

# Efficient Integration of Machine Learning and Optimization Techniques into Passenger Ships Structural Design

Ciro BUSIELLO <sup>a,1</sup>, Mauro SICCHIERO <sup>a</sup>, Matteo SIDARI <sup>a</sup>, Lorenzo FABRIS <sup>b</sup> and Gianluigi ROZZA <sup>b</sup>.

<sup>a</sup>*Fincantieri S.p.A., Merchant Ships Business Unit, Fincantieri S.p.A., Trieste, Italy*

<sup>b</sup>*Mathematics Area, mathLab, SISSA, via Bonomea 265, I-34136 Trieste, Italy*

**Abstract.** The structural design of passenger ships is undergoing a transformative shift towards more sustainable and efficient solutions. This study focuses on practical applications of Machine Learning (ML) and optimization techniques to streamline the design process. By addressing challenges in structural optimization, such as high-dimensional parameter spaces and expensive simulations, we propose an advanced framework that integrates ML-driven parameterization and optimization workflows directly into the design pipeline.

Our methodology emphasizes the development of an efficient optimization strategy to balance conflicting objectives, such as minimizing material usage while ensuring structural resilience under extreme conditions. The construction of efficient surrogate models reduces dependency on computationally expensive Finite Element Analysis, leveraging predictive models to identify promising configurations before final validation. In addition, multi-objective optimization enables exploration of trade-offs between steel usage, safety constraints, and manufacturing costs. The methodology includes a parameterization refinement step, in which the addition of a limited number of new decision variables ensures an increased ability in capturing the structural behavior and stress distribution, with the consequent achievement of a better value of the objective function.

The framework has been successfully applied to multiple passenger ship projects, demonstrating significant improvements in the early design phase. This practical integration not only speeds up the design process but also facilitates innovative ship designs that meet stringent environmental and performance criteria. The approach highlights the potential of ML and optimization to support the maritime industry, bridging the gap between computational advancements and real-world applications.

**Keywords.**

machine learning, optimization, parameterization, design process, passenger ship

## 1. Introduction

The design of passenger ships is a complex engineering challenge that must balance comfort, safety, and performance while addressing environmental regulations and sus-

---

<sup>1</sup>Corresponding Author: [ciro.busiello@fincantieri.it](mailto:ciro.busiello@fincantieri.it).

tainability through innovative solutions for energy efficiency and minimal resource consumption [10]. Finite Element Analysis (FEA) remains the primary verification methodology in industrial structural design, ensuring compliance with safety regulations and performance requirements. However, the ship design process is subject to strict time constraints, requiring efficient computational strategies that can deliver accurate results within limited timeframes. For this reason, traditional design methods are no longer sufficient to achieve an optimal balance among these factors.

In this context, emerging technologies such as machine learning and optimization techniques are proving to be key tools in addressing these challenges. These approaches allow for the analysis of large datasets and complex simulations, facilitating the design of more sustainable and efficient ship structures. This paper will explore the role of these technologies in passenger ship design, highlighting potential benefits, innovative applications, and how these new techniques have been successfully integrated into the design process. In particular, we will present a computational pipeline for the optimization of the structural mass of passenger cruise ships in the early design phase. This pipeline was developed as a collaboration between Fincantieri S.p.A. and SISSA mathLab, and is extensively described in [11,3].

The paper is structured as follows: in Section 2 the optimization pipeline is described, and the engineering considerations on the design problem are detailed in Section 3. Section 4 describes the construction of reduced order models, the multi- and single-objective optimization procedures, and the automatic reparameterization procedure. Section 5 provides the considerations about the integration of the methodology into the design process. The application of the pipeline to a representative test case is discussed in Section 6, and Section 7 draws the final considerations.


## 2. Proposed Methodology

The proposed methodology relies on an advanced optimization pipeline that is fully integrated into the structural design process. This pipeline has been built to ensure that its inputs and tools align with those typically used by naval designers, facilitating a seamless integration into the existing design workflow. It leverages numerical optimization techniques to optimize the distribution of thicknesses in the primary structural elements of the ship, in order to minimize the structural mass while ensuring that the maximum deformation and structural weaknesses remain within safety limits.

The complete procedure is represented in Figure 1. The process begins with the creation of the mechanical model using a Finite Element Method (FEM) software that defines the geometry, meshing, and material properties of the structure (Section A).

Following the initial structural analysis, a random sampling of the parametric domain is generated, to create a database of high-fidelity solutions that reflects the potential design variations. Reduced Order Models (ROMs) are then developed from the solutions in the high-fidelity database, aiming to capture the essential structural behavior, while reducing computational complexity (Section B).

The next phase involves the use of a multi-objective optimization algorithm to explore the trade-offs between different quantities of interest (QoIs), such as material usage, safety constraints, and structural performance under extreme conditions. A subsequent cycle of single-objective optimization is then performed using Bayesian optimiza-



**Figure 1.** Flowchart of optimization pipeline. Each block represents a different module of the tool.

tion and a local search. These methods refine the design iteratively, focusing on minimizing the mass of the structure while maintaining structural performance requirements (Section D).

At this point, the parameterization of the model can be refined by generating new decision variables. The refined model becomes more expressive, allowing for the discovery of better optima that achieve superior design configurations (Section E).

The result obtained from the proposed pipeline represents an excellent starting point for standard design, significantly reducing development time and providing a solid foundation for further optimizations and refinements in the design process

3. Problem Definition

The structural optimization of cruise ships is constrained by rigid general arrangement requirements, limiting design flexibility. The geometry of the hull is dictated by functional and commercial needs, leaving the optimization to focus on variables impacting structural performance within these constraints.

A key design variable is the thickness of steel plates in the hull, decks, and bulkheads, constituting approximately 70% of the hull’s steel weight. Optimal allocation of plate thickness must balance weight reduction, structural integrity, and manufacturability. Unlike secondary reinforcements, plate thickness directly affects global ship behavior.

3.1. Quantities of Interest

The optimization process evaluates structural performance through a set of objective functions or QoIs that guide the search for an optimal configuration. These include total weight, which affects both performance and cost, and the vertical center of gravity (VCG), which plays a crucial role in ensuring the vessel’s stability. Other critical factors include transversal center of gravity (TCG), which influences transversal weight distribution, and the global deflection, which helps prevent excessive stress concentrations in localized structural assessments.

Additional QoIs monitor the number of plasticized elements, ensuring structural yielding remains confined to manageable regions that can be manually refined in a later

design phase. Similarly, the number of unstable elements is considered to maintain buckling stability while minimizing the need for extensive stiffening reinforcements.

Constraint violations are allowed within certain thresholds to promote a broader exploration of the design space, thereby improving optimization efficiency. Once an optimal configuration is identified, final refinements are applied, including localized thickness increases and anti-buckling reinforcements.


Some QoIs, such as total weight and VCG, can be computed directly, while others require FEM analyses. To enhance computational efficiency, model reduction techniques and surrogate modeling approaches accelerate FEM evaluations while maintaining accuracy.

### 3.2. Design Space

The selection of steel plate thickness occurs within a discrete design space, constrained by commercially available thickness increments (typically 0.5–1 mm) and regulatory requirements. These constraints include minimum thickness values mandated by classification societies, corrosion allowances to ensure long-term structural integrity, and local strength checks accounting for deck loads, watertight boundaries, and reinforcement needs. Upper bounds are also imposed to prevent excessive structural weight.

Structural optimization is conducted using a FEM, where elements are grouped based on similar stress behavior in order to define the parameters of the problem. Defining parameters appropriately is crucial for capturing stress variations while managing computational complexity. Common parameter definitions include central versus end zones for decks and bulkheads, longitudinal strips and transverse regions for external shell plating, and major internal structures for longitudinal bulkheads. Stiffened panels often maintain a fixed thickness to reduce complexity.

If too fine a parameterization is chosen, the number of possible configurations can be extremely large (e.g.,  $20^{50}$  for 50 parameters with 20 thickness levels), so careful parameter selection is necessary to ensure computational feasibility.



**Figure 2.** The FEA mesh of a cruise ship model on the left, and the initial parameterized sections chosen for the structural optimization on the right.

## 4. Methods


### 4.1. Reduced Order Models

Reduced order modelling are computational techniques used to simplify complex, high-dimensional systems while retaining the essential behavior of the original model. These

models are built using various techniques, but we focus on the combination of Proper Orthogonal Decomposition and Gaussian Process Regression (GPR)[6,8,5]. The main advantage of ROMs is their ability to drastically reduce computational costs, making it possible to perform fast simulations and optimizations, especially when dealing with large-scale problems like those in structural design. These models are particularly useful in engineering fields, where the goal is to evaluate multiple design alternatives quickly and efficiently without sacrificing accuracy. ROMs are built by extracting the most important features from high-fidelity simulations, such as those obtained from FEA. Together, these techniques allow ROMs to reduce the number of independent variables in a model while maintaining the accuracy needed for design decisions. The reduced models can then be used in optimization workflows to efficiently evaluate a wide range of alternative designs, supporting faster decision-making and enabling the design of more sustainable and cost-effective systems.

#### 4.2. Multi and Single Objective Optimization of the Surrogate Models

It is not possible to define an absolute optimal configuration in the structural design of a ship. Instead, the final solution must be a compromise between the various QoIs of the project. Only the designer can determine the most suitable trade-off between QoIs based on the specific requirements and constraints of the design. To support this decision-making process, it is essential to start with a multi-objective optimization study, aimed at identifying the Pareto frontier of the problem. The Pareto frontier is a set of optimal solutions, consisting of all points for which there is no point that is simultaneously better for all objectives considered in the optimization function. The multi-objective optimization process is carried out using a genetic algorithm [9], specifically the Non-dominated Sorting Genetic Algorithm (NSGA-III) [2,1]. This algorithm explores the design space and generates an approximate Pareto frontier through random mutation and selection of the best configurations. The Pareto frontier allows designers to evaluate various solutions and select the most appropriate one based on the design requirements. The objective functions considered in the optimization process include the ones presented in section 3.1.



**Figure 3.** The evolution of projected Pareto frontiers. On the left, the trade-off between the vertical deflection and total mass; on the right, the trade-off between the number of buckled elements and total mass. Between subsequent runs of the genetic algorithm, well-chosen candidates are added to the high-fidelity database and the ROMs are updated, so that the frontier is refined and more accurate.

After completing the multi-objective optimization phase, the next step in the design process is focused on refining the structural design through a single-objective optimiza-

tion approach, starting from a specific region of Pareto frontier that fits best with the priorities of the project. In this phase, we leverage a Bayesian optimization (BO) [4,7] algorithm to explore the design space more efficiently. This probabilistic approach to black-box optimization is particularly valuable as it reduces the number of FEA evaluations needed compared to traditional optimization techniques, making it balance exploration and exploitation. Constraints on the QoIs are enforced with a quadratic penalty term added to the mass computation. The designer decides how many elements can fail yielding or buckling criteria to satisfy a certain ship weight while maintaining the desired structural integrity and deformation constraints.


#### 4.3. Reparameterization

The cruise ship industry is constantly expanding, and recently, the design of ships exceeding 200,000 GRT has become increasingly common. As ships become larger, coupled with the desire to capture the optimization process as accurately as possible, the number of parameters necessary to optimize a model has increased significantly (80-100 parameters). The issue is further aggravated by the influence of parameter selection on the final result, which could leave the optimization stuck in a local optimum.

To limit the number of parameters, reduce simulation times and select parameters more effectively, our approach is to automatically generate a parameterization tailored to the stress state emerging from the optimization process. Since it is important to have an understanding of the structural behavior, even if approximate, a model with a minimal parameterization (called macro-parameters) is initially built. After an optimization cycle, a more refined parameterization is proposed by analyzing the structural responses of the elements, clustering those that exhibit similar behaviors. By separating different behaviors and associating the corresponding elements to different decision variables, the model adapts to the previously unknown stress configuration and future optimizations are able to exploit the increased versatility.


#### 4.4. Macro Parameters, Patches and Clustering

On average, a complete model is built with about ten macro-parameters, and the subdivision must consider the structural orientation (decks, longitudinal bulkheads, transverse bulkheads) and the regulatory minimum thickness requirements. By respecting these constraints, for example, an external shell will be part of a different macro-parameter than an internal longitudinal bulkhead, as they have different regulatory minimum thickness requirements. The concept of patches is related to the way steel plates are distributed. In fact, the size of a patch is equal to the size of a steel plate, approximately 12 meters in length and 3 meters in width. Several patches can be grouped to form new parameters.



**Figure 4.** The FEA mesh of a cruise ship model on the left, center the macro-parameters and patches definition on the right.

Grouping of the patches is performed by solving an optimization problem defined as an Integer Linear Program (ILP), which considers the sensitivity of the structural responses of each patch in relation to its thickness. The patches with the same optimal thickness are grouped together to form a single parameter.



**Figure 5.** Details of reparameterization of deck 14. Top left is the  $\sigma_x$  component of the stress tensor in sagging condition, red represent the larger values and blue is zero. Center left is the patch subdivision, each color denotes a different patch. Bottom left is the parameterization proposed by the designer, with 3 sections each. The right column represents the subsequent parameterizations from the clustering procedure, from top to bottom: 16 parameters for the entire model, 27 and 40.

## 5. Integration into the Design Process

The traditional approach to designing using finite element models involves manually modifying the thickness of plates to meet the regulations' criteria. This involves repeating multiple simulations, with a large expenditure of computational time and resources.

Compared to a traditional approach, the proposed pipeline requires minimal manual adjustments. The designer's main task is to define macro parameters and set up the problem, including thickness variability, constraints, and objective function selection. This setup process typically takes about one day. The initial sampling for the construction of ROMs takes 1 night, multi-objective and single-objective optimization takes 1 day, and the reparameterization takes 1 day. The last 2 sections can be repeated at least 1 time. While total computational time is approximately one week, the designer's direct involvement is minimal after the initial setup. The automated process significantly reduces manual effort while achieving a more efficient and accurate optimization outcome.


This approach is fully compatible with standard design timelines. Moreover, only requiring a global FEA model allows the methodology to be employed even at the bid design stage, providing valuable insights early in the project life cycle.

## 6. Experimental Results

In this chapter, we present the results of applying the proposed method to a real-world case study involving a cruise ship with an approximate gross tonnage of 80,000 TSL, focusing on macroscopic outcomes in terms of QoIs [3].

Ten macro parameters, assigned to specific regions as described in Table 1, were used as the initial subdivision of the FE model. Subsequently, optimization was performed through three cycles of parameterization, resulting in a final subdivision into 40 parameters.

The graph in Figure 6 highlights how three cycles of optimization with automatic parameterization (base 10p → our 16p → our 27p → our 40p) lead to a better optimum than the model generated by a naval designer (designer 40p) and a single reparameterization (base 10p → one-shot 40p).



**Figure 6.** Parameterization History. This figure shows the evolution of the objective function—selected for single-objective optimization—as a function of the number of high-fidelity simulations. Horizontal segments represent the process of expanding the high-fidelity FE results database, while the steps indicate the discovery of improved optimal configurations.

The effect of the parametrization is described in Figure 7, where the vertical axis shows the difference between the penalized mass of the incumbent configuration and its lower bound, as a percentage of the latter. The expression used is

$$m_{\text{gap}}(\mathbf{x}) = 100 \frac{\mathbf{d} \cdot \mathbf{x} + m_{\text{bar}} n_b(\mathbf{x}) + f_{\text{pen}}(\mathbf{x}) - \mathbf{d} \cdot \mathbf{x}_{\text{LB}}}{\mathbf{d} \cdot \mathbf{x}_{\text{LB}}}, \quad (1)$$


where  $\mathbf{d}$  is a vector of linear densities associated to the parameters,  $\mathbf{x}$  and  $\mathbf{x}_{\text{LB}}$  are respectively the vectors of parameter values and their lower bounds,  $m_{\text{bar}}$  is the mass of a stiffener,  $n_b(\mathbf{x})$  is the number of stiffeners, and  $f_{\text{pen}}(\mathbf{x})$  is the penalization due to the violation of constraints on black-box QoIs.

In this test case, the optimization mainly worked on the parameters *Decks in public areas*, *Bottom*, and *Internal bulkheads*, highlighting their importance as the primary structures for longitudinal strength. The mass of these three groups of parameters decreased as the number of parameters increased, while the distribution of the yielded and unstable elements was gradually modified, with most of them shifting to the *Internal bulkheads* parameter.

In optimization, focus is often placed on solution accuracy and avoiding local minima, assuming the problem is correctly formulated. However, the selection of parameters plays a crucial role in achieving meaningful results. To evaluate the tool's performance, the optimal solutions obtained through three cycles of automated parameterization were compared with a configuration consisting of 40 parameters manually selected by an experienced designer. We also analyze a single "one-shot" reparameterization, from 10 to 40 parameters. Comparing the results shown in Table 2 the automated parameterization made choices similar to those of an expert, but with greater efficiency and an enhanced ability to capture structural behavior. The 40 parameters model obtained from repeated reparameterization achieved an optimized mass closer to the lower bound and fully uti-

**Table 1.** Parameters description of the full ship test case. The regions where the initial subdivision into macro parameters is defined are highlighted. For each region, the table shows the number of patches involved, the minimum and maximum allowed thickness (in mm), and the number of possible values to be considered in the optimization.

| Macro Parameter | Region                                  | Patches | Min. | Max. | # values |
|-----------------|-----------------------------------------|---------|------|------|----------|
| 1               | Decks in public areas                   | 3038    | 5.0  | 25.0 | 16       |
| 2               | Decks in machinery areas                | 548     | 6.5  | 25.0 | 13       |
| 3               | Lifeboat exposed deck                   | 90      | 13.0 | 25.0 | 5        |
| 4               | Inner bottom                            | 252     | 12.0 | 25.0 | 6        |
| 5               | Bottom                                  | 312     | 15.0 | 25.0 | 3        |
| 6               | Shell plating, above waterline          | 60      | 14.0 | 25.0 | 4        |
| 7               | Shell plating, below waterline          | 118     | 15.0 | 25.0 | 3        |
| 8               | Internal bulkheads                      | 367     | 5.0  | 25.0 | 16       |
| 9               | External bulkheads, in super-structures | 190     | 8.0  | 25.0 | 10       |
| 10              | External bulkheads                      | 31      | 6.5  | 25.0 | 13       |



**Figure 7.** Effects of the reparameterizations on the distribution of mass and buckled elements across the initial macro-parameters. The four bars for each macro-parameter correspond to parameterizations with 10, 16, 27 and 40 parameters, respectively. Different colors correspond to different decision variables, with red denoting the most recent additions and darker shades of blue the original subdivision.

**Table 2.** Optimization results for different parameterizations of the full ship. The columns report the optimum mass value computed as in (1), followed by the number of yielded and unstable elements, and the VCG, expressed as a percentage of the respective penalized threshold values. The first four rows report the evolution of the automatic parameterization, while the last row shows the results obtained with the parameterization defined by an experienced designer.

| Parameterization   | $m_{\text{gap}}$ [% LB] | # yielded [% crit] | # buckled [% crit] | VCG [% crit] |
|--------------------|-------------------------|--------------------|--------------------|--------------|
| Base 10 p.         | 34.21                   | 46.25              | 100.38             | 99.35        |
| Our method 16 p.   | 19.01                   | 56.00              | 99.88              | 99.76        |
| Our method 27 p.   | 14.12                   | 98.25              | 99.29              | 99.71        |
| Our method 40 p.   | 13.03                   | 98.25              | 100.01             | 99.88        |
| Our one-shot 40 p. | 15.07                   | 51.50              | 99.87              | 99.85        |
| Designer 40 p.     | 23.55                   | 55.00              | 100.15             | 99.99        |

lized the thresholds set for yielded and buckled elements, as well as the position of the VCG. This suggests that the optimization was limited by the actual constraints and not by the definition of the parameters, while other approaches reached their best  $m_{\text{gap}}$  utilizing only around 50% of the budget of yielded elements.

The distribution of obtained thickness values demonstrates sound engineering logic as shown in Figure 5. The optimizer identifies broad areas with constant thickness and

gradual variations, ensuring an efficient yet structurally robust design.

## 7. Conclusions

The work presented in this paper is the result of five years of collaboration between Fincantieri and SISSA mathLab. The outcome is an efficient computational pipeline that seamlessly integrates advanced computational techniques with practical engineering experience, ensuring compatibility with standard design timelines and processes. In the traditional approach, we estimate that a designer would require at least three weeks of continuous work to reach a configuration of a quality comparable with the results of just a week of pipeline execution, which instead involves minimal user interaction.

The adoption of such a process has the benefit of ensuring a holistic approach by forcing designers to evaluate all conflicting requirements at the beginning, rather than integrating them one at a time and evaluating the results incrementally. Another positive aspect is the autonomous exploration of vast parametric spaces, not limited by the possibly biased and incomplete experience of the designer, thus unlocking new opportunities for ship structural optimization. This is of great value when tackling innovative projects, where prior knowledge cannot serve as a reliable guide and the best solution may not be intuitive. An indirect consequence of this approach is that designers working on separate projects are provided with initial configurations generated using a consistent methodology, increasing standardization and preventing the introduction of inefficient choices at the early stages of development.

The current pipeline is already being used in production without major limitations. In particular, automated parameterization has solved the key challenge of identifying the smallest set of necessary parameters, and future challenges will focus on improving computational and optimization efficiency to further reduce processing times.

Promising strategies include applying machine learning techniques, leveraging models built from past optimizations, both for re-parameterization and during optimization. At a higher level, future research could also explore the analysis of multi-objective optimization results, which is currently based on engineering experience but could instead be guided by more advanced methodologies.

## References

- [1] K. Deb and H. Jain. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. *IEEE Transactions on Evolutionary Computation*, 18(4):577–601, 2014. doi:10.1109/TEVC.2013.2281535.
- [2] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. doi:10.1109/4235.996017.
- [3] L. Fabris, M. Tezzele, C. Busiello, M. Sicchiero, and G. Rozza. Data-driven parameterization refinement for the structural optimization of cruise ship hulls, 2024. URL: <https://arxiv.org/abs/2411.09525>, doi:10.48550/ARXIV.2411.09525.
- [4] R. B. Gramacy. *Surrogates - Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. Chapman & Hall/CRC, Boca Raton, 2021.
- [5] M. Guo and J. S. Hesthaven. Reduced order modeling for nonlinear structural analysis using Gaussian process regression. *Computer Methods in Applied Mechanics and Engineering*, 341:807–826, Nov. 2018. doi:10.1016/j.cma.2018.07.017.

- [6] J. S. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. SpringerBriefs in Mathematics. Springer International Publishing, Cham, 2016. [doi:10.1007/978-3-319-22470-1](https://doi.org/10.1007/978-3-319-22470-1).
- [7] M. J. Kochenderfer and T. A. Wheeler. *Algorithms for Optimization*. The MIT Press, 2019.
- [8] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, Massacchuttes, 2005. [doi:10.7551/mitpress/3206.001.0001](https://doi.org/10.7551/mitpress/3206.001.0001).
- [9] S. N. Sivanandam and S. N. Deepa. *Introduction to Genetic Algorithms*. Springer, Berlin, Heidelberg, 2008. [doi:10.1007/978-3-540-73190-0](https://doi.org/10.1007/978-3-540-73190-0).
- [10] M. Tadros, M. Ventura, and C. G. Soares. Review of current regulations, available technologies, and future trends in the green shipping industry. *Ocean Engineering*, 280:114670, July 2023. [doi:10.1016/j.oceaneng.2023.114670](https://doi.org/10.1016/j.oceaneng.2023.114670).
- [11] M. Tezzele, L. Fabris, M. Sidari, M. Sicchiero, and G. Rozza. A multifidelity approach coupling parameter space reduction and nonintrusive POD with application to structural optimization of passenger ship hulls. *International Journal for Numerical Methods in Engineering*, 124(5):1193–1210, Nov. 2023. [doi:10.1002/nme.7159](https://doi.org/10.1002/nme.7159).