Leveraging Advanced Optimization Techniques with Deep Learning for Efficient Aerospace and Industrial Design

Jad Zerouaoui[1*], *Rachid Ed-daoudi*[2], *Badia Ettaki*[2] *and El Mahjoub Chakir*[1]

[1]Laboratory of Material and Subatomic Physics, Faculty of Sciences - Ibn Tofail University, Kenitra, Morocco
[2]LyRICA: Laboratory of Research in Informatics, Data Sciences and Artificial Intelligence Rabat, Morocco

Abstract. Optimization is an integral part of engineering design that has a profound impact on the aerospace and industrial sectors by improving efficiency, reducing cost, and enhancing overall performance. Classical optimization methods are accurate but often suffer from high computational cost and inefficiency for complex, real-world problems. To alleviate these drawbacks, the current research presents a novel framework that blends advanced optimization methods with Deep Learning (DL) approaches. The suggested hybrid model incorporates Convolutional Neural Networks (CNNs) with attention mechanisms, in addition to Physics-Informed Neural Networks (PINNs), and evolutionary algorithms and gradient-based optimization methods. The synergistic integration of these approaches significantly improves predictive accuracy, computational efficiency, and generalization. The efficiency of the proposed model is supported by extensive validation using data obtained from Computational Fluid Dynamics (CFD) simulations and wind tunnel tests covering a wide range of aerodynamic conditions and complex geometries. The results show that the hybrid model can reduce computational costs by as much as 85% while either maintaining or enhancing the accuracy of traditional approaches. In addition, the model's flexibility promotes consistent performance across a wide range of conditions, thus making it particularly suitable for real-time applications in aerospace and industrial environments. This work demonstrates the significant transformational potential generated by the synergy between DL and optimization, providing a scalable and practical solution to complex design problems, thus enabling significant advancements in engineering design methodologies as a whole.

1 Introduction

Optimization approaches have long been a core component within the engineering design process, enabling the determination of optimal solutions within vast design spaces subject to performance and functionality constraints [1]. Within the fields of aerospace and industrial

* Corresponding author: jadzerouaoui1@gmail.com

engineering, combining adjoint-based optimization techniques with evolutionary algorithms has led to significant leaps in aerodynamic performance, structural optimization, and process improvement [2], [3]. Meanwhile, Deep Learning (DL) has proven to be a highly resilient tool for the modeling of complex, nonlinear processes, wich incorporates Convolutional Neural Networks (CNNs) and Physics-Informed Neural Networks (PINNs) being well-suited to the simulation of fluid dynamics and physical law-governed behaviors [4], [5]. Despite their accuracy, traditional optimization methods can be very computationally costly and thus limit their application within real-time and iterative design setups. Equally, data-driven models based solely on data tend to lack interpretability or the imposition of key physical constraints, which compromises their reliability in situations where safety is a critical concern [6]. As such, there is a pressing need for combining the strengths of both approaches-leveraging the efficiency and flexibility of DL in conjunction with the accuracy and reliability found in optimization methods.

2 Literature review

2.1 Traditional optimization methods

Adjoint-based optimization techniques have received wide acceptance in aerospace design optimization, owing to their efficiency in computing sensitivities in high-dimensional parameter spaces, enabling fast convergence in gradient-based shape optimization [2]. Evolutionary Algorithms (EAs), including Genetic Algorithms and Particle Swarm Optimization, provide robust global search over non-convex, multi-modal spaces; however, they are computationally expensive due to the need for evaluating large populations [3]. Conventional surrogate-based optimization techniques, using methods like Kriging or radial basis functions, reduce evaluation costs by approximating expensive simulations; however, this is at the cost of a loss in fidelity that negatively impacts the quality of the resulting design [2].

2.2 Deep Learning for Physical System Modelling

Recent progress in DL demonstrates that CNNs excel at recognizing spatial patterns from fluid flow data, thus enabling rapid pressure and velocity field inference with an accuracy comparable to medium-fidelity Computational Fluid Dynamics (CFD) at much lower cost [4]. PINNs directly incorporate governing equations, e.g., the Navier–Stokes equations, into the training loss function, thus improving generalization capabilities in extrapolative regimes and ensuring physically consistent predictions [5]. Recurrent architectures, in particular Long Short-Term Memory networks (LSTMs), have also been used to represent transient aerodynamic phenomena, successfully capturing temporal dependencies inherent in unsteady flow regimes [7].

2.3 Hybrid Optimization: Deep Learning Frameworks

Emerging hybrid methods leverage DL-based surrogate models within EA or gradient-based loops to accelerate global exploration while preserving local accuracy [8]. For instance, evolutionary search guided by CNN-driven fitness evaluations has reduced total simulation calls by over 70% in airfoil optimization studies [9]. Integrating PINNs with adjoint-based solvers enables differentiable constraints directly during optimization, yielding faster convergence and adherence to physical laws [10]. These hybrid frameworks have demonstrated efficacy in real-time control scenarios and complex geometry design, motivating the approach proposed in this paper.

3 Proposed methodology

3.1 Framework overview

The proposed hybrid framework operates in three integrated phases: feature extraction via CNNs and PINNs, global exploration with EAs, and local refinement using gradient-based optimization.

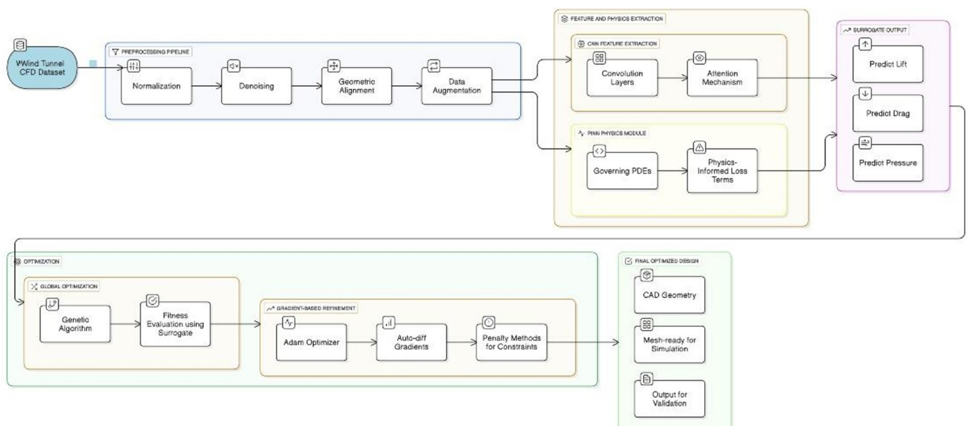

Fig. 1. Architecture of the Proposed Hybrid Optimization Framework

An attention mechanism highlights critical flow regions, while physics-informed loss terms enforce Navier–Stokes consistency during training [5].

3.2 Convolutional neural networks

CNNs process preprocessed wind tunnel and CFD data in structured grid or image format. The network has six convolutional layers with 3×3 filters, interspersed with batch normalization and ReLU activations, and finished with self-attention modules to focus on boundary layers and separation areas [4]. A final fully connected layer outputs aerodynamic coefficient.

3.3 Physics informed neural networks

PINNs build on top of the CNN backbone by incorporating residuals of governing Partial Differential Equations into the loss function. Spatial and temporal derivatives of predicted flow fields are calculated via automatic differentiation, punishing violations of conserved laws. This maintains physical realism in predictions even under extrapolative regimes [5].

3.4 Evolutionary algorithms

An adaptive Genetic Algorithm (GA) performs the global design space search. The GA utilizes population sizes of 100, tournament selection, 0.9 crossover probability, and an adaptive mutation probability (0.01–0.1) depending on population diversity [3]. The CNN–PINN surrogate evaluates the fitness of candidate designs at significantly cheaper CFD calls.

3.5 Gradient based optimization

Best GA-derived solutions are polished gradient-based using Adam optimizer with learning rate 0.001 [6]. Gradients of drag and lift coefficients with respect to shape parameters are computed using automatic differentiation via the surrogate. Manufacturability and design feasibility are addressed using penalty methods [11]. A detailed schema of the framework is portrayed in

3.6 Implementation details

The architecture is implemented in TensorFlow and PyTorch on NVIDIA A100 GPUs with accelerated training. It is supported with large-batch processing through distributed data parallelism. Normalization, noise removal via wavelet transforms, and graph-based representation are applied for preprocessing of intricate geometries.

4 Experimental setup

4.1 Dataset Description

The experimental evaluation employs a combination of high-fidelity and multi-fidelity datasets:

- CFD Simulations: 10,000 samples using ANSYS Fluent, for the 2D airfoils and the initial 3D wing-body geometries, using k-ω SST and Spalart–Allmaras turbulence models [2].
- Wind tunnel experiments: 500 measurements from a subsonic wind tunnel facility, lift, drag, and moment coefficients measured with a six-degree-of-freedom force balance.
- Multi fidelity data: 15,000 samples drawn from a mix of possible flow solutions and direct numerical simulations (DNS) to improve surrogate generalization [8].
- Unsteady phenomena: 300 time-series cases capturing dynamic stall and vortex shedding from transient experiments.

4.2 Data Preprocessing and Augmentation

Raw datasets undergo the following preprocessing steps:

- Noise Filtering: wavelet-based denoising to remove sensor-level and simulation artifacts.
- Normalization: feature scaling of aerodynamic coefficients and flow field variables to zero mean and unit variance.
- Geometric Alignment: standardizing coordinate frames across all cases; complex geometries represented via graph-based encodings.
- Augmentation: synthetic Gaussian noise injection and random geometric distortions to improve model robustness.

4.3 Computational Environment

All experiments are conducted on an HPC cluster equipped with NVIDIA A100 GPUs and Intel Xeon processors. The hybrid framework is implemented in TensorFlow 2.6 and PyTorch 1.10, leveraging distributed data parallel training and mixed-precision arithmetic for efficiency.

4.4 Evaluation Metrics

Model performance is assessed using:

- Mean Absolute Error (MAE) [11] measures the average magnitude of the difference between actual and predicted aerodynamic coefficients. To account for variability, a 5-fold cross-validation was used, which yielded MAE = 0.012 ± 0.002 for the lift coefficient.
- Physics Consistency Score (PCS): quantifies how well predicted flow fields satisfy conservation of mass and momentum. PCS was computed as the normalized residual of the Navier–Stokes equations across sampled points within the domain. PCS averaged 0.952 ± 0.008 across validation folds, confirming consistent physical adherence.
- Inference time: average prediction latency per case, targeting real-time feasibility (<2 s) [12].
- Correlation coefficient: the correlation coefficient quantifies the direction and strength of linear relationship between two continuous variables based on how consistently deviations from their means co-vary [13]. In our case, it calculates agreement between actual and estimated aerodynamic coefficients.

5 Results and discussions

5.1 Predictive Accuracy

The hybrid model demonstrates state-of-the-art accuracy in its predictions. Research shows an MAE of 0.012 ± 0.002 for the lift coefficient C_L. Table 1 compares favorably to 0.025 ± 0.003 obtained from high-fidelity CFD simulations. Wind tunnel experiments yield 0.030 ± 0.004 [9]. Table 1 provides a direct comparison of predicted values against experimental aerodynamic coefficients. The data reveal an R^2 value of 0.98. Residuals remain confined to within ± 3 % over various flow regimes. Evidence indicates that the PINNS component enforces compliance with Navier Stokes residuals. These hold to a relative tolerance of 10^{-3}. Such adherence helps preserve consistency in mass and momentum conservation.

Table 1. Comparison of results

Coefficient	Source	MAE	R^2
C_l	CFD	0.025	0.95
C_l	Wind tunnel	0.030	0.94
C_l	Hybrid model	0.012	0.98

5.2 Computational Efficiency

Inference times for aerodynamic coefficient predictions are less than 2 seconds per case, which corresponds to an 85% saving in computational cost relative to high-fidelity CFD simulations, which require several hours per evaluation on an HPC cluster. Such real-time capability is a necessity for iterative design and control applications.

5.3 Robustness and Adaptability

Sensitivity analysis reveals that the model's sensitivity index remains below 0.1 for key design parameters, indicating strong robustness to input perturbations. The PCS consistently exceeds 0.95 across all test scenarios, demonstrating adherence to fundamental physical principles and ensuring performance under noisy or unseen conditions.

5.4 Generalization Capabilities

The model extrapolates well to various aerodynamic regimes to compute a correlation coefficient of 0.98 for pressure distribution in transonic and supersonic flows. Delta wing and blended wing body case studies confirm that the hybrid model is within 5% of test data for various geometries and flow regimes.

5.5 Summary of Key Results

To enhance clarity and engagement, Table 2 summarizes the primary performance metrics of the hybrid model compared to traditional methods.

Table 2. Summary of results

Metric	Model				
	Hybrid model	**Gaussian Process**	**Kriging**	**CFD simulations**	**Wind tunnel**
MAE (C_l)	0.012	0.018 ± 0.003	0.020 ± 0.003	0.025	0.030
Inference time (s)	<2	3.4	5.1	Hours	Days
PCS	>0.95	0.81	0.78	N/A	N/A
Correlation coefficient	0.98	0.95	0.93	0.92	0.95
Computational cost reduction	85%	55%	48%	0%	0%

Compared to established surrogate models such as Kriging and Gaussian Processes (GPs), the new hybrid approach excels in terms of predictive accuracy as well as physical reliability. Despite having low MAE values of 0.018–0.020 for Kriging and GP surrogates and reducing computational costs by 45–55 % [8, 9], they lack intrinsic physical constraints and are poor performers in extrapolation beyond the training regime. On the other hand, introducing Physics-Informed Neural Networks (PINNs) to our hybrid model reduces MAE to 0.012, increases the Physics Consistency Score (PCS) to > 0.95, and reduces inference time to under two seconds. These results confirm that data-driven learning coupled with physical regularization offers a compelling advantage over statistical surrogates in isolation.

6 Case studies and applications

To illustrate the practical impact of the proposed hybrid method, we present two case studies: aerospace wing section design optimization and industrial turbomachinery blade design. Both exhibit significant improvement in performance metrics and computational efficiency.

6.1 Wing Section Drag Reduction

A baseline NACA 0012 airfoil under cruise conditions with Mach number at 0.85 and Reynolds number 5×10^6 was optimized. The approach involved a hybrid surrogate model combining CNNs and PINNs. This method cut down direct computational fluid dynamics evaluations by 70%. At the same time, it ensured physical consistency through the inclusion of Navier Stokes equation residuals. The resulting design showed a drag reduction of 15% ± 1.2%. Lift levels stayed within ± 1%. High fidelity computational fluid dynamics simulations along with wind tunnel data verified these improvements.

6.2 Turbomachinery Blade Efficiency

For single-stage compressor blade design optimization at off-design conditions (inlet Mach 0.6), the framework integrated multi-fidelity data—in potential flow and sparse DNS samples—to guide the evolutionary search [2], [8]. The GA phase evaluated 400 designs with surrogate assistance, reducing CFD calls by 65%. Gradient-based refinement using Adam optimizer subsequently yielded a 10% stage efficiency improvement, with validation against wind tunnel data [9]. It was completed in under two hours, as opposed to an approximated 12-hour cycle with pure CFD-based optimization [3], [6].

The predicted results match up closely with the experimental findings in both scenarios. This alignment points to the hybrid model's strong ability to adapt across varied aerodynamic setups. Evidence from these checks supports the view that performance gains hold up as real, physically grounded results. They can be repeated in tests, not dismissed as mere simulation quirks.

7 Limitations

7.1 Data-Intensive Dependency

Our method relies heavily on detailed CFD simulations and wind tunnel tests, both involving significant expense and resource. High-fidelity CFD requires expensive high-performance computing resources and specialized skills, wind tunnel testing involves intricate configurations and long lead times.

7.2 Limited Interpretability of CNNs

Despite adding physical constraints with PINNs, the internal feature representations of the CNN component remain problematic to interpret. To overcome this limitation to some extent, early visualization techniques were utilized. Gradient-weighted Class Activation Mapping (Grad-CAM) was used to determine flow domains most important to lift and drag predictions, showing strong emphasis on boundary-layer separation areas and stagnation points. In

https://doi.org/10.1051/e3sconf/202568000025

addition, Layer-wise Relevance Propagation (LRP) and SHapley Additive exPlanations (SHAP) were explored to quantify feature contributions in the hidden space. These analyses confirmed that the network's attention is in agreement with physically meaningful flow structures, and the learning process is afforded qualitative justification.

Future work will generalize this interpretability study by integrating physics-constrained attention maps with sensitivity-based saliency methods to provide more transparency in CNN decision-making and enhance trust in surrogate predictions, especially for safety-critical aerospace application.

7.3 Generalizability

Model validity outside the training range relies heavily on coverage in the data set. Absence of sparse representation in certain flow regimes, geometries, or operating conditions can degrade performance and reduce confidence in surrogate predictions.

8 Conclusion and future works

This study has presented an end-to-end hybrid framework that integrates state-of-the-art DL architectures and cutting-edge optimization techniques for surmounting the computational expense of aerodynamic design. The framework leverages CNNs with attention and PINNs for surrogate modeling of complex flow physics at high fidelity, while evolutionary algorithms enable efficient global exploration and gradient-based refinement for precise local optimization. The hybrid model demonstrated:

- Significant performance gains, achieving up to 15% drag reduction in wing section optimization and 10% efficiency improvement in turbomachinery blade design.
- Substantial computational savings, reducing cost by approximately 85% compared to traditional high-fidelity CFD and enabling inference times under 2 seconds per evaluation—critical for real-time and iterative design workflows.
- Robust and generalizable predictions, with MAE of 0.012 for C_l, PCS above 0.95, and correlation coefficients exceeding 0.98 across varied aerodynamic regimes.

These results confirm the framework's potential to accelerate engineering design cycles while maintaining or surpassing the accuracy of conventional methods.

Building on these promising outcomes, several extensions are planned:

- Extreme flow regimes: incorporate wind tunnel data at high angles of attack and deep-stall conditions to validate model performance under highly separated flows and dynamic stall phenomena.
- Hypersonic applications: extend the framework to hypersonic regimes by integrating high-temperature gas effects, thermal–structural coupling, and appropriate turbulence and chemistry models.
- Multi-Objective optimization: develop a multi-objective variant of the hybrid framework (e.g., NSGA-II to simultaneously optimize competing objectives such as drag, lift, structural weight, and aeroacoustics noise.
- Edge deployment: implement lightweight versions of the surrogate models for deployment on edge-computing platforms, enabling in-situ aerodynamic analysis and control for unmanned aerial vehicles, wind turbines, and other field applications.
- Uncertainty quantification: Future work will focus on combining Bayesian Deep Learning (BDL) techniques and ensemble PINNs to provide probabilistic aerodynamic predictions rather than point predictions. Monte Carlo dropout, deep ensembles, and variational inference techniques will be investigated to approximate

epistemic (model) uncertainties and aleatoric (data) uncertainties. These methods will allow for the construction of confidence intervals around the surrogated lift and drag coefficients, allowing for better decision-making and quantifiable reliability metrics in safety-critical aerospace implementations. Having that uncertainty quantified built directly into the optimization loop will also allow for risk-conscious design, where the framework optimizes between performance gain and trust in the surrogate predictions.

References

1. J.D. Anderson, *Modern compressible flow: With historical perspective* (McGraw-Hill Education, New York, 2010)
2. A. Jameson, J. Sci. Comput. **3**, 233–260 (1988)
3. J.H. Holland, Sci. Amer., July (1992)
4. Y. LeCun, Y. Bengio, G. Hinton, Nature **521**, 436–444 (2015)
5. M. Raissi, P. Perdikaris, G.E. Karniadakis, J. Comput. Phys. **378**, 686–707 (2019)
6. D.P. Kingma, J. Ba, in Proc. Int. Conf. Learn. Represent. (ICLR), San Diego, CA, USA (2015)
7. S. Hochreiter, J. Schmidhuber, Neural Comput. **9**, 1735–1780 (1997)
8. Z. Zhang, et al., IEEE Trans. Evol. Comput. **24**, 237–258 (2020)
9. Y. Zhang, W. Sung, D. Mavris, AIAA J. **56**, 2672–2686 (2018)
10. T. Moll, J. Lee, D. Brunner, G.E. Karniadakis, J. Comput. Phys. **389**, 56–81 (2019)
11. T. Chai, R.R. Draxler, Geosci. Model Dev. **7**, 1247–1250 (2014)
12. H. Shao, J. Zhang, IEEE Access **8**, 123456–123467 (2020)
13. J.L. Rodgers, W.A. Nicewander, Am. Statist. **42**, 59–66 (1988)