Bayesian shape optimisation in high dimensional design spaces using isogeometric boundary element analysis

Shahroz Khan* and Panagiotis Kaklis †
Department of Naval Architecture, Ocean and Marine Engineering, University of Strathclyde, Glasgow, United Kingdom.

Konstantinos Kostas ‡
School of Engineering and Digital Sciences, Nazarbayev University, Nur-Sultan, Kazakhstan

Andrea Serani § and Matteo Diez ¶
CNR-INM, National Research Council-Institute of Marine Engineering, Rome, Italy

In this work, we employ dimensionality reduction and a Bayesian optimisation approach in an isogeometric analysis (IGA) setting to reduce the design space's dimensionality and ease its exploration while reducing the number of required design evaluations. In the first step, statistical dependencies implicit in the shape modification function encode essential latent features of the underlining shape while maintaining the maximum geometric variance. These latent features are used to form a low-dimensional design subspace with a correspondingly low-dimensional representation of the shape modification function. The subspace is then employed in design optimisation with a Bayesian approach. During space exploration, smooth surface representations are reconstructed from the discrete design instances of the subspace and evaluated with an IGA-enabled hydrodynamic solver. The proposed approach is demonstrated for a design optimisation of a naval ship-hull model, originally parameterised by 27 parameters, aiming at the minimisation of its wave-making resistance. The benefits of the proposed approach are contrasted with conventional optimisation procedures.

I. Introduction

A tremendous scientific effort over the last decade resulted in establishing the isogeometric analysis (IGA) [1] approach in the field of computer-aided analysis/engineering (CAA/CAE). IGA's popularity stems from its ability to employ standard CAD model representations, such as B-splines and Non-Uniform Rational B-splines (NURBS), as well as their recent extensions [2–4], in computational analysis without the need of any preprocessing step. This approach does not require any conversion of a typical CAD model into an analysis-suitable one and therefore bridges the gap between CAD and CAE technologies by circumventing the mesh generation bottleneck daunting the CAE community for many decades. To further illustrate this, we can mention here that in the typical CAE workflow, around 80% of the design time is spent on converting accurate and smooth CAD models into analysis-suitable meshes, and only 20% is used for performing the actual design analysis; see [5].

Furthermore, common CAD geometry representations, such as NURBS and B-Splines, can be efficiently employed in the construction of parametric free-form models [6], comprising one or several surface patches, at any required level of accuracy and smoothness (fairness), and via the adoption of IGA approach can maintain their appropriateness for analysis. Control points are commonly employed when performing local and/or global shape modifications of the constituent surface patches and quite naturally they become the obvious selection when considering design variables in the context of design optimisation. However, control points usage as design variables does not come without problems, and higher level parameters are preferable and generally more efficient in formulating and solving shape optimisations problems; see [7–9].

*Marie Skłodowska-Curie Research Fellow, Department of Naval Architecture, Ocean and Marine Engineering, University of Strathclyde, 16 Richmond St, Glasgow G1 1XQ, United Kingdom.

†Reader, Department of Naval Architecture, Ocean and Marine Engineering, University of Strathclyde, 16 Richmond St, Glasgow G1 1XQ, United Kingdom.

‡School of Engineering and Digital Sciences, Nazarbayev University, Qabanbay Batyr Ave 53, Nur-Sultan 010000, Kazakhstan.

§Research Scientist, CNR-INM Rome, Via di Vallerano 139, 00128.

¶Senior Research Scientist, CNR-INM Rome, Via di Vallerano 139, 00128, AIAA Member.

Therefore, NURBS-based IGA has become a favourable approach for simulation-driven shape optimisation as all involved stages, i.e., design, parameterisation, and analysis, employ the same geometry representation. Despite that, during optimisation, one still has to circumvent the problems of inappropriate/invalid designs with wiggly and/or irregular shapes, and cure the notorious curse of dimensionality. The former is often related to the increased number of control points and/or higher smoothness requirements appearing in complex free-form shapes with many intricate local features. This increased number of control points can lead to extremely high-dimensional design spaces, and shape oscillations and irregularities, which require exhaustive design exploration and numerous design constraints to attain meaningful optimum designs. Moreover, shape modification via control points can pose an additional challenge when models comprise multiple patches connected with a geometric continuity that needs to be retained; otherwise, design modifications could result in invalid or irregular geometries. A notable departure from the optimisation of control point is found in Kostas et al. [7] in which IGA-suitable procedural parameterisation (PP) is proposed. Typically, in PP, a set of surface control points defining a certain free-form feature is coupled together with a functional relation creating a versatile and high-level design parameter set with a significantly reduced size. Thus, using a relatively small set of high-level parameters, the procedural relations systematically modify the control net, which not only performs a plausible and fair surface modification but also ensures geometric continuity and automatic satisfaction of constraints. However, the construction of such parametric models requires a significant understanding of the baseline design, the employed model representation, and the correlation between model features and its performance. Even then, one can guarantee neither the identification of all relevant parameters nor exclude redundancies that increase unnecessarily the design space dimensions [10].

In this work, we employ dimensionality reduction techniques [11, 12] and a Bayesian optimisation [13] approach to cure the curse of dimensionality and ease the design space exploration by reducing the number of required design evaluations. In the first step, statistical dependencies implicit in the design parameters encode essential latent features of the underlining shape and form a lower-dimensional subspace while maintaining the maximum geometric variance of the original space. The feature extraction typically commences with design parameters used to define the baseline design. However, if the original space created by these parameters is an orthotope, then latent features are only defining a new orientation of the original design space without capturing any geometric variance. Therefore, dimensionality reduction is performed on a discretised version of the shape modification vector, which results in a lower-dimensional subspace that maintains geometric variance of the original space. Subsequently, subspace exploration, with the aim of shape optimisation, is performed via a Bayesian approach with design instances being evaluated by an Isogeometric Analysis (IGA) hydrodynamic solver [14], which guides the process towards optimum designs. Evaluating a design instance with the IGA solver requires a smooth surface representation of the design, which necessitates the solution of the inverse problem, i.e., constructing an appropriate surface representation from the lower-dimensional subspace vector instance. This is a two-step process comprising first the reconstruction of the high-dimensional discretised version of the design from the subspace, followed by a fitting process that generates a corresponding smooth surface representation. The resulting NURBS model is commonly a surface with a very large number of control points and consequently the same large number of degrees of freedom (DoFs) for the IGA-enabled analysis tool. Therefore, evaluation of such surface reconstructions, with mainly redundant control points (and DoFs), is time-consuming and defies the benefit of the enhanced accuracy with a low number of DoFs offered by IGA-enabled approaches. Simultaneously, the benefit of the performed dimensionality reduction can easily disappear and therefore, an iterative procedure is introduced for reducing the number of redundant control points before passing the design to the solver. This process uses a predefined, problem-specific geometric tolerance that guarantees the required accuracy for the application at hand while significantly reducing control points/DoFs.

II. Design Space Dimensionality Reduction Method

Let a parent design, $\mathcal{G}$, lying in a three-dimensional ambient space $\mathcal{A} \subset \mathbb{R}^3$, which is parameterised with a set of n design parameters $\mathbf{x} = \{x_1, x_2, x_3, \ldots, x_n\}$. The associated parameterisation of $\mathcal{G}$ creates a shape modification vector $\boldsymbol{\mathcal{S}}(\boldsymbol{\gamma}, \mathbf{x}) \in \mathbb{R}^3$, which for any $\mathbf{x} \in \mathcal{X}$ modifies each coordinate set $\boldsymbol{\gamma} \in \mathcal{G}$ of the parent design to a new coordinate set $\boldsymbol{\gamma}'$ of the new design $\mathcal{G}'$ as follows:

$$\boldsymbol{\gamma}' = \boldsymbol{\gamma} + \boldsymbol{\mathcal{S}}(\boldsymbol{\gamma}, \mathbf{x}). \tag{1}$$

The bounding limits of $\mathbf{x}$ define the design/search space $\mathcal{X} := \{x_i^l \leq x_i \leq x_i^u,\ i = 1, 2, \ldots, n\} \subset \mathbb{R}^n$, where $\mathbf{x}^l$ and $\mathbf{x}^u$ are the side constraints that determine the allowed variability. The optimum parametric set $\mathbf{x}^* \in \mathcal{X} \subset \mathbb{R}^n$ is found by maximising/minimising a function $g(\boldsymbol{\mathcal{S}}(\boldsymbol{\gamma}, \mathbf{x})) \in \mathbb{R}$ representing the performance criterion of interest for $\mathcal{G}$. In

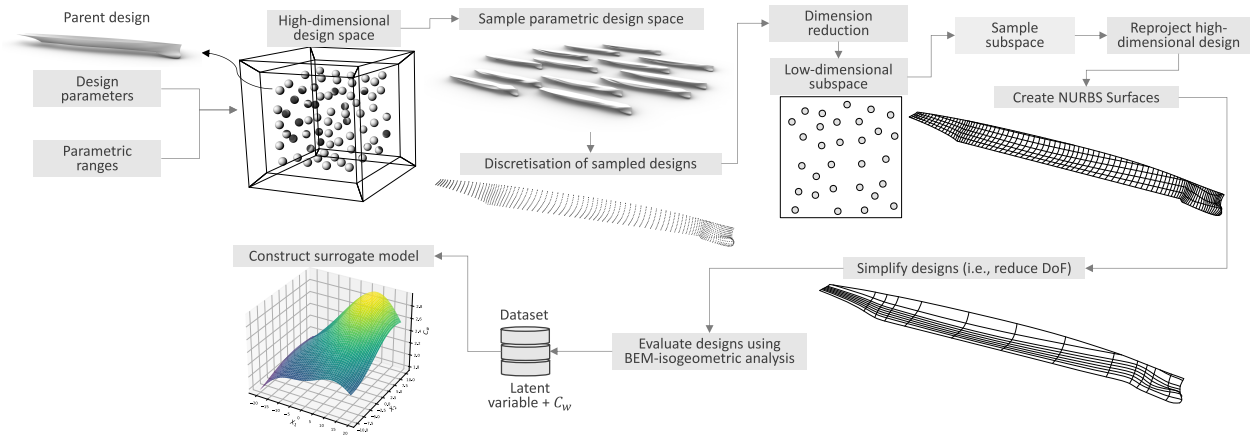

Fig. 1 Workflow of the proposed approach.

summary, and without loss of generality since any maximisation problem can be converted into a minimisation one, the following optimisation problem is solved: Find $\mathbf{x}^* \in \mathbb{R}^n$ such that

$$\mathbf{x}^* = \text{argmin}_{x \in \mathcal{X}}\, g\left(\boldsymbol{\mathcal{S}}(\boldsymbol{\gamma}, \mathbf{x})\right) \tag{2}$$

The present work aims to extract geometrically-active latent features for the formulation of a dimensionally-reduced representation of $\boldsymbol{\mathcal{S}}(\boldsymbol{\gamma}, \mathbf{x})$, namely $\hat{\boldsymbol{\mathcal{S}}}(\boldsymbol{\gamma}, \mathbf{u})$, where $\mathbf{u} = \{u_1, u_2, \ldots, u_m\} \in \mathcal{U} \subset \mathbb{R}^m$, with $m << n$, is a *Geometrically-Active Latent Variable* (GALV) vector. This vector pertains to the geometrically-active features of $\mathcal{X}$ and serves as the new dimensional coordinate set for the m-dimensional subspace GAS($\mathcal{U}$). Therefore, the initial optimisation problem (2) can be replaced by a computationally more efficient optimisation in $\mathcal{U}$ that yields the optimum vector $\mathbf{u}^* \in \mathcal{U}$, which ultimately, via $\hat{\boldsymbol{\mathcal{S}}}(\boldsymbol{\gamma}, \mathbf{u})$, results in the optimal design $\mathcal{G}^*$.

A. Extracting geometrically active features

Construction of $\hat{\boldsymbol{\mathcal{S}}}(\boldsymbol{\gamma}, \mathbf{u})$ commences with the eigendecomposition of a symmetric and positive definite covariance matrix, $\mathbf{C}$. For GAS, this matrix is constructed numerically by discretising $\mathcal{G}$ into E elements of equal measure $\Delta\mathcal{G}$, which in 3D ambient space yields the spatial discretisation of $\boldsymbol{\mathcal{S}}(\boldsymbol{\gamma}, \mathbf{x})$ as:

$$\boldsymbol{\mathcal{S}}(\boldsymbol{\gamma}, \mathbf{x}) = \mathbf{d}(\mathbf{x}) = \begin{bmatrix} \mathbf{d}_1(\mathbf{x}) \\ \mathbf{d}_2(\mathbf{x}) \\ \mathbf{d}_3(\mathbf{x}) \end{bmatrix} = \begin{bmatrix} d_{1,1}(\mathbf{x}) \\ \vdots \\ d_{1,E}(\mathbf{x}) \\ d_{2,1}(\mathbf{x}) \\ \vdots \\ d_{3,E}(\mathbf{x}) \end{bmatrix}, \tag{3}$$

where $\mathbf{d}(\mathbf{x})$ is a column matrix of size $[L \times 1]$ and $L = 3E$. Using this discretisation, $\mathbf{C}$ can be obtained as:

$$\mathbf{C} = \int_{\mathcal{X}} \mathbf{d}(\mathbf{x})\mathbf{d}(\mathbf{x})^T \rho(\mathbf{x}) d\mathbf{x}. \tag{4}$$

For the complex problem, evaluation of $\mathbf{C}$ requires solving high order integrals, which, if the dimensionality of the design space is sufficiently small, can be solved with techniques like tensor product Gauss-Legendre quadrature. Therefore, in the present work, it is approximated while sampling $\mathcal{X}$ with a statistically convergent number of Monte Carlo realisations, N, which produces the data-set $\mathbf{D}$ as

$$\mathbf{D} = \begin{bmatrix} \mathbf{d}(\mathbf{x_1}) \\ \mathbf{d}(\mathbf{x_2}) \\ \vdots \\ \mathbf{d}(\mathbf{x_N}) \end{bmatrix} = \begin{bmatrix} d_{1,1}(\mathbf{x}_1) & \dots & d_{1,E}(\mathbf{x}_1) & d_{2,1}(\mathbf{x}_1) & \dots & d_{2,E}(\mathbf{x}_1) & d_{3,1}(\mathbf{x}_1) & \dots & d_{3,E}(\mathbf{x}_1) \\ d_{1,1}(\mathbf{x}_2) & \dots & d_{1,E}(\mathbf{x}_2) & d_{2,1}(\mathbf{x}_2) & \dots & d_{2,E}(\mathbf{x}_2) & d_{3,1}(\mathbf{x}_2) & \dots & d_{3,E}(\mathbf{x}_2) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ d_{1,1}(\mathbf{x}_N) & \dots & d_{1,E}(\mathbf{x}_N) & d_{2,1}(\mathbf{x}_N) & \dots & d_{2,E}(\mathbf{x}_N) & d_{3,1}(\mathbf{x}_N) & \dots & d_{3,E}(\mathbf{x}_N) \end{bmatrix}. \tag{5}$$

At the discrete level, $\mathbf{C}$ can be approximated by $\hat{\mathbf{C}}$:

$$\mathbf{C} \approx \hat{\mathbf{C}} = \frac{1}{N} \sum_{r=1}^{N} \mathbf{d}(\mathbf{x}_r)\mathbf{d}(\mathbf{x}_r)^T, \tag{6}$$

whose eigendecomposition is obtained by the solution of the following eigenproblem:

$$\hat{\mathbf{C}}\mathbf{W} = \Lambda\mathbf{W}. \tag{7}$$

Herein, $\mathbf{W}$ is the $[L \times L]$ feature matrix whose columns are orthogonal eigenvectors $\mathbf{W} = \{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_L\}$ with $(\mathbf{w}_k^T\mathbf{w}_k = 1)$. This spans the basis of an eigenspace, which creates rotations of $\mathbb{R}^L$. $\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_L)$ are the eigenvalues sorted in descending order $\lambda_1 \geq \lambda_2 \geq, \dots, \lambda_L \geq 0$ and represents the geometric variance along the corresponding eigenvectors. The eigenvectors have the properties of maximising the geometric variance of designs projected on them and minimising the mean squared distance between the original points and the relative projections. Equation 7 can be solved via different techniques; however, the most commonly used one is Singular Value Decomposition (SVD), which is also employed work; see more details on SVD in [15]. The eigenvector with the highest eigenvalue captures the maximum geometric variance of designs in $\mathcal{X}$. Therefore, for dimension reduction, columns of $\mathbf{W}$ are rearranged in descending order, based on their eigenvalues, and are partitioned into two sets containing dominant and non-dominant features.

$$\mathbf{\Lambda} = \begin{bmatrix} \mathbf{\Lambda}_1 & \\ & \mathbf{\Lambda}_2 \end{bmatrix}, \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_1 & \mathbf{W}_2 \end{bmatrix}. \tag{8}$$

The columns of $\mathbf{W}_1 = \{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_m\}$ and $\mathbf{W}_2 = \{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{L-m}\}$ are the dominant and non-dominant eigenvectors, respectively, while $\mathbf{\Lambda}_1 = \{\lambda_1, \lambda_2, \dots \lambda_m\}$ and $\mathbf{\Lambda}_2 = \{\lambda_1, \lambda_2, \dots \lambda_{L-m}\}$ are their corresponding eigenvalues. The separation between $\mathbf{W}_1$ and $\mathbf{W}_2$ is done based on the eigenvalues and hence, $\mathbf{W}_1$ contains the first m columns of $\mathbf{W}$, corresponding to the first largest m eigenvalues. These m eigenvalues constitute $\mathbf{\Lambda}_1$ and we pick m so that $\mathbf{W}_1$ retains a minimum of 95% of the geometric variance (σ^2):

$$\sigma^2 = \frac{\sum_{i=1}^{m} \lambda_i}{\sum_{i=1}^{L} \lambda_i}. \tag{9}$$

Among the geometrically active $\mathbf{u}^1 = \{u_i^1, i = 1, 2, \dots, m\} = \mathbf{W}_1^T\mathbf{d}(\mathbf{x})$ and the inactive $\mathbf{u}^2 = \{u_1^2, u_2^2, \dots, u_{L-m}^2\} = \mathbf{W}_2^T\mathbf{d}(\mathbf{x})$ latent variables, we are only interested in $\mathbf{u}^1$ as their basis $\mathbf{W}_1$ covers the largest geometric variability and designs are only negligibly influenced by $\mathbf{u}^2$. Finally, the reduced-order representation of $\boldsymbol{\mathcal{S}}(\boldsymbol{\gamma}, \mathbf{x})$, i.e., $\hat{\boldsymbol{\mathcal{S}}}(\boldsymbol{\gamma}, \mathbf{u})$ in discrete form can be calculated via $\hat{\mathbf{d}}(\mathbf{u}) = \mathbf{W}_1\mathbf{u} \in \mathbb{R}^L$, where $\hat{\mathbf{d}}(\mathbf{t})$ can be also considered as the minimum-squared-error approximation of $\mathbf{d}(\mathbf{x})$, i.e.,

$$MSE = \frac{1}{N} \sum_{r=1}^{N} \left\| \hat{\mathbf{d}}(\mathbf{u}_r) - \mathbf{d}(\mathbf{x}_r) \right\|^2. \tag{10}$$

B. Design surface reconstruction and simplification

As described above, the GALV, i.e., $\mathbf{u}$, corresponds to a lower-dimensional parameterisation of a discretised design. However, in the context of isogeometric analysis, a smooth representation of the design is needed; therefore, during optimisation in $\mathcal{U}$, the corresponding NURBS surface for each design is reconstructed to evaluate its performance. Surface fitting to a point grid/cloud is a well-studied topic with applications in reverse engineering, in CAD, and Computer Graphics communities. However, for complex geometries and performance indices that are highly sensitive

to shape modifications, the common surface fitting approaches may result in undesired surface irregularities that greatly affect the analysis process. This is especially true for ship hulls and wave resistance estimations that can be greatly affected by even small and local shape irregularities. Therefore, it is of uttermost importance to accurately reconstruct the continuous ship hull model with a sufficiently fair and smooth (G^2 continuity is commonly required) surface representation.

In the present work, we exploit the structure and geometric information we have embedded in our grid during the discretisation phase and firstly construct a set of appropriately spaced section curves that are then interpolated via a surface skinning scheme to generate a smooth surface representation of the ship hull; see Figure 2. This process stems from the traditional ship-lines design approach and forms the basis of many modern design approaches used in naval architecture & engineering. The resulting surface model is evaluated with an in-house developed IGABEM solver.

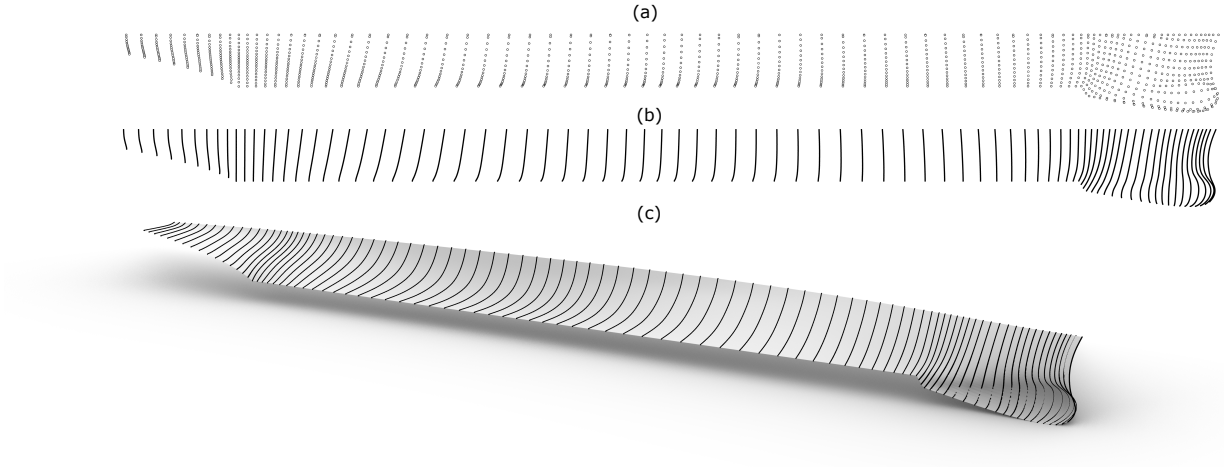

Fig. 2 Surface generation from the design grid used in PCA. (a) Design grid, (b) section curves and (c) NURBS surface representation of the design.

III. BEM-Isogeometric Analysis

The isogeometric solver adopted in the optimisation for the evaluation of the wave resistance coefficient, C_w, is based on a linear approximation which avoids the treatment of the non-linear condition on the free surface, i.e., the so-called Neumann-Kelvin problem. Following the formulation in Baar and Price [16], the wave-resistance problem is formulated as a Boundary Integral Equation (BIE) defined on the ship's wetted surface S and the corresponding waterline l, as

$$\frac{\mu(\mathbf{P})}{2} - \int_S \mu(\mathbf{Q}) \frac{\partial G(\mathbf{P},\mathbf{Q})}{\partial \mathbf{n}(\mathbf{P})} \mathrm{d}S(\mathbf{Q}) - \frac{1}{k} \int_l \mu(\mathbf{Q}) \frac{\partial G(\mathbf{P},\mathbf{Q})}{\partial \mathbf{n}(\mathbf{P})} n_x(\mathbf{Q}) \tau_y(\mathbf{Q}) \mathrm{d}l(\mathbf{Q}) = -\mathbf{U} \cdot \mathbf{n}(\mathbf{P}) \tag{11}$$

where μ is the density of the source-sink distribution on S, $G(\mathbf{P},\mathbf{Q})$ the Neumann-Kelvin Green's function, $\mathbf{U}$ denotes the steady forward speed of the ship, and $k = g/||U||^2$ is the characteristic wave number. From the solution of the above integral equation, several quantities of interest, such as velocity/pressure distributions, wave profile and ship wave resistance can be obtained.

In the present work, a high-order Boundary Element Method (BEM), based on the isogeometric analysis approach, is applied for the numerical solution of the BIE in Eq. (11); see [14] for relevant details. In this setting, the field quantities are represented by the very same basis that is being used for representing the geometry of the body boundary, i.e., the submerged surface of the ship hull. For the latter, we assume that it can be accurately represented as a multi-patch NURBS surface $\mathbf{x}_p(t_1,t_2)$, as follows:

$$\mathbf{x}_p(t_1,t_2) = \sum_{i_1=0}^{n_1^p} \sum_{i_2=0}^{n_2^p} \mathbf{d}^p_{i_1 i_2} R^p_{i_1 i_2, k_1^p k_2^p}(t_1,t_2) := \sum_{\mathbf{i}=0}^{\mathbf{n^p}} \mathbf{d}^p_{\mathbf{i}} R^p_{\mathbf{i},\mathbf{k}^p}(t_1,t_2), \ (t_1,t_2) \in I_1^p \times I_2^p, \tag{12}$$

where $\mathbf{d}_{\mathbf{i}}^{p}$ are the control points, $R_{\mathbf{i},\mathbf{k}^p}^{p}(t_1,t_2)$ are the rational basis functions:

$$R_{\mathbf{i},\mathbf{k}^p}^{p}(t_1,t_2) := R_{i_1 i_2,k_1^p k_2^p}^{p}(t_1,t_2) = \frac{w_{i_1 i_2}^{p} N_{i_1,k_1^p}^{p}(t_1) N_{i_2,k_2^p}^{p}(t_2)}{\sum_{l_1=0}^{n_1^p}\sum_{l_2=0}^{n_2^p} w_{l_1 l_2}^{p} N_{l_1,k_1^p}^{p}(t_1) N_{l_2,k_2^p}^{p}(t_2)}, p = 1,2,\ldots,N, \tag{13}$$

with $N_{i,k}(t)$ being the i-th k order B-Spline basis function and p the patch index. Furthermore, we consider that each interval $I_{\bullet}^{p}$ is partitioned appropriately by the corresponding knot vector $J_{\bullet}^{p}$. As already mentioned above, in IGA context, we employ the same representation for the unknown source-sink density distribution, i.e.,

$$\mu(t_1,t_2) = \sum_{\mathbf{i}=0}^{\mathbf{n}^p+\boldsymbol{\ell}^p} \mu_{\mathbf{i}}^{p} R_{\mathbf{i},\mathbf{l}^p}^{p}(t_1,t_2),\ (t_1,t_2) \in I_1^p \times I_2^p,\ p = 1,2,\ldots,N. \tag{14}$$

where $\boldsymbol{\ell}^p$ are the additional knots that can be inserted to refine the original spline space for improving our field quantity approximations. representation, sequence of nested finite-dimensional spaces. Substituting Eq. (14) into Eq. (11) and using an appropriate set of collocation points, Greville points in our case, we arrive at a discrete form of Eq. (11) producing a linear system that is solved for the unknown coefficients $\mu_{\mathbf{i}}^{p}$.

IV. Bayesian Optimisation

There are many optimisation problems in which the objective function may be considered a black-box function that does not assume any functional form and which might be computationally expensive to evaluate. Therefore, evaluation of the objective function, $g\left(\hat{\boldsymbol{\mathcal{S}}}(\boldsymbol{\gamma},\mathbf{u})\right)$ in our case, may be restricted to few sample instances in the design space resulting in inadequate response sampling. If g was computationally inexpensive, one could sample many designs, e.g. via grid search, random search etc, however, if it is expensive, then it is imperative to find an optimal solution with the least number of design evaluations with g. In Bayesian strategy we treat the objective function as a random function and place a prior over it. The prior captures beliefs about the behaviour of the function. After gathering function evaluations, which are treated as data, the prior is updated to form the posterior distribution over the objective function. Bayesian optimisation is therefore a useful tool for computationally expensive black-box functions, as in our case. There are three main steps: initialisation, objective function evaluations and Bayesian update. In the initialisation step, the prior belief about g in the form of a surrogate model is constructed using a set of pseudo-random samples evaluated directly with the high-fidelity solver. In the Bayesian update, optimisation uses acquisition functions to update the prior with a *new sample* to get a posterior that better approximates g. The position of the new sample is determined with an acquisition function, e.g., expected improvement, which is a trade-off between exploitation and exploration.

V. Experimentation: Shape reparameterisation and optimisation

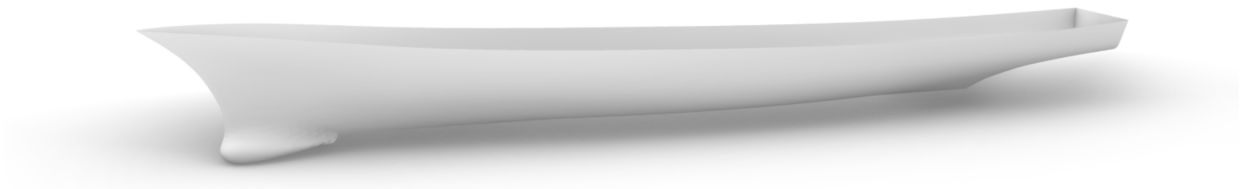

Fig. 3 Ship hull geometry of DTMB 5415 naval ship model used as a test case for the proposed approach.

The DTMB 5415 hull model (see Fig 3) is a widely used benchmark ship model employed in the shape optimisation research community and broader. In the present work, this hull model is considered for the minimisation of its wave-making resistance coefficient, C_w, at a Froude number equal to 0.350. The wave-making resistance constitutes a significant part of the ship's total resistance, when considering high Froude numbers, which are common for naval and large container ships. It corresponds to the energy consumed to generate the free-surface waves appearing when a body moves on or near the free surface of oceans, rivers or lakes. Moreover, it is sensitive to local features of the hull form, such as the bulbous bow or sonar dome; thus, a significant reduction can be achieved without significantly

modifying the ship's overall dimensions or capacity, which are often critical design constraints. Although minimising wave-making resistance can be extremely important for several ship types, its evaluation is computationally intensive. Table 1 summarises the main characteristics of the hull and test conditions.

Table 1 DTMB 5415 original (model scale) hull main particulars.

Quantity	Symbol	Unit	Value
Displacement	∇	m^3	0.549
Length between perpendiculars	L_{pp}	m	5.720
Beam	B	m	0.760
Draft	T	m	0.248
Longitudinal center of gravity	LCG	m	2.884
Vertical center of gravity	VCG	m	0.056
Water density	ρ	kg/m^3	998.5
Kinematic viscosity	ν	m^2/s	1.09E-6
Gravity acceleration	g	m/s^2	9.803
Froude number	Fr	–	0.350
Wave resistance coefficient	C_w	–	9.1476E-04

A. Shape Modification Method

Shape modification is determined by a recursive combination of $M = 27$ shape modification vectors over a hyper-rectangle embedding the semi-hull:

$$\boldsymbol{\psi}_i(\boldsymbol{\zeta}) \;:\; \mathcal{A} = [0, L_{\zeta_1}] \times [0, L_{\zeta_2}] \times [0, L_{\zeta_3}] \in \mathbb{R}^3 \longrightarrow \mathbb{R}^3, \tag{15}$$

with $i = 1, \ldots, M$. Specifically, $\boldsymbol{\delta}(\boldsymbol{\zeta}, \mathbf{x}) = \boldsymbol{\delta}_M$, where

$$\boldsymbol{\delta}_i(\boldsymbol{\zeta}, \mathbf{x}) = x_i \boldsymbol{\psi}_i(\boldsymbol{\zeta}), \quad \text{with} \quad \begin{cases} \boldsymbol{\zeta} = \boldsymbol{\zeta} + \boldsymbol{\delta}_{i-1} \\ \boldsymbol{\delta}_1 = 0 \end{cases} \tag{16}$$

The coefficients $\{x_i, i = 1, \ldots, M \in \mathbb{R}\}$ are the design parameters and form a 27−dimensional initial/original design space $\mathcal{X}$. For modification, the shape functions are defined as:

$$\boldsymbol{\psi}_i(\boldsymbol{\zeta}) := \prod_{j=1}^{3} \sin\left(\frac{a_{ij}\pi\zeta_j}{L_{\zeta_j}} + r_{ij}\right) \mathbf{e}_{q(i)}. \tag{17}$$

In Eq. (17), $\{a_{ij}, j = 1, 2, 3\} \in \mathbb{R}$ define the order of the function along j-th axis, $\{r_{ij}, j = 1, 2, 3\} \in \mathbb{R}$ are the corresponding spatial phases, $\{L_{\zeta_j}, j = 1, 2, 3\} \in \mathbb{R}$ are the hyper-rectangle edge lengths, and $\mathbf{e}_{q(i)}$ is the corresponding unit vector. Modifications are applied along ζ_1, ζ_2, and ζ_3, with $q(i) = 1, 2,$ or 3 respectively; see additional details in [**?**].

B. Shape reparameterisation with GALV

Extraction of geometric features for the GALV vector $\mathbf{u} \in \mathcal{U}$ that constructs $\hat{\boldsymbol{\mathcal{S}}}(\boldsymbol{\gamma}, \mathbf{u})$ commences with the eigendecomposition of $\mathbf{C}$, which accumulates the geometric variability of designs into eigenvalues and eigenvectors. The former captures the original geometric variability of the full design space, whereas the latter determines the order of importance of the shape deformation modes. Here, $\mathbf{u} \in \mathbb{R}^m$ is set to capture 95% variability of the original design space, i.e., $\sigma^2 = 95\%$

As mentioned before, the eigendecomposition is performed on a discrete version of $\mathbf{C}$. The process begins by randomly sampling $N = 9000$ hull-form designs from the original design space $\mathcal{X}$. The sampling follows a uniform random

distribution performed by Monte Carlo sampling. Afterwards, the discretisation of the sampled designs is generated using $E = [90 \times 25]$ elements that correspond to $L = 3 \times 90 \times 25$ coordinate values. This creates a training data-set $\mathbf{D} \in \mathbb{R}^{6750\times9000}$, which is used in Eq. (6). Figure 4 depicts (a) the accumulative percentage of geometric variance retained with each added eigenvector and (b) MSE calculated in Eq. (10).

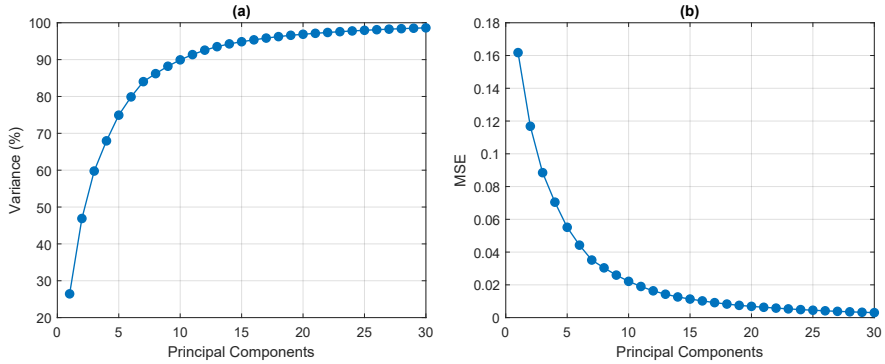

Fig. 4 Graphs showing (a) the percentage of geometric variance retained and (b) the convergence of MSE as a function of the number of dimensions (m) considered for the lower-dimensional design space.

One may observe in Fig. 4 that the first 15 eigenvectors, $\mathbf{W}_1 = \{\mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_{15}\}$ obtained from the eigendecomposition of $\mathbf{C}$ retain 95.18% of the geometric variance. This result indicate that $m = 15$ is sufficient to capture more than 95% of the geometric variability, which can be achieved with the original $n = 27$ parameters. In other words, there are only 15 latent directions that capture the great majority of the geometric variation in $\mathcal{X}$, which results in an approximately 44% reduction of the original design space's dimensionality. Figures 5 depicts the first three eigenvectors, $\{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\}$, on the hull surface, with blue/yellow areas corresponding to small/large surface deformations, respectively. Visualising the eigenmodes directly on the design surface offers an insight of great practical value as it demonstrates the corresponding type and order of variance captured in each case.

Fig. 5 Shape deformation of the hull model corresponding to the first three eigenvectors.

C. Bayesian shape optimisation

As mentioned before, ship-hull shape optimisation targets the minimisation of the wave-resistance coefficient, C_w. The achieved values at each iteration of the optimisation process are shown in Fig. 6. As explained earlier, Bayesian optimisation is initiated by constructing a surrogate model as a prior belief about g, i.e., C_w in our case. Therefore, at the zeroth iteration, 200 designs were sampled using the technique proposed in [17], which ensures uniformly distributed and diverse samples. During optimisation, the acquisition function is minimised instead of C_w, which drives the exploration toward global optima. Therefore, C_w values of design instances over optimisation iteration may not always decrease due to the exploratory part of the optimisation. We can easily see in Fig. 6 that the objective value reduced rapidly in the first few iterations and the current best was updated only after iteration 29. After iteration 29, the reduction slows down and reaches the estimated optimum (see Fig. 7) at iteration 73. Relatively few design evaluations were used to reach the optimised design with significant improvement in C_w, from 9.1476E-04 to 2.0670E-04.

VI. Conclusions

This work utilised a free-surface IGA-BEM for the shape optimisation of the ship hull. To reduce the overall computational cost, we first reduce the dimensionality of the design space to extract the latent features. These features are then used to create a subspace whose dimensional is approximately 44% lower than the original design space, but it still retains the 95.18% of the geometric variance present in the original design space. Afterwards, this subspace is explored using Bayesian optimisation with the aim of finding the optimal design (i.e., a design with minimal wave resistance coefficient) with the least number of design evaluations. The design obtained at the end of the optimisation showed approximately 77% reduction in the wave resistance coefficient compared to the original design space.

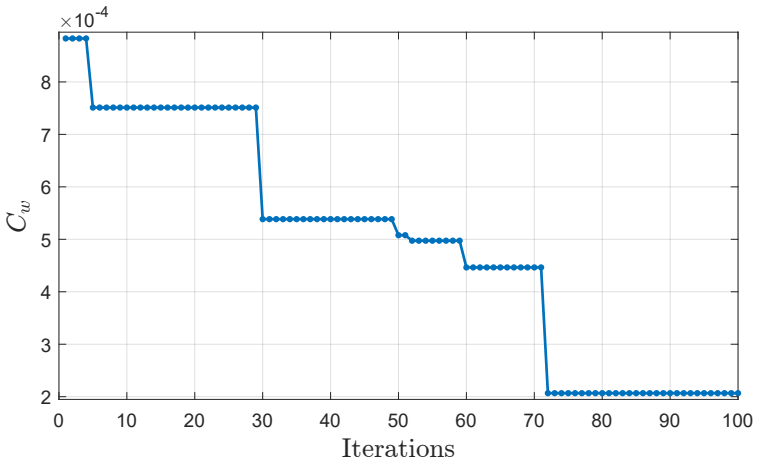

Fig. 6 Optimisation history depicting the current best (C_w) updates vs optimisation iterations.

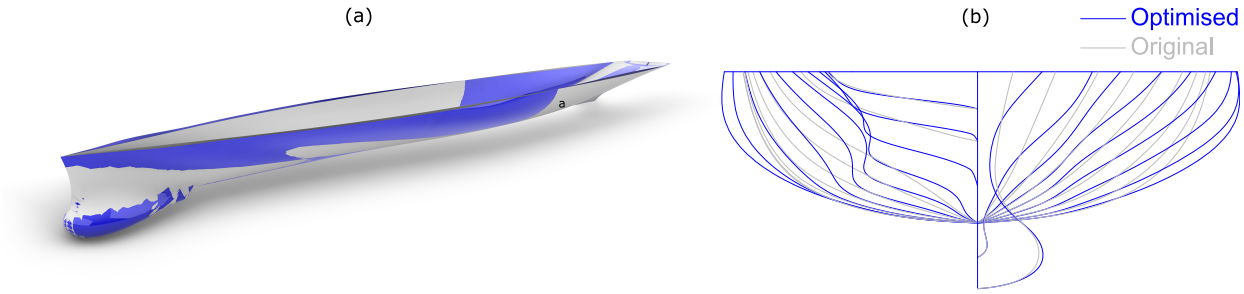

Fig. 7 Comparison between the original and optimised hull shapes.

Acknowledgements

The first two authors are thankful for the support and funding received from the EU Horizon-2020 Research and Innovation Programme under the Marie Skłodowska–Curie grant agreement No. 860843 – GRAPES: *Learning, Processing and Optimising Shapes*. The last author is supported by The Nazarbayev University FDCRGP 2022–24 funded project: SOFFA - PHYS: *Shape Optimisation of Free-form Functional surfaces using isogeometric Analysis and Physics-Informed Surrogate Models*, Grant No. 11022021FD2927. The fourth and fifth authors are grateful to the US Office of Naval Research Global for its support through grants N62909-11-1-7011 and N62909-21-1-2042.

References

[1] Hughes, T. J., Cottrell, J. A., and Bazilevs, Y., "Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement," *Computer methods in applied mechanics and engineering*, Vol. 194, No. 39-41, 2005, pp. 4135–4195.

[2] Sederberg, T. W., Cardon, D. L., Finnigan, G. T., North, N. S., Zheng, J., and Lyche, T., "T-spline simplification and local refinement," *ACM transactions on graphics (TOG)*, Vol. 23, No. 3, 2004, pp. 276–283.

[3] Bressan, A., "Some properties of LR-splines," *Computer Aided Geometric Design*, Vol. 30, No. 8, 2013, pp. 778–794.

[4] Giannelli, C., Jüttler, B., and Speleers, H., "THB-splines: The truncated basis for hierarchical splines," *Computer Aided Geometric Design*, Vol. 29, No. 7, 2012, pp. 485–498.

[5] Herrema, A. J., Wiese, N. M., Darling, C. N., Ganapathysubramanian, B., Krishnamurthy, A., and Hsu, M.-C., "A framework for parametric design optimization using isogeometric analysis," *Computer Methods in Applied Mechanics and Engineering*, Vol. 316, 2017, pp. 944–965.

[6] Khan, S., Gunpinar, E., Mert Dogan, K., Sener, B., and Kaklis, P., "ModiYacht: intelligent CAD tool for parametric, generative, attributive and interactive modelling of yacht hull forms," *SNAME 14th International Marine Design Conference*, OnePetro, 2022.

[7] Kostas, K., Ginnis, A., Politis, C., and Kaklis, P., "Ship-hull shape optimization with a T-spline based BEM–isogeometric solver," *Computer Methods in Applied Mechanics and Engineering*, Vol. 284, 2015, pp. 611–622.

[8] Kostas, K., Amiralin, A., Sagimbayev, S., Massalov, T., Kalel, Y., and Politis, C., "Parametric model for the reconstruction and representation of hydrofoils and airfoils," *Ocean Engineering*, Vol. 199, 2020, p. 107020.

[9] Khan, S., Kaklis, P., Serani, A., and Diez, M., "Geometric moment-dependent global sensitivity analysis without simulation data: application to ship hull form optimisation," *Computer-Aided Design*, 2022, p. 103339.

[10] Khan, S., and Kaklis, P., "From regional sensitivity to intra-sensitivity for parametric analysis of free-form shapes: Application to ship design," *Advanced Engineering Informatics*, Vol. 49, 2021, p. 101314.

[11] Khan, S., Serani, A., Diez, M., and Kaklis, P., "Physics-informed feature-to-feature learning for design-space dimensionality reduction in shape optimisation," *AIAA Scitech 2021 Forum*, 2021, p. 1235.

[12] Khan, S., Kaklis, P., Serani, A., Diez, M., and Kostas, K., "Shape-supervised Dimension Reduction: Extracting Geometry and Physics Associated Features with Geometric Moments," *Computer-Aided Design*, 2022, p. 103327.

[13] Mockus, J., "Application of Bayesian approach to numerical methods of global and stochastic optimization," *Journal of Global Optimization*, Vol. 4, No. 4, 1994, pp. 347–365.

[14] Belibassakis, K., Gerostathis, T. P., Kostas, K., Politis, C., Kaklis, P., Ginnis, A., and Feurer, C., "A BEM-isogeometric method for the ship wave-resistance problem," *Ocean Engineering*, Vol. 60, 2013, pp. 53–67.

[15] Abdi, H., and Williams, L. J., "Principal component analysis," *Wiley interdisciplinary reviews: computational statistics*, Vol. 2, No. 4, 2010, pp. 433–459.

[16] Baar, J., and Price, W. G., "Developments in the calculation of the wavemaking resistance of ships," *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, Vol. 416, No. 1850, 1988, pp. 115–147.

[17] Khan, S., and Gunpinar, E., "Sampling CAD models via an extended teaching–learning-based optimization technique," *Computer-Aided Design*, Vol. 100, 2018, pp. 52–67.