Session 14

Turbomachinery and Airfoils

Optimal Input Design for Dynamic Stall Modelling of a Pitching Aerofoil Using Reinforcement Learning

Witmeur Charly[1], Decuyper Jan[1], Schoukens Maarten[2], Runacres Mark[1].

[1]*Thermo and Fluid Dynamics (FLOW), Vrije Universiteit Brussel (VUB).*

[2]*Control Systems Group, Eindhoven University of Technology.*

Abstract

Pitching wings, with a time-varying angle of attack, are encountered in many applications such as helicopter blades, vertical-axis wind turbines, and hydrokinetic cross-flow turbines. For engineering purposes, including design optimisation and real-time control, fast inference of the highly nonlinear and unsteady input-output relationship between angle of attack and lift coefficient is required. In this context, recent data-driven modelling techniques emerge as promising solutions. However, despite the plethora of available model structures and parameter estimation methods in the literature, the question of how to generate an informative and low-cost training dataset for nonlinear system identification remains largely open.

In this paper, we design an input signal with the objective of achieving space-fillingness in the system's state space. To the authors' knowledge, this is the first time such an input design strategy has been applied to a pitching wing system. Reinforcement learning (RL) is employed to accomplish this task, demonstrating significant advantages in flexibility and computational efficiency compared to existing optimisation-based techniques. The RL agent is trained by interacting with a surrogate-model environment, enabling efficient training without costly interactions with either an experimental setup or CFD simulations. The surrogate is trained on CFD data generated beforehand using a state-of-the-art swept-sine input.

The RL-generated input achieves higher space-filling performance than the swept-sine while reducing the experiment duration by an order of magnitude. Importantly, the agent is free to generate any input within predefined constraints and is not restricted to known excitation classes, resulting in a first-of-its-kind input for pitching wing systems. Finally, we validate that the designed input forms a high-quality dataset for model training, producing a model comparable to one trained on swept-sine data.

Keywords: Pitching aerofoil, Unsteady aerodynamics, Nonlinear system identification, Input design, Reinforcement learning.

Modelling NACA 0018 Unsteady Aerodynamics using State-Space Kolmogorov-Arnold Networks

Gonçalo Granjal Cruz[1,2,*]; Mark Runacres[1,2]; Jan Decuyper[1,2]

[1]*Thermo and Fluid Dynamics (FLOW), Faculty of Engineering, Vrije Universiteit Brussel (VUB), Brussels, 1050, Belgium.*

[2]*Brussels Institute for Thermal-Fluid Systems and Clean Energy (BRITE), Vrije Universiteit Brussel (VUB) and Université libre de Bruxelles (ULB), Brussels, 1050, Belgium.*

Abstract

The accurate prediction of unsteady aerodynamic loads is critical for the design and control of systems operating in complex flow regimes, such as Vertical Axis Wind Turbines (VAWTs) and Unmanned Aerial Vehicles (UAVs). Dynamic stall, characterized by severe hysteresis and sharp transient lift fluctuations, typically requires high-fidelity Computational Fluid Dynamics (CFD) or case specific semi-empirical models. While deep learning offers a data-driven alternative, standard recurrent architectures often struggle to efficiently capture the sharp nonlinearities of stall with a reduced number of parameters. This work introduces the use of State-Space Kolmogorov-Arnold Networks (SS-KAN) [1], a novel architecture that integrates KANs in the structured dynamical framework of state-space models, for modelling unsteady aerodynamics.

The proposed SS-KAN model is trained on experimental data of a NACA 0018 aerofoil at a Reynolds number of $Re = 2.8 \times 10^5$ with 0.3% turbulence intensity. The training dataset consists of "artificial aerodynamic" pitch manoeuvers designed to excite a broad range of frequencies and amplitudes coupled with the correspoding pitch rates, enabling the prediction of the lift coefficient (C_L).

We demonstrate in Figure 1 that the SS-KAN architecture successfully captures the full envelope of unsteady aerodynamics, including the reconstruction of deep stall hysteresis loops. The model maintains predictive accuracy on unseen picthing maneuvers while utilizing significantly fewer parameters than comparable MLP-based models. This combination of accuracy, compactness, and dynamical structure positions SS-KAN as an option for real-time applications, such as model predictive control (MPC) and aeroelastic analysis in low-Reynolds number flight regimes.

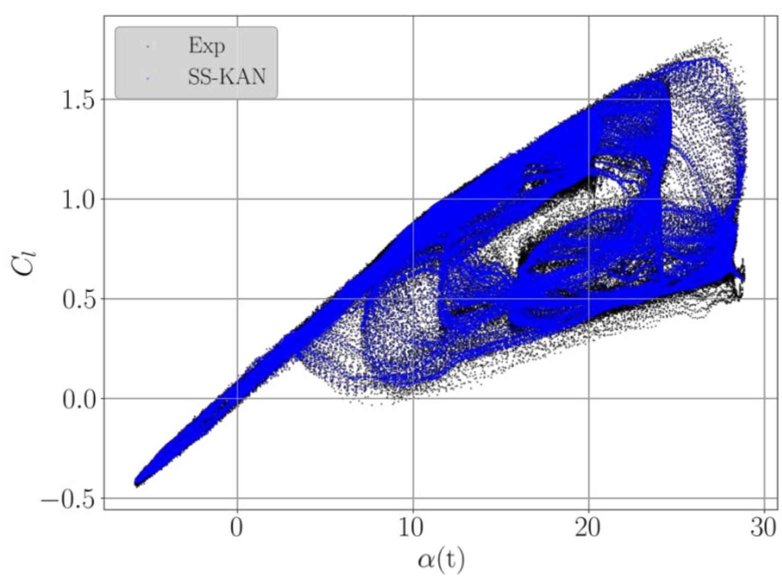

Figure 1 - SS-KAN model comparison prediction against reference data on the pitching envelope of a NACA 0018 aerofoil

Keywords: unsteady aerodynamics; nonlinear system identification; state-space kolgomorov-arnold networks

[1] G. G. Cruz, B. Renczes, M. C. Runacres and J. Decuyper, "State-Space Kolmogorov Arnold Networks for Interpretable Nonlinear System Identification," in *IEEE Control Systems Letters*, vol. 9, pp. 847-852, 2025, doi: 10.1109/LCSYS.2025.3578019.

Data-Driven Sparse Reconstruction of Three-Dimensional Steady and Unsteady Flow Fields in a Centrifugal Impeller under Nominal and Disturbed Operating Conditions

[1]Donghyuk Kang *

[1]Department of Mechanical Engineering/Saitama University Shimo-okubo 255, Sakura-ku, Saitama-shi, Saitama 338-8570, Japan 1

Abstract

Unsteady flow phenomena in turbomachinery, such as surge and rotating stall, impose severe operational constraints and may lead to performance degradation, vibration, noise, and structural failure. These instabilities are particularly prominent in low-flow-rate operating regimes and are strongly associated with complex three-dimensional internal flow structures. Accurate characterization of such flow fields is therefore essential not only for understanding instability mechanisms but also for enabling flow-based monitoring and control strategies.

However, direct measurement or high-fidelity numerical simulation of internal three-dimensional unsteady flow fields remains computationally or experimentally prohibitive in practical systems. Full-order computational fluid dynamics (CFD) simulations require millions of degrees of freedom and extensive computational time, while direct experimental access to internal flow structures is severely limited by sensor placement constraints and spatial resolution issues. Consequently, a framework that enables accurate and rapid reconstruction of turbomachinery flow fields from limited measurements is of substantial importance.

This study proposes a unified data-driven reconstruction framework for turbomachinery flow fields under steady, unsteady, and externally forced operating conditions. The approach integrates reduced-order modeling (ROM), proper orthogonal decomposition (POD), sparse sensing, Bayesian estimation, and dynamic mode decomposition with control (DMDc) to achieve accurate flow reconstruction with drastically reduced computational cost.

Keywords: Proper orthogonal decomposition, Bayesian estimation, Dynamic mode decomposition with control

Data-Driven Analysis of Cavitation Erosion Geometry

[1]Szu Yung Chen*; [1]Izuru Kambayashi; [1]Ning Yang; [1]Motohiko Nohmi; [1]Ryota Yoshihara

[1]Ebara Corporation

Abstract

The prediction and prevention of cavitation erosion are crucial for the reliable design and maintenance of modern industrial systems. While recent advancements in computational fluid dynamics (CFD) have improved erosion risk prediction technologies, forecasting the actual erosion geometry remains highly challenging due to its complex, multi-scale nature and the intricate fluid-material interactions involved. Consequently, relying on the accumulation and thorough analysis of experimental data remains essential for understanding these mechanisms.

In this study, a data-scientific approach was employed to analyze the geometric evolution of cavitation erosion on a pure aluminum test specimen. The specimen was subjected to a 150-hour cavitation erosion test, and surface shape changes were periodically measured using a 3D scanner. These detailed time-series displacement data were then analyzed using Proper Orthogonal Decomposition (POD). This method aims to generate a Reduced Order Model (ROM) capable of representing the complex spatial-temporal growth of erosion, reducing the high computational costs of exhaustive internal material stress calculations

Experimental observations revealed that the erosion initially showed plastic deformation, characterized by the continuous growth of adjacent dimples and protrusions. While the POD analysis successfully reproduced the growth in its primary mode, the higher-order modes exhibited significant nonlinear complexity.

Although POD provides valuable insights into the erosion process, further research is required to refine the ROM to accurately generalize these shape changes. Ultimately, these findings lay the groundwork for integrating localized shape-change models, such as cellular automata, into fluid analyses to streamline and enhance industrial cavitation erosion predictions.

Keywords: Cavitation, Erosion, Data Science, Proper Orthogonal Decomposition, Reduced Order Model

Multi-fidelity Aerodynamic Transfer Modeling Considering Heterogeneous Shapes Based on Point Cloud Variational Autoencoder for Efficient Propeller Shape Design

[1]Xiaojing Wu*; [1]Zijun Zuo 2; [2]Weiwei Zhang 3;

[1]*Institute of Al for Industries,Chinese Academy of Sciences, Nanjing, 211135, China*

[2]*School of aeronautics, Northwestern Polytechnical University, Xi'an, 710072, China*

Abstract

The complexity of propeller geometry and flow makes aerodynamic modeling a high-dimensional and computationally expensive problem. Multi-fidelity data fusion modeling is an effective approach for achieving efficient aerodynamic modeling of propellers. Due to the inability of low-fidelity Blade Element Momentum Theory (BEMT) to account for the continuous variation of cross-sectional airfoils, the geometric heterogeneity poses challenges for multi-fidelity data fusion modeling. To address this issue, this paper proposes a heterogeneous shape aerodynamic multi-fidelity transfer modeling method for propellers from the perspective of point cloud. By learning the propeller shape point clouds represented by different parameters with Point Cloud Variational Auto-encoder (PCVAE), a unified low-dimensional representation of the propeller's heterogeneous shape is achieved. In the unified latent feature space, multi-fidelity aerodynamic transfer modeling is realized based on a Point Cloud-Transfer Learning Deep Neural Network (PC-TLDNN), which significantly reduces the reliance on high-fidelity data. On this basis, an optimization framework for transfer modeling is constructed. The effective-ness of the proposed method is verified through a specific propeller case study. The results demonstrate that the method achieves high-fidelity aerodynamic modeling and optimization of propellers with full parameters under small sample, high-fidelity data conditions.

Keywords: Aerodynamic modeling; Geometric parameterization; Point cloud; Variational autoencoder; Feature extraction; Optimization;

Proceedings of the 2nd international Symposium on AI and Fluid Mechanics
Paper No [leave blank]

Physics-Guided Operator Learning for Integrated Aerodynamic Optimization of Transonic Fan Blades

[1]Xinyue LAN; [1]Sun GANG*;

[1]*Department of Aeronautics & Astronautics, Fudan University, Shanghai 200433, China*

Keywords: Transonic fan blades; Tip leakage vortex; Multi-objective optimization; Surface bump; Shock control

Abstract. In transonic fan blades, shock waves and complex gap flows are critical factors leading to aerodynamic performance loss. Conventional bump-based optimization primarily targets steady performance metrics, offering limited control over unsteady dynamics. To overcome this limitation, a near-tip surface morphing framework is proposed, introducing localized bumps to actively regulate shock–tip flow interactions and associated unsteady phenomena. Within this framework, operator learning constructs a high-dimensional mapping from bump geometries to flow responses, capturing complex geometry–flow couplings and enabling efficient exploration of the high-dimensional design space. An embedded attention mechanism adaptively identifies and emphasizes geometric features most influential to shock intensity, tip leakage vortex evolution, and separation onset, enhancing sensitivity to critical flow-driving structures. Coordinated multi-objective optimization within this architecture redistributes strong normal shocks into weaker oblique shocks, stabilizes tip leakage vortices, delays flow separation, and suppresses leading-edge unsteady fluctuations. Notable improvements in aerodynamic efficiency and surge margin are achieved, while maintaining similar pressure ratio and flow rate, particularly under high-speed, high-pressure-ratio conditions. These results establish a systematic, physics-informed framework for local geometric refinement and simultaneous regulation of steady and unsteady aerodynamic behaviours, offering a new pathway for high-efficiency transonic fan blade design.

1. Introduction

High-performance transonic fan blades are widely used in modern aero-engines because of their capability to provide high pressure ratio and compact structural configuration[1-3]. However, under transonic operating conditions, strong shock waves, tip leakage flow, and complex vortex interactions frequently occur near the blade tip region[4, 5]. These flow phenomena produce considerable aerodynamic loss and often lead to flow instability, reduced efficiency, and deterioration of compressor operating margin[6-8]. Previous studies have shown that the evolution of tip leakage vortices is strongly influenced by blade loading, pressure distribution, and shock-wave structures[9]. Increasing flow rate and pressure ratio usually strengthen the pressure difference across the blade tip and intensify leakage-flow transport. In addition, interactions between leakage vortices and passage shocks can significantly affect local flow stability and aerodynamic performance[10].

Various passive and active flow-control approaches have been proposed to improve tip-region flow behaviour[11]. Surface bump control has attracted considerable attention because localized geometric perturbations can modify near-wall pressure distribution and regulate shock-wave structures without introducing complex mechanical systems[12]. Most existing bump-based optimization studies mainly focus on steady aerodynamic indicators such as efficiency or pressure ratio. Their capability in controlling unsteady vortex evolution and coupled shock–vortex interactions remain limited. At the same time, recent developments in machine learning provide new opportunities for high-dimensional aerodynamic optimization[13]. Conventional surrogate models often struggle to accurately capture highly nonlinear geometry–flow relationships in transonic turbomachinery environments[14]. Operator learning methods offer improved capability in learning functional mappings between geometric configurations and flow responses[15]. Combined with physics-guided mechanisms, operator learning can efficiently identify critical flow-driving structures and improve optimization robustness[16].

In the present work, a physics-guided operator learning framework is proposed for integrated aerodynamic optimization of transonic fan blades. Localized surface bumps are introduced near the blade tip region to regulate shock-wave structures, leakage-flow evolution, and tip separation behaviour. The study focuses on coordinated regulation of both steady and unsteady aerodynamic phenomena under transonic operating conditions.

2. Numerical simulation

Three-dimensional numerical simulations were conducted for a transonic fan rotor. The governing equations were based on the unsteady Reynolds-averaged Navier–Stokes equations (URANS). The Spalart–Allmaras turbulence model[17] was employed for the unsteady aerodynamic calculations. The computational domain consisted of a single blade passage representing 1/18 of the full annulus, with periodic boundary conditions applied in the circumferential direction. The computational grid consisted of approximately 2.08 million cells, with 109 grid points in the radial direction. The first near-wall grid height was set to 5×10^{-6} m to ensure sufficient near-wall resolution. For the URANS simulations, the inlet boundary conditions corresponded to sea-level standard atmospheric conditions, with a total temperature of 288.15 K, a total pressure of 101325 Pa, and axial inflow direction. Under the design operating condition, the fan rotor was operated at a rotational speed of 2400 rpm, and the outlet static pressure was set to 101325 Pa. As illustrated in figure 1, the initial fan blade exhibited strong passage shock structures within the 80%–95% span region. Relative Mach number contours showed that the flow upstream of the shock remained supersonic, while downstream regions rapidly decelerated into subsonic flow. The shock wave imposed a strong adverse pressure gradient on the blade surface boundary layer. As the shock interacted with the near-wall viscous region, the boundary-layer thickness increased significantly downstream of the shock location. This thickened boundary layer became increasingly sensitive to downstream pressure recovery, promoting flow separation near the trailing edge region.

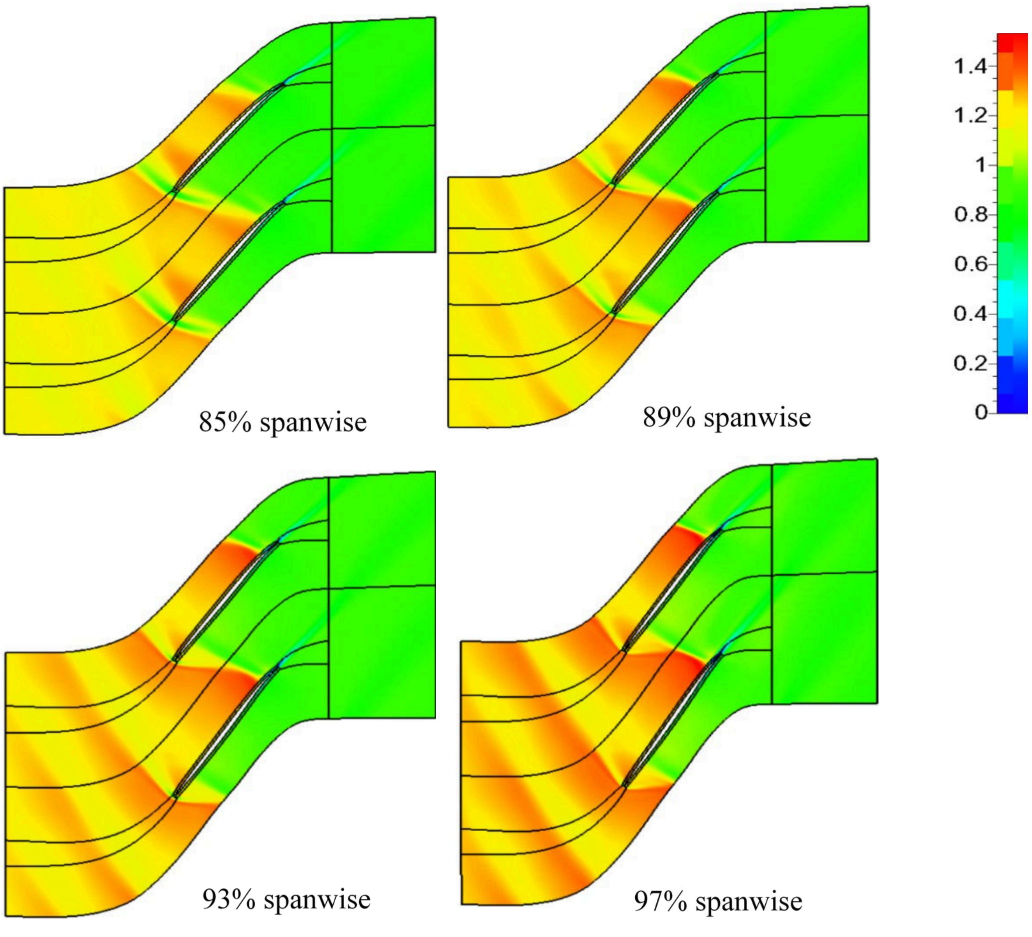

Figure 1. Relative Mach number contours at 85%–97% spanwise sections.

For detailed analysis of shock-wave and boundary-layer interaction mechanisms, unsteady simulations were further performed on a representative blade cascade extracted from the 95% span section. At the cascade inlet, the static pressure was 78400 Pa, the total temperature was 293 K, the inlet Mach number was 1.25, and the inlet flow angle was 60.73°. At the outlet, the static pressure was set to 116000 Pa. The steady Reynolds-averaged Navier–Stokes (RANS) simulations were first conducted using the k–ω SST turbulence model[18]. Based on the converged RANS solution, large-eddy simulations (LES) were then performed. The Smagorinsky–Lilly subgrid-scale model[19] was employed for LES. The LES computational grid contained approximately 400,000 nodes, with a first-layer wall spacing of 1×10^{-6} m. The grid consisted of a single layer in the spanwise direction. The LES results further revealed that the pressure-side shock interaction generated stronger boundary-layer growth than the suction-side shock. On the pressure surface, the shock appeared around 47% chord location and induced a substantial increase in local viscous thickness. In contrast, the suction-side shock occurred closer to the trailing edge and exhibited a weaker influence on boundary-layer development. These results indicate that the pressure-side shock interaction region is the dominant source of aerodynamic loss in the near-tip flow field and therefore becomes the primary target for local geometric control.

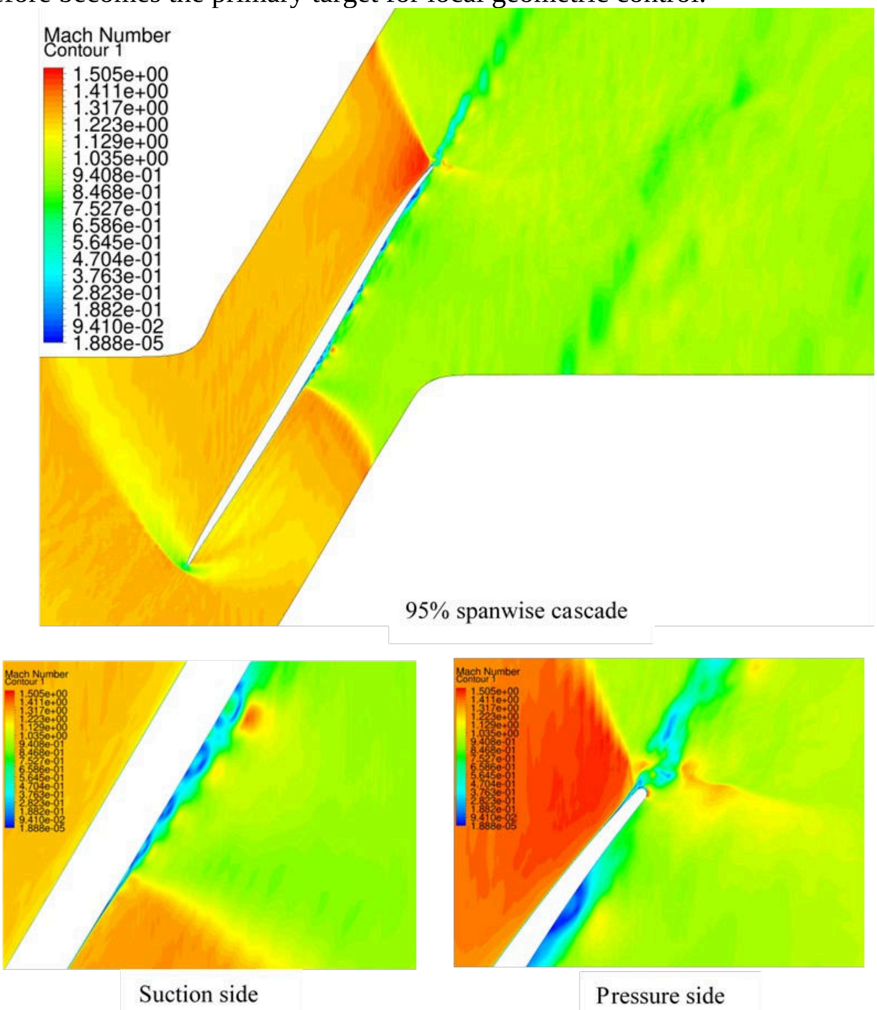

Figure 2. Unsteady LES Mach number contours of blade cascade at 95% span location.

3. Near-Tip Surface Morphing Strategy

To regulate the shock structure and suppress shock-induced separation, localized surface morphing based on shock control bumps was introduced near the pressure-side shock interaction region. The bump geometry was parameterized using the Hicks–Henne method. The bump deformation can be expressed as:

$$b_i(x) = \alpha_i \sin^w(\pi x^{\ln 0.5/\ln x_i}) \tag{1}$$

where α_i controls bump amplitude, x_i specifies bump location, and w determines bump width. Three bump amplitudes were initially investigated, corresponding to 0.003, 0.005, and 0.01 times the blade chord length. The bump center was located near the pressure-side shock position at 47% chord. Initial investigations were conducted to evaluate the effects of three bump amplitudes, namely 0.003, 0.005, and 0.01 times the blade chord length. The results indicate that the most significant aerodynamic improvement is achieved when the bump is centered at the pressure-side shock location (47% chord) with an amplitude of 0.005 chord length, leading to a reduction in the total pressure loss coefficient of approximately 1.56% relative to the baseline cascade.

Table 1. Comparison of bump effects.

	Baseline (No bump)	Bump 1	Bump 2	Bump 3
Bump height (chord length)		0.003	0.005	0.01
Total pressure loss coefficient	0.08126	0.08026	0.07999	0.08061
Relative change		-1.24%	-1.56%	-0.81%

A physics-informed latent representation strategy was developed, where isentropic Mach number distributions were employed to characterize steady shock structures. Four representative spanwise sections at 80%, 85%, 90%, and 95% blade height were selected, and 125 chordwise sampling points were extracted for each section, resulting in a 500-dimensional aerodynamic representation space. An attention-enhanced autoencoder was introduced for dimensionality reduction and feature extraction. The autoencoder compresses each 125-dimensional sectional Mach number distribution into a 10-dimensional latent vector. The embedded attention mechanism adaptively emphasizes regions associated with shock fronts and pressure recovery gradients, enabling the latent space to preserve essential flow features, including shock position, shock strength, and near-tip pressure recovery behavior. Based on this representation, a DeepONet-based surrogate model was constructed to learn the nonlinear operator mapping between latent variables and aerodynamic responses, thereby avoiding repeated high-fidelity CFD simulations. The model accounts for both steady and unsteady aerodynamic characteristics, where the steady shock structure is represented by isentropic Mach number distributions, while wake unsteadiness is quantified using the root-mean-square (RMS) of velocity fluctuations in the downstream region, capturing wake vortex instability and unsteady mixing intensity.

To further improve optimization efficiency, gradient information extracted from the DeepONet surrogate model was incorporated into the evolutionary optimization process. Within the NSGA-III framework, candidate solutions were first evaluated using the surrogate model and then updated along gradient directions associated with improved aerodynamic efficiency prior to the selection operation. The optimization was performed using a population size of 50 individuals over 40 generations.

4. Results

Comparisons of isentropic Mach number distributions demonstrated that the optimized blade delayed the shock position and reduced shock intensity within the 80%–95% span region. Outside the shock-interaction region, the optimized and initial blades exhibited nearly identical aerodynamic distributions, indicating that the localized geometric perturbation primarily influenced shock structures without significantly altering the overall blade loading distribution. Static pressure contours on the suction surfaces of the baseline and optimized fan blades are shown in figure 4. It can be observed that the originally strong passage shock near the blade tip was interrupted and redistributed after optimization, resulting in an overall reduction in shock intensity. The optimized blade configuration exhibited significant flow-improvement characteristics. The oblique shock structures induced by the localized bump effectively moderated the incoming flow acceleration and reduced the abrupt pressure rise across the shock. The modified pressure distribution further weakened the tip leakage vortex and stabilized its spatial evolution, preventing premature vortex breakdown within the blade passage. As a result, leading-edge unsteady fluctuations were substantially suppressed, leading to a more stable transonic flow field. Further analysis of the blade surface boundary-layer behavior revealed that the optimized geometric deformation not only improved the shock structure but also favorably influenced the laminar-to-turbulent transition process. By regulating the local favorable pressure gradient, the optimized blade delayed the premature onset of transition and enlarged the laminar-flow region near the blade surface. This reduction in skin-friction loss provided an additional physical mechanism contributing to the observed improvement in aerodynamic efficiency.

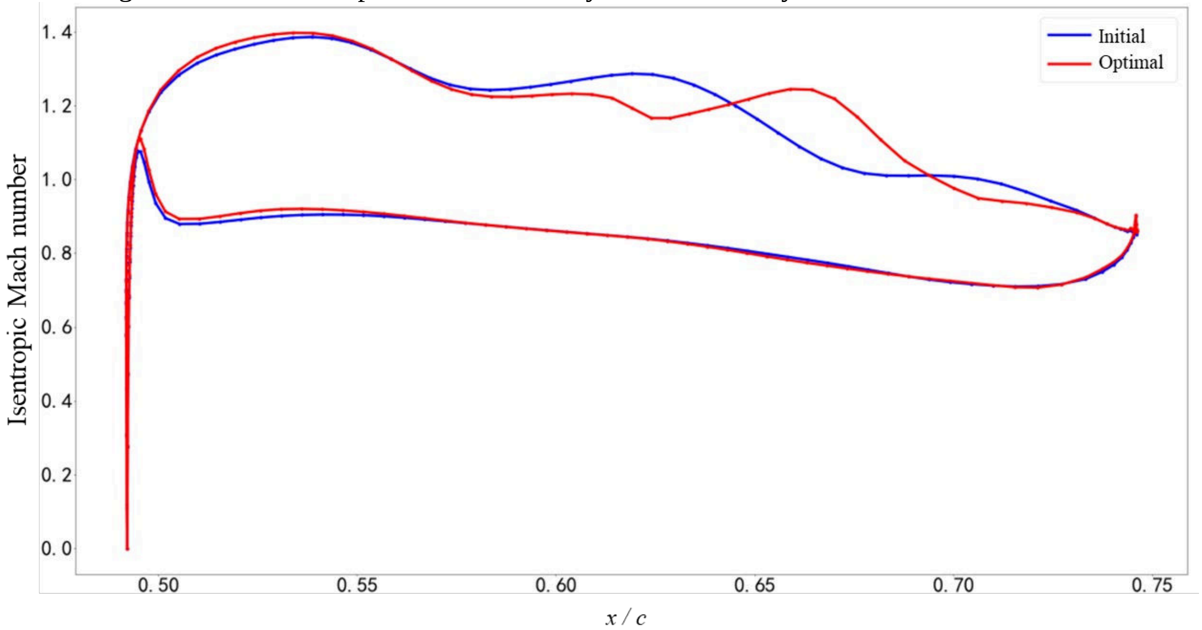

Figure 3. Isentropic Mach number distributions at 85% span for the initial and optimized blades.

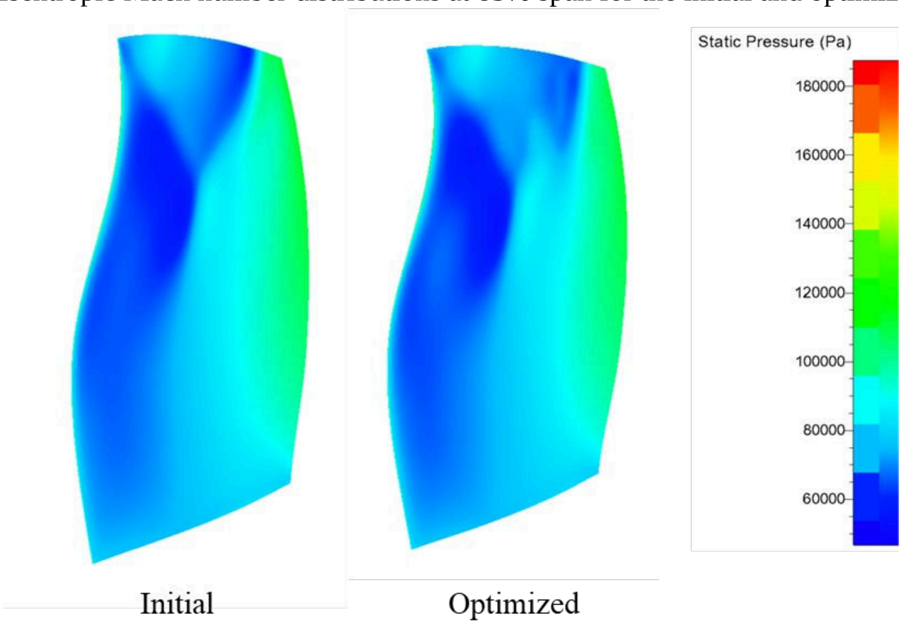

Figure 4. Suction-surface static pressure distributions of the initial and optimized blades.

The optimized blade achieved an efficiency improvement of 0.81 percentage points within the 80%–95% span region. Meanwhile, mass flow rate increased by 0.44%, and pressure ratio changed by only −0.23%.

Table 2. Performance of the initial and shock-control blades.

Blade configuration	Single-passage mass flow rate (kg/s)	Total pressure ratio	Isentropic efficiency (%)
Initial blade	72.50	1.497	91.74
Shock-control blade	72.82	1.494	92.55

5. Conclusion

A data-driven near-tip surface morphing framework was developed for transonic fan optimization. Localized shock-control bumps were introduced near the pressure-side shock interaction region, and an aerodynamic latent-space optimization framework combining attention-enhanced autoencoders, DeepONet surrogate models, and gradient-assisted NSGA-III optimization was established. The optimized blade achieved significant aerodynamic performance improvement under high-speed and high-pressure-ratio conditions, confirming the effectiveness of localized shock-control strategies for transonic fan applications.

6. References

[1] G. Karadimas, Design of high-performance fans using advanced aerodynamics codes, (1988).
[2] S. Sun, S. Wang, L. Zhang, L. Ji, Design and performance analysis of a two-stage transonic low-reaction counter-rotating aspirated fan/compressor with inlet counter-swirl, Aerospace Science and Technology, 111 (2021) 106519.
[3] X. Tang, J. Luo, F. Liu, Aerodynamic shape optimization of a transonic fan by an adjoint-response surface method, Aerospace Science and Technology, 68 (2017) 26–36.
[4] L. Xu, Shockwave and noise abatement of transonic fans, Turbo Expo: Power for Land, Sea, and Air, 2004, pp. 1715–1722.
[5] J. Du, F. Lin, H. Zhang, J. Chen, Numerical investigation on the self-induced unsteadiness in tip leakage flow for a transonic fan rotor, (2010).
[6] X. Jia, X. Dong, Q. Zhang, Y. Zhang, X. Lu, Mechanisms and evolution of flutter stability for a transonic fan under linear reynolds number variations, Aerospace Science and Technology, (2025) 111264.
[7] J. Luo, Z. Fu, Y. Zhang, W. Fu, J. Chen, Aerodynamic optimization of a transonic fan rotor by blade sweeping using adaptive Gaussian process, Aerospace Science and Technology, 137 (2023) 108255.
[8] Q. Zhang, D. Meng, S. Xu, D. Wang, X. Hang, Efficient aerodynamic uncertainty quantification for a three-dimensional transonic fan rotor, Journal of Turbomachinery, 147 (2025) 081013.
[9] Z. Li, X. Zhu, Z. Yan, T. Pan, X. Zheng, Numerical study on the Re effects on the tip flow structures of transonic fan under inlet distortion, Aerospace Science and Technology, 141 (2023) 108530.
[10] G. An, R. Zhou, B. Liu, X. Yu, G. Wu, X. Nan, Effects of the loading coefficient on the adjustable operating range of a single-stage transonic fan, Physics of Fluids, 37 (2025).
[11] S. Lee, Y. Zhao, J. Luo, J. Zou, J. Zhang, Y. Zheng, Y. Zhang, A review of flow control strategies for supersonic/hypersonic fluid dynamics, Aerospace Research Communications, 2 (2024) 13149.
[12] D. Di Pasquale, S. Prince, Passive transonic shock control on bump flow for wing buffet suppression, Aerospace, 10 (2023) 569.
[13] D.I. Lopez, T. Ghisu, S. Shahpar, Global optimization of a transonic fan blade through AI-enabled active subspaces, Journal of Turbomachinery, 144 (2022) 011013.
[14] X. Yang, G. Sun, Y. Jin, A. Ju, Sensitivity analysis and ANN-based surrogate optimization for transonic fan rotor: Multi-condition aerodynamic design incorporating fan-intake interaction, Aerospace Science and Technology, (2025) 111488.
[15] X. Liu, S. Yang, H. Sun, Z. Wang, X. Guan, Y. Gu, Y. Wang, Review of deep learning-based aerodynamic shape surrogate models and optimization for airfoils and blade profiles, Physics of Fluids, 37 (2025).
[16] M. Rauseo, F. Zhao, M. Vahdati, Physics guided machine learning modelling of compressor stall flutter, Journal of the Global Power and Propulsion Society, 8 (2024) 295–309.
[17] A. Bueno-Orovio, C. Castro, F. Palacios, E. Zuazua, Continuous adjoint approach for the Spalart-Allmaras model in aerodynamic optimization, AIAA journal, 50 (2012) 631–646.
[18] F.R. Menter, Two-equation eddy-viscosity turbulence models for engineering applications, AIAA Journal, 32 (1994) 1598–1605.
[19] K.E. Lilly, The representation of small-scale turbulence in numerical simulation experiments, 1966.

IRMA: Impeller Representation Model for Aerodynamics

Alessandro Pela[1], **Michele Marconcini**[1], **Andrea Arnone**[1], **Andrea Agnolucci**[2], **Marco Bicchi**[2], **Elisabetta Belardini**[2], **Roberto Valente**[2], **Angelo Grimaldi**[2]

[1]Dept. of Industrial Engineering, via Santa Marta 3, University of Florence, Firenze, 50139, Italy

[2]Baker Hughes, via Felice Matteucci 2, Firenze, 50127, Italy

Keywords: CFD, CNN, Machine Learning, Centrifugal Compressors

Abstract Accurate aerodynamic modeling requires representations capable of capturing both the smooth spatial variations of blade geometry and the complex behavior of flowfields. Impeller geometry varies continuously along chordwise and spanwise directions and contains detailed features that significantly affect performance. Flowfields, in contrast, can exhibit sharp gradients and localized discontinuities, such as shock waves in transonic conditions, which must be resolved to accurately represent aerodynamic behavior. Effectively linking geometric features with resulting aerodynamic response is therefore essential for reliable surrogate modeling and design exploration. This work introduces IRMA (Impeller Representation Model for Aerodynamics), a multimodal neural network that learns a shared latent representation capturing the mapping between three-dimensional impeller geometry and its aerodynamic response. The model simultaneously encodes discrete geometric parameters and image-based representations of blade-angle distributions, from which coupled decoders predict both CFD-derived flowfields and global performance metrics. By explicitly modeling the relationship between geometry and aerodynamic outputs, IRMA provides a compact and expressive surrogate capable of linking continuous design features to both spatially distributed and integral aerodynamic quantities, enabling efficient analysis and optimization of centrifugal impeller geometries.

1 Introduction

Radial machines play a key role in many energy conversion applications. Deep learning techniques have recently enabled data-driven design and optimization of turbomachinery stages. Most existing approaches rely on supervised learning methods, particularly neural networks, to establish relationships between geometric input parameters and global stage performance metrics [1–3] by leveraging CFD-based databases.

In a previous study by the authors [4], it was shown that such approaches can be extended beyond global quantities to include local flow features, such as blade loading. This is particularly valuable from an engineering perspective, as it allows local flow information to be directly incorporated into the optimization process, enabling the imposition of specific loading characteristics as design constraints. Additional examples of flow field prediction using convolutional neural networks in turbomachinery applications are reported in [5, 6].

This work extends previous research by addressing the generalization of geometric inputs. Although surrogate models exhibit strong predictive capabilities, geometry representation remains a key limitation, as no standard formulation exists. Most approaches depend on specific parameterizations, with model accuracy closely tied to this choice.

In this context, several strategies have been proposed for geometric parameterization. The most widely adopted relies on discrete parameter sets describing local features of the blade and meridional channel [3, 7]. These parameters are typically selected based on engineering knowledge, targeting quantities with the strongest impact on aerodynamic performance.

Alternative approaches rely on constructing new geometries through linear combinations of reference shapes available in literature or existing databases, as proposed in [8]. This strategy enables the exploration of a relatively broad design space by incorporating significantly different baseline configurations. However, the generated database inherently retains the geometric features of the original reference shapes, limiting the possibility of performing fully general perturbations of the design space.

A further relevant class of methods is based on Free-Form Deformation (FFD), as discussed in [9–11]. In this framework, the geometry is embedded within a lattice of control points, and modifications are introduced by displacing these points, with smooth deformations propagated through interpolation functions. This allows the generation of shape variations while preserving geometric consistency. Nevertheless, these approaches remain fundamentally perturbative, as they deform an existing baseline geometry, and the achievable design space may still be influenced by the initial reference shape.

More recent approaches aim to overcome explicit parameterization altogether by leveraging advanced neural architectures, such as generative adversarial networks and variational autoencoders, to learn geometric representations directly from datasets [12]. Although these methods offer high flexibility, the lack of explicit engineering parameters can reduce their practical usability within design workflows, where interpretability and direct control over physically meaningful quantities remain essential.

This work introduces a hybrid input representation that overcomes conventional parameterizations by combining geometric generality with engineering interpretability, enabling more flexible turbomachinery optimization.

2 Numerical Models

In this work, a multimodal neural network is employed to handle heterogeneous input data. The model combines discrete parameters describing the meridional channel with distributed inputs representing the blade angle of an impeller along both chordwise and spanwise directions. These inputs are jointly encoded into an N-dimensional latent space, which is subsequently used to predict both global performance quantities and spatial distributions of blade loading on the pressure and suction sides. The development of the model relies on a database generated through CFD simulations. This dataset constitutes the basis for the training and validation of the neural network. The compressor geometries included in the database are generated using an in-house parameterization tool, shown in Figure 1.

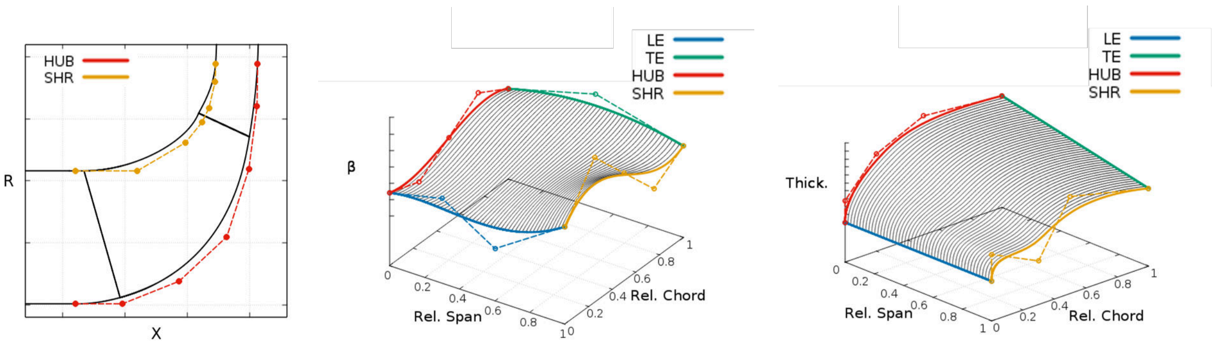

(a) Meridional channel (b) Blade angle (c) Blade thickness

Figure 1: Impeller geometry parameterization

The tool is used to generate the full three-dimensional geometry of the centrifugal compressor. The meridional channel is defined through Bézier control points, which govern the curvature at both hub and tip. Once the channel is established, the blade sections at hub and tip are constructed based on chordwise distributions of blade angle and thickness. As illustrated in Figure 1, Bézier control points are employed both to control the meridional channel geometry (1a) and to define the boundary values of the blade angle (1b) and thickness (1c) distributions. In the present work, the thickness distribution is kept fixed; however, the same framework could be extended to include thickness variations by introducing additional perturbations. The database is generated starting from a reference geometry, which is systematically modified by applying perturbations to the Bézier control points, enabling the exploration of a range of geometrical configurations.

A structured viscous H-type mesh, validated through a prior sensitivity analysis [13], is adopted (see Figure 2). Simulations assume circumferential periodicity, modeling a single blade passage.

The CFD simulations are performed using the TRAF code, a three-dimensional viscous RANS solver capable of both steady and unsteady analyses, developed at the University of Florence [14]. The code has been extensively applied to the investigation of centrifugal compressors in both industrial and turbocharging contexts [15]. A detailed description of the numerical setup adopted in the present work is available in [13].

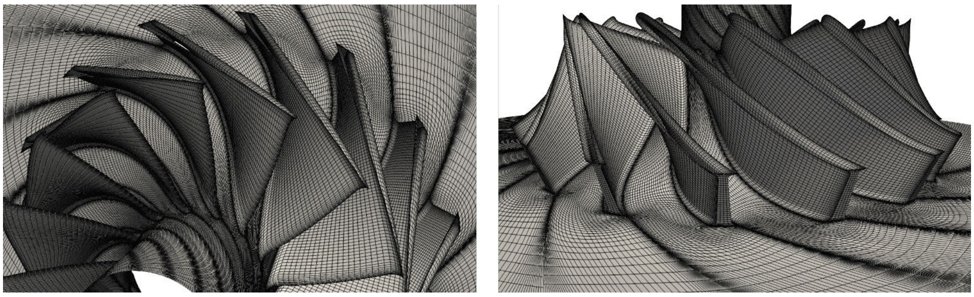

Figure 2: Computational domain

For each generated mesh, a single operating condition is simulated by prescribing the backpressure. The solution is considered converged when variations in the main performance indicators, such as flow coefficient, polytropic efficiency, and work coefficient, fall below 0.5%. Once convergence is achieved, a post-processing procedure is carried out to extract the relevant performance metrics. In particular, the global quantities used to characterize the compressor behavior include the flow coefficient ϕ, work coefficient τ, polytropic efficiency η_{p}, and impeller exit flow angle YAW.

In addition to global performance metrics, blade loading is extracted from the static pressure distribution on both the pressure and suction sides. This information is represented as pressure maps, stored as images over normalized chord-wise and span-wise coordinates, with pixel intensity encoding the local pressure value (Fig. 3a). Further details are provided in [4]. A similar representation is also adopted for the blade angle distribution, which is encoded as an image in a manner analogous to the pressure maps. In this case, the pixel intensity corresponds to the local value of the blade angle, defined over the same chord-wise and span-wise coordinate system (Fig.3b).

This approach preserves engineering interpretability, as local blade angle values remain directly accessible. At the same time, it reduces the dependence of the surrogate model on a specific functional parameterization, allowing it to operate directly on the distributions rather than the predefined design variables. This improves generalization, enabling the integration of datasets generated with different parameterization strategies. Overall, the proposed representation decouples the learning process from the choice of geometric parameterization, enhancing flexibility and extensibility.

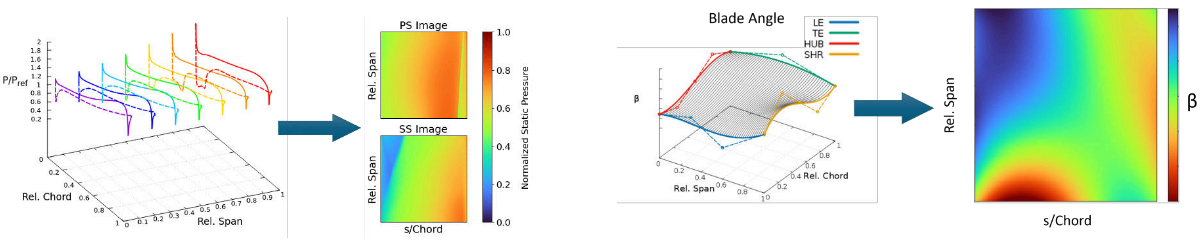

(a) Blade loading

(b) Blade angle

Figure 3: Image conversion of input and output data

The neural network adopted in this work is a multimodal architecture designed to process heterogeneous inputs. It consists of two separate branches. The first processes the geometric variables X, corresponding to the Bézier control points defining the meridional channel, through a fully connected network. The second branch processes the blade angle distribution, provided as a single-channel image, using a convolutional architecture.

The two input streams are independently encoded into a shared N-dimensional latent space. From this representation, the network splits into two output branches. The first branch, based on fully connected layers, predicts the global performance quantities y. The second branch adopts a convolutional decoder to reconstruct the blade loading maps on the pressure and suction sides.

A schematic representation of the network architecture is shown in Figure 4.

The network is trained using the Adam optimization algorithm [16]. Nonlinear activation functions between layers are implemented using the ReLU function. The dataset consists of approximately 2000 CFD-generated

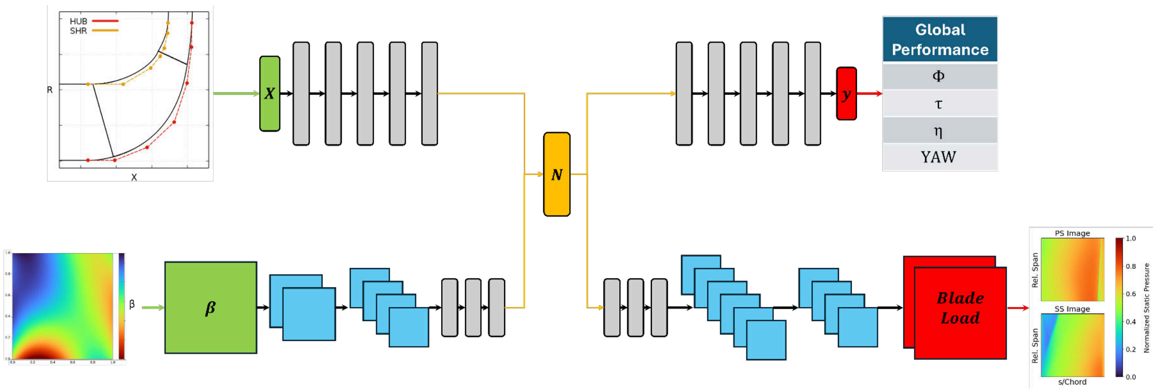

Figure 4: IRMA Neural Network scheme. The color legend is: green= input data, azure= $2D$ data (conv/transpose conv layers), grey= $1D$ data (dense layers), orange= N-dimensional latent space, red=output data

cases. It is splitted into training and test sets using an 80/20 partition to evaluate model generalization.

During training, the formulation of the loss function plays a key role in ensuring a balanced learning process between global performance variables and image-based outputs. The total loss is defined as a weighted sum of two contributions:

$$Loss = MSE_y + \alpha \cdot MSE_{img} \tag{1}$$

where MSE_y represents the mean squared error associated with the global performance vector y, and MSE_{img} denotes the pixel-wise mean squared error between predicted and reference blade loading images. The parameter α acts as a scaling factor to balance the relative contribution of the two loss terms, ensuring comparable orders of magnitude. An excessively large value of α may bias the training toward image reconstruction at the expense of accurate prediction of global quantities, whereas an overly small value may lead to poor learning of the image-based outputs.

3 Training Results

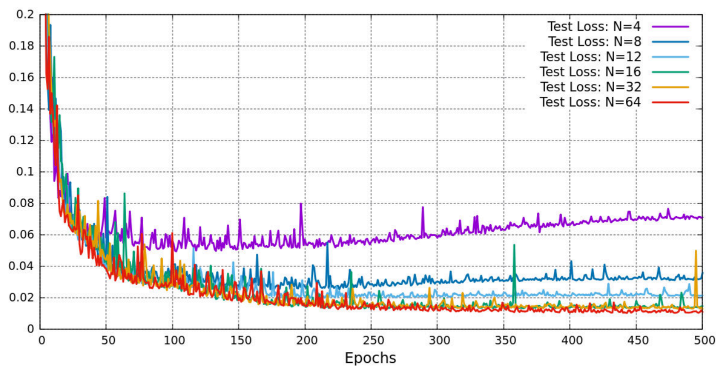

Figure 5: Test losses comparison for different latent space dimensions

The training procedure is carried out over 500 epochs. In order to investigate the minimum latent space dimensionality required to adequately represent the dataset, the same architecture is trained using different values of N. The corresponding results, reported in Figure 5, compare the test loss for $N = 4, 8, 12, 16, 32, 64$. The evaluation is based on the test loss in order to assess the model's ability to generalize to unseen data not used during training.

The comparison shows that increasing the latent space dimension beyond a certain threshold does not lead to significant improvements in predictive performance. In particular, a latent space dimension of $N = 16$ (the green line) appears sufficient to capture the information contained in the combined representation of discrete geometric parameters and image-based inputs, as the test error remains comparable to that obtained with higher-dimensional configurations.

Based on this analysis, the model with $N = 16$ is selected as the final architecture. The reported results are therefore obtained considering this configuration and evaluated exclusively on the test set.

The prediction performance of the global outputs is summarized in Figure 6, which shows the test set results for the 4 output variables of the MLP branch. In each plot, the predicted values are compared against the corresponding ground truth values. The majority of the samples lie within a $\pm 1.0\%$ error band, indicating a satisfactory level of accuracy for all considered performance metrics.

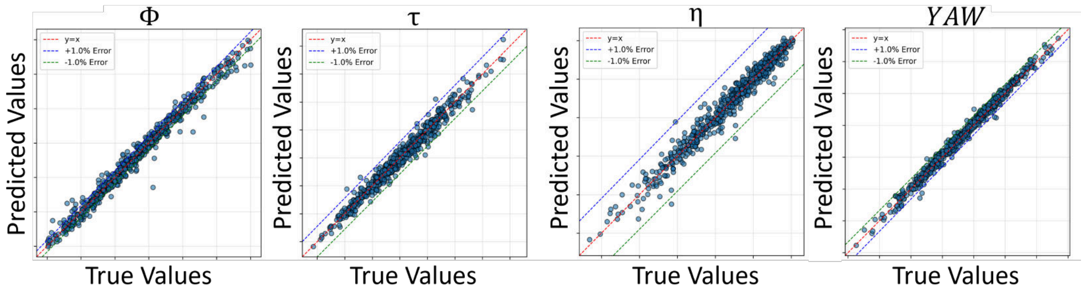

Figure 6: Global performance results of the test set

The prediction of blade loading is illustrated in Figure 7. The decoder performance is evaluated by comparing the predicted and reference pressure maps on both the pressure side (PS) and suction side (SS) of the blades for the test set. In addition, absolute percentage error maps are computed to highlight local discrepancies between predicted and true fields. The blade loading patterns exhibit significant variability across the test dataset; nevertheless, the decoder is able to consistently reproduce the overall pressure distribution with good accuracy on both blade sides. By considering a fixed span-wise position, the chord-wise variation of the pressure difference between PS and SS can be further analyzed. In transonic operating conditions, strong gradients may appear due to the presence of shock waves, which locally increase the complexity of the flow field. Despite these challenging regions, the model maintains a robust predictive capability, with only limited deviations from the reference solutions.

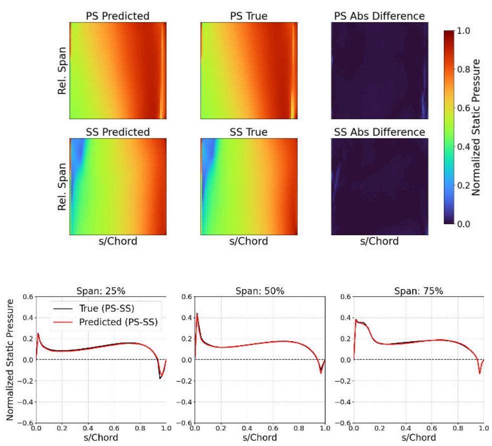

(a) Positive incidence case

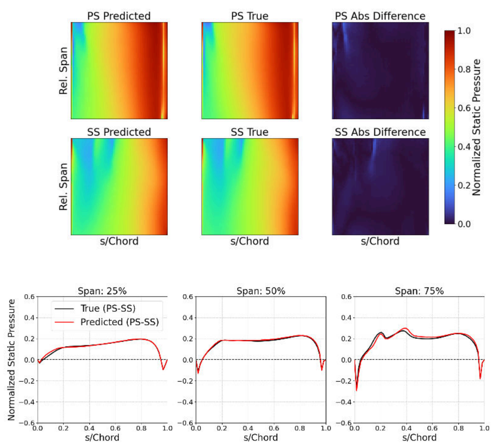

(b) Negative incidence case

Figure 7: Blade Loading results of the test set

4 Conclusions

In the present work, a neural network is proposed and trained on CFD data of centrifugal impellers. The approach addresses the limitations of classical geometric parameterizations, which typically restrict the design space to a fixed number of degrees of freedom. Instead, the model directly uses a continuous representation of the blade geometry, here expressed through the blade angle distribution in spatial form. The results show that the geometric description can be effectively compressed into a 16-dimensional latent space without significant loss of predictive accuracy. Within this representation, the learned features jointly encode both local blade loading and global performance parameters of the impeller.The methodology is general and can be extended to other turbomachinery configurations, offering a flexible surrogate modeling strategy.

References

[1] Hehn, A., Mosdzien, M., Grates, D. and Jeschke, P. "Aerodynamic optimization of a transonic centrifugal compressor by using arbitrary blade surfaces." *ASME J. Turbomach.*, Vol. 140. 5. 2018. doi:10.1115/1.4038908.

[2] Agnolucci, A., Marconcini, M., Arnone, A., Toni, L., Grimaldi, A. and Giachi, M. "Centrifugal compressor stage efficiency and rotor stiffness augmentation via artificial neural networks." *Turbo Expo: Power for Land, Sea, and Air*, Vol. 2D: p. V02DT37A014. 2021. doi:10.1115/GT2021-59998.

[3] Pela, A., Marconcini, M., Arnone, A., Agnolucci, A., Toni, L., Valente, R. and Belardini, E. "Exploiting Deep Learning for the Optimization of Transonic Centrifugal Impellers." *J. Phys.: Conf. Ser, The 79th ATI Annual Congress, Genoa, Italy*, Vol. 2893: p. 012126. 2024. doi:10.1088/1742-6596/2893/1/012126.

[4] Pela, A, Marconcini, M, Agnolucci, A, Bicchi, M, Belardini, E, Grimaldi, A, Toni, L, Valente, R and Arnone, A. "Convolutional Neural Network Approach for Impeller Blade Loading Inference." *Computer and Fluids* (2026) under review.

[5] Bruni, G., Maleki, S. and Krishnababu, S. "C(NN)FD - A Deep Learning Framework for Turbomachinery CFD Analysis." *IEEE Trans. Ind. Inform.* (2024)doi:10.1109/TII.2024.3393140.

[6] Bruni, G., Maleki, S. and Krishnababu, S. "C(NN)FD - Deep Learning Modelling of Multi-Stage Axial Compressors Aerodynamics." *Proc. of the ASME Turbo Expo*: p. V012T34A002. 2025. doi:10.1115/GT2025-151295.

[7] Casey, M. V. "A Computational Geometry for the Blades and Internal Flow Channels of Centrifugal Compressors." *J. Eng. Power* (1983)doi:10.1115/1.3227414.

[8] Fesquet, J., Bauerheim, M., Rojda, L., Bousquet, Y. and Binder, N. "Application of Deep Learning for Fan Rotor Blade Performance Prediction in Turbomachinery." *ASME J Turbomach* Vol. 147 (2025): p. 111002. doi:10.1115/1.4068438.

[9] Adjei, R. A., Wang, W. and Liu, Y. "Aerodynamic Design Optimization of an Axial Flow Compressor Stator Using Parameterized Free-Form Deformation." *J. Eng. Gas Turbines Power* Vol. 141 No. 10 (2019): p. 101015. doi:10.1115/1.4044692.

[10] Li, X., Meng, T., Li, W., Zhou, L. and Ji, L. "Integrated passage design based on extended free-form deformation and adjoint optimization." *Chin. J. Aeronaut.* (2023)doi:10.1016/j.cja.2022.11.021.

[11] Wang, Y., Sun, Z., Liu, J., Liu, M. and Zhou, Y. "Optimization design of centrifugal impeller based on Bezier surface and FFD space grid parameterization." *PLOS ONE* (2024)doi:10.1371/journal.pone.0310792.

[12] Wang, Y., Shimada, Kenji and Barati Farimani, Amir. "Airfoil GAN: Encoding and synthesizing airfoils for aerodynamic shape optimization." *J. Comput. Des. Eng.* (2023)doi:10.1093/jcde/qwad046.

[13] Pela, A., Marconcini, M., Arnone, A., Toni, L., Valente, R., Agnolucci, A., Grimaldi, A. and Pacciani, R. "On the Feasibility of Helicity-Based Corrections of Turbulence Models for Improving RANS Predictions of Centrifugal Compressor Stages." *Proc. of the ASME Turbo Expo*, Vol. 13D: p. V13DT35A026. 2023. doi:10.1115/GT2023-103852.

[14] Arnone, A. "Viscous analysis of three-dimensional rotor flow using a multigrid method." *ASME J. Turbomach.* Vol. 116 (1994): pp. 435–445. doi:10.1115/1.2929430.

[15] Marconcini, M., Rubechini, F., Arnone, A. and Ibaraki, S. "Numerical investigation of a transonic centrifugal compressor." *ASME J. Turbomach.* Vol. 130 No. 1 (2008): p. 011010. doi:10.1115/1.2752186.

[16] Kingma, D. P. and Ba, J. "Adam: A Method for Stochastic Optimization." *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*. 2015. San Diego, CA, USA. doi:10.48550/arXiv.1412.6980.

AI-Driven Prediction of Pressure Loss assisting Draft Tube Design in Hydraulic Turbines

[2]Sobieczky Florian; [1]Kontoleontos Evgenia*; [2]Roßbory Michael

[1] *Andritz Hydro GmbH*

[2] *Software Competence Center Hagenberg GmbH*

Abstract

This paper presents our research on hydraulic turbine component design, enabling the evaluation of candidate configurations through a predictive model. The predictive capability enables faster and more informed hydraulic design decisions by providing estimates of component performance prior to a detailed feature analysis. This method consists of training a Deep Neural Network (DNN) on a systematically selected features set on a database of existing draft tubes, which includes their geometrical parameters and corresponding hydraulic behaviour. The prediction of pressure loss is driven by a deliberately balanced selection of features describing the draft tube's geometry as well as its area-expansion diffusion curve characteristics.

Given the technical challenge of a limited size of available training data in a high dimensional feature space, this approach is aiming at a useful balance between a highly descriptive feature set for satisfactory model accuracy and a feasible degree of generalizability of the deep learning model. Using a mixture-of-experts strategy that clusters the data into domains of similar hydraulic performance metrics, the study shows that systematic feature selection can successfully overcome the challenge of the limited size of available training datasets.

Keywords: Draft tube design, AI-driven pressure loss prediction, limited database size, mixture-of-experts