



UNIVERSITÀ
DEGLI STUDI
FIRENZE

DOTTORATO DI RICERCA IN ENERGETICA E
TECNOLOGIE INDUSTRIALI E AMBIENTALI
INNOVATIVE

CICLO XXXVIII

COORDINATOR Prof. Lapo Governi

Deep learning techniques for the design and optimization of centrifugal compressor families

Settore scientifico disciplinare ING-IND/08

Ph.D. Candidate

Alessandro Pela

*Department of Industrial
Engineering, University of
Florence*

Tutor

Prof. Michele Marconcini

*Department of Industrial
Engineering, University of
Florence*

Firenze, 01/11/2022 - 31/10/2025

*"The oldest and strongest emotion of mankind is fear, and
the oldest and strongest kind of fear is fear of the unknown"*

- H.P. Lovecraft

To all those who made me who I am

Abstract

The growing urgency of mitigating global warming has accelerated the demand for innovative solutions in energy conversion systems, directly impacting the design of turbomachinery. In particular, centrifugal compressors require systematic optimization procedures capable of ensuring high performance across a wide range of operating conditions.

Due to the high cost of three-dimensional CFD simulations, optimization must rely on surrogate models that efficiently approximate aerodynamic behavior. The methodology is based on four main steps: (i) the definition of a suitable parameterization of the impeller geometry, (ii) the generation of a numerical database through reliable CFD simulations at selected operating points, (iii) the training of a surrogate model based on artificial intelligence to capture the correlation between geometry and performance, and (iv) the application of an optimization strategy to identify optimal solutions.

This procedure is then extended from single impeller optimization to the generation of entire compressor families. By introducing rules for continuous geometry generation, intermediate designs can be automatically derived while preserving both performance and geometric consistency.

The thesis is organized to provide a clear and comprehensive development of this framework. The first chapter introduces the context of the work, highlighting the challenges and opportunities in high-efficiency turbomachinery design. The following chapter presents the

numerical models employed, including the CFD solver and the validation of the numerical setup, ensuring reliable performance predictions. The subsequent section focus on the artificial intelligence tools used to define the surrogate model. In the fifth chapter a detailed description of the multi-objective optimization strategy applied to identify optimal impeller geometries is presented. Methods to reduce optimization time, including advanced sampling techniques and efficient training procedures, are then discussed in the sixth chapter. Finally, the framework is extended to the generation of compressor families, illustrating a systematic approach for producing intermediate designs with preserved performance and geometric coherence.

This integrated framework provides not only a systematic approach for optimizing individual impellers, but also a scalable strategy for designing entire families of centrifugal compressors. The result is a methodology that combines accuracy, computational efficiency, and flexibility, offering a concrete contribution towards the next generation of high-performance, energy-efficient turbomachinery.

Keywords: CFD, Centrifugal Compressors, Artificial Intelligence, Turbomachinery, Neural Networks

Publications

- A. Pela, M. Marconcini, A. Arnone, L. Toni, R. Valente, A. Agnolucci, A. Grimaldi, and R. Pacciani. “On the feasibility of helicity-based corrections of turbulence models for improving RANS predictions of centrifugal compressor stages”. In: vol. Volume 13D. Turbo Expo. June 2023, V13DT35A026. doi: 10.1115/GT2023-103852
- A. Pela, M. Marconcini, A. Arnone, A. Agnolucci, L. Toni, R. Vaente, and E. Belardini. “Exploiting Deep Learning for the Optimization of Transonic Centrifugal Impellers”. In: J. Phys.: Conf.Ser, The 79th ATI Annual Congress, Genoa, Italy. Vol. 2893.2024, p. 012126. doi: 10.1088/1742-6596/2893/012126.
- A. Pela, M. Marconcini, A. Agnolucci, E. Belardini, A. Grimaldi, L. Toni, R. Valente, and A. Arnone. “Convolutional Neural Network Approach for Impeller Blade Loading Inference”. In: 1st International Symposium on AI and Fluid Mechanics (AIFLU-IDs). May 2025.
- A. Pela, M. Marconcini, A. Agnolucci, E. Belardini, A. Grimaldi, L. Toni, R. Valente, and A. Arnone. “Towards the Automatic Generation of Transonic Centrifugal Impellers by Deep Learning and Feature Importance Techniques”. In: vol. Volume 12. Turbo Expo. June 2025. doi: 10.1115/GT2025-153947.

Contents

Abstract	iii
Publications	v
List of Figures	xi
List of Tables	xiii
Nomenclature	xv
1 Introduction	1
2 Centrifugal Compressor aerodynamics	5
3 Numerical Framework	11
3.1 Geometry parameterization	11
3.2 Mesh	14
3.3 CFD - TRAF code	17
3.4 CFD - helicity correction model	19
3.5 Helicity correction model - validation on centrifugal com- pressor stage	22
3.5.1 Experimental results	23
3.5.2 Compressor 1	26
3.5.3 Compressor 2	33

4	Artificial intelligence and neural networks	41
4.1	MLP code - BaguetteNet	44
4.2	Optimization Code	47
4.3	Example of Optimization procedure using Surrogate model	50
5	Optimization speedup	61
5.1	Database size	62
5.2	Convolutional neural network approach for blade load- ing inference	64
5.3	Geometrical throat area estimation tool for maximum swallowing capacity estimation	71
6	Continuous family generation	77
6.1	Feature importance analysis	80
6.2	Geometry Generation	85
6.3	Generation Results	88
	Conclusions and future work	90
	Acknowledgments	97
	Bibliography	101

List of Figures

1.1	CO_2 Emission Pathways and Global Temperature Projections	1
2.1	Example of centrifugal compressor [11]	6
2.2	Comparison between single-stage and multi-stage compressors [12]	7
2.3	Example of turbocharger [13]	8
2.4	Baker Hughes's High Pressure Ratio Compressor (HPRC) [14]	8
3.1	Impeller Meridional Channel	12
3.2	Return Channel Meridional Channel (multiple spanwise sections)	13
3.3	Chord-wise/span-wise blade angle/thickness distribution	14
3.4	Contours of surface static pressure and vectors of relative velocity close to the suction surface [5]	15
3.5	Computational domain of a impeller with radial inlet and vaneless diffuser	15
3.6	Computational domain of a return channel with U-bend	16
3.7	Test rig layout	24
3.8	Compressor 1, performance comparison in terms of polytropic efficiency	26
3.9	Compressor 1, performance comparison in terms of work coefficient	28

3.10 Compressor 1, performance comparison in terms of polytropic head	28
3.11 Compressor 1, spanwise distribution comparison in terms of normalized pressure ratio (Section 60)	30
3.12 Compressor 1, relative Mach number contour at impeller 85% span at $\phi/\phi_{choke} = 0.806$	32
3.13 Compressor 1, normalized relative total pressure contours at impeller exit (Section 20) at $\phi/\phi_{choke} = 0.806$.	33
3.14 Compressor 1, normalized turbulent kinetic energy production at $\phi/\phi_{choke} = 0.806$	34
3.15 Compressor 2, performance comparison in terms of polytropic efficiency	35
3.16 Compressor 2, performance comparison in terms of work coefficient	35
3.17 Compressor 2, performance comparison in terms of polytropic head	36
3.18 Compressor 2, spanwise distribution comparison in terms of normalized pressure ratio (Section 60)	37
3.19 :	38
3.20 Compressor 2, normalized relative total pressure contours at impeller exit (Section 20) at $\phi/\phi_{choke} = 0.691$.	39
4.1 Machine learning classification	42
4.2 Example of MLP architecture	45
4.3 Example of loss history of train and test datasets . . .	46
4.4 Example of output values comparison between true values and predicted values for each output feature	48
4.5 Iterative optimization procedure scheme	49
4.6 Pareto front (in red) superimposed to the starting database (3 operating points in purple, orange and green)	53
4.7 Performance comparison in terms of polytropic efficiency	54
4.8 Performance comparison in terms of work coefficient . .	54

4.9	Performance comparison in terms of polytropic head coefficient	55
4.10	Impeller blade load comparisons at maximum polytropic efficiency point.	59
4.11	Normalized relative Mach number contours at peak efficiency point at near impeller shroud	60
5.1	Mean Absolute Error value of test set vs database size	63
5.2	Blade loading data collection example	67
5.3	Schematic representation of the decoder neural network	68
5.4	CNN training results for different incidence examples (Test set).	70
5.5	Example of minimum area research. Black area = initial guess. Red Area = final area	75
5.6	Geometrical throat area vs maximum flow coefficient of a transonic impeller database	76
6.1	Design space	78
6.2	Correlation Matrix of Φ_1 database	84
6.3	Geometry generation process scheme for a generic parameter x_i	87
6.4	Performance comparisons	91
6.5	Peak Efficiency Trend Along Input Flow Coefficient . .	92
6.6	Performance comparison with master geometries . . .	92
6.7	Performance comparison using normalized flow coefficient	92
6.8	Impeller blade load comparisons	93
6.9	Normalized relative Mach number contours at peak efficiency point at 80% span	94

List of Tables

3.1	Mesh sensitivity analysis of Compressor 1	23
3.2	Measurement Sections	24

Nomenclature

Latin Symbols

Symbol	Description	Unit
A^*	Throat Area	m^2
D_2	Impeller outlet diameter	m
H	Polytropic head coefficient	-
h	Helicity	-
h_0	Total enthalpy	m^2/s^2
k	Turbulent kinetic energy	m^2/s^2
M_{1r}	Impeller eye relative Mach number	-
Mu_2	Peripheral Mach Number	-
$\dot{m}$	Mass flow Rate	kg/s
P_0	Total pressure	Pa
P_{dyn}	Dynamic Pressure	Pa
P_k	Turbulent kinetic energy production term	m^2/s^3
P_s	Static Pressure	Pa
R	Gas constant	-
T_0	Total Temperature	K
u	Velocity	m/s
u_2	Impeller outlet peripheral speed	m/s

Greek Symbols

Symbol	Description	Unit
α	Learning rate	-
γ	Heat capacity ratio	-
η_p	Polytropic efficiency	-
μ	Dynamic viscosity	$Pa\cdot s$
μ_t	Turbulent viscosity	$Pa\cdot s$
Φ, ϕ, φ	Flow coefficient	-
ρ	Static Density	kg/m^3
ρ_0	Total Density	kg/m^3
τ	Work coefficient	-
ω	Specific turbulent dissipation rate	$1/s$

Acronyms

Acronym	Meaning
AI	Artificial Intelligence
ANN	Artificial Neural Network
CNN	Convolutional Neural Network
CO_2	Carbon dioxide
DL	Deep Learning
DOF	Degree Of Freedom
CFD	Computational Fluid Dynamics
LASSO	Least Absolute Shrinkage and Selection Operator
LES	Large Eddy Simulation
LNG	Liquefied Natural Gas
MAE	Mean Absolute Error
ML	Machine Learning

Acronym	Meaning
MLP	Multi Layer Perceptron
MSE	Mean Squared Error
ORC	Organic Rankine Cycle
PCA	Principal component analysis
PS	Pressure Side
RQEVM	Realizable quadratic eddy viscosity
RANS	Reynolds Averaged Navier Stokes (equations)
ReLU	Rectified Linear Unit
SGD	Stochastic Gradient Descent
SS	Suction Side
SST	Shear Stress Transport
TKE	Turbulent Kinetic Energy
TVD	Total Variation Diminishing
VAE	Variational Auto Encoder

Chapter 1

Introduction

The urgency of addressing the global warming crisis has significantly intensified over the last decade, as emphasized during the Glasgow COP26 summit [1]. One of the primary targets outlined in this context is the achievement of carbon neutrality by 2050 (see fig. 1.1). This objective represents a critical driver for innovation in energy conversion systems and, consequently, for turbomachinery design.




Figure 1.1: CO_2 Emission Pathways and Global Temperature Projections

Among the sectors that have received increasing attention are liq-

uefied natural gas (LNG) and the organic Rankine cycle (ORC)[2]. These systems, which involve unconventional working fluids, are expected to play a central role in the ongoing energy transition. Their growing importance has directly influenced the demand for reliable, efficient, and flexible turbomachinery solutions capable of adapting to the requirements of such processes.

In this context, the design of centrifugal compressors is undergoing a substantial transformation. Rather than focusing on a single optimized compressor configuration, the objective is the development of entire families of compressors, each tailored to specific operating conditions and applications. Given the limited time frames imposed by the energy transition and the market demand for rapid deployment, it becomes imperative to significantly reduce the design time required for a complete compressor family, while still ensuring robustness, efficiency, and adaptability.

Compressor design can be addressed using several approaches. On one side, less computationally expensive tools such as one-dimensional [3] or two-dimensional methods can be employed for preliminary design. A typical example is provided by throughflow models [4], which are capable of generating fast results. However, these approaches cannot predict cross-flows perpendicular to the mean streamsurface, while the flow inside a centrifugal compressor has been demonstrated to be strongly three-dimensional. On the other hand, high-fidelity analyses are becoming increasingly affordable thanks to the growth in computational power. As of today, their suitability remains limited to the performance study of selected components, when a deep understanding of the flow structures and behavior is sought. For everyday aerodynamic design activity, industry still relies on RANS-based CFD codes, as they represent the best trade-off between computational cost and accuracy. This aspect is particularly critical in the context of optimization, where a wide range of geometries must be analyzed to explore the full design

space.

Flow topologies occurring in centrifugal compressors can be very complex and are often dominated by strong secondary flows, flow separations, and shock wave–boundary layer interactions, especially in transonic conditions [5]. In such scenarios, the proper selection of the turbulence model becomes a key issue in order to guarantee a sufficient confidence level in the predicted results, while ensuring robustness and numerical stability with reasonably fast simulations. For this reason, it is essential that CFD predictions are as accurate and reliable as possible, and whenever feasible, they should be validated against experimental data.

The design of centrifugal compressors is increasingly driven by the need to perform systematic optimization studies across a wide range of operating conditions. Due to the inherently high computational cost of three-dimensional simulations, optimization procedures are often supported by surrogate models, which provide an efficient approximation of the aerodynamic performance. These surrogate models rely on the construction of accurate databases, generated through carefully selected CFD simulations, and their predictive capability strongly depends on the quality of this underlying dataset. Ensuring the accuracy of surrogate models is therefore of paramount importance, as they directly affect the reliability of the optimization outcomes.

Another critical aspect is that optimization cannot be limited to the design point alone. In several applications compressors are required to operate reliably over a broad range of conditions, and therefore the off-design performance must be explicitly considered in the optimization process. Accounting for off-design behavior, however, inevitably increases the computational effort, as it requires the evaluation of multiple operating points for each candidate geometry. The need to balance design-point optimization with off-design performance evaluation, while managing the associated computational effort, represents one of

the primary challenges in compressor optimization.

In recent years, artificial intelligence (AI) techniques have emerged as a promising approach to address some of these challenges. Deep learning methodologies, in particular, have shown immense potential across a variety of scientific disciplines, including turbomachinery design. Among these methods, artificial neural networks (ANNs) [6] have been employed for multi-objective optimization, enabling the capture of complex, non-linear relationships between geometric inputs and performance outputs [7] [8] [9] [10]. Unlike traditional methods, ANNs do not rely on explicit mathematical models; instead, they learn from data, reducing the designer's burden in model selection and allowing efficient exploration of diverse design configurations. AI-based approaches thus offer a pathway to accelerate design processes and support the creation of accurate surrogate models. While previous discussion focused on the optimization of a single compressor, modern design requirements demand the creation of entire families of machines. To achieve this, information obtained during the optimization of individual compressors must be transferred and generalized to support the design of multiple geometries.

The main objective of this thesis is to develop a methodology for the design and optimization of a transonic centrifugal compressor family. The tools developed are required to be fast, robust, and reliable. To ensure the reliability of the framework, a validation was first carried out using available experimental data and geometries from Baker Hughes Florence. Numerical tools for surrogate modeling and optimization were then developed based on existing techniques, followed by a discussion of strategies to accelerate the optimization process. Finally, to improve the robustness of the methodology, a procedure for the automatic generation of centrifugal impellers is presented, allowing the designer to explore the design space freely while obtaining impeller geometries with performance consistent with other family members.

Chapter 2

Centrifugal Compressor aerodynamics

Centrifugal compressors, also referred to as radial compressors, are turbomachines designed to increase the pressure of a compressible fluid. Unlike axial compressors, in which the flow mainly moves parallel to the axis of rotation, radial compressors guide the flow radially outward from the impeller. This fundamental difference influences the geometry of both the rotor and stator blades, as well as the overall aerodynamic behavior of the machine.

It is important to distinguish between compressors and pumps. While compressors are used for compressible fluids, pumps are intended for incompressible fluids. Despite the similarity in their general objective of increasing fluid pressure, the physical properties of the working medium dictate the choice of the machine and the corresponding aerodynamic design considerations.

A stage of a centrifugal compressor consists of an impeller followed by a diffuser. The rotor and stator blades exhibit strongly different geometries to effectively manage the energy transfer and guide the flow. Diffusers can take various forms depending on the application, including vaned, vaneless, and pipe diffusers. The choice of the diffuser




Figure 2.1: Example of centrifugal compressor [11]

affects the pressure recovery, flow uniformity, and overall efficiency of the stage.

Centrifugal compressors can be arranged in either single-stage or multi-stage configurations 2.2. In a single-stage configuration 2.2(a), the diffuser is typically followed by a volute, or scroll, which collects the flow radially over 360° and redirects it tangentially. In multi-stage configurations 2.2(b), the diffuser is followed by a U-bend (or crossover) and a return channel. The U-bend collects the flow exiting the diffuser and redirects it radially inward, while the return channel reduces the swirl velocity of the flow and recovers static pressure before directing it to the next stage impeller. These components collectively ensure ef-




Figure 2.2: Comparison between single-stage and multi-stage compressors [12]

ficient energy transfer and maintain favorable flow conditions throughout the compressor. Multi-stage compressors are generally used in industrial applications where higher discharge pressures are required.

Centrifugal compressors find application in a variety of fields. Single-stage compressors are widely used in turbocharging 2.3 for automotive and marine engines, as well as in energy production, often driven by either a radial in-flow turbine or an electric motor. Industrial applications, such as oil and gas 2.4, chemical processing, energy generation, mining, and metallurgy, employ both single and multi stage compressors depending on the required discharge pressure. The selection of configuration is closely tied to the desired operational characteristics and performance objectives.

The use of non-conventional fluids in LNG and ORC applications introduces aerodynamic conditions that differ significantly from those encountered with conventional gases. In particular, the higher molecular weight of these fluids lowers the speed of sound, resulting in elevated peripheral Mach numbers at the impeller tip. Under these conditions,



Figure 2.3: Example of turbocharger [13]



Figure 2.4: Baker Hughes's High Pressure Ratio Compressor (HPRC) [14]

flows that were previously subsonic in conventional compressors can become transonic, with local supersonic regions and associated shock waves forming within the impeller passages.

Existing centrifugal compressors, originally designed for conventional fluids, are therefore unsuitable for these applications. Their geometries and blade profiles are optimized for subsonic operation, and they cannot efficiently handle the transonic flow conditions now required. As a result, new compressor designs are necessary, specifically tailored to manage the aerodynamic challenges of transonic operation. This includes careful consideration of blade shaping, diffusion channels, and flow turning to minimize shock-induced losses and maintain stable operation across the required operating range.

The design of these transonic impellers is critical to achieving high efficiency and robustness. Accurate prediction and management of compressibility effects, secondary flows, and potential flow separation are therefore essential aspects of the aerodynamic design process.

For the similarity analysis, it is common practice to describe the stage performance using dimensionless groups. The most important parameters that characterize the machine inside the design space are:

- The flow coefficient, which is the adimensional the mass flow:

$$\phi = \frac{\dot{m}_1}{\rho_{01} \frac{\pi}{4} D_2^2 u_2} \quad (2.1)$$

- The Peripheral Mach number, which is the adimensional form of the impeller speed:

$$Mu_2 = \frac{u_2}{\sqrt{\gamma R T_{t1}}} \quad (2.2)$$

A set of global performance parameters are often used, to describe the machine's effect on the working fluid:

- For a perfect gas, polytropic efficiency can be expressed as:

$$\eta_p = \frac{\gamma - 1}{\gamma} \frac{\ln(P_{02}/P_{01})}{\ln(T_{02}/T_{01})} \quad (2.3)$$

- Work coefficient, which expresses the work input of the impeller:

$$\tau = \frac{\Delta h_0}{u_2^2} \quad (2.4)$$

- Polytropic Head coefficient, which expresses the effective work input transmitted inside the flow:

$$H = \eta_p \tau \quad (2.5)$$

The described coefficients will be used in all the following chapters and their definition will be repeated, to facilitate the comprehension of the reader.

Chapter 3

Numerical Framework

3.1 Geometry parameterization

Geometric parametrization plays a fundamental role in the design and optimization of centrifugal compressors. By providing a systematic description of the impeller shape, parametrization allows designers to explore variations of the geometry and perform targeted modifications or perturbations to improve aerodynamic performance. The choice of geometrical parameters is pivotal to the success of the method. A low number of geometrical features does not explore a wide variety of solutions. At the same time, an excessive number of degrees of freedom increases the risk of studying non feasible geometries.

The impeller is primarily described in the meridional plane, where the hub and shroud curves define the main flow channel. For fully three-dimensional geometries, additional intermediate meridional sections are introduced to capture the spanwise variation of the blade and to refine the representation of the flow path. These curves are typically represented using a finite set of parameters, such as splines or Bézier curves, which enable a flexible and compact definition of the meridional geometry. An example of parameterization of a "ruled" impeller (only hub and tip sections are parametrized) is shown in 3.1, while an

example of multi-stream return channel is shown in 3.2.




Figure 3.1: Impeller Meridional Channel

Within each meridional section, the leading edge and trailing edge of the impeller blade are identified. These points serve as references for defining the camber line of the blade section, which provides the basis for the detailed blade shape. The geometry of the blade is specified independently for each section along the span. Two main characteristics are defined per section: the blade angle distribution and the thickness distribution. Both distributions can be described using a limited number of parameters, allowing for a compact yet flexible representation of the three-dimensional blade geometry 3.3.

Once the blade sections are defined, the full 3D shape of the impeller is generated by interpolating between the sections along the span. This approach allows the creation of either simple, "ruled" blades or more complex shaped blades, depending on the number of sections and the desired level of aerodynamic refinement. By systematically controlling the geometry in this way, designers can efficiently explore design variations and optimize the impeller for performance, robust-




Figure 3.2: Return Channel Meridional Channel (multiple spanwise sections)

ness, and operational requirements.

Regarding the parameterization of the impeller, a design based only on hub and tip sections can be reasonable. However, for high Mach number configurations, it is often advantageous to increase the complexity of the parameterization by adding control sections along the span.

Hazby et al. [5] present an example of these benefits by comparing three different transonic impeller blade configurations. The main difference lies in the shape of the leading edge. The study analyzes the




Figure 3.3: Chord-wise/span-wise blade angle/thickness distribution

position of the shock wave and plots iso-pressure lines along the blade. The results show that a more detailed blade shape provides clear advantages. A forward-swept impeller moves the shock wave toward the inner region, improving stability. A “barreled” shape reduces shock intensity at intermediate span levels. The comparison is illustrated in Figure 3.4.

3.2 Mesh

Once the impeller geometry has been defined, the computational domain must be discretized by means of a mesh. In the case of three-dimensional configurations, structured viscous H-type grids are widely employed, as they provide adequate resolution of the boundary layers and a flexible adaptation to different inlet and outlet configurations. This approach can be applied to both axial and radial impeller inlets, as well as to geometries with vaneless or vaned diffusers, and even to multi-stage arrangements including return channels. An example of meshed impeller, with radial inlet and vaneless diffuser, is represented in figure 3.5, while a mesh of a return channel is presented in figure




Figure 3.4: Contours of surface static pressure and vectors of relative velocity close to the suction surface [5]

3.6.




Figure 3.5: Computational domain of a impeller with radial inlet and vaneless diffuser

The different mesh blocks are typically connected using a mixing plane approach, which allows the averaging of flow properties between




Figure 3.6: Computational domain of a return channel with U-bend

adjacent components and the preservation of flow quantities between systems with different relative velocity.

A critical aspect of mesh generation lies in the definition of the appropriate grid resolution. If the mesh is too coarse, the geometric features of the impeller and flow path, and the resulting flow field, may not be accurately captured, leading to unreliable simulation results. On the other hand, an excessively fine mesh drastically increases computational cost without necessarily improving the accuracy of the predicted performance. To achieve this balance, a mesh sensitivity study is commonly performed. In this process, CFD simulations are carried out using different levels of mesh refinement, and the resulting post-processed quantities—such as pressure ratio, efficiency, or velocity distributions—are compared. Convergence of these results with mesh refinement provides confidence that the chosen grid resolution is sufficient to ensure both accuracy and computational efficiency.

3.3 CFD - TRAF code

The CFD is carried out using the TRAF code, a 3D steady/unsteady, viscous, multi-row, multi-block RANS solver developed at University of Florence [15, 16]. The code has been already used to study centrifugal compressors over the years for industrial and turbocharger applications [17, 18, 19, 20].

The flow solver incorporates several convective flux schemes that range from central ones with artificial dissipation, to TVD strategies built on top of upwind schemes [21]. For the present study, the central scheme proposed by Jameson, Schmidt, and Turkel, with artificial dissipation has been used [22].

The governing equations are solved in conservative form and mapped into a curvilinear, body-fitted coordinate system. A finite volume approach is employed to discretize the computational domain, ensuring accurate representation of complex geometries and robust numerical performance.

Several turbulence model for turbulence closure are implemented, ranging from algebraic to more complex one and two equation approaches:

- Baldwin-Lomax algebraic model [23];
- Degani-Schiff algebraic model [24];
- Mixing length algebraic model [25];
- One-equation Spalart-Allmaras model [26];
- One-equation Spalart-Allmaras model with Spalart-Shur correction [27];
- Two-equation $k - \omega$ Wilcox Low-Reynolds model [28];
- Two-equation $k - \omega$ Wilcox High-Reynolds model 1988 version [28];

- Two-equation $k - \omega$ Menter SST model [29];
- Two-equation $k - \omega$ Wilcox High-Reynolds model 2006 version [30];
- Two-equation $k - \omega$ Wilcox High-Reynolds model modified by Rung (LEA) [31];
- Realizable quadratic eddy viscosity model RQEVM by Rung [31][32].

In particular, the $k - \omega$ family of models (such as the Wilcox formulations and the Menter SST model) is commonly used in centrifugal compressor applications, as it provides a suitable trade-off between computational cost and predictive capability in the presence of strong secondary flows and shock-boundary layer interactions.

For steady-state calculations, the interaction between adjacent blade rows is handled through a mixing plane approach, which enables the exchange of circumferentially averaged flow properties across interfaces. For unsteady analyses, the solver adopts a full annulus approach with sliding interfaces, ensuring a more detailed resolution of the time-dependent interaction between impeller and diffuser.

Boundary conditions for subsonic inlet and subsonic outlet are imposed to the computational domain. At the inlet, four flow quantities must be prescribed (typically the total pressure, total temperature, flow direction), and turbulence intensity or length scale. At the outlet, a single variable, usually the static pressure, is imposed to define the overall operating point.

The code has been parallelized for both shared and distributed memory architectures, allowing efficient use of modern high-performance computing platforms. Using MPI and OpenMP for parallelization, a GPU version of the code is also available [33].

3.4 CFD - helicity correction model

In the numerical investigation of turbomachinery flows, it is essential to employ modeling strategies capable of capturing the relevant physical phenomena with sufficient accuracy. Since experimental data remain the ultimate reference for validation, numerical models must be carefully selected and developed to ensure consistency with measurements, thereby enabling a reliable interpretation of complex flow structures.

Helicity is the mathematical term used for describing the projection of the vorticity vector onto the velocity vector. Its use is quite common in particle physics where it is employed for describing the projection of the particle spin in the direction of its motion. In fluid dynamics, it was first introduced by Moffatt [34] as a value for measuring the knottedness of vortex lines in the flow.

The use of helicity as a correction term for turbulence models was already tested by Liu et al. [35] for the validation of a numerical model for predicting corner separation using Spalart Allmaras (SA) turbulence model [26]. As the original SA model predicted a larger than expected separation zone, a helicity-based correction of turbulence production and dissipation was implemented to account for energy backscatter, which plays a critical role in corner separation regions. Liu et al. [36] proposed another application of helicity-based correction in the Menter's shear stress transport model (SST) [29], analyzing once again the problem of overestimation of corner flow separations. Such a correction results in an increased turbulent kinetic energy production in corner regions, according to the helicity field, and this is reflected in a reduction of the separation zone size. Improved agreement with experimental data has been observed when the correction is applied.

Helicity however is not Galilean Invariant, meaning that its definition requires the choice of the reference in which the velocity and vorticity components are considered. The aforementioned works stud-

ied a linear PVD (prescribed velocity distribution) cascade, so the closure modifications were considered in an inertial frame of reference only. In another study of Lee et al. [37], the SA model with a helicity correction-based model is used for predicting the performance of low-speed fan rotors. The study showed that, even with rotating parts, the correction model improves the flow predictions, especially in near stall conditions, and consequently contributes to a more realistic calculated operating range. As previously said, the choice of the reference frame for the helicity terms in this case is not unambiguous.

The mentioned studies [35][35][37] proved that a correction of turbulence production and destruction based on helicity improves the flow prediction of complex and tridimensional flows, which is an appealing feature for centrifugal compressor studies.

Wilcox's k - ω turbulence model is chosen as the baseline closure [30]. It is a two equation model where the transported quantities are the turbulent kinetic energy k and ω , defined as the ratio between the dissipation and production of turbulent energy. The production term of the first equation is:

$$P_k = -\rho \overline{\mathbf{u}' \otimes \mathbf{u}'} \cdot \nabla \mathbf{u} \quad (3.1)$$

where $\overline{\mathbf{u}' \otimes \mathbf{u}'}$ represents the Reynolds stress tensor, and for the second equation:

$$P_\omega = \frac{\alpha \rho}{\mu_t} P_k \quad (3.2)$$

where α is a model coefficient and μ_t is the eddy viscosity. Based on the Boussinesq hypothesis, (3.1) is modeled as:

$$P_k = \mu_t S^2 \quad (3.3)$$

where S is the mean strain rate tensor:

$$S^2 = 2S \cdot S \quad (3.4)$$

$$S = S_{\text{ym}}(\nabla \mathbf{u}) \quad (3.5)$$

when helicity is introduced in the model, the original production term in (3.1) is now replaced with:

$$P_{kh} = P_k (1 + C_{1h} h^{C_{2h}}) \quad (3.6)$$

where h is the absolute value of helicity density, defined as:

$$h = \frac{|\mathbf{u} \cdot \boldsymbol{\omega}|}{|\mathbf{u}| |\boldsymbol{\omega}|} \quad (3.7)$$

where $\mathbf{u}$ is the velocity vector and $\boldsymbol{\omega}$ its vorticity. The two constants of the model are C_{1h} and C_{2h} , which determine respectively the maximum increase of P_k^h and its rate of growth. There is no absolute correct choice of these values. For instance, In Liu et al.'s work [35][36] the constants are set as $C_{1h} = 0.71$ and $C_{2h} = 0.6$, while in Lee et al.'s [37] work the chosen values are $C_{1h} = 0.5$ and $C_{2h} = 0.7$.

A higher value of C_{1h} results in an increase of the predicted stall margin, thanks to a higher amount on turbulent kinetic energy, and enhances the losses due to turbulence, on the contrary a small value of C_{1h} results in an underestimation of the stall margin. After calibration the two coefficients were held constant for all cases presented in this study [38].

As already pointed out, there is no physical argument suggesting to consider velocity and vorticity in the absolute rather than in the relative frame of reference for the evaluation of the helicity. Any particular combination yields a different prediction, and no generality can be claimed for any option. It is worth mentioning that the proposed model, which does not introduce new transport equations, does not increase the computational cost.

3.5 Helicity correction model - validation on centrifugal compressor stage

This section describes part of the validation campaign conducted by the author. The primary objective is to fine-tune the model coefficients $C1_h$ and $C2_h$ in order to obtain performance curves that closely match the available experimental data provided by Baker Hughes.

Two intermediate stages for high flow coefficient centrifugal compressors are presented. The first one is characterized by higher peripheral Mach number as compared to the second one, so a different level of complexity in the flow structures is to be anticipated. Two H-type grids are used for the stage discretization: the first one comprises the impeller and a vaneless diffuser (Fig. 3.5), the second one discretizes the U-bend and return channel (Fig. 3.6). The two computational blocks are linked with a mixing plane.

As a mesh sensitivity analysis is done using the baseline $k - \omega$ model, the same analysis was repeated for the helicity-corrected model providing very similar results. The tested grids are a coarse mesh G1, a medium mesh G2, and a fine mesh G3. The medium one, which is the one targeted for present study, ensures the best compromise between accuracy and computational cost. The objective is to find a mesh density with less than 0.5% discrepancy compared to G3 for the most important performance parameters, such as polytopic efficiency and work coefficient. The near-wall mesh resolution set to ensure $y^+ = O(1)$. The sensitivity analysis data is reported on table 3.1.

The rotational speed is kept constant. The boundary conditions are adiabatic walls with no slip condition, prescribed spanwise distributions of total pressure, total temperature and flow angles at inlet, outlet static pressure value is imposed at the hub with a distribution based on the radial equilibrium. Boundary conditions for the turbulence quantities are imposed in terms of Turbulence intensity and tur-

Table 3.1: Mesh sensitivity analysis of Compressor 1

Grid	ϕ/ϕ_d	No. of Elements Error with G3	Polytropic Efficiency Error with G3	Work Coefficient Error with G3
G1	1.00	1.00 M	1.06%	1.17%
G2	1.00	2.30 M	0.18%	0.32%
G3	1.00	4.50 M	-	-
G1	0.86	1.00 M	1.00%	1.16%
G2	0.86	2.30 M	0.48%	0.41%
G3	0.86	4.50 M	-	-

bulence length scale.. Convergence acceleration techniques are used to decrease convergence time, such as local time stepping, residual smoothing and a 2-level multigrid strategy.

The experimental activity was performed on a dedicated model test bench at the Turbomachinery Technology Laboratory (TTL) in Baker Hughes facility in Florence, Italy. The test bench is a closed loop rotating single stage rig, which is normally employed for performance measurements of different centrifugal compressor stage families; further details can be found in [39][40][41].

3.5.1 Experimental results

The test campaigns were carried out using carbon dioxide as the working gas and approximately ambient conditions at the stage inlet. The impeller has been tested in an intermediate stage arrangement. Indeed, a multistage compressor was simulated through a “pseudo-stage” in the flow-path upstream of the impeller, including a set of pre-swirl vanes followed by a return channel. A vaneless diffuser and a downstream return channel are included in the stage as well (see fig. 3.7).

The test rig was instrumented with total temperature and total pressure rakes to allow flange-to-flange measurements. Measurements were also taken at various locations throughout the test article to derive single component (i.e. impeller, diffuser, return channel) performance.




Figure 3.7: Test rig layout

The measurements of total pressure were performed with Kiel probes, total temperature with shielded J-type thermocouples and flow angles with three-hole and five-hole probes, for the yaw and pitch angles respectively. At each of the measurement locations throughout the compressor, static pressure taps were positioned around the circumference of both the hub and the shroud sides; see also Table 3.2.

Table 3.2: Measurement Sections

Section No.	Measurement
10 Impeller Inlet	P_0, P_s, T_0 , Yaw Angle, P_{dyn}
20 Impeller Outlet	P_0, P_s, T_0 , Yaw and Pitch Angles, P_{dyn}
40 Diffuser Outlet	P_0, P_s, T_0 , Yaw and Pitch Angles, P_{dyn}
60 Stage Outlet	P_0, P_s, T_0 , Flow Angle

The mass flow was measured according to EN ISO 5167 standard with an orifice plate; whereas the rig rotating speed was measured using a magnetic pick-up-based key-phaser. Dynamic pressure probes, having a resonance frequency over 250 kHz, have been installed at sections 10, 20 and 40 to detect stall inception. In the diffuser sections

(20 and 40) 3 probes are installed at the same radius and at different tangential positions, in order to determine phase shift which will be required in data processing; the reader is addressed to [42][43][44][45] for more details about procedure, criteria, and data analysis. From an experimental standpoint, stall investigation is carried out by reducing compressor volumetric flow rate from right limit up to surge occurrence. Unsteady pressure signals are acquired, and stall is identified by the presence of sub-synchronous phenomena. Once stall region is identified, a detailed analysis of rotating stall is then performed; in particular the stall pattern and number of lobes is reconstructed from signals phase shift.

The non-dimensional parameters used for describing the swallowing capacity and the impeller rotational speed are, respectively, the flow coefficient ϕ and the peripheral Mach number Mu_2 :

$$\phi = \frac{\dot{m}_1}{\rho_{01} \frac{\pi}{4} D_2^2 u_2} \quad (3.8)$$

$$Mu_2 = \frac{u_2}{\sqrt{\gamma R T_{t1}}}$$

where the subscripts 1 and 2 indicate the impeller inlet and outlet sections respectively (sections 10 and 20 of Fig. 3.7). For high flow coefficient and peripheral Mach number, a high impeller inlet relative Mach number is expected, which leads to a transonic region near the impeller shroud. In transonic inducers, the entropy generation due to shocks has a strong impact on impeller efficiency. In order to correctly predict the stage performance, CFD needs to correctly simulate the interaction between the shock wave and boundary layer on the blade suction side.

The numerical framework has been validated against experiments on a large number of centrifugal compressor stages over a wide range of flow coefficients and peripheral Mach numbers. For sake of brevity only two cases will be presented. Such cases have been selected to demon-

strate model capability to enhance CFD predictability when required and to not jeopardize standard model capabilities.

3.5.2 Compressor 1

This machine was chosen for the model validation because of its high flow coefficient and peripheral Mach number, its transonic behavior and the availability of experimental data. Figure 3.8 shows the results in terms of stage polytropic efficiency curve, at the design rotational speed. Polytropic efficiency for a perfect gas can be expressed as:

$$\eta_p = \frac{\gamma - 1}{\gamma} \frac{\ln(P_{02}/P_{01})}{\ln(T_{02}/T_{01})} \quad (3.10)$$

where the terms with subscripts 1 and 2 belong respectively to the stage inlet and outlet section (sections 10 and 60 of Fig. 3.7).




Figure 3.8: Compressor 1, performance comparison in terms of polytropic efficiency

Also in this plot, a comparison between the use and the suppression of helicity-based correction model is shown (to neglect the effects of

the model C_{1h} is set to 0). The improvement of the flow prediction using helicity is clear, especially at low flow coefficients. It is noticeable how the CFD without helicity correction predicts the stall boundary at higher flow coefficients than expected. The curve shape is also different from the experimental one. If helicity correction is added, the extension of the operating range on the left side of the curve is enhanced, while the choke margin is unchanged, which means that the model does not influence flow blockage at choke in a significant way. If we compare the absolute values of polytropic efficiency between the two CFDs, we see that the addition of turbulent energy, given by helicity, tends to decrease the efficiency of the whole stage. The correction model has an influence at every flow coefficient, in the considered range, by increasing the turbulent kinetic energy. Predicted efficiency absolute values are greater than the experimental ones. One of the main reasons is that some external losses, such as windage, are not being accounted for in the numerical model. Figure 3.9 shows the comparison of work coefficient, defined as:

$$\tau = \frac{\Delta h_0}{u_2^2} \quad (3.11)$$

where Δh_0 is the total enthalpy variation through the impeller.

The comparison shows that the slope given by the model with the helicity-based correction has a better agreement with the one shown by experimental results. In contrast, at low flow coefficients, the CFD without helicity model tends to decrease the curve slope.

Figure 3.10 shows the polytropic head coefficient results, defined as the product of 2.3 and 2.4:

$$H = \eta_p \tau \quad (3.12)$$

Since RANS calculations are not suited to identify compressor stall, the authors identify a “numerical” stall with the last operating point




Figure 3.9: Compressor 1, performance comparison in terms of work coefficient




Figure 3.10: Compressor 1, performance comparison in terms of polytropic head

reaching convergence. Convergence is considered to be attained when


changes lower than 0.5% occur in terms of flow coefficient, polytropic efficiency, work coefficient and head coefficient, after successive runs with a prescribed number of steps. The approaching of the stall margin can be visualized by the slope of the polytropic head becoming horizontal. It can be seen that CFD with helicity model reaches a zero slope at the same flow coefficient of experimental data, while calculations without correction model see a premature stall.

Further insights can be gained by comparing computed and experimental spanwise distributions of flow quantities in several stage locations, and at different flow coefficients. A comparison of the normalized total-to-total pressure ratio at stage exit (section 60 in fig. 3.7), at design rotational speed, is shown in figs. 3.11(a)-3.11(b) respectively at $\phi/\phi_{choke} = 0.806$ (near stall condition) and 0.948 (design condition).


In near stall conditions (fig. 3.11(a)) it is possible to observe that the increase in turbulent kinetic energy, due to helicity corrections, gives a better match with experimental data, in terms of absolute values and profile shape.

A similar behavior can be observed at design flow coefficients (fig. 3.11(b)), where the profile of CFD with helicity model is closer to the experimental data. It is also interesting to look at the comparison between the two CFD analyses in fig. 3.11(b): between midspan and hub, the predicted total-to-total pressure ratio is lower with the addition of helicity, and the opposite is true near shroud. This can be explained by remembering that helicity increases turbulence, which is directly linked to losses. At the same time, a more energetic fluid has a smaller separation where a shock wave is encountered, so the pressure losses due to the interaction between shock wave and boundary layer separation are smaller.

In order to investigate even further on the effects of helicity corrections, in particular at low flow coefficients, a study of the flow field near



((a)) Near stall condition



((b)) Design point

Figure 3.11: Compressor 1, spanwise distribution comparison in terms of normalized pressure ratio (Section 60)


the impeller shroud is made. Figure 3.12 shows the normalized Mach number contours in the relative reference frame, at $\phi / \phi_{\text{choke}} = 0.806$

comparing the CFD with and without the helicity correction model. As expected, the high peripheral Mach number, combined with a high incidence, generates a transonic region and a strong interaction between the shock wave and the boundary layer. The result is a separated shear layer near the blade suction side, which significantly increases losses.


Downstream of the shock wave, a region of low momentum fluid is created, which tends to separate while subjected to the adverse pressure gradient. The impact of helicity correction is highlighted in figs. 3.12(a)3.12(b), where it is apparent that the low momentum region extends for a larger portion of the vane for the standard $k - \omega$ model while it is significantly reduced when the correction is exploited.

Contours of non-dimensional relative stagnation pressure are used to highlight the regions where losses are located in a section just downstream the impeller trailing edge. Figure 3.13 compares the standard and the helicity-corrected turbulence model for the near stall condition. The secondary flows distribution is in good agreement with the typical “jet & wake” pattern. The area of low relative stagnation pressure fluid, corresponding to the so-called “wake”, is located in the suction side-shroud corner. For the standard turbulence model fig. 3.13(b) the wake increases and occupies a larger pitchwise extension compared to fig 3.13(a) where the helicity correction is used. The increased production of turbulent kinetic energy in regions of the flow field where helicity is dominant reduces the extent of low-momentum fluid and thus the formation of secondary flows at the impeller outlet. This effect delays the onset of numerical instability corresponding to the last converged point in the predicted characteristic (see fig. 3.10).

The evolution of the turbulent kinetic energy (TKE) along the vaned passage is highlighted on figure 3.14. As expected, the case with helicity correction model has a greater turbulent kinetic energy production term.



((a)) With helicity correction



((b)) Without helicity correction

Figure 3.12: Compressor 1, relative Mach number contour at impeller 85% span at $\phi/\phi_{choke} = 0.806$





Figure 3.13: Compressor 1, normalized relative total pressure contours at impeller exit (Section 20) at $\phi/\phi_{choke} = 0.806$

3.5.3 Compressor 2

In order to have a broader evaluation of the $k - \omega$ model, another machine is discussed. In this case, the design peripheral Mach number is lower than that of compressor 1, so a less critical flow configuration is expected, especially at the impeller shroud. The comparison of poly-




Figure 3.14: Compressor 1, normalized turbulent kinetic energy production at $\phi/\phi_{choke} = 0.806$

tropic efficiency, work coefficient, and head coefficient is presented in figures 3.15 3.16 and 3.17.

It can be seen how the standard $k - \omega$ model is already able to provide a good prediction of the measured characteristics, in terms of shape and operating range. In particular the polytropic efficiency (fig. 3.15) is overestimated by about 2-3% and the work coefficient (fig. 3.16) is in pretty good agreement with the measurements. It is important to verify whether the introduction of helicity correction model does not worsen the flow prediction, so that the introduced model can be profitably adopted for a wider combination of flow coefficients at peripheral Mach numbers. Comparing the performance curves in figures 3.15 3.16 and 3.17, there are no major changes in the predicted shape and operating range, while a difference between the CFD results in terms of absolute efficiency, and consequently in absolute polytropic head, values are detected. This is consistent with the higher production of turbulent kinetic energy of the $k - \omega$ model resulting in a higher




Figure 3.15: Compressor 2, performance comparison in terms of polytropic efficiency




Figure 3.16: Compressor 2, performance comparison in terms of work coefficient

level of loss.

Spanwise distributions of normalized total-to-total pressure ratio





Figure 3.17: Compressor 2, performance comparison in terms of polytropic head

are shown in Fig. 13(a) for the near stall conditions ($\phi/\phi_{choke} = 0.691$) at stage exit where experimental data are available.


It can be seen how the two numerical predictions are close to each other (note that the x -resolution of fig. 3.18 is twice as the one in fig. 3.11), and the helicity-corrected turbulence model better matches the experimental results. Figure 3.18(b) shows the same comparison when the flow coefficient is near choke ($\phi/\phi_{choke} = 0.959$). In this case, it is worthwhile to notice that a side effect of helicity correction is recognizable in a smoothing of the spanwise profile of total pressure ratio, probably due to the higher turbulent mixing resulting in a slightly different shape of the distribution.

Once again, the impact of the helicity-correction on the local prediction of the flow field is verified in fig. 3.19 where the contours of relative Mach number near the impeller shroud are compared for the near stall condition.

The global picture is quite similar in this case, even if the standard turbulence model predicts a larger extension of the low-momentum re-




((a)) Near stall condition




((b)) Near choke condition

Figure 3.18: Compressor 2, spanwise distribution comparison in terms of normalized pressure ratio (Section 60)

gion near the suction side (Fig. 3.19(b)). The secondary flows pattern is visible in figure 3.20, where the relative total pressure contours are




((a)) With helicity correction




((b)) Without helicity correction

Figure 3.19: :
Compressor 2, relative Mach number contours at impeller 85% span
at $\phi/\phi_{choke} = 0.691$

shown at the impeller exit for the same operating condition.



((a)) With helicity correction



((b)) Without helicity correction

Figure 3.20: Compressor 2, normalized relative total pressure contours at impeller exit (Section 20) at $\phi/\phi_{choke} = 0.691$

The high-loss region in this case moves from the corner of the suction side while remaining close to the shroud. Compared to Compressor 1 (Fig. 3.13) the loss regions are of comparable size indicating a lower impact of the helicity correction for this case.

In conclusion, the introduction of the helicity correction model has been demonstrated to better match the centrifugal compressor experimental results. Further details of this work can be found in [38].

Chapter 4

Artificial intelligence and neural networks

Artificial Intelligence (AI) is a field of computer science that aims to create systems capable of performing tasks that normally require human intelligence, such as recognizing patterns, solving complex problems, or making decisions. Within AI, Machine Learning (ML) is a key subset. ML allows computers to learn from data and improve performance over time without being explicitly programmed for each task. In other words, AI is the overall goal of simulating intelligent behavior, while ML provides the practical methods to achieve it by learning from data.

Machine Learning can be divided into two main types: supervised learning and unsupervised learning. In supervised learning, models are trained using datasets where both inputs and outputs are known, enabling the system to learn the relationship between data and make predictions for new cases. Typical examples include regression, used to predict continuous values, and classification, used to assign inputs to discrete categories. In unsupervised learning, models do not have labeled outputs and focus on discovering hidden structures in data, such as clustering or dimensionality reduction.

A further subset of Machine Learning is Deep Learning (DL), which uses deep neural networks inspired by the human brain. Unlike traditional machine learning, which often requires manual selection of relevant features (feature engineering), deep learning can automatically learn complex representations of data through multiple layers of transformations. This makes it highly effective in tasks such as image recognition, natural language processing, and detecting complex patterns in large datasets.




Figure 4.1: Machine learning classification

Deep learning methodologies show great potential across a myriad of scientific disciplines, and the field of turbomachinery design is no exception. In certain scenarios, researchers employ techniques like genetic algorithms for optimization [46][47]. The technique is structured as follows: a set of solutions is derived from simulations. The top-performing solutions are then selected to form a new set for the next iteration. However, there are several downsides. The method necessitates substantial computational resources, and numerous solutions are discarded as they fail to meet the criteria for the subsequent round. Additionally, genetic algorithms do not allow modification of constraints and output parameters during operation. This limitation can hinder adaptability and flexibility in optimization processes.

Another method approach relies on response surfaces [48], built from polynomials, which approximate CFD solutions. Despite the utilization of the entire database in this method, the presence of a large number of degrees of freedom required to describe the geometry renders the approach less accurate. This is due to a phenomenon commonly referred to as the "curse of dimensionality", wherein an excessive number of degrees of freedom increases prediction error. The only way to diminish this error is to significantly increase the dimensions of the database.

Another method, known as the adjoint method [49], provides an alternative approach to turbomachinery optimization. This method computes the derivatives of design parameters with respect to desired performance objectives. This enables engineers to swiftly identify which modifications in geometry or operating conditions maximize the performance of interest. One of its key features is the ability to assess the impact of multiple design objectives simultaneously, facilitating enhanced multi-objective optimization. Moreover, it overcomes the obstacle mentioned in the response surface, as its computational cost is independent of the number of design variables. However, a limitation of this method is its limited flexibility in selecting output parameters after the database creation, which may constrain adaptability during the optimization process. Additionally, the adjoint method requires significant code modification to incorporate the derivative computations within the solver. Furthermore, with presence of multiple local optima, the method may converge to a suboptimal solution.

In recent years, Artificial Neural Networks (ANN) have proven particularly effective in capturing complex, non-linear relationships between geometric inputs and overall performance outputs. These correlations are essential for enabling efficient design and optimization. ANNs are a class of supervised learning algorithms widely used for multi-objective optimization tasks. They operate by learning from la-

beled datasets, which are split into training and testing sets to evaluate model performance. ANNs consist of interconnected nodes, or neurons, organized into layers, including an input layer, one or more hidden layers, and an output layer. Unlike traditional optimization techniques, ANNs do not rely on explicit mathematical models but instead learn complex relationships between input and output data through the adjustment of weights in the network. The automatic adjustment of trainable parameters in ANNs alleviates the designer's responsibility regarding model selection, while there are still hyperparameters to consider, such as the size of the entire network. However, the flexibility of the model allows for the training of any combination or subset of input parameters and labels once the database is created. This versatility enables the exploration of various design configurations and performance objectives without the need to rebuild the model from scratch.

4.1 MLP code - BaguetteNet

In order to exploit the potential of Artificial Neural Networks for turbomachinery optimization, a dedicated training code was developed as part of the present work. The tool is based on a classical Multi-Layer Perceptron (MLP) architecture and requires as input a database generated from CFD simulations. The input layer corresponds to the degrees of freedom defined in the geometric parameterization of the impeller, while the output layer is associated with the global performance quantities of interest, which serve as labels for the supervised learning task.

An example of MLP architecture is reported in Figure 4.2. The network is structured as a feed-forward architecture, in which the information propagates unidirectionally from the input layer to the output layer through the hidden layers, without any feedback connections.

The framework is designed to provide a high degree of flexibility to




Figure 4.2: Example of MLP architecture

the user. The number of hidden layers and the number of neurons in each layer can be freely specified, allowing the creation of architectures tailored to the complexity of the problem. In addition, the user can define the loss metric, the number of training epochs, and the percentage of data assigned to training and testing. In order to assess the predictive capability of the model, the database is divided into two subsets: a training set and a testing set. Only the training data are used during the learning process, while the testing data are employed to evaluate the prediction error on unseen configurations. This procedure ensures that the model is not only capable of reproducing the training data but also of generalizing to new cases.

The training process is iterative and extends over the number of epochs specified by the user. During each epoch, the network parameters (weights and biases) are automatically updated through the back-propagation algorithm, which minimizes the loss function by adjusting the parameters in the direction of the error gradient. Over successive epochs, this procedure progressively reduces the discrepancy between predicted and target values, improving the accuracy of the model. An example of loss history of train and test sets are reported in Figure 4.3.

The implementation of the code was carried out using the PyTorch




Figure 4.3: Example of loss history of train and test datasets

library, which provides a robust and efficient environment for defining neural network architectures, handling automatic differentiation, and performing optimization. The code, called BaguetteNet, takes as input the labelled database, and a set of different hidden layer architectures. The code serially trains the different Neural Networks. The architecture with the most accurate predictions on the test set is saved as the final model.

To check the quality of the training on the test set, several loss metrics are available:

- **Mean Squared Error (MSE):**

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (4.1)$$

- **Mean Absolute Error (MAE):**

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (4.2)$$

- **Maximum Absolute Error (MaxAbs):**

$$\text{MaxAbs} = \max_{i=1,\dots,n} |y_i - \hat{y}_i| \quad (4.3)$$

where:

- y_i denotes the true value of the i -th sample,
- $\hat{y}_i$ denotes the predicted value of the i -th sample,
- n is the total number of samples.

To further check the quality of the training, a comparison between true values and predicted ones of every output feature can be plotted. An example is shown in Figure 4.4. In this plot, the x -axis represents the predicted values for each feature, while the y -axis represents the corresponding true values. The closer the points lie to the line $y = x$, the more accurate the prediction.

4.2 Optimization Code

Once the neural network has been properly trained and validated, it can predict the performance of new geometries within seconds, effectively serving as a surrogate for costly CFD evaluations. This capability opens the way to a systematic exploration of the design space. Nevertheless, such predictive power alone does not guide the designer toward an optimized geometry unless the results are framed within an optimization framework. For this reason, a second code was developed to establish Pareto fronts based on user-defined performance objectives, ensuring that the role of each output parameter is carefully selected according to industrial and product requirements, and that the database contains the relevant variables for post-processing.

The optimization process is iterative. At each iteration, a set of N candidate designs is randomly generated within the design space and




Figure 4.4: Example of output values comparison between true values and predicted values for each output feature

evaluated by the trained ANN. The network rapidly provides the corresponding performance outputs, thus avoiding the need for additional costly CFD evaluations. A series of filters is then applied to remove solutions that do not satisfy predefined feasibility or performance constraints. For the remaining cases, a scoring system is employed in order to enable ranking even under multi-objective conditions. Each performance metric (e.g., efficiency, pressure ratio) can be set to be maximized or minimized, and a weight is assigned to each objective according to the designer's priorities. This approach reduces the multi-objective optimization problem to a single scalar score per candidate solution, facilitating straightforward comparison. The solutions are then sorted in descending order of score. The top M candidates constitute the Pareto front, representing the best trade-offs among the objectives for that iteration. The process is repeated over successive iterations, progressively refining the exploration of the design space

and moving towards optimal geometries.

In this work, the Pareto front is defined with respect to multiple objectives that can be evaluated either at the same operating condition or across different operating points. The optimization framework allows the inclusion of several performance metrics, each of which can be independently specified as a quantity to be maximized or minimized. For example, the optimization problem may be formulated to maximize efficiency at the design point while simultaneously minimizing the flow coefficient at near-stall conditions, with the objective of increasing the left side operating range.

A scheme of the iterative optimization procedure is presented in Figure 4.5




Figure 4.5: Iterative optimization procedure scheme

This framework provides two main advantages. First, it leverages the surrogate model to drastically reduce the computational cost compared to direct CFD-based optimization. Second, it generates a diverse set of high-performing geometries, offering the designer multiple alternatives from which to choose according to practical or application-specific requirements.

4.3 Example of Optimization procedure using Surrogate model

To perform aerodynamical optimization via ANN, a database correlating geometry with performance metrics is created. In this study, an internally developed parameterization scheme is leveraged. It consists of 33 degrees of freedom (DOFs), which are used to describe:

- Meridional channel: this includes the impeller radial inlet, vaneless diffuser, and blade corner point locations, considering the shape of the meridional channel for a multi-stage compressor (see Fig. 3.1);
- Free-form blade: this covers chordwise and spanwise blade angle distributions at hub and tip sections, chordwise and spanwise thickness distributions at hub and tip sections, leading edge spanwise shape in the meridional channel and spanwise airfoil stacking law. Moreover, along the span, the chordwise thickness and blade angle distributions are obtained by nonlinear interpolation of corresponding values at hub and tip (see Fig. 3.3). Another degree of freedom describes the magnitude of the nonlinear operator.

The Sobol sequence [50] is employed to generate a random distribution of perturbed data. Unlike a normal random distribution, the Sobol sequence ensures uniform coverage of the given data range, hence equally representing the perturbation space.

In order to map the design space, the performance of 1500 geometries are calculated using 3D CFD at three operating points, for a total of 4500 runs. The choice of using this number of observations is based on the required prediction accuracy. With less data, the Neural Network increases its prediction error.

After postprocessing, it's necessary to identify parameters that are indicative of machine performance. Some parameters are selected for maximization, such as efficiency, while others are monitored to ensure that the machine meets some technical specifications, such as work coefficient. Those parameters will be the labels, or output parameters, of the database. Three different operating points are computed for each perturbed geometry to estimate the impeller operating range. Operating points are initially selected based on the baseline performance curve, after which each perturbation is executed by imposing the same outlet static pressure. The three operating points are:

- near stall: the impeller operates under conditions close to stall;
- design point: the perturbed solutions should all approximate $\phi/\phi_{design} = 1$. This ensures that the impeller operates near its design flow coefficient;
- choke point: the focus is solely on the maximum flow coefficient. This operating point provides crucial information regarding the impeller's performance at maximum flow conditions.

The CFD database serves as the foundation for the Multi-Layer Perceptron (MLP) architecture. The optimization algorithm employed for backpropagation is Adam (Adaptive Moment Estimation) [51], which uses Stochastic Gradient Descent (SGD), with a batch of 32 observations, and root mean square propagation (RMSprop). Mean Squared Error serves as the loss metric, nonlinearities are computed using ReLU (Rectified Linear Unit) as activation function between each layer, except for the last one, which remains linear. The database is partitioned into training and testing datasets via an 80-20 split. Training proceeds for a specified number of epochs, monitoring the loss to prevent overfitting. It has been observed that beyond 300 epochs, there are no further improvements in prediction errors. Similar network architectures have been used in the past for aerodynamic optimizations [7][8].

Each output parameter plays specific role in the optimization procedure:

- for the near stall operating point, all the output parameters are utilized as filters to ensure that solutions have a low flow coefficient, providing insights into performance at low flow rate, while maintaining high efficiency to better handle high pressure ratio conditions;
- at the design point, the primary objective is to maximize polytropic efficiency. Additionally, other output parameters such as flow coefficient and work coefficient are employed as narrow band-pass filters to ensure that the technical specifications of the compressor are preserved;
- since low-choke limit solutions are undesirable, a high-pass filter is imposed for ϕ_{choke} .

The Neural Network training and optimization phases do not require a huge amount of time, compared to database generation time. 60 minutes are required from the ANN training (while testing 10 different hidden layer architectures), while 120 minutes are required for the optimization phase (100 iterations).

The resulting Pareto front is represented, in red, in figure 4.6. In that image, the low-pass, band-pass and high-pass filters on flow coefficients at different operating points are also displayed.

The best 3 geometries are validated via CFD, in order to make the entire operating curve. Each solution is chosen based on specific performance criteria, embodying different aspects of optimal design.

- GEOM01 emerges as a robust option, boasting an exceptionally wide operating range. Moreover, among the various solutions within the Pareto front, GEOM01 stands out for its comparatively minimal efficiency decay at low flow coefficients;




Figure 4.6: Pareto front (in red) superimposed to the starting database (3 operating points in purple, orange and green)

- GEOM02 prioritizes the maximization of the flow coefficient at choke, making it the optimal solution for handling expected increases in flow rates beyond the initial design parameters during its operational lifespan;
- GEOM03 is distinguished by its high value of the polytropic efficiency at design point, though at the expense of operating range.

In Figure 4.7, 4.8 and 4.9 the operating curves of the 3 selected designs are presented, in terms of polytropic efficiency, defined in 2.3, work coefficient, defined in 2.4, and polytropic head, defined in 2.5.

Flow coefficient is expressed as ϕ/ϕ_{max} , in order to better compare operating ranges. As observed, the solution with the highest efficiency, GEOM03, exhibits the lowest operating range. It is also worth mentioning that by transitioning to lower flow coefficients from the point of peak efficiency, this machine experiences the steeper decline in effi-




Figure 4.7: Performance comparison in terms of polytropic efficiency




Figure 4.8: Performance comparison in terms of work coefficient

ciency. GEOM1 exhibits a lower peak efficiency than GEOM03, yet it demonstrates a higher operating range, being able to reach lower nor-




Figure 4.9: Performance comparison in terms of polytropic head coefficient

malized flow coefficients. Notably, at low flow coefficients, GEOM01 presents the lowest performance decline, compared to its peak value. GEOM02 stands as an intermediary between the two other solutions, exhibiting a peak efficiency value between GEOM01 and GEOM03, and an operational range almost equal to GEOM01.

RANS calculations are not suited to identify compressor stall. Therefore, the stall is identified with the last operating point reaching convergence. Convergence is considered achieved when there is less than 0.5% change in flow coefficient, polytropic efficiency, work coefficient, and head coefficient after successive runs with a prescribed number of steps. An alternative method to assess the stall margin is given by the flow coefficient at which the slope of polytropic head becomes horizontal, as on Figure 4.9. In Figure 4.8 the effect of narrow band-pass filter of work coefficient at the second operating point can be highlighted, where the same value of baseline is imposed.

The determination of which solution to prioritize, considering both

industrial and technological requirements, lies with the designer choice.

It is also interesting to compare the three optimized solutions in terms of impeller blade load distribution. In Figures 4.10(a), 4.10(b) and 4.10(c), the normalized chordwise pressure distributions on blade pressure and suction sides are compared. For every geometry, the peak efficiency operating point is plotted, at 3 different span levels. A strongly different load distribution is present close to impeller shroud, as shown in Figure 4.10(a). This region is of particular interest as it significantly influences the performance, given the transonic nature of the flow, due to the high value of the inlet relative Mach number at casing, and the consequent interaction between shock waves and boundary layer, which intensifies with the shock magnitude. Therefore, it is worth to discuss how the three different geometries handle this design condition:

- GEOM01 exhibits the highest shock wave intensity, resulting in a more pronounced interaction between the shock wave and the boundary layer, leading to increased losses and lower efficiency;
- GEOM02 decreases the shock intensity, compared to GEOM01, by reducing flow acceleration before the shock wave;
- GEOM03 presents a different configuration, by substituting the normal shock wave with two oblique shocks. These shocks can potentially improve performance because, at equivalent upstream Mach numbers, they generate lower shock intensity. This results in a reduced adverse pressure gradient for the boundary layer, which is generally the primary source of aerodynamic losses.

Examining the flowfields at the peak efficiency points, at 80% of relative span, provides further insight. In Fig. 4.11(a), contours of the normalized relative Mach number for GEOM01 reveal the expected presence of a bow shock and passage shock structure ahead of the

leading edge. Flow acceleration occurs in the uncovered section of the impeller blades, followed by significant compression within the passage shock. In the contour plot for GEOM02, as depicted in Figure 4.11(b), the reduction in shock intensity is evident due to the negative camber of the uncovered portion of the blade passage. The negative camber induces converging iso-Mach lines, promoting flow compression and consequently reducing its velocity ahead of the passage shock. Consequently, it is expected that the increase of boundary layer thickness after the shock also diminishes, compared to GEOM01. GEOM03 in Figure 4.11(c) is in a negative incidence configuration, which changes the passage shock configuration. Instead of presenting a normal shock wave, the passage shock is made by two oblique shocks.


To summarize, the presented methodology demonstrates an aerodynamic optimization approach for high Mach number centrifugal impellers exploiting deep learning. The main objective is to optimize the impeller aerodynamics leveraging a comprehensive CFD simulation database, while employing artificial neural networks (ANNs) to scrutinize a number of viable solutions. The database includes information about three operating points, ranging from near choke to near stall. A Pareto front is built using ANN predictions by applying the proper boundaries on each output parameter that needs to be monitored. One notable advantage of this methodology is its ability to allow the selection of multiple solutions tailored to different design priorities, such as maximizing the flow rate, the peak efficiency, or the operating range, from a database encompassing a wide variety of geometries. It is also possible to define multiple Pareto fronts without changing the database or the trained neural network.

The success of this approach is given by several key factors. Firstly, the selection of the parameterization scheme plays a critical role in defining the design space. Additionally, the quality and robustness of the CFD database are paramount, as they directly impact the accuracy


and reliability of the optimized solutions generated by the ANN. Lastly, the careful selection of output parameters is crucial for guiding the optimization process towards achieving the desired design goals.

The proposed methodology offers a streamlined approach to aerodynamic optimization, enabling designers to explore a wide range of possibilities efficiently. Moreover, the flexibility of the approach allows for the adaptation of solutions to specific design requirements and constraints, further enhancing its applicability.


Further details of this work can be found in [52].



((a)) 80% span




((b)) 50% span




((c)) 20% span


Figure 4.10: Impeller blade load comparisons at maximum polytropic efficiency point.



((a)) GEOM01



((b)) GEOM02



((c)) GEOM03

Figure 4.11: Normalized relative Mach number contours at peak efficiency point at near impeller shroud

Chapter 5

Optimization speedup

Although the procedure described in the previous sections demonstrates both its strength and potential, it still presents some limitations and can certainly be improved. In particular, accelerating the optimization process is of great importance, especially when the goal is to generate entire families of machines rather than a single optimized design. Enhancing computational efficiency becomes a key enabler in making such approaches viable for industrial applications.

The main bottleneck in the optimization procedure is represented by the generation of the database. The minimum database size required to properly train a surrogate model cannot be established a priori, making the process particularly time-consuming.

Moreover, the database must contain information not only at the design point but also across a sufficiently wide operating range, since the impeller is expected to guarantee satisfactory performance under variable conditions. For this reason, multiple databases are usually generated, each corresponding to different operating points.

A further complication arises from the numerical stability of the geometries under CFD simulations. While even design point calculations may occasionally diverge, this risk becomes more significant in off-design conditions. In such cases, the computational effort invested

in CFD may not translate into usable data, reducing the overall efficiency of the database generation process.

While CFD simulations generate valuable flow information, this data is often underutilized in aerodynamic design. In contrast, human designers typically integrate both global performance metrics and detailed flow data, such as velocity fields and blade loading distributions, to assess and improve design quality.

5.1 Database size

A critical aspect of the optimization framework concerns the definition of the database size. One potential strategy for reducing computational costs could consist in exploiting the experience gained from previous training campaigns. For instance, if a neural network reaches a satisfactory accuracy level after being trained on N geometries, it may appear reasonable to assume that a similar number of cases would be sufficient for subsequent databases. However, the effectiveness of this approach remains questionable.

First, the flow regime may vary significantly from one machine to another. In particular, the presence of nonlinear phenomena, for example associated with high Mach numbers inside the impeller passage, may require a larger database than in subsonic cases, where the flow is generally more predictable. As a result, the required dataset size is strongly influenced by the specific aerodynamic conditions under consideration.

Second, the amplitude of the geometric perturbations applied during the parametrization also plays a decisive role. When the perturbation range is too wide, a larger number of geometries becomes necessary to ensure adequate coverage of the design space. This, however, leads to longer computational times and an increased probability of generating geometries that are far from the optimized ones, which may

result in non-convergent CFD simulations. Conversely, when the perturbation range is too narrow, the design space may not be sufficiently explored, limiting the generalization capability of the trained model and reducing the overall effectiveness of the optimization process.

The best practice consists in building the database incrementally, while constantly monitoring the progress of the training. In practice, after an initial generation of geometries, the network is trained and the loss function is evaluated. If the error curve no longer shows significant improvements, it indicates that the amount of data is sufficient to represent the design space under consideration. Conversely, if the loss remains high or unstable, the database must be enlarged with additional geometries until a reliable predictive model is achieved.

An example of that is represented on figure 5.1 , where the training loss for each database size (also called number of observations) is compared. The chosen loss is Mean Absolute Error (MAE) and its definition can be found in 4.2.




Figure 5.1: Mean Absolute Error value of test set vs database size

As can be seen, after 1200 observations there is no improvement of training loss, hence it is not necessary to increase the database size.

5.2 Convolutional neural network approach for blade loading inference

A significant limitation of AI lies in their frequent disregard for the underlying flow physics. While computational fluid dynamics (CFD) simulations provide detailed insights into the flow field, this information is typically unexploited in conventional neural network models. Developing a methodology that can leverage this rich dataset effectively is crucial for combining global performance metrics with detailed local flow characteristics, offering a comprehensive evaluation of the system's behavior without increasing the time required for database generation.

The following work presents a methodology that leverages neural network architectures to predict flow characteristics in a transonic impeller. Geometric inputs are still generated using a parametric tool, and the outputs are categorized into two types: global performance parameters, which are used to train the multilayer perceptron (MLP), and blade loading data represented as pressure fields, which are collected as images. These images are then employed to train a reverse convolutional neural network (CNN).

In recent years, several studies have focused on training neural networks to incorporate flow information into the design process. These methods typically use convolutional neural networks (CNNs), introduced by LeCun [53], which are capable of processing and analyzing images, enabling the integration of detailed flow data into the design workflow.

In ref. [54], Guo et al. use CNNs to predict 2D steady laminar flow fields around vehicles, achieving faster results than traditional solvers while maintaining accuracy. This enables rapid aerodynamic design iterations. A similar approach is discussed in ref. [55]. Deng et al. [56] propose an encoder-decoder method using a CFD database of 2D airfoils. The encoder compresses flow field images into feature vectors,

which are correlated with design variables using an MLP, integrating detailed flow information into the design process. Zalger[57] applies variational autoencoders (VAEs) to optimize aircraft turbomachinery components. The VAE encodes high dimensional CFD flow data into a lower-dimensional latent space, aiding efficient design space exploration while reducing computational costs. Wang et al. [58] combine VAEs with deep learning to predict flow fields around supercritical airfoils. Their method uses an encoder-decoder architecture, with an MLP mapping airfoil geometry to latent vectors, effectively predicting corresponding flow fields. In ref. [59], a VAE is applied to CFD data of multi-wedge hypersonic inlets. A hybrid model with an MLP reconstructs and predicts flow fields, correlating geometric parameters with latent space representations. Sekar et al. [60] introduce an inverse design method using CNNs to correlate 2D airfoil loading distributions with geometric features, enabling optimal airfoil geometry determination based on desired loading characteristics.

The reviewed studies highlight the potential of deep learning, particularly Convolutional Neural Networks (CNNs), in aerodynamic design and flow prediction. However, most of these works focus on 2D cases and simpler flow configurations, such as airfoils or vehicle bodies. While promising in general fluid dynamics, their direct application to 3D turbomachinery remains limited. The majority of studies examine relatively simple geometries and flow conditions, which may not fully capture the complexity of turbomachinery flows. Thus, while CNNs show potential in flow field prediction and design optimization, their use in more complex 3D turbomachinery systems requires further exploration.

This work involves training a database of centrifugal compressor impeller geometries, which includes global performance parameters and local flow data, in the form of blade load distributions along the chord and span. The goal is to develop a surrogate CFD model that can

provide accurate predictions of both the machine's performance and the blade loading, enabling faster and more efficient design iterations while capturing the complex flow characteristics of the impeller.

To achieve this, the first step is the creation of a database derived from CFD simulations. This database will then serve as the foundation for training the neural network. The geometry of the compressor is generated using a parameterization tool, as discussed before. The tool generates the 3D geometry of the centrifugal compressor, which is then automatically meshed with a viscous H-type grid. For each meshed geometry, a single operating point, imposing back-pressure, is run with CFD. As discussed before, convergence is considered reached when there is less than a 0.5% change in global performance values. Once convergence is reached, post-processing takes place.

In addition to the global parameters, detailed information regarding the blade loading is also collected. The blade loading is determined by the static pressure at each point on the blade surface, both on the pressure side and the suction side. To collect this data, a pressure map is created for each side of the blade, saved as an image. To generate the pressure maps, an averaging process is applied once the resolution of the pressure image is defined, ensuring that the pressure values are properly distributed across the image grid.

An example of image generation is reported on figure 5.2. These pressure maps are then stored in the database for each geometry, providing a clear representation of the blade loading distribution.

There is no strict requirement to use blade surface pressure as the sole feature. This choice is guided by human design experience, which recognizes blade loading as a critical factor in compressor performance. However, nothing prevents the inclusion, or even substitution, of other flow features that are equally important for machine design. Similarly, there is no restriction on employing three-dimensional flow fields instead of two-dimensional projections. However, this approach is gen-




Figure 5.2: Blade loading data collection example

erally avoided for multiple reasons: convolutional neural networks tend to achieve higher accuracy when trained on 2D data, and, at the same time, it is desirable that the extracted features retain a clear engineering meaning. Capturing the entire 3D flow field would provide excessive information without offering additional practical value for the design process.

The surrogate model input data consist of images with two channels, one corresponding to the pressure side (PS) and the other to the suction side (SS). The decoder network adopts a hybrid architecture that combines Multi-Layer Perceptron (MLP) and Convolutional Neural Network (CNN) layers. In the first stage, a fully connected MLP is employed, starting with an input layer followed by several hidden layers, all using the ReLU activation function. The output of this stage is then reshaped and passed into a series of CNN layers. These convolutional layers progressively refine the spatial resolution while reducing the number of channels, ultimately producing an output with two channels at maximum resolution.

The training procedure is carried out in a manner similar to the MLP: the dataset is split into training and test sets; the Mean Squared Error (MSE) is adopted as the loss function, computed at each pixel value; the weight updates are performed via backpropagation. The use

of GPU acceleration is suggested to significantly reduce training time and speed up convergence, but its use is not mandatory.





Figure 5.3: Schematic representation of the decoder neural network


The Figure 5.3 illustrates a schematic overview of the network architecture. Decoder results are presented in figure 5.4. The comparison is made between the true PS and SS images of the test set and their corresponding predicted images. The absolute difference is also computed to generate the corresponding difference images for PS and SS. As shown, the blade load configurations exhibit considerable variation across the test set; however, the decoder consistently predicts the blade behavior with high accuracy. This work presents cases involving negative incidence (fig. 5.4(a), 5.4(b)), low incidence (fig. 5.4(c), 5.4(d)), and positive incidence (fig. 5.4(e), 5.4(f)), highlighting the method's robustness in handling diverse scenarios.

By fixing a specific span, the chord-wise difference of pressure distributions (PS-SS) can be analyzed. The largest discrepancies are observed in regions with high gradients, such as the leading and trailing edges, where pressure variations are more pronounced. The trailing edge, in particular, shows a sharp variation due to its blunt shape. Additionally, in transonic conditions, strong gradients arise due to shock waves, especially in the presence of shocks. However, even in these re-


gions, the decoder’s performance remains impressive, with only minor deviations from the true values. The overall consistency of the predictions further demonstrates the decoder’s capability in capturing the complex aerodynamic behavior of the compressor blades across various operating conditions. Given the high accuracy in predicting blade loading, the method can be effectively employed to assess a priori the loading distribution of geometries obtained at the end of the optimization process, simply by querying the trained CNN. This additional information provides a valuable criterion for selecting which optimized geometries should be run with CFD across the operating curve. For instance, a solution exhibiting high efficiency at the design point but characterized by excessively positive incidence may only have a limited stable operating range on the left side of the map. An alternative approach consists in actively incorporating the predicted blade loading into the optimization loop itself. During the scoring phase, where each geometry is assigned a performance index, the evaluation can be extended beyond global performance metrics to include criteria related to loading quality, such as proximity to zero incidence or the presence of shock structures. By penalizing unfavorable loading patterns, this strategy can guide the optimization toward geometries with inherently better aerodynamic behavior. Such an approach has the potential to eliminate the need for explicitly including near-stall operating points in the database, thereby reducing the computational cost of optimization dramatically. Further details of this work can be found in [61].




((a)) Negative Incidence 1




((b)) Negative Incidence 2




((c)) Low Incidence 1



((d)) Low Incidence 2



((e)) Positive Incidence 1



((f)) Positive Incidence 2

Figure 5.4: CNN training results for different incidence examples (Test set).

5.3 Geometrical throat area estimation tool for maximum swallowing capacity estimation

In an optimization procedure based on neural networks and CFD databases, the identification of the right-hand limit of the operating curve is a critical task. The determination of the right margin condition requires CFD simulations for the database, performed only to evaluate the maximum mass flow rate. Other performance parameters, such as efficiency or pressure ratio, cannot be reliably used for optimization in this region. This is because, close to choke, the slope of the characteristic curves becomes almost vertical, and the relationship between flow rate and global performance is no longer unique.

For compressors operating at high flow coefficients and high Mach numbers, the right-hand limit of the characteristic curve is not determined by blade incidence, but rather by the onset of choke. This phenomenon occurs when the flow inside the blade passage accelerates up to Mach one at the throat section. Once this condition is reached, no further increase in mass flow rate is possible, since the flow cannot accelerate beyond this point.

Choke can be described using the classical one-dimensional theory of converging-diverging nozzles. For a perfect gas, the mass flow rate at choke is given by:

$$\dot{m}_{\max} = A_t P_0 \sqrt{\frac{\gamma}{RT_0}} \left(\frac{2}{\gamma + 1} \right)^{\frac{\gamma+1}{2(\gamma-1)}} \quad (5.1)$$

where:

- $\dot{m}_{\max}$ is the maximum mass flow rate at choke,
- A_t is the cross-sectional area at the throat (minimum area) of the blade passage,

- P_0 is the total (stagnation) pressure upstream of the throat,
- T_0 is the total (stagnation) temperature upstream of the throat,
- γ is the ratio of specific heats of the gas (c_p/c_v),
- R is the specific gas constant.

Although simplified, this formulation is directly analogous to what occurs inside the blade passage of a centrifugal compressor. The minimum cross-sectional area between blades behaves like the throat of a nozzle, and the maximum achievable mass flow rate is directly linked to this geometric feature. Once the working fluid and the rotational speed of the impeller are fixed, a nearly linear dependence emerges between the throat area and the maximum mass flow rate. This opens the way for estimating the choke condition from purely geometric information, without the need for dedicated CFD simulations.

It is important to clarify that a distinction exists between the minimum geometric area, which lies in a fixed reference plane, and the minimum effective throat area of the blade passage. Since the compressor flow field is inherently three-dimensional, the critical area that limits the mass flow is defined locally as the cross-section normal to the flow direction. Consequently, equating the minimum geometric throat area to the critical flow area represents an approximation, whose validity depends on how uniform and well-aligned the flow direction is within the throat region. In cases where the flow remains largely aligned, this simplification provides a sufficiently accurate estimate for engineering purposes, whereas strong three-dimensional effects may reduce its reliability.

The developed approach is based on a gradient descent approach [62]. In its standard formulation, gradient descent is an iterative optimization algorithm aimed at minimizing a given cost function. At each iteration, the cost is first evaluated based on the current set of

input variables. The gradient of the cost function is then approximated by perturbing each variable one at a time and observing the corresponding change in the cost value. This provides an estimate of the partial derivatives with respect to each variable, typically obtained through a forward-difference approximation. The gradient vector is thus assembled as the collection of these partial derivatives.

The gradient descent procedure can be summarized as follows:

1. Initialize the variables $\mathbf{x} = [x_1, x_2, \dots, x_n]$ and set the learning rate α .
2. Repeat until convergence:
 - Evaluate the cost function $L(\mathbf{x})$.
 - For each variable x_i :
 - Perturb x_i by a small amount δ .
 - Compute the approximate gradient component:

$$g_i \approx \frac{L(x_1, \dots, x_i + \delta, \dots, x_n) - L(\mathbf{x})}{\delta}$$

- Assemble the gradient vector $\mathbf{g} = [g_1, g_2, \dots, g_n]$.
- Update the variables:

$$\mathbf{x} = \mathbf{x} - \alpha \mathbf{g}$$

In this work, the same principle is applied to the estimation of the throat area within the blade passage. The cost function is defined as the geometric area resulting from the intersection between a candidate plane and the passage, while the variables to be optimized are the parameters describing the plane itself. By iteratively updating these variables in the opposite direction of the gradient, the method drives the solution toward the plane that corresponds to the minimum intersection area, which is taken as the throat.

Although gradient descent can, in general, converge to local minima, in this specific case the problem formulation is such that the throat area should be unique. Therefore, choosing an appropriate initial position for the plane provides a practical safeguard to ensure convergence toward the correct minimum.

An example of the method applied to a centrifugal impeller is shown in figure 5.5. The black plane represents the initial guess, while the red plane highlights the minimum area automatically identified by the algorithm. It is worth mentioning that the procedure operates independently of the passage complexity and has been successfully tested on three-dimensional impeller geometries, purely radial blade rows such as nozzles and vaned diffusers, as well as axial machines. The method has proven to be robust across these different configurations. The most critical aspect remains the definition of a suitable initial guess, which determines the accuracy and reliability of the final result.

To validate the effectiveness of the proposed method, both the choke mass flow rate obtained from CFD simulations and the corresponding geometric throat area were computed for an entire database of impellers. This allowed for a direct assessment of the accuracy of the approach. Figure 5.6 illustrates the resulting correlation, where the throat area is reported on the horizontal axis and the flow coefficient on the vertical axis. The agreement is remarkably strong, confirming that the geometric throat area can reliably replace CFD simulations for choke prediction.

The time savings are substantial: while CFD requires several hours to reach convergence, the throat area estimation is completed in a few minutes.

It should be underlined that the goal is not to precisely calculate the maximum flow coefficient by determining the throat area. As shown in 5.1, the maximum mass flow also depends on other factors related to the flow properties and the speedline, since total temperature and




Figure 5.5: Example of minimum area research. Black area = initial guess. Red Area = final area

pressure vary with the reference frame. During the optimization process, the throat area information is used as a filter for the Pareto front according to the target maximum flow coefficient. For instance, if the baseline flow coefficient is lower than the desired target, once its throat area is computed, a high-pass filter can be applied during optimization to keep only configurations with a throat area larger than the baseline value.




Figure 5.6: Geometrical throat area vs maximum flow coefficient of a transonic impeller database

Chapter 6

Continuous family generation

In the previous chapters, a step-by-step procedure for identifying an aerodynamically optimized impeller geometry has been described. In summary, the process begins with the definition of a geometric parametrization, followed by the creation of a numerical setup. Once the boundary conditions are imposed, CFD simulations are performed on perturbed geometries. The resulting post-processed data, combined with the corresponding geometric inputs, form a comprehensive database. This database is then employed to train a surrogate model based on neural networks. Finally, an optimization framework is applied to identify the best-performing geometries, which satisfy the imposed constraints while maximizing or minimizing a cost function, potentially incorporating multiple objectives.

An important limitation of the previously described methodology is that it is designed to optimize a single machine. While the approach is well-suited for generating an individual optimized impeller, extending it to the creation of entire families of machines presents additional challenges. In particular, it is not possible to directly reuse the same database or neural network for different machines. The perturbed geometries generated for one machine may be significantly different from the optimal geometries of another. For instance, a database created for

a machine operating at a given flow coefficient may not be sufficient for a new machine designed for a flow coefficient twice as large; the geometric perturbations applied to the inlet area are unlikely to capture the new design space. A similar consideration applies to boundary conditions, as the database is generated assuming a specific rotational speed. Variations in rotational speed can make the original database unrepresentative of the new operating conditions.

When generating a new family of machines, it is necessary to apply the previously described optimization procedure to a certain number of master designs, depending on the extent of the design space. The design space can be represented as shown in figure 6.1.




Figure 6.1: Design space

In general, the family may span a specific range of flow coefficients ϕ and peripheral Mach numbers at the impeller exit (Mu_2). It should be noted, however, that, as highlighted by Rusch and Casey [63], the complexity of the machine strongly depends on a third parameter: the relative Mach number at the impeller eye (M_{1r}). This parameter characterizes the nature of the flow near the tip, and depending on

whether the flow is subsonic, transonic, or supersonic, the resulting optimized geometries can exhibit significantly different shapes. For this reason, in the present work, the family has been defined along two dimensions: flow coefficient and impeller eye relative Mach number. In this study, a family at constant M_{1r} is considered, although the methodology can be extended to multiple values of M_{1r} .

The proposed methodology aims to generate a continuous set of geometries. Instead of relying solely on the optimized master designs and obtaining intermediate geometries through simple interpolation, a set of rules is defined for each geometrical feature. These rules allow for the automatic generation of intermediate geometries while maintaining the key characteristics of the master designs. The designer provides the design conditions, and the procedure generates the corresponding geometry, supporting the preservation of the performance attributes of the original databases.

The chosen example has supersonic inlet Mach number greater than one, thereby indicating transonic flow conditions at the impeller blade tip. Additionally, all generated machines have the same outlet diameter and the same target work coefficient at their own design flow coefficient.

Because the generation of this family is designed to be continuous, the designer is allowed to specify any value of the flow coefficient within the given range. As a result, adjacent geometries must smoothly connect to ensure a continuous transition. The aim is to produce machines that not only achieve similar performance levels but also exhibit consistent flow characteristics throughout the entire family. This approach ensures that each design meets specific aerodynamic requirements while providing the flexibility to adapt to varying flow conditions. To accomplish this, three distinct databases are employed, each corresponding to a family master designed with a specific flow coefficient: Φ_1 , Φ_2 and Φ_3 .

In order to define a set of rules for geometric generation, it is not

sufficient to simply rely on the geometric information of the optimized masters as the flow coefficient varies. In general, multiple geometric solutions exist along the Pareto front that exhibit excellent performance, but have significantly different geometric features. A straightforward interpolation of these solutions may produce intermediate geometries that are highly irregular, compromising the performance preservation. Therefore, it is necessary to conduct a study to identify which geometric features have the greatest impact on performance and which are relatively less influential. To do so, feature importance analysis can be leveraged.

6.1 Feature importance analysis

Feature importance analysis is a fundamental technique in machine learning and data science, designed to assess the contribution of individual input variables to a model's predictive performance. This method is essential for gaining insights into the behavior of complex models and plays a key role in improving model robustness by identifying and mitigating the influence of irrelevant or redundant features. This helps reduce dimensionality without compromising the model's ability to generalize effectively to new, unseen data.

One of the most widely used techniques for dimensionality reduction is Principal Component Analysis (PCA) [64]. PCA transforms the original dataset into a set of orthogonal components, known as principal components, which are linear combinations of the original features. These components are ordered such that the first few capture the maximum variance in the dataset. PCA is particularly effective in high-dimensional data settings, where it reduces the number of features while retaining the most critical information. However, in this context, PCA does not contribute to feature reduction due to the nature of the input dataset. Since the dataset is generated using a

Sobol sequence, which ensures a uniform and random distribution, no significant feature reduction is expected.

LASSO (Least Absolute Shrinkage and Selection Operator) [65] is a regularization technique derived from linear regression that incorporates an $L1$ penalty to shrink the coefficients of less important features to zero. By imposing a constraint on the sum of the absolute values of the regression coefficients, LASSO reduces model complexity and helps prevent overfitting. The LASSO estimator is formally defined as follows:

$$\hat{\beta} = \arg \min_{\beta} \left\{ \frac{1}{2n} \sum_{i=1}^n (y_i - \mathbf{x}_i^{\top} \beta)^2 + \lambda \sum_{j=1}^p |\beta_j| \right\} \quad (6.1)$$

where:

- n = number of observations in the dataset.
- p = number of features (independent variables).
- y_i = observed value of the dependent variable for the i -th observation.
- $\mathbf{x}_i$ = feature vector for the i -th observation.
- β = vector of regression coefficients.
- $\hat{\beta}$ = estimated coefficients obtained with LASSO.
- λ = regularization parameter controlling the strength of the penalty ($\lambda > 0$).
- $\sum_{j=1}^p |\beta_j|$ = $L1$ penalty term, which encourages sparsity by shrinking many coefficients to zero.

In this method, each output of the dataset is fitted using linear regression with LASSO, where the absolute value of the slope serves as an indicator of feature importance. However, both simple linear

regression and LASSO share a common limitation: they are sensitive to the scale of the input features. The magnitude of the coefficients in both models is influenced by the range of feature values, meaning that features with larger ranges may dominate the model, regardless of their actual importance. Consequently, feature scaling or normalization is typically required to ensure that all features contribute equally to the model and that the coefficients are comparable.

The shuffling method, also known as permutation importance, [66] is another approach for evaluating the significance of individual features in a trained model, particularly in artificial neural networks (ANNs). This technique involves randomly shuffling the values of a specific feature across all observations, thereby breaking its relationship with the target variable. The model's performance is then reassessed with the shuffled feature. A significant decline in performance indicates that the feature was crucial to the model's predictions. This method is particularly advantageous because it captures both linear and non-linear dependencies, making it highly effective for complex models. Unlike traditional linear regression techniques, the shuffling method can assess the importance of features in capturing intricate, non-linear relationships with the target variable. However, its effectiveness is heavily reliant on the quality and accuracy of the trained model; if the model is poorly trained or overfitting, the feature importance scores derived from this method may be misleading.

In this work the correlation matrix [67] is leveraged. It is an effective method for assessing feature importance, particularly when dealing with regression tasks. This approach evaluates the relationships between pairs of features in the dataset by calculating Pearson's correlation coefficient, which quantifies the degree of linear association between two variables.

The Pearson's correlation coefficient [68] between two variables X and Y is given by the formula:

$$r_{XY} = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2 \sum_{i=1}^n (Y_i - \bar{Y})^2}} \quad (6.2)$$

Where:


- X_i and Y_i are the individual data points for features X and Y ,
- $\bar{X}$ and $\bar{Y}$ are the means of features X and Y ,
- n is the number of data points.

The interpretation of the correlation value is as follows:

- a value close to +1 indicates a strong positive linear correlation, meaning that as X increases, Y tends to increase as well.
- a value close to 0 implies no linear correlation between X and Y .
- a value close to -1 indicates a strong negative linear correlation, meaning that as X increases, Y tends to decrease.

A key advantage of the correlation matrix is its independence from the range of input features. Unlike linear regression, Pearson's correlation coefficients are standardized within the range of -1 to +1. Moreover, the correlation matrix does not require a trained model, in contrast to the shuffling method. This characteristic makes the correlation matrix a straightforward yet effective tool for feature importance analysis, as it is both easy to compute and interpret. The correlation matrix directly reflects the linear dependencies between features, thereby facilitating the identification of redundant features and supporting feature reduction processes.

In this study, each of the three datasets (Φ_1 , Φ_2 and Φ_3) is analyzed using the correlation matrix. By arranging the input features in columns and the output features in rows, the resulting matrix is obtained in Fig. 6.2. The results are presented as a heatmap, where

Figure 6.2: Correlation Matrix of Φ_1 database

cells with higher absolute values of the correlation coefficients, hence stronger correlation, are represented with brighter colors.

By analyzing the matrix column-wise, it becomes apparent that certain geometric features have a greater impact than others, indicating that some features are likely to be reduced. Additionally, it is noteworthy that the majority of the significant geometric features influence all performance parameters to a considerable extent, rather than affecting only a subset of features. Conversely, by analyzing the matrix row-wise, insights can be gained regarding which geometric feature most strongly influences each individual performance metric. For instance, the flow coefficient at the choke point is expected to be primarily influenced by all degrees of freedom associated with the geometric throat areas of the impeller, whereas features such as the airfoil stacking shape appear to have minimal influence on this parameter.

The three datasets were compared, and the results, in terms of the magnitude of the correlation coefficients, were found to be similar. This supports the consistency of the method. Additionally, the presence of similar dependencies across different design flow coefficients allows for the establishment of geometry generation rules based on feature importance.





6.2 Geometry Generation

A geometry generation procedure, given a design flow coefficient, is now described. It is outlined that a preliminary phase was conducted to select the master geometries for each design flow coefficient Φ_1 , Φ_2 , and Φ_3 . The selected geometries serve as the basis for extracting the necessary geometric information to generate intermediate geometries. The three geometries, chosen based on the predictions from the ANNs and the optimization strategy, were subsequently validated through computational fluid dynamics (CFD) simulations to obtain performance curves and to ensure that the selected geometries exhibited similar performance characteristics. Additionally, since the impellers operating at the same impeller eye Mach number, it is expected that the three machines also exhibit comparable blade loadings.

The geometry generation process follows the scheme outlined below in Fig. 6.3. For each degree of freedom, a value is assigned that is dependent on the design flow coefficient. The function that maps the flow coefficient to the degree of freedom is derived from the three master geometries, taking into account its influence on performance as determined by the feature importance analysis. The effect of each feature is shown in the corresponding column of the feature importance table.

If a degree of freedom does not exhibit a strong correlation, feature reduction can be applied, and a constant value is assigned, calculated as the average of the three master geometries (Fig. 6.3(b)). Conversely, if a degree of freedom shows a strong correlation, the geometric parameter is determined using a spline interpolation of the three master geometries (Fig. 6.3(c)). In some cases, due to the independent selection of the master geometries, the spline function does not exhibit a monotonic trend within the domain of the design flow coefficient. When this occurs, linear regression of the three master values is applied to determine the geometric parameter in substitution of spline interpolation,

ensuring a more consistent and regular behavior (Fig. 6.3(d)). This procedure ensures that a geometry corresponding to an intermediate flow coefficient will have geometric features that are a hybrid of the features from the three master geometries.

((a)) Decision Chart for the generic parameter x_i ((b)) Case where x_i is not important - feature reduction((c)) Case where x_i is important - spline interpolation((d)) Case where x_i is important, but monotonic trend is required - linear regressionFigure 6.3: Geometry generation process scheme for a generic parameter x_i

6.3 Generation Results

After establishing the rules for assigning each degree of freedom, the performance of the generated geometries can be tested. A total of ten geometries are created, each corresponding to a different design flow coefficient. Among these, three geometries are set to Φ_1 , Φ_2 and Φ_3 , matching the design flow coefficients of the master geometries. It should be noted that, due to the feature reduction process and the use of linear regression instead of spline interpolation for certain geometric features, the geometries at these three flow coefficients differ slightly from the original designs. For consistency, Φ represents the original master, while φ is a member of the resulting family. Hence, $\Phi_1 = \varphi_1$, $\Phi_2 = \varphi_7$ and $\Phi_3 = \varphi_{10}$. Each performance curve is computed from choke to numerical stall.

The performance of the generated geometries is presented in Fig. 6.4 which illustrates the polytropic efficiency, work coefficient, and polytropic head, respectively.

It is noteworthy that, as shown in Fig. 6.4(a), all generated geometries exhibit a similar curve shape, indicating consistent aerodynamic behavior across the family. The efficiency values are remarkably close among the different geometries. However, a noticeable decrease in peak efficiency is observed as the flow coefficient increases. This trend can be attributed to the fact that all geometries share the same outer diameter, leading to a larger impeller eye diameter for higher Φ values. Consequently, the contribution of the isentropic centrifugal component to the overall work decreases, impacting the efficiency.

Moreover, blade thickness plays a significant role in influencing performance. Given that the impeller design features a blunt trailing edge, the intensity of the wake increases with trailing edge thickness. For mechanical integrity, a higher Φ necessitates a thicker blade profile, which in turn results in stronger wakes at the trailing edge. This trade-off between mechanical strength and aerodynamic performance becomes

more pronounced at higher flow coefficients, contributing to the observed reduction in efficiency.

The trend is highlighted in Fig. 6.5, where only the maximum efficiency operating point of each curve is displayed.

The comparison between the performance curves of the master geometries, Φ_1 , Φ_2 and Φ_3 , and their generated counterparts, φ_1 , φ_7 and φ_{10} , in Fig. 6.6 confirms that there is no significant decline in overall performance or any noticeable shift in the centering of the performance curves.

A more detailed comparison of the performance curves is presented in Fig. 6.7, where the flow coefficient is normalized to allow for a clearer evaluation of the operating range. It is noteworthy that all members of the generated family display similar curve shapes and operating ranges, indicating consistent performance across different geometries.

As previously discussed, all members of the generated family share the same impeller eye Mach number, implying a similar blade loading behavior across the geometries. Figure 6.8 illustrates the distribution of the normalized impeller pressure along the chord at three different span-wise locations, 20%, 50%, and 80%, with each geometry evaluated at its point of maximum efficiency. As expected, the transonic flow conditions lead to the formation of a passage shock at higher span levels, which is evident in the characteristic slope change observed on the suction side of Fig. 6.8(c).

Notably, despite variations in flow coefficients and geometric features among the different family members, they exhibit remarkably consistent pressure distribution profiles. Specifically, at 80% of relative span (Fig. 6.8(c)), the passage shock appears at nearly the same chord-wise location and with similar intensity across all geometries. Furthermore, the optimal incidence value for peak efficiency, indicated by the distance between the pressure and suction side curves at the leading edge, is consistently similar for each member of the family.

While shock compression is an effective method for increasing pressure, the interaction between the passage shock and the boundary layer leads to an increase in aerodynamic losses. As the intensity of the losses rises with the strength of the shock, it becomes crucial from an aerodynamic design perspective to avoid overly strong shocks.

Another valuable comparison between the geometries can be made by examining the flowfields. Figure 6.9 presents the normalized relative Mach number contours for four cases: φ_1 , φ_4 , φ_7 and φ_{10} . The comparison is made at 80% of relative span. As anticipated, the flowfield exhibits typical transonic flow behavior, characterized by the presence of a bow shock and passage shock structure ahead of the leading edge. Flow acceleration occurs in the uncovered part of the suction side followed by significant compression through the passage shock.

As observed in the comparison, the flowfields exhibit striking similarity across all four geometries. Despite variations in design flow coefficients and geometric features, the overall flow behavior remains consistent, highlighting the robustness of the generated family in terms of aerodynamic performance.

In conclusion, the procedure have proven to be extremely robust to generate geometries that all have similar performance characteristics, bot for operating range and flow features. The proposed procedure ensures that the generated geometries not only fit within the specified flow coefficient range but also retain aerodynamic efficiency and robustness. The methodology is based by the hypothesis that unifying geometries based on a consistent impeller eye relative Mach number will yield comparable aerodynamic behavior and efficiency across the family. The success of the method lies in its flexibility: it allows for the selection of different master geometries and the adaptation of feature reduction rules, providing a robust framework for optimizing design parameters while maintaining high performance. Further details of this work can be found in [69].




((a)) Polytropic efficiency η_p ((b)) Work Coefficient τ ((c)) Polytropic head H

Figure 6.4: Performance comparisons




Figure 6.5: Peak Efficiency Trend Along Input Flow Coefficient




Figure 6.6: Performance comparison with master geometries







Figure 6.7: Performance comparison using normalized flow coefficient



((a)) 20% span



((b)) 50% span



((c)) 80% span

Figure 6.8: Impeller blade load comparisons





((a)) φ_1 ((b)) φ_4 ((c)) φ_7 ((d)) φ_{10}

Figure 6.9: Normalized relative Mach number contours at peak efficiency point at 80% span

Conclusions and future work

In this thesis, several numerical tools have been developed for the design and optimization of high-Mach impellers.

A contribution has been made to turbulence modeling through the modification of a $k - \omega$ model based on helicity. This approach has demonstrated improved agreement between CFD simulations and experimental data, which were made available from Baker Hughes test cases. Additionally, the robustness increase of the model reduces the number of discarded solutions during database data collection.

In addition, deep learning methodologies have been presented and integrated into the design framework. Among these, a multilayer perceptron (MLP) neural network developed with PyTorch has been employed to correlate geometric features (degrees of freedom, DOF) with the global performance of the machine. Once a surrogate model with satisfactory accuracy was obtained, it was embedded into a multi-objective optimization tool to accelerate the exploration of the design space.

Furthermore, additional tools were introduced to enhance the optimization process. A convolutional neural network was trained to predict blade loading at the design point, and a dedicated tool was developed to identify the throat area of the blade channel. The first tool shows potential for maximizing the operating range at low flow rates, thereby replacing the need for generating a near-stall database. The second tool, assuming a strong correlation between throat area

and maximum mass flow capacity, can be used as an alternative to a near-choke database.

A method for the continuous generation of impeller geometries has also been introduced. Instead of relying on a finite number of family masters, rules have been defined for the assignment of geometric parameters as a function of the desired design point. This process is based on feature importance analysis and leverages previously generated databases, which had been used to obtain the optimized masters.

Overall, the proposed methodologies demonstrate the effectiveness of combining CFD and machine learning techniques for advancing the design and optimization of turbomachinery components.

Future developments of this research may proceed along several directions. Enhancement of surrogate modeling techniques by exploring more advanced architectures. Extension of the optimization framework to account for manufacturing constraints and structural considerations, in order to bring the design closer to industrial application. Exploration of additional flow features that can be predicted using Convolutional Neural Networks (CNNs), to assess whether they provide further insights for the optimization process beyond blade loading. Extension of the optimization framework to stator components, such as vaned diffusers and return channels, in order to perform full-stage optimization rather than focusing solely on individual components. Expansion of the continuous impeller geometry generation tool to cover a larger design space, beyond a single line of $M_{1r} = \text{const}$. An extension of this methodology is planned for machines with varying impeller eye relative Mach numbers. This will involve applying the same framework and performing interpolation between two sets of curves, enabling the generation of geometries based on two input parameters, rather than solely relying on the design flow coefficient. This enhancement could further broaden the applicability of the method, allowing for a wider range of operating conditions.

An additional promising improvement, which could significantly reduce the generation time of the family, concerns the reuse of a database created during the optimization of a primary master for the optimization of other masters within the family. To the authors' knowledge, such techniques, which are related to transfer learning, have not yet been reported in the literature for similar applications.

Furthermore, the information collected during database construction could be leveraged to estimate the impact of production non-conformities on performance, providing valuable insights for design robustness and quality control.

Acknowledgments

The University of Florence is kindly acknowledged for the opportunity to develop this thesis and the related work.

Baker Hughes is acknowledged for the collaboration and for granting permission to publish these results. Special thanks are extended to 3Blue1Brown [70] and Daniel Bourke [71] for their open educational content, which inspired the development of the deep learning tools.

Finally, the Tgroup is most deeply acknowledged for their fundamental support and guidance throughout these years, which has been essential for the completion of this work.

Bibliography

- [1] <https://ukcop26.org/the-glasgow-climate-pact/>.
- [2] A. Giovannelli. “Development of turbomachines for renewable energy systems and energy-saving applications”. In: vol. 153. 2018, pp. 10–15. DOI: <https://doi.org/10.1016/j.egypro.2018.10.069>.
- [3] M. Casey and C. Robinson. “Some Properties of the Exit Velocity Triangle of a Radial Compressor Impeller”. In: *Journal of Turbomachinery* 145.5 (Nov. 2022). DOI: 10.1115/1.4056088.
- [4] M. Casey and C. Robinson. “A New Streamline Curvature Through-flow Method for Radial Turbomachinery”. In: *Journal of Turbomachinery* 132.3 (Apr. 2010). DOI: 10.1115/1.3151601.
- [5] H. Hazby, C. Robinson, M. Casey, D. Rusch, and R. Hunziker. “Free-Form Versus Ruled Inducer Design in a Transonic Centrifugal Impeller”. In: Turbo Expo Volume 2C: Turbomachinery (June 2017). DOI: 10.1115/1.4038176.
- [6] R. Uhrig. “Introduction to artificial neural networks”. In: vol. 1: Proceedings of IECON '95 - 21st Annual Conference on IEEE Industrial Electronics. 1995, pp. 33–37. DOI: 10.1109/IECON.1995.483329.
- [7] A. Agnolucci, M. Marconcini, A. Arnone, L. Toni, A. Grimaldi, and M. Giachi. “Centrifugal Compressor Stage Efficiency and Rotor Stiffness Augmentation via Artificial Neural Networks”. In:

- vol. Volume 2D. Turbo Expo. June 2021. DOI: 10.1115/GT2021-59998.
- [8] A. Agnolucci, M. Marconcini, D. Biliotti, L. Toni, A. Baroni, and L. Romani. “Robust Design of a Vaned Diffuser in a Centrifugal Compressor Stage With a High Load Impeller”. In: vol. Volume 13D. Turbo Expo. June 2023, V13DT35A022. DOI: 10.1115/GT2023-103585.
- [9] A. Hehn, M. Mosdzien, D. Grates, and P. Jeschke. “Aerodynamic optimization of a transonic centrifugal compressor by using arbitrary blade surfaces”. In: *ASME J. Turbomach.* Vol. 140. 5. DOI:10.1115/1.4038908. 2018.
- [10] S. Pierret and R. Van den Braembussche. “Turbomachinery blade design using a Navier–Stokes solver and artificial neural network”. In: *ASME J. Turbomach.* Vol. 121. 2. DOI:10.1115/1.2841318. 1999.
- [11] S. L. Dixon. *Fluid Mechanics and Thermodynamics of Turbomachinery: Seventh Edition*. Oct. 2013.
- [12] R. H. Aungier. *Centrifugal Compressors: A Strategy for Aerodynamic Design and Analysis*. ASME Press, Jan. 2000. DOI: 10.1115/1.800938.
- [13] *Garret products website*. <https://www.garrettmotion.com/racing-and-performance/performance-turbos/>.
- [14] *Baker Hughes products website*. <https://www.bakerhughes.com/centrifugal-axial-compressors/high-pressure-ratio-compressor-hprc>.
- [15] A. Arnone, M.S. Liou, and L. A. Povinelli. “Multigrid calculation of three-dimensional viscous cascade flows”. In: *J. Propul. Power* 9.4 (1993), pp. 605–614. DOI: 10.2514/3.23664.

- [16] A. Arnone. “Viscous analysis of three-dimensional rotor flow using a multigrid method”. In: *ASME J. Turbomach.* 116 (1994), pp. 435–445. DOI: 10.1115/1.2929430.
- [17] M. Marconcini, F. Rubechini, A. Arnone, and S. Ibaraki. “Numerical investigation of a transonic centrifugal compressor”. In: *ASME J. Turbomach.* 130.1 (2008), p. 011010. DOI: 10.1115/1.2752186.
- [18] M. Marconcini, F. Rubechini, A. Arnone, and S. Ibaraki. “Numerical analysis of the vaned diffuser of a transonic centrifugal compressor”. In: *ASME J. Turbomach.* 132.4 (2010), p. 041012. DOI: 10.1115/1.2988481.
- [19] M. Marconcini, A. Bianchini, M. Checcucci, G. Ferrara, A. Arnone, L. Ferrari, D. Biliotti, and D. T. Rubino. “A three-dimensional time-accurate computational fluid dynamics simulation of the flow field inside a vaneless diffuser during rotating stall conditions”. In: *ASME J. Turbomach.* 139.2 (2017), p. 021001. DOI: 10.1115/1.4034633.
- [20] M. Batisti, M. Marconcini, A. Agnolucci, E. Belardini, A. Grimaldi, L. Toni, R. Valente, and A. Arnone. “Experimental and Numerical Rotating Stall Investigation of a Multistage Centrifugal Compressor”. In: *Turbo Expo: Power for Land, Sea, and Air.* 2025. DOI: 10.1115/GT2025-153766.
- [21] R. Pacciani, M. Marconcini, and A. Arnone. “Comparison of the AUSM⁺-up and Other Advection Schemes for Turbomachinery Applications”. In: *Shock Waves* 29.5 (2019), pp. 705–716. DOI: 10.1007/s00193-018-0883-4.
- [22] A. Jameson, W. Schmidt, and E. Turkel. “Numerical solution of the Euler equations by finite volume methods using Runge Kutta time stepping schemes”. In: *14th Fluid and Plasma Dynamics Conference.* 1981. DOI: 10.2514/6.1981-1259.

- [23] B. Baldwin and H. Lomax. “Thin-layer approximation and algebraic model for separated turbulent flows”. In: *16th Aerospace Sciences Meeting*. 1978. DOI: 10.2514/6.1978-257.
- [24] D. Degani and L. B. Schiff. “Computation of turbulent supersonic flows around pointed bodies having crossflow separation”. In: *Journal of Computational Physics* 66.1 (1986), pp. 173–196. DOI: [https://doi.org/10.1016/0021-9991\(86\)90059-8](https://doi.org/10.1016/0021-9991(86)90059-8).
- [25] A. Arnone and R. Pacciani. “IGV–Rotor Interaction Analysis in a Transonic Compressor Using the Navier–Stokes Equations”. In: *Journal of Turbomachinery* 120.1 (Jan. 1998), pp. 147–155. DOI: 10.1115/1.2841376.
- [26] P. Spalart and S. Allmaras. “A One-Equation Turbulence Model for Aerodynamic Flows”. In: *AIAA* 439 (Jan. 1992). DOI: 10.2514/6.1992-439.
- [27] P. R. Spalart and M. Shur. “On the sensitization of turbulence models to rotation and curvature”. In: *Aerospace Science and Technology* 1.5 (1997), pp. 297–302. DOI: [https://doi.org/10.1016/S1270-9638\(97\)90051-1](https://doi.org/10.1016/S1270-9638(97)90051-1).
- [28] D. C. Wilcox et al. *Turbulence modeling for CFD*. Vol. 2. La Canada, CA: DCW Industries, 1998.
- [29] F. R. Menter. “Two-equation eddy-viscosity turbulence models for engineering applications”. In: *AIAA Journal* 32.8 (1994), pp. 1598–1605. DOI: 10.2514/3.12149.
- [30] D. C. Wilcox. “Formulation of the $k-\omega$ turbulence model revisited”. In: *AIAA Journal* 46.11 (2008), pp. 2823–2838. DOI: 10.2514/1.36541.
- [31] T. Rung et al. “Assessment of Explicit Algebraic Stress Models in Transonic Flows”. In: *Engineering Turbulence Modelling and Experiments* 4. 1999, pp. 659–668. DOI: <https://doi.org/10.1016/B978-008043328-8/50063-1>.

- [32] M. Franke, S. Wallin, and F. Thiele. “Assessment of explicit algebraic Reynolds-stress turbulence models in aerodynamic computations”. In: *Aerospace Science and Technology* 9.7 (2005), pp. 573–581. DOI: <https://doi.org/10.1016/j.ast.2005.06.002>.
- [33] F. Poli, M. Marconcini, R. Pacciani, D. Magarielli, E. Spano, and A. Arnone. “Exploiting GPU-Based HPC Architectures to Accelerate an Unsteady CFD Solver for Turbomachinery Applications”. In: vol. Volume 10C: Turbomachinery — Design Methods and CFD Modeling for Turbomachinery; Ducts, Noise, and Component Interactions. Turbo Expo. June 2022, V10CT32A035. DOI: 10.1115/GT2022-82569.
- [34] H. K. Moffatt. “The Degree of Knottedness of Tangled Vortex Lines”. In: *Journal of Fluid Mechanics* 35.1 (1969), pp. 117–129. DOI: 10.1017/S0022112069000991.
- [35] Y. W. Liu, L. P. Lu, L. Fang, and F. Gao. “Modification of Spalart-Allmaras Model with Consideration of Turbulence Energy Backscatter Using Velocity Helicity”. In: *Physics Letters A* 375.24 (2011), pp. 2377–2381. DOI: <https://doi.org/10.1016/j.physleta.2011.05.023>.
- [36] Y. Liu, Y. Tang, A. D. Scillitoe, and P. G. Tucker. “Modification of Shear Stress Transport Turbulence Model Using Helicity for Predicting Corner Separation Flow in a Linear Compressor Cascade”. In: *Journal of Turbomachinery* 142.2 (Jan. 2020), p. 021004. DOI: 10.1115/1.4045658.
- [37] K. B. Lee, M. Wilson, and M. Vahdati. “Validation of a Numerical Model for Predicting Stalled Flows in a Low-Speed Fan - Part I: Modification of Spalart-Allmaras Turbulence Model”. In: *Journal of Turbomachinery* 140.5 (Apr. 2018), p. 051008. DOI: 10.1115/1.4039051.

- [38] A. Pela, M. Marconcini, A. Arnone, L. Toni, R. Valente, A. Agnolucci, A. Grimaldi, and R. Pacciani. “On the feasibility of helicity-based corrections of turbulence models for improving RANS predictions of centrifugal compressor stages”. In: vol. Volume 13D. Turbo Expo. June 2023, V13DT35A026. DOI: 10.1115/GT2023-103852.
- [39] L. Toni, V. Ballarini, S. Cioncolini, G. Persico, and P. Gaetani. “Unsteady Flow Field Measurements in an Industrial Centrifugal Compressor”. In: *Proceedings of the 39th Turbomachinery Symposium*. Houston, USA, 2010.
- [40] A. Bianchini, G. Ferrara, L. Ferrari, V. Ballarini, L. Bianchi, L. Tapinassi, and L. Toni. “Effects Due to the Temperature Measurement Section on the Performance Estimation of a Centrifugal Compressor Stage”. In: *Journal of Engineering for Gas Turbines and Power* 134.3 (Jan. 2012), p. 032402. DOI: 10.1115/1.4004722.
- [41] E. Guidotti, L. Tapinassi, L. Toni, L. Bianchi, P. Gaetani, and G. Persico. “Experimental and Numerical Analysis of the Flow Field in the Impeller of a Centrifugal Compressor Stage at Design Point”. In: vol. Volume 7: Turbomachinery, Parts A, B, and C. Turbo Expo. June 2011, pp. 1845–1856. DOI: 10.1115/GT2011-45036.
- [42] A. Bianchini, D. Biliotti, G. Ferrara, L. Ferrari, E. Belardini, M. Giachi, and L. Tapinassi. “Some Guidelines for the Experimental Characterization of Vaneless Diffuser Rotating Stall in Stages of Industrial Centrifugal Compressors”. In: vol. Volume 2D: Turbomachinery. Turbo Expo. June 2014, V02DT42A028. DOI: 10.1115/GT2014-26401.
- [43] A. Bianchini, D. Biliotti, D. T. Ferrara Rubino, L. Ferrari, and G. Ferrara. “Experimental Analysis of the Pressure Field Inside

- a Vaneless Diffuser from Rotating Stall Inception to Surge”. In: *Journal of Turbomachinery* 137.11 (Sept. 2015), p. 111007. DOI: 10.1115/1.4031354.
- [44] D. Biliotti, A. Bianchini, G. Ferrara, L. Ferrari, E. Belardini, M. Giachi, L. Tapinassi, and G. Vannini. “Analysis of the Rotordynamic Response of a Centrifugal Compressor Subject to Aerodynamic Loads Due to Rotating Stall”. In: *Journal of Turbomachinery* 137.2 (Sept. 2014), p. 021002. DOI: 10.1115/1.4028246.
- [45] D. Biliotti, L. Toni, E. Belardini, G. Vannini, M. Giachi, and D. T. Rubino. “Stall Induced Aerodynamic Forcing and Rotor Vibrations in a Multistage Centrifugal Compressor”. In: *3rd Middle East Turbomachinery Symposium (METS III)*. 2015.
- [46] W. Wahba and A. Tourlidakis. “A Genetic Algorithm Applied to the Design of Blade Profiles for Centrifugal Pump Impellers”. In: *15th AIAA Computational Fluid Dynamics Conference*. 2001. DOI: 10.2514/6.2001-2582.
- [47] D. Bonaiuti and V. Pediroda. “Aerodynamic Optimization of an Industrial Centrifugal Compressor Impeller Using Genetic Algorithm”. In: *Proc. Eurogen 2001, Evolutionary Methods for Design, Optimization and Control*. 2001, pp. 467–472.
- [48] D. Bonaiuti, A. Arnone, M. Ermini, and L. Baldassarre. “Analysis and Optimization of Transonic Centrifugal Compressor Impellers Using the Design of Experiments Technique”. In: *Journal of Turbomachinery* 128.4 (Sept. 2006), pp. 786–797. DOI: 10.1115/1.1579507.
- [49] T. Verstraete, L. Müller, and J. Müller. “Adjoint-Based Design Optimisation of an Internal Cooling Channel U-Bend for Minimised Pressure Losses”. In: *International Journal of Turbomachinery, Propulsion and Power* 2 (June 2017). DOI: 10.3390/ijtp2020010.

- [50] I. M. Sobol. “Distribution of Points in a Cube and Approximate Evaluation of Integrals”. In: *USSR Computational Mathematics and Mathematical Physics* 7.4 (1967), pp. 86–112. DOI: [https://doi.org/10.1016/0041-5553\(67\)90144-9](https://doi.org/10.1016/0041-5553(67)90144-9).
- [51] D. P. Kingma and J. Ba. “Adam: A Method for Stochastic Optimization”. In: *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*. San Diego, CA, USA, 2015. DOI: 10.48550/arXiv.1412.6980.
- [52] A. Pela, M. Marconcini, A. Arnone, A. Agnolucci, L. Toni, R. Valente, and E. Belardini. “Exploiting Deep Learning for the Optimization of Transonic Centrifugal Impellers”. In: *J. Phys.: Conf. Ser, The 79th ATI Annual Congress, Genoa, Italy*. Vol. 2893. 2024, p. 012126. DOI: 10.1088/1742-6596/2893/1/012126.
- [53] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. “Backpropagation Applied to Handwritten Zip Code Recognition”. In: *Neural Computation* 1 (Dec. 1989), pp. 541–551. DOI: 10.1162/neco.1989.1.4.541.
- [54] X. Guo, W. Li, and F. Iorio. “Convolutional Neural Networks for Steady Flow Approximation”. In: Association for Computing Machinery, Aug. 2016, pp. 481–490. DOI: 10.1145/2939672.2939738.
- [55] S. Chickerur and P. Ashish. “A Convolutional Neural Network Based Approach for Computational Fluid Dynamics”. In: *2021 Second International Conference on Smart Technologies in Computing, Electrical and Electronics (ICSTCEE)*. Dec. 2021, pp. 1–5. DOI: 10.1109/ICSTCEE54422.2021.9708548.
- [56] K. Deng, H. Chen, and Y. Zhang. “Flow structure oriented optimization aided by deep neural network”. In: *Proceedings of the Tenth International Conference on Computational Fluid Dynamics (ICCFD10)*. Barcelona, Spain, 2018, pp. 9–13.

- [57] Jonathan M Zalger. “Application of variational autoencoders for aircraft turbomachinery design”. In: 2017.
- [58] J. Wang, C. He, R. Li, H. Chen, C. Zhai, and M. Zhang. “Flow field prediction of supercritical airfoils via variational autoencoder based deep learning framework”. In: *Physics of Fluids* 33 (Aug. 2021). DOI: 10.1063/5.0053979.
- [59] Z. Tan, R. Li, and Y. Zhang. “Flow Field Reconstruction of 2D Hypersonic Inlets Based on a Variational Autoencoder”. In: *Aerospace* 10.9 (2023). DOI: 10.3390/aerospace10090825.
- [60] V. Sekar, M. Zhang, C. Shu, and B. Khoo. “Inverse Design of Airfoil Using a Deep Convolutional Neural Network”. In: *AIAA Journal* 57.3 (2019), pp. 993–1003. DOI: 10.2514/1.J057894.
- [61] A. Pela, M. Marconcini, A. Agnolucci, E. Belardini, A. Grimaldi, L. Toni, R. Valente, and A. Arnone. “Convolutional Neural Network Approach for Impeller Blade Loading Inference”. In: 1st International Symposium on AI and Fluid Mechanics (AIFLU-IDs). May 2025.
- [62] Sebastian Ruder. “An overview of gradient descent optimization algorithms”. In: *ArXiv* (2016). DOI: 10.48550/arXiv.1609.04747.
- [63] Daniel Rusch and Michael Casey. “The Design Space Boundaries for High Flow Capacity Centrifugal Compressors”. In: *Journal of Turbomachinery* 135.3 (Mar. 2013). DOI: 10.1115/1.4007548.
- [64] A. Maćkiewicz and W. Ratajczak. “Principal components analysis (PCA)”. In: *Computers & Geosciences* 19.3 (1993), pp. 303–342. DOI: 10.1016/0098-3004(93)90090-R.
- [65] R Tibshirani. “Regression Shrinkage and Selection via the Lasso”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 58.1 (1996), pp. 267–288.

-
- [66] L. Breiman. “Random Forests”. In: *Machine Learning* 45 (Oct. 2001), pp. 5–32. DOI: 10.1023/A:1010950718922.
- [67] J. Han, M. Kamber, and J. Pei. *Data Mining: Concepts and Techniques*. 3rd. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2011. DOI: 10.1016/C2009-0-61819-5.
- [68] *Correlation Coefficient: Simple Definition, Formula, Easy Steps*. Available at: <https://www.statisticshowto.com/probability-and-statistics/correlation-coefficient-formula/>. 2024.
- [69] A. Pela, M. Marconcini, A. Agnolucci, E. Belardini, A. Grimaldi, L. Toni, R. Valente, and A. Arnone. “Towards the Automatic Generation of Transonic Centrifugal Impellers by Deep Learning and Feature Importance Techniques”. In: vol. Volume 12. Turbo Expo. June 2025. DOI: 10.1115/GT2025-153947.
- [70] <https://www.3blue1brown.com/>.
- [71] <https://www.mrdbourke.com/>.