



UNIVERSITÀ
DEGLI STUDI
DI PADOVA



MSC PROGRAMME IN MECHANICAL ENGINEERING

BIANCHIN GIACOMO

EVOLUTIONARY MULTI-OBJECTIVE OPTIMIZATION OF A TRANSONIC AEROENGINE
GAS-TURBINE VANE

MASTER'S THESIS

SUPERVISORS:

PROF. F. DE VANNA
DR. F. AVANZI

ACADEMIC YEAR 2025-2026

Contents

- 1 Introduction** **5**

- 2 Methods** **9**
 - 2.1 Geometry and Discretization 9
 - 2.2 Numerical scheme and boundary conditions 11
 - 2.3 Model sensitivity and accuracy 13
 - 2.4 Parametric model 17
 - 2.5 Optimization problem formulation 19

- 3 Results and Discussion** **25**

- 4 Conclusions** **39**

- Bibliography** **41**

- 5 Acknowledgements** **47**

Introduction

The environmental directives regulating pollution in the aviation field have gradually become more restrictive, imposing increasing reductions in fuel consumption to reduce CO₂ and NO_x emissions. The major effects of these policies have stimulated the research of high efficiency solutions for the propulsion system, recognized as the primary contributor to aircraft-derived pollution. Specifically, improving machine design can result in relevant advancements, beneficial for both the environmental footprint and the mission costs. To achieve this objective, it has been necessary to increase the operating temperatures—a design choice that benefits thermal cycle efficiency but requires high speed flows to rapidly eject hot gases, thus reducing thermal exposure. Consequently, transonic gas-turbines have gradually become a common configuration for aeroengines, with an increasing interest towards even faster solutions. However, as the operating Mach number raises, compressible effects strengthen, introducing augmented blade loading and more complex flow phenomena, like shockwaves, that destabilize boundary layer dynamics and compromise wake evolution [3, 11]. Despite representing challenging operating conditions for testing, significant research effort has been dedicated to reproduce and characterize fluid mechanisms through transonic gas-turbine cascades. Dissipative effects mainly result from suction side viscous flows dynamics, where the boundary layer may interact with local and passage shocks, further compromising the wake development downstream with direct detrimental effects on the cascade total pressure losses [22, 6, 14]. Therefore, several strategies have been endeavored to control and prevent critical interactions, many of these relying on the evident dependency of these phenomena on loading

variations, different Mach numbers and manipulated geometrical parameters [4, 12, 20, 25, 16]. However, the collection of dataset has been limited by the experimental resolution demands, set by the highly unsteady nature and strong local gradients typical of transonic regimes. This explains why, while relevant correlations have been investigated for Mach, Reynolds number and incidence angle, an extensive understanding is still an open research field [1]. In this landscape, the Von Karman Institute LS89 cascade remains a canonical benchmark thanks to detailed wall and wake data that support validation and design-oriented analyses [3, 11]. Off-design behavior is equally relevant for aero-engines: systematic variations of incidence, Reynolds number, and inflow conditions have been shown to trigger separation and rapidly increase losses, with transonic cascade campaigns linking trends to loading indicators such as the Zweifel coefficient [31, 13, 10, 21, 5]. Scale-resolving simulations based on LES, including wall-modeled formulations, can reproduce key experimental features of transonic turbine cascades more faithfully than RANS, as shown in recent studies [8, 33]. However, their computational cost is still prohibitive for routine CFD-in-the-loop design and optimization. Consequently, systematic shape design is most often pursued via RANS. For transonic turbines, the adequacy of RANS is strongly influenced by the turbulence and transition models. Michelassi et al. showed that introducing transition models is necessary to obtain a realistic description of the flow field in a transonic turbine rotor cascade, and that such models can provide adequate predictions of shock-induced separation trends [23]. Consequently, validated RANS setups remain the workhorse for parametric investigations and optimization studies, albeit still expensive when embedded in CFD-in-the-loop design. Because CFD-in-the-loop optimization remains expensive, prior studies have mainly acted on two levers. The first lever is addressed through compact and physically meaningful parametrizations that reduce the effective number of decision variables while preserving manufacturability and aerodynamic plausibility [2]. A complementary route is to derive constraints from pre-existing shape databases so as to prevent unrealistic geometries, reducing the trial-and-error of the optimization process. The main limitation is the need for a representative database to extract the relevant features [18]. The second lever concerns the optimization strategy by reducing the number of expensive CFD calls. Adjoint and gradient-based methods enable the use of optimization algorithms that often lead to faster convergence compared to gradient-free optimization algorithms [34, 17]. Their practical application, however, may require a significant implementation ef-

fort. To combine fast local improvement with robust global exploration, several hybrid approaches have been proposed in which adjoint information is coupled with gradient-free optimization algorithms [35, 27, 26]. In parallel, surrogate-assisted and data-driven strategies reduce objective function calculation time by replacing a fraction of the expensive CFD evaluations with fast models trained on previously computed designs [29, 32, 19]. Within this landscape, evolutionary algorithms remain attractive for multi-objective problems and complex constraint handling due to their ease of deployment and robustness, but they typically require a large number of CFD evaluations to approach a converged Pareto front [30]. On the Von Karman Institute LS89 transonic cascade, Sanchez Torreguitart et al. presented a single-point, single-objective optimization aimed at loss reduction, based on a CAD parametrization and an adjoint and gradient-driven approach [28]. Hottois et al. then compared gradient-free and gradient-based strategies in a multi-objective setting, still targeting improved aerodynamic performance through reduced losses, and showed that gradient-based methods can achieve comparable Pareto-optimal solutions at a reduced computational effort [15]. Finally, hybrid optimization frameworks have also been investigated on LS89 to combine global exploration with fast local improvement within a single-objective formulation, again with the common goal of reducing cascade losses and improving overall performance [7]. Overall, these studies are primarily loss-driven and show that extending LS89 optimisation from scalar loss minimisation to multi-objective formulations is inherently more demanding and therefore increases both computational effort and the complexity of the design trade-off. Within this context, evolutionary algorithms remain attractive for multi-objective problems and complex constraints, but typically require a considerable number of CFD evaluations to approach a converged Pareto front [30]. Exploiting the physics-based nature of evolutionary techniques and the computational efficiency of the two-dimensional cascades assumption, the present work advances the current stage of optimization studies dedicated to the LS89 transonic nozzle vanes. Additional local accuracy in geometry parameterization is introduced adopting a hybrid curve topology approach. The numerical setup is validated under a single-channel periodic configuration, and systematic simulations are performed for several geometrical variants, generated and controlled by a genetic algorithm. A two-objective minimization problem is formulated aiming at improvements in the expansion ratio and the cascade loss coefficient. The resulting dataset is analyzed to extract relevant correlations between geometric

properties and objective functions. After selecting optimized individuals with opposite aerodynamic performance, a systematic investigation is conducted to characterize the general flow dynamics and the impact of local interactions between compressible and viscous phenomena.

The remainder of the paper is organized as follows. Chapter 2 describes the computational setup, boundary conditions, and the optimization loop. Chapter 3 presents and discusses the results. Chapter 4 summarizes the main conclusions.

Methods

2.1 Geometry and Discretization

The blade coordinates, the pitch t , and the chord c are taken directly from the Von Karman Institute MUR47 test case [3] and are used without modification. The main geometric parameters of the LS89 cascade are reported in Table 2.1.

Table 2.1: Geometric parameters of the LS89 cascade from the Von Karman Institute MUR47 test case.

Parameter	Variable	Value	Unit
Chord	c	67.647	mm
Axial chord	c_{ax}	38.8	mm
Stagger angle	γ_{stag}	55	deg
Pitch	t	57.46	mm

Under the 2D cascade assumption, a Cartesian frame is adopted, with the x -axis aligned along the axial direction and the y -axis parallel to the pitch direction. Unless otherwise stated, coordinates are expressed in the cascade plane and normalized by the chord c . All geometric quantities used in this work are illustrated in Fig. 2.1a. The reference camber line is parameterized by a non-dimensional curvilinear coordinate $u \in [0, 1]$ measured between the leading and the trailing edge. The computational domain is discretized, considering a single blade passage, with pitchwise translational periodicity enforced on the upper




Figure 2.1: Relevant geometrical parameters of the case study a and near-wall details of the computational grid around the leading b and trailing c edges.

and lower boundaries of the control volume. A structured multi-block mesh is employed, starting from a near-wall O-grid generated by normal extrusion from the blade surface. The first cell height is set to achieve a target $y^+ < 0.5$ over the entire blade surface. The nodes are progressively clustered near the leading and trailing edges to improve accuracy in capturing shape curvature and flow field

gradients. The mesh size is systematically increased to generate three refinements, coarse, medium and fine, respectively accounting for 1.0×10^5 , 2.0×10^5 , and 4.0×10^5 cells. Compared to the coarse grid, the medium grid leverages increased resolution in the suction side region near the trailing edge, where the shockwave occurs, with limited increment of nodes along the O-grid block. The fine grid is obtained from the medium one through a uniform refinement of both the O-grid and the shock region. The corresponding mesh quality indicators are summarized in Table 2.2.

Table 2.2: Mesh quality indicators for the three grid levels. Lower skewness and area ratios closer to 1 indicate better quality.

Grid	Skewness _{avg}	Area ratio _{avg}
Coarse	0.072	1.07
Medium	0.066	1.06
Fine	0.067	1.05

2.2 Numerical scheme and boundary conditions

Steady-state RANS equations are solved using the pressure-based finite-volume solver ANSYS Fluent[®] with a coupled pressure/velocity scheme, assuming a compressible working fluid, with density and viscosity varying according to ideal-gas and Sutherland laws, respectively. Conservation variables and turbulence terms are discretized with second-order accuracy. A pressure-inlet boundary condition is set at the inlet, thus enforcing stagnation quantities, while pressure-outlet is adopted at the outflow boundary. Blade surfaces are treated as no-slip walls. At the inlet, the turbulence quantities are prescribed by imposing a turbulence intensity of $T_u = 5\%$ together with a turbulent viscosity ratio $\mu_t/\mu = 10$. Simulations are initialized with a hybrid technique and a first set of 400 iterations using first-order accuracy is adopted to stabilize the solution. Then, schemes resolution is switched to second-order up to convergence. All under-relaxation factors are kept at default values and gradient reconstruction is approached with a least-squares cell-based method. Two convergence criteria are chosen to stop iterations. First, all normalized residuals are required to fall below 10^{-3} . Second, a physics-based criterion is introduced: the latter monitors the blade drag ratio

residuals until they become lower than 1×10^{-2} , within a sliding window of 300 consecutive iterations. Here, the drag ratio is defined as:

$$D_{\text{ratio}} = \frac{D}{D_{\text{baseline}}}, \quad (2.1)$$

where D is the axial force of the blade and D_{baseline} is the reference values of the baseline blade. The flow configuration follows the experimental conditions of the MUR47, as reported in Table 2.3, whose relevant quantities are described below. The inlet total temperature and total pressure are defined as:

$$T_1^0 = T_1 \left(1 + \frac{\gamma - 1}{2} M_1^2 \right), \quad p_1^0 = p_1 \left(1 + \frac{\gamma - 1}{2} M_1^2 \right)^{\frac{\gamma}{\gamma - 1}} \quad (2.2)$$

where T_1 and p_1 are the static temperature and static pressure at the inlet, respectively; M_1 is the inlet Mach number and γ is the ratio of specific heats, here assumed as a constant $\gamma = 1.4$. The isentropic Mach number, M_{is} , is defined as:

$$M_{\text{is}} = \sqrt{\frac{2}{\gamma - 1} \left[\left(\frac{p_1^0}{p} \right)^{\frac{\gamma - 1}{\gamma}} - 1 \right]}, \quad (2.3)$$

where p is the local static pressure. The outlet Reynolds number definition used in this work follows:

$$\text{Re}_2 = \frac{\bar{\rho}_2 \bar{U}_2 c}{\bar{\mu}_2}, \quad (2.4)$$

where $\bar{\rho}_2$, \bar{U}_2 and $\bar{\mu}_2$ indicate, respectively, the area-weighted averages of density, velocity magnitude, and dynamic viscosity at the outlet plane. The specific heat at constant pressure and the molecular weight are kept constant, with $c_p = 1006.43 \text{ J (kg K)}^{-1}$ and $M = 28.996 \text{ g mol}^{-1}$ respectively.

Table 2.3: Von Karman Institute MUR47 operating conditions.

Case	p_1^0 (Pa)	p_2 (Pa)	T_1^0 (K)	$M_{2,\text{is}}$
MUR47	159600	83661	420	1.020

A computed flow field, representative of the simulations under the given boundary conditions, is reported through Mach contours in Figure 2.2 along

with numeric colored labels (1)–(6) to better identify relevant flow features. Near the leading edge (1) a stagnation region forms, followed by a rapid acceleration along the suction side. Over this surface, the flow on the suction side becomes locally supersonic (2), inducing an extended region with $M > 1$ that suddenly reverts to the subsonic regime as a shockwave occurs near the trailing edge (3). Although this phenomenon slightly increases the boundary layer thickness, no evident interaction is detected, thus preventing potential separations. Conversely, the acceleration near the pressure side is smoother, never leading to high-supersonic evolution (4). The flow downstream is dominated by wakes and wake/shockwave interactions dynamics. Specifically, the wake released behind the trailing edge (5) tends to be uniformly distributed in the flow field. However, further downstream the interaction with the shockwaves generated by the other blades (6) modifies its evolution, inducing a shear layer thickening and a localized sudden deceleration.

2.3 Model sensitivity and accuracy

The definition of a suitable numerical setup for the optimization analysis is achieved by performing a sensitivity study. Combining the three mesh refinements and four turbulence models with increasing complexity of viscous terms treatment ($k-\omega$ -SST, $k-\omega$ -SST- γ , $k-\varepsilon$ -realizable, and Transition-SST), twelve simulations are conducted under the experiments' flow configuration. Relevant global statistics are compared in Figure 2.3. The spatial resolution exhibits no significant influence, while turbulence modeling has different effects depending on the variable taken into consideration. The prediction of the outlet flow angle, β_2 , shows a negligible dependency on both the mesh refinement and viscous terms treatment (Fig. 2.3a). On the other hand, the $k-\omega$ SST and, even more, the $k-\varepsilon$ depict a tendency to overestimate wake losses, leading to underpredicted overall total pressure ratios ($\overline{p_2^0}/p_1^0$) compared to the other models (Fig. 2.3b). Here, $\overline{p_2^0}$ denotes the area-averaged value at the exit station and the same notation is adopted for this metric when applied to other variables. The isentropic Mach number at the outlet station is estimated with no relevant difference among the different numerical configurations (Fig. 2.3c), all in good alignment with the available experiments. Differently, the average outlet Reynolds number (Fig. 2.3d) exhibits similar dependencies to those of the total pressure ratio, mixing over- and under-prediction tendencies.




Figure 2.2: Mach number distribution through the cascade from simulations of the reference test conditions.

To further investigate the sensitivity at the local level, the distribution of the y^+ along the cascade walls is compared among the different computational setups (Fig. 2.4). Although the incremental refinement from coarse (Figs. 2.4a 2.4b 2.4c 2.4d) to medium (Figs. 2.4e 2.4f 2.4g 2.4h) and from medium to fine (Figs. 2.4i 2.4j 2.4k 2.4l) allows the reduction of the maximum value, it never exceeds 0.25, demonstrating a sufficient resolution already at the coarsest level.

The model accuracy is evaluated by comparing the numerical results with the available measurements from MUR47 (Fig. 2.5). An additional numerical dataset, obtained by Fornasari et al. [11] using RANS techniques and adiabatic




Figure 2.3: Grid resolution and turbulence model sensitivity depicted through global metrics distributions and experimental references. The statistics include: the outlet flow angle a, the total pressure expansion ratio b, the isentropic outlet Mach number c, and the outlet Reynolds number d.

conditions, is included as a computational reference. Along the cascade walls, both the pressure (Fig. 2.5a) and the isentropic Mach number (Fig. 2.5b) align with measurements and existing numerical data, with no clear distinction due to the grid resolution or turbulence model. Especially along the pressure side, all the datasets recover similar flow evolution, without distinct discrepancies. Conversely, on the suction side, the numerical data tend to depart from the tests




Figure 2.4: Grid resolution (top to bottom) and turbulence model sensitivity (left to right) depicted through y^+ distribution and cumulative frequency on the cascade walls for any computational setup.

behavior starting from $s/c = 0.5$, with comparable trends between the present model and the one by Fornasari et al. [11]. The maximum discrepancy occurs near $s/c = 1.0$, before recovering better alignment towards the trailing edge.

The previous discussion indicates that the different numerical configurations perform similarly. Following the existing literature, the $k-\omega$ SST is adopted as turbulence model for the following analyses, while the choice of the grid resolution is guided by the Grid Convergence Index (GCI). The procedure outlined by Celik et al. [24] is applied to both an integral and a local variable: respectively,




Figure 2.5: Pressure ratio a and isentropic Mach number b distributions along the cascade walls, comparing the present model’s sensitivity with reference experiments and other adiabatic simulations [11].

the cascade drag, D , and the wall pressure at the stagnation point. The results reported in Table 2.4 confirm the coarse discretization sufficiently accurate, with a GCI_{medium}^{32} stably below 0.007% for any variable. Therefore, smaller grid sizes are retained to favor faster simulations, thus reducing the computational cost of the optimization loop.

2.4 Parametric model

The geometry parameterization is reported in Figure 2.6. Similar to Agromayor et al. [2], the reference camber line is represented by a cubic Bézier curve with 4 control points: $\{\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2, \mathbf{B}_3\}$. The endpoints, \mathbf{B}_0 and \mathbf{B}_3 , are fixed at the baseline coordinates, reported by the Von Karman Institute: $\mathbf{B}_0 = (0.052094, 0.105932)$ and $\mathbf{B}_3 = (0.538989, -0.773205)$, while the two interior points, \mathbf{B}_1 and \mathbf{B}_2 , are varied in all spatial directions. The reference camber line is parameterized as $\mathbf{x}_{\text{ref}}(u)$, and $\mathbf{n}_{\text{ref}}(u)$ denotes the corresponding unit normal vector at station u . The pres-

Table 2.4: GCI summary for blade drag D and stagnation-point wall static pressure $p(0;0)$. Grid size h and refinement ratios follow the ASME procedure.

	Global ($\phi = D$)	Local ($\phi = p(0;0)$)
ϕ_1	6626.2608 N	159386 Pa
ϕ_2	6626.2570 N	159388 Pa
ϕ_3	6626.2286 N	159394 Pa
p	4.90	1.60
GCI_{fine}^{21}	0.0000548%	0.005%
e_a^{32}	0.000429%	0.00376%
GCI_{medium}^{32}	0.000127%	0.00657%

sure and suction sides are described by two independent fourth-order B-splines, each defined by ten control points. The B-spline endpoints coincide with the Bézier anchors \mathbf{B}_0 and \mathbf{B}_3 and are kept fixed. The eight interior points are evenly distributed along the camberline, adopting equally spaced stations $u_i \in (0, 1)$, $i = 1, \dots, 8$. The points locations are regulated through specific offsets in the camberline-normal direction, resulting in the following expressions:

$$\mathbf{P}_i^{\text{PS}} = \mathbf{x}_{\text{ref}}(u_i) - \alpha_i^{\text{PS}} \mathbf{n}_{\text{ref}}(u_i) \quad (2.5)$$

$$\mathbf{P}_i^{\text{SS}} = \mathbf{x}_{\text{ref}}(u_i) + \alpha_i^{\text{SS}} \mathbf{n}_{\text{ref}}(u_i) \quad (2.6)$$

where α_i^{PS} and α_i^{SS} are the local half-thickness offsets on the pressure and suction sides, respectively. Compared to the original formulation by Agromayor et al. [2], the present parameterization offers a wider exploration range, avoiding the constraint that fixes the same half-thickness at the pressure and suction sides by imposing $\alpha_1^{\text{PS}} = \alpha_1^{\text{SS}}$ and $\alpha_8^{\text{PS}} = \alpha_8^{\text{SS}}$. Here, they are treated as independent decision variables, which increases the flexibility of the leading and trailing edge shapes, while preserving B-spline smoothness. The resulting parameterization uses a total of 20 design parameters: four for the Bézier camber line (the x and y coordinates of \mathbf{B}_1 and \mathbf{B}_2) and sixteen for the surface offsets α_i^{PS} and α_i^{SS} , $i = 1, \dots, 8$, constrained to vary only along the local normal direction to the camber line section. The variation ranges of the decision variables are reported in Table 2.5.

Table 2.5: Decision variables ranges allowed in the optimization, expressed in non-dimensional units.

Decision variable	Lower bound	Baseline	Upper bound
B_{1x}	0.311913	0.388230	0.464174
B_{1y}	-0.266087	-0.223549	-0.177391
B_{2x}	0.362174	0.453651	0.542522
B_{2y}	-0.603131	-0.503684	-0.402087
α_1^{PS}	0.106435	0.133044	0.159652
α_2^{PS}	0.087217	0.109391	0.131565
α_3^{PS}	0.059130	0.073913	0.088696
α_4^{PS}	0.037696	0.047304	0.056913
α_5^{PS}	0.017739	0.022174	0.026609
α_6^{PS}	0.017739	0.021537	0.026609
α_7^{PS}	0.011826	0.014783	0.017739
α_8^{PS}	0.015522	0.019217	0.022913
α_1^{SS}	0.106435	0.133044	0.159652
α_2^{SS}	0.122696	0.153739	0.184783
α_3^{SS}	0.093131	0.116783	0.140435
α_4^{SS}	0.057652	0.072435	0.087217
α_5^{SS}	0.032522	0.041724	0.050261
α_6^{SS}	0.019957	0.025278	0.045087
α_7^{SS}	0.013304	0.016261	0.019217
α_8^{SS}	0.015522	0.019217	0.022913

2.5 Optimization problem formulation

The optimization analysis is formulated as a two-objective minimization problem, targeting the loss coefficient, ω , and the inverse expansion ratio, r , which respectively read as:

$$\omega = \frac{p_1^0 - \overline{p_2^0}}{p_1^0 - p_1} \quad (2.7)$$

$$r = \frac{p_2}{p_1} \quad (2.8)$$

where p_1^0 denotes the inlet total pressure, $\overline{p_2^0}$ the area-weighted average outlet total pressure, and p_1 and p_2 correspond to static pressures. To consider the variation of the two parameters with respect to the baseline values, they are normalized by ω_{baseline} and r_{baseline} , accordingly, resulting in the corresponding normalized quantities, ω^* and r^* . Table 2.6 reports the objective function values of the baseline across all grid refinements and turbulence models adopted in the validation process.

Table 2.6: Baseline’s objective function values across all grid refinement and turbulence models

Refinement	Turbulence model	ω	r
Coarse	k - ω -SST	1.726	0.532
Medium	k - ω -SST	1.727	0.532
Fine	k - ω -SST	1.743	0.532
Coarse	k - ω -SST- γ	1.295	0.532
Medium	k - ω -SST- γ	1.291	0.532
Fine	k - ω -SST- γ	1.310	0.532
Coarse	k - ε -realizable	2.011	0.532
Medium	k - ε -realizable	2.010	0.532
Fine	k - ε -realizable	1.986	0.532
Coarse	Transition-SST	1.296	0.532
Medium	Transition-SST	1.292	0.532
Fine	Transition-SST	1.308	0.532

The problem solution is approached using an evolutionary technique, the Non-dominated Sorting Genetic Algorithm (NSGA-II), as implemented in the open-source Python Multi-Objective optimization library PyMOO. An initial population of 50 individuals is generated with a mixed-nature sampling of the decision variables: half is copied directly from the baseline set, while the remainder is sampled randomly. The analysis is carried out through 50 generations, adopted a Simulated Binary Crossover (SBX) with probability 0.9 to mix parents characters. Additionally, a polynomial mutation is introduced, leveraging a distribution index of $\eta = 20$ [9] and a per-variable mutation probability of $1/n_{\text{var}}$, where $n_{\text{var}} = 20$ is the number of decision variables. Two aerodynamic feasibility constraints are imposed: the outlet flow angle must remain within $\pm 1\%$ of

the baseline value and the mass flow rate within $\pm 5\%$. In addition, a geometric constraint is enforced on blade thickness: candidates whose maximum thickness falls below 50% of the baseline maximum thickness are considered unacceptable designs. Any violation of these constraints is handled by penalizing objective function values. With the chosen numerical setup, each individual evaluation requires approximately 2.5 minutes of wall-clock time when executed on four CPU cores of a AMD Ryzen 7 6800HS, 3.20 GHz. The flow chart of the optimization algorithm is reported in Figure 2.7.




Figure 2.6: Sketch of the cascade's parametric model with splines topology and control points degrees of freedom.





Figure 2.7: Flow chart of the optimization setup.


Results and Discussion

The results of the simulations are reported in Figure 3.1, showing the distribution of the objective functions space. Individuals tend to gather in the lower left region of the plot as generations advance, indicating the conflicting behavior of the two objectives that leads to the formation of a Pareto front (Fig. 3.1a). The shape of this region is strongly affected by the feasibility constraints adopted to align generated designs' performance to baseline operations. To facilitate the understanding of their effect, rejected individuals are reported in the background. Analyzing their distribution, the resulting Pareto front exhibits a clear smooth trend, while the effective front obtained from retained designs features a segmented shape, composed of two distinct, straight branches, with a lower exploration broadness. Consequently, feasibility constraints do not only shape the Pareto front, but also limit further progress of the optimization analysis, potentially binding the benefits of additional iterations. Although the evolution is not monotonic, allowing the exploration of the entire space even at the latest steps, younger offspring are clearly clustered at improved designs. The outcome is further investigated to consider how objectives correlate with the cascade performance. Specifically, the analysis focuses on the blade axial force, or drag, compared to the baseline value (Fig. 3.1b). The results indicate an evident and monotonic correlation between this parameter and the improvement of the expansion ratio, while the blade axial loading appears mainly insensitive to the reduction of overall losses.

To extract potential design correlations, the objective functions are analyzed, examining the dependencies on two relevant cascade geometric parameters: the maximum thickness $d_{\max}^* = d_{\max} / d_{\max, \text{baseline}}$ and the camber angle $\alpha^* = (\beta_{2b} -$



(a)



(b)

★ Baseline ● Opt 1 ● Opt 2 ● Opt 3

Figure 3.1: Objective functions space distribution, depicting the generational evolution a and the correlation with baseline drag variations b.

$\beta_{1b})/(\beta_{2b} - \beta_{1b})_{\text{baseline}}$ (Fig. 3.2). Selecting only the Pareto solutions, a least-square fit is adopted to evaluate the existence of linear correlation. Although the correlation coefficients are significantly low, never exceeding the limit of ± 0.65 obtained between losses and thickness (Fig. 3.2a), the trends consistently confirm conflicting behavior of the two objectives. Specifically, thicker geometries favor improved expansion ratios (Fig. 3.2b) while introducing additional losses. Conversely, reducing the camber angle results as a beneficial design driver for losses minimization (Fig. 3.2c) even though it compromises the cascade ability to expand the flow (Fig. 3.2d). Additionally, analyzing the axial force distribution, it is possible to observe how optimizing the two objective leads to opposite outcomes: lower losses tend to reduce the cascade drag, but improved expansion ratios increase it.

To further investigate this aspect, the objective functions variations are analyzed with respect to the cascade drag (Fig. 3.3). The expansion ratio and the axial force exhibit a strong linear correlation, with a clear monotonic trend, a reduced band dispersion and the absence of outliers (Fig. 3.3a). The negative coefficient confirms a conflicting dependency: improving the expansion ratio always determines an increment of the cascade drag. On the other hand, the broadness in the loss-drag space (Fig. 3.3b) prevents the identification of a clear correlation, despite a slight tendency to weak positive dependency, partially confirming the previous discussion. Individuals distribute uniformly in the plot, with no evident directionality, except for the segmented boundary outlined by the Pareto solutions—an effect induced by the feasibility constraints and not related to potential correlations.

To properly characterize the properties of optimized designs, three individuals are selected from opposite sides of the front. They are labeled Opt 1, Opt 2 and Opt 3 and respectively represent solutions with: minimum losses, minimum expansion ratio and trade-off performance. Compared to the baseline, Opt 1 reduces the normalized loss coefficient by about 30% ($\omega^* = 0.702$) with a negligible deterioration of the inverse expansion ratio ($r^* \approx 1.001$), corresponding to an increase of about 0.7% in outlet total pressure p_2^0 . Minimizing r^* , Opt 2 improves the inverse expansion ratio by approximately 0.2%, while introducing a 26% increase in overall losses ($\omega^* = 1.261$). In contrast, the third design blends the properties of the other two, mixing intermediate improvements of both objectives. Specifically, ω^* is reduced by 12% while inducing a limited optimization of the expansion ratio ($r^* \approx 0.999$). The objective functions of the baseline and




Figure 3.2: Pareto optimal designs' dependencies on relevant geometric properties. Maximum thickness and cascade camberline, normalized with the baseline value, are correlated with normalized losses a–c and total pressure ratio b–d using a linear least-squares fit.

optimized designs are quantitatively summarized in Table 3.1.

Table 3.1: Baseline and three representative Pareto optimal designs: normalized objectives and outlet total pressure ratio with respect to the baseline.

Case	ω^*	r^*	$\bar{p}_2^0 / \bar{p}_{2,\text{baseline}}^0$
Baseline	1.000	1.000	1.0000
Opt 1 (min ω^*)	0.702	1.001	1.0067
Opt 2 (min r^*)	1.261	0.998	0.9963
Opt 3 (trade-off)	0.883	0.999	1.0053

The geometries of the four selected solutions are compared in (Fig. 3.4). The analysis of the entire profile (Fig. 3.4a) indicates that, for any individual, the main shape differences affect the region before mid-chord, keeping the section towards the trailing edge mostly the same as the baseline. Thus, details around the leading edge, on the suction (Fig. 3.4b) and pressure side (Fig. 3.4c) respectively, are reported for enhanced understanding. Along the upper surface, Opt 1 is consistently thinner, while keeping the lower outline similar to the baseline. Conversely, Opt 2 exhibits a considerably thicker geometry on the suction side but features mixed characters on the pressure side. Specifically, the initial section of the leading edge is smoother, until the same thickness as the previous design is reached near $x/c = 0.1$: henceforth, the profile tends to remain stably larger than the others. The third optimized solution is characterized by an intermediate thickness near the suction side leading edge, but near mid-chord it aligns with Opt 2 geometry. Along the pressure side, this design exhibits the lowest thickness among all the geometries, determining a considerably reduced bluntness of the fore section.

The flow field evolution is analyzed comparing the pressure coefficient contours of the four geometries (Fig. 3.5), where the parameter is defined as $C_p = 2(p - p_1) / (\rho_1 U_1^2)$. Compared to the baseline (Fig. 3.5a), Opt 1 (Fig. 3.5b) generally exhibits higher C_p values along the suction side, with a clear reduction of local strong gradients. On the other hand, Opt 2 performs a smoother acceleration, inducing a flow evolution mostly under higher pressure values (Fig. 3.5c). However, three localized pressure drops occur over the suction side, where the one in common with the baseline depicts a weaker jump. Additionally, a relevant gradient can be identified near mid-chord, while a limited pressure drop emerges

near the trailing edge. Concerning Opt 3 (Fig. 3.5d), the flow field through the cascade shares a similar evolution to that of Opt 1, with a generally minimized pressure decay on the whole domain and the absence of strong localized gradients along the suction side. For all the results, the main differences in the pressure field are related to the latter profile region, while along the pressure side the flow evolution is comparable.

To properly characterize shockwaves evolution, the flow field is further investigated observing the distribution of the Mach number (Fig. 3.6). This analysis underscores the correlation between the pressure gradient previously observed in the baseline domain and the formation of a shockwave in the ending section of the profile (Fig. 3.6a). Differently, contours in Opt 1 confirm a reduced acceleration capability of this cascade design (Fig. 3.6b): the flow evolution induces in general lower Mach number, mitigating the supersonic regions and preventing localized evident shocks. The stream diffusion results more gradual, promoting the boundary layer stability over the suction side and limiting the wake thickness downstream. The analysis of the Mach number distribution on the Opt 2 field (Fig. 3.6c) corroborates the observation regarding the critical compressibility effects. The shockwave near the throat involves the maximum Mach value reached in the examined fields, making the phenomenon significantly strong. The resulting flow deceleration partially mitigates the jump amplitude across the second shockwave, reducing its intensity. This lowers the impact of the phenomenon, inducing decreased interactions with the other blades wakes. The trailing edge shock has a direct outcome on the local boundary layer evolution, increasing its thickness and resulting in the release of a thicker wake. The Mach number distribution for Opt 3 (Fig. 3.6d) confirms a flow evolution similar to that of Opt 1. Despite a slightly greater expansion, the flow behavior remains smooth, with no evident shockwaves occurring on the suction side and no significant deterioration of the boundary layer stability. Again, the analysis does not highlight clear variations in flow evolution along the pressure side among the geometries.

The specific evolution with local accuracy is analyzed reporting wall-flow distributions (Fig. 3.7). The isentropic Mach number along the blade surfaces confirms the similarity of the flow behavior on the pressure side, except for a localized gradient near the leading edge of Opt 2 (Fig. 3.7a). On the other hand, a higher variability is present on the suction side, especially downstream of coordinate $s/c = 0.5$. Here, the three optimized geometries exhibit a comparable performance, with M_{is} exceeding 1, thus differing from the baseline behavior

that maintains a smoother trend. At this location they reach the peak conditions, anticipating the flow phenomenon observed in the starting design. The parameter then decays, with different patterns for the three geometries: Opt 1 and Opt 3 maintain similar distributions, smoothly aligning to the trailing edge flow, while Opt 2 exhibits the highest peak before a local sudden drop that anticipates a progressive increment towards values comparable with the other solutions. After coordinate $s/c = 1.0$, corresponding to the baseline maximum M_{is} , the four geometries tend to a gradual decrease, even though Opt 2 shows the occurrence of a narrow localized peak, followed by a sudden drop that re-aligns the statistics values with the others. Analogous considerations hold for the wall pressure distribution (Fig. 3.7b). Smoother pressure recoveries indicate more gradual diffusion processes along the pressure side. However, improved expansion ratios result because the pressure minima of optimized designs become lower than those of the baseline, even though the wake evolution also contributes to determining r —a fact that explains the similarity between Opt 1 and Opt 3 local statistics, despite their different positioning in the objective function space. To characterize viscous fluxes evolution, the local analysis introduces the distribution of the skin friction coefficient, $C_f = \tau_w / (0.5 \rho_1 U_1^2)$, where τ_w is the local wall shear stress and ρ_1 and U_1 are the density and magnitude of the absolute velocity at the cascade inlet (Fig. 3.7c). The distribution along the pressure side shows significant similarity, aligning with previous discussions. However, minor discrepancies can be detected and they indicate a stable tendency of Opt 2 to reach higher values over the entire wall surface, after depicting the most intense peak. Conversely, Opt 1 mostly exhibits the lowest wall shear, which is consistent with its ability to minimize losses, reducing the impact of viscous effects. Along the suction side, Opt 3 shares a similar behavior with the baseline, except for a higher peak near $s/c = 0.5$, after which the two solutions almost re-align. At the same location, Opt 2 exhibits the strongest wall shear, followed by a sudden drop related to the occurrence of the shockwave. A similar effect can be observed near the trailing edge, confirming this design as affected by significant shock/boundary layer interactions. Considering Opt 1, the distribution of C_f exhibits a peak that slightly exceeds the baseline value. After this location the skin friction decreases, remaining almost stably below the other solutions. Performing a smoother flow deceleration, this design enables a higher boundary layer stability, minimizing the impact of viscous fluxes on the walls.

The wake structure is analyzed at an outlet cross-section located $1.2 c_{ax}$ down-

stream of the cascade trailing edge (Fig. 3.8). To measure the intensity of viscous fluxes, the analysis adopts the normalized turbulent kinetic energy, $k = 2K/(U_1^2)$, where K denotes the specific turbulent kinetic energy (Fig. 3.8a). The distributions clearly indicate that Opt 2 introduces higher dissipative effects, resulting in a stronger turbulence production and in a larger wake extension. On the other hand, compared to baseline, the other two designs reduce wake dissipation, with Opt 1 also ensuring a significant reduction of the shear layer band. Additionally, turbulence production has a direct impact on total pressure drops (Fig. 3.8b). Consequently, Opt 2 results as the worst design in terms of losses, with a pressure drop higher than -20% and the largest wake spatial extension among the chosen solutions. In contrast, the baseline guarantees a significant improvement while also reducing the wake bandwidth. Considering Opt 1 and Opt 3, the minimum total pressure is similar. However, the optimality of the former is explained by the considerably lower spatial extension of the shear layer. To corroborate this analysis, the results are further investigated with a quantitative approach. Thus, the total pressure is integrated along the one-pitch cross-sectional plane. Each individual evaluation, I_j , is computed as:

$$I_j = \int_0^1 \left(\frac{p^0}{p_1^0} \right)_j d\left(\frac{y}{t}\right), \quad (3.1)$$

where y/t denotes the pitch-normalized vertical coordinate. The corresponding results are reported in Table 3.2, normalized by the baseline integral value, as I_j/I_{baseline} . The computed values further confirm the previous discussions regarding cascade losses, indicating Opt 1 and Opt 3 as superior to the baseline, while underscoring the augmented dissipation introduced by Opt 2.

Table 3.2: Integral of normalized total pressure along the wake cut at $x = 1.2 c_{ax}$, normalized by the baseline value.

Case	I_j/I_{baseline}
Baseline	1.000
Opt 1 (min ω^*)	1.004
Opt 2 (min r^*)	0.994
Opt 3 (trade-off)	1.003

Overall, this analysis supports further concluding remarks regarding cascades

design principles. Specifically, when losses are a major concern, geometries have to reduce the evolution of shock/boundary-layer interactions on the suction side to limit wake thickness and turbulence production. Conversely, to improve the expansion ratio, higher accelerations are required, but this behavior may introduce shockwaves, resulting in boundary layer destabilization and higher dissipation.




Figure 3.3: Correlation between Pareto individuals' normalized cascade drag and related objectives, including normalized pressure ratio a and loss coefficient b. Least-squares fit reported in red.




Figure 3.4: Shape comparison between baseline and optimized individuals in the Pareto front as reported in Figure 3.1a, with magnified views around the leading edge suction b and pressure c sides.




Figure 3.5: Pressure coefficient distributions for the baseline a and the chosen optimized geometries, respectively: Opt 1 b, Opt 2 c and Opt 3 d.




Figure 3.6: Mach number distributions for the baseline a and the chosen optimized geometries, respectively: Opt 1 b, Opt 2 c and Opt 3 d. The common colour scale is shown in e.




Figure 3.7: Comparison of wall statistics between baseline and optimized geometries. The plots include the isentropic Mach number a, the pressure ratio b and the skin-friction coefficient c.




Figure 3.8: Comparison of wake statistics between baseline and optimized geometries. The quantities, extracted at $x = 1.2 c_{ax}$, include the normalized turbulent kinetic energy a and the total pressure ratio b.

Conclusions

This study presents a novel evolutionary approach to transonic gas-turbine vanes design, optimizing the parameterized profile of the LS89 cascade. The model leverages translational periodicity, enabling the generation of a two-dimensional single-passage domain. Thus, steady-state Reynolds-Averaged Navier-Stokes (RANS) equations are adopted to initially validate the numerical setup and then to compute the solutions for the geometrical variants. To this end, a two-objective genetic optimization is used to advance generations, satisfying the simultaneous minimization of the expansion ratio and the loss coefficient.

The procedure enables the improvement of baseline performance, reducing profile losses by 30% and the expansion ratio by 0.2%. However, the two objectives exhibit conflicting behavior, requiring dedicated geometrical properties and design principles, except for a common optimization driver that suggests to place the maximum channel restriction near mid-chord. Improving the expansion ratio requires adopting thicker profiles, which induces higher accelerations but introduces additional shockwaves that compromise boundary layer stability and increase losses. Conversely, to reduce wake dissipation thinner blades with sharper leading edges are favorable, even though this design mitigates shock/boundary layer interactions and wake extension at the cost of lower flow expansion.

Future studies are expected to further advance the present outcomes. Specifically, increasing the local accuracy of the parameterization and enlarging the exploration ranges will allow for improved control over localized flow phenomena, thus leading to better optimization results and additional design correlations.

Bibliography

- [1] Santosh Abraham et al. "Effect of Turbine Airfoil Shape on Aerodynamic Losses for Turbine Airfoils Operating Under Transonic Conditions". In: *Proceedings of ASME Turbo Expo 2011: Power for Land, Sea and Air*. GT2011-45188. Vancouver, British Columbia, Canada, June 2011, pp. 511–521. doi: 10.1115/GT2011-45188.
- [2] Roberto Agromayor et al. "A Unified Geometry Parametrization Method for Turbomachinery Blades". In: *Computer-Aided Design* 133 (2021), p. 102987. doi: 10.1016/j.cad.2020.102987.
- [3] T. Arts and M. Lambert de Rouvroit. "Aero-Thermal Performance of a Two-Dimensional Highly Loaded Transonic Turbine Nozzle Guide Vane: A Test Case for Inviscid and Viscous Flow Computations". In: *J. Turbomach.* 114.1 (1992), pp. 147–154. doi: 10.1115/1.2927978.
- [4] N. C. Baines et al. "The Aerodynamic Development of a Highly Loaded Nozzle Guide Vane". In: *J. Turbomach.* 108.2 (Oct. 1986), pp. 261–268. issn: 0889-504X. doi: 10.1115/1.3262046. url: <https://doi.org/10.1115/1.3262046>.
- [5] David Baumgärtner, John J. Otter, and Andrew P. S. Wheeler. "The Effect of Isentropic Exponent on Transonic Turbine Performance". In: *J. Turbomach.* 142.8 (Oct. 2020), p. 081007. doi: 10.1115/1.4046528.
- [6] James L. Bettner. *Experimental Investigation in an Annular Cascade Sector of Highly Loaded Turbine Stator Blading. Volume 3: Performance of Tandem Blade*. Contractor Report NASA CR-1341. National Aeronautics and Space Administration (NASA), 1969.

- [7] Arnaud Châtel, Tom Verstraete, and Grégory Coussement. “Multipoint Optimization of an Axial Turbine Cascade Using a Hybrid Algorithm”. In: *J. Turbomach.* 142.2 (Feb. 2020), p. 021017. DOI: 10.1115/1.4046231.
- [8] Francesco De Vanna and Ernesto Benini. “Wall-Modeled Large-Eddy Simulation of a Transonic Gas Turbine Vane—Part II: Mach Number Effect and Losses Prediction”. In: *J. Turbomach.* 147.11 (2025), p. 111016. ISSN: 0889-504X. DOI: 10.1115/1.4069384.
- [9] Kalyanmoy Deb, Karthik Sindhya, and Tatsuya Okabe. “Self-adaptive simulated binary crossover for real-parameter optimization”. In: *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation.* 2007, pp. 1187–1194. DOI: 10.1145/1276958.1277190.
- [10] Ashlie B. Flegel-McVetta, Paul W. Giel, and Gerard E. Welch. *Aerodynamic Measurements of a Variable-Speed Power-Turbine Blade Section in a Transonic Turbine Cascade at Low Inlet Turbulence.* Technical Memorandum NASA/TM–2013-218069. ASME Turbo Expo paper GT2013-94695. Cleveland, OH: NASA Glenn Research Center, Aug. 2013.
- [11] Davide Fornasari, Davide Regazzo, and ... “RANS-Based Aerothermal Database of LS89 Transonic Turbine Cascade Under Adiabatic and Cooled Wall Conditions”. In: *Energies* 18.19, 5321 (2024). DOI: 10.3390/en18195321.
- [12] H. E. Gallus, K.-D. Broichhausen, and J. M. Henne. “Experimental Unsteady Shock–Boundary Layer Interaction at Single Blades and in Linear Cascades”. In: *Proceedings of the ASME 1986 International Gas Turbine Conference and Exhibit.* 86-GT-218. Düsseldorf, West Germany, June 1986. DOI: 10.1115/86-GT-218.
- [13] H. P. Hodson and R. G. Dominy. “The Off-Design Performance of a Low-Pressure Turbine Cascade”. In: *J. Turbomach.* 109.2 (1987), pp. 201–209. ISSN: 0889-504X. DOI: 10.1115/1.3262086.
- [14] H. Hoheisel et al. “Influence of Free-Stream Turbulence and Blade Pressure Gradient on Boundary Layer and Loss Behavior of Turbine Cascades”. In: *J. Turbomach.* 109.2 (1987), pp. 210–219. ISSN: 0889-504X. DOI: 10.1115/1.3262087.

- [15] Romain Hottot et al. "Comparing Gradient-Free and Gradient-Based Multi-Objective Optimization Methodologies on the VKI-LS89 Turbine Vane Test Case". In: *J. Turbomach.* 145.3 (Feb. 2023), p. 031001. doi: 10.1115/1.4055577.
- [16] J. F. Kline and S. M. Nosek. *Two-Dimensional Cascade Investigation of a Turbine Tandem Blade Design: Flow and Wake Characteristics on Turbine Tandem Blades in a Two-Dimensional Cascade Tunnel*. Technical Memorandum NASA TM X-1836. Cleveland, OH: NASA Lewis Research Center, July 1969.
- [17] Haitao Li et al. "Inverse Problem for Isentropic Mach-Number on Blade Wall in Aerodynamic Shape Design of Turbomachinery Cascades by Using Adjoint Method". In: *Proc. ASME Turbo Expo 2011: Turbomachinery, Parts A, B, and C*. 2011, pp. 1229–1240. doi: 10.1115/GT2011-45808.
- [18] Jichao Li, Sicheng He, and Joaquim R. R. A. Martins. "Data-Driven Constraint Approach to Ensure Low-Speed Performance in Transonic Aerodynamic Shape Optimization". In: *Aerosp. Sci. Technol.* 92 (2019), pp. 536–550. doi: 10.1016/j.ast.2019.06.008.
- [19] Zhi Li and Yan Liu. "Linear Combination Design Method for Turbine Cascades". In: *Aerosp. Sci. Technol.* 168 (2026), p. 111072. issn: 1270-9638. doi: 10.1016/j.ast.2025.111072.
- [20] F. Martelli and A. Boretti. "Development of an Experimental Correlation for Transonic Turbine Flow". In: *J. Turbomach.* 109.2 (1987), pp. 246–250. issn: 0889-504X. doi: 10.1115/1.3262092.
- [21] Ashlie B. McVetta, Paul W. Giel, and Gerard E. Welch. *Aerodynamic Investigation of Incidence Angle Effects in a Large Scale Transonic Turbine Cascade*. Technical Memorandum NASA/TM-2013-218070. Cleveland, OH: NASA Glenn Research Center, Aug. 2013.
- [22] D. J. Mee et al. "An Examination of the Contributions to Loss on a Transonic Turbine Blade in Cascade". In: *J. Turbomach.* 114.1 (1992), pp. 155–162. issn: 0889-504X. doi: 10.1115/1.2927979.
- [23] V. Michelassi, W. Rodi, and P.-A. Gieß. "Experimental and Numerical Investigation of Boundary-Layer and Wake Development in a Transonic Turbine Cascade". In: *Aerosp. Sci. Technol.* 2.3 (1998), pp. 191–204. issn: 1270-9638. doi: 10.1016/S1270-9638(98)80053-9.

- [24] “Procedure for Estimation and Reporting of Uncertainty Due to Discretization in CFD Applications”. In: *J. Fluids Eng.* 130.7 (July 2008), p. 078001. DOI: 10.1115/1.2960953.
- [25] A. D. Rossiter, G. Pullan, and A. P. Melzer. “The Influence of Boundary Layer State and Trailing Edge Wedge Angle on the Aerodynamic Performance of Transonic Turbine Blades”. In: *J. Turbomach.* 145.4 (Apr. 2023), p. 041008. DOI: 10.1115/1.4055868.
- [26] Antonio Rubino, Piero Colonna, and Matteo Pini. “Adjoint-Based Unsteady Optimization of Turbomachinery Operating with Nonideal Compressible Flows”. In: *J. Propuls. Power* 37.6 (2021), pp. 910–918. DOI: 10.2514/1.837920.
- [27] Antonio Rubino et al. “Fully-Turbulent Adjoint Method for the Unsteady Shape Optimization of Multi-Row Turbomachinery”. In: *Aerosp. Sci. Technol.* 106 (2020), p. 106132. DOI: 10.1016/j.ast.2020.106132.
- [28] Ismael Sanchez Torreguitart, Tom Verstraete, and Lasse Mueller. “Optimization of the LS89 Axial Turbine Profile Using a CAD and Adjoint Based Approach”. In: *Int. J. Turbomach. Propuls. Power* 3.3 (Aug. 2018), p. 20. DOI: 10.3390/ijtp3030020.
- [29] Andrea Serani, Matteo Diez, and Domenico Quagliarella. “Aerodynamic Shape Optimization in Transonic Conditions Through Parametric Model Embedding”. In: *Aerosp. Sci. Technol.* 155 (2024), p. 109611. DOI: 10.1016/j.ast.2024.109611.
- [30] Toyotaka Sonoda et al. “A Study of Advanced High-Loaded Transonic Turbine Airfoils”. In: *J. Turbomach.* 128.4 (2004), pp. 650–657. ISSN: 0889-504X. DOI: 10.1115/1.2221325.
- [31] M. Stastny and P. Safarik. “Experimental Analysis Data on the Transonic Flow Past a Plane Turbine Cascade”. In: *Proceedings of the ASME Gas Turbine and Aeroengine Congress and Exposition*. 90-GT-313. Brussels, Belgium, June 1990. DOI: 10.1115/90-GT-313.
- [32] Fernando Tejero et al. “Neural Network-Based Multi-Point, Multi-Objective Optimisation for Transonic Applications”. In: *Aerosp. Sci. Technol.* 136 (2023), p. 108208. ISSN: 1270-9638. DOI: 10.1016/j.ast.2023.108208.

- [33] Patrick Tene Hedje et al. "Large Eddy Simulations of a High-Speed Low-Pressure Turbine Cascade at Subsonic and Transonic Mach Numbers". In: *J. Turbomach.* 147.10 (Oct. 2025), p. 101010. doi: 10.1115/1.4068124.
- [34] Chaolei Zhang and Zhenping Feng. "Aerodynamic Shape Design Optimization for Turbomachinery Cascade Based on Discrete Adjoint Method". In: *Proc. ASME Turbo Expo 2011: Turbomachinery, Parts A, B, and C.* 2011, pp. 1219–1228. doi: 10.1115/GT2011-45805.
- [35] Jianshe Zhang et al. "A Discrete Adjoint Framework Coupled with Adaptive PCE for Robust Aerodynamic Optimization of Turbomachinery under Flow Uncertainty". In: *Aerosp. Sci. Technol.* 142 (2023), p. 108592. issn: 1270-9638. doi: 10.1016/j.ast.2023.108592.

Chapter 5

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Prof. Francesco De Vanna, for the continuous guidance, valuable feedback, and encouragement throughout this work.

I am also grateful to my co-supervisor, Dr. Filippo Avanzi, for the insightful discussions and practical support that helped shape and improve this thesis.

I would like to thank my classmates and colleagues for the stimulating environment, the shared effort during our courses, and the many constructive exchanges along the way.

Finally, I am deeply thankful to my family and friends for their constant support, patience, and understanding during the past months.

Their encouragement has been essential in completing this journey.