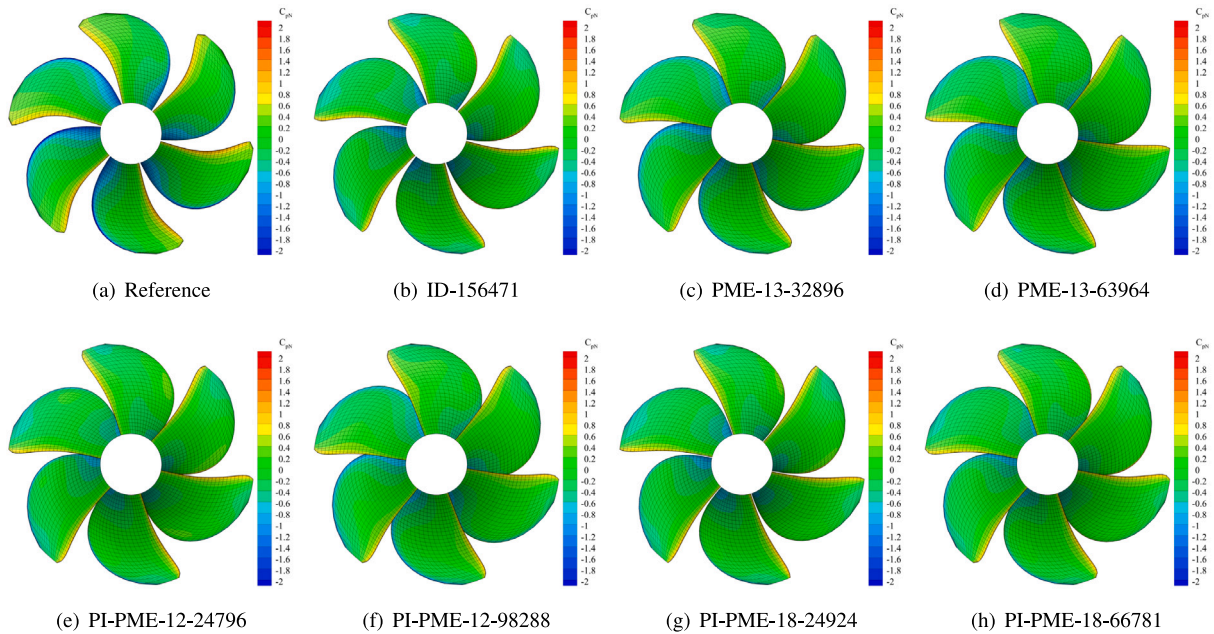


Fig. 29. Unsteady pressure distributions (C_{pN}) on the pressure side under the non-uniform nominal wake at the design advance coefficient $J_A = 0.87$; Comparison between the reference geometry (a) and the optimized configurations from the various design spaces.

of results and handling of geometries (generation and processing of results) in the optimization-based design process. The results of the optimizations and the analyses of the design spaces associated with the various reduced models, however, show some fundamental differences (and advantages) between the two methodologies. While PME ensures a rigorous compression of the design space, its purely geometrical nature may lead to the neglect of critical features that, although marginal from a geometric variance perspective, have a relevant influence on those hydrodynamic performances that are crucial for an effective design. By incorporating physical information into the reduction process, instead, PI-PME achieves a more balanced and robust representation of the design space. The method preserves not only the global shape

variability but also local geometric features, thanks to the importance that physical information attributes to them. It is the case, for instance, of the sectional camber, which gains weight in the definition of the reduced parameters through its fundamental influence on the pressure distribution. This coupling between geometry and physics yields a more reliable optimization framework, where the reduced space accurately reflects the underlying shape-performance relationships. Specifically, the comparative analysis highlights that PME requires the reconstruction of approximately 99% of the geometrical variance to yield feasible improvements, with a minimum of 13 modes to approach acceptable performance. However, even at this threshold, important shortcomings persist, particularly regarding local phenomena relevant to an

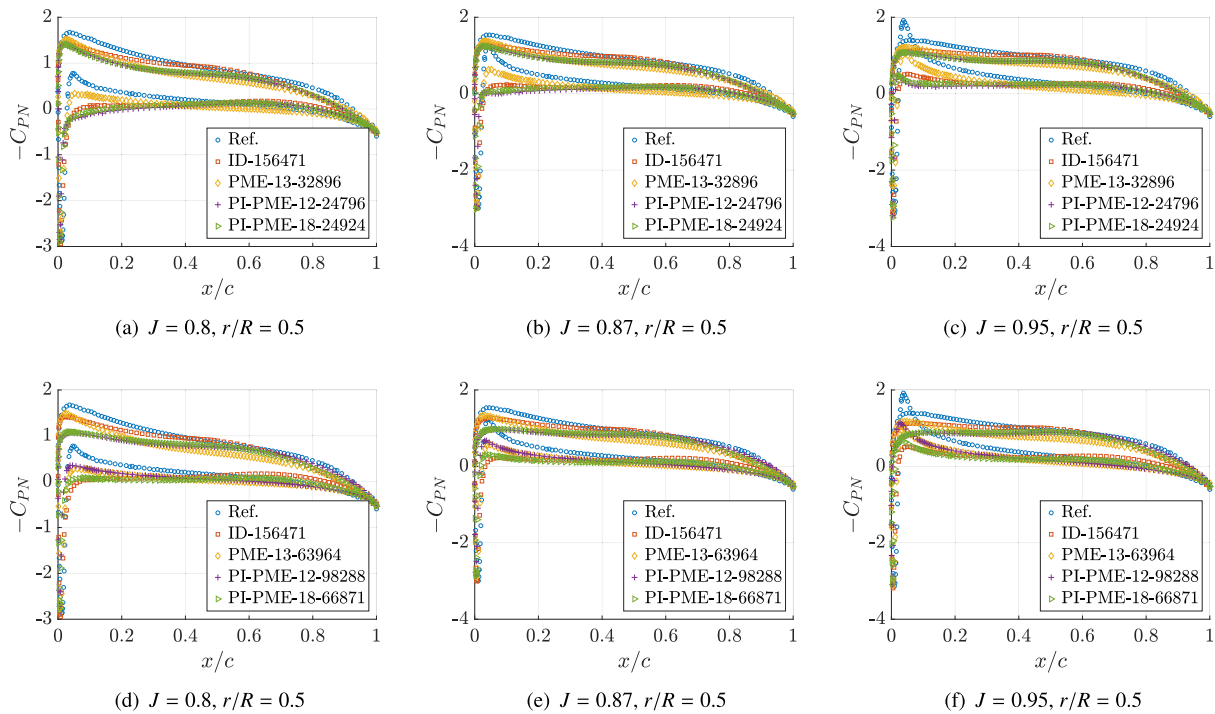


Fig. 30. Sectional pressure distributions using RANS at equivalent advance coefficients; root section comparison with the reference propeller and ID-156471 from the original design space.

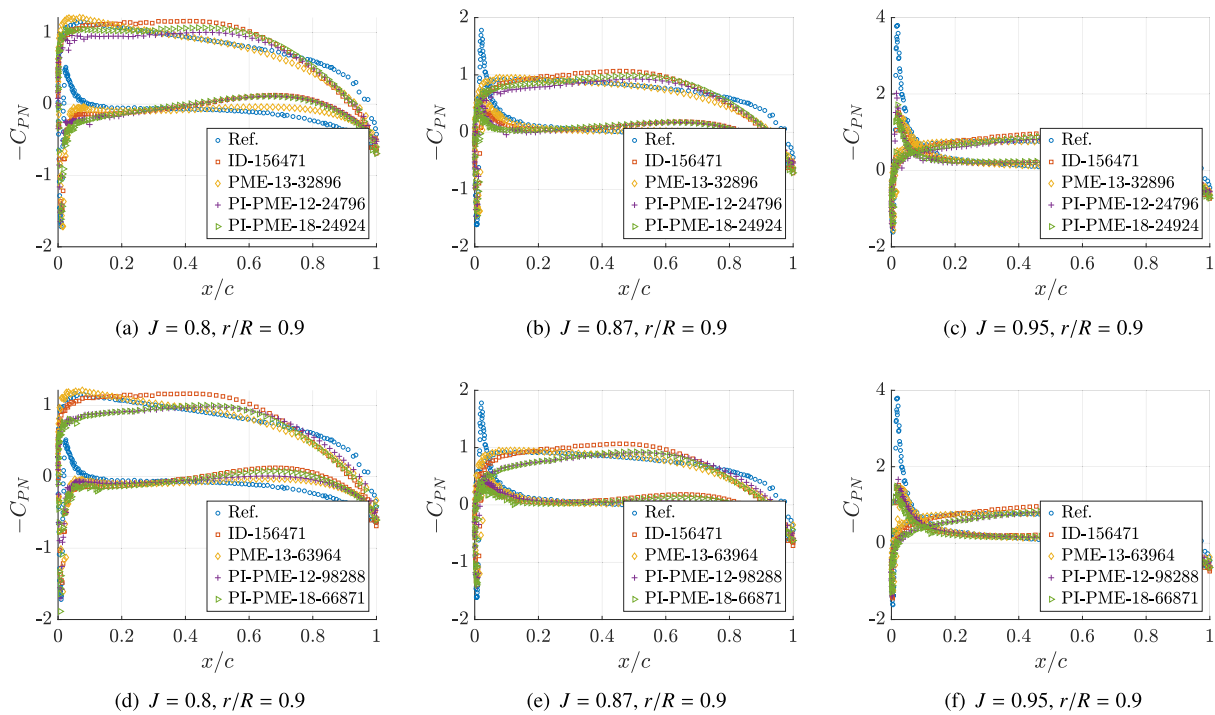


Fig. 31. Sectional pressure distributions using RANS at equivalent advance coefficients; Tip section comparison with the reference propeller and ID-156471 from the original design space.

optimal design. At 99% of the geometric variance, no information on the sectional hydrofoil shape is retained by the reduced parameters, preventing any local shape modifications capable of balancing global changes (e.g., pitch) with sectional adjustments (e.g., camber and thickness). This is the reason why the avoidance of pressure side cavitation in the case of the PME model is achieved mainly through loading of the pitch, with the obvious consequence of reduced (or not so improved)

margins against suction side cavitation and higher tip vortex strength with respect to the optimal designs in the original non-reduced space.

In contrast, PI-PME attains comparable or better optimization outcomes with fewer reduced parameters. Already with 12 modes, the explicit inclusion of physical information – in the present case the distributed pressure field representing cavitation risk – enables the reduced space to retain many local geometric features that strongly

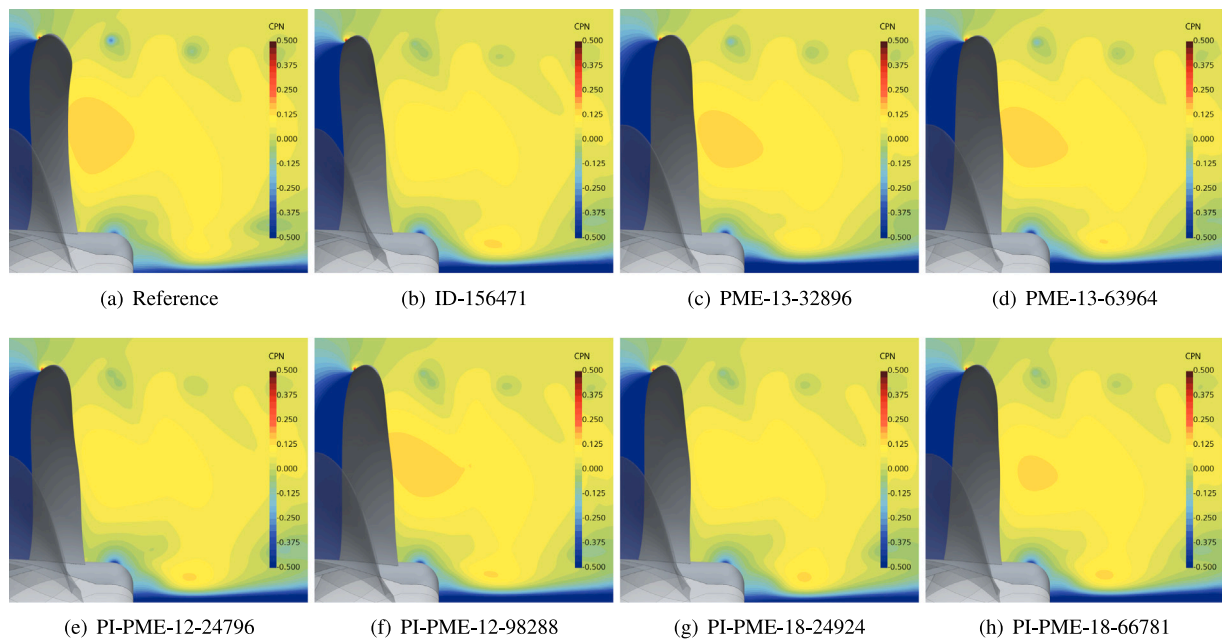


Fig. 32. Pressure drop in the tip vortex.

influence the optimization objectives and constraints. This strongly enlarges, at no additional cost compared to the almost equivalent PME-13 model, the design space and, consequently, the space of solutions, enabling the selection of better (or better balanced) geometries that, for instance, are those capable of a more effective reduction of the cavitation inception simultaneously at any of the functioning conditions included in the design, or of a further reduction of the tip vortex strength without risks in terms of pressure side cavitation. PI-PME-12-24796, for instance, achieves a 30% weaker tip vortex compared to PME-13-32896 thanks to a further pitch reduction at the blade tip. Crucially, this unloading does not increase the risk of pressure-side cavitation; cavitation levels remain comparable between the two geometries and are markedly lower than in the reference propeller, owing to the leading-edge camber-line modification captured by the PI-PME reduced parameters. Using 18 modes, the reduced-order model reproduces a design space that closely matches the original one, preserving the confidence intervals and, for most of the geometrical features considered, the absolute minimum and maximum values, while delivering solutions with performance comparable to that of the full-order space and reducing both dimensionality and convergence time by approximately 50%. At almost no additional cost compared to PI-PME-12 (identical initial training set, very similar convergence trend), this makes, for the current design problem, the PI-PME-18 model the recommended choice for a preliminary faster but reliable exploration of possible optimal configurations.

For the sake of fair comparison, indeed, all optimization processes were initialized with the same number of samples (2500) from any design space, and each initial population was evolved for 50 generations. However, the qualitative analysis of the Pareto fronts at different stages of the evolution revealed, obviously, a faster convergence rate when the number of design parameters was reduced. This outcome suggests the potential to decrease the total number of evolutionary generations required, thereby achieving significant savings in computational cost.

These results emphasize the strategic value of embedding physics into dimensionality-reduction techniques, as this ensures that the retained variations in the reduced spaces correspond to the geometrical features most relevant to the design, because they are the result of a selection process based on their actual influence on objectives and constraints. This capability is particularly critical in high-fidelity marine hydrodynamics applications, where local flow details govern efficiency,

robustness, and cavitation control. In this context, PI-PME emerges as a promising framework for accelerating design cycles while maintaining accuracy, thereby bridging the gap between reduced-order modeling and practical engineering requirements.

It may be noted that, from a computational standpoint, PI-PME introduces an additional offline phase associated with the construction of the physical snapshot database. Nevertheless, this phase is conceived to rely primarily on low-fidelity information, whose purpose is to capture dominant physical trends rather than exact high-fidelity quantities. As shown in previous investigations (Serani et al., 2025), the number of samples required for a stable physics-informed embedding is moderate, and the corresponding cost is typically negligible compared with the overall effort of the subsequent high-fidelity optimization. In the present applications, this cost is comparable to that required for the characterization of a couple of generations among the fifty considered in the optimization process and is therefore fully acceptable, given the substantial reduction in effective design dimensionality and the resulting improvement in the practical tractability of the optimization process. Furthermore, ongoing research is exploring alternative PI-PME formulations based on geometric features strongly correlated with the underlying physics, with the objective of eliminating the need for preliminary flow simulations altogether (see, e.g., Khan et al. 2022a,b). These developments are expected to further improve the cost-benefit ratio of the method and will be addressed in future studies.

Future research will further explore physics-driven PME approaches (Serani et al., 2025), where the reduced spaces are guided entirely by physical descriptors rather than geometric variance. Such an extension could further enhance the interpretability and predictive power of reduced-order models, paving the way for more adaptive, physics-centered design methodologies applicable not only to marine propellers but also to a broader class of hydrodynamic shape optimization problems. Finally, given the inherent nonlinearities of the physics-informed formulation, which PME and its variants are not able to properly capture, further research should pursue nonlinear extensions of the methodology (D'Agostino et al., 2017), such as autoencoders or manifold learning.

CRedit authorship contribution statement

Stefano Gaggero: Writing – review & editing, Writing – original draft, Visualization, Validation, Investigation, Software, Formal

analysis, Conceptualization. **Andrea Serani:** Writing – review & editing, Writing – original draft, Visualization, Investigation, Supervision, Software, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

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.

References

- Ballarin, F., D'Amario, A., Perotto, S., Rozza, G., 2019. A POD-selective inverse distance weighting method for fast parametrized shape morphing. *Internat. J. Numer. Methods Engrg.* 117 (8), 860–884. <http://dx.doi.org/10.1002/nme.5982>.
- Bellman, R.E., et al., 1957. *Dynamic Programming*. In: Cambridge Studies in Speech Science and Communication, Princeton University Press, Princeton.
- Bertetta, D., Brizzolara, 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.
- Bosschers, J., 2018. A semi-empirical prediction method for broadband hull-pressure fluctuations and underwater radiated noise by propeller tip vortex cavitation. *J. Mar. Sci. Eng.* 6 (2), 49.
- Bosschers, J., 2019. Mechanisms of low-frequency broadband noise by cavitating tip-vortices on marine propellers. In: *Sixth International Symposium on Marine Propulsors SMP*, vol. 19.
- Çelik, C., Danışman, D.B., Khan, S., Kaklis, P., 2021. A reduced order data-driven method for resistance prediction and shape optimization of hull vane. *Ocean Eng.* 235, 109406.
- 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. Mar. Sci. Eng.* 11 (1), 230. <http://dx.doi.org/10.3390/jmse11010230>.
- Chen, X., Diez, M., Kandasamy, M., Zhang, Z., Campana, E.F., Stern, F., 2015. High-fidelity global optimization of shape design by dimensionality reduction, metamodels and deterministic particle swarm. *Eng. Optim.* 47 (4), 473–494. <http://dx.doi.org/10.1080/0305215X.2014.895340>.
- D'Agostino, D., Serani, A., Campana, E.F., Diez, M., 2017. Nonlinear methods for design-space dimensionality reduction in shape optimization. In: *3rd International Conference on Machine Learning, Optimization, and Big Data. MOD 2017*, Volterra, Italy.
- 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. <http://dx.doi.org/10.1016/j.oceaneng.2019.106852>.
- Diez, M., Campana, E.F., Stern, F., 2015. Design-space dimensionality reduction in shape optimization by Karhunen-Loève expansion. *Comput. Methods Appl. Mech. Engrg.* 283, 1525–1544.
- Esteco, 2022. modeFRONTIER. Technical Report, <https://engineering.esteco.com/modefrontier/>.
- Fine, N.E., 1992. *Nonlinear Analysis of Cavitating Propellers in Nonuniform Flow* (Ph.D. thesis). Massachusetts Institute of Technology.
- Gaggero, S., 2018. Design of PBCF energy saving devices using optimization strategies: A step towards a complete viscous design approach. *Ocean Eng.* 159, 517–538.
- Gaggero, S., 2020. Numerical design of a RIM-driven thruster using a RANS-based optimization approach. *Appl. Ocean Res.* 94, 101941.
- Gaggero, S., 2022. Influence of laminar-to-turbulent transition on model scale propeller performances. Part II: cavitating conditions. *Ships Offshore Struct.* 17 (4), 772–791.
- Gaggero, S., 2025. Robust simulation-based design optimization of marine propellers. *Ocean Eng.* 321, 120397.
- Gaggero, S., Tani, G., Villa, D., Viviani, M., Ausonio, P., Travi, P., Bizzarri, G., Serra, F., 2017. Efficient and multi-objective cavitating propeller optimization: An application to a high-speed craft. *Appl. Ocean Res.* 64, 31–57.
- Gaggero, S., Vernengo, G., Villa, D., Bonfiglio, L., 2020. A reduced order approach for optimal design of efficient marine propellers. *Ships Offshore Struct.* 15 (2), 200–214.
- Gaggero, S., Villa, D., 2017. Steady cavitating propeller performance by using OpenFOAM, StarCCM+ and a boundary element method. *Proc. Inst. Mech. Eng. Part M: J. Eng. Marit. Environ.* 231 (2), 411–440.
- Gaggero, S., Villa, D., 2018. Cavitating propeller performance in inclined shaft conditions with OpenFOAM: PPTC 2015 test case. *J. Mar. Sci. Appl.* 17 (1), 1–20.
- Gaggero, S., Villa, D., Brizzolara, S., 2010. RANS and PANEL method for unsteady flow propeller analysis. *J. Hydrodyn. Ser. B* 22 (5), 564–569.
- Harries, S., Abt, C., 2019. Faster turn-around times for the design and optimization of functional surfaces. *Ocean Eng.* 193, 106470. <http://dx.doi.org/10.1016/j.oceaneng.2019.106470>.
- 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. Mar. Sci. Eng.* 9 (10), 1069.
- Kerwin, J., Leopold, R., 1964. A design theory for subcavitating propellers. *Trans. SNAME* 72, 294–335.
- Khan, S., Gaggero, S., Kaklis, P., Vernengo, G., Villa, D., 2023. Accelerating simulation-driven optimisation of marine propellers using shape-supervised dimension reduction. *arXiv preprint arXiv:2305.07876*.
- Khan, S., Kaklis, P., Serani, A., Diez, M., 2022a. Geometric moment-dependent global sensitivity analysis without simulation data: application to ship hull form optimisation. *Comput.-Aided Des.* 103339.
- Khan, S., Kaklis, P., Serani, A., Diez, M., Kostas, K., 2022b. Shape-supervised dimension reduction: Extracting geometry and physics associated features with geometric moments. *Comput.-Aided Des.* 103327.
- Kinnas, S.A., Hsin, C.Y., 1992. Boundary element method for the analysis of the unsteady flow around extreme propeller geometries. *AIAA J.* 30 (3), 688–696. <http://dx.doi.org/10.2514/3.10973>.
- Lee, H., Kinnas, S.A., 2004. Application of a boundary element method in the prediction of unsteady blade sheet and developed tip vortex cavitation on marine propellers. *J. Ship Res.* 48 (01), 15–30.
- Lidtko, A.K., Lloyd, T., Lafeber, F.H., Bosschers, J., 2022. Predicting cavitating propeller noise in off-design conditions using scale-resolving CFD simulations. *Ocean Eng.* 254, 111176.
- Masood, Z., Khan, S., Qian, L., 2021. Machine learning-based surrogate model for accelerating simulation-driven optimisation of hydropower Kaplan turbine. *Renew. Energy* 173, 827–848. <http://dx.doi.org/10.1016/j.renene.2021.04.005>.
- Serani, A., Diez, M., 2023. Parametric model embedding. *Comput. Methods Appl. Mech. Engrg.* 404, 115776.
- Serani, A., Diez, M., 2025. A survey on design-space dimensionality reduction methods for shape optimization. *Arch. Comput. Methods Eng.* 1–28.
- Serani, A., Diez, M., Quagliarella, D., 2024a. Aerodynamic shape optimization in transonic conditions through parametric model embedding. *Aerosp. Sci. Technol.* 155, 109611.
- Serani, A., Palma, G., Wackers, J., Quagliarella, D., Gaggero, S., Diez, M., 2025. Extending parametric model embedding with physical information for design-space dimensionality reduction in shape optimization. *Eng. Comput.* 1–21.
- Serani, A., Scholcz, T.P., Vanzi, V., 2024b. A scoping review on simulation-based design optimization in marine engineering: Trends, best practices, and gaps. *Arch. Comput. Methods Eng.* 31 (8), 4709.
- Serani, A., Stern, F., Campana, E.F., Diez, M., 2022. Hull-form stochastic optimization via computational-cost reduction methods. *Eng. Comput.* 38 (Suppl 3), 2245–2269.
- Siemens Digital Industries Software, 2023. STAR-CCM+, version 18.06.007. Technical Report, <https://starccm.smartcae.com/>.
- Villa, D., Gaggero, S., Coppede, A., Vernengo, G., 2020. Parametric hull shape variations by reduced order model based geometric transformation. *Ocean Eng.* 216, 107826.
- 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. <http://dx.doi.org/10.1016/j.oceaneng.2024.117604>.