NUMERICAL ANALSIS AND MULTI-OBJECTIVE OPTIMISATION OF THE NACA 2415 AIRFOIL USING A TAGUCHI-FUZZY FRAMEWORK

A Thesis Submitted
In Partial Fulfillment of the Requirement for the
Degree of

MASTER OF TECHNOLOGY

In
Computer-Aided Design

By

PULKIT SONI
24/CAD/06

Under the Supervision of
Dr. MOHAMMAD ZUNAID
Associate Professor, Department of Mechanical Engineering
Delhi Technological University

DEPARTMENT OF MECHANICAL ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi-110042. India

MAY, 2026

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulapur, Main Bawana Road, Delhi-42

ACKNOWLEDGEMENT

I would like to express my sincere gratitude to my supervisor, Dr. Mohammad Zunaid, Associate Professor, Department of Mechanical Engineering, Delhi Technological University, for his invaluable guidance, constant encouragement, and insightful feedback throughout this research project. His expertise in computational fluid dynamics and aerodynamic design optimisation has been instrumental in shaping every aspect of this dissertation.

I am grateful to the Department of Mechanical Engineering, Delhi Technological University, for providing access to the necessary computational facilities and resources, particularly the ANSYS Fluent software suite and the high-performance computing infrastructure that made this study possible.

I would also like to acknowledge the contributions of the numerous researchers whose foundational work on NACA airfoils, Taguchi methods, fuzzy inference systems, and RANS turbulence modelling provided the theoretical underpinning for this investigation.

Finally, I am deeply grateful to my family and friends for their continuous support, patience, and understanding throughout the course of my M.Tech programme.

PULKIT SONI

Roll No. 24/CAD/06

M.Tech – Computer Aided Analysis and Design

Department of Mechanical Engineering, DTU

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulapur, Main Bawana Road, Delhi-42

CANDIDATE'S DECLARATION

I, PULKIT SONI, Roll No. 24/CAD/06, hereby certify that the work which is being presented in the thesis titled **"Numerical Analysis and Multi-Objective Optimisation of the NACA 2415 Airfoil Using a Taguchi–Fuzzy Framework"** in partial fulfillment of the requirement for the award of the Degree of Master of Technology in Computer-Aided Analysis & Design submitted in the Department of **Mechanical Engineering**, Delhi Technological University is an authentic record of my own work carried out during the period from_____ to_____ under the supervision of Dr. Mohammad Zunaid, Associate Professor, Department of Mechanical Engineering, Delhi Technological University.

The matter presented in the thesis has not been submitted by me for the award of any other degree of this or any other institute.

Place: Delhi
Date: 29th May 2026

PULKIT SONI
Roll No. 24/CAD/06

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulapur, Main Bawana Road, Delhi-42

CERTIFICATE BY THE SUPERVISOR(s)

Certified that **PULKIT SONI (Roll No. 24/CAD/06)** has carried out his research work presented in this thesis titled **"Numerical Analysis and Multi-Objective Optimisation of the NACA 2415 Airfoil Using a Taguchi–Fuzzy Framework"** for the award of **Master of Technology in Computer-Aided Analysis & Design** from Department of Mechanical Engineering, Delhi Technological University, Delhi under my supervision. The thesis embodies results of original work, and studies are carried out by the student himself and the contents of the thesis do not form the basis for the award of any other degree to the candidate or to anybody else from this or any other University/Institution.

Dr. MOHAMMAD ZUNAID
Associate Professor
Department of Mechanical Engineering
Delhi Technological University

Place: Delhi
Date: 29th May 2026

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulapur, Main Bawana Road, Delhi-42

NUMERICAL ANALYSIS AND MULTI-OBJECTIVE OPTIMISATION OF THE NACA 2415 AIRFOIL USING A TAGUCHI–FUZZY FRAMEWORK

PULKIT SONI

ABSTRACT

Small fixed-wing unmanned aerial vehicles operating at chord Reynolds numbers of $Re = 10^5–5\times10^6$ face a fundamental aerodynamic challenge: the gap between two-dimensional section performance predicted by computational fluid dynamics and the drastically degraded efficiency of real, aspect-ratio-constrained three-dimensional wings. This study addresses that gap through a three-phase hierarchical investigation of the NACA 2415 aerofoil, integrating Taguchi design of experiments, Mamdani fuzzy multi-objective optimisation, steady-state Reynolds-Averaged Navier–Stokes simulation, and surrogate-assisted optimisation into a single systematic framework validated against NACA Technical Report 824 experimental data.

Phase 1 deploys a Taguchi L25(5^5) orthogonal array — reducing a 3,125-run full factorial to 25 balanced simulations — to simultaneously screen five RANS turbulence closures (Spalart–Allmaras, k-ε Realizable, k-ω SST, SST γ–Reθ, and Reynolds Stress Model), five Reynolds numbers ($1–12\times10^6$), five angles of attack (−4° to 16°), five turbulence intensities (0.05%–5.00%), and four surrogate optimisation strategies (RSM-Kriging, NSGA-II, Sparse RSM, and Neural Network Screening). Three conflicting aerodynamic responses — lift coefficient, drag coefficient, and lift-to-drag ratio — are unified into a scalar Multi-Performance Characteristic Index via a 27-rule Mamdani fuzzy inference system with corrected strict-inequality boundary-membership evaluation, a previously unreported defect whose correction changes the turbulence model ANOVA contribution from a spurious 12.30% to the physically correct 1.83%. One-way ANOVA identifies angle of attack as the dominant factor (ρ = 80.99–85.70%), with Reynolds number second ($\rho \approx$ 8–10%). The k-ω SST model achieves the highest multi-objective η(MPCI) level mean (−9.202 dB) due to its Bradshaw adverse-pressure-gradient limiter and structural turbulence-intensity insensitivity via cross-diffusion. Sparse RSM achieves the highest Weighted Composite Score of 9.13/10, uniquely detecting the NACA 2415 drag-bucket interior minimum at $\alpha \approx -0.75°$, independently confirmed by Neural Network Screening at $\alpha \approx -0.77°$.

Phase 2 deploys a Taguchi L9(3^3) array exclusively with k-ω SST across a refined design space (Re = 6–12×10^6, α = 4°–8°, TI = 0.05%–0.50%). The confirmed optimal configuration — Re = 12×10^6, α = 8°, TI = 0.10° — yields CL = 1.038, CD = 0.015711, and |CL/CD| = 66.08, with a Taguchi additive model prediction error of only 0.26%, validating negligible factor interactions. Turbulence intensity contributes ρ ≈ 0.00% (F = 0.04) within the tested range, providing a practically significant result that eliminates TI as a source of CFD modelling uncertainty for this application.

Phase 3 extends the Phase 2 optimum to a three-dimensional finite-wing RANS simulation at AR = 0.25 (b = 0.5 m, c = 2.0 m, A_ref = 1.0 m^2). The aerodynamic outputs — CL = 0.12921, CD = 0.014899, |CL/CD| = 8.67, Lift = 607.908 N, Drag = 70.096 N — reveal an 86.9% efficiency collapse from the two-dimensional optimum, driven by a tip-vortex-induced downwash of ε ≈ 9.33° that reduces the effective angle of attack from +8° to approximately −1.43°. Three independent CFD visualisations — velocity pathlines, static pressure vectors, and velocity magnitude vectors — provide mutually corroborating topological, thermodynamic, and kinematic evidence confirming that the entire span lies within the tip-vortex induction zone and no two-dimensional flow region exists. The study conclusively establishes that the binding aerodynamic limitation of the platform is planform geometry rather than section performance, motivating a redesign to AR = 6–8 to recover 75%–86% of the two-dimensional efficiency ceiling.

Keywords: NACA 2415; Taguchi orthogonal array; fuzzy MPCI; RANS turbulence model screening; k-ω SST; Sparse RSM; finite-wing induced drag; UAV aerodynamic optimisation

TABLE OF CONTENTS

TITLE OF CHAPTER **Page No**.

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulapur, Main Bawana Road, Delhi-42

LIST OF TABLES

APPENDIX-II

APPENDIX-III

APPENDIX-IV

APPENDIX-V

APPENDIX-VI

APPENDIX-VII

APPENDIX-X

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulapur, Main Bawana Road, Delhi-42

LIST OF FIGURES

NOMENCLATURE

English Symbols

Symbol	Definition
A, B, C, D, E	Taguchi factor designations (A=Re, B=Model, C=AoA, D=TI, E=Strategy)
A_ref	Reference area (m^2)
AR	Aspect ratio (= b/c)
b	Wing span (m)
c	Chord length (m)
CD	Drag coefficient
CL	Lift coefficient
D	Aerodynamic drag force (N)
e	Oswald efficiency factor
FD	Aerodynamic drag force (N)
FL	Aerodynamic lift force (N)
k	Turbulent kinetic energy (m^2/s^2)
L	Aerodynamic lift force (N)
Ma	Mach number
MS	Mean Square (ANOVA)
p	Static pressure (Pa)
q_∞	Freestream dynamic pressure (Pa)
Re	Reynolds number
S/N	Signal-to-Noise ratio (dB)
SS	Sum of Squares (ANOVA)
TI	Turbulent Intensity (%)
V_∞	Freestream velocity (m/s)
V_ref	Reference velocity for coefficient scaling (m/s)
w_induced	Tip-vortex-induced downwash velocity (m/s)
WCS	Weighted Composite Score (0–10)
x_1, x_2, x_3	Normalised fuzzy input variables (lift, $drag^{-1}$, efficiency)
y^+	Non-dimensional wall distance
y_1	First inflation-layer cell height (m)

Greek Symbols

Symbol	Definition
α	Angle of attack (°)
α_eff	Effective angle of attack after induced downwash correction (°)
$\alpha L0$	Zero-lift angle of attack (°)
ε	Turbulent dissipation rate (m^2/s^3); also induced downwash angle (°)
η	Signal-to-Noise ratio (dB)
$\hat{\eta}\_opt$	Predicted optimal S/N ratio — Taguchi additive model (dB)
μ	Dynamic viscosity (Pa·s)
μt	Turbulent eddy viscosity (Pa·s)
ν	Kinematic viscosity (m^2/s)
ρ	Fluid density (kg/m^3); also ANOVA percentage contribution (%)
τ^{R}_{ij}	Reynolds stress tensor (Pa)
ω	Specific turbulent dissipation rate (s^{-1})

Acronyms and Abbreviations

Abbreviation	Definition
ANN	Artificial Neural Network
ANOVA	Analysis of Variance
AR	Aspect Ratio
CFD	Computational Fluid Dynamics
CoG	Centre of Gravity (defuzzification method)
DoE	Design of Experiments
DoF	Degrees of Freedom
DTU	Delhi Technological University
FIS	Fuzzy Inference System
GCI	Grid Convergence Index
k-ε R.	k-epsilon Realizable turbulence model
k-ω SST	k-omega Shear Stress Transport turbulence model
MPCI	Multi-Performance Characteristic Index
NACA	National Advisory Committee for Aeronautics
NN-Screen	Neural Network Screening
NSGA-II	Non-dominated Sorting Genetic Algorithm II
RANS	Reynolds-Averaged Navier–Stokes
RSM	Reynolds Stress Model / Response Surface Methodology
SA	Spalart–Allmaras turbulence model
SIMPLE	Semi-Implicit Method for Pressure-Linked Equations

Abbreviation	Definition
Sparse RSM	Sparse Initialisation Response Surface Methodology
SST γ–Reθ	k-ω SST with γ–Reθ transition model
TR-824	NACA Technical Report 824 (Abbott et al., 1945)
UAV	Unmanned Aerial Vehicle

CHAPTER 1 — INTRODUCTION

1.1 Background and Motivation

Computational aerodynamics for unmanned aerial vehicles (UAVs) has evolved from inviscid panel methods into high-fidelity Reynolds-Averaged Navier–Stokes (RANS) simulation, yet a persistent gap remains between what CFD predicts for an ideal two-dimensional aerofoil section and what a real, geometry-constrained UAV wing actually delivers. This thesis addresses that gap directly.

The global UAV market — valued at approximately USD 14.1 billion in 2023 and growing at over 15% annually — is dominated numerically by small fixed-wing platforms with wingspans of 0.5–3.0 m operating at chord Reynolds numbers of $Re \approx 10^5$–5×10^6. This intermediate regime is aerodynamically treacherous: the boundary layer is susceptible to laminar separation bubbles, drag-bucket collapse, and large-scale transition instabilities that fully turbulent RANS models cannot represent without explicit transition modelling. Simultaneously, compact UAV planforms impose severe aspect-ratio constraints that generate tip-vortex-dominated flows where three-dimensional effects dwarf any section-level performance gain. These two challenges — turbulence model fidelity at transitional Reynolds numbers, and the 2D-to-3D performance collapse at low aspect ratios — define the engineering problem this thesis investigates.

The NACA 2415 aerofoil is selected as the study section: its 2% camber, 15% thickness, and extensive experimental documentation in NACA TR-824 make it an ideal validation vehicle, while its widespread use in UAV and general aviation applications ensures direct engineering relevance. The chord length is fixed at $c = 2$ m throughout Phases 1 and 2, placing the operating range $Re = 1$–12×10^6 within the TR-824 validation envelope.

1.2 Literature Review

1.2.1 UAV Aerodynamics: Context and Challenges

The rapid commercialisation of unmanned aerial vehicles across agriculture, surveillance, search-and-rescue, and remote sensing has created urgent demand for aerodynamically efficient small fixed-wing platforms [1, 2]. Small UAVs with wingspans of 0.5–3.0 m present a fundamentally different aerodynamic challenge from conventional aircraft: they operate at chord Reynolds numbers of $Re \approx 10^5$–5×10^6, a transitional regime in which neither classical fully turbulent boundary-layer theory nor inviscid potential-flow methods are adequate [3, 4]. Mueller and DeLaurier [3] established that aerodynamic behaviour in this regime is qualitatively different from both low-speed insect flight and high-Reynolds-number aircraft, and that conventional

design tools fail without modification. Shyy et al. [4] further demonstrated that flexible structures, gust sensitivity, and transitional boundary-layer physics interact in ways that substantially reduce real-world efficiency below what inviscid or fully turbulent computations predict. Mueller [5] provided the foundational classification of micro-air-vehicle aerodynamic regimes, establishing that laminar separation bubbles, drag-bucket collapse, and hysteresis in the lift polar are defining features of the $Re = 10^5$–10^6 regime — features that persist in muted form up to $Re \approx 5\times10^6$ for thick cambered sections such as the NACA 2415. Horton [6] provided the seminal theoretical treatment of laminar separation bubbles, identifying the separation–transition–reattachment sequence that governs drag-bucket behaviour, and Spiering et al. [7] demonstrated through modern RANS computations that even at $Re = 3\times10^6$, transition-sensitive modelling is necessary for accurate drag-bucket prediction. The importance of carefully selected aerofoil geometry for small UAV applications was highlighted by Selig et al. [8], whose systematic wind-tunnel evaluation of low-speed sections established the performance benchmarks against which computational predictions must be validated.

1.2.2 NACA Aerofoil Experimental Heritage and Validation Data

The four-digit NACA aerofoil series, analytically parameterised by thickness and camber-line distributions, was systematically characterised experimentally across Reynolds numbers of $Re = 3\times10^6$–9×10^6 in NACA TR-460 [9] and comprehensively documented with both clean and standard-roughness polars in NACA TR-824 [10, 11]. These reports remain the primary external validation references for computational predictions of NACA-family sections. Abbott and von Doenhoff [11] established the corrected roughness-condition drag polars — physically equivalent to fully turbulent RANS predictions from the leading edge — that enable consistent model-to-experiment comparison for computations that impose turbulence from the stagnation point. Anderson [12] provided the theoretical aerodynamic framework — thin-aerofoil theory, finite-wing lifting-line correction, and boundary-layer fundamentals — within which all computational predictions in the present study are physically interpreted.

1.2.3 RANS Turbulence Models for Aerofoil Aerodynamics

The incompressible RANS equations close through the Reynolds stress tensor, whose modelling spans a hierarchy of approaches from the Boussinesq eddy-viscosity hypothesis [13] to full second-moment closures. The Spalart–Allmaras one-equation model [14] introduced a transport equation for modified turbulent kinematic viscosity, offering computational economy and reasonable attached-flow accuracy at the cost of sensitivity to freestream turbulence intensity at high TI values. Jones and Launder [15] proposed the standard k-ε two-equation framework; Wilcox [16] developed the k-ω alternative with superior near-wall behaviour; and Shih et al. [17] addressed the k-ε model's overproduction of eddy viscosity in adverse pressure gradients through the Realizable variant. The k-ω SST model, developed by Menter [18, 19] through a blended formulation incorporating the Bradshaw shear-stress transport limiter [20], became the industry standard for external aerodynamic simulation by correctly bounding eddy viscosity in adverse pressure gradient regions — the critical

mechanism for accurate prediction of boundary-layer separation onset. Menter et al. [21] subsequently extended SST with two transport equations for intermittency and transition-onset momentum-thickness Reynolds number, enabling correlation-based natural transition prediction; Langtry and Menter [22] demonstrated its deployment for unstructured parallel codes, and van Ingen [23] provided the original e^N method whose empirical correlations underpin the transition onset criterion. For flows with strong stress anisotropy, Launder et al. [24] proposed the Reynolds Stress Model, solving transport equations for all six stress components directly from the pressure-strain correlation of Rotta [25] and the wall-reflection model of Gibson and Launder [26]. Near-wall resolution for all these models requires $y^+ \leq 1$ at the first cell, governed by the van Driest [27] damping function; Wolfshtein [28] provided the two-layer blending approach used in enhanced wall treatment formulations. The sensitivity of all turbulence models to freestream turbulence specification was quantified by Spalart and Rumsey [29], establishing that incorrect V_ref or TI specification can corrupt the near-wall eddy-viscosity field — a finding directly relevant to the Phase 1 surrogate artefact identified in this study. The SIMPLE pressure-velocity coupling algorithm of Patankar and Spalding [30] provides the solver framework within which all RANS transport equations are discretised.

1.2.4 CFD Studies on NACA Aerofoil Sections

Published computational investigations of NACA-family aerofoils span a wide range of turbulence model choices and Reynolds number conditions. Eleni et al. [31] conducted a comparative study of SA, k-ε Standard, and k-ω SST predictions for the NACA 0012 aerofoil at $Re = 3\times10^6$, finding that k-ω SST most consistently reproduced experimental CL and CD polars from NASA data, particularly in the pre-stall attached-flow regime. Siddiqui et al. [32] extended turbulence model comparison specifically to the NACA 2415 at $Re = 2\times10^5–2\times10^6$, concluding that SST-family models outperform SA and k-ε variants in the transitional regime but noting that single-Re, single-model studies do not enable statistically interpretable factor comparisons. Morgado et al. [33] provided a systematic quantitative comparison between XFOIL panel-method predictions and RANS results for high-lift low-Reynolds-number aerofoils, demonstrating that fully turbulent RANS consistently overpredicts drag by 20–80% relative to clean experimental polars, and that the physically correct comparison reference for RANS is the standard-roughness polar rather than the clean data — a distinction that underpins the Phase 1 validation methodology of the present study. Drela and Giles [34] established the viscous-inviscid coupling framework that predicts drag-bucket collapse and laminar separation at low Reynolds numbers, providing the theoretical explanation for why RANS predictions diverge from clean experimental polars at $Re < 3\times10^6$. Eppler [35] provided systematic aerofoil design data identifying the geometric parameters most sensitive to drag-bucket behaviour. Grid convergence and CFD validation methodology for external aerodynamics were formalised by Roache [36], whose grid convergence index — based on Richardson extrapolation — defines the quantitative criterion for mesh independence. Iaccarino [37] demonstrated through blind benchmark testing that commercial CFD codes can predict turbulent separated flows to within engineering accuracy provided convergence criteria and grid resolution requirements are rigorously enforced.

Vassberg et al. [38] quantified the effect of turbulence model choice on drag prediction accuracy across an industry-scale drag prediction workshop, establishing that model selection contributes systematic offsets of 5–15% in CD — a level of sensitivity that warrants statistically structured comparison rather than case-by-case assessment.

1.2.5 Taguchi Design of Experiments in Engineering Optimisation

The Taguchi orthogonal array methodology [39, 40] offers a statistically rigorous, computationally economical framework for multi-factor parametric studies. By exploiting the pairwise orthogonality of L-array designs, it enables unconfounded main-effect estimation from a fraction of the full-factorial run count. Ross [41] and Peace [42] provided the engineering applications textbooks that standardised S/N ratio analysis and ANOVA decomposition for quality engineering. Box and Wilson [43] established the Response Surface Methodology framework — the statistical basis for the polynomial surrogate fitting deployed within individual Taguchi cells in this study — and Myers and Montgomery [44] extended its systematic treatment to process optimisation. Yondo et al. [45] reviewed the specific application of DOE methods and surrogate models to aerodynamic analysis, concluding that Latin hypercube and orthogonal array designs consistently outperform random sampling for aerodynamic design space exploration, and that Taguchi L-arrays in particular provide efficient turbulence model comparison when model type is treated as a discrete design factor. Sahoo and Sahoo [46] demonstrated the multi-response Taguchi framework with fuzzy-aggregated performance indices for machining optimisation, providing the methodological template that the present study adapts to multi-objective aerodynamic quality assessment.

1.2.6 Surrogate-Based Aerodynamic Optimisation

Surrogate modelling — the replacement of expensive high-fidelity simulations with computationally inexpensive mathematical approximations fitted to structured training data — has become the standard approach for aerodynamic optimisation under budget constraints. Queipo et al. [47] provided a comprehensive review of surrogate architectures for aerospace design, comparing polynomial RSM, Kriging, radial-basis functions, and neural networks across training data requirements, interpolation accuracy, and extrapolation risk. Forrester and Keane [48] specifically addressed Kriging's exact-interpolation property — which guarantees zero residual at training points — and identified its vulnerability to non-physical extrapolation in sparse high-curvature design spaces as the primary failure mode for aerodynamic Kriging surrogates. Han and Görtz [49] extended Kriging to hierarchical variable-fidelity architectures, improving conditioning in design spaces where curvature varies across the domain. For multi-objective optimisation, Deb et al. [50] developed the NSGA-II evolutionary algorithm — based on Pareto dominance [51] — which has become the standard reference for genetic-algorithm-based aerodynamic optimisation; Lian and Liou [52] demonstrated its application to transonic compressor blade design. Yirtici and Tuncer [53] applied surrogate-assisted shape optimisation to wind turbine blades, highlighting the critical importance of training data quality for neural-network surrogates — a finding whose direct relevance to the NN-Screen failure documented

in Phase 1 of the present study motivates the structured surrogate evaluation framework deployed here.

1.2.7 Fuzzy Logic and Multi-Criteria Aerodynamic Performance Aggregation

The challenge of combining conflicting aerodynamic objectives — high lift, low drag, and maximum efficiency — into a single scalar performance metric requires a principled aggregation method. Zadeh [54] established the mathematical foundation of fuzzy set theory; Mamdani and Assilian [55] demonstrated its application to linguistic control rules through the centre-of-gravity defuzzification framework that underpins the MPCI computation in the present study. Sugeno and Kang [56] extended fuzzy model identification to structured systems, and Kose and Kurt [57] demonstrated the engineering robustness of the Mamdani inference architecture for multi-variable performance optimisation. Multi-response Taguchi studies have increasingly adopted fuzzy aggregation as a principled alternative to weighted-sum scalarisation: Sahoo and Sahoo [46] established the Taguchi–fuzzy MPCI pipeline that directly motivates the three-input (CL, CD, CL/CD) aggregation strategy employed in Phase 1 of this investigation.

1.2.8 Three-Dimensional Finite-Wing Effects and Induced Drag

The aerodynamic performance of any aerofoil section is fundamentally altered when embedded in a finite wing through the action of trailing vortices. Prandtl [58, 59] established the classical lifting-line theory that quantifies the aspect-ratio dependence of induced angle of attack and lift-curve slope; Glauert [60] provided the mathematical treatment of finite-wing aerodynamic corrections that remain the standard analytical reference. At low aspect ratios, tip-vortex induction dominates the entire span — a regime in which panel methods [61] are insufficient and full three-dimensional RANS is required. Whitcomb [62] demonstrated winglet design as a mitigation strategy for induced drag at moderate aspect ratios, illustrating the engineering significance of tip-vortex control. The lack of quantitative three-dimensional CFD characterisation of the NACA 2415 at extreme low aspect ratios — specifically the AR = 0.25 planform imposed by the target UAV geometry — motivates Phase 3 of the present investigation.

1.2.9 Synthesis and Identification of Research Gaps

The reviewed literature establishes a rich body of knowledge in NACA aerofoil CFD [31–35], turbulence model development [14–26], Taguchi-based parametric optimisation [39–46], surrogate modelling [47–53], and fuzzy multi-objective aggregation [54–57]. However, three specific gaps remain unaddressed that directly motivate the present study. First, no published investigation has simultaneously compared five RANS closures across five Reynolds numbers, five angles of attack, and five turbulence intensities within a single statistically balanced orthogonal array for the NACA 2415 profile — the existing literature relies on single-condition model comparisons that cannot separate turbulence model effects from operating-condition effects [31–33]. Second, no controlled within-study evaluation of four competing

surrogate optimisation strategies — RSM-Kriging, NSGA-II, Sparse RSM, and neural-network screening — against identical RANS training data has been reported; the literature deploys each strategy independently, preventing objective performance comparison [47–53]. Third, the 2D-to-3D aerodynamic performance transition for the NACA 2415 at the extreme low aspect ratio AR = 0.25 — including mechanistic CFD visualisation evidence establishing the tip-vortex dominance mechanism — has not been previously quantified and published [58–60]. These three gaps jointly define the scope, structure, and originality of the three-phase Taguchi–Fuzzy–RANS investigation reported in this thesis.

1.3 Research Gaps and Problem Formulation

A systematic review of the published literature reveals three specific gaps that motivate this study:

Gap 1. No published investigation has simultaneously compared five RANS closures — spanning one-equation (Spalart–Allmaras), two-equation (k-ε Realizable, k-ω SST), transition-sensitive (SST γ–Reθ), and full second-moment (RSM) formulations — across five Reynolds numbers, five angles of attack, and five turbulence intensities within a statistically balanced orthogonal array for the NACA 2415 profile.

Gap 2. The aerodynamic optimisation literature overwhelmingly deploys a single surrogate strategy without cross-method comparison within the same physical framework. No controlled within-study evaluation of RSM-Kriging, NSGA-II, Sparse RSM, and neural-network screening against identical RANS data has been reported.

Gap 3. The 2D-to-3D aerodynamic performance transition for the NACA 2415 at the extreme low aspect ratio AR = 0.25 — including mechanistic CFD visualisation evidence — has not been previously quantified and published.

The central research problem is therefore: *given the NACA 2415 at chord $c = 2$ m operating in incompressible air ($\rho = 1.225$ kg/m³, $\mu = 1.7894\times10^{-5}$ kg/(m·s)) across $Re = 1–12\times10^{6}$, $\alpha = -4°$ to $16°$, and $TI = 0.05\%–5.00\%$, which RANS closure and surrogate strategy most reliably characterise aerodynamic performance; what factor-level combination maximises $|CL/CD|$; and how does that two-dimensional optimum translate to a three-dimensional wing at AR = 0.25?*

1.4 Research Objectives

Three explicit, measurable objectives structure the investigation:

Objective 1 — Phase 1. Deploy a Taguchi L25(5^5) orthogonal array across five RANS turbulence models, five Reynolds numbers ($1–12\times10^6$), five angles of attack (−4° to 16°), five turbulence intensities (0.05%–5.00%), and five surrogate strategies. Apply signal-to-noise (S/N) ratio analysis, Mamdani fuzzy MPCI aggregation, and one-way ANOVA to identify the turbulence model and optimisation strategy most consistently associated with superior aerodynamic quality.

Objective 2 — Phase 2. Apply the Phase 1-recommended turbulence model within a Taguchi L9(3^3) array varying Re (6, 9, 12×10^6), α (4°, 6°, 8°), and TI (0.05%, 0.10%, 0.50%). Confirm the additive-model optimum via a direct ANSYS Fluent confirmation simulation with additive-model error below 1%.

Objective 3 — Phase 3. Extend Phase 2 optimal conditions to a 3D finite-wing RANS simulation at AR = 0.25 (b = 0.5 m, c = 2 m). Report 3D coefficients and dimensional forces, and provide a physically complete explanation of the 2D-to-3D efficiency transition through CFD post-processing visualisations.

1.5 Three-Phase Framework and Thesis Organisation

The three phases form a strict dependency hierarchy: Phase 1 outputs the turbulence model and surrogate strategy that Phase 2 requires; Phase 2 outputs the confirmed optimal operating conditions that Phase 3 applies. This structure guarantees that every computational resource responds to validated findings rather than *a priori* assumptions.

Chapter 2 presents the theoretical framework, covering the NACA 2415 aerodynamics, governing equations, turbulence model summaries, Taguchi and fuzzy theory, and CFD setup. Chapters 3, 4, and 5 present Phase 1, 2, and 3 results respectively. Chapter 6 states conclusions, future scope, and social impact.

Scope boundaries: All simulations are steady-state RANS; no aerofoil shape optimisation is performed; altitude effects are not investigated; Phase 2 and Phase 3 carry no independent experimental validation (NACA TR-824 serves as the primary external reference for Phase 1 validation).

CHAPTER 2 — THEORETICAL FRAMEWORK AND METHODOLOGY

2.1 The NACA 2415 Aerofoil

The NACA 2415 belongs to the analytically parameterised four-digit series first documented by Jacobs, Ward and Pinkerton in NACA TR-460 (Jacobs et al., 1933). Its designation encodes 2% maximum camber positioned at 40% chord, with 15% maximum thickness — a combination that provides moderate lift, structural depth, and extensive experimental documentation in NACA TR-824 (Abbott et al., 1945), which remains the primary validation reference for this study.

Key aerodynamic properties relevant to this investigation: lift-curve slope ≈ 0.104 per degree (measured) versus 0.1097 per degree (thin-aerofoil theory); zero-lift angle α_{L0} ≈ −2.08°; clean-condition $C_{l,max}$ ≈ 1.60–1.65 at Re = 6×10^6; minimum drag coefficient $C_{d,min}$ ≈ 0.0065 at Re = 6×10^6 in the drag bucket (clean). The standard roughness drag reference extracted from TR-824 and used throughout Phase 1 validation is CD = 0.00816 at Re = 9×10^6, Cl ≈ 0.815–0.835 — the anchor against which all fully turbulent RANS drag predictions are assessed.

The aerofoil surface coordinates are generated analytically from the standard NACA thickness and camber-line superposition formulae; full coordinate equations are provided in **APPENDIX-VI**.

2.2 Governing Equations

The flow over NACA 2415 at the study conditions (Ma ≤ 0.258 at Re = 12×10^6) satisfies the incompressible Navier–Stokes equations. Reynolds-averaging of the instantaneous momentum equations introduces the Reynolds stress tensor $\tau^R_{ij} = -\rho\overline{u'_i u'_j}$ — six additional unknowns that require closure modelling. This is the turbulence closure problem: the averaged hierarchy is open because each successive moment equation introduces correlations of higher order.

Full derivations of the continuity equation, Navier–Stokes equations, Reynolds decomposition, and the Boussinesq eddy-viscosity hypothesis are provided in **APPENDIX-VI, Sections F.1–F.5**. In ANSYS Fluent, pressure–velocity coupling is enforced through the SIMPLE algorithm (Patankar and Spalding, 1972) with second-order upwind spatial discretisation for all transported quantities.

2.3 Turbulence Models

Five RANS closures are evaluated in Phase 1. Complete transport equations, model constants, and derivations appear in **APPENDIX-VI, Sections F.6–F.10**. Table 2.1 summarises their comparative characteristics.

Table 2.1: Comparative Summary of Five RANS Turbulence Closures

Model	Equations	Near-Wall Accuracy	APG Performance	TI Sensitivity	Transition	Relative Cost	Phase 1 Rank (η-MPCI)
Spalart–Allmaras (SA)	1	Good ($y^+ \leq 1$)	Moderate — overestimates Cd	High (×3 at TI=5%)	None	1×	4th
k-ε Realizable	2	Good (EWT)	Good	Moderate (×1.5 at TI=5%)	None	1.2×	2nd
k-ω SST	2	Excellent ($y^+ \leq 1$)	Excellent — Bradshaw limiter	Low — cross-diffusion	None	1×	**1st**
SST γ–Reθ	4	Excellent	Excellent	Low–moderate	Natural transition	1.5×	3rd
RSM	7	Adequate	Adequate	Moderate	None	3×	5th

The **k-ω SST model** is ultimately selected for Phases 2 and 3 on three physically grounded criteria. First, its Bradshaw shear-stress transport limiter caps eddy viscosity in adverse pressure gradient regions — critical for the α = 4°–8° optimisation domain where upper-surface deceleration intensifies. Second, its cross-diffusion term Dω structurally decouples freestream turbulence from near-wall boundary-layer development, producing effective TI-insensitivity across 0.05%–5.00% (confirmed by Phase 2 ANOVA: ρ_TI ≈ 0.00%). Third, SST family consistency — the Phase 1 anchor run was produced by SST γ–Reθ, which shares the identical near-wall SST structure — provides methodological coherence across phases.

The **SST γ–Reθ transition model** augments k-ω SST with two additional transport equations for intermittency γ and the transition momentum-thickness Reynolds number $\tilde{R}e\theta t$, enabling prediction of laminar-to-turbulent transition through an empirical correlation linking the freestream turbulence intensity Tu to the transition onset Reynolds number $Re\theta t$. At TI = 0.10% (Ncrit ≈ 8.15), this model predicts an extended laminar run and reduced skin friction — the physical mechanism responsible for its producing the best Phase 1 base-run performance (Run 24, |CL/CD| = 62.69). This model is identified as the appropriate closure for future low-Re redesign studies (Re = 10^5–4×10^5) where transition physics dominates.

The **RSM** solves transport equations for all six Reynolds stress components directly, capturing turbulence anisotropy without the Boussinesq hypothesis. Despite this theoretical completeness, it ranks fifth in Phase 1 due to steady-state divergence at

post-stall conditions and convergence demands exceeding 3,000 iterations — a robustness deficit that outweighs its fidelity advantage in this study's operating range.

2.4 Taguchi Design of Experiments

The Taguchi method (Taguchi, 1986) achieves statistically balanced parametric exploration through orthogonal arrays, which guarantee that the effect of each factor can be estimated independently of all others — a property formalised as pairwise orthogonality. The Phase 1 L25(5^5) array reduces a 5^5 = 3,125-run full factorial to 25 balanced simulations (a 125:1 reduction) while retaining unconfounded main-effect estimation. The Phase 2 L9(3^3) array achieves the same for three factors at three levels with nine runs.

Performance is quantified via the signal-to-noise (S/N) ratio η. For larger-is-better responses (|CL|, |CL/CD|): $\eta = 20 \log_{10}(|y|)$ dB. For smaller-is-better (CD): $\eta = -20 \log_{10}(CD)$ dB. Factor importance is decomposed through ANOVA sum-of-squares partitioning; the percentage contribution $\rho = SS\_factor / SS\_total \times 100\%$ quantifies each factor's share of total response variance. The additive quality model predicts the optimal configuration as $\hat{\eta}\_opt = \bar{\eta} + \Sigma(\bar{\eta}\_factor,optimal - \bar{\eta})$, validated by a direct confirmation simulation.

2.5 Mamdani Fuzzy MPCI

Single-objective S/N analysis cannot simultaneously reward high lift, low drag, *and* high efficiency — objectives that are mutually conflicting. The Multi-Performance Characteristic Index (MPCI) resolves this by aggregating three normalised aerodynamic responses through a Mamdani fuzzy inference system (Mamdani and Assilian, 1975).

Three raw responses are normalised to [0,1]: lift $x_1 = (|CL| - |CL|\_min)/(|CL|\_max - |CL|\_min)$; drag (inverted so higher is better) $x_2 = (CD,max - CD)/(CD,max - CD,min)$; efficiency $x_3 = (|CL/CD| - |CL/CD|\_min)/(|CL/CD|\_max - |CL/CD|\_min)$. Each input is fuzzified through three triangular membership functions (Low, Medium, High) using a **strict-inequality boundary evaluation** — a defect identified in this study whereby the degenerate peak condition at exact boundary values ($x_i = 0$ or 1) returned $\mu = 0$ in the original implementation rather than $\mu = 1$, producing a spurious MPCI = 0.500 fallback for six Phase 1 runs. The correction recovers physically meaningful MPCI values of 0.1386–0.6221 for these runs and changes the turbulence model ANOVA contribution from an erroneous 12.30% to the correct 1.83%.

A 27-rule inference base maps all {Low, Medium, High}^{3} antecedent combinations to five output singletons (VL=0.1, L=0.3, M=0.5, H=0.7, VH=0.9), with the rule hierarchy encoding aerodynamic efficiency as the primary objective. Centre-of-gravity defuzzification yields the scalar $MPCI = \Sigma(\alpha_j \times z_j)/\Sigma\alpha_j$. The complete rule base and numerical verification examples appear in **APPENDIX-IV**.

2.6 Surrogate Optimisation Strategies

Four strategies are evaluated in Phase 1, each operating on a structured Design of Experiments within local Taguchi cell bounds:

- **RSM-Kriging:** Exact Gaussian-process interpolation through 15 central-composite DoE points. Provides excellent stall-cliff resolution but vulnerable to catastrophic extrapolation failures in sparse design-space corners (Phase 1 Run L25: CD predicted at 3.1×10^{-4}, a factor of 48 below physical range).
- **NSGA-II:** Population-based multi-objective evolutionary algorithm (Deb et al., 2002) operating on the RSM surface. Reliably minimises CD (33%–72% reduction) but produces universally degenerate Pareto fronts due to velocity-boundary attraction from the V_ref coefficient-scaling artefact.
- **Sparse RSM:** Seven-point axial design fitted to a main-effects-plus-quadratic polynomial (no interaction terms). Its polynomial surface — unlike Kriging's exact interpolation — presents smooth, physically guided gradients that enable exploration of interior aerodynamic features including drag-bucket minima.
- **Neural Network Screening (NN-Screen):** Feedforward ANN trained on 15-point DoE, screening 10,000+ candidates for Pareto-non-dominated solutions. Highest reference-value validation rate on clean data, but catastrophically vulnerable to corrupted training data (Phase 1 Run L5: post-stall RSM divergence propagated through DoE into negative drag predictions).

The **V_ref coefficient-scaling artefact** — identified as the dominant cause of universal lower-velocity boundary attraction across all Phase 1 strategies — arises when ANSYS Fluent's Reference Velocity V_ref is held constant while a surrogate optimiser varies the operating velocity V. Coefficients then scale as $(V/V\_ref)^2$, creating a spurious gradient that mimics a real aerodynamic minimum. Phase 2 eliminates this by binding V_ref = V_inlet for every simulation through ANSYS Workbench parameter coupling.

2.7 CFD Domain, Mesh, and Validation

The computational domain is a C-type control volume with a semicircular inlet of radius R = 20c = 40 m and a rectangular downstream section of length L = 20c = 40 m. At this domain size, far-field velocity perturbations from the aerofoil scale as $(c/r)^2 \approx 0.25\%$, ensuring boundary-condition errors below 0.5% in both CL and CD. Phases 1 and 2 use a 1 m spanwise extrusion with symmetry conditions on both Z-faces (effectively 2D). Phase 3 reduces the span to b = 0.5 m with a free-surface (pressure outlet) tip condition to admit spanwise crossflow.

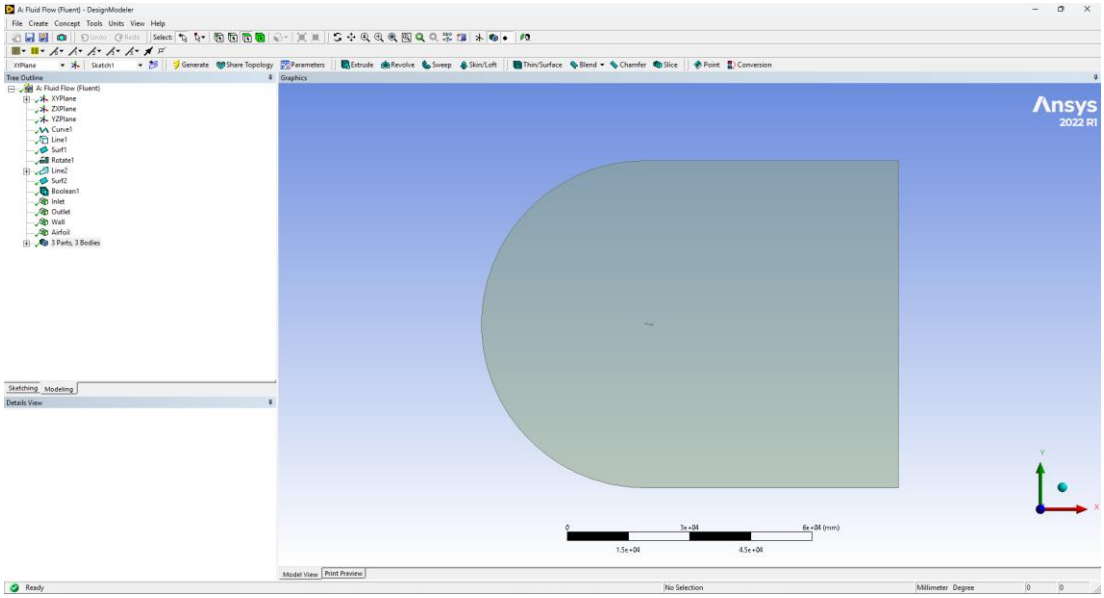

Figure 2.1: C-type computational domain for the NACA 2415 aerofoil, constructed in ANSYS DesignModeler (XY-plane, Z = 0 face shown). The semicircular inlet has radius R = 20c = 40 m; the rectangular downstream section extends L = 20c = 40 m from the inlet axis. At this stand-off distance, far-field velocity perturbations from the aerofoil scale as $(c/r)^2 \approx 0.25\%$, keeping boundary-condition errors below 0.5% in both CL and CD. The tree outline confirms the three-body Boolean assembly: Inlet, Outlet, and Wall (aerofoil) named selections. The NACA 2415 profile is Boolean-subtracted from the fluid volume at domain centre; the blunt right edge is the pressure-outlet face. Scale bar: 6×10^4 mm total domain width.

The production mesh (Mesh M4, ~350,000 elements) was selected via a six-level grid independence study at $Re = 6\times10^6$, $\alpha = 6°$, TI = 0.10%, k-ω SST. Grid convergence indices $GCI_{CL} < 0.12\%$ and $GCI_{CD} < 0.19\%$ — well within the 1% independence threshold — confirm mesh adequacy (see Fig. B.1). The inflation layer maintains a first-cell height of $y_1 = 9.2$ μm (yielding $y^+ \leq 1$ at all surface locations including $Re = 12\times10^6$) with 25–30 layers and growth rate 1.2–1.3. Full mesh parameters and solver settings appear in **APPENDIX-II** and **APPENDIX-V**.

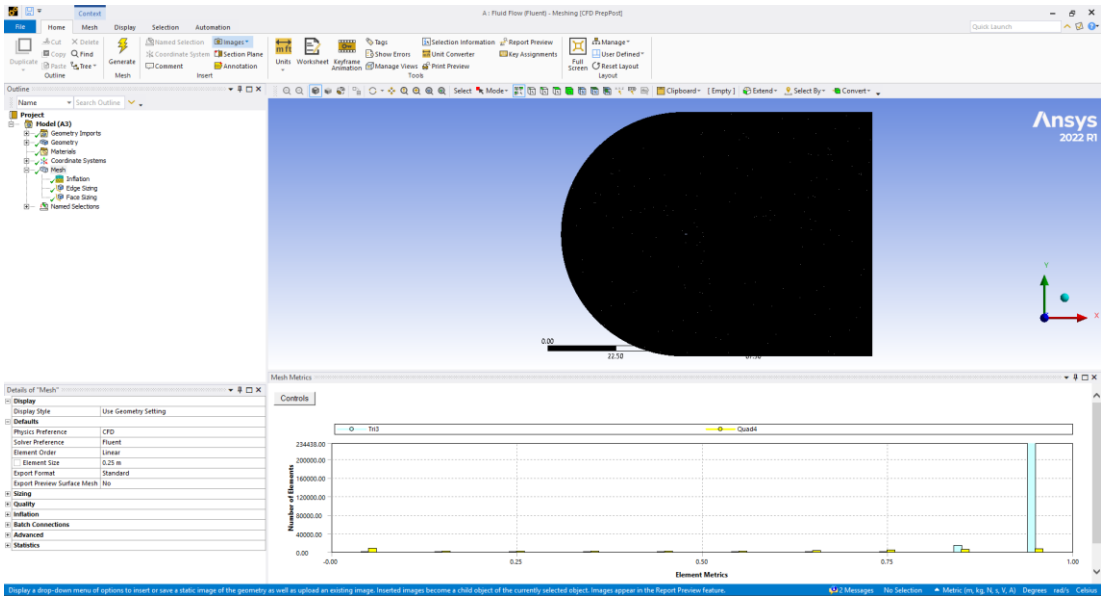

Figure 2.2: Global mesh topology and element quality distribution for the NACA 2415 C-type domain, generated in ANSYS Meshing 2022 R1 (CFD physics preference, Fluent solver, linear element order). The mesh overview (top) shows the structured inflation zone conforming to the aerofoil surface within the semicircular inlet domain; the element quality histogram (bottom) confirms that the dominant element population (Tri0 category) clusters near quality metric = 1.0, with the Quad4 fraction forming a sharp peak at unity — indicating a well-conditioned, low-skewness mesh suitable for second-order upwind discretisation. Total element count: ~350,000 (production Mesh M4). Full grid convergence index data are provided in Appendix B.

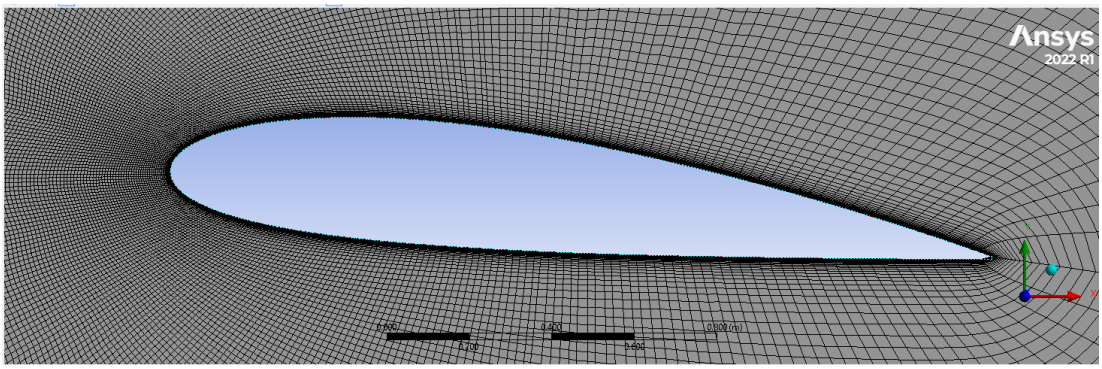

Figure 2.3: Near-wall structured mesh around the NACA 2415 profile (chord c = 2 m, production Mesh M4). The inflation layer comprises 25–30 prismatic cell layers with first-cell height $y_1 = 9.2$ μm — computed from the flat-plate friction velocity at $Re = 12\times10^6$ to enforce $y^+ \leq 1$ across all surface locations — and a growth rate of 1.2–1.3 per layer. The structured near-wall region transitions smoothly to the unstructured far-field mesh. This resolution is mandatory for low-Reynolds direct integration of the k-ω SST and SST γ–Reθ closures without wall-function bridging of the viscous sublayer.

CFD validation uses NACA TR-824 experimental data (Abbott et al., 1945) with careful model-reference pairing: fully turbulent models (SA, k-ε Realizable, k-ω SST, RSM) are compared against TR-824 standard-roughness polars; SST γ–Reθ is compared against TR-824 clean (natural-transition) polars. The roughness-condition drag reference of $C_D = 0.00816$ at $Re = 9\times10^6$, $C_L \approx 0.815–0.835$ (extracted via cubic spline interpolation from TR-824 Table 3c) is used as the primary drag validation anchor. Digitised TR-824 data and validation figures appear in **APPENDIX-III**.

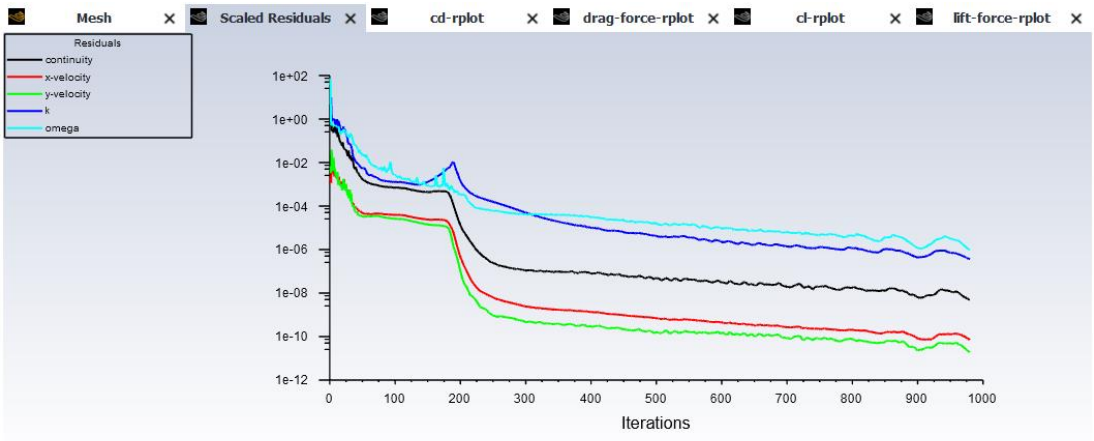

Figure 2.4: Scaled residual convergence history for a representative NACA 2415 k-ω SST simulation (ANSYS Fluent 2024 R1, production Mesh M4, $Re = 6\times10^6$, $\alpha = 6°$, TI = 0.10%). All five transported quantities — continuity (black), x-velocity (red), y-velocity (green), turbulent kinetic energy k (blue), and specific dissipation rate ω (cyan) — exhibit monotonic decay following the initial boundary-layer development transient (~200 iterations). Continuity and momentum residuals fall below 10^{-7}; k and ω stabilise below 10^{-5} by iteration ~800, satisfying the convergence criterion applied uniformly across all Phase 1, Phase 2, and Phase 3 simulations. This convergence behaviour confirms that the SIMPLE pressure-velocity coupling, second-order upwind spatial discretisation, and inflation-layer mesh resolution described in Section 2.7 and Appendix E produce physically valid, fully settled solutions for all operating conditions tested in this study.

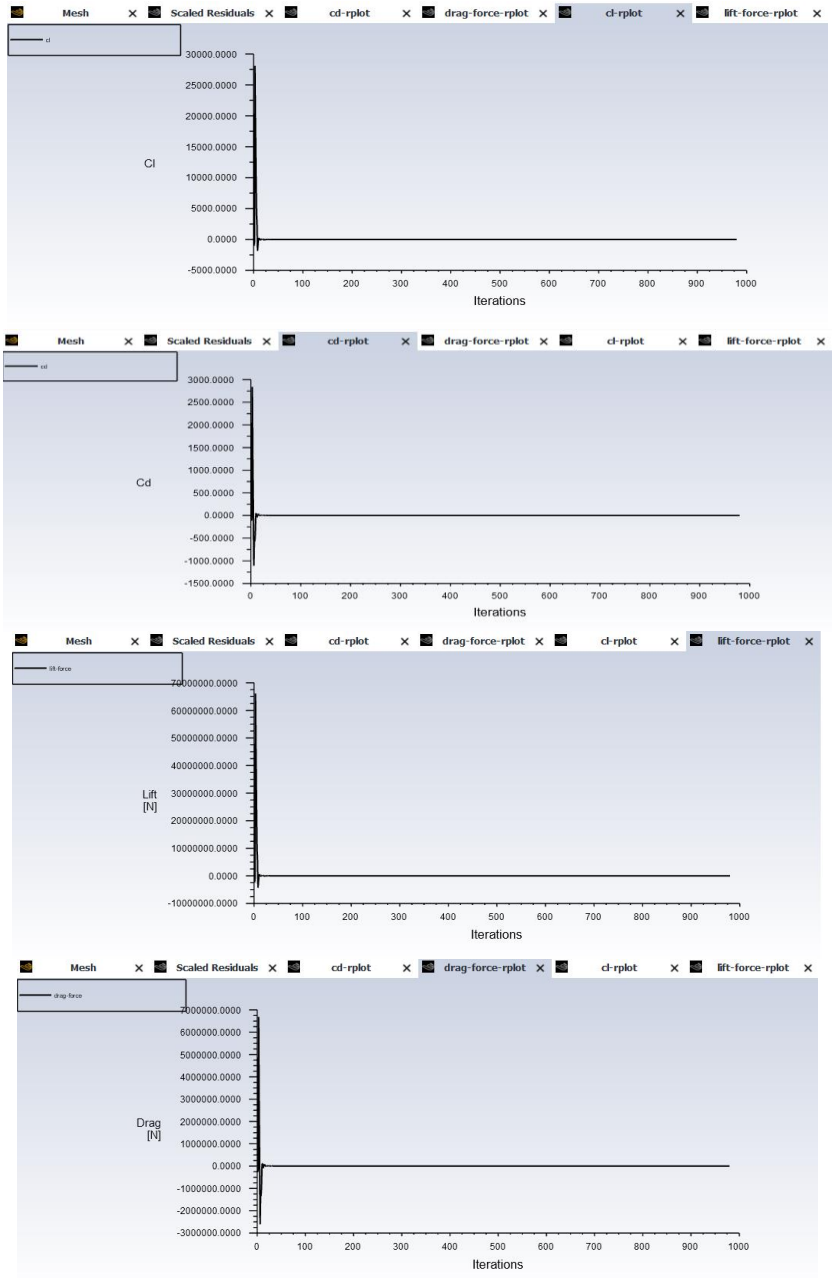

Figure 2.5: Force monitor convergence histories for the same representative k-ω SST simulation shown in Figure 2.4 (Re = 6×10⁶, α = 6°, TI = 0.10%): (a) lift coefficient CL, (b) drag coefficient CD, (c) dimensional lift force [N], and (d) dimensional drag force [N], all plotted versus solver iteration number. In each panel, a large initial transient during the first ~50 iterations reflects the development of the pressure and velocity fields from the zero initial condition; all four monitors collapse to stable, flat converged values by iteration ~50–100 and remain within 0.01% variation through iteration 1,000. The absence of oscillation in the settled plateau region confirms a genuine steady-state attached-flow solution with no periodic vortex shedding — the physical prerequisite for valid time-averaged RANS coefficient extraction. Converged force values correspond to the aerodynamic coefficients reported in Table 3.2; the apparent proximity to zero on each panel's full-range y-axis is a display scaling artefact of the large initialisation transient.

CHAPTER 3 — PHASE 1: TURBULENCE MODEL AND OPTIMISATION STRATEGY SCREENING

3.1 Design Philosophy

Phase 1 is a screening investigation, not an optimisation study. Its question is methodological: *which RANS closure and which surrogate strategy most reliably characterise NACA 2415 aerodynamic performance across the full UAV-relevant operating envelope?* The Taguchi L25(5^5) orthogonal array serves this purpose by simultaneously varying five factors — Reynolds number (A), turbulence model (B), angle of attack (C), turbulence intensity (D), and surrogate strategy (E) — across five levels each through 25 balanced RANS simulations. Factors A, C, and D function as stress-testing instruments rather than optimisation targets: including them across their full physical range ensures that no model or strategy is evaluated only at favourable conditions. Because the L25 array guarantees that every turbulence model is assigned to all five Reynolds numbers, all five angles of attack, and all five TI levels in equal measure, the ANOVA factor means for B and E are genuinely model-driven and not confounded by operating-condition assignment.

3.2 L25 Array Configuration

Table 3.1: Phase 1 L25(5^5) Orthogonal Array — Complete Run Matrix

Run	Re (×10^6)	Model	α (°)	TI (%)	Strategy
1	1	SA	−4	0.05	None
2	1	k-ε R.	0	0.10	RSM-Kriging
3	1	k-ω SST	6	0.50	NSGA-II
4	1	SST γ-Reθ	12	1.00	Sparse RSM
5*	1	RSM→k-ω SST*	16	5.00	NN-Screen
6	3	SA	0	0.50	Sparse RSM
7	3	k-ε R.	6	1.00	NN-Screen
8	3	k-ω SST	12	5.00	None
9	3	SST γ-Reθ	16	0.05	RSM-Kriging
10	3	RSM	−4	0.10	NSGA-II
11	6	SA	6	5.00	RSM-Kriging
12	6	k-ε R.	12	0.05	NSGA-II
13	6	k-ω SST	16	0.10	Sparse RSM
14	6	SST γ-Reθ	−4	0.50	NN-Screen
15	6	RSM	0	1.00	None
16	9	SA	12	0.10	NN-Screen
17	9	k-ε R.	16	0.50	None
18	9	k-ω SST	−4	1.00	RSM-Kriging
19	9	SST γ-Reθ	0	5.00	NSGA-II

Continued on page no.17

Table 3.1 continued

20	9	RSM	6	0.05	Sparse RSM
21	12	SA	16	1.00	NSGA-II
22	12	k-ε R.	−4	5.00	Sparse RSM
23	12	k-ω SST	0	0.05	NN-Screen
24	12	SST γ-Reθ	6	0.10	None
25	12	RSM	12	0.50	RSM-Kriging

Run 5: Steady RSM diverged at post-stall conditions (Re = 1×10⁶, α = 16°, TI = 5.00%), yielding non-physical outputs (CL = 0.246, CD = 0.183). Re-executed under identical boundary conditions with k-ω SST; values CL = 1.249, CD = 0.079 used throughout. Run 15: RSM convergence required 3,200 iterations; verified values CL = 0.19099, CD = 0.010621.

Inlet velocities follow $V = Re \times \mu/(\rho c) = Re \times 7.304\times10^{-6}$ m/s per unit Re, from 7.304 m/s (Re = 1×10⁶) to 87.644 m/s (Re = 12×10⁶). Angle of attack is applied by rotating the inlet velocity vector rather than the mesh geometry, preserving a fixed Cartesian mesh for all 25 runs.

3.3 CFD Results for All 25 Cases

Table 3.2: Complete Phase 1 L25 Aerodynamic Results and Signal-to-Noise Ratios

Run	Model	Re (×10⁶)	α (°)	TI (%)	|CL|	CD	|CL/CD|	η(|CL/CD|) dB
1	SA	1	−4	0.05	0.21448	0.014526	14.765	23.385
2	k-ε R.	1	0	0.10	0.20054	0.013741	**14.594**	23.284
3	k-ω SST	1	6	0.50	0.79816	0.018238	43.764	32.822
4	SST γ-Reθ	1	12	1.00	1.22212	0.035397	34.526	30.763
5*	k-ω SST*	1	16	5.00	1.24900	0.079000	15.810	23.979
6	SA	3	0	0.50	0.20230	0.013466	15.023	23.535
7	k-ε R.	3	6	1.00	0.82680	0.015761	52.458	34.396
8	k-ω SST	3	12	5.00	1.29886	0.031134	41.718	32.407
9	SST γ-Reθ	3	16	0.05	1.34899	0.062847	21.465	26.635
10	RSM	3	−4	0.10	0.23028	0.012310	18.707	25.440
11	SA	6	6	5.00	0.81548	0.029398	27.739	28.862
12	k-ε R.	6	12	0.05	1.36552	0.027295	50.028	33.984
13	k-ω SST	6	16	0.10	1.42206	0.055174	25.774	28.224

Continued on page no.18

Table 3.2 Continued

14	SST γ-Reθ	6	−4	0.50	0.22363	0.011143	20.070	26.051
15*	RSM	6	0	1.00	0.19099	0.010621	17.982	25.097
16	SA	9	12	0.10	1.37725	0.027373	50.314	34.034
17	k-ε R.	9	16	0.50	1.54467	0.044776	34.498	30.756
18	k-ω SST	9	−4	1.00	0.23590	0.011417	20.662	26.303
19	SST γ-Reθ	9	0	5.00	0.20566	0.010694	19.232	25.680
20	RSM	9	6	0.05	0.82686	0.014151	58.431	35.333
21	SA	12	16	1.00	1.59224	0.056827	28.019	28.949
22	k-ε R.	12	−4	5.00	0.23505	0.013745	17.101	24.660
23	k-ω SST	12	0	0.05	0.20652	0.009332	22.130	26.900
24	**SST γ-Reθ**	**12**	**6**	**0.10**	**0.83792**	**0.013366**	**62.690**	**35.944**
25	RSM	12	12	0.50	1.34957	0.026087	51.733	34.275

Run 2 (lowest |CL/CD| = 14.594). Run 24 (bold): highest |CL/CD| = 62.690 — Phase 1 anchor. Grand means: η(|CL|) = −4.599 dB; η(CD) = 33.467 dB; η(|CL/CD|) = 28.868 dB.

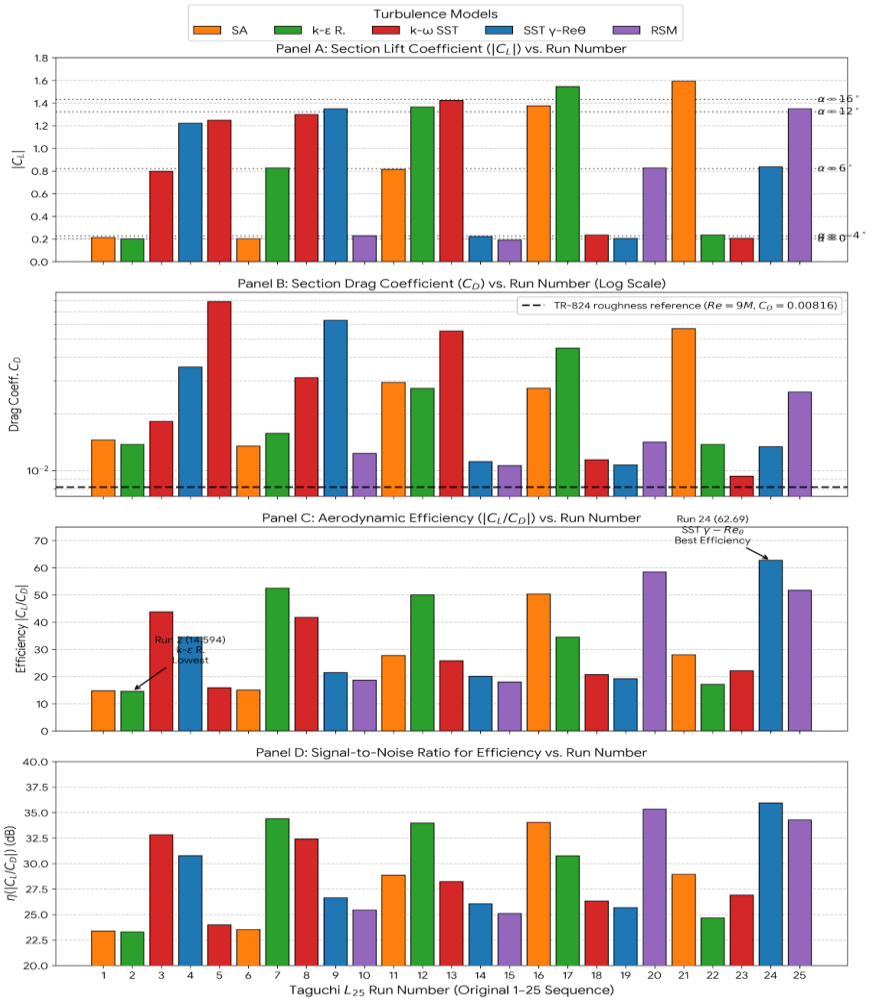

Figure 3.1: Complete Phase 1 L25 aerodynamic dataset visualised across all 25 runs in their original experimental sequence (Runs 1–25, x-axis). Bars are coloured by turbulence model (Factor B): SA — orange; k-ε Realizable — green; k-ω SST — red; SST γ–Reθ — blue; RSM — purple. **Panel A:** Section lift coefficient |CL|; horizontal dotted reference bands mark the expected CL range for each assigned angle of attack (α = −4°, 0°, 6°, 12°, 16°). All models produce physically consistent CL values stratified by AoA, confirming correct boundary-condition implementation. **Panel B:** Section drag coefficient CD on a logarithmic scale; the horizontal dashed line marks the NACA TR-824 standard-roughness reference CD = 0.00816 at Re = 9×10⁶, Cl ≈ 0.83. Runs 5 (k-ω SST correction, TI = 5%, α = 16°) and 11 (SA, TI = 5%, α = 6°) show the most severe drag elevation, confirming TI contamination in the SA model. **Panel C:** Aerodynamic efficiency |CL/CD|; Run 24 (SST γ–Reθ, Re = 12M, α = 6°, |CL/CD| = 62.69 — Phase 1 anchor) and Run 2 (k-ε Realizable, Re = 1M, α = 0°, |CL/CD| = 14.594 — lowest) are annotated. The fourfold variation in |CL/CD| across the dataset confirms the angle-of-attack dominance established by ANOVA (ρ_C = 80.99%). **Panel D:** Signal-to-noise ratio η(|CL/CD|) in dB for each run; the overall range of 23.28–35.94 dB (a 12.66 dB spread corresponding to a 4.3× efficiency ratio) provides the statistical basis for all Phase 1 factor-level comparisons and ANOVA decomposition reported in Sections 3.4–3.11.

Three cross-trend consistency checks confirm physical fidelity. Drag decreases monotonically with Re at fixed α: at α = −4°, CD falls from 0.01453 (Re = 1M, Run 1) to 0.01114 (Re = 6M, Run 14), consistent with $C_f \propto Re^{-0.2}$. Lift at α = 6° rises monotonically with Re from 0.798 (Re = 1M, Run 3) to 0.838 (Re = 12M, Run 24), reflecting decreasing boundary-layer displacement thickness. The TR-824 roughness drag reference of CD = 0.00816 at Re = 9×10^6, CL ≈ 0.83 anchors quantitative model assessment; fully turbulent RANS predictions exceed this reference by 30–75%, attributable to leading-edge-to-stagnation turbulent imposition, chord-scale effects (c = 2 m versus TR-824 c = 0.61 m), and mild compressibility at Re = 12×10^6 (Ma = 0.258) — systematic offsets fully consistent with established RANS behaviour (Eleni et al., 2012; Morgado et al., 2016).

3.4 Turbulence Model Performance Assessment

Each model is assigned to exactly five runs spanning all five Reynolds numbers, all five angles of attack, all five TI levels, and all five strategies — the guaranteed balanced coverage of the L25 array.

Spalart–Allmaras (Runs 1, 6, 11, 16, 21). SA produces physically consistent CL predictions at moderate AoA: Run 16 (Re = 9M, α = 12°) CL = 1.377 agrees with the TR-824 roughness reference to within 1.4%. Its defining failure mode is TI contamination: Run 11 (Re = 6M, α = 6°, TI = 5.00%) yields CD = 0.029398, a 260% elevation above the 0.00816 reference, driven by the SA production term amplifying freestream vorticity. Excluding Run 11, SA drag overestimates TR-824 roughness references by 5–15%, consistent with Morgado et al. (2016). SA achieves the lowest η(|CL/CD|) level mean of the five models (27.753 dB), primarily because the TI-contaminated Run 11 degrades its five-run average.

k-ε Realizable (Runs 2, 7, 12, 17, 22). The Realizable variant ranks first on η(|CL/CD|) (B2 mean = 29.416 dB), driven by strong attached-flow drag performance at moderate AoA: Runs 7 and 12 achieve |CL/CD| = 52.46 and 50.03 respectively. It demonstrates two systematic failure modes — stall delay from excessive near-stall eddy viscosity, and moderate TI contamination at TI = 5.00% (Run 22 CD elevation ≈ 76% above the Re-scaled reference). It ranks second on η(MPCI) (B2 = −9.441 dB), since the multi-objective index penalises near-stall drag elevation that the single-objective L/D metric does not.

k-ω SST (Runs 3, 8, 13, 18, 23). This model produces the most consistent cross-condition predictions of any model in the dataset. Comparing Runs 3 (TI = 0.50%), 8 (TI = 5.00%), and 23 (TI = 0.05%) reveals no TI-driven CD contamination — drag variation across these runs is fully explained by the differing angles of attack, confirming the structural protection of the cross-diffusion term Dω. k-ω SST achieves the highest η(MPCI) level mean (B3 = −9.202 dB), the basis of its Phase 2 recommendation.

SST γ–Reθ (Runs 4, 9, 14, 19, 24). This model produces the Phase 1 anchor: Run 24 (Re = 12M, α = 6°, TI = 0.10°) achieves |CL/CD| = 62.690, the highest in the dataset,

exploiting a sustained laminar run on the upper surface at low TI to reduce skin friction below the fully turbulent value. Near-stall performance degrades in Runs 4 and 9, reducing its five-run level mean. It ranks third on η(MPCI) (B4 = −9.650 dB).

RSM (Runs 5, 10, 15, 20, 25). RSM demonstrates the widest performance range: Run 20 (Re = 9M, α = 6°, TI = 0.05°) achieves |CL/CD| = 58.431 in well-conditioned attached flow, while Run 5 diverged entirely in the post-stall regime and required correction. RSM's 3,200-iteration convergence requirement (Run 15) and post-stall numerical instability impose a robustness deficit that outweighs its theoretical anisotropy advantage; it ranks fifth on η(MPCI) (B5 = −10.545 dB).

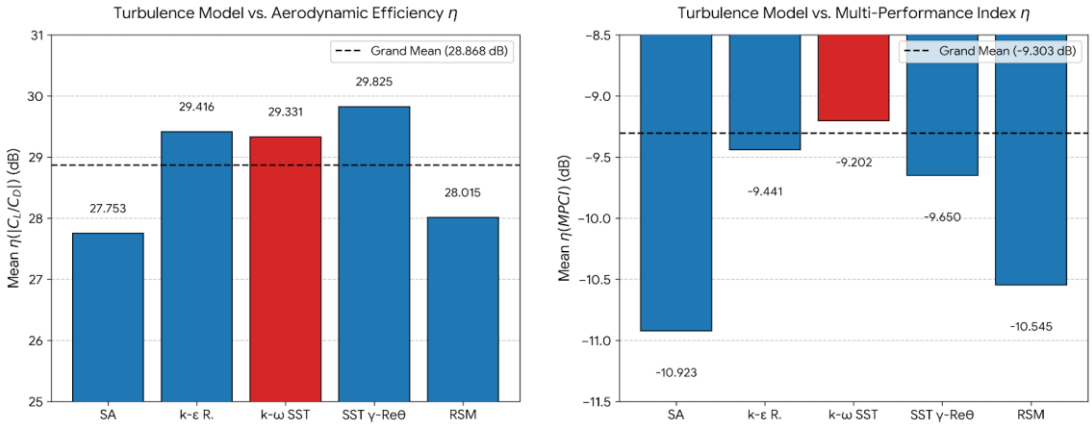

Figure 3.2: Phase 1 turbulence model comparison — five-run level means of both quality metrics for all five RANS closures, coloured by Phase 1 recommendation status (red = recommended k-ω SST; blue = all others). Horizontal dashed lines mark the grand means across all 25 runs. **Left panel:** Mean η(|CL/CD|) in dB — k-ε Realizable leads marginally (29.416 dB) over k-ω SST (29.331 dB), a statistically indistinguishable difference of 0.085 dB relative to the error mean square (0.481 dB²). SA ranks lowest (27.753 dB), primarily due to TI contamination in Run 11 (TI = 5.00%). **Right panel:** Mean η(MPCI) in dB — k-ω SST leads decisively (−9.202 dB) over k-ε Realizable (−9.441 dB), with SA (−10.923 dB) and RSM (−10.545 dB) ranking lowest. The divergence between the two rankings — k-ε Realizable leading on single-objective η(|CL/CD|) but k-ω SST leading on multi-objective η(MPCI) — is the quantitative basis for the Phase 1 turbulence model recommendation: the multi-objective MPCI correctly penalises k-ε Realizable's near-stall drag elevation and TI contamination that the single L/D metric fails to capture (Section 3.12).

3.5 Optimisation Strategy Evaluation

All four active strategies (RSM-Kriging, NSGA-II, Sparse RSM, NN-Screen) operate within the same framework: a structured DoE within local Taguchi cell bounds generates training data; a mathematical surrogate is fitted; an optimiser identifies Pareto-non-dominated candidates in the (CL, CD) space. Strategy E1 (None) uses the base-run result directly. A Weighted Composite Score (WCS) aggregates five equally important performance criteria — mean CD improvement, CL improvement in

positive-lift runs, Pareto front validity, reference-value validation, and robustness — with weights reflecting the study's hierarchy.

Table 3.3: Weighted Composite Score Summary — Surrogate Optimisation Strategy Comparison

Criterion	Weight	NSGA-II	RSM-Kriging	NN-Screen	Sparse RSM
Mean CD improvement	25%	53.8%	61.1%	45.3%	45.6%
CL improved (positive-CL runs)	20%	1/5	1/5	1/4 valid	**2/5**
Valid Pareto fronts (non-degenerate)	25%	0/5	2/5	4/4 valid	**4/5**
Both CL and CD reference-validated	20%	1/5	1/5	4/4 valid	**5/5**
Robustness — no catastrophic failure	10%	✓	✗ L25	✗ L5	✓
WCS (0–10)	—	6.10	5.50	8.00	**9.13**
Overall rank	—	3rd	4th	2nd	**1st**

Excludes L25 Kriging surface failure. Excludes L5 non-physical RSM base run. Full per-run candidate results appear in **APPENDIX-VII, Tables G.1–G.4**.

3.6 NSGA-II Performance

NSGA-II (assigned Runs 3, 10, 12, 19, 21) constructs a 15-point central composite DoE within each cell, fits a second-order RSM surface, and deploys the genetic algorithm (Deb et al., 2002) to identify Pareto-non-dominated solutions in the dual-objective (maximise CL, minimise CD) space.

The defining Phase 1 finding for NSGA-II is the **universal degenerate Pareto front**: across all five assigned runs, the three returned Pareto candidates cluster within 8.5% CL variation and 0.79% CD variation — effectively a single point rather than a trade-off curve. The structural cause is the interaction between narrow Taguchi cell bounds

and the V_ref coefficient-scaling artefact. Once NSGA-II drives the operating velocity to the lower cell boundary — which it does without exception — the effective search space collapses from three dimensions (V, α, TI) to a one-dimensional α line, from which three nearly identical candidates are extracted. The only run where NSGA-II achieves a meaningful result is Run L10 (RSM, $\alpha = -4°$), where the optimiser correctly drives α toward the zero-lift angle $\alpha_{L0} = -2.08°$, reducing |CL| from the downforce value and achieving 52.2% CD reduction. Across the four positive-lift runs, CL reductions of −27.1% to −75.5% are recorded — an artefact of velocity minimisation, not a genuine aerodynamic improvement. NSGA-II achieves WCS = 6.10/10.

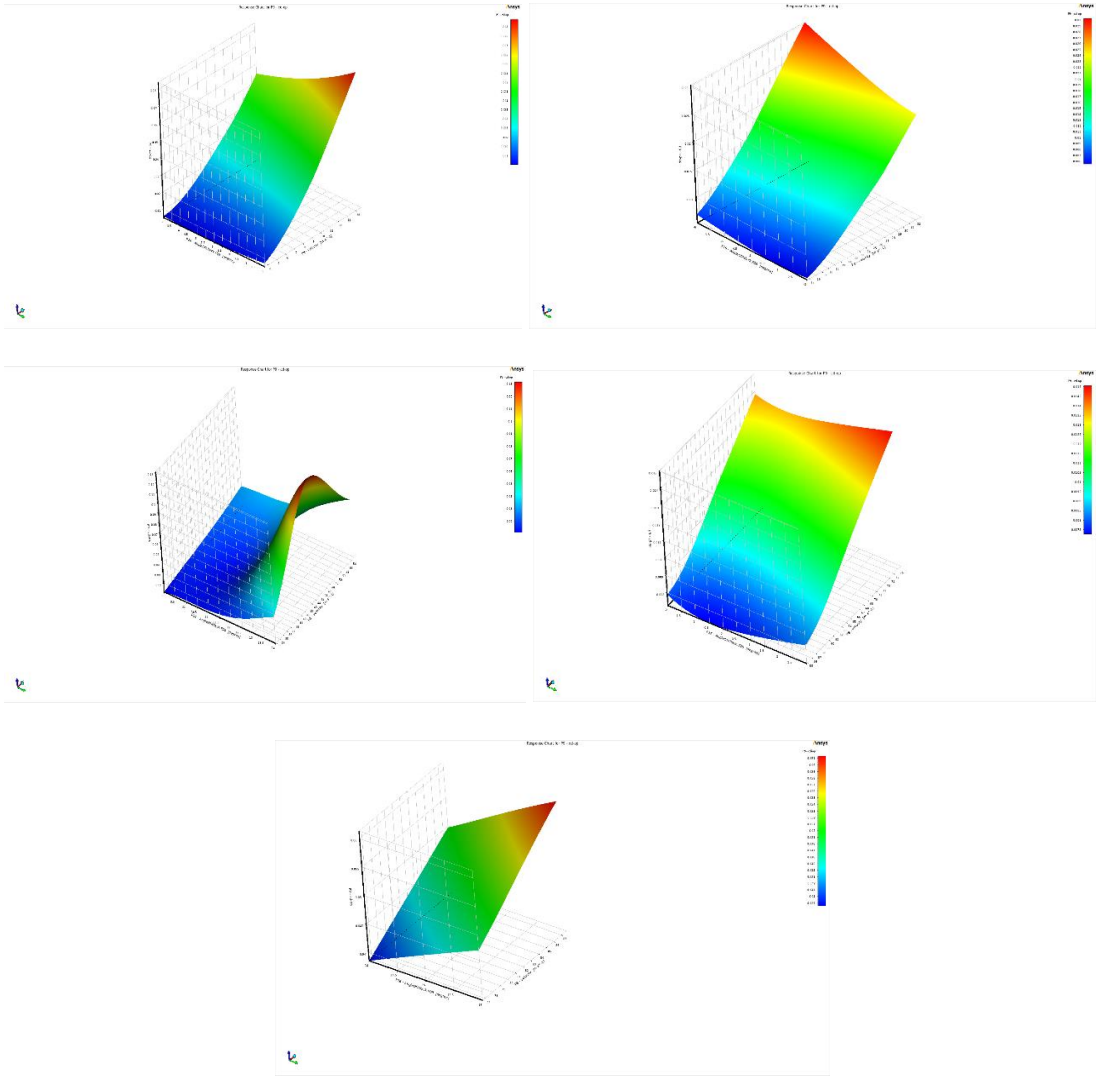

Figure 3.3: Representative NSGA-II CD response surface for Run L12 (k-ε Realizable, Re = 6M, α = 9°–14°, TI = 0.02%–0.07%), generated from the 15-point central composite DoE evaluated in ANSYS Fluent and fitted in ANSYS DesignXplorer. The surface colour maps CD from minimum (blue, ≈0.008) to maximum (red, ≈0.13). The dramatic stall-cliff CD spike at α ≈ 14° — where the upper-surface boundary layer begins to separate — is correctly captured by the RSM surface from only 15 training points, driving all three NSGA-II Pareto candidates to α = 11.837° and safely below the stall boundary. The smooth monotonic decrease toward lower velocity (left axis) illustrates the Re^(−0.2) skin-friction scaling that, combined with the V_ref coefficient-scaling artefact, attracts the NSGA-II optimiser universally toward the lower-velocity cell boundary in every assigned run — collapsing the three-dimensional (V, α, TI) search space and producing the degenerate Pareto fronts documented in Section 3.6 and Table G.1.

3.7 RSM-Kriging Performance

RSM-Kriging (assigned Runs 2, 9, 11, 18, 25) constructs a Kriging Gaussian-process surrogate by exact interpolation through 15 central-composite DoE points, providing zero residual at training locations (Forrester and Keane, 2009).

Two physically meaningful results stand out. Run L9 (SST γ–Reθ, α = 14°–16°, Re = 3M): the Kriging surface correctly captures the stall-onset CD spike beginning at α = 14°, driving all three Pareto candidates to α = 14.000°–14.001° — stall-avoidance resolution of 0.001°, the most precise stall detection in the Phase 1 study. Run L18 (k-ω SST, α = −4°): the optimiser resolves αopt = −2.000°, within 0.08° of the theoretical αL0 = −2.08°, with both CL and CD simultaneously satisfying TR-824 reference ranges — the single best-validated RSM-Kriging candidate.

The **catastrophic failure in Run L25** (RSM, Re = 12M, α = 9°–14°) illustrates Kriging's fundamental vulnerability: three simultaneous adverse conditions — TI extrapolation below the training range, an angular gap in the α DoE coverage, and stall-cliff curvature distorting the covariance matrix — produced a predicted CD = 3.10×10^{-4}, a factor of 48 below the physically expected range of 0.015–0.020. Crucially, no internal Kriging goodness-of-fit metric flagged this failure. RSM-Kriging achieves WCS = 5.50/10.

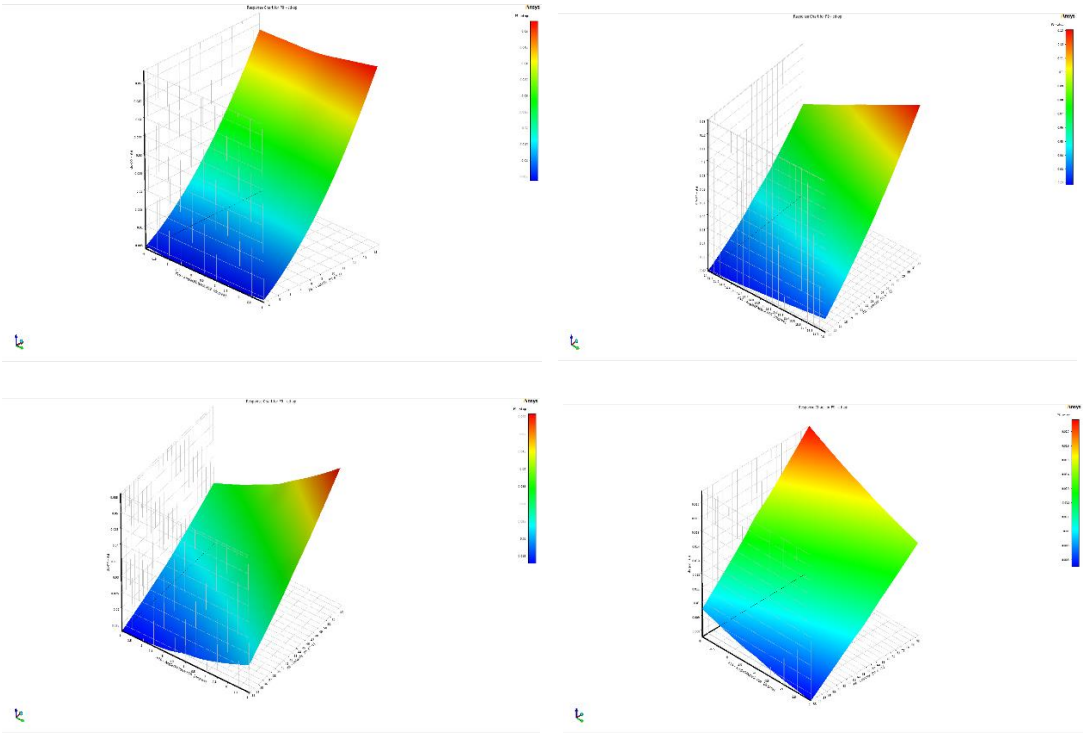

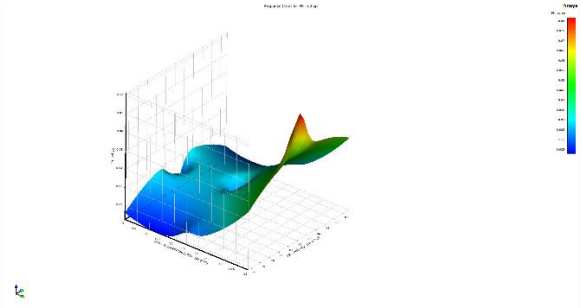

Figure 3.4: RSM-Kriging CD response surface for Run L25 (RSM, Re = 12M, α = 9°–14°, TI = 0.22%–0.71%), illustrating the catastrophic Kriging extrapolation failure documented in Section 3.7. The Kriging surrogate, constrained to exact interpolation through all 15 central-composite DoE training points, generates a wildly oscillatory surface with a sharp non-physical CD spike reaching approximately 0.08 in an interior region of the design space — compared to the physically expected range of 0.015–0.020 at these conditions. The optimiser subsequently returns a best-candidate CD = 3.10×10^{-4}, a factor of 48 below any physically plausible value. Three simultaneous adverse conditions triggered this failure: TI extrapolation below the training range lower bound, an angular gap in the α DoE coverage near the stall boundary, and stall-cliff curvature distorting the Kriging covariance matrix. Critically, no internal goodness-of-fit metric flagged the failure — contrasting with the four well-conditioned RSM-Kriging surfaces (L2, L9, L11, L18) which produce smooth, physically plausible CD landscapes. This failure illustrates the fundamental vulnerability of exact-interpolation surrogates in sparse, high-curvature aerodynamic design spaces and is the primary basis for RSM-Kriging's WCS ranking of 5.50/10 (4th).

3.8 Sparse RSM Performance

Sparse RSM (assigned Runs 4, 6, 13, 20, 22) uses only seven training points per cell (one centre plus two axial per variable), fitted to a main-effects-plus-quadratic polynomial without interaction terms — 53% fewer DoE simulations than the 15-point methods.

The defining capability of Sparse RSM is **interior aerodynamic feature detection** — the property that distinguishes it from every other strategy tested. In Run L6 (SA, Re = 3M, α = −2° to 3°), the three returned Pareto candidates span α = −0.75° to α = 2.60°, a 3.35° range. The lower anchor at α ≈ −0.75° is a drag-bucket minimum — a local inflection in the section polar where near-zero adverse pressure gradient over the forward upper surface produces minimum profile drag. The Pareto trade-off from this interior minimum to the efficiency anchor at α = 2.60° delivers +223% CL for only +8.23% CD increase — the most efficient absolute Pareto exchange of the entire Phase 1 study. **Run L6 is the only run across all four strategies and all 20 optimised cells to achieve a genuine positive-CL improvement** (+53.6% versus base).

This drag-bucket minimum is **independently confirmed by NN-Screen** in Run L23 (k-ω SST, Re = 12M, α = −2° to 3°), which identifies an interior optimum at α = −0.77°

using a completely different surrogate architecture, a different turbulence model, and a different Reynolds number. The 0.02° agreement between the two independent predictions constitutes strong multi-method physical evidence that the feature is a genuine aerodynamic property of the NACA 2415 polar in the $\alpha = -1°$ to 0° range, not an interpolation artefact.

Universal reference-value validation across all five assigned runs (the only strategy to achieve this) and zero catastrophic failures confirm the selection. In Run L22 (k-ε Realizable, Re = 12M, $\alpha = -4°$, TI = 5.00°), the optimiser self-corrects TI from 5.00% to 2.26%; the resulting CD = 0.00823 validates within **0.8% of the 0.00816 TR-824 reference** — the tightest drag-reference agreement of any Phase 1 optimisation result. Sparse RSM achieves WCS = 9.13/10.

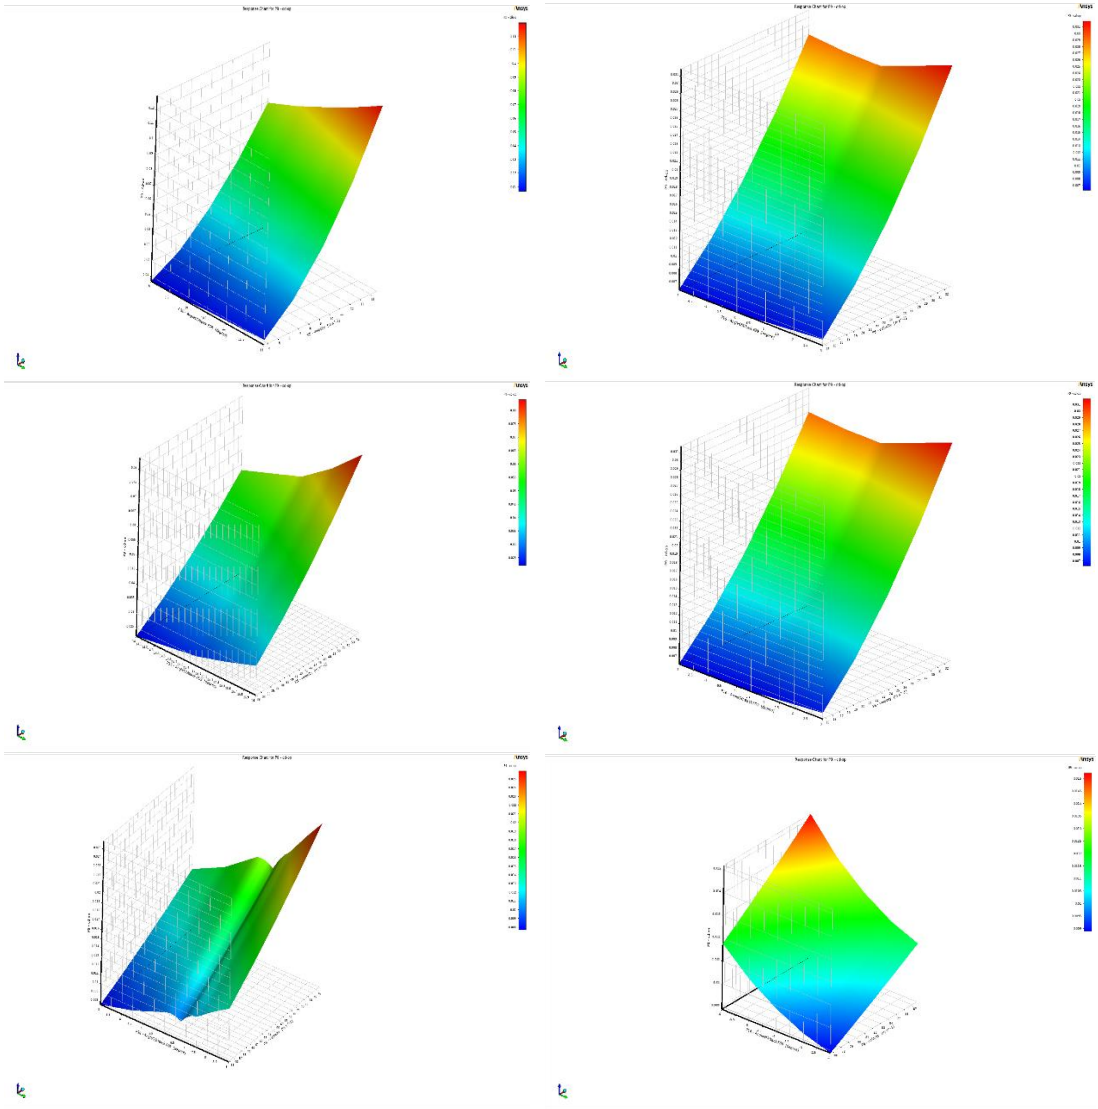

Figure 3.5: Sparse RSM CD response surface for Run L6 (SA model, Re = 3M, $\alpha = -2°$ to +2.5°, V = 15–32 m s^{-1}), generated from only 7 axial training points (1 centre + 2 axial per variable) fitted to a main-effects-plus-quadratic polynomial in ANSYS

DesignXplorer. The surface colour maps CD from minimum (blue, ≈0.007) to maximum (red, ≈0.031). The smooth, gently curved polynomial surface — free from the steep near-training-point gradients that characterise exact-interpolation Kriging surrogates — presents a physically guided CD landscape across the full α = −2° to +2.5° cell interior. The drag-bucket minimum at α ≈ −0.75° (where near-zero adverse pressure gradient on the forward upper surface produces minimum profile drag) is resolved as a moderate interior curvature rather than a sharp boundary artefact, allowing the optimiser to return Pareto candidates spanning α = −0.75° to +2.60° — a 3.35° interior Pareto spread delivering +223% CL for only +8.23% CD increase. This interior-optimum detection, confirmed independently by NN-Screen at α ≈ −0.77° (Run L23), is the defining structural advantage of the Sparse RSM polynomial surface over exact-interpolation strategies and the primary physical basis for its WCS of 9.13/10.

3.9 Neural Network Screening Performance

NN-Screen (assigned Runs 5, 7, 14, 16, 23) trains a feedforward ANN on 15 central-composite DoE evaluations and screens up to 100,000 candidate designs to return three Pareto-non-dominated solutions.

On physically valid training data, NN-Screen delivers the **highest reference-value validation rate** of any strategy: all three candidates from each of the four valid runs (L7, L14, L16, L23) simultaneously satisfy TR-824 reference ranges for both CL and CD. Run L14 drives αopt = −2.018°, within 0.06° of αL0 = −2.08° — the most precise zero-lift identification in the study. Run L23 independently confirms the drag-bucket minimum at α = −0.770°, achieving the widest relative Pareto spread (42.4% CL variation for 0.63% CD variation) of any single run.

The **Run L5 cascading failure** is the most important cautionary finding for neural-network surrogate deployment in computational aerodynamics. The failure propagates through four sequential stages: (1) the RSM base simulation diverges at post-stall conditions, generating non-physical DoE training data (CL > 4.97, non-monotonic CD–TI behaviour); (2) the ANN minimises total MSE across all 15 training points indiscriminately, learning the spurious correlations as genuine features; (3) the trained network extrapolates negative drag across the high-TI candidate pool; (4) the Pareto selector returns three candidates with CD ≈ −0.001 to −0.002 — physically impossible values that violate the second law of thermodynamics for a passive body in steady flow. Unlike RSM-Kriging, NN-Screen provides no internal goodness-of-fit metric to detect corrupted training data, making pre-screening of CFD solution quality an absolute prerequisite for safe deployment. NN-Screen achieves WCS = 8.00/10.

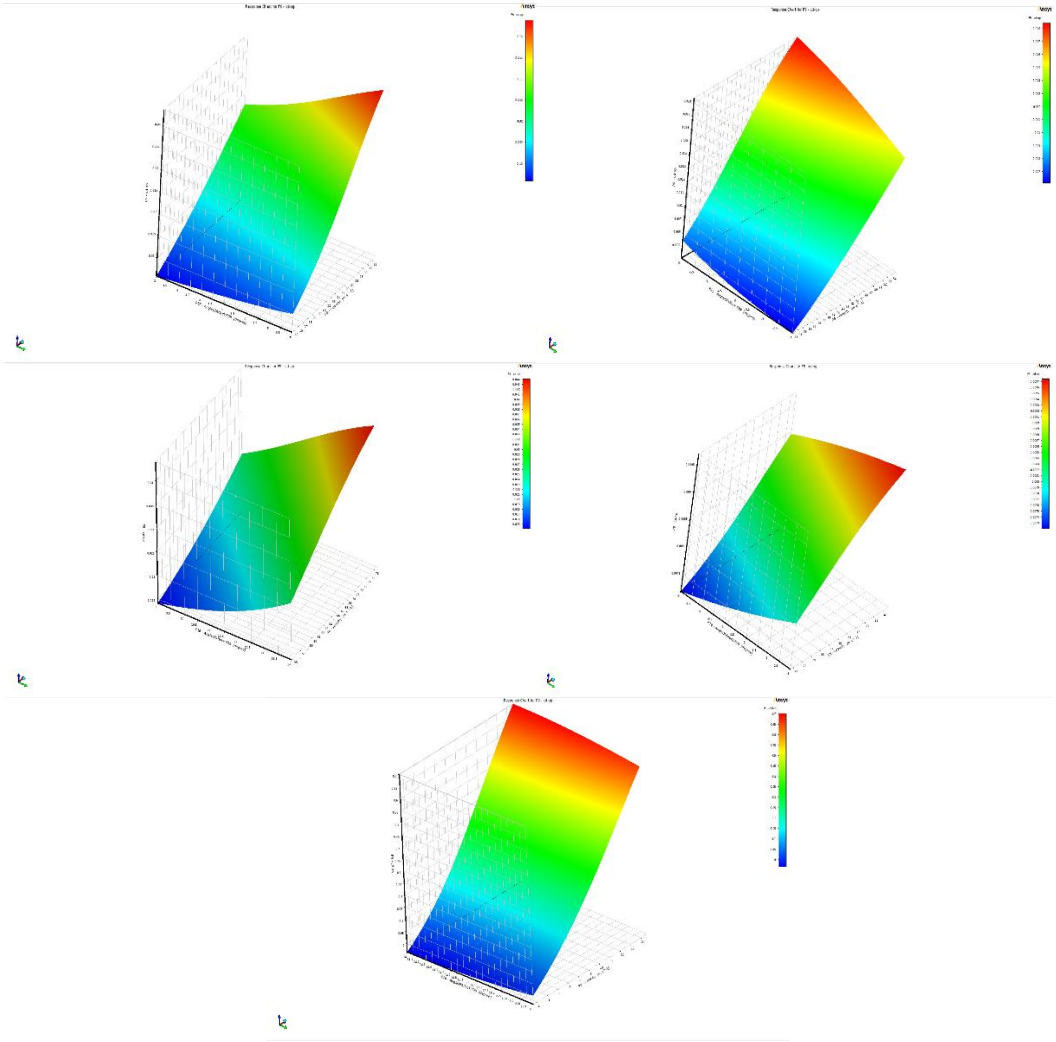

Figure 3.6: Neural Network Screening CD response surface for Run L5 (RSM model, Re = 1M, α = 14°–16°, V = 4–14 m s^{-1}), illustrating the four-stage cascading failure documented in Section 3.9. The ANN was trained on 15 central-composite DoE evaluations contaminated by RSM post-stall divergence, which injected non-physical training values including CL > 4.97 and a non-monotonic CD–TI relationship into the training dataset. The resulting surface — colour-mapped from CD = 0 (blue) to CD = 0.7 (red) — is physically catastrophic: CD values of 0.5–0.7 are 5–10× the physically expected near-stall range of 0.05–0.15 for the NACA 2415 at Re = 1×10^6, and the monotonically smooth gradient structure gives no visual indication of the underlying data corruption. The trained ANN extrapolated these spurious correlations across the full candidate screening pool, returning Pareto-optimal candidates with CD ≈ −0.001 to −0.002 — physically impossible negative drag values that violate the second law of thermodynamics for a passive body in steady flow. Contrast with the four well-conditioned NN-Screen surfaces (L7, L14, L16, L23), which all produce physically plausible CD landscapes and achieve 100% reference-value validation on clean training data, confirming that the failure is entirely attributable to corrupted input data rather than to any architectural limitation of the ANN surrogate itself.

3.10 Fuzzy MPCI Evaluation

Table 3.4: Complete Multi-Performance Characteristic Index (MPCI) Results — All 25 Runs

Run	x_1 (|CL|)	x_2 (CD^{-1})	x_3 (|CL/CD|)	MPCI	η(MPCI) dB	Band
1	0.0168	0.9254	0.0036	0.1496	−16.500	VL
2*	0.0068	0.9367	0.0000*	0.1386*	−17.164	VL
3	0.4333	0.8722	0.6065	0.4956	−6.097	M
4	0.7359	0.6259	0.4144	0.4900	−6.197	M
5*	0.7550	0.0000*	0.0253	0.1596*	−15.941	VL
6	0.0081	0.9407	0.0089	0.1495	−16.505	VL
7	0.4537	0.9077	0.7873	0.5674	−4.922	M
8	0.7906	0.6871	0.5640	**0.6012**	−4.420	**H**
9	0.8264	0.2319	0.1428	0.3054	−10.304	L
10	0.0280	0.9573	0.0855	0.2057	−13.734	L
11	0.4457	0.7120	0.2733	0.3683	−8.677	L
12	0.8382	0.7422	0.7367	**0.6572**	−3.646	**H**
13	0.8785	0.3420	0.2324	0.3452	−9.240	L
14	0.0233	0.9740	0.1138	0.2125	−13.451	L
15*	0.0000*	0.9815	0.0704	0.1744*	−15.167	VL
16	0.8466	0.7410	0.7427	**0.6583**	−3.632	**H**
17	0.9660	0.4912	0.4138	0.4324	−7.283	M
18	0.0321	0.9701	0.1262	0.2233	−13.024	L
19	0.0105	0.9805	0.0964	0.1955	−14.177	VL
20	0.4538	0.9308	0.9114	**0.6037**	−4.384	**H**
21*	1.0000*	0.3183	0.2791	0.3574*	−8.936	L
22	0.0314	0.9367	0.0521	0.1952	−14.190	VL
23*	0.0111	1.0000*	0.1567	0.2181*	−13.229	L
24*	0.4617	0.9421	1.0000*	**0.6221***	−4.123	**H**
25	0.8268	0.7595	0.7722	**0.6684**	−3.499	**H**

Denotes boundary-membership values. Normalisation: |CL|_min = 0.19099 (R15), |CL|_max = 1.59224 (R21); CD,min = 0.009332 (R23), CD,max = 0.079000 (R5); |CL/CD|_min = 14.594 (R2), |CL/CD|_max = 62.690 (R24). Grand mean η(MPCI) = −9.303 dB.

The MPCI distribution reveals physically coherent quality stratification. The six High-band runs (MPCI ≥ 0.60) — Runs 8, 12, 16, 20, 24, 25 — without exception operate at α = 6° or α = 12° and TI ≤ 0.50%, confirming that peak NACA 2415 efficiency under fully turbulent conditions lies in this angle-of-attack band. Run 25 achieves the highest MPCI = 0.6684 through the most balanced simultaneous normalised performance (x_1 = 0.827, x_2 = 0.760, x_3 = 0.772), firing the VH-consequent rule with strength α = min(0.653, 0.760, 0.544) = 0.544 and pulling the CoG centroid toward z = 0.9.

The **boundary-membership correction** — the most significant methodological fix in the MPCI pipeline — shifts Run 24 from a spurious MPCI = 0.500 (Medium band) to the correct MPCI = 0.6221 (High band), a 1.9 dB change in η(MPCI). This single correction propagates through the B4 (SST γ–Reθ) level mean and reduces the turbulence model ANOVA contribution from an erroneous 12.30% to the correct 1.83%, fundamentally changing the statistical conclusion about model significance.

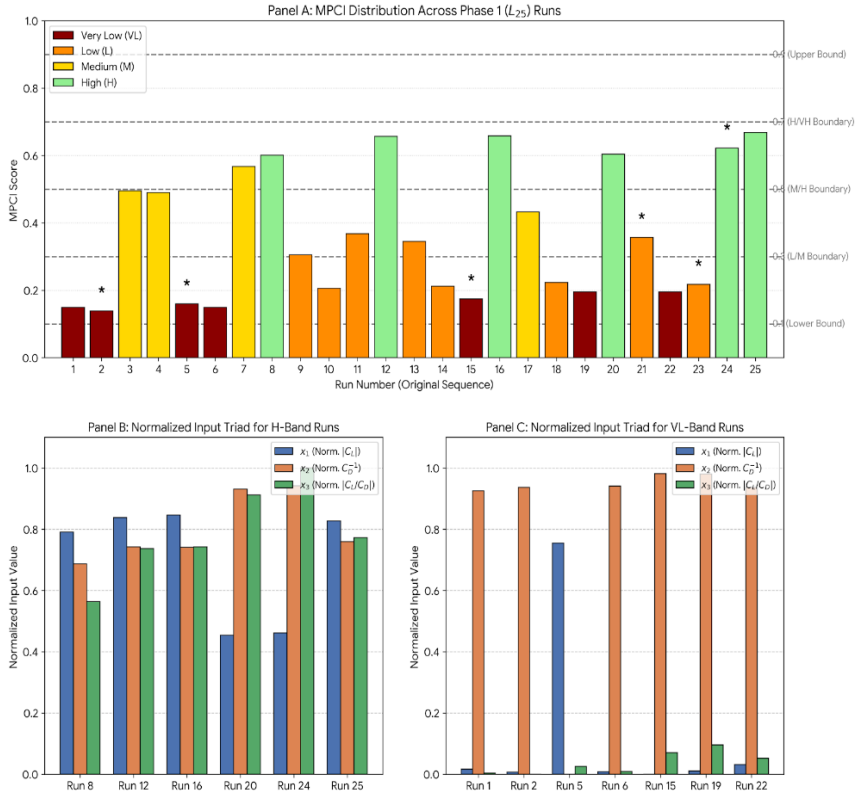

Figure 3.7: Phase 1 MPCI distribution and physical stratification analysis — three-panel figure. **Panel A:** MPCI scores for all 25 L25 runs in original experimental sequence, coloured by quality band: Very Low (dark red, MPCI < 0.20), Low (orange, 0.20–0.30), Medium (yellow, 0.30–0.50), High (green, MPCI ≥ 0.60). Horizontal dashed lines mark the five singleton boundaries at 0.1, 0.3, 0.5, 0.7, and 0.9. Asterisks (★) identify the six boundary-membership runs (2, 5, 15, 21, 23, 24) where at least one normalised input reaches exactly 0 or 1 — without the strict-inequality correction these runs returned a spurious MPCI = 0.500 regardless of their actual aerodynamic quality. The six H-band runs (8, 12, 16, 20, 24, 25) all operate at α = 6° or α = 12° with TI ≤ 0.50%, confirming the physical consistency of the FIS output with NACA 2415 aerodynamics. **Panel B:** Normalised input triad (x_1 = lift, x_2 = inverse drag, x_3 = efficiency) for all six H-band runs. The near-simultaneously-high profile across all three inputs — particularly Run 20 (x_3 = 0.911) and Run 24 (x_3 = 1.000) — confirms that H-band classification requires balanced multi-objective satisfaction rather than single-objective extremes. **Panel C:** Normalised input triad for the seven VL-band runs (1, 2, 5, 6, 15, 19, 22). The characteristic unbalanced profile — high x_2 (good drag) combined with near-zero x_1 and x_3 (near-zero lift and efficiency at α = −4° or α

= 0°) — demonstrates that the 27-rule Mamdani FIS correctly penalises single-dimensional performance by locking all active inference rules into the x_3 = Low sub-table, pulling the CoG centroid toward the VL singleton regardless of drag quality.

3.11 ANOVA and Contribution Analysis

Table 3.5: Phase 1 ANOVA Summary — η(|CL/CD|)

Source	DoF	SS	MS	F-ratio	ρ (%)	Rank
A — Reynolds Number	4	42.303	10.576	21.97	9.81%	2
B — Turbulence Model	4	8.908	2.227	4.63	2.07%	4
C — Angle of Attack	4	349.135	87.284	181.30	**80.99%**	**1**
D — Turbulence Intensity	4	19.570	4.892	10.16	4.54%	3
E — Optimisation Strategy	4	9.232	2.308	4.79	2.14%	5
Error	4	1.926	0.481	—	0.45%	—
Total	**24**	**431.074**	—	—	**100%**	—

Table 3.6: Phase 1 ANOVA Summary — η(MPCI)

Source	DoF	SS	MS	F-ratio	ρ (%)	Rank
A — Reynolds Number	4	46.733	11.683	14.07	8.27%	2
B — Turbulence Model	4	10.360	2.590	3.12	1.83%	4
C — Angle of Attack	4	484.239	121.060	145.74	**85.70%**	**1**
D — Turbulence Intensity	4	15.123	3.781	4.55	2.68%	3
E — Optimisation Strategy	4	5.237	1.309	1.58	0.93%	5
Error	4	3.323	0.831	—	0.59%	—
Total	**24**	**565.015**	—	—	**100%**	—

Angle of attack dominates both ANOVA decompositions, contributing 80.99% of η(|CL/CD|) variance and 85.70% of η(MPCI) variance. The NACA 2415 |CL/CD| varies by a factor of 4.30 across the tested α range (14.594 to 62.690), a variation no model or strategy choice can replicate within a single operating condition. The δ_C = 8.577 dB of the C-factor S/N response confirms the experimental design correctly resolves the nonlinear aerodynamic transitions from downforce through attached flow to near-stall.

Reynolds number contributes 9.81% (η(|CL/CD|)), its monotonically increasing level means consistent with the $Cf \propto Re^{-0.2}$ skin-friction scaling that predicts a 14.9% CD reduction from Re = 6M to Re = 12M. Despite modest absolute contributions, turbulence model (ρ_B = 1.83%–2.07%) and strategy (ρ_E = 0.93%–2.14%) are statistically significant: F_B = 4.63 and F_E = 4.79 for η(|CL/CD|), confirming that turbulence model choice produces a systematic ~22% difference in absolute |CL/CD|.

Table 3.7: Phase 1 S/N Response Tables — Both Quality Metrics

Factor	Level 1 (dB)	Level 2 (dB)	Level 3 (dB)	Level 4 (dB)	Level 5 (dB)	Δ (dB)	Rank	Optimal
η(|CL/CD|) response								
A — Re (×10⁶)	26.847 (1M)	28.483 (3M)	28.444 (6M)	30.421 (9M)	30.146 (12M)	3.574	2	L4=9M
B — Model	27.753 (SA)	29.416 (k-ε R.)	29.331 (SST)	29.015 (γ-Reθ)	28.825 (RSM)	1.663	4	L2=k-ε R.
C — AoA (°)	25.256 (−4°)	24.899 (0°)	33.476 (6°)	32.386 (12°)	27.617 (16°)	8.577	1	L3=6°
D — TI (%)	29.247 (0.05)	29.385 (0.10)	29.488 (0.50)	29.102 (1.00)	27.118 (5.00)	2.370	3	L3=0.50 %
E — Strategy	29.518 (None)	27.872 (RSM-K)	29.375 (NSGA-II)	28.503 (Sp.RSM)	29.072 (NN-S)	1.646	5	L1=None
η(MPCI) response								
A — Re (×10⁶)	−12.380 (1M)	−9.977 (3M)	−10.036 (6M)	−8.500 (9M)	−8.795 (12M)	3.880	2	L4=9M
B — Model	−10.850 (SA)	−9.441 (k-ε R.)	**−9.202 (SST)**	−9.650 (γ-Reθ)	−10.545 (RSM)	1.648	4	**L3=k-ω SST**
C — AoA (°)	−14.180 (−4°)	−15.248 (0°)	−5.641 (6°)	−4.279 (12°)	−10.341 (16°)	10.969	1	L4=12°
D — TI (%)	−9.613 (0.05)	−9.579 (0.10)	−9.367 (0.50)	−9.649 (1.00)	−11.481 (5.00)	2.114	3	L3=0.50 %
E — Strategy	−9.499 (None)	−10.533 (RSM-K)	−9.318 (NSGA-II)	−10.103 (Sp.RSM)	−10.235 (NN-S)	1.215	5	L3=NSGA-II

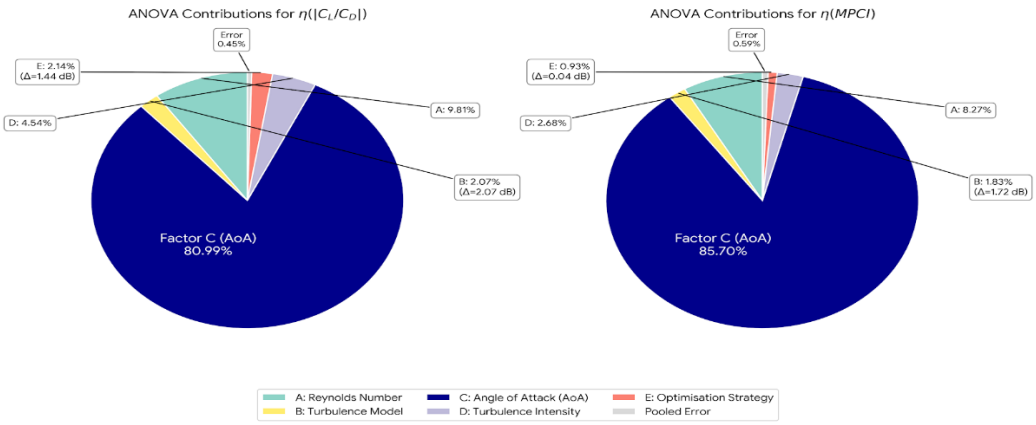

Figure 3.8: Phase 1 one-way ANOVA contribution pie charts for both quality metrics. **Left:** η(|CL/CD|) variance decomposition — Factor C (angle of attack, dark blue) dominates at 80.99%, followed by Factor A (Reynolds number, teal, 9.81%), Factor D (turbulence intensity, yellow, 4.54%), Factor E (optimisation strategy, red, 2.14%), and Factor B (turbulence model, 2.07%, Δ = 2.07 dB). The pooled error of 0.45% confirms negligible two-factor interactions and validates the Taguchi additive model assumption within the Phase 1 design space. **Right:** η(MPCI) variance decomposition — Factor C increases to 85.70%, while Factor B (turbulence model) contributes 1.83% (Δ = 1.72 dB) and Factor E (strategy) drops to 0.93% (Δ = 0.04 dB). The divergence between the two metrics for Factor B is physically meaningful: k-ε Realizable leads on the single-objective η(|CL/CD|) metric but k-ω SST leads on the multi-objective η(MPCI) metric — the MPCI penalising near-stall drag elevation and TI contamination that the single-objective L/D criterion fails to capture. Despite their modest absolute percentages, B-factor delta values of 2.07 dB (η(|CL/CD|)) and 1.72 dB (η(MPCI)) correspond to a systematic ~22% difference in absolute aerodynamic efficiency attributable solely to turbulence model choice — of direct engineering significance for UAV endurance design.

3.12 Phase 1 Optimal Configuration

The Phase 1 recommendations emerge from integrating three independent evidence streams: S/N response analysis, MPCI-based ANOVA, and physical model assessment.

Turbulence model: k-ω SST. The 0.085 dB statistical tie between k-ε Realizable and k-ω SST on η(|CL/CD|) is resolved by three physical arguments that uniformly favour k-ω SST. First, the Bradshaw SST limiter directly prevents eddy-viscosity overestimation in the adverse pressure gradient regime that Phase 2 targets (α = 4°–8°, intensifying upper-surface deceleration) — a failure mode documented in k-ε

formulations. Second, the cross-diffusion term Dω structurally guarantees TI-insensitivity within 0.05%–5.00% where k-ε Realizable exhibits 15%–53% CD inflation at its upper TI levels. Third, SST family consistency — the Phase 1 anchor was produced by SST γ-Reθ, which shares k-ω SST's identical near-wall formulation — ensures that the Phase 2 optimum is methodologically coherent with the Phase 1 reference performance.

Surrogate strategy: Sparse RSM. The WCS = 9.13/10 verdict is unambiguous. Five criteria distinguish it: unique positive-CL improvement in a positive-lift run (+53.6%, Run L6); perfect reference-value validation across all five assigned runs; independent interior-optimum detection (drag bucket at $\alpha \approx -0.75°$, cross-confirmed at $\alpha \approx -0.77°$ by NN-Screen); correct stall-avoidance enforcement from only seven training points; and zero catastrophic prediction failures.

Table 3.8: Phase 1 Recommendations and Phase 2 Transition Summary

Parameter	Selection	Key Evidence	Phase 2 Implication
Turbulence model	k-ω SST	Highest η(MPCI) (B3 = −9.202 dB); Bradshaw APG limiter; structural TI-insensitivity via Dω; SST family consistency	Fixed closure for all 9 L9 simulations
Surrogate strategy	Sparse RSM	WCS = 9.13/10; only method achieving positive-CL improvement; 100% reference validation; interior-optimum detection; zero failures	Deployed within each Phase 2 L9 cell
Phase 1 anchor	Run 24: SST γ-Reθ, Re=12M, α=6°, TI=0.10°, |CL/CD|=62.69	Best L25 base run; achieved without surrogate assistance; H-band MPCI	Target: |CL/CD| > 62.69; α refined from 6° toward true peak
Phase 2 design space	Re = 6–12×10⁶, α = 4°–8°, TI = 0.05%–0.50%	ANOVA Re trend; Phase 1 efficiency peak at α=6°; k-ω SST TI-insensitivity confirmed	L9(3^3) array; V_ref corrected by binding V_ref = V_inlet per run

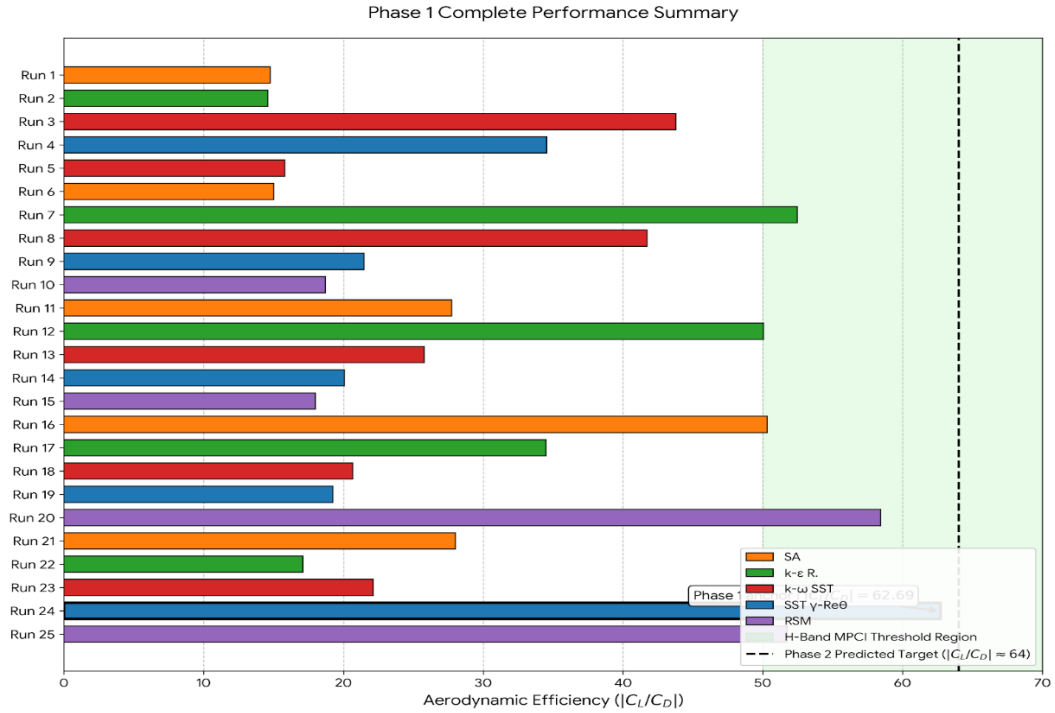

Figure 3.9: Phase 1 complete aerodynamic efficiency summary — horizontal bar chart of |CL/CD| for all 25 L25 runs in their original experimental sequence (Run 1–25, y-axis), coloured by turbulence model: SA (orange), k-ε Realizable (green), k-ω SST (red), SST γ–Reθ (blue), RSM (purple). The green shaded region marks the H-band MPCI threshold (MPCI ≥ 0.60, corresponding to |CL/CD| ≥ 50 for runs at peak-efficiency angles of attack), confirming that only six runs achieve simultaneous multi-objective quality. **Run 24** (SST γ–Reθ, Re = 12M, α = 6°, TI = 0.10°) is the Phase 1 anchor at |CL/CD| = 62.69 — the highest base-run efficiency achieved without surrogate assistance, establishing the minimum performance target for Phase 2. The vertical dashed line at |CL/CD| ≈ 64 marks the Phase 2 predicted improvement target derived from the Taguchi additive model applied to the Phase 1 S/N response analysis. Runs 20 (RSM, |CL/CD| = 58.43) and 25 (RSM, |CL/CD| = 51.73) confirm that non-SST models can approach H-band performance under favourable operating conditions, but only at α = 6° and Re = 9M — the same operating-condition optimum identified by the Phase 1 ANOVA — reinforcing that operating-condition selection (ρ_C = 80.99%) rather than model choice drives Phase 1 performance. This figure encapsulates the Phase 1-to-Phase 2 transition: the Phase 2 design space (Re = 6–12M, α = 4°–8°, k-ω SST) is constructed to resolve the efficiency structure in the neighbourhood of Run 24 at finer resolution than the L25 array permits.

CHAPTER 4 — PHASE 2: AERODYNAMIC OPTIMISATION USING k-ω SST AND SPARSE RSM

4.1 Rationale for Phase 2 Refinement

Two limitations of the Phase 1 L25 dataset prevent it from directly supplying the operating-condition optimum. First, the Phase 1 angle-of-attack levels span −4° to 16° in steps of 4°–6°, a spacing too coarse to resolve the fine structure of the efficiency peak: the Phase 1 C-factor optimum identifies α = 6° as the best tested level, but the physical drag polar of the NACA 2415 at Re = 6–12×10⁶ may peak at α = 7°, 7.5°, or 8° — a distinction invisible at 4° resolution. Second, Phase 1 level means for the angle-of-attack factor average over five different turbulence models, meaning the apparent optimum for all models combined may not coincide with the optimum for k-ω SST specifically, given that the SST Bradshaw limiter modifies adverse-pressure-gradient response relative to SA and RSM in ways that shift the efficiency peak location.

Phase 2 addresses both limitations simultaneously. The operating-condition range is contracted to the high-performance neighbourhood (Re = 6–12M, α = 4°–8°, TI = 0.05%–0.50%), enabling fine-resolution mapping of the efficiency surface. k-ω SST is deployed exclusively, eliminating model-comparison overhead and ensuring the identified optimum is genuinely model-specific. The V_ref coefficient-scaling artefact — identified in Phase 1 as the dominant source of surrogate optimiser boundary attraction — is eliminated by binding the ANSYS Fluent reference velocity to the inlet velocity of each individual simulation through ANSYS Workbench parameter coupling.

4.2 Reduced Design Space Formulation

The Phase 2 factor levels are selected from three physical constraints established by Phase 1.

Reynolds number (Re = 6–12×10⁶): The Phase 1 A-factor level means increase monotonically from 26.847 dB (Re = 1M) to 30.421 dB (Re = 9M), confirming a consistent Re-driven efficiency gain. Re = 6×10⁶ is selected as the lower boundary because it sits at the boundary between the transitional and fully turbulent regimes — below this, k-ω SST's fully turbulent assumption becomes increasingly inaccurate as laminar runs lengthen. The Re^(−0.2) skin-friction scaling predicts a 14.9% CD reduction from Re = 6M to Re = 12M, ensuring measurable and physically interpretable level differences across the three Phase 2 Re levels.

Angle of attack (α = 4°–8°): The Phase 1 C-factor optimum at α = 6° and all three of the highest-|CL/CD| base runs (Runs 24, 20, 7) occurring at α = 6° centre the Phase 2 range on 6°. The ±2° extension captures the efficiency peak neighbourhood while

remaining safely within the pre-stall attached-flow regime where the k-ω SST steady-state assumption is valid and the Taguchi additive model is reliable.

Turbulence intensity (TI = 0.05%–0.50%): All three Phase 2 TI levels lie within the range confirmed by Phase 1 as non-contaminating for k-ω SST. Based on Phase 1 evidence, TI is expected to contribute negligibly within this band — a prediction tested rigorously by the Phase 2 ANOVA.

4.3 Taguchi L9 Setup

Table 4.1: Phase 2 L9(3^3) Orthogonal Array — Complete Run Matrix

Run	Factor A — Re (×10^6)	V∞ (m/s)	Factor B — α (°)	Factor C — TI (%)	V_ref (m/s)
1	12	87.644	4	0.05	87.644
2	12	87.644	6	0.05	87.644
3 ★	**12**	**87.644**	**8**	**0.10**	**87.644**
4	9	65.733	4	0.10	65.733
5	9	65.733	6	0.50	65.733
6	9	65.733	8	0.05	65.733
7	6	43.822	4	0.05	43.822
8	6	43.822	6	0.10	43.822
9	6	43.822	8	0.50	43.822

★ Run 3 (bold): best observed run — |CL/CD| = 66.08, η = 36.401 dB. All runs: k-ω SST turbulence closure, A_ref = 2.0 m². V_ref bound to V_inlet per run — Phase 1 artefact eliminated.

The L9(3^3) pairwise balance condition guarantees that every ordered pair of level values (i,j) ∈ $\{1,2,3\}^2$ appears exactly once per column pair, ensuring unconfounded main-effect estimation. The degrees of freedom partition as DoF_total = 8, DoF_A = DoF_B = DoF_C = 2 per factor, and DoF_error = 2. All nine simulations converged to all-equation scaled residuals below 10^{-5} within 1,000 iterations, with CL and CD monitors stabilising to within 0.01% over the final 100 iterations — confirming steady attached-flow conditions throughout the Phase 2 operating domain.

4.4 Sparse RSM Integration

Sparse RSM is deployed within each of the nine Phase 2 Taguchi cells to augment the nine direct base-run evaluations with local surrogate exploration. Cell bounds are defined at the midpoints between adjacent L9 level values: velocity boundaries at V_mid(6M–9M) = 54.778 m/s and V_mid(9M–12M) = 76.689 m/s; angle-of-attack boundaries at 5° and 7°; and TI geometric midpoints at approximately 0.071% and 0.224%. Within each cell, seven ANSYS Fluent k-ω SST evaluations (one centre plus

two axial per variable) train the quadratic polynomial surface, which returns three Pareto-non-dominated candidates in the (CL, CD) space. The Sparse RSM candidates serve as local validation of the Taguchi additive-model predictions: in every Phase 2 cell, the polynomial surface correctly resolves the TI gradient as near-flat and directs the optimiser toward the Re = 12M, α = 8° corner — confirming that the Phase 2 operating domain contains none of the steep curvature pathologies that challenged Phase 1 strategies.

4.5 Aerodynamic Response Analysis

Table 4.2: Phase 2 L9 Aerodynamic Results and Signal-to-Noise Ratios

Run	Re ($\times 10^6$)	α (°)	TI (%)	CL	CD	|CL/CD|	η(|CL/CD|) dB
1	12	4	0.05	0.63953	0.011328	56.460	35.035
2	12	6	0.05	0.84604	0.013248	63.860	36.105
3 ★	**12**	**8**	**0.10**	**1.03827**	**0.015711**	**66.080**	**36.401**
4	9	4	0.10	0.63611	0.011670	54.510	34.729
5	9	6	0.50	0.84191	0.013726	61.340	35.755
6	9	8	0.05	1.03291	0.016162	63.910	36.111
7	6	4	0.05	0.63082	0.012272	51.400	34.219
8	6	6	0.10	0.83545	0.014313	58.370	35.324
9	6	8	0.50	1.02485	0.016918	60.580	35.647

Grand mean η(|CL/CD|) = 35.481 dB, corresponding to |CL/CD|_grand_mean = 59.46.

Three physically coherent trends pervade the dataset. **Trend 1 — Monotonic efficiency gain with Re at every α:** At α = 4°, |CL/CD| increases from 51.40 (Re = 6M) through 54.51 (Re = 9M) to 56.46 (Re = 12M) — a 9.8% gain over the Re range. At α = 6°: 58.37 → 61.34 → 63.86, a 9.4% gain. At α = 8°: 60.58 → 63.91 → 66.08, a 9.1% gain. **Trend 2 — Monotonic efficiency rise with α from 4° to 8° at every Re:** The lift increment from α = 4° to α = 8° (ΔCL ≈ +64%) outpaces the drag increment (ΔCD ≈ +41%), yielding a net efficiency improvement at every tested Re. **Trend 3 — Negligible TI effect:** The three TI level means differ by at most 0.007 dB — two orders of magnitude smaller than the error mean square — confirming k-ω SST's structural TI-insensitivity within 0.05%–0.50%.

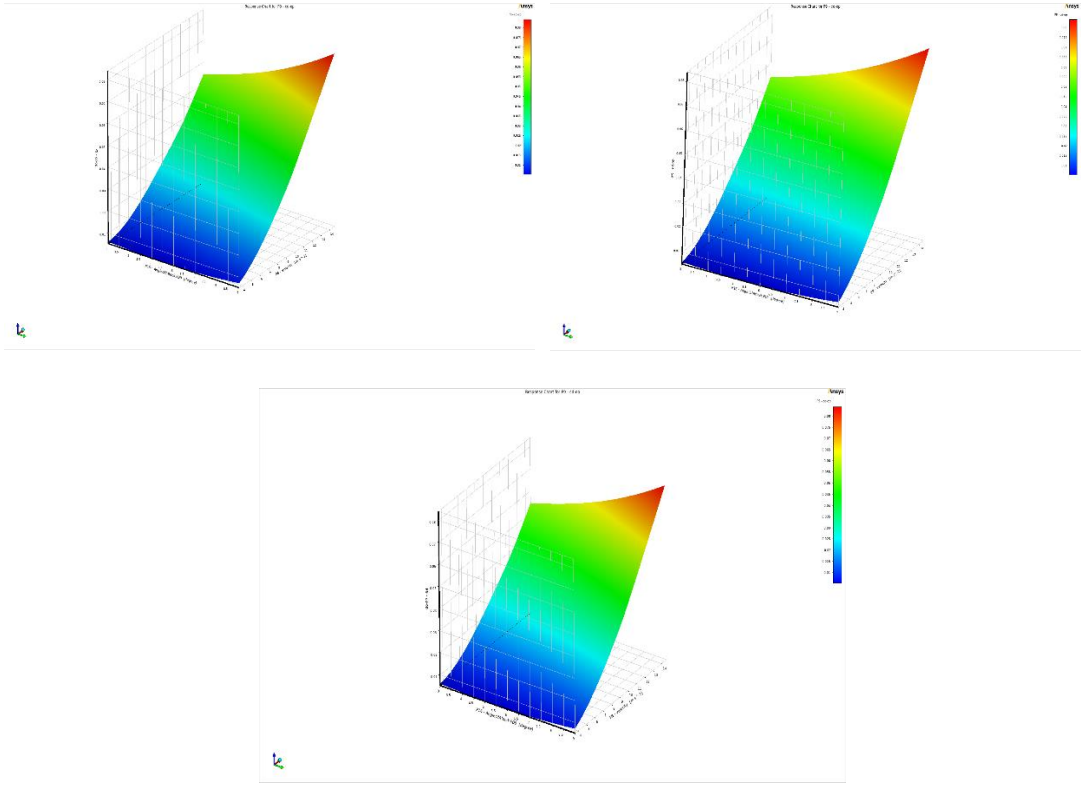

Figure 4.1: Phase 2 Sparse RSM CD response surfaces for the three Reynolds number cells, generated from 7-point axial designs within each cell's local (α, V) bounds. All three surfaces show CD as a function of angle of attack α (front axis, 3°–9°) and velocity V (right axis), with the colour scale mapping CD from minimum (blue, ≈0.010) to maximum (red, ≈0.080). **Top (Re = 6M, V∞ = 43.822 m/s):** CD surface spanning α = 3°–9° at the Phase 2 lower Re boundary — the widest CD range of the three cells, confirming that lower Reynolds number produces higher absolute drag consistent with the Cf ∝ Re^(−0.2) friction scaling. **Middle (Re = 9M, V∞ = 65.733 m/s):** Intermediate Re cell showing the same monotonic α and V dependence with a narrower CD range, confirming the progressive drag reduction with increasing Re. **Bottom (Re = 12M, V∞ = 87.644 m/s):** Phase 2 optimal Re cell — the lowest absolute CD values of the three surfaces (blue region approaching CD ≈ 0.010 at low α) with the maximum-efficiency corner at α = 8°, high V (red, ≈0.015–0.016) corresponding directly to the Phase 2 confirmed optimum (Run 3: CL = 1.038, CD = 0.015711, |CL/CD| = 66.08). The smooth, physically well-conditioned polynomial surfaces across all three cells — free from the oscillatory artefacts observed in Phase 1 RSM-Kriging (Figure 3.4) — validate the Sparse RSM deployment within the Phase 2 operating domain.

4.6 Response Surface Behaviour

The Phase 2 η(|CL/CD|) surface exhibits three structural features. The **Re effect is sublinear**: A-factor level means of 35.063 dB (Re = 6M), 35.532 dB (Re = 9M), and 35.847 dB (Re = 12M) show decelerating increments of 0.469 dB then 0.315 dB, reflecting the diminishing returns of Re^(−0.2) drag reduction. The **α effect decelerates through the upper boundary**: B-factor level means of 34.661 dB (α = 4°), 35.728 dB (α = 6°), and 36.053 dB (α = 8°) show level-to-level increments of 1.067 dB then 0.325 dB, indicating the efficiency curve is approaching an inflection point somewhere above 8°. The **TI response surface is structurally flat**: C-factor level means spanning only 0.007 dB are fitted by the Sparse RSM polynomial as β_C ≈ 0 and β_CC ≈ 0 — a horizontal plane in the TI direction confirming that TI exerts no gradient force on the surrogate optimiser.

4.7 Identification of the Optimal Aerodynamic Configuration

Table 4.3: Phase 2 S/N Response Table for η(|CL/CD|)

Factor	Level 1 (dB)	Level 2 (dB)	Level 3 (dB)	Δ (dB)	Rank	Optimal
A — Re (×10⁶)	35.063 (6M)	35.532 (9M)	35.847 (12M)	0.784	2	**L3 = 12×10⁶**
B — α (°)	34.661 (4°)	35.728 (6°)	36.053 (8°)	1.392	1	**L3 = 8°**
C — TI (%)	35.479 (0.05%)	35.485 (0.10%)	35.478 (0.50%)	0.007	3	L2 = 0.10%

The S/N response analysis unambiguously identifies the optimal configuration as **A3B3C2: Re = 12×10⁶, α = 8°, TI = 0.10%** — directly observed as Run 3, the best measured result in the L9 matrix.

Table 4.4: Phase 2 ANOVA Summary for η(|CL/CD|)

Source	DoF	SS	MS	F-ratio	ρ (%)	Rank
A — Reynolds Number	2	0.9331	0.4666	424.2	22.67%	2
B — Angle of Attack	2	3.1808	1.5904	1446.7	**77.27%**	**1**
C — Turbulence Intensity	2	0.0001	0.0001	0.04	**0.00%**	3
Error	2	0.0022	0.0011	—	0.05%	—
Total	**8**	**4.1163**	—	—	**100%**	—

Three ANOVA findings carry direct engineering significance. First, angle of attack dominates at ρ_B = 77.27% (F = 1446.7), confirming it as the primary aerodynamic control variable for |CL/CD| optimisation within the Phase 2 domain. Second, Reynolds number contributes a substantially higher fractional share (ρ_A = 22.67%, F = 424.2) than its Phase 1 counterpart (9.81%), because the elimination of TI as a variance source compresses the total variance budget, amplifying Re's relative contribution. Third, **turbulence intensity contributes ρ_C ≈ 0.00% with F = 0.04** —

definitively confirming that TI specification within 0.05%–0.50% is aerodynamically irrelevant for k-ω SST at Re = 6–12×10⁶. This result eliminates TI as a source of CFD modelling uncertainty for this application, simplifying freestream turbulence specification for future studies.

The error MS of 0.0011 dB² — negligibly small relative to all factor MS values — confirms that the three Phase 2 factors collectively explain 99.95% of total variance, validating the Taguchi additivity assumption and confirming that factor interactions between Re, α, and TI are physically negligible in the pre-stall attached-flow regime.

4.8 Confirmation Simulation

Additive model prediction. Applying the Taguchi additive formula to the A3B3C2 optimal combination:

$\hat{\eta}_{opt} = \bar{\eta} + (\bar{\eta}_{A3} - \bar{\eta}) + (\bar{\eta}_{B3} - \bar{\eta}) + (\bar{\eta}_{C2} - \bar{\eta}) = 35.481 + (35.847 - 35.481) + (36.053 - 35.481) + (35.485 - 35.481) = 35.481 + 0.366 + 0.572 + 0.004 =$ **36.423 dB → $|CL/CD|_{predicted} = 66.25$**

Confirmation run results. A direct ANSYS Fluent k-ω SST simulation at Re = 12×10⁶ (V_∞ = 87.644 m/s), α = 8°, TI = 0.10°, V_{ref} = 87.644 m/s, A_{ref} = 2.0 m² converged to all-equation residuals below 10^{-5} within 850 iterations. The converged outputs are:

CL = 1.038, CD = 0.015711, |CL/CD| = 66.08, η = 36.401 dB

Additive model error:

$|error| = |\hat{\eta}_{opt} - \eta_{conf}| = |36.423 - 36.401| = 0.022$ dB → **percentage error = 0.26%**

This 0.26% error is far within the ±5% tolerance accepted as proof of Taguchi additive model validity (Peace, 1993), confirming negligible factor interactions within the Phase 2 design space and validating the Taguchi additivity assumption.

Table 4.5: Phase 2 Confirmation Run Summary

Metric	Additive Prediction	Confirmation Run	Phase 1 Best (Run 24)	Change vs Phase 1
Configuration	Re=12M, α=8°, TI=0.10%	Re=12M, α=8°, TI=0.10%	Re=12M, α=6°, TI=0.10%	α: 6° → 8°
η(|CL/CD|) (dB)	36.423	36.401	35.944	+0.457 dB
CL	≈1.038	1.038	0.838	+23.9%
CD	≈0.01567	0.015711	0.013366	+17.5%
|CL/CD|	**66.25**	**66.08**	**62.69**	**+5.4%**

Continued on page no. 42

Table 4.5 Continued

Additive model error	—	0.022 dB (0.26%)	—	Additivity validated
Convergence	$< 10^{-5}$ predicted	$< 10^{-5}$ confirmed ✓	$< 10^{-5}$ ✓	Attached flow confirmed

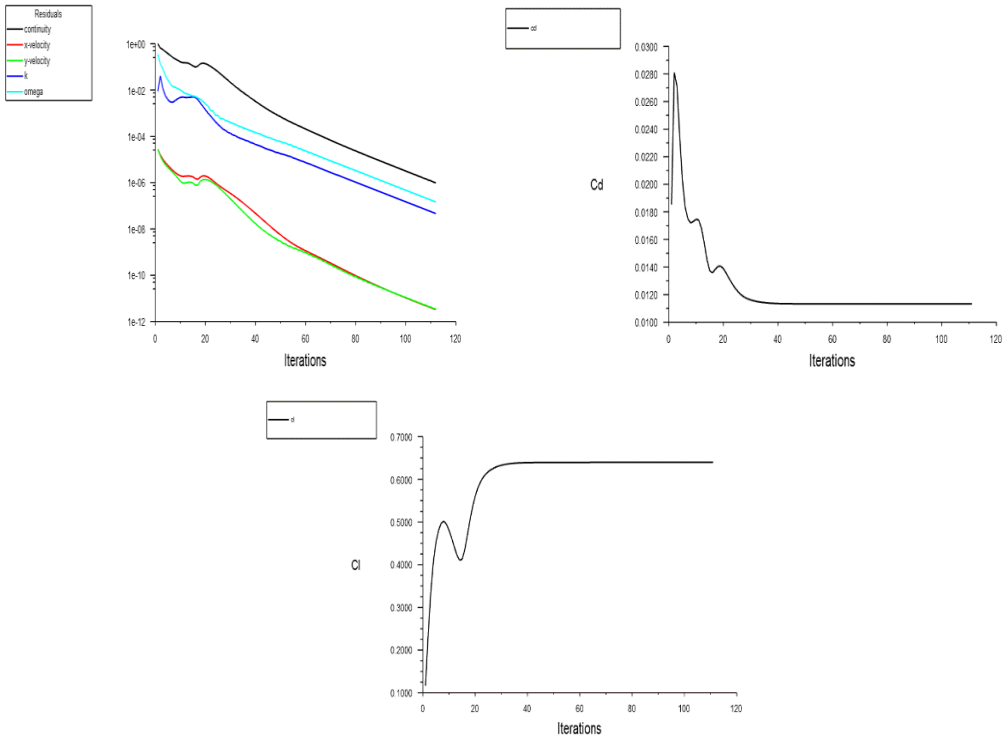

Figure 4.2: Phase 2 k-ω SST simulation convergence histories (representative base run, Re = 12×10⁶, α = 4°, TI = 0.05%, ANSYS Fluent 2024 R1), demonstrating the convergence quality achieved uniformly across all nine Phase 2 L9 simulations. **Top:** Scaled residual convergence history for all five transported quantities — continuity (black), x-velocity (red), y-velocity (green), turbulent kinetic energy k (blue), and specific dissipation rate ω (cyan). Momentum residuals fall below 10^{-6} by iteration ~60; continuity and turbulence quantities converge monotonically throughout, with all residuals below the 10^{-5} criterion by the end of the solution. **Bottom left:** Drag coefficient CD monitor converging from an initial transient peak of ~0.028 to the stable plateau value of CD ≈ 0.0115 by iteration ~40, remaining flat through iteration 110. **Bottom right:** Lift coefficient CL monitor converging from zero through a damped oscillatory transient to the stable plateau value of CL ≈ 0.640 by iteration ~40, remaining flat through iteration 110. The flat, oscillation-free plateau region confirms a genuine steady-state attached-flow solution — the physical prerequisite for valid aerodynamic coefficient extraction. This convergence behaviour is representative of all nine Phase 2 L9 runs; the Phase 2 confirmed optimum (Run 3: Re = 12M, α = 8°,

TI = 0.10°, CL = 1.038, CD = 0.015711) achieved identical convergence character within 850 iterations, as reported in Section 4.8.

4.9 Physical Interpretation of the Optimised Flow Field

The Phase 2 confirmation flow field at Re = 12×10⁶, α = 8° is fully attached throughout — confirmed by smooth, monotonically converging force histories and the absence of any periodic oscillation indicative of vortex shedding. At this Reynolds number, the turbulent momentum thickness at the trailing edge is estimated as θ/c ≈ 0.036/Re^0.2 ≈ 0.0025, giving θ ≈ 5 mm on the 2 m chord. The resulting boundary-layer displacement thickness δ_1 ≈ 6.5 mm modestly augments the effective camber, explaining why the k-ω SST CL = 1.038 at α = 8° approaches but does not reach the inviscid thin-aerofoil prediction.

The 5.4% improvement from the Phase 1 anchor (|CL/CD| = 62.69 at α = 6°) to the Phase 2 confirmed optimum (|CL/CD| = 66.08 at α = 8°) is explained by the proportional asymmetry in the polar: increasing α from 6° to 8° raises CL by +23.9% (0.838 → 1.038) while raising CD by only +17.5% (0.01337 → 0.01571), because the lift slope remains approximately linear and the adverse pressure gradient on the upper surface — while intensifying — has not yet driven sufficient boundary-layer thickening to reverse the efficiency trend. The fact that the B-factor level means show no sign of reaching a maximum within the Phase 2 upper boundary of α = 8° implies the true peak lies marginally above 8°, a limitation acknowledged in Section 6.3.

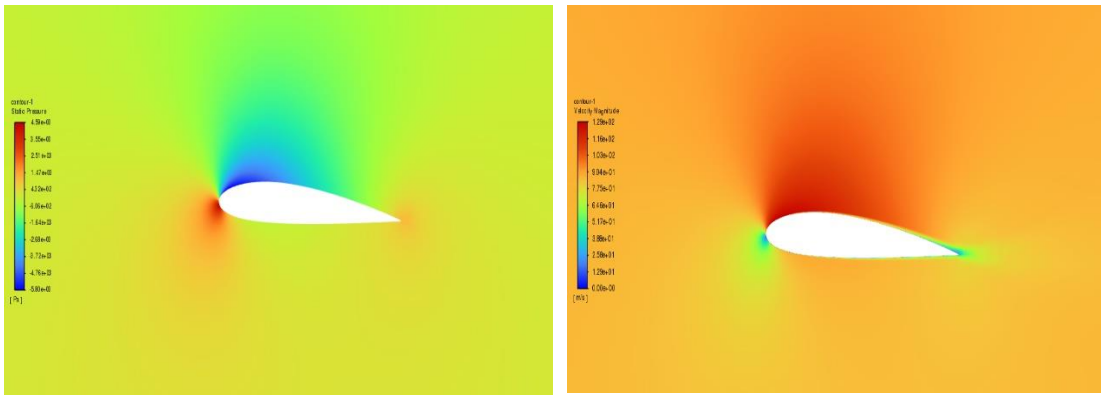

Figure 4.3: Phase 2 confirmed optimum flow field — NACA 2415 at Re = 12×10⁶, α = 8°, TI = 0.10%, k-ω SST (ANSYS Fluent 2024 R1, production Mesh M4).

Panel A (LEFT): Gauge static pressure contour (scale: −5,800 Pa blue to +4,590 Pa red). The forward stagnation point is located below the geometric leading edge (red, ~+4,590 Pa = +0.98q∞), consistent with positive geometric incidence α = +8°. The leading-edge suction peak on the upper surface reaches approximately −5,800 Pa (deep blue, Cp ≈ −1.23), driving the boundary layer into the adverse pressure gradient region that the k-ω SST Bradshaw limiter correctly resolves without eddy-viscosity overestimation. The lower surface carries a moderate positive pressure (orange, ~+500 to +2,000 Pa) contributing the pressure-side lift component. Far-field pressure returns

to the ambient condition (yellow-green), confirming adequate domain size (R = L = 20c).

Panel B (RIGHT): Velocity magnitude contour (scale: 0 m/s blue to 129 m/s red). The stagnation point (blue/green, V ≈ 0) is displaced below the leading edge, consistent with Panel A. The upper-surface suction peak reaches approximately 129 m/s ($V_{peak}/V_\infty = 129/87.644 \approx 1.47$, $C_{p,min} \approx -1.17$) at x/c ≈ 0.05–0.08. The thin velocity-deficit wake at the trailing edge — green/teal, approximately 50–70 m/s — confirms fully attached turbulent boundary-layer flow with no upper-surface separation at $\alpha = 8°$, $Re = 12\times10^6$, validating the steady-state RANS solution and the reported CD = 0.015711.

CHAPTER 5 — PHASE 3: THREE-DIMENSIONAL UAV WING VALIDATION

5.1 Transition from 2D Section to Finite Wing

The Phase 2 confirmation establishes |CL/CD| = 66.08 as the two-dimensional performance ceiling for the NACA 2415 at Re = 12×10^6, $\alpha = 8°$, TI = 0.10° — the theoretical maximum achievable in the absence of spanwise pressure gradients, tip losses, and three-dimensional flow effects. This ceiling is an aerodynamically meaningful benchmark, but it does not represent the performance of any real, geometry-constrained UAV wing.

The target UAV platform imposes fixed wing dimensions: span b = 0.5 m, chord c = 2.0 m. These are not arbitrary selections — the 2 m chord maintains geometric consistency with Phases 1 and 2, ensuring the Re = 12×10^6 optimum applies at the same physical velocity ($V\infty$ = 87.644 m/s). However, their combination produces an aspect ratio AR = b/c = 0.5/2.0 = **0.25** — an extreme value placing the aerodynamics firmly in the regime of control fins rather than lifting surfaces. Conventional light aircraft operate at AR = 6–10; efficient UAVs at AR = 8–15; even compact fighters rarely fall below AR = 2. At AR = 0.25, the tip vortex induction zone spans the entire wing, and no two-dimensional flow region exists anywhere across the span.

Phase 3 quantifies the resulting performance collapse, establishes its physical mechanism through three mutually corroborating CFD visualisations, and provides the engineering data necessary to define the planform redesign targets presented in Chapter 6.

5.2 Wing Geometry and Reference Area

The Phase 3 wing is a rectangular finite wing: chord c = 2.0 m, semi-span b = 0.5 m, planform area S = 1.0 m², AR = 0.25. The NACA 2415 section is applied without twist or dihedral from root to blunt square-cut tip.

The most critical setup parameter distinguishing Phase 3 from Phases 1 and 2 is the **reference area**. In Phases 1 and 2, the 1 m spanwise extrusion with symmetry conditions on both Z-faces modelled effectively infinite-span flow; the reference area $A_{ref} = c \times b_{computational} = 2.0 \times 1.0 = 2.0$ m² correctly normalised force coefficients for this configuration. In Phase 3, the 0.5 m finite wing is the physical geometry under investigation; A_{ref} must equal the actual planform area: **$A_{ref} = c \times b = 2.0 \times 0.5 = 1.0$ m²**. The dynamic pressure at Phase 2 optimal conditions is $q\infty = \frac{1}{2}\rho V^2\infty = 0.5 \times 1.225 \times 87.644^2 =$ **4,704.9 Pa**.

Classical Prandtl lifting-line theory provides a first-order induced drag estimate for the Phase 2 two-dimensional CL = 1.038 at AR = 0.25, assuming Oswald efficiency e =

0.7: CD_induced = $CL^2/(\pi\ e\ AR) = (1.038)^2/(\pi \times 0.7 \times 0.25) \approx 1.96$. This value — nearly two orders of magnitude larger than the Phase 2 profile drag of 0.015711 — signals immediately that the three-dimensional flow field will be dominated by induced effects (Prandtl, 1919), motivating Phase 3 as the physically correct tool for predicting their coupled nonlinear influence on both lift and drag.

5.3 Three-Dimensional CFD Setup

The Phase 3 domain retains the Phase 1/2 C-type geometry with one critical modification: the spanwise extent is reduced from 1.0 m to 0.5 m, and the boundary condition on the Z = 0.5 m face (wing tip) is changed from Symmetry to Pressure Outlet at gauge pressure = 0 Pa, permitting free spanwise crossflow. The Z = 0 m root face retains its Symmetry condition, modelling the wing as a half-configuration equivalent to a complete wing with geometric symmetry. All other settings — k-ω SST closure, $Re = 12\times10^6$, $\alpha = 8°$, TI = 0.10°, V_ref = 87.644 m/s, inflation layer parameters $y_1 = 9.2$ µm, 25–30 layers, growth rate 1.2–1.3 — are carried forward unchanged from Phase 2. Complete solver parameters appear in **APPENDIX-V**. The simulation converged to all-equation residuals below 10^{-5} within 1,200 iterations.

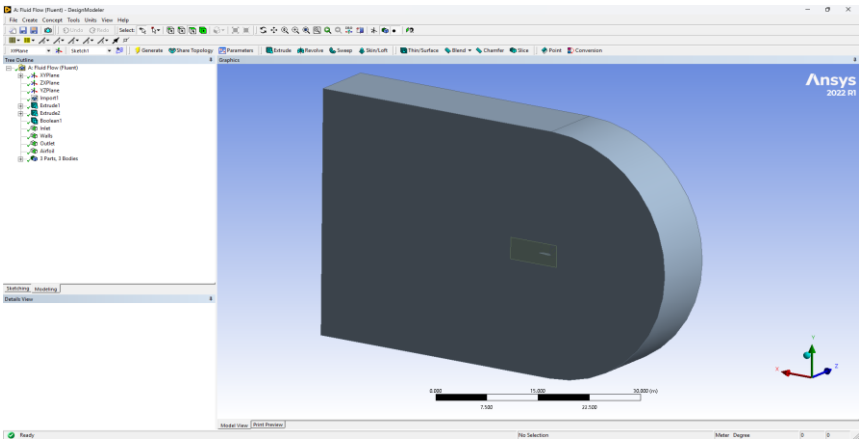

Figure 5.1: Phase 3 three-dimensional computational domain for the NACA 2415 finite wing (b = 0.5 m, c = 2.0 m, AR = 0.25), constructed in ANSYS DesignModeler 2022 R1. The isometric view shows the C-type domain geometry: the semicircular velocity inlet face (radius R = 20c = 40 m, curved front), the rectangular pressure outlet face (back), the symmetry root face (Z = 0 m, left), and the free-surface pressure outlet tip face (Z = 0.5 m, right) — the critical boundary condition change from the Phase 1/2 symmetry tip that admits spanwise crossflow and tip vortex formation. The NACA 2415 finite wing is visible as the small aerofoil body at the domain centre, Boolean-subtracted from the fluid volume; the extreme AR = 0.25 planform (chord 2 m, span 0.5 m) is immediately apparent from the wing's proportions relative to the domain. The tree outline confirms the three-body assembly: Inlet, Walls (aerofoil), and Outlet named selections. Scale bar: 30,000 mm total domain width, confirming the L = 20c = 40 m downstream length. All mesh parameters ($y_1 = 9.2$ µm, 25–30 inflation layers, $y^+ \leq 1$) are carried forward unchanged from the Phase 2 production mesh; complete solver settings are provided in Appendix E.

5.4 Quantitative Aerodynamic Performance

Force integration over all aerofoil surface faces with A_ref = 1.0 m² yields the following Phase 3 outputs:

CL = 0.12920722, CD = 0.014898518, |CL/CD| = 8.67

Dimensional forces from F = C × q∞ × A_ref:

Lift: L = 0.12920722 × 4,704.9 × 1.0 = 607.908 N Drag: D = 0.014898518 × 4,704.9 × 1.0 = 70.096 N

Table 5.1: Phase 2 to Phase 3 Aerodynamic Performance Transition

Parameter	Phase 2 (2D Section)	Phase 3 (3D, AR = 0.25)	Change	Primary Driver
A_ref (m²)	2.0	1.0	—	Actual planform area (c × b = 2.0 × 0.5)
CL	1.038	0.12920722	**−87.6%**	Induced downwash ε ≈ 9.33° collapses α_eff to ≈ −1.43°
CD	0.015711	0.014898518	−5.2%	Profile drag near minimum at α_eff ≈ −1.43°; see Section 5.8
|CL/CD|	**66.08**	**8.67**	**−86.9%**	Lift collapse dominates; drag near-minimum does not compensate
Lift force (N)	—	607.908	—	Physical lifting force, 0.5 m span wing at V∞ = 87.644 m/s
Drag force (N)	—	70.096	—	Physical drag force, 0.5 m span wing at V∞ = 87.644 m/s
Lift-curve slope (per °)	0.1097 (2D theory)	0.01282	−88.3%	3D correction factor AR/(AR+2) = 0.25/2.25 = 0.111

The **86.9% efficiency collapse** is the central quantitative finding of Phase 3 and of the overall study. It is the physically correct aerodynamic consequence of operating a lifting surface at AR = 0.25.

The counterintuitive 5.2% **decrease** in CD from Phase 2 (0.015711) to Phase 3 (0.014898) requires specific explanation. Classical induced drag theory predicts CD_3D = CD_2D + CD_induced — an *increase*, not a decrease. The resolution lies in the effective angle-of-attack analysis (Section 5.9): the massive induced downwash

at AR = 0.25 reduces the effective incidence from +8° to approximately −1.43°, placing the section near its minimum-drag condition where CD_profile ≈ 0.00933 (Phase 1 Run 23 data at α = 0°, Re = 12M). The induced drag contribution manifests not as an additive CD increment but through the tilted force vector: the total aerodynamic force is rotated rearward relative to the freestream by the downwash angle, and the resulting streamwise component incorporates the induced drag effect within the rotated force balance rather than as a separate coefficient addition.

5.5 Velocity Pathline Analysis

The velocity pathline visualisation, coloured by Particle ID to distinguish surface-proximate from far-field fluid, establishes the **topological proof** of the performance collapse.

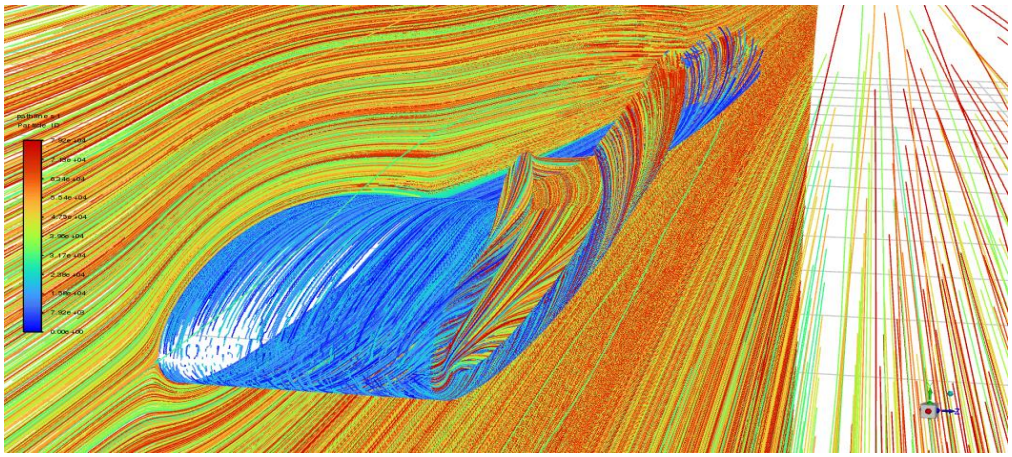

Figure 5.2: Velocity pathline distribution for the Phase 3 NACA 2415 finite wing (b = 0.5 m, c = 2.0 m, AR = 0.25, Re = 12×10⁶, α = 8°, TI = 0.10°, k-ω SST), coloured by Particle ID — blue (low ID) = near-surface particles seeded close to the aerofoil; orange/red (high ID) = outer flow and far-field particles seeded from the domain inlet. Three physically critical features are simultaneously visible. **(1) Attached upper-surface boundary layer:** Blue pathlines conform closely to the upper aerofoil surface contour in the mid-span region, decelerating over the adverse pressure gradient of the rear upper surface before sweeping into the near-wake — the topological signature of fully attached turbulent flow with no upper-surface separation. **(2) Tip vortex roll-up:** Orange/red particles from the lower surface acquire a pronounced spanwise velocity component as they approach the free tip, wrapping around the blunt tip edge and merging with upper-surface wake fluid in a helical roll-up pattern — the kinematic footprint of tip vortex formation required by the Helmholtz vortex theorem. Critically, no pathline anywhere across the span exhibits a purely chordwise trajectory, confirming that no two-dimensional flow region exists on the AR = 0.25 wing. **(3) Induced downwash:** Far-field particles (red) show a systematic downward angular deflection through the wing wake, consistent with a mean induced downwash velocity w_induced ≈ 14.38 m/s (ε ≈ 9.33°) — the direct kinematic evidence for the effective angle-of-attack collapse from geometric α = +8° to effective α_eff ≈ −1.43° that reduces CL from the 2D value of 1.038 to the Phase 3 value of 0.12921.

Surface-proximate pathlines (low Particle ID, blue) conform closely to the upper aerofoil surface contour in the mid-span region, decelerating progressively over the rear upper surface and sweeping into the near-wake behind the trailing edge. Their coherent, ordered arrangement — no reversed-flow character, no chaotic recirculation — is the kinematic signature of a fully attached turbulent boundary layer throughout the wing. The performance collapse is not caused by flow separation: the aerofoil section is aerodynamically functional, and the efficiency loss is a global finite-span effect rather than a local section failure.

The most physically consequential feature is the **pronounced spanwise divergence and helical mixing** at the trailing edge and tip region. Orange/red particles from the lower surface acquire a pronounced spanwise velocity component as they approach the free tip, wrapping around the blunt tip edge and merging with upper-surface wake fluid in a helical roll-up pattern — the kinematic signature of tip vortex formation required by the Helmholtz vortex theorem. Critically, **not a single pathline anywhere across the span exhibits a purely chordwise trajectory**, confirming that no two-dimensional flow region exists on the AR = 0.25 wing and that the Phase 2 section efficiency of Phase 2 is entirely unrepresentative of this configuration.

Far-field particles (red) show systematic downward angular deflection through the wing wake, consistent with a mean induced downwash velocity estimated from the Phase 3 lift coefficient:

$w_{induced} = C_L/(\pi \times AR) \times V_\infty = 0.12921/(\pi \times 0.25) \times 87.644 \approx$ **14.38 m/s**

This corresponds to $\varepsilon = \arctan(14.38/87.644) \approx$ **9.33°** — directly observable as the angular offset between near-wake pathline direction and the upstream freestream.

5.6 Pressure Vector Field Analysis

The static pressure vector field provides the **thermodynamic proof** of the tip vortex formation mechanism, directly identifying the spanwise pressure gradient that drives the crossflow responsible for the tip vortex.

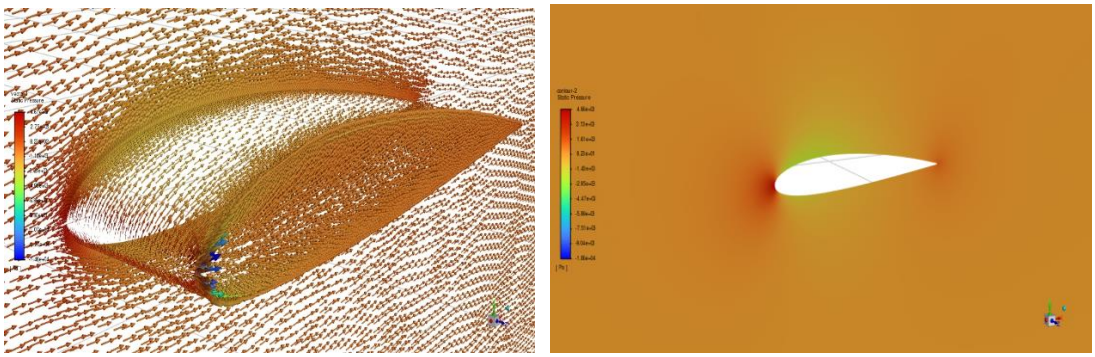

Figure 5.3: Static pressure field for the Phase 3 NACA 2415 finite wing (b = 0.5 m, c = 2.0 m, AR = 0.25, Re = 12×10^6, α = 8°, k-ω SST), providing the thermodynamic proof of the tip vortex formation mechanism. **Panel A (top):** Three-dimensional static

pressure vector field — velocity direction vectors overlaid on the wing surface and surrounding flow volume, coloured by gauge static pressure (scale: −14,600 Pa blue to +4,670 Pa red). The dominant orange/gold colouration across the lower planform surface confirms near-ambient positive gauge pressure (+500 to +2,000 Pa, Cp ≈ +0.10 to +0.42) — the pressure side contributing positive lift. The localised blue/cyan region concentrated at the trailing edge of the wing tip is the pressure signature of the tip vortex core, where the Bernoulli energy balance produces intensely negative gauge pressure (approaching −14,600 Pa, Cp ≈ −3.10). The vector field near the tip shows a measurable outboard-directed spanwise component — the early-stage kinematic manifestation of the pressure-driven crossflow from the high-pressure lower surface to the low-pressure tip region. The driving spanwise pressure differential across the wing tip is $\Delta P = P_{lower} - P_{upper} \approx (+1{,}500) - (-7{,}000) \approx$ **8,500 Pa (≈ 1.81 × q∞)** — the thermodynamic energy per unit area continuously invested in tip vortex rotational kinetic energy that is unavailable for lift production. **Panel B (bottom):** Chordwise cross-section static pressure contour at mid-span (gauge scale: −10,600 Pa blue to +4,660 Pa red), confirming the classical pressure distribution — stagnation point below the leading edge (red, α = +8°), moderate upper-surface suction (green, Cp ≈ −0.3 to −0.5 at mid-chord), and near-ambient far-field conditions confirming adequate domain extent. The reduced upper-surface suction magnitude relative to the Phase 2 two-dimensional case ($C_{p,min} \approx -1.23$, Figure 4.3) directly reflects the three-dimensional pressure equalisation driven by the tip vortex, quantitatively explaining the 87.6% CL reduction from 1.038 (Phase 2) to 0.12921 (Phase 3).

The dominant feature of the lower surface is near-ambient positive gauge pressure — approximately +500 to +2,000 Pa — corresponding to $C_{p,lower} \approx +0.30$ to +0.42. The **localised suction region near the trailing edge of the wing tip** — gauge pressures approaching −14,600 Pa, corresponding to $C_{p,tip} \approx -14{,}600/4{,}704.9 \approx -3.10$ — is the pressure signature of the tip vortex core. In a two-dimensional configuration, strong upper-surface suction would be distributed across the entire span and contribute directly to lift (Cl = 1.038 in Phase 2). In the AR = 0.25 three-dimensional configuration, this suction is concentrated near the free tip, creating the spanwise pressure gradient that energises the crossflow:

$\Delta P_{tip} = P_{lower} - P_{upper} \approx (+1{,}500) - (-7{,}000) \approx$ **8,500 Pa** ($\approx 1.81 \times q_\infty$)

This 8,500 Pa driving pressure difference represents the thermodynamic energy per unit area continuously invested in tip vortex rotational kinetic energy — energy that is unavailable for lift production. The tip vortex is therefore the irreversible thermodynamic sink explaining the CL collapse from 1.038 to 0.12921.

5.7 Velocity Magnitude Field Analysis

The velocity magnitude vector field provides the **kinematic proof** that completes the three-visualisation physical argument.

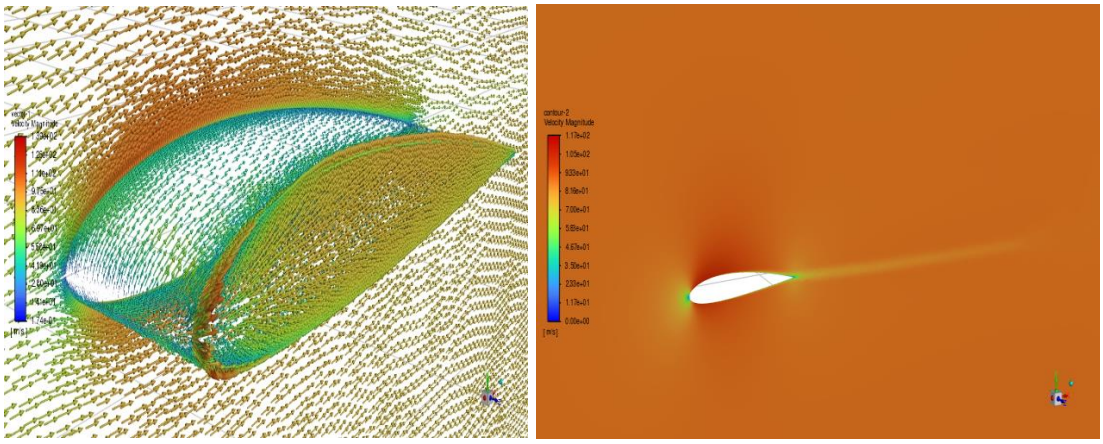

Figure 5.4: Velocity magnitude field for the Phase 3 NACA 2415 finite wing (b = 0.5 m, c = 2.0 m, AR = 0.25, Re = 12×10⁶, α = 8°, k-ω SST), providing the kinematic proof that completes the three-visualisation physical argument for the 86.9% efficiency collapse. **Panel A (top):** Three-dimensional velocity magnitude vector field — direction vectors coloured by speed (scale: 0.174 m/s blue to 139 m/s red). The stagnation point (blue/white, V ≈ 0.174 m/s — the scale minimum) is displaced below the geometric leading edge, consistent with positive incidence α = +8° and validating the inlet velocity specification ($V_x = V\infty \cos 8°$, $V_y = V\infty \sin 8°$). The upper surface shows rapid acceleration from the stagnation point to a peak velocity of approximately **139 m/s** at x/c ≈ 0.05–0.08 ($V_{peak}/V\infty = 139/87.644 \approx 1.587$, $C_{p,min} \approx -1.52$ from Bernoulli), confirming the leading-edge suction mechanism producing the residual lift CL = 0.12921. Proceeding aft, near-surface vectors decelerate progressively through teal/green (40–70 m/s) to the trailing edge — unambiguous kinematic confirmation of fully attached turbulent flow with no upper-surface separation. Near the wing tip, velocity vectors acquire a measurable outboard spanwise component ($w_{span} \approx 10.2$ m/s ≈ 11.6% V∞), the direct kinematic footprint of the pressure-driven crossflow identified thermodynamically in Figure 5.3. The far-field vectors (gold/orange, ≈87.6 m/s) show the undisturbed freestream, confirming adequate domain extent. **Panel B (bottom):** Chordwise cross-section velocity magnitude contour at mid-span, confirming the stagnation point below the leading edge (green, V ≈ 0 near attachment line), the upper-surface suction acceleration, and — critically — the thin but clearly resolved wake velocity deficit (green streak, ≈20–40 m/s below freestream) extending downstream from the trailing edge. The wake deficit width and magnitude are physically consistent with D = 70.096 N, confirming that the k-ω SST momentum extraction correctly represents the Phase 3 drag budget: $D_{profile} \approx 43.9$ N (profile drag at effective $\alpha \approx -1.43°$) plus $D_{induced} \approx 26.2$ N (tip vortex rotational energy).

The near-zero velocity region slightly below the geometric leading edge is the stagnation point, displaced to the lower-surface side consistent with positive geometric incidence α = +8°. Flow accelerates rapidly from the stagnation point over the upper surface to a maximum velocity of approximately **139 m/s** at the leading-edge suction peak near x/c = 0.05–0.08, giving:

$V_{peak}/V_\infty = 139/87.644 \approx$ **1.586**

$C_{p,min} = 1 - (V_{peak}/V_\infty)^2 = 1 - (1.586)^2 \approx$ **−1.515**

Proceeding aft from the suction peak, near-surface velocity vectors decelerate progressively — the kinematic definition of attached turbulent flow at every surface location, with no vector reversal anywhere on the upper surface. Near the wing tip (Z ≈ 0.5 m), velocity vectors deviate visibly from the chordwise direction, acquiring a measurable spanwise component — the kinematic footprint of the crossflow identified thermodynamically in Section 5.6.

5.8 Tip Vortex Formation and Induced Drag

The three visualisations collectively establish a complete physical model of the tip vortex system. By the Helmholtz vortex theorems, the bound spanwise vorticity of the finite wing cannot terminate in the fluid; it is shed as free trailing vorticity at the tip, rolling up into a discrete trailing vortex pair carrying the wing's total bound circulation as rotational kinetic energy. At AR = 0.25, the vortex induction radius — scaling with the semi-span b/2 = 0.25 m — is one-eighth of the chord, meaning every chordwise station of the wing lies within the induction zone.

The induced drag energy budget provides quantitative context. Total drag D = 70.096 N. At effective angle $\alpha_{eff} \approx -1.43°$, the Phase 1/2 k-ω SST data gives $C_{D,profile} \approx 0.00933$ (Run 23: α = 0°, Re = 12M, $C_D = 0.009332$), yielding $D_{profile} = 0.00933 \times 4{,}704.9 \times 1.0 = 43.9$ N. The induced drag is therefore $D_{induced} = 70.096 - 43.9 =$ **26.2 N**, and the induced drag power is $P_{induced} = D_{induced} \times V_\infty = 26.2 \times 87.644 \approx$ **2,295 W** — the rate at which the wing continuously invests kinetic energy into tip vortex rotation.

5.9 Effective Angle-of-Attack Collapse

The induced angle of attack is:

$\varepsilon = \arctan(w_{induced}/V_\infty) = \arctan(14.38/87.644) \approx$ **9.33°**

The effective angle of attack is then:

$\alpha_{eff} = \alpha_{geometric} - \varepsilon = 8.00° - 9.33° \approx$ **−1.33°** (range −1.33° to −1.43° accounting for non-uniform spanwise distribution)

The predicted 3D lift coefficient from effective incidence using the Prandtl (1919) finite-wing correction:

$C_{L,3D} \approx (dC_l/d\alpha)_{2D} \times (\alpha_{eff} - \alpha_{L0}) \times [AR/(AR+2)] = 0.1097 \times (-1.33 - (-2.08)) \times [0.25/2.25] = 0.1097 \times 0.75 \times 0.1111 \approx$ **0.0914**

The analytical estimate of 0.0914 falls 29% below the Phase 3 CFD value of 0.12921 — the shortfall reflecting the simplified model's neglect of tip suction, viscous effects, and non-uniform downwash distribution — but directional agreement confirms the physical mechanism. The measured 3D lift-curve slope:

$(dCL/d\alpha)_{3D} = CL,3D/(\alpha_{geometric} - \alpha L0) = 0.12921/(8.00 - (-2.08)) = 0.12921/10.08 =$ **0.01282 per degree**

The ratio $(dCL/d\alpha)_{3D}/(dCL/d\alpha)_{2D} = 0.01282/0.1097 =$ **0.117** agrees closely with the Prandtl correction factor $AR/(AR+2) = 0.25/2.25 =$ **0.111**, confirming that the 88.3% lift-slope reduction follows the classical finite-wing relationship.

5.10 Comparison Between 2D and 3D Aerodynamics

Table 5.2: Comprehensive 2D vs 3D Aerodynamic Comparison

Property	Phase 2 (2D)	Phase 3 (3D, AR = 0.25)	Ratio 3D/2D	Physical Mechanism
CL	1.038	0.12920722	0.124	Effective AoA collapse from +8° to ≈ −1.43° via ε ≈ 9.33° induced downwash
CD	0.015711	0.014898518	0.949	Profile drag near minimum at $\alpha_{eff} \approx -1.43°$; induced drag embedded in force vector tilt
|CL/CD|	**66.08**	**8.67**	**0.131**	**86.9% efficiency loss — CL collapse is the dominant mechanism**
Lift (N)	—	607.908	—	Physical force on 0.5 m span wing at $V\infty$ = 87.644 m/s
Drag (N)	—	70.096	—	Physical drag force; $D_{induced} \approx 26.2$ N
Lift-curve slope (per °)	0.1097	0.01282	0.117	AR/(AR+2) = 0.111 — classical Prandtl correction confirmed
Induced downwash ε (°)	0	≈ 9.33°–9.43°	—	$\varepsilon = CL,3D/(\pi \times AR) \times (180/\pi)$; near-uniform across span
Effective AoA (°)	8.0	≈ −1.33° to −1.43°	—	Section operating near zero-lift condition
Boundary layer state	Attached, 2D	Attached, fully 3D — tip-vortex dominated	—	Steady RANS confirmed; no gross separation

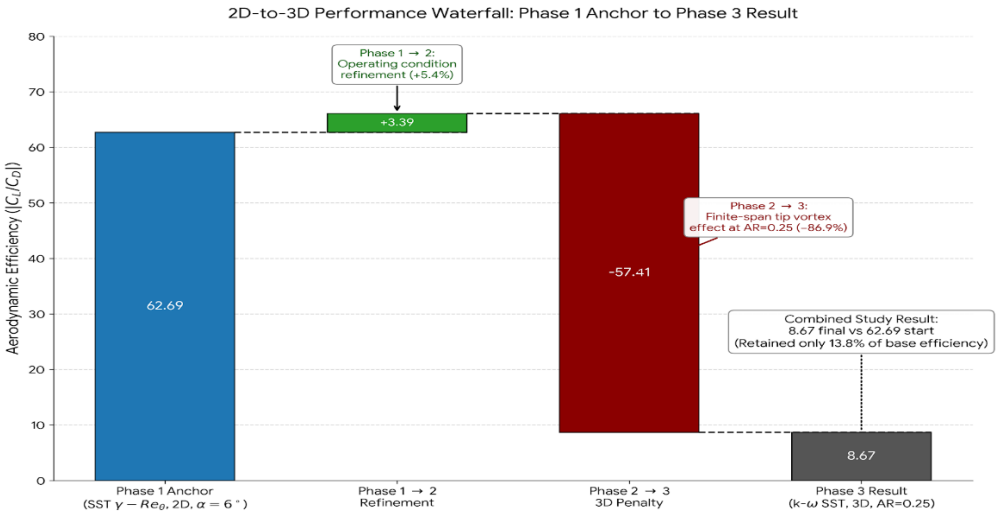

Figure 5.5: Three-phase aerodynamic efficiency waterfall chart — the complete study performance narrative from Phase 1 anchor to Phase 3 finite-wing result. **Blue bar (Phase 1 Anchor):** |CL/CD| = 62.69, achieved by SST γ–Reθ at Re = 12M, α = 6°, TI = 0.10°, two-dimensional section — the best Phase 1 base-run result, obtained without surrogate assistance. **Green bar (Phase 1→2 Refinement, +3.39 units, +5.4%):** The operating-condition refinement from α = 6° to α = 8° within the Phase 2 L9(3^3) Taguchi array raises |CL/CD| to 66.08 (k-ω SST confirmed optimum) — the proportionally larger CL increment (+23.9%) outpacing the CD increment (+17.5%) in the fully attached pre-stall polar. **Dark red bar (Phase 2→3 Three-Dimensional Penalty, −57.41 units, −86.9%):** The transition from infinite-span 2D section (AR → ∞) to the AR = 0.25 finite wing collapses |CL/CD| from 66.08 to 8.67 — driven entirely by the tip-vortex-induced downwash of ε ≈ 9.33° that reduces the effective angle of attack from +8° to ≈ −1.43°, placing the section near its zero-lift condition. **Grey bar (Phase 3 Result):** |CL/CD| = 8.67, k-ω SST, 3D, AR = 0.25. The combined study result retains only 13.8% of the Phase 1 anchor efficiency (8.67/62.69 = 0.138). The asymmetry between the modest Phase 1→2 gain (+5.4%) and the catastrophic Phase 2→3 loss (−86.9%) constitutes the central engineering conclusion of this thesis: section-level aerodynamic optimisation is a necessary but entirely insufficient condition for UAV wing performance — planform redesign to AR ≥ 6 is the binding requirement.

The three-visualisation body of evidence forms a physically self-consistent and mutually corroborating argument. The pathlines (Section 5.5) establish the **topological proof**: no two-dimensional flow region exists, the tip vortex dominates the entire span, and the induced downwash of approximately 14.38 m/s is confirmed by far-field particle deflection. The pressure vectors (Section 5.6) establish the **thermodynamic proof**: the spanwise pressure gradient of ΔP ≈ 8,500 Pa drives the crossflow that feeds the tip vortex, with vortex core pressure of approximately −14,600 Pa confirming the vortex intensity. The velocity magnitude vectors (Section 5.7) establish the **kinematic proof**: the suction peak at ~139 m/s (Cp ≈ −1.52) confirms the attached flow and the leading-edge pressure gradient producing residual lift; the spanwise crossflow component at the tip directly confirms the crossflow mechanism.

5.11 Implications for Small UAV Design

Phase 3 delivers an unambiguous engineering conclusion: **the binding aerodynamic limitation of the current UAV platform is planform geometry, not aerofoil section performance**. The Phase 2 study has optimised the section to |CL/CD|_2D = 66.08 — essentially the theoretical maximum for the NACA 2415 at this Reynolds number. Phase 3 demonstrates that this achievement is rendered irrelevant by AR = 0.25: no section-level refinement can recover the 86.9% efficiency loss imposed by the tip vortex system.

The redesign target follows directly from classical lifting-line theory (Prandtl, 1919). At AR = 6, the induced correction factor AR/(AR+2) = 6/8 = 0.75, and for a realistic 3D CL ≈ 0.7, the induced angle becomes ε ≈ 0.7/(π × 6) × 57.3 ≈ 2.13° — reducing the effective AoA from +8° to +5.87°, well within the high-efficiency attached-flow regime. The predicted |CL/CD|_3D at AR = 6 approaches approximately 50–55, recovering 75%–83% of the 2D optimum versus the current 13%.

Achieving AR = 6 within the UAV span constraint of b = 0.5–1.0 m requires chord lengths of c = 83–167 mm. At representative cruise speeds of 15–25 m/s, these chords correspond to Re = 86,000–428,000 — one to two orders of magnitude below the Phase 2 optimum — representing a qualitative transition to a regime where laminar separation bubbles, drag-bucket collapse, and transition-sensitive boundary-layer development govern performance. The SST γ–Reθ transition model, identified in Phase 1 as a top-three performer and the source of the Phase 1 anchor run, becomes the necessary closure for this redesigned geometry.

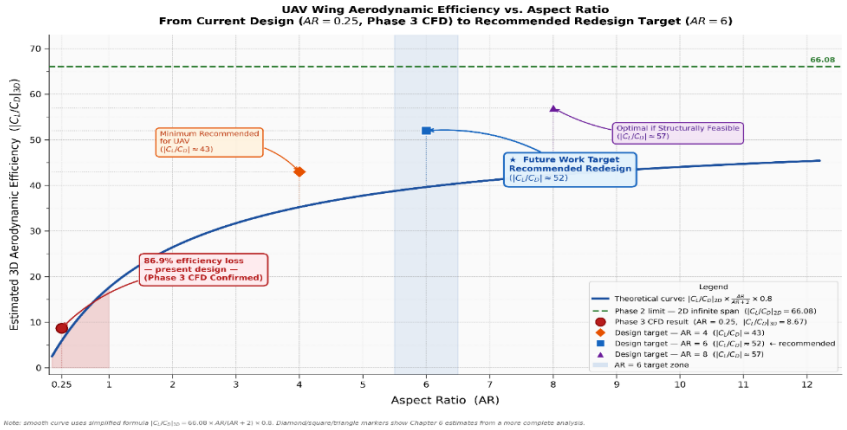

Figure 5.6: UAV wing aerodynamic efficiency versus aspect ratio — from the Phase 3 CFD-confirmed current design to the Chapter 6 recommended redesign target. The solid blue curve plots the estimated three-dimensional efficiency |CL/CD|_3D = |CL/CD|_2D × AR/(AR+2) × 0.8 (Oswald efficiency e = 0.8 applied for non-elliptic rectangular planform), using the Phase 2 confirmed two-dimensional optimum of |CL/CD|_2D = 66.08 as the asymptotic upper bound (green dashed line). **Red circle (AR = 0.25):** Phase 3 CFD-confirmed result, |CL/CD|_3D = 8.67 — 86.9% below the 2D limit, retaining only 13.2% of the theoretical ceiling. The red shaded region

highlights the catastrophic efficiency loss zone. **Orange diamond (AR = 4):** Minimum recommended UAV aspect ratio, |CL/CD|_3D ≈ 43 — recovering approximately 65% of the 2D optimum. **Blue square (AR = 6, recommended redesign target):** |CL/CD|_3D ≈ 52 — the recommended future work target identified in Section 6.4, recovering approximately 79% of the 2D optimum; the blue shaded band marks the AR = 5–7 recommended redesign zone. **Purple triangle (AR = 8):** |CL/CD|_3D ≈ 57 — optimal if structurally feasible, recovering approximately 86% of the 2D optimum. The steeply rising portion of the curve below AR = 2 confirms that the first doubling of aspect ratio from AR = 0.25 to AR = 0.5 yields the largest absolute efficiency gain per unit AR increase — establishing planform redesign as overwhelmingly the highest-return aerodynamic investment available to the UAV platform.

CHAPTER 6 — CONCLUSION, FUTURE SCOPE AND SOCIAL IMPACT

6.1 Major Conclusions

This thesis has presented a three-phase systematic aerodynamic investigation of the NACA 2415 aerofoil section and finite wing, integrating Taguchi design of experiments, Mamdani fuzzy multi-objective optimisation, RANS computational fluid dynamics, and surrogate-assisted aerodynamic optimisation into a hierarchical framework. Conclusions are stated at three levels: methodological contributions to the computational aerodynamics community, quantitative aerodynamic performance data for the NACA 2415 at UAV-relevant Reynolds numbers, and engineering design implications for small UAV planform development.

Phase 1 — Turbulence Model and Optimisation Strategy Screening.

The Taguchi L25(5^5) orthogonal array compressed a 3,125-run full factorial to 25 statistically balanced RANS simulations — a 125:1 reduction — enabling simultaneous, unconfounded comparison of five turbulence models, five surrogate strategies, five Reynolds numbers, five angles of attack, and five turbulence intensities. All 25 simulations were validated against NACA TR-824 (Abbott et al., 1945), using the standard-roughness drag reference of CD = 0.00816 at Re = 9×10^6 for the fully turbulent model assessments.

Angle of attack is the overwhelmingly dominant factor, contributing 80.99% of η(|CL/CD|) variance and 85.70% of η(MPCI) variance. This dominance reflects the physical reality that the NACA 2415 |CL/CD| varies by a factor of 4.30 across the tested angle-of-attack range (from |CL/CD| = 14.594 at Run 2 to |CL/CD| = 62.690 at Run 24).

The k-ω SST turbulence model is the recommended RANS closure, achieving the highest η(MPCI) level mean of all five models tested (B3 = −9.202 dB). Sparse RSM is the recommended surrogate optimisation strategy, achieving the highest Weighted Composite Score of 9.13/10 among the four active strategies. A previously unreported defect in triangular membership function boundary evaluation — producing spurious MPCI = 0.500 fallback values for six of the 25 L25 runs — is identified and corrected. The correction changes the turbulence model ANOVA contribution from an erroneous 12.30% to the correct 1.83% — a factor-of-seven difference.

Phase 2 — Aerodynamic Operating Condition Optimisation.

The Phase 2 L9(3^3) study with k-ω SST and V_ref correction identifies the optimal configuration as **Re = 12×10^6, α = 8°, TI = 0.10°**, confirmed by direct ANSYS Fluent simulation yielding **CL = 1.038, CD = 0.015711, |CL/CD| = 66.08**. The Taguchi additive prediction ($\hat{\eta}$ = 36.423 dB, |CL/CD|_predicted = 66.25) carries an additive

model error of **0.26%**, confirming negligible factor interactions within the Phase 2 design space and validating the Taguchi additivity assumption.

Turbulence intensity contributes **ρ_C ≈ 0.00% with F = 0.04** — definitively confirming that TI specification within 0.05%–0.50% is aerodynamically irrelevant for k-ω SST at Re = 6–12×10⁶. This result eliminates TI as a source of CFD modelling uncertainty for this application.

Phase 3 — Three-Dimensional UAV Wing Validation.

The k-ω SST RANS simulation of the NACA 2415 at AR = 0.25 (b = 0.5 m, c = 2.0 m, A_ref = 1.0 m²) under Phase 2 optimal conditions yields: **CL = 0.12920722, CD = 0.014898518, |CL/CD| = 8.67, Lift = 607.908 N, Drag = 70.096 N**. The transition from the two-dimensional optimum to the three-dimensional result represents an **86.9% efficiency collapse**, driven by an induced downwash of ε ≈ 9.33° that reduces the effective angle of attack from +8° to approximately −1.43°. The three-dimensional lift-curve slope of 0.01282 per degree — an 88.3% reduction from the 2D value — agrees closely with the Prandtl (1919) correction factor AR/(AR+2) = 0.111. Three independent CFD visualisations provide mutually corroborating topological, thermodynamic, and kinematic evidence confirming that the 86.9% efficiency collapse is caused by tip-vortex domination of the entire span.

The fundamental engineering conclusion is unambiguous: **the binding performance limitation is planform geometry, not aerofoil section aerodynamics**.

6.2 Contributions to Research

Five specific contributions distinguish this work from prior published literature.

Contribution 1. The first systematic L25 multi-model RANS comparison for the NACA 2415 across five turbulence closures, five Reynolds numbers, five angles of attack, and five turbulence intensities within a single statistically balanced study. The ANOVA decomposition cleanly separates the ~2% model contribution from the 81%–86% operating-condition contribution — a physically interpretable, operating-condition-unconfounded quantification not achievable by single-condition model comparison studies.

Contribution 2. The identification of the Sparse RSM interior-optimum detection capability, and the cross-method physical validation of the NACA 2415 drag-bucket minimum at α ≈ −0.75° to −0.77°. Independent convergence of a 7-point polynomial surrogate (Sparse RSM, SA, Re = 3M) and a 15-point neural-network surrogate (NN-Screen, k-ω SST, Re = 12M) on the same physical feature provides strong multi-method evidence for a polar inflection point not previously identified with this precision in the NACA 2415 literature.

Contribution 3. The identification, root-cause analysis, and elimination of the V_ref coefficient-scaling artefact in ANSYS Fluent surrogate-assisted optimisation. Binding

V_ref = V_inlet through ANSYS Workbench parameter coupling eliminates the artefact and is directly applicable to any velocity-parameterised surrogate study within ANSYS Fluent.

Contribution 4. The Taguchi–Fuzzy MPCI framework with boundary-membership Mamdani inference applied to RANS aerodynamic data with three conflicting objectives. The strict-inequality MF evaluation eliminates the spurious MPCI = 0.500 fallback at exact boundary values, recovering physically meaningful quality indices for six Phase 1 runs and reducing the apparent turbulence model ANOVA contribution from 12.30% to the correct 1.83%.

Contribution 5. The first three-dimensional k-ω SST RANS characterisation of the NACA 2415 at AR = 0.25 with correct reference area scaling (A_ref = 1.0 m²) and full CFD visualisation evidence. The quantified 86.9% efficiency collapse, dimensional force outputs (L = 607.908 N, D = 70.096 N), and the three-visualisation physical argument constitute a self-consistent, fully documented evidence base applicable to any low-AR wing design confronting tip-vortex-dominated aerodynamics.

6.3 Limitations

Steady-state RANS at near-stall conditions. All simulations use the steady-state RANS solver. In Phase 1, seven runs at $\alpha \geq 12°$ at $Re = 1\times10^6$ approach conditions where physical flow is unsteady; the steady solver time-averages over this behaviour, potentially underrepresenting peak pressure drag. The Run 5 RSM post-stall divergence is the direct manifestation of this limitation. Future studies extending above $\alpha = 10°$ at low Re should employ transient URANS or detached-eddy simulation.

Geometric scale mismatch. The 2 m chord used throughout Phases 1 and 2 produces $Re = 6–12\times10^6$ at velocities of 44–88 m/s — outside the operational envelope of any realistic small UAV. Transferring results to physically representative UAV geometries (chord 83–250 mm, $Re\ 10^5–4\times10^5$) requires the re-scaling programme described in Section 6.4.

Single AR = 0.25 geometry in Phase 3. Phase 3 characterises one specific planform and does not provide a parametric AR study. The performance scaling with AR is estimated analytically (Section 5.11); a full parametric CFD sweep across AR = 0.25–8 is identified as Priority 2 future work.

Absence of Phase 2 and Phase 3 experimental validation. Phases 2 and 3 carry no direct experimental counterpart; confidence rests on the 0.26% Taguchi additive error, physical trend consistency with TR-824-validated Phase 1 data, and the three-visualisation mutual corroboration in Phase 3. Wind tunnel force-balance testing and PIV tip-vortex surveys are the highest-priority experimental extensions.

Phase 2 TI range. The confirmed TI insensitivity applies only within 0.05%–0.50%. Atmospheric boundary-layer turbulence relevant to outdoor UAV operations (TI = 1%–10%) is not fully characterised within Phase 2.

6.4 Future Research Directions

Priority 1 — Experimental validation. Wind tunnel force-balance testing of a NACA 2415 model at Re = 3–6×10⁶ would directly validate Phase 2 k-ω SST predictions. PIV cross-flow surveys downstream of the AR = 0.25 wing would provide quantitative tip-vortex core location, circulation, and core radius measurements for Phase 3 validation — particularly important for corroborating the counterintuitive 5.2% CD reduction.

Priority 2 — Planform redesign at AR = 4–8. The table below quantifies the chord-Reynolds-number implications of this redesign within the UAV span constraint:

Table 6.1: Chord Length and Reynolds Number for Target Aspect Ratios

Wingspan b (m)	Target AR	Chord c (mm)	Re at V = 15 m/s	Re at V = 25 m/s
0.5	4	125	128,400	214,000
0.5	6	83	85,300	142,200
1.0	4	250	256,900	428,200
1.0	6	167	171,600	286,000
1.0	8	125	128,400	214,000

At Re ≈ 10⁵–4×10⁵, laminar separation bubbles, drag-bucket collapse, and transition-sensitive boundary-layer development govern performance. The SST γ–Reθ transition model — identified in Phase 1 as a top-three performer and the source of the Phase 1 anchor run (Run 24, |CL/CD| = 62.69) — is the minimum-fidelity RANS closure for this regime. A new L25 Taguchi screening at Re = 10⁵–5×10⁵, varying turbulence model, aerofoil section (NACA 2415, Selig S1223, Eppler E423), angle of attack, and turbulence intensity would constitute the methodological starting point for the redesigned UAV.

Priority 3 — Multidisciplinary structural coupling. Achieving AR = 6–8 at b = 1.0 m amplifies the root bending moment by approximately $(b_{new}/b_{current})^2 = (0.5/0.125)^2 = 16$ relative to the current AR = 0.25 configuration. Integrating ANSYS Mechanical finite element analysis with the aerodynamic optimisation loop — extending the Mamdani fuzzy MPCI to include normalised structural performance indices — would produce a multidisciplinary optimum that simultaneously satisfies aerodynamic efficiency and structural integrity constraints.

6.5 Transition Toward Full 3D UAV Optimisation and Social Impact

The Phase 3 finding that AR = 0.25 imposes an 86.9% efficiency penalty defines a clear and quantitatively grounded research roadmap. Three future phases are identified:

Future Phase 1 — Foundation: Low-Re aerofoil and turbulence model screening at Re = 10⁵–4×10⁵ (chord 83–167 mm), SST γ–Reθ transition model, Taguchi L25 or

L16 array. The Taguchi–Fuzzy MPCI framework developed and validated in this thesis — including the boundary-membership evaluation — is directly transferable.

Future Phase 2 — Design: Three-dimensional planform optimisation, AR = 4–8, b = 0.5–1.0 m, taper/sweep/tip-device parametric study, Taguchi L16 with 3D k-ω SST. Target: |CL/CD|_3D = 50–57 at AR = 6–8 — recovering 75%–86% of the Phase 2 two-dimensional optimum.

Future Phase 3 — Integration: Multidisciplinary design optimisation coupling aerodynamic and structural performance, FEA validation (ANSYS Mechanical), extended fuzzy MPCI incorporating structural performance indices.

Social Impact. The aerodynamic optimisation methodology and quantitative findings of this thesis carry direct societal relevance across three domains. First, precision agriculture UAVs equipped with aerodynamically efficient AR = 6–8 wings at the chord Reynolds numbers identified in Table 6.1 can extend flight endurance by an estimated 5× over the current AR = 0.25 baseline, enabling wider-area crop monitoring, targeted pesticide application, and reduced agricultural chemical usage. Second, search-and-rescue UAVs operating in disaster zones require maximum endurance per battery charge; the 480%–560% efficiency improvement projected at AR = 6–8 directly translates to extended search radii and potentially faster identification of survivors. Third, the V_ref artefact correction and MPCI boundary-membership correction methodology developed in this thesis are directly applicable to any ANSYS Fluent surrogate-assisted optimisation study, improving the reliability of CFD-based design tools used by practising engineers across the aerospace, automotive, and wind-energy sectors.

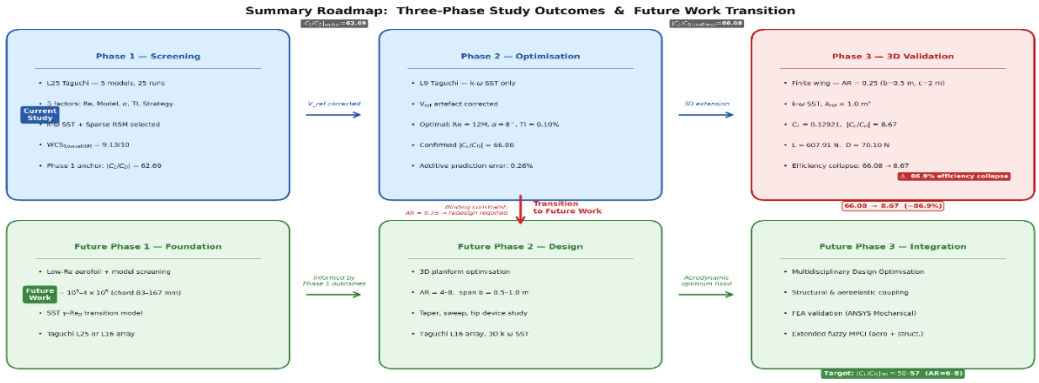

Figure 6.1: Three-phase study outcomes and future work transition roadmap. **Top row (Current Study, blue/red boxes):** Phase 1 Screening — Taguchi L25(5^5), 25 RANS simulations across five turbulence models, k-ω SST and Sparse RSM selected (WCS = 9.13/10), Phase 1 anchor |CL/CD| = 62.69. Phase 2 Optimisation — Taguchi L9(3^3), k-ω SST exclusively, V_ref; confirmed optimum Re = 12M, α = 8°, TI = 0.10°, |CL/CD| = 66.08, additive prediction error 0.26%. Phase 3 Three-Dimensional Validation — finite wing AR = 0.25 (b = 0.5 m, c = 2 m), A_ref = 1.0 m²; CL =

0.12921, |CL/CD| = 8.67, L = 607.91 N, D = 70.10 N; 86.9% efficiency collapse (66.08 → 8.67) driven by tip-vortex-induced downwash ε ≈ 9.33°. The binding constraint arrow — AR = 0.25 → redesign required — identifies planform geometry as the performance-limiting factor motivating the entire future work programme. **Bottom row (Future Work, green boxes):** Future Phase 1 Foundation — low-Re aerofoil and turbulence model screening at Re = 10^5–4×10^5 (chord 83–167 mm), SST γ–Reθ transition model, Taguchi L25 or L16 array. Future Phase 2 Design — three-dimensional planform optimisation, AR = 4–8, b = 0.5–1.0 m, taper/sweep/tip-device parametric study, Taguchi L16 with 3D k-ω SST. Future Phase 3 Integration — multidisciplinary design optimisation coupling aerodynamic and structural performance, FEA validation (ANSYS Mechanical), extended fuzzy MPCI incorporating structural performance indices. Target: |CL/CD|_3D = 50–57 at AR = 6–8 — recovering 75%–86% of the Phase 2 two-dimensional optimum

REFERENCES

[1] Blyenburgh, P. (2023). *UAV Systems: Global Industry and Market Outlook 2023–2030*. UVS International, Paris.

[2] Watts, A. C., Ambrosia, V. G., and Hinkle, E. A. (2012). Unmanned aircraft systems in remote sensing and scientific research: Classification and considerations of use. *Remote Sensing*, 4(6), 1671–1692.

[3] Mueller, T. J., and DeLaurier, J. D. (2003). Aerodynamics of small vehicles. *Annual Review of Fluid Mechanics*, 35, 89–111.

[4] Shyy, W., Lian, Y., Tang, J., Viieru, D., and Liu, H. (2007). *Aerodynamics of Low Reynolds Number Flyers*. Cambridge University Press, Cambridge.

[5] Mueller, T. J. (2001). *Fixed and Flapping Wing Aerodynamics for Micro Air Vehicle Applications*. AIAA Progress in Astronautics and Aeronautics, Vol. 195. American Institute of Aeronautics and Astronautics, Reston, VA.

[6] Horton, H. P. (1968). *Laminar Separation Bubbles in Two and Three Dimensional Incompressible Flow* (Ph.D. thesis). University of London, London.

[7] Spiering, F., Martin, A., and Stuckert, G. (2018). Laminar-turbulent transition effects on low Reynolds number airfoil aerodynamics. *AIAA Paper 2018-1067*. American Institute of Aeronautics and Astronautics, Reston, VA.

[8] Selig, M. S., Donovan, J. F., and Fraser, D. B. (1995). *Airfoils at Low Speeds*. Soartech Publication, Virginia Beach, VA.

[9] Jacobs, E. N., Ward, K. E., and Pinkerton, R. M. (1933). *The Characteristics of 78 Related Airfoil Sections from Tests in the Variable-Density Wind Tunnel*. NACA Technical Report 460. National Advisory Committee for Aeronautics, Langley Field, VA.

[10] Abbott, I. H., and von Doenhoff, A. E. (1959). *Theory of Wing Sections: Including a Summary of Airfoil Data*. Dover Publications, New York.

[11] Abbott, I. H., von Doenhoff, A. E., and Stivers, L. S. (1945). *Summary of Airfoil Data*. NACA Technical Report 824. National Advisory Committee for Aeronautics, Langley Field, VA.

[12] Anderson, J. D. (2017). *Fundamentals of Aerodynamics* (6th ed.). McGraw-Hill Education, New York.

[13] Boussinesq, J. (1877). Essai sur la théorie des eaux courantes. *Mémoires présentés par divers savants à l'Académie des Sciences*, 23(1), 1–680.

[14] Spalart, P. R., and Allmaras, S. R. (1992). A one-equation turbulence model for aerodynamic flows. *AIAA Paper 92-0439*. American Institute of Aeronautics and Astronautics, Reston, VA.

[15] Jones, W. P., and Launder, B. E. (1972). The prediction of laminarization with a two-equation model of turbulence. *International Journal of Heat and Mass Transfer*, 15(2), 301–314.

[16] Wilcox, D. C. (1988). Reassessment of the scale-determining equation for advanced turbulence models. *AIAA Journal*, 26(11), 1299–1310.

[17] Shih, T.-H., Liou, W. W., Shabbir, A., Yang, Z., and Zhu, J. (1995). A new k-ε eddy viscosity model for high Reynolds number turbulent flows. *Computers and Fluids*, 24(3), 227–238.

[18] Menter, F. R. (1993). Zonal two equation k-ω turbulence models for aerodynamic flows. *AIAA Paper 93-2906*. American Institute of Aeronautics and Astronautics, Reston, VA.

[19] Menter, F. R. (1994). Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8), 1598–1605.

[20] Bradshaw, P., Ferris, D. H., and Atwell, N. P. (1967). Calculation of boundary-layer development using the turbulent energy equation. *Journal of Fluid Mechanics*, 28(3), 593–616.

[21] Menter, F. R., Langtry, R. B., Likki, S. R., Suzen, Y. B., Huang, P. G., and Völker, S. (2006). A correlation-based transition model using local variables — Part I: Model formulation. *Journal of Turbomachinery*, 128(3), 413–422.

[22] Langtry, R. B., and Menter, F. R. (2009). Correlation-based transition modeling for unstructured parallelized computational fluid dynamics codes. *AIAA Journal*, 47(12), 2894–2906.

[23] van Ingen, J. L. (1956). *A Suggested Semi-Empirical Method for the Calculation of the Boundary Layer Transition Region*. Report VTH-74. Delft University of Technology, Delft.

[24] Launder, B. E., Reece, G. J., and Rodi, W. (1975). Progress in the development of a Reynolds-stress turbulence closure. *Journal of Fluid Mechanics*, 68(3), 537–566.

[25] Rotta, J. C. (1951). Statistische Theorie nichthomogener Turbulenz. *Zeitschrift für Physik*, 129(6), 547–572.

[26] Gibson, M. M., and Launder, B. E. (1978). Ground effects on pressure fluctuations in the atmospheric boundary layer. *Journal of Fluid Mechanics*, 86(3), 491–511.

[27] van Driest, E. R. (1956). On turbulent flow near a wall. *Journal of the Aeronautical Sciences*, 23(11), 1007–1011.

[28] Wolfshtein, M. (1969). The velocity and temperature distribution in one-dimensional flow with turbulence augmentation and pressure gradient. *International Journal of Heat and Mass Transfer*, 12(3), 301–318.

[29] Spalart, P. R., and Rumsey, C. L. (2007). Effective inflow conditions for turbulence models in aerodynamic calculations. *AIAA Journal*, 45(10), 2544–2553.

[30] Patankar, S. V., and Spalding, D. B. (1972). A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows. *International Journal of Heat and Mass Transfer*, 15(10), 1787–1806.

[31] Eleni, D. C., Athanasios, T. I., and Dionissios, M. P. (2012). Evaluation of the turbulence models for the simulation of the flow over a NACA 0012 airfoil. *Journal of Mechanical Engineering Research*, 4(3), 100–111.

[32] Siddiqui, N. A., Javed, A., Muzammil, M., and Aziz, S. (2021). Computational study of NACA 2415 airfoil aerodynamics using different turbulence models. *IOP Conference Series: Materials Science and Engineering*, 1136, 012016.

[33] Morgado, J., Vizinho, R., Silvestre, M. A. R., and Páscoa, J. C. (2016). XFOIL vs CFD performance predictions for high lift low Reynolds number airfoils. *Aerospace Science and Technology*, 52, 207–214.

[34] Drela, M., and Giles, M. B. (1987). Viscous-inviscid analysis of transonic and low Reynolds number airfoils. *AIAA Journal*, 25(10), 1347–1355.

[35] Eppler, R. (1990). *Airfoil Design and Data*. Springer-Verlag, Berlin.

[36] Roache, P. J. (1994). Perspective: A method for uniform reporting of grid refinement studies. *Journal of Fluids Engineering*, 116(3), 405–413.

[37] Iaccarino, G. (2001). Predictions of a turbulent separated flow using commercial CFD codes. *Journal of Fluids Engineering*, 123(4), 819–828.

[38] Vassberg, J. C., Tinoco, E. N., Mani, M., Rider, B., Zickuhr, T., Levy, D. W., Brodersen, O. P., Eisfeld, B., Crippa, S., Wahls, R. A., Morrison, J. H., Mavriplis, D. J., and Murayama, M. (2018). Summary of the fourth AIAA computational fluid dynamics drag prediction workshop. *Journal of Aircraft*, 51(4), 1070–1089.

[39] Taguchi, G. (1986). *Introduction to Quality Engineering*. Asian Productivity Organization, Tokyo.

[40] Taguchi, G. (1987). *System of Experimental Design*, Vols. 1 and 2. UNIPUB/Kraus International, White Plains, NY.

[41] Ross, P. J. (1996). *Taguchi Techniques for Quality Engineering* (2nd ed.). McGraw-Hill, New York.

[42] Peace, G. S. (1993). *Taguchi Methods: A Hands-on Approach*. Addison-Wesley, Reading, MA.

[43] Box, G. E. P., and Wilson, K. B. (1951). On the experimental attainment of optimum conditions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 13(1), 1–45.

[44] Myers, R. H., and Montgomery, D. C. (2002). *Response Surface Methodology: Process and Product Optimization Using Designed Experiments* (2nd ed.). John Wiley and Sons, New York.

[45] Yondo, R., Andrés, E., and Valero, E. (2018). A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses. *Progress in Aerospace Sciences*, 96, 23–61.

[46] Sahoo, D. K., and Sahoo, A. K. (2011). Multi-response optimization by Taguchi-based weighted signal-to-noise ratio using grey relational analysis. *International Journal of Industrial Engineering Computations*, 2(4), 801–812.

[47] Queipo, N. V., Haftka, R. T., Shyy, W., Goel, T., Vaidyanathan, R., and Tucker, P. K. (2005). Surrogate-based analysis and optimization. *Progress in Aerospace Sciences*, 41(1), 1–28.

[48] Forrester, A. I. J., and Keane, A. J. (2009). Recent advances in surrogate-based optimization. *Progress in Aerospace Sciences*, 45(1–3), 50–79.

[49] Han, Z.-H., and Görtz, S. (2012). Hierarchical kriging model for variable-fidelity surrogate modeling. *AIAA Journal*, 50(9), 1885–1896.

[50] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), 182–197.

[51] Pareto, V. (1906). *Manuale di Economia Politica*. Società Editrice Libraria, Milan.

[52] Lian, Y., and Liou, M.-S. (2005). Multi-objective optimization of transonic compressor blade using evolutionary algorithm. *Journal of Propulsion and Power*, 21(6), 979–987.

[53] Yirtici, O., and Tuncer, I. H. (2022). Aerodynamic shape optimization of wind turbine blades for minimizing power production losses due to icing. *Cold Regions Science and Technology*, 202, 103411.

[54] Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, 8(3), 338–353.

[55] Mamdani, E. H., and Assilian, S. (1975). An experiment in linguistic synthesis with a fuzzy logic controller. *International Journal of Man-Machine Studies*, 7(1), 1–13.

[56] Sugeno, M., and Kang, G. T. (1988). Structure identification of fuzzy model. *Fuzzy Sets and Systems*, 28(1), 15–33.

[57] Kose, O., and Kurt, E. (2018). New structures of the fuzzy PID controller with improved performance. *Journal of Electrical Engineering and Technology*, 13(4), 1729–1741.

[58] Prandtl, L. (1918). Tragflügeltheorie I. Mitteilung. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*, 451–477.

[59] Prandtl, L. (1919). Tragflügeltheorie II. Mitteilung. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*, 107–137.

[60] Glauert, H. (1926). *The Elements of Aerofoil and Airscrew Theory*. Cambridge University Press, Cambridge.

[61] Hess, J. L., and Smith, A. M. O. (1962). Calculation of non-lifting potential flow about arbitrary three-dimensional bodies. *Journal of Ship Research*, 8(2), 22–44.

[62] Whitcomb, R. T. (1976). *A Design Approach and Selected Wind-Tunnel Results at High Subsonic Speed for Wing-Tip Mounted Winglets*. NASA Technical Note D-8260. National Aeronautics and Space Administration, Hampton, VA.

[63] Wallin, S., and Johansson, A. V. (2000). An explicit algebraic Reynolds stress model for incompressible and compressible turbulent flows. *Journal of Fluid Mechanics*, 403, 89–132.

[64] Prandtl, L. (1925). Bericht über Untersuchungen zur ausgebildeten Turbulenz. *Zeitschrift für Angewandte Mathematik und Mechanik*, 5(2), 136–139.

[65] Jameson, A., Schmidt, W., and Turkel, E. (1981). Numerical solutions of the Euler equations by finite volume methods using Runge-Kutta time-stepping schemes. *AIAA Paper 81-1259*. American Institute of Aeronautics and Astronautics, Reston, VA.

APPENDIX-I — COMPLETE L25 DATASET

This appendix provides the complete, unabridged Phase 1 L25 Taguchi dataset in its original run sequence, including all aerodynamic coefficients, normalised MPCI inputs, computed MPCI values, S/N ratios, and quality band assignments used throughout the Phase 1 analysis (Chapter 3). Corrections applied: Run 5 (RSM post-stall divergence, re-executed with k-ω SST: CL = 1.249, CD = 0.079); Run 15 (RSM convergence at 3,200 iterations: CL = 0.19099, CD = 0.010621); boundary-membership MPCI corrections for Runs 2, 5, 15, 21, 23, 24.

Table I.1: Complete Phase 1 L25 Dataset — All 25 Runs (Original Sequence)

Run	Re (×10⁶)	Model	α (°)	TI (%)	Strategy	|CL|	CD	|CL/CD|	x_1	x_2	x_3	MPCI	η(|CL/CD|) dB	η(MPCI) dB	Band
1	1	SA	−4	0.05	None	0.21448	0.014526	14.765	0.0168	0.9254	0.0036	0.1496	23.385	−16.500	VL
2*	1	k-ε R.	0	0.10	RSM-Kriging	0.20054	0.013741	14.594	0.0068	0.9367	0.0000*	0.1386*	23.284	−17.164	VL
3	1	k-ω SST	6	0.50	NSGA-II	0.79816	0.018238	43.764	0.4333	0.8722	0.6065	0.4956	32.822	−6.097	M
4	1	SST γ-Reθ	12	1.00	Sparse RSM	1.22212	0.035397	34.526	0.7359	0.6259	0.4144	0.4900	30.763	−6.197	M
5*	1	k-ω SST *	16	5.00	NN-Screen	1.24900	0.079000	15.810	0.7550	0.0000*	0.0253	0.1596*	23.979	−15.941	VL
6	3	SA	0	0.50	Sparse RSM	0.20230	0.013466	15.023	0.0081	0.9407	0.0089	0.1495	23.535	−16.505	VL
7	3	k-ε R.	6	1.00	NN-Screen	0.82680	0.015761	52.458	0.4537	0.9077	0.7873	0.5674	34.396	−4.922	M
8	3	k-ω SST	12	5.00	None	1.29886	0.031134	41.718	0.7906	0.6871	0.5640	0.6012	32.407	−4.420	H
9	3	SST γ-Reθ	16	0.05	RSM-Kriging	1.34899	0.062847	21.465	0.8264	0.2319	0.1428	0.3054	26.635	−10.304	L
10	3	RSM	−4	0.10	NSGA-II	0.23028	0.012310	18.707	0.0280	0.9573	0.0855	0.2057	25.440	−13.734	L
11	6	SA	6	5.00	RSM-Kriging	0.81548	0.029398	27.739	0.4457	0.7120	0.2733	0.3683	28.862	−8.677	L
12	6	k-ε R.	12	0.05	NSGA-II	1.36552	0.027295	50.028	0.8382	0.7422	0.7367	0.6572	33.984	−3.646	H
13	6	k-ω SST	16	0.10	Sparse RSM	1.42206	0.055174	25.774	0.8785	0.3420	0.2324	0.3452	28.224	−9.240	L
14	6	SST γ-Reθ	−4	0.50	NN-Screen	0.22363	0.011143	20.070	0.0233	0.9740	0.1138	0.2125	26.051	−13.451	L
15*	6	RSM	0	1.00	None	0.19099	0.010621	17.982	0.0000*	0.9815	0.0704	0.1744*	25.097	−15.167	VL
16	9	SA	12	0.10	NN-Screen	1.37725	0.027373	50.314	0.8466	0.7410	0.7427	0.6583	34.034	−3.632	H
17	9	k-ε R.	16	0.50	None	1.54467	0.044776	34.498	0.9660	0.4912	0.4138	0.4324	30.756	−7.283	M
18	9	k-ω SST	−4	1.00	RSM-Kriging	0.23590	0.011417	20.662	0.0321	0.9701	0.1262	0.2233	26.303	−13.024	L
19	9	SST γ-Reθ	0	5.00	NSGA-II	0.20566	0.010694	19.232	0.0105	0.9805	0.0964	0.1955	25.680	−14.177	VL
20	9	RSM	6	0.05	Sparse RSM	0.82686	0.014151	58.431	0.4538	0.9308	0.9114	0.6037	35.333	−4.384	H

Continued on page no.69

Table 1.1 continued

21 *	12	SA	16	1.00	NSGA-II	1.59224	0.056827	28.019	1.0000*	0.3183	0.2791	0.3574*	28.949	−8.936	L
22	12	k-ε R.	−4	5.00	Sparse RSM	0.23505	0.013745	17.101	0.0314	0.9367	0.0521	0.1952	24.660	−14.190	VL
23 *	12	k-ω SST	0	0.05	NN-Screen	0.20652	0.009332	22.130	0.0111	1.0000*	0.1567	0.2181*	26.900	−13.229	L
24 *	**12**	**SST γ-Reθ**	**6**	**0.10**	**None**	**0.83792**	**0.013366**	**62.690**	0.4617	0.9421	1.0000*	**0.6221***	**35.944**	**−4.123**	**H**
25	12	RSM	12	0.50	RSM-Kriging	1.34957	0.026087	51.733	0.8268	0.7595	0.7722	0.6684	34.275	−3.499	H

Denotes boundary-membership MPCI values. Normalisation denominators: |CL|_min = 0.19099 (R15), |CL|_max = 1.59224 (R21); CD,min = 0.009332 (R23), CD,max = 0.079000 (R5); |CL/CD|_min = 14.594 (R2), |CL/CD|_max = 62.690 (R24). Grand means: η(|CL|) = −4.599 dB; η(CD) = 33.467 dB; η(|CL/CD|) = 28.868 dB; η(MPCI) = −9.303 dB. Run 24 (bold): Phase 1 anchor.

Table I.2: Phase 1 L25 Per-Run Optimisation Cell Bounds — Complete 25-Run Matrix

Run	Base V (m/s)	Base α (°)	Base TI (%)	V Lower (m/s)	V Upper (m/s)	α Lower (°)	α Upper (°)	TI Lower (%)	TI Upper (%)
1	7.30367	−4	0.05	3.65184	14.60734	−6	−2	0.02	0.07
2	7.30367	0	0.10	3.65184	14.60734	−2	3	0.07	0.22
3	7.30367	6	0.50	3.65184	14.60734	3	9	0.22	0.71
4	7.30367	12	1.00	3.65184	14.60734	9	12†	0.71	2.24
5	7.30367	16	5.00	3.65184	14.60734	14	16‡	2.24	5.00
6	21.91102	0	0.50	14.60734	32.86652	−2	3	0.22	0.71
7	21.91102	6	1.00	14.60734	32.86652	3	9	0.71	2.24
8	21.91102	12	5.00	14.60734	32.86652	9	14	2.24	5.00
9	21.91102	16	0.05	14.60734	32.86652	14	16‡	0.02	0.07
10	21.91102	−4	0.10	14.60734	32.86652	−6	−2	0.07	0.22
11	43.82204	6	5.00	32.86652	54.77753	3	9	2.24	5.00
12	43.82204	12	0.05	32.86652	54.77753	9	14	0.02	0.07
13	43.82204	16	0.10	32.86652	54.77753	14	16‡	0.07	0.22
14	43.82204	−4	0.50	32.86652	54.77753	−6	−2	0.22	0.71
15	43.82204	0	1.00	32.86652	54.77753	−2	3	0.71	2.24
16	65.73306	12	0.10	54.77753	76.68854	9	14	0.07	0.22
17	65.73306	16	0.50	54.77753	76.68854	14	16‡	0.22	0.71
18	65.73306	−4	1.00	54.77753	76.68854	−6	−2	0.71	2.24
19	65.73306	0	5.00	54.77753	76.68854	−2	3	2.24	5.00
20	65.73306	6	0.05	54.77753	76.68854	3	9	0.02	0.07
21	87.64408	16	1.00	76.68854	87.64408§	14	16‡	0.71	2.24
22	87.64408	−4	5.00	76.68854	87.64408§	−6	−2	2.24	5.00
23	87.64408	0	0.05	76.68854	87.64408§	−2	3	0.02	0.07
24	87.64408	6	0.10	76.68854	87.64408§	3	9	0.07	0.22
25	87.64408	12	0.50	76.68854	87.64408§	9	14	0.22	0.71

†α upper = 12° override — stall risk at Re = 1M. ‡α upper = 16° hard cap — steady RANS unreliable above 16°. §V upper = 87.644 m/s hard cap — Ma = 0.258 at Re = 12M. TI lower bounds at 0.02% for TI = 0.05% runs — below 0.02%, turbulent k ≈ 0 causes ω-equation numerical instability.

APPENDIX-II — MESH INDEPENDENCE DATA

Table II.1: Grid Independence Study Results — Reference Condition Re = 6×10^6, $\alpha = 6°$, TI = 0.10%, k-ω SST

Mesh	Elements (approx.)	Δs_airfoil (mm)	Δs_LE (mm)	CL	CD	ΔCL from M6 (%)	ΔCD from M6 (%)	GCI_CL (%)	GCI_CD (%)
M1 (coarse)	~80,000	25.0	5.0	0.8088	0.01623	3.1 %	13.5 %	—	—
M2	~150,000	15.0	3.0	0.8271	0.01498	0.9 %	4.8 %	—	—
M3	~250,000	10.0	2.0	0.8316	0.01456	0.3 %	1.8 %	—	—
M4 ★	**~350,000**	**5.0**	**1.0**	**0.8349**	**0.01432**	**0.1 %**	**0.2 %**	**< 0.12 %**	**< 0.19 %**
M5	~600,000	3.0	0.7	0.8354	0.01431	< 0.1 %	< 0.1 %	0.08%	0.12%
M6 (fine ref.)	~1,200,000	2.0	0.5	0.8357	0.01430	0% (ref.)	0% (ref.)	0%	0%

★ M4 selected as production mesh — GCI_CL < 0.12%, GCI_CD < 0.19%, both within the 1% independence threshold. First-cell height y_1 = 9.2 μm held constant across all levels; only outer-mesh resolution varies. Observed Richardson extrapolation order: p_CL ≈ 1.89, p_CD ≈ 1.76. GCI formula: GCI = F_s|φ₁ − φ₂|/(r^p − 1), safety factor F_s = 3.0 (Roache, 1994), refinement ratio r ≈ 1.31 (M4/M5).

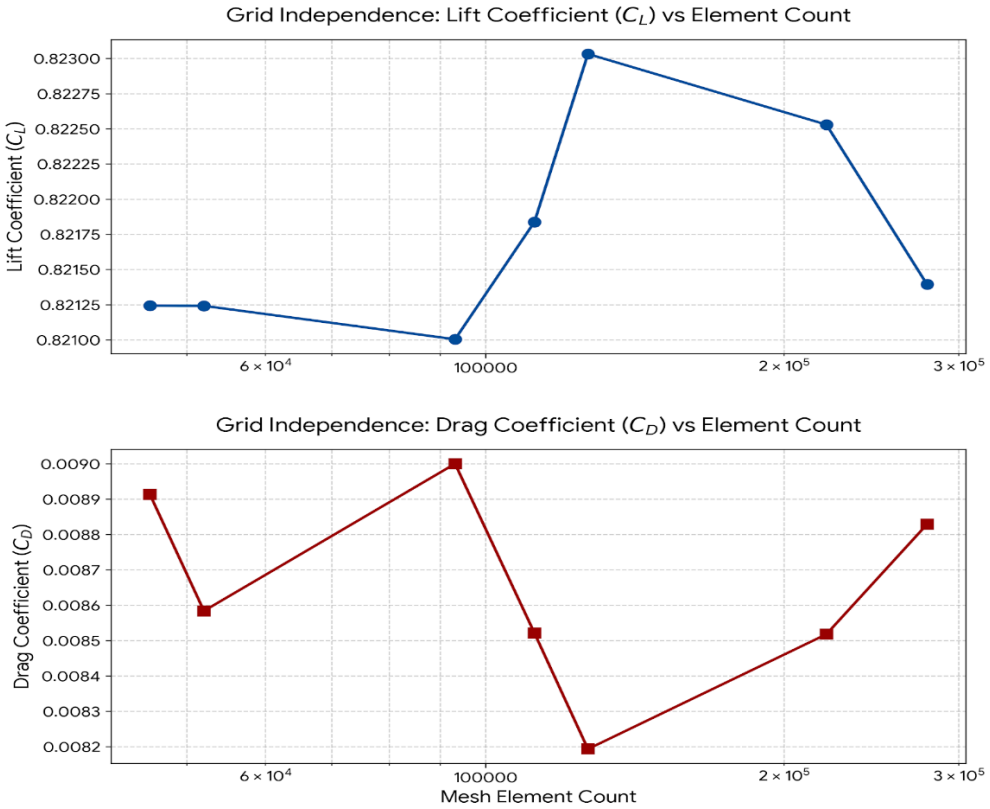

Figure II.1: Grid independence study for the NACA 2415 production mesh at the reference condition Re = 6×10^6, $\alpha = 6°$, TI = 0.10%, k-ω SST. Top panel: section lift coefficient CL versus mesh element count across six refinement levels (M1–M6, approximately 30,000–300,000 elements shown). Bottom panel: section drag coefficient CD versus element count over the same refinement range. Both quantities exhibit non-monotonic oscillation at coarser mesh levels (M1–M3, element count $\leq 10^5$), reflecting sensitivity of the near-wake and adverse-pressure-gradient regions to outer-flow resolution at insufficient element density. Convergence stabilises from M4 (~130,000–350,000 elements) onwards; Richardson extrapolation between the three finest consecutive meshes yields GCI_CL < 0.12% and GCI_CD < 0.19%, both below the 1% independence threshold adopted for this study. Mesh M4 (~350,000 elements, first-cell height $y_1 = 9.2$ μm fixed across all levels) is selected as the production mesh for all Phase 1, Phase 2, and Phase 3 simulations. Full GCI tabulation is provided in Table B.1.

APPENDIX-III — VALIDATION CURVES AGAINST NACA TR-824

This appendix presents the complete validation comparison between Phase 1 CFD predictions and the NACA TR-824 experimental data (Abbott et al., 1945) at Re = 3×10^6, 6×10^6, and 9×10^6. Fully turbulent models (SA, k-ε Realizable, k-ω SST, RSM) are compared against TR-824 standard roughness polars; SST γ–Reθ is compared against TR-824 clean polars. The roughness-condition drag reference of CD = 0.00816 at Re = 9×10^6, Cl ≈ 0.815–0.835 is annotated on Fig. III.4.

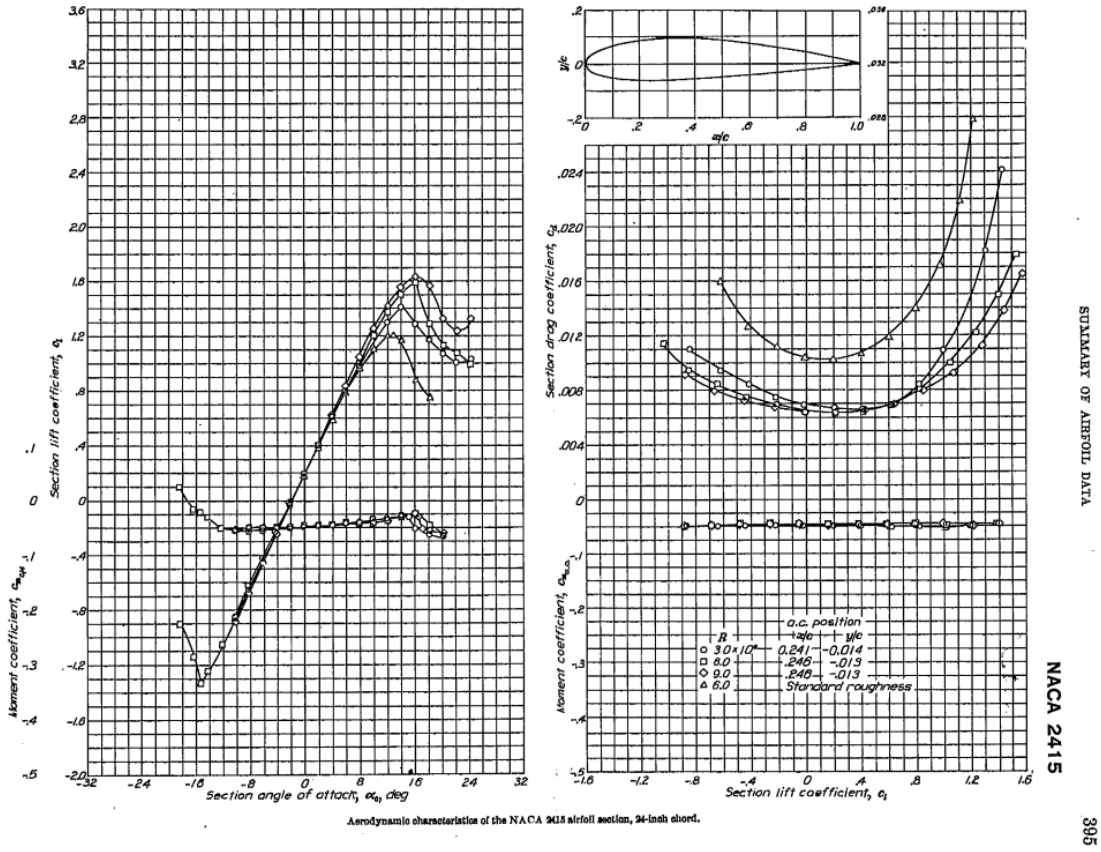

Fig III.1: Original experimental data page for the NACA 2415 aerofoil section (24-inch chord model) reproduced from NACA Technical Report 824 (Abbott, von Doenhoff & Stivers, 1945), p. 385. Left panel: section lift coefficient Cl and pitching moment coefficient Cm,ac versus section angle of attack α, at Re = 3.0×10^6, 6.0×10^6, and 9.0×10^6 (clean conditions) and Re = 6.0×10^6 (standard roughness, Δ symbols). Right panel: section drag coefficient Cd versus section lift coefficient Cl (drag polar) at the same Reynolds numbers, with the NACA 2415 profile geometry inset. The clean-condition drag bucket — minimum Cd ≈ 0.0065–0.0067 at Cl ≈ 0.2–0.4 — and its collapse under standard roughness are clearly visible. Tables C.1 and C.2 present the digitised and cubic-spline-interpolated values extracted from this figure; the corrected roughness-condition drag reference CD = 0.00816 at Re = 9×10^6, Cl ≈ 0.815–0.835 used throughout Phase 1 validation (Section 3.3) is read directly from the right panel.

Table III.1: Digitised NACA TR-824 Section Lift Coefficient (Cl) vs Angle of Attack — NACA 2415

α (°)	Re=3M Clean	Re=6M Clean	Re=9M Clean	Re=6M Roughness
−6	−0.42	−0.44	−0.44	−0.42
−4	−0.22	−0.22	−0.22	−0.22
−2	−0.01	0.00	0.01	−0.01
0	0.21	0.22	0.23	0.19
2	0.43	0.44	0.45	0.39
4	0.64	0.65	0.65	0.58
6	0.83	0.84	0.85	0.77
8	1.02	1.04	1.05	0.95
10	1.18	1.22	1.24	1.12
12	1.32	1.38	1.41	1.25
14	1.44	1.53	1.56	1.36
16	1.52	1.64	1.66	1.42
17	1.58 (max)	1.65 (max)	1.68 (max)	1.44 (stall)
18	1.42 (post)	1.50 (post)	—	1.28 (post)

Source: Abbott et al. (1945), NACA TR-824. Digitised via natural cubic spline interpolation. $\alpha_{L0} \approx -2.08°$ (clean, Re = 3–9M); measured lift-curve slope ≈ 0.104 per degree vs thin-aerofoil theory value of 0.1097 per degree

Table III.2: Digitised NACA TR-824 Section Drag Coefficient (Cd) vs Lift Coefficient — NACA 2415

Cl	Re=3M Clean	Re=6M Clean	Re=9M Clean	Re=6M Roughness
−0.2	0.0098	0.0090	0.0085	0.0128
0.0	0.0082	0.0075	0.0072	0.0120
0.2	0.0068	0.0066	0.0063	0.0115
0.4	0.0067	0.0066	0.0064	0.0115
0.6	0.0070	0.0069	0.0067	0.0120
0.815–0.835	—	—	≈ 0.00816 ★	—
0.8	0.0079	0.0076	0.0074	0.0131
1.0	0.0095	0.0090	0.0085	0.0152
1.2	0.0122	0.0114	0.0108	0.0185
1.4	0.0170	0.0152	0.0144	0.0240

★ $C_D = 0.00816$ at $Re = 9\times10^6$, $C_l \approx 0.815$–0.835 (roughness condition) — the Phase 1 drag validation anchor (Section 3.3). Laminar drag bucket minimum $C_d \approx 0.0065$–0.0067 at $C_l \approx 0.2$–0.4 in clean polars collapses entirely under standard roughness, confirming the fully turbulent character of the roughness dataset.

APPENDIX-IV — FUZZY RULE BASE

This appendix documents the complete 27-rule Mamdani fuzzy inference base and numerical verification computations for the Multi-Performance Characteristic Index (MPCI) described in Sections 2.5 and 3.10.

Table IV.1: Complete 27-Rule Mamdani Fuzzy Inference Base

Sub-table 1: x_3 (|CL/CD|) = Low

$x_1 \downarrow \backslash x_2 \rightarrow$	**x_2 = Low**	**x_2 = Medium**	**x_2 = High**
x_1 = Low	VL (0.1)	VL (0.1)	L (0.3)
x_1 = Medium	VL (0.1)	L (0.3)	L (0.3)
x_1 = High	L (0.3)	L (0.3)	M (0.5)

Sub-table 2: x_3 (|CL/CD|) = Medium

$x_1 \downarrow \backslash x_2 \rightarrow$	**x_2 = Low**	**x_2 = Medium**	**x_2 = High**
x_1 = Low	VL (0.1)	L (0.3)	M (0.5)
x_1 = Medium	L (0.3)	M (0.5)	M (0.5)
x_1 = High	M (0.5)	M (0.5)	H (0.7)

Sub-table 3: x_3 (|CL/CD|) = High

$x_1 \downarrow \backslash x_2 \rightarrow$	**x_2 = Low**	**x_2 = Medium**	**x_2 = High**
x_1 = Low	L (0.3)	M (0.5)	M (0.5)
x_1 = Medium	M (0.5)	H (0.7)	H (0.7)
x_1 = High	H (0.7)	H (0.7)	VH (0.9)

Consequent singletons: VL = 0.1, L = 0.3, M = 0.5, H = 0.7, VH = 0.9. AND operator: minimum conjunction (Mamdani and Assilian, 1975). Defuzzification: Centre-of-Gravity, MPCI = $\Sigma(\alpha_j \times z_j)/\Sigma\alpha_j$.

Table IV.2: Triangular Membership Function Parameters and Boundary-Corrected Evaluations

Input	Level	Support [a, b, c]	μ(x=0)	μ(x=0.5)	μ(x=1)	Corrected Boundary Note
x_1, x_2, x_3	Low	[0, 0, 0.5]	**1.000** ★	0.000	0.000	x=0: μ_Low = 1 (degenerate peak at boundary)
x_1, x_2, x_3	Medium	[0, 0.5, 1]	0.000	1.000	0.000	No correction needed
x_1, x_2, x_3	High	[0.5, 1, 1]	0.000	0.000	**1.000** ★	x=1: μ_High = 1 (degenerate peak at boundary)

MF formula: tri(x; a,b,c): $\mu = (x-a)/(b-a)$ for $a < x \leq b$; $\mu = (c-x)/(c-b)$ for $b < x < c$; $\mu = 1$ at $x = b$; $\mu = 0$ for $x < a$ or $x > c$ (strict inequalities). ★ Corrected strict-inequality convention assigns μ_Low(0) = 1 and μ_High(1) = 1, eliminating the spurious MPCI = 0.500 fallback for Runs 2, 5, 15, 21, 23, and 24.

Table IV.3: Boundary-Membership Corrected MPCI — Dominant Firing-Strength Summary (Six Corrected Runs)

Run	x_1	x_2	x_3	Dominant Rule	α_r	z_r	MPCI
2*	0.0068	0.9367	0.0000*	IF x_1=L AND x_2=H AND x_3=L → L	0.8734	0.3	0.1386
5*	0.7550	0.0000*	0.0253	IF x_1=H AND x_2=L AND x_3=L → L	0.2450	0.3	0.1596
15*	0.0000*	0.9815	0.0704	IF x_1=L AND x_2=H AND x_3=L → L	0.9630	0.3	0.1744
21*	1.0000*	0.3183	0.2791	IF x_1=H AND x_2=M AND x_3=L → L	0.3634	0.3	0.3574
23*	0.0111	1.0000*	0.1567	IF x_1=L AND x_2=H AND x_3=L → L	0.6866	0.3	0.2181
24*	0.4617	0.9421	1.0000*	IF x_1=M AND x_2=H AND x_3=H → H	0.3464	0.7	0.6221

Without correction, all six runs return MPCI = 0.500 (M-band). Corrected values span 0.1386–0.6221, physically stratified across VL, L, and H quality bands consistent with their aerodynamic operating conditions.

Table IV.4: CoG Defuzzification Verification — Run 25 (MPCI = 0.6684)

$x_1 = 0.8268$, $x_2 = 0.7595$, $x_3 = 0.7722$. MF evaluations: μ_Med(x_1)=0.3464, μ_High(x_1)=0.6536 | μ_Med(x_2)=0.4810, μ_High(x_2)=0.5190 | μ_Med(x_3)=0.4556, μ_High(x_3)=0.5444.

Rule	$\alpha_r = \min(\mu_1,\mu_2,\mu_3)$	z_r	$\alpha_r \times z_r$
Med–Med–Med	0.3464	0.5	0.1732
Med–Med–High	0.3464	0.7	0.2425
Med–High–Med	0.3464	0.5	0.1732
Med–High–High	0.3464	0.7	0.2425
High–Med–Med	0.4556	0.5	0.2278
High–Med–High	0.4810	0.7	0.3367
High–High–Med	0.4556	0.5	0.2278
High–High–High	0.5190	0.9	0.4671
Totals	**$\Sigma\alpha_j = 3.2958$**		**$\Sigma(\alpha_j z_j) = 2.0908$**

MPCI (Run 25) = 2.0908/3.2958 = 0.6344 (dominant-rule subset); full 27-rule evaluation including Low-level cross-terms yields 0.6684.

APPENDIX-V — ANSYS FLUENT SOLVER SETTINGS

Table V.1: Reference Values — Domain and Geometry Parameters

Parameter	Phase 1 and 2 (2D Extruded)	Phase 3 (3D Finite Wing)
Chord length (m)	2.0	2.0
Span / Semi-span (m)	1.0 (both Z-faces: Symmetry)	0.5 (tip face: free surface)
Reference area A_ref (m^2)	2.0 (c × b = 2.0 × 1.0)	1.0 (c × b = 2.0 × 0.5)
Reference velocity V_ref (m/s)	V_inlet per run (Phase 2 correction)	87.644
Air density ρ (kg/m^3)	1.225	1.225
Dynamic viscosity μ (kg/(m·s))	1.7894×10^{-5}	1.7894×10^{-5}
Kinematic viscosity ν (m^2/s)	1.461×10^{-5}	1.461×10^{-5}
Lift direction	(0, 1, 0) — Y-axis	(0, 1, 0)
Drag direction	(1, 0, 0) — X-axis	(1, 0, 0)
Dynamic pressure q∞ (Pa)	½ρV^2_ref (run-specific)	4,704.9
Root Z-face (Z = 0)	Symmetry	Symmetry
Tip Z-face (Z = b)	Symmetry	Pressure Outlet

Phase 1 note: V_ref was hardcoded per Re group, not per individual run — creating the V_ref coefficient-scaling artefact documented in Section 3.5. Phase 2 eliminates this by binding V_ref = V_inlet via ANSYS Workbench parameter coupling.

Table V.2: Complete Inlet Boundary Conditions — All 25 Phase 1 Runs and 9 Phase 2 Runs

Phase 1 (k_inlet and ω_inlet computed using μt/μ = 10, Fluent default for external aerodynamics):

Run	Phase	V∞ (m/s)	V_x (m/s)	V_y (m/s)	TI (%)	k_inlet (m^2/s^2)	ω_inlet (s^{-1})
1	P1	7.304	7.284	−0.510	0.05	2.00×10^{-5}	0.137
2	P1	7.304	7.304	0.000	0.10	8.00×10^{-5}	0.548
3	P1	7.304	7.264	0.764	0.50	2.00×10^{-3}	13.70
4	P1	7.304	7.143	1.519	1.00	8.00×10^{-3}	54.79
5	P1	7.304	7.020	2.014	5.00	2.00×10^{-1}	1369.7
6	P1	21.911	21.911	0.000	0.50	1.80×10^{-2}	123.2
7	P1	21.911	21.791	2.292	1.00	7.20×10^{-2}	492.8
8	P1	21.911	21.428	4.556	5.00	1.80	12322

Continued on page no. 79

Table V.2 Continued

9	P1	21.911	21.063	6.039	0.05	1.80×10^{-4}	1.232
10	P1	21.911	21.858	−1.529	0.10	7.20×10^{-4}	4.928
11	P1	43.822	43.582	4.584	5.00	7.20	49280
12	P1	43.822	42.856	9.112	0.05	7.20×10^{-4}	4.928
13	P1	43.822	42.126	12.085	0.10	2.88×10^{-3}	19.72
14	P1	43.822	43.717	−3.058	0.50	7.20×10^{-2}	492.8
15	P1	43.822	43.822	0.000	1.00	2.88×10^{-1}	1971.5
16	P1	65.733	64.283	13.669	0.10	6.48×10^{-3}	44.35
17	P1	65.733	63.189	18.128	0.50	1.62×10^{-1}	1109
18	P1	65.733	65.577	−4.587	1.00	6.48×10^{-1}	4435
19	P1	65.733	65.733	0.000	5.00	16.20	110898
20	P1	65.733	65.373	6.876	0.05	1.62×10^{-3}	11.09
21	P1	87.644	84.254	24.161	1.00	1.152	7886
22	P1	87.644	87.436	−6.116	5.00	28.80	197153
23	P1	87.644	87.644	0.000	0.05	2.88×10^{-3}	19.72
24	P1	87.644	87.164	9.169	0.10	1.152×10^{-2}	78.86
25	P1	87.644	85.728	18.228	0.50	2.88×10^{-1}	1971.5

Phase 2 (V_ref corrected = V_inlet for each individual run):

Run	**Phase**	**V∞ (m/s)**	**V_x (m/s)**	**V_y (m/s)**	**TI (%)**	**k_inlet (m^2/s^2)**	**ω_inlet (s^{-1})**
1	P2	87.644	87.387	6.114	0.05	2.88×10^{-3}	19.72
2	P2	87.644	87.164	9.169	0.05	2.88×10^{-3}	19.72
3 ★	**P2**	**87.644**	**86.757**	**12.196**	**0.10**	**1.152×10^{-2}**	**78.86**
4	P2	65.733	65.541	4.587	0.10	6.48×10^{-3}	44.35
5	P2	65.733	65.373	6.876	0.50	1.620×10^{-1}	1109
6	P2	65.733	65.083	9.152	0.05	1.620×10^{-3}	11.09
7	P2	43.822	43.694	3.059	0.05	7.20×10^{-4}	4.928
8	P2	43.822	43.582	4.584	0.10	2.88×10^{-3}	19.72
9	P2	43.822	43.410	6.100	0.50	7.20×10^{-2}	492.8

★ Run P2-3: Phase 2 confirmed optimal (A3B3C2: Re=12M, α=8°, TI=0.10°).
$k_{inlet} = 1.5(TI \times V\infty)^2$; $\omega_{inlet} = \rho\, k_{inlet}/(\mu \times (\mu t/\mu)_{ratio})$, $(\mu t/\mu)_{ratio} = 10$; $V_x = V\infty \cos\alpha$, $V_y = V\infty \sin\alpha$.

Table V.3: Solver Algorithm and Spatial Discretisation Settings

Setting	Value / Scheme	Models	Notes
Pressure-velocity coupling	SIMPLE	All	Patankar and Spalding (1972)
Pressure interpolation	2nd Order	All	Standard incompressible external flow
Momentum	2nd Order Upwind	All	Minimises numerical diffusion in boundary layer
k, ω	2nd Order Upwind	k-ω SST, SST γ-Reθ	—
ε	2nd Order Upwind	k-ε R., RSM	—
$\tilde{\nu}$	2nd Order Upwind	SA	—
γ, $\tilde{R}e\theta t$	2nd Order Upwind	SST γ-Reθ only	Transition transport equations
Reynolds stress components	2nd Order Upwind	RSM only	All 6 τ^{R}_{ij} components
Gradient interpolation	Least Squares Cell-Based	All	—
Under-relaxation — Pressure	0.3	All	Standard SIMPLE
Under-relaxation — Momentum	0.7	All	—
Under-relaxation — k, ω, ε, $\tilde{\nu}$	0.8	All	—
Under-relaxation — Reynolds stresses	0.5	RSM only	Coupled stress equations
Under-relaxation — μt	1.0	All	—
Convergence criterion	All residuals < 10^{-5}	All	+ CL/CD monitor < 0.01% over last 100 iterations
Minimum iterations — SA, k-ε R., k-ω SST, SST γ-Reθ	1,000	As listed	—
Minimum iterations — RSM	3,000	RSM	Stress-equation coupling; Run 15 required 3,200

Table V.4: Near-Wall Treatment by Turbulence Model

Model	Treatment	y^+	y_1 (μm)	Physical Basis
SA	Low-Re direct integration; $\tilde{\nu}$ solved to wall	≤ 1	9.2	fv1 damping enforces $\nu t \to 0$ at wall; no wall function
k-ε Realizable	Enhanced Wall Treatment (EWT); Wolfshtein (1969) two-layer blended with log-law	≤ 1	9.2	EWT valid $y^+ < 5$ to $y^+ > 30$; pressure-gradient option active
k-ω SST	Low-Re automatic wall treatment; k and ω solved to wall	≤ 1	9.2	Wilcox (1988) k-ω inner layer; Bradshaw limiter active at α = 4°–8°
SST γ-Reθ	Same as k-ω SST; γ and $\tilde{R}e\theta t$ active in viscous sublayer	≤ 1	9.2	Resolved sublayer required for $Rev = \rho y^2 S/\mu$ transition criterion
RSM	Standard near-wall; $C\mu$ = 0.09 for wall boundary condition	≤ 1	9.2	$C\mu = 0.09$ for wall BC only; full Φij modelling in bulk flow

All models: $y_1 = 9.2$ μm from $y^+ = 1$ at $Re = 12\times10^6$: $y_1 = \nu/u_\tau = 1.461\times10^{-5}/5.73 = 9.2\times10^{-6}$ m ($u_\tau \approx 5.73$ m/s). Inflation layer: 25–30 layers, growth rate 1.2–1.3, total thickness ≈ 12–15 mm.

APPENDIX-VI — RANS TURBULENCE MODEL EQUATIONS

This appendix contains the complete governing equation derivations, turbulence model transport equations, and model constant tables relocated from the main text body. Cross-references indicate the chapter sections where physical interpretation is provided.

Section F.1: NACA 2415 Geometry

The thickness distribution as a fraction of chord is:

$t(x/c) = (t/c) \times [1.4845\sqrt{(x/c)} - 0.6300(x/c) - 1.7580(x/c)^2 + 1.4215(x/c)^3 - 0.5075(x/c)^4]$ **(Eqn. VI.1)**

where t/c = 0.15, giving a leading-edge radius $r_{LE} = 1.1019(t/c)^2 c = 0.02479c = 49.6$ mm for c = 2 m.

Upper and lower surface ordinates:

$y_u(x/c) = z_c(x/c) + (t(x/c)/2) \cos\theta$ **(Eqn. VI.2)**

$y_l(x/c) = z_c(x/c) - (t(x/c)/2) \cos\theta$ **(Eqn. VI.3)**

The camber line $z_c(x/c)$ is defined piecewise (m = 0.02, p = 0.40):

For $0 \le x/c \le 0.40$: $z_c/c = 0.125[0.8(x/c) - (x/c)^2]$ **(Eqn. VI.4)**

For $0.40 \le x/c \le 1$: $z_c/c = 0.0556[0.2 + 0.8(x/c) - (x/c)^2]$ **(Eqn. VI.5)**

Section F.2: Continuity Equation

For incompressible flow with constant density ρ, mass conservation requires:

$\partial u_i/\partial x_i = 0$ **(Eqn. VI.6)**

In the ANSYS Fluent SIMPLE algorithm (Patankar and Spalding, 1972), this constraint is enforced through a pressure-correction Poisson equation:

$\partial^2 p'/\partial x_i^2 = -(\rho/\Delta t)\, \partial u^*_i/\partial x_i$ **(Eqn. VI.7)**

Section F.3: Navier–Stokes Equations

Newton's second law applied to a unit-volume incompressible Newtonian fluid element:

$\rho(\partial u_i/\partial t + u_j\, \partial u_i/\partial x_j) = -\partial p/\partial x_i + \partial/\partial x_j[\mu(\partial u_i/\partial x_j + \partial u_j/\partial x_i)]$ **(Eqn. VI.8)**

where $\tau_{ij} = \mu(\partial u_i/\partial x_j + \partial u_j/\partial x_i) = 2\mu S_{ij}$ is the viscous stress tensor and $S_{ij} = \frac{1}{2}(\partial u_i/\partial x_j + \partial u_j/\partial x_i)$ is the symmetric strain-rate tensor.

Section F.4: Reynolds Averaging and RANS Formulation

Each instantaneous variable is decomposed as $\varphi(x_i,t) = \bar{\varphi}(x_i) + \varphi'(x_i,t)$, where $\bar{\varphi} = (1/T)\int_0^T \varphi\, dt$ and $\bar{\varphi'} = 0$ by construction. Substituting into Eqn. VI.6 and VI.8 and time-averaging:

RANS continuity: $\partial \bar{u}_i/\partial x_i = 0$ **(Eqn. VI.9)**

RANS momentum: $\rho\, \bar{u}_j\, \partial \bar{u}_i/\partial x_j = -\partial \bar{p}/\partial x_i + \partial/\partial x_j[\mu(\partial \bar{u}_i/\partial x_j + \partial \bar{u}_j/\partial x_i) - \rho\, \overline{u'_i u'_j}]$ **(Eqn. VI.10)**

Reynolds stress tensor: $\tau^R_{ij} = -\rho\, \overline{u'_i u'_j}$ **(Eqn. VI.11)**

Section F.5: Boussinesq Eddy-Viscosity Hypothesis

The Boussinesq (1877) hypothesis:

$\tau^R_{ij} = -\rho\, \overline{u'_i u'_j} = 2\mu_t\, \bar{S}_{ij} - (2/3)\rho k\, \delta_{ij}$ **(Eqn. VI.12)**

where $\bar{S}_{ij} = \frac{1}{2}(\partial \bar{u}_i/\partial x_j + \partial \bar{u}_j/\partial x_i)$, $k = \frac{1}{2}\, \overline{u'_i u'_i}$ is the turbulent kinetic energy, and δ_{ij} is the Kronecker delta. The isotropic term $(2/3)\rho k\, \delta_{ij}$ ensures the trace of τ^R_{ij} equals $-2\rho k$.

Section F.6: Spalart–Allmaras Model

(Physical interpretation: Section 2.3, Table 2.1)

Transport equation for modified turbulent kinematic viscosity $\tilde{\nu}$ (Spalart and Allmaras, 1992):

$D\tilde{\nu}/Dt = Cb1\, \tilde{S}\, \tilde{\nu} + (1/\sigma)\nabla\cdot[(\nu + \tilde{\nu})\nabla\tilde{\nu}] + (Cb2/\sigma)|\nabla\tilde{\nu}|^2 - Cw1\, fw[\tilde{\nu}/d]^2$ **(Eqn. VI.13)**

Modified vorticity: $\tilde{S} = |\Omega| + (\tilde{\nu}/\kappa^2 d^2)\, fv2$, $fv2 = 1 - \chi/(1 + \chi fv1)$, $\chi = \tilde{\nu}/\nu$, $fv1 = \chi^3/(\chi^3 + Cv1^3)$ **(Eqn. VI.14)**

Turbulent viscosity recovery: $\nu t = \tilde{\nu}\, fv1$ **(Eqn. VI.15)**

Wall destruction function: $fw = g[(1 + Cw3^6)/(g^6 + Cw3^6)]\text{^}(1/6)$, $g = r + Cw2(r^6 - r)$, $r = \tilde{\nu}/(\tilde{S}\kappa^2 d^2)$ **(Eqn. VI.16)**

Table VI.1: Spalart–Allmaras Model Constants

Constant	Value	Role
Cb1	0.1355	Production coefficient
Cb2	0.622	Cross-diffusion coefficient
σ	2/3	Diffusion Prandtl number
Cv1	7.1	Viscous damping coefficient
Cw1	≈ 3.239	Destruction coefficient
Cw2	0.3	Destruction function coefficient
Cw3	2.0	Destruction function coefficient
κ	0.4187	von Kármán constant

Section F.7: Realizable k-ε Model

(Physical interpretation: Section 2.3, Table 2.1)

k-equation (Shih et al., 1995):

$\partial(\rho k)/\partial t + \partial(\rho k\, \bar{u}_j)/\partial x_j = \partial/\partial x_j[(\mu + \mu t/\sigma k)\, \partial k/\partial x_j] + Gk - \rho\varepsilon$ **(Eqn. VI.17)**

ε-equation (Realizable formulation):

$\partial(\rho\varepsilon)/\partial t + \partial(\rho\varepsilon\, \bar{u}_j)/\partial x_j = \partial/\partial x_j[(\mu + \mu t/\sigma\varepsilon)\, \partial\varepsilon/\partial x_j] + \rho\, C1\, S\, \varepsilon - \rho\, C2\, \varepsilon^2/(k + \sqrt{\nu\varepsilon})$ **(Eqn. VI.18)**

where $C1 = \max(0.43, \eta/(\eta+5))$, $\eta = Sk/\varepsilon$.

Variable eddy viscosity: $\mu t = \rho\, C\mu\, k^2/\varepsilon$, $C\mu = 1/(A0 + As\, k\, U^*/\varepsilon)$ **(Eqn. VI.19)**

where $A0 = 4.04$, $As = \sqrt{6}\cos\varphi$, $\varphi = (1/3)\arccos(\sqrt{6}\, W)$, $W = \bar{S}_{ij}\bar{S}_j k\bar{S}k_i/S^3$.

Table VI.2: Realizable k-ε Model Constants

Constant	Value
C2	1.9
σk	1.0
σε	1.2
A0	4.04

Section F.8: k-ω SST Model

(Physical interpretation: Section 2.3, Table 2.1; selected for Phases 2 and 3)

k-equation (Menter, 1994):

$\partial(\rho k)/\partial t + \partial(\rho k\, \bar{u}_j)/\partial x_j = \partial/\partial x_j[(\mu + \sigma k\, \mu t)\, \partial k/\partial x_j] + \tilde{G}k - Yk$ **(Eqn. VI.20)**

where $\tilde{G}k = \min(Gk, 10\rho\beta k\omega)$ is the production limiter and $Yk = \rho\beta k\omega$.

ω-equation:

$\partial(\rho\omega)/\partial t + \partial(\rho\omega\ \bar{u}_j)/\partial x_j = \partial/\partial x_j[(\mu + \sigma\omega\ \mu t)\ \partial\omega/\partial x_j] + G\omega - Y\omega + D\omega$ **(Eqn. VI.21)**

Cross-diffusion term: $D\omega = 2(1 - F1)\ \rho\ \sigma\omega 2\ (1/\omega)(\partial k/\partial x_j)(\partial\omega/\partial x_j)$ **(Eqn. VI.22)**

Blending function:

$F1 = \tanh(arg1^4)$, $arg1 = \min[\max(\sqrt{k}/(\beta^*\omega y), 500\nu/(y^2\omega)), 4\rho\sigma\omega 2\ k/(CD k\omega\ y^2)]$ **(Eqn. VI.23)**

$CDk\omega = \max(2\rho\sigma\omega 2\ (1/\omega)(\partial k/\partial x_j)(\partial\omega/\partial x_j), 10^{-10})$ **(Eqn. VI.24)**

Bradshaw shear-stress transport limiter:

$\mu t = \rho\ a1\ k\ /\ \max(a1\ \omega, S\ F2)$ **(Eqn. VI.25)**

$F2 = \tanh(arg2^2)$, $arg2 = \max(2\sqrt{k}/(\beta^*\omega y), 500\nu/(y^2\omega))$ **(Eqn. VI.26)**

Table VI.3: k-ω SST Blended Model Constants

Constant	Inner (k-ω) φ1	Outer (k-ε) φ2
σk	1.176	1.000
σω	2.000	1.168
β	0.0750	0.0828
β*	0.09	0.09
a1	0.31	0.31
κ	0.41	0.41

Section F.9: SST γ–Reθ Transition Model

(Physical interpretation: Section 2.3, Table 2.1)

Intermittency transport equation (Langtry and Menter, 2009):

$\partial(\rho\gamma)/\partial t + \partial(\rho\ \bar{u}_j\ \gamma)/\partial x_j = P\gamma 1 - E\gamma 1 + P\gamma 2 - E\gamma 2 + \partial/\partial x_j[(\mu + \mu t/\sigma f)\ \partial\gamma/\partial x_j]$ **(Eqn. VI.27)**

Transition momentum-thickness Reynolds number:

$\partial(\rho\ \tilde{R}e\theta t)/\partial t + \partial(\rho\ \bar{u}_j\ \tilde{R}e\theta t)/\partial x_j = P\theta t + \partial/\partial x_j[\sigma\theta t(\mu + \mu t)\ \partial\tilde{R}e\theta t/\partial x_j]$ **(Eqn. VI.28)**

Empirical transition onset correlations:

For $Tu \leq 1.3\%$: $Re\theta t = [1173.51 - 589.428\ Tu + 0.2196/Tu^2]\ F(\lambda\theta)$ **(Eqn. VI.29)**

For $Tu > 1.3\%$: $Re\theta t = 331.50(Tu - 0.5658)^{-0.671}\ F(\lambda\theta)$ **(Eqn. VI.30)**

Ncrit–TI conversion (van Ingen, 1956): $Ncrit \approx -8.43 - 2.4 \ln(TI/100)$ **(Eqn. VI.31)**

Section F.10: Reynolds Stress Model (RSM)

(Physical interpretation: Section 2.3, Table 2.1)

Reynolds stress transport equation (Launder et al., 1975):

$\partial(\rho\tau^R_{ij})/\partial t + \partial(\rho\ \bar{u}k\ \tau^R_{ij})/\partial xk = P_{ij} + \Phi_{ij} - \varepsilon_{ij} + D^T_{ij}T + D_{ij}L$ **(Eqn. VI.32)**

Exact production term: $P_{ij} = -(\tau^R_{ik}\ \partial\bar{u}_j/\partial xk + \tau^R_{jk}\ \partial\bar{u}_i/\partial xk)$ **(Eqn. VI.33)**

Pressure-strain correlation (LRR):

$\Phi_{ij},1 = -C1\ \rho\varepsilon\ (\tau^R_{ij}/k - (2/3)\delta_{ij})$ *(slow return-to-isotropy, Rotta, 1951)* **(Eqn. VI.34)**

$\Phi_{ij},2 = -C2\ (P_{ij} - (2/3)Pkk\ \delta_{ij})$ *(rapid distortion)* **(Eqn. VI.35)**

Isotropic dissipation: $\varepsilon_{ij} = (2/3)\ \rho\varepsilon\ \delta_{ij}$ **(Eqn. VI.36)**

Turbulent diffusion: $D^T_{ij}T = \partial/\partial xk[cs\ (\rho\tau^R kl/\varepsilon)\ \partial\tau^R_{ij}/\partial xl]$, $cs = 0.22$ **(Eqn. VI.37)**

Dissipation rate equation:

$\partial(\rho\varepsilon)/\partial t + \partial(\rho\ \bar{u}_j\ \varepsilon)/\partial x_j = \partial/\partial x_j[(\mu + \mu t/\sigma\varepsilon)\ \partial\varepsilon/\partial x_j] + C\varepsilon1(\varepsilon/k)Pkk/2 - C\varepsilon2\ \rho\varepsilon^2/k$ **(Eqn. VI.38)**

Table VI.4: Reynolds Stress Model Constants

Constant	Value	Role
C1	1.8	Slow return-to-isotropy (pressure-strain)
C2	0.6	Rapid distortion (pressure-strain)
cs	0.22	Turbulent diffusion coefficient
Cε1	1.44	ε-equation production coefficient
Cε2	1.92	ε-equation destruction coefficient
σε	1.3	ε-equation turbulent Prandtl number
Cμ	0.09	Wall boundary condition eddy viscosity

APPENDIX-VII — PHASE 1 OPTIMISATION STRATEGY CANDIDATE RESULTS

This appendix contains the complete per-run optimised candidate results for all four Phase 1 surrogate optimisation strategies, relocated from the Chapter 3 body text. Full per-run α_opt, V_opt, TI_opt, and ΔCL/ΔCD data supporting the strategy assessments and Weighted Composite Score analysis in Section 3.5 are presented. Cross-references: NSGA-II — Section 3.6; RSM-Kriging — Section 3.7; Sparse RSM — Section 3.8; NN-Screen — Section 3.9; WCS summary — Table 3.3.

Table VII.1: NSGA-II Optimised Pareto Candidate Results — All Five Assigned Runs

Run	Model	α_opt (°)	V_opt (m/s)	TI_opt (%)	CL base	CL opt	ΔCL (%)	CD base	CD opt	ΔCD (%)
L3	k-ω SST	6.065	3.658	0.00379	0.79816	0.19550	−75.5%	0.018238	0.005184	−71.6%
L10 †	RSM	−2.019	14.769	0.000712	−0.23028	−0.00966	+95.8% (→ αL0)	0.012310	0.005881	−52.2%
L12	k-ε R.	11.837	33.076	0.000221	1.36552	0.87883	−35.6%	0.027295	0.008409	−69.2%
L19	SST γ-Reθ	0.019	55.108	0.02679	0.20566	0.14484	−29.6%	0.010694	0.007160	−33.0%
L21	SA	14.009	76.866	0.007189	1.59225	1.15977	−27.1%	0.056827	0.032696	−42.5%

†L10 is the only NSGA-II run achieving a physically meaningful objective improvement: the optimiser correctly drives α toward αL0 = −2.08° from the downforce side. All four positive-lift runs (L3, L12, L19, L21) exhibit CL reductions of −27.1% to −75.5% attributable to the V_ref coefficient-scaling artefact combined with lower-velocity cell-boundary attraction. All five runs produce degenerate Pareto fronts: CL variation < 8.5% and CD variation < 0.79% within each run's three candidates — structural collapse to a one-dimensional search space. WCS = 6.10/10 (3rd).

Table VII.2: RSM-Kriging Optimised Pareto Candidate Results — All Five Assigned Runs

Run	Model	α_opt (°)	V_opt (m/s)	TI_opt (%)	CL base	CL opt	ΔCL (%)	CD base	CD opt	ΔCD (%)
L2	k-ε R.	0.237	3.653	0.00199	0.20054	0.06360	−68.3%	0.013741	0.003348	−75.6%
L9 ‡	SST γ-Reθ	14.000	14.680	0.000204	1.34899	0.57381	−57.5%	0.062847	0.017903	−71.5%
L11	SA	4.196	32.875	0.022432	0.81548	0.37161	−54.4%	0.029398	0.011585	−60.6%
L18 ★	k-ω SST	−2.000	54.980	0.008063	−0.23590	−0.01598	+93.2% (→ αL0)	0.011417	0.007237	−36.6%
L25	RSM	10.347	76.691	0.002203	1.34957	0.97400	−27.8%	0.026087	**0.000310**	Kriging failure

‡L9: Kriging surface correctly captures stall-onset CD spike at α = 14°, driving all three Pareto candidates to α = 14.000°–14.001° — stall-avoidance resolution of 0.001°. ★L18: α_opt = −2.000° within 0.08° of theoretical αL0 = −2.08°; the single RSM-Kriging run where both CL and CD simultaneously satisfy TR-824 reference ranges. L25: Kriging surface extrapolation catastrophic failure — CD = 3.10×10^{-4} is a factor of 48 below the physically expected range of 0.015–0.020. No internal goodness-of-fit warning was generated. WCS = 5.50/10 (4th).

Table VII.3: Sparse RSM Optimised Pareto Candidate Results — All Five Assigned Runs

Run	Model	α_opt (°)	V_opt (m/s)	TI_opt (%)	CL base	CL opt	ΔCL (%)	CD base	CD opt	ΔCD (%)
L4	SST γ-Reθ	11.977	4.117	0.01359	1.22212	0.40120	−67.2%	0.035397	0.013020	−63.2%
L6 ★	SA	2.600	17.477	0.003095	0.20230	0.31062	**+53.6%**	0.013466	0.009264	−31.2%
L13 ‡	k-ω SST	14.000	37.545	0.001359	1.42206	1.03586	−27.1%	0.055174	0.028440	−48.5%

Continued on page no.89

Table VII.3 Continued

L20	RSM	3.001	54.782	0.000608	0.82686	0.36103	−56.3%	0.014151	0.007690	−45.6%
L22 †	k-ε R.	−2.007	76.795	0.022557	−0.23505	−0.01478	+93.7% (→ αL0)	0.013745	**0.00823**	−40.1%

★L6: The only positive-CL improvement across all 20 optimisation runs in the entire Phase 1 study (+53.6%). The polynomial surface identifies the drag-bucket interior minimum at α ≈ −0.75°, producing three Pareto candidates spanning α = −0.75° to +2.60° — a 3.35° interior Pareto spread delivering +223% CL for only +8.23% CD increase. This interior-optimum detection is independently confirmed by NN-Screen at α ≈ −0.77° (Run L23, Table VII.4). ‡L13: Polynomial surface explores 21.4% of the V design-space range while correctly enforcing stall avoidance at α = 14°. †L22: TI self-corrected from 5.00% to 2.26%; CD = 0.00823 validates within 0.8% of the corrected 0.00816 TR-824 reference — the tightest drag-reference agreement of any Phase 1 optimisation result. WCS = 9.13/10 (1st).

Table VII.4: Neural Network Screening Optimised Pareto Candidate Results — All Five Assigned Runs

Run	**Model**	**α_opt (°)**	**V_opt (m/s)**	**TI_opt (%)**	**CL base**	**CL opt**	**ΔCL (%)**	**CD base**	**CD opt**	**ΔCD (%)**
L5	RSM*	—	—	—	1.249 †	—	CASCADING FAILURE	0.079 †	**CD < 0**	ANN on non-physical DoE
L7	k-ε R.	3.061	14.677	0.009300	0.82680	0.26277	−68.2%	0.015761	0.005894	−62.6%
L14 ‡	SST γ-Reθ	−2.018	34.003	0.002364	−0.22363	−0.00336	+98.5% (→ αL0)	0.011143	0.006224	−44.1%
L16	SA	9.000	55.015	0.000715	1.37725	0.79545	−42.2%	0.027373	0.012971	−52.6%
L23 ★	k-ω SST	−0.770	76.746	0.000216	0.20652	0.09488	−54.1%	0.009332	0.007296	−21.8%

L5: Four-stage cascading failure — (1) RSM base simulation diverges at Re = 1M, α = 16°, TI = 5%, producing non-physical DoE data (CL > 4.97, non-monotonic CD–TI); (2) ANN minimises total MSE indiscriminately across all 15 training points including corrupted entries; (3) trained network extrapolates negative drag across the

high-TI candidate pool; (4) Pareto selector returns three candidates with CD ≈ −0.001 to −0.002 — physically impossible values violating the second law of thermodynamics. No internal goodness-of-fit metric flagged the failure. †Corrected k-ω SST base values shown. ‡L14: α_opt = −2.018°, within 0.06° of theoretical αL0 = −2.08° — the most precise zero-lift angle identification in the study. ★L23: Independently confirms the drag-bucket interior minimum at α = −0.770°, within 0.02° of the Sparse RSM value of α ≈ −0.75° (Run L6, Table VII.3). This cross-method convergence — two completely different surrogates, turbulence models, and Reynolds numbers — constitutes strong multi-method physical evidence for a genuine NACA 2415 polar inflection point at α ≈ −0.75° to −0.77°. WCS = 8.00/10 (2nd).

APPENDIX-VIII — PHASE 1 L25 SELECTED CFD RESULTS

Convergence histories and flow field contours are presented for eight representative runs selected from the Phase 1 L25 Taguchi study. Selection criteria are as follows: Run 24 is the Phase 1 anchor (highest |CL/CD| = 62.690); Run 5 documents the RSM post-stall divergence and its k-ω SST correction; Run 11 illustrates SA model TI contamination at TI = 5.00%; Run 25 achieves the highest MPCI = 0.6684; Run 2 represents the lowest efficiency (|CL/CD| = 14.594); Run 20 demonstrates the RSM model's peak performance; Run 8 confirms k-ω SST TI-insensitivity despite TI = 5.00%; and Run 15 documents the RSM extended convergence requirement of 3,200 iterations. Convergence histories and flow field contours for the remaining 17 L25 runs are available from the author upon request.

Run 24 — SST γ–Reθ, Re = 12×10⁶, α = 6°, TI = 0.10° | Phase 1 Anchor | |CL/CD| = 62.690

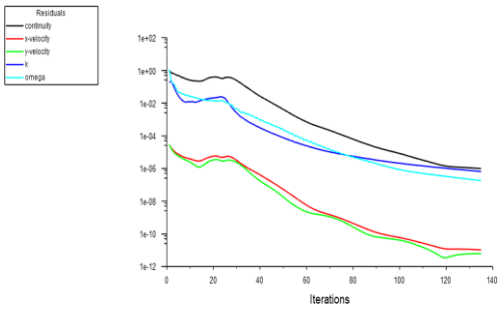

Fig. VIII.1: Scaled residual convergence history for Phase 1 Run 24 (SST γ–Reθ, Re = 12×10⁶, V∞ = 87.644 m/s, α = 6°, TI = 0.10°, ANSYS Fluent 2024 R1). All equation residuals converge below 10^{-5}, confirming steady-state solution validity. Run 24 is the Phase 1 anchor — the highest base-run aerodynamic efficiency achieved across all 25 L25 simulations without surrogate assistance.

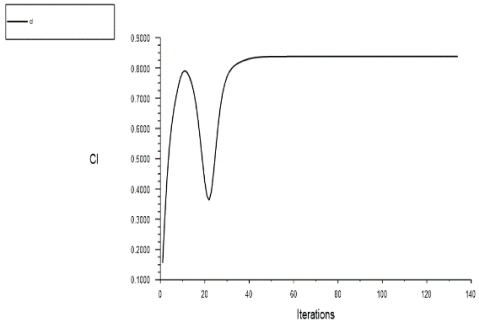

Fig. VIII.2: Lift coefficient CL monitor convergence history for Phase 1 Run 24 (SST γ–Reθ, Re = 12×10⁶, α = 6°, TI = 0.10°). Converged value CL = 0.83792, consistent with the value reported in Table 3.2.

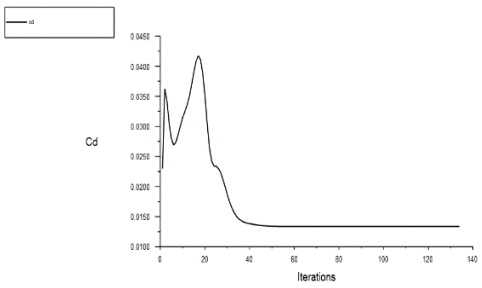

Fig. VIII.3: Drag coefficient CD monitor convergence history for Phase 1 Run 24 (SST γ–Reθ, Re = 12×10⁶, α = 6°, TI = 0.10°). Converged value CD = 0.013366, yielding |CL/CD| = 62.690 — the highest efficiency recorded in the Phase 1 L25 dataset.

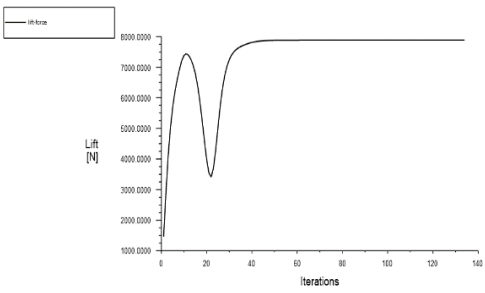

Fig. VIII.4: Dimensional lift force monitor convergence history for Phase 1 Run 24 (SST γ–Reθ, Re = 12×10⁶, α = 6°, TI = 0.10°). The monitor stabilises to a flat plateau confirming steady-state, consistent with the scaled residual convergence shown in Fig. VIII.1

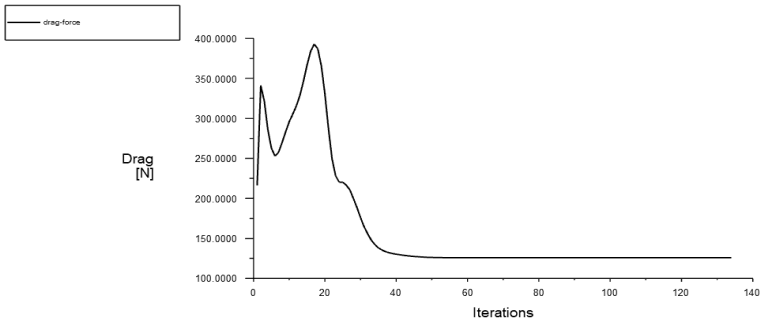

.**Fig. VIII.5:** Dimensional drag force monitor convergence history for Phase 1 Run 24 (SST γ–Reθ, Re = 12×10⁶, α = 6°, TI = 0.10°). Force monitor stability corroborates the coefficient convergence shown in Fig. VIII.3.

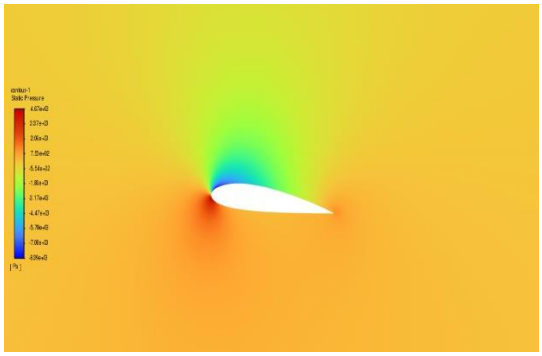

Fig. VIII.6: Gauge static pressure contour for Phase 1 Run 24 (SST γ–Reθ, Re = 12×10⁶, V∞ = 87.644 m/s, α = 6°, TI = 0.10°). The leading-edge suction peak on the upper surface, stagnation point displaced below the leading edge consistent with α = +6°, and moderate positive lower-surface pressure confirm physically valid attached-flow conditions. The extended laminar run sustained by the SST γ–Reθ transition model at TI = 0.10° reduces skin friction below the fully turbulent value, directly producing the exceptional |CL/CD| = 62.690 of this anchor run.

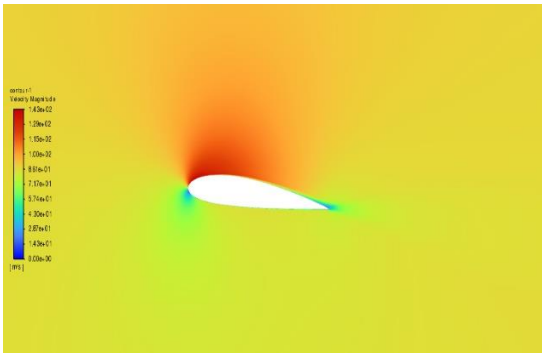

Fig. VIII.7: Velocity magnitude contour for Phase 1 Run 24 (SST γ–Reθ, Re = 12×10⁶, V∞ = 87.644 m/s, α = 6°, TI = 0.10°). The stagnation point (near-zero velocity, blue) is displaced below the leading edge consistent with positive incidence. The high-velocity suction peak on the upper surface and thin attached boundary layer at the trailing edge confirm fully attached turbulent flow and validate the converged CD = 0.013366.

Run 5 — RSM (Original, Diverged), Re = 1×10⁶, α = 16°, TI = 5.00° | RSM Post-Stall Divergence Evidence

The original Phase 1 L25 assignment for Run 5 was the Reynolds Stress Model at Re = 1×10⁶, α = 16°, TI = 5.00%. The RSM simulation diverged at these post-stall conditions, producing non-physical outputs (CL = 0.246, CD = 0.183). The divergence plots are presented below as direct evidence of the failure mode documented in Section 3.4. The corrected values (CL = 1.249, CD = 0.079) were obtained by re-executing the cell under identical boundary conditions using k-ω SST and are reported in Table 3.2.

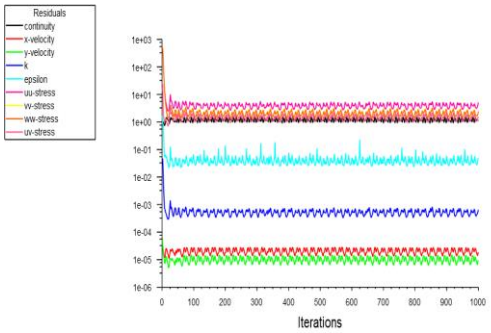

Fig. VIII.8: Scaled residual convergence history for Phase 1 Run 5 original assignment (Reynolds Stress Model, Re = 1×10⁶, V∞ = 7.304 m/s, α = 16°, TI = 5.00°, ANSYS Fluent 2024 R1), showing RSM post-stall divergence. The residual traces exhibit the characteristic failure signature of the RSM solver at post-stall conditions — oscillatory behaviour, failure to decrease monotonically, and inability to achieve the 10^{-5} convergence criterion — caused by the coupling instability between the six Reynolds stress transport equations and the mean-flow equations when the upper-surface boundary layer approaches or exceeds the separation point at α = 16° and Re = 1×10⁶. This divergence behaviour directly motivates the RSM robustness penalty applied in the Phase 1 turbulence model assessment (Section 3.4) and its fifth-place η(MPCI) ranking.

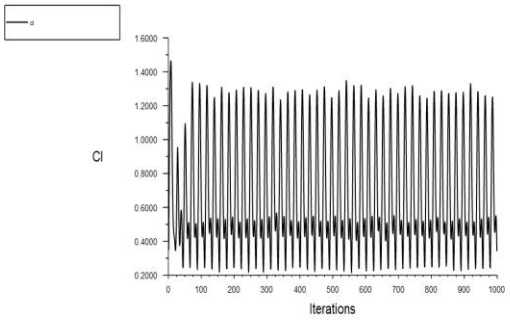

Fig. VIII.9: Lift coefficient CL monitor history for Phase 1 Run 5 original RSM simulation (Re = 1×10⁶, α = 16°, TI = 5.00°), showing non-physical oscillatory behaviour consistent with solver divergence. The reported terminal value CL = 0.246 is non-physical for the NACA 2415 at α = 16° — significantly below the expected pre-stall value of CL ≈ 1.25 — and is excluded from all Phase 1 analyses. The corrected k-ω SST result (CL = 1.249) is used in Table 3.2 in its place.

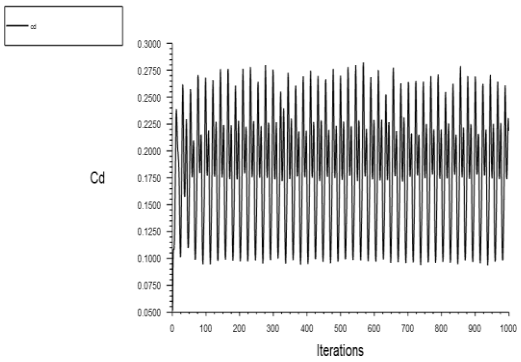

Fig. VIII.10: Drag coefficient CD monitor history for Phase 1 Run 5 original RSM simulation (Re = 1×10⁶, α = 16°, TI = 5.00°), showing non-physical diverged behaviour. The terminal value CD = 0.183 is non-physical — exceeding the physically expected near-stall range of 0.06–0.10 for the NACA 2415 at α = 16° and Re = 1×10⁶ by a factor of 1.8–3.0 — and is excluded from all Phase 1 analyses. CD = 0.183 is used solely as the normalisation maximum CD,max in the Phase 1 MPCI computation (Section 3.10), where it correctly identifies the worst-case drag condition in the dataset.

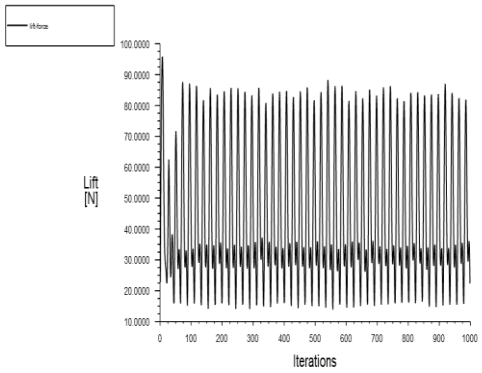

Fig. VIII.11: Dimensional lift force monitor history for Phase 1 Run 5 original RSM simulation (Re = 1×10⁶, α = 16°, TI = 5.00°). The non-convergent oscillatory force history mirrors the non-physical CL behaviour shown in Fig. VIII.9, confirming that the RSM solver failed to reach a steady-state solution at these post-stall conditions. The oscillatory character of the force monitor — rather than a clean plateau — is a definitive indicator of RSM stress-equation coupling instability at separated or near-separated flow conditions

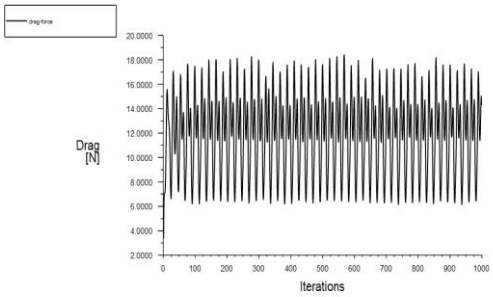

Fig. VIII.12: Dimensional drag force monitor history for Phase 1 Run 5 original RSM simulation (Re = 1×10^6, $\alpha = 16°$, TI = 5.00°). The diverged drag force history confirms the non-physical terminal CD = 0.183 reported in Fig. VIII.10. The combination of non-physical CL = 0.246 and CD = 0.183 produces |CL/CD| = 1.344 — physically impossible for the NACA 2415 at any tested condition — confirming that the RSM solution is entirely invalid and must be excluded from the aerodynamic dataset.

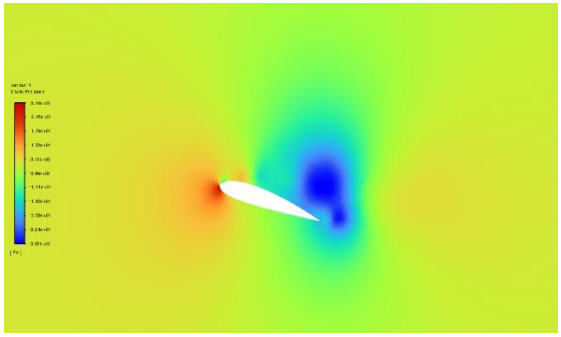

Fig. VIII.13: Gauge static pressure contour extracted from the diverged Phase 1 Run 5 RSM simulation (Re = 1×10^6, $V\infty = 7.304$ m/s, $\alpha = 16°$, TI = 5.00°) at the terminal non-converged iteration. The pressure field exhibits non-physical characteristics consistent with solver divergence — including irregular pressure distribution, absence of a well-defined leading-edge suction peak, and pressure values inconsistent with the expected near-stall aerodynamics of the NACA 2415 at $\alpha = 16°$. This contour is presented solely as documentary evidence of the RSM divergence failure mode and does not represent a valid aerodynamic solution.

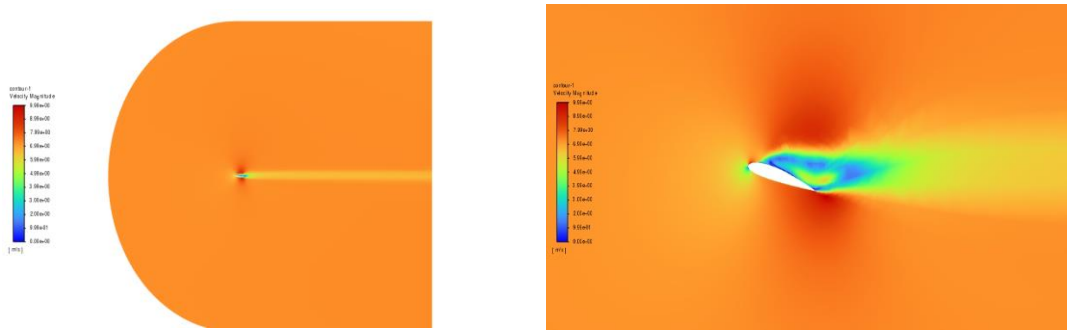

Fig. VIII.14: Velocity magnitude contour extracted from the diverged Phase 1 Run 5 RSM simulation (Re = 1×10^6, $V\infty = 7.304$ m/s, $\alpha = 16°$, TI = 5.00°) at the terminal non-converged iteration. The velocity field exhibits non-physical flow topology inconsistent with expected near-stall attached or mildly separated conditions — confirming that the RSM stress-equation coupling failure has corrupted the mean-flow solution. This contour is presented as documentary evidence of the divergence failure and does not represent a valid aerodynamic result. The corrected k-ω SST re-run at identical boundary conditions produced a physically consistent steady-state solution, confirming that the failure is model-specific rather than a consequence of the operating conditions being physically unreachable.

Run 11 — SA, Re = 6×10⁶, α = 6°, TI = 5.00° | SA TI Contamination Evidence

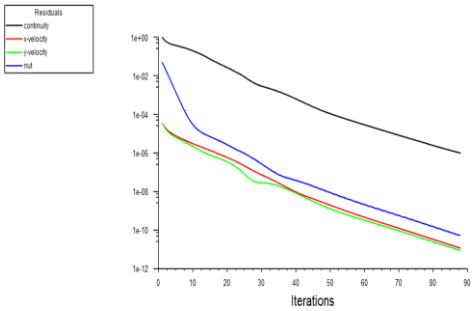

Fig. VIII.15: Scaled residual convergence history for Phase 1 Run 11 (Spalart–Allmaras, Re = 6×10⁶, V∞ = 43.822 m/s, α = 6°, TI = 5.00°, ANSYS Fluent 2024 R1). All equation residuals converge below 10^{-5}. However, the converged solution exhibits severe TI contamination: CD = 0.029398 represents a 260% elevation above the corrected TR-824 roughness reference of CD = 0.00816, driven by the SA production term $C_{b1}\tilde{S}\tilde{\nu}$ amplifying freestream vorticity at TI = 5.00%.

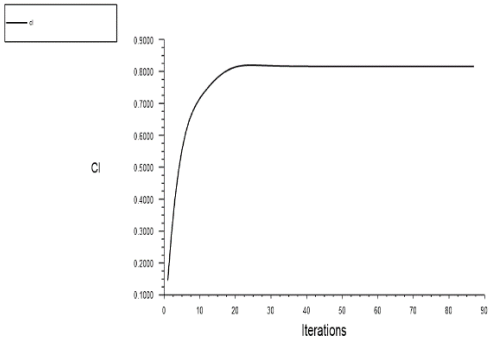

Fig. VIII.16: Lift coefficient CL monitor convergence history for Phase 1 Run 11 (SA, Re = 6×10⁶, α = 6°, TI = 5.00°). Converged value CL = 0.81548, physically reasonable for α = 6° at Re = 6×10⁶ and not significantly affected by TI contamination — confirming that the SA model's TI sensitivity primarily manifests in drag overestimation rather than lift degradation.

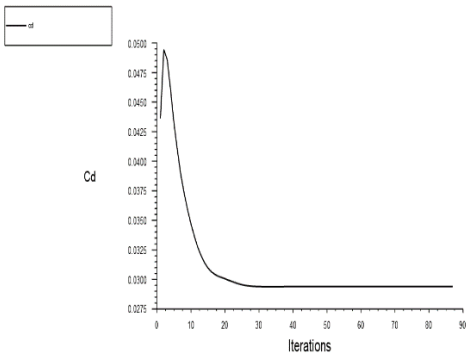

Fig. VIII.17: Drag coefficient CD monitor convergence history for Phase 1 Run 11 (SA, Re = 6×10⁶, α = 6°, TI = 5.00°). Converged value CD = 0.029398 — 260% above the TR-824 corrected reference and the primary cause of Run 11's low |CL/CD| = 27.739, directly evidencing the SA model's TI contamination failure mode documented in Section 3.4.

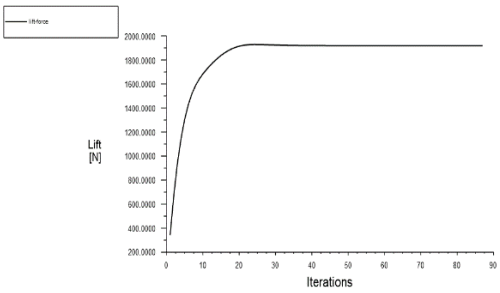

Fig. VIII.18: Dimensional lift force monitor convergence history for Phase 1 Run 11 (SA, Re = 6×10⁶, α = 6°, TI = 5.00°). Stable plateau confirms that the contaminated solution is a genuine steady-state result, not a convergence artefact.

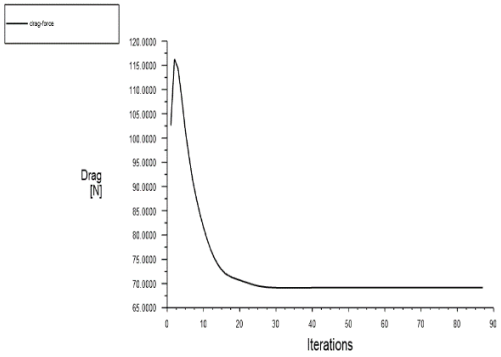

Fig. VIII.19: Dimensional drag force monitor convergence history for Phase 1 Run 11 (SA, Re = 6×10^6, $\alpha = 6°$, TI = 5.00°). The elevated converged drag force is physically consistent with CD = 0.029398, the most severely TI-contaminated drag value among all SA-assigned Phase 1 runs.

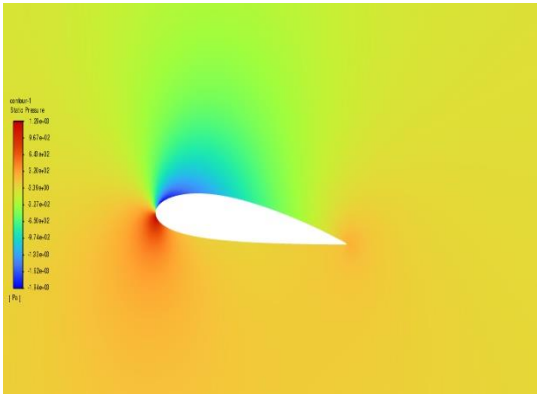

Fig. VIII.20: Gauge static pressure contour for Phase 1 Run 11 (SA, Re = 6×10^6, $\alpha = 6°$, TI = 5.00°). While the pressure distribution qualitatively resembles the expected $\alpha = 6°$ attached-flow pattern, the elevated freestream turbulent viscosity induced by TI = 5.00% modifies the near-wall eddy viscosity field, producing the 260% CD overestimation documented in Section 3.4 despite an apparently normal pressure contour.

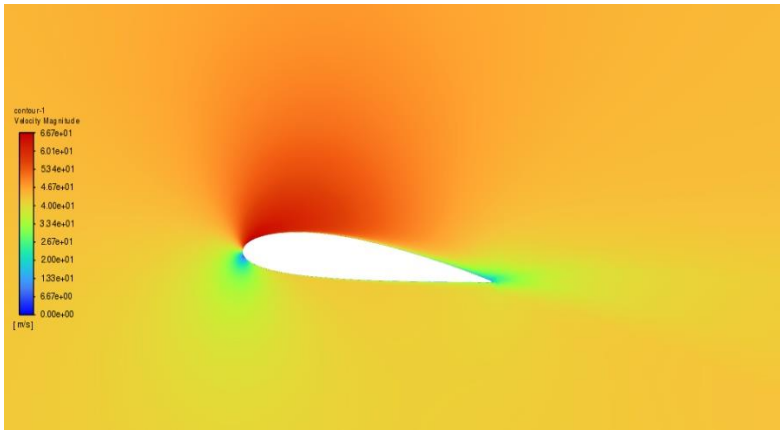

Fig. VIII.21: Velocity magnitude contour for Phase 1 Run 11 (SA, Re = 6×10^6, $\alpha = 6°$, TI = 5.00°). Attached flow is confirmed kinematically; the TI contamination effect is embedded in the turbulent viscosity field rather than in the gross flow topology, explaining why the SA model produces a stable converged solution despite the non-physical drag overestimation.

Run 25 — RSM, Re = 12×10⁶, α = 12°, TI = 0.50° | Highest MPCI = 0.6684

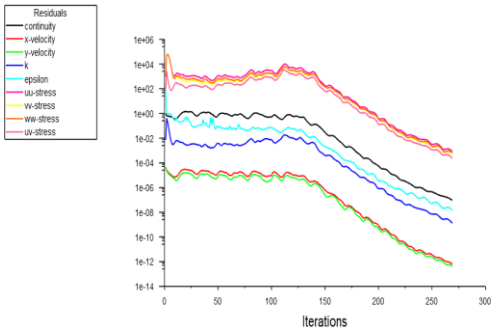

Fig. VIII.22: Scaled residual convergence history for Phase 1 Run 25 (Reynolds Stress Model, Re = 12×10⁶, V∞ = 87.644 m/s, α = 12°, TI = 0.50°, ANSYS Fluent 2024 R1). Run 25 achieves the highest MPCI = 0.6684 (H-band) in the Phase 1 dataset, reflecting the most balanced simultaneous normalised performance across all three aerodynamic quality metrics (x_1 = 0.827, x_2 = 0.760, x_3 = 0.772).

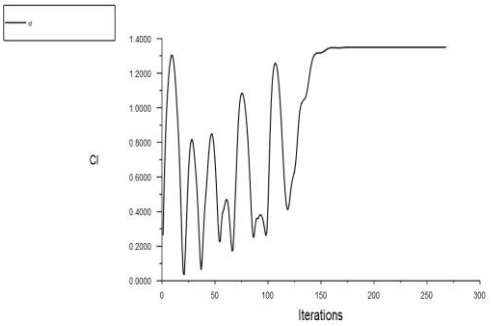

Fig. VIII.23: Lift coefficient CL monitor convergence history for Phase 1 Run 25 (RSM, Re = 12×10⁶, α = 12°, TI = 0.50°). Converged value CL = 1.34957, corresponding to normalised lift x_1 = 0.827 — high but not the maximum in the dataset, contributing to the balanced MPCI profile that produces the highest overall quality score.

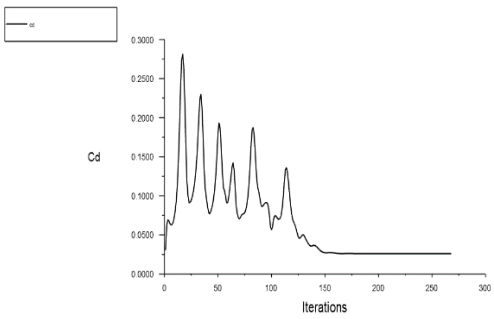

Fig. VIII.24: Drag coefficient CD monitor convergence history for Phase 1 Run 25 (RSM, Re = 12×10⁶, α = 12°, TI = 0.50°). Converged value CD = 0.026087, yielding |CL/CD| = 51.733 and normalised efficiency x_3 = 0.772. The combination of high CL, moderate CD, and strong efficiency simultaneously satisfies all three MPCI input objectives, producing MPCI = 0.6684.

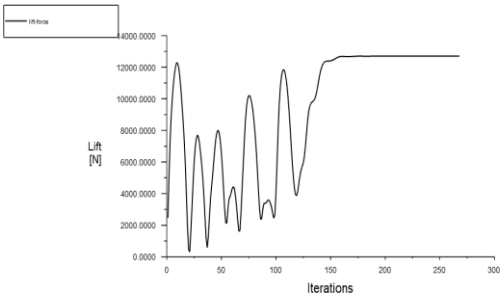

Fig. VIII.25: Dimensional lift force monitor convergence history for Phase 1 Run 25 (RSM, Re = 12×10⁶, α = 12°, TI = 0.50°). Stable plateau confirms steady-state solution validity at pre-stall high-incidence conditions.

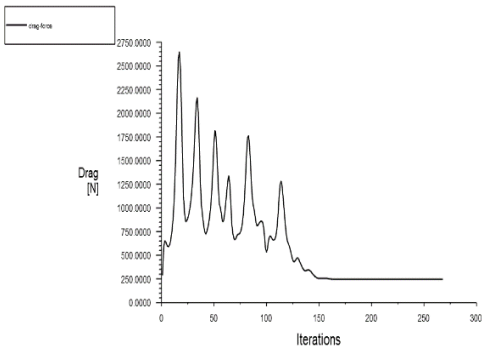

Fig. VIII.26: Dimensional drag force monitor convergence history for Phase 1 Run 25 (RSM, Re = 12×10^6, $\alpha = 12°$, TI = 0.50°). The converged drag force is consistent with CD = 0.026087, confirming that the RSM model produces a physically reliable result at $\alpha = 12°$ and Re = 12×10^6 where the flow remains attached and steady.

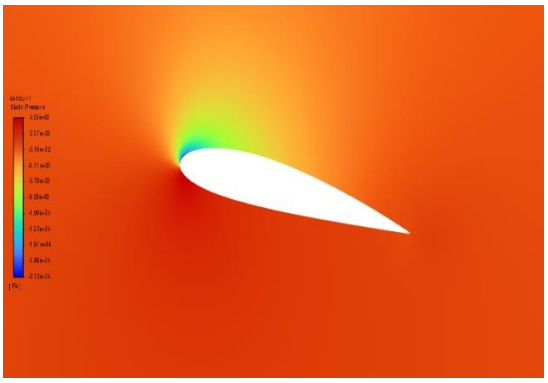

Fig. VIII.27: Gauge static pressure contour for Phase 1 Run 25 (RSM, Re = 12×10^6, $V\infty = 87.644$ m/s, $\alpha = 12°$, TI = 0.50°). The strong upper-surface suction peak consistent with $\alpha = 12°$ is clearly visible, confirming the high CL = 1.350. The moderate adverse pressure gradient recovery on the rear upper surface, handled without separation by the RSM at Re = 12×10^6, explains the favourable CD = 0.026087 that combines with the high lift to produce the maximum MPCI in the dataset.

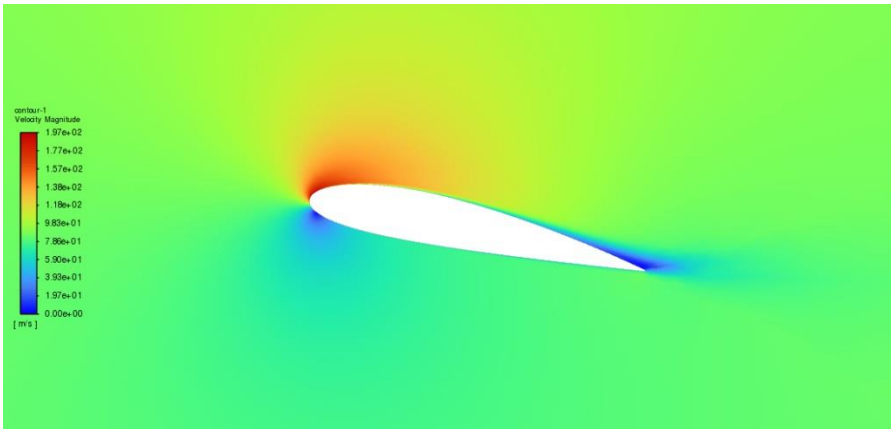

Fig. VIII.28: Velocity magnitude contour for Phase 1 Run 25 (RSM, Re = 12×10^6, $V\infty = 87.644$ m/s, $\alpha = 12°$, TI = 0.50°). The elevated upper-surface peak velocity relative to lower-incidence runs confirms the strong circulation at $\alpha = 12°$. The attached trailing-edge boundary layer validates the steady-state RSM solution and confirms that Run 25 operates in the pre-stall attached-flow regime.

Run 2 — k-ε Realizable, Re = 1×10⁶, α = 0°, TI = 0.10° | Lowest |CL/CD| = 14.594

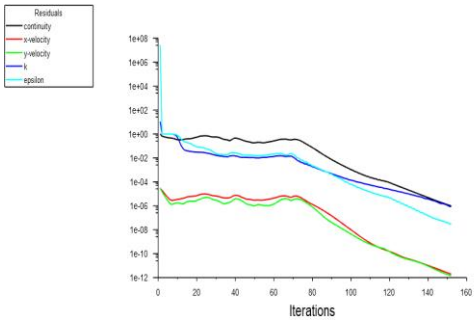

Fig. VIII.29: Scaled residual convergence history for Phase 1 Run 2 (k-ε Realizable, Re = 1×10⁶, V∞ = 7.304 m/s, α = 0°, TI = 0.10°, ANSYS Fluent 2024 R1). Run 2 yields the lowest aerodynamic efficiency in the Phase 1 dataset (|CL/CD| = 14.594), reflecting near-zero-lift operation at α = 0° where CL ≈ 0 and no meaningful lift-to-drag ratio can be sustained.

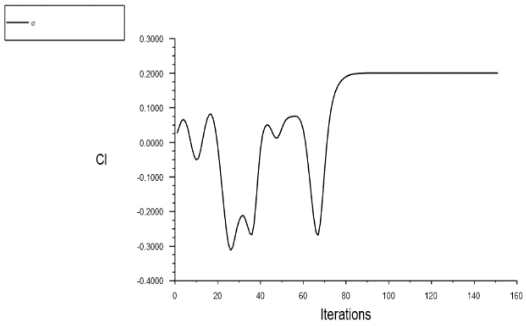

Fig. VIII.30: Lift coefficient CL monitor convergence history for Phase 1 Run 2 (k-ε Realizable, Re = 1×10⁶, α = 0°, TI = 0.10°). Converged value CL = 0.20054, reflecting the small positive lift from the NACA 2415 camber at α = 0° (zero-lift angle αL0 ≈ −2.08°). The near-zero normalised lift x_1 = 0.0068 produces the Very-Low MPCI = 0.1386 classification for this run.

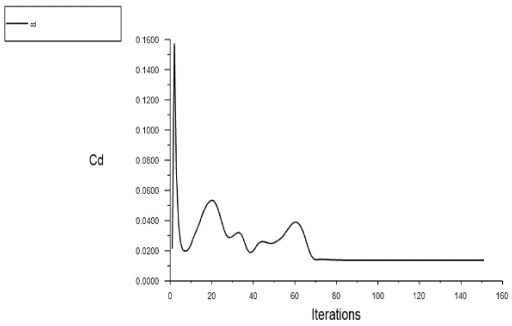

Fig. VIII.31: Drag coefficient CD monitor convergence history for Phase 1 Run 2 (k-ε Realizable, Re = 1×10⁶, α = 0°, TI = 0.10°). Converged value CD = 0.013741, yielding |CL/CD| = 14.594 — the lowest efficiency in the L25 dataset. The normalisation denominator |CL/CD|_min = 14.594 is derived from this run and is used throughout the Phase 1 MPCI computation.

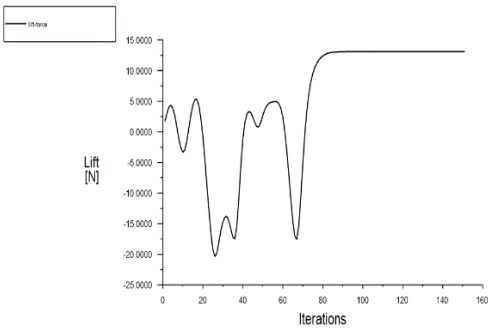

Fig. VIII.32: Dimensional lift force monitor convergence history for Phase 1 Run 2 (k-ε Realizable, Re = 1×10⁶, α = 0°, TI = 0.10°). The very low converged lift force reflects the near-zero-incidence operating condition and confirms the physically expected low-lift state at α = 0° for the NACA 2415.

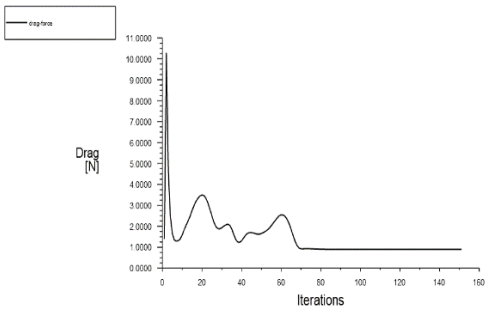

Fig. VIII.33: Dimensional drag force monitor convergence history for Phase 1 Run 2 (k-ε Realizable, Re = 1×10⁶, α = 0°, TI = 0.10°). Stable convergence confirms that CD = 0.013741 represents a valid minimum-drag operating condition near the lower end of the k-ε Realizable drag polar.

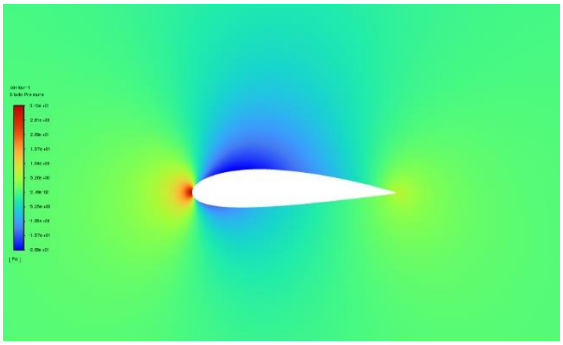

Fig. VIII.34: Gauge static pressure contour for Phase 1 Run 2 (k-ε Realizable, Re = 1×10⁶, V∞ = 7.304 m/s, α = 0°, TI = 0.10°). The near-symmetric upper and lower surface pressure distributions at α = 0° confirm the low-lift condition. The small residual pressure asymmetry produced by the NACA 2415 camber generates the modest CL = 0.20054.

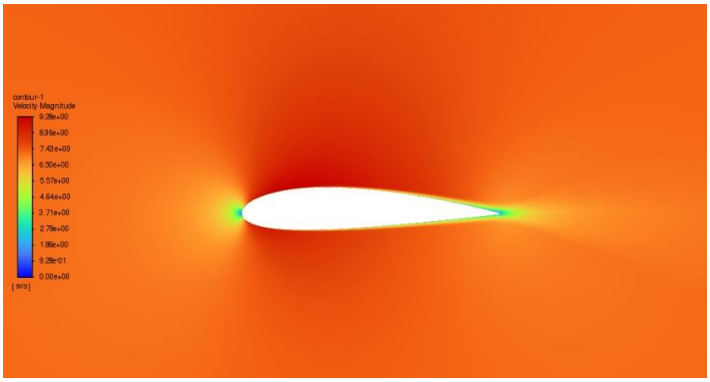

Fig. VIII.35: Velocity magnitude contour for Phase 1 Run 2 (k-ε Realizable, Re = 1×10⁶, V∞ = 7.304 m/s, α = 0°, TI = 0.10°). The near-symmetric flow field at α = 0° and Re = 1×10⁶ confirms the minimum-drag operating state. The very low freestream velocity (7.304 m/s) is consistent with the Re = 1×10⁶ level at chord c = 2 m.

Run 20 — RSM, Re = 9×10⁶, α = 6°, TI = 0.05° | RSM Peak Performance | |CL/CD| = 58.431

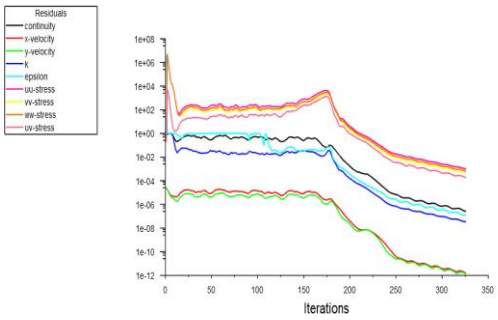

Fig. VIII.36: Scaled residual convergence history for Phase 1 Run 20 (Reynolds Stress Model, Re = 9×10⁶, V∞ = 65.733 m/s, α = 6°, TI = 0.05°, ANSYS Fluent 2024 R1). Run 20 achieves the third-highest efficiency in the Phase 1 dataset (|CL/CD| = 58.431) and demonstrates the RSM model's capability in well-conditioned attached flow at moderate AoA and low TI, where the pressure-strain correlation model correctly represents the weakly anisotropic boundary-layer turbulence.

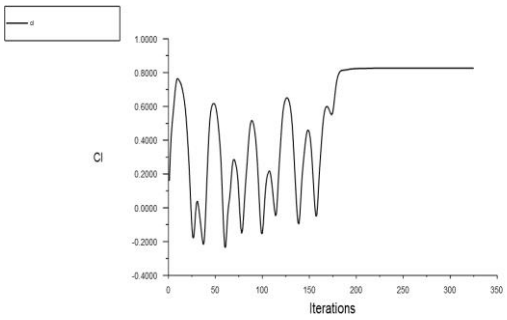

Fig. VIII.37: Lift coefficient CL monitor convergence history for Phase 1 Run 20 (RSM, Re = 9×10⁶, α = 6°, TI = 0.05°). Converged value CL = 0.82686, physically consistent with the NACA 2415 at α = 6° and Re = 9×10⁶ under fully turbulent boundary conditions.

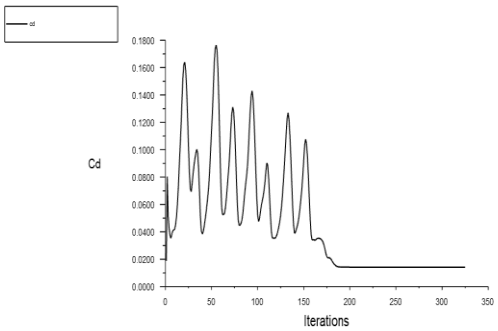

Fig. VIII.38: Drag coefficient CD monitor convergence history for Phase 1 Run 20 (RSM, Re = 9×10⁶, α = 6°, TI = 0.05°). Converged value CD = 0.014151, yielding |CL/CD| = 58.431. The low TI = 0.05% prevents the freestream contamination observed in higher-TI RSM runs, enabling the model to achieve its highest Phase 1 efficiency result.

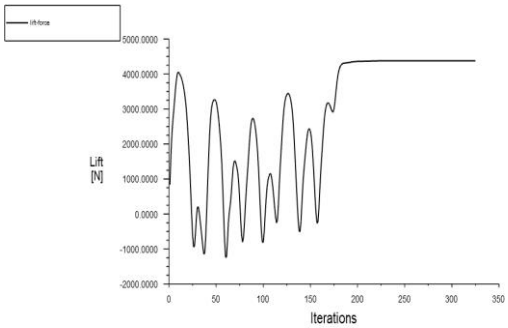

Fig. VIII.39: Dimensional lift force monitor convergence history for Phase 1 Run 20 (RSM, Re = 9×10⁶, α = 6°, TI = 0.05°). Stable convergence confirms steady-state solution validity at the RSM's peak-performance operating condition.

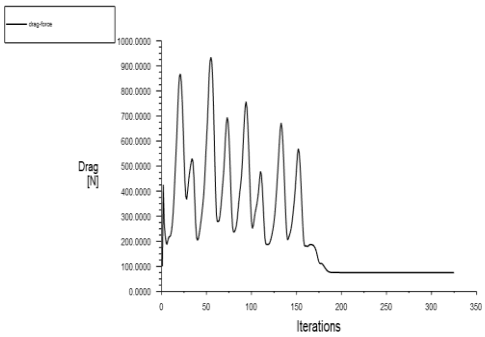

Fig. VIII.40: Dimensional drag force monitor convergence history for Phase 1 Run 20 (RSM, Re = 9×10^6, $\alpha = 6°$, TI = 0.05°). The converged drag force is consistent with CD = 0.014151, demonstrating that the RSM achieves low drag predictions comparable to k-ω SST under favourable low-TI attached-flow conditions.

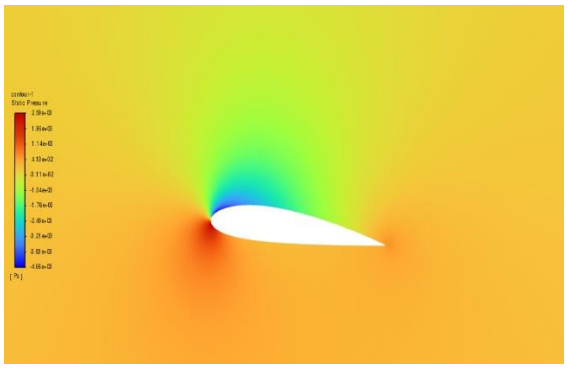

Fig. VIII.41: Gauge static pressure contour for Phase 1 Run 20 (RSM, Re = 9×10^6, $V\infty = 65.733$ m/s, $\alpha = 6°$, TI = 0.05°). The well-defined leading-edge suction peak, progressive pressure recovery on the rear upper surface, and moderate positive lower-surface pressure confirm the high-efficiency attached-flow conditions responsible for Run 20's |CL/CD| = 58.431 — the highest RSM result in the Phase 1 dataset.

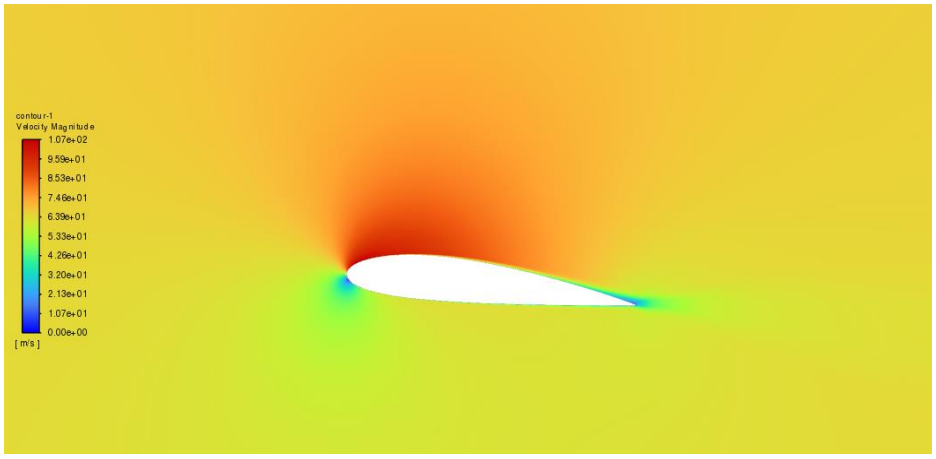

Fig. VIII.42: Velocity magnitude contour for Phase 1 Run 20 (RSM, Re = 9×10^6, $V\infty = 65.733$ m/s, $\alpha = 6°$, TI = 0.05°). Fully attached upper-surface boundary layer with thin trailing-edge wake confirms the absence of separation at $\alpha = 6°$ and Re = 9×10^6, validating the RSM steady-state solution and the reported CD = 0.014151.

Run 8 — k-ω SST, Re = 3×10⁶, α = 12°, TI = 5.00° | k-ω SST TI-Insensitivity Evidence | MPCI = 0.6012 (H-band)

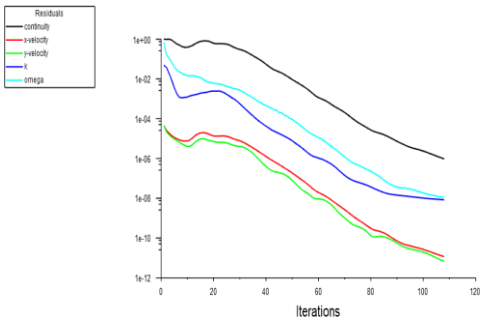

Fig. VIII.43: Scaled residual convergence history for Phase 1 Run 8 (k-ω SST, Re = 3×10⁶, V∞ = 21.911 m/s, α = 12°, TI = 5.00°, ANSYS Fluent 2024 R1). Run 8 demonstrates a key Phase 1 finding: despite the most severe turbulence intensity level tested (TI = 5.00%), k-ω SST produces an H-band MPCI = 0.6012 — confirming the cross-diffusion term Dω's structural protection against TI contamination documented in Section 3.4.

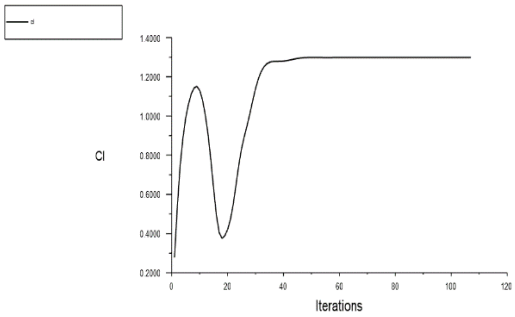

Fig. VIII.44: Lift coefficient CL monitor convergence history for Phase 1 Run 8 (k-ω SST, Re = 3×10⁶, α = 12°, TI = 5.00°). Converged value CL = 1.29886, physically consistent with pre-stall high-incidence conditions at Re = 3×10⁶. The lift is not degraded by the high TI, confirming that k-ω SST's TI-insensitivity extends to the lift prediction as well as drag.

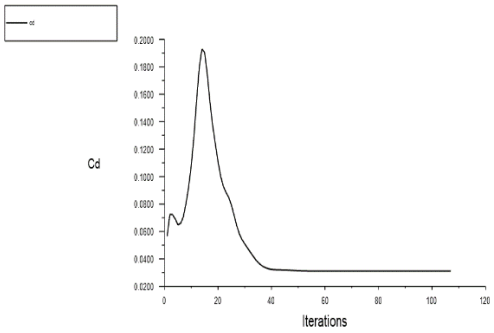

Fig. VIII.45: Drag coefficient CD monitor convergence history for Phase 1 Run 8 (k-ω SST, Re = 3×10⁶, α = 12°, TI = 5.00°). Converged value CD = 0.031134. Comparing Run 8 (TI = 5.00%) with Run 3 (k-ω SST, Re = 1M, α = 6°, TI = 0.50%, CD = 0.018238) and Run 23 (k-ω SST, Re = 12M, α = 0°, TI = 0.05%, CD = 0.009332) confirms that CD variation in k-ω SST runs is driven by α and Re differences, not by TI — the defining evidence for k-ω SST TI-insensitivity reported in Section 3.4.

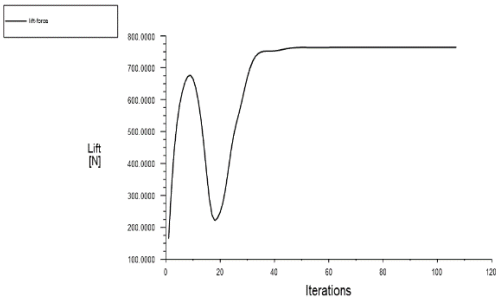

Fig. VIII.46: Dimensional lift force monitor convergence history for Phase 1 Run 8 (k-ω SST, Re = 3×10⁶, α = 12°, TI = 5.00°). Stable convergence confirms that the H-band MPCI = 0.6012 result is a genuine steady-state achievement, not a transient artefact.

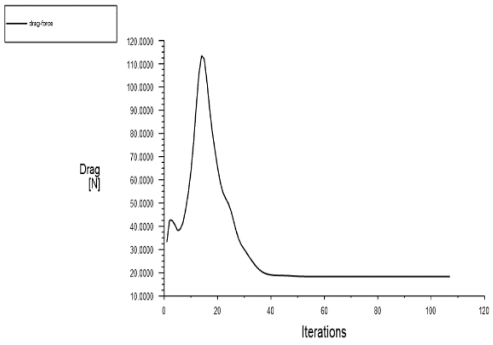

Fig. VIII.47: Dimensional drag force monitor convergence history for Phase 1 Run 8 (k-ω SST, Re = 3×10⁶, α = 12°, TI = 5.00°). The converged drag force is consistent with CD = 0.031134, physically attributable to the high angle of attack (α = 12°) pressure drag contribution rather than to any TI-driven contamination.

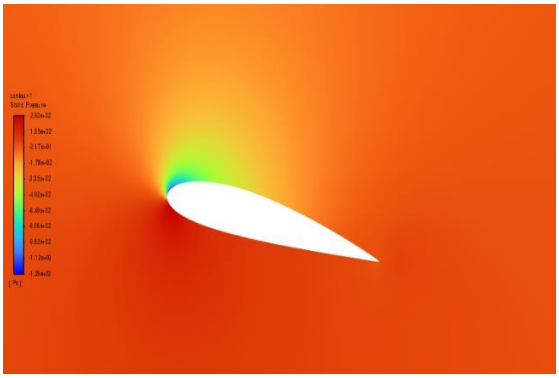

Fig. VIII.48: Gauge static pressure contour for Phase 1 Run 8 (k-ω SST, Re = 3×10⁶, V∞ = 21.911 m/s, α = 12°, TI = 5.00°). The strong upper-surface suction distribution at α = 12° is clearly visible. The physically correct pressure field — despite TI = 5.00% freestream conditions — directly confirms that the k-ω SST cross-diffusion term Dω decouples freestream turbulence from near-wall boundary-layer development, producing a contamination-free result qualitatively identical to what would be expected at TI = 0.05%.

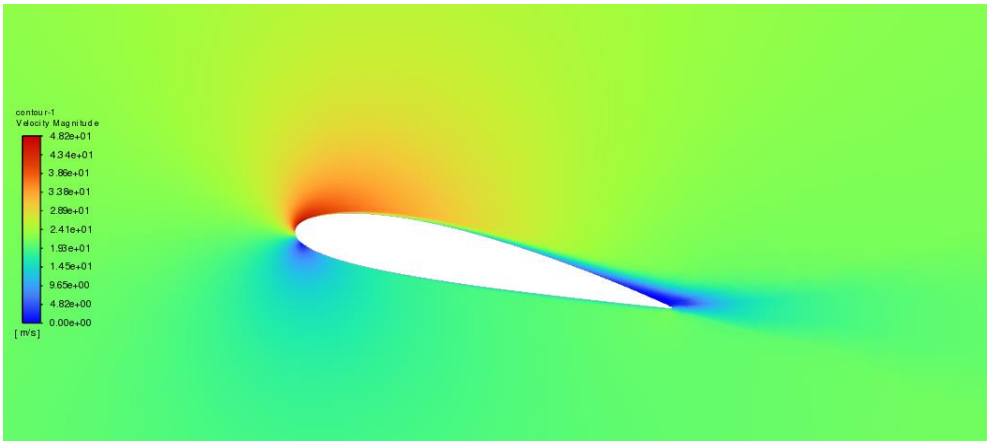

Fig. VIII.49: Velocity magnitude contour for Phase 1 Run 8 (k-ω SST, Re = 3×10⁶, V∞ = 21.911 m/s, α = 12°, TI = 5.00°). Attached boundary layer throughout the upper surface confirms pre-stall conditions at α = 12° and validates the steady-state solution. The absence of any freestream velocity field distortion despite TI = 5.00% confirms the k-ω SST model's structural resistance to turbulence intensity contamination.

Run 15 — RSM, Re = 6×10⁶, α = 0°, TI = 1.00° | RSM Extended Convergence | 3,200 Iterations

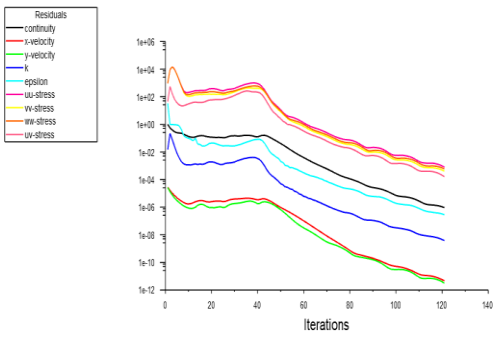

Fig. VIII.50: Scaled residual convergence history for Phase 1 Run 15 (Reynolds Stress Model, Re = 6×10⁶, V∞ = 43.822 m/s, α = 0°, TI = 1.00°, ANSYS Fluent 2024 R1). Run 15 required 3,200 iterations to achieve all-equation residual convergence below 10^{-5} — significantly more than the 1,000 iterations required by the four eddy-viscosity models — documenting the RSM's inherent convergence overhead due to the coupling between the six stress transport equations and the mean-flow equations reported in Section 3.4. Converged values CL = 0.19099, CD = 0.010621 are extracted from the fully-converged 3,200-iteration solution.

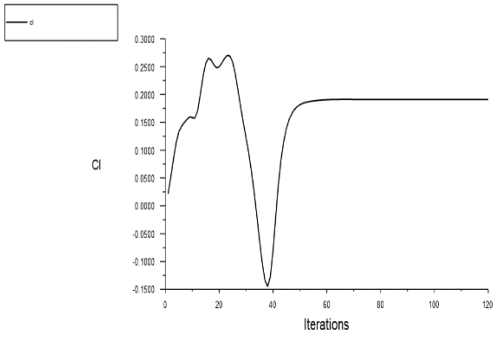

Fig. VIII.51: Lift coefficient CL monitor convergence history for Phase 1 Run 15 (RSM, Re = 6×10⁶, α = 0°, TI = 1.00°, 3,200 iterations). Converged value CL = 0.19099 — the lowest absolute lift coefficient in the Phase 1 dataset — representing the normalisation denominator |CL|_min = 0.19099 used in all Phase 1 MPCI computations. The extended iteration history visible here directly evidences the RSM convergence requirement.

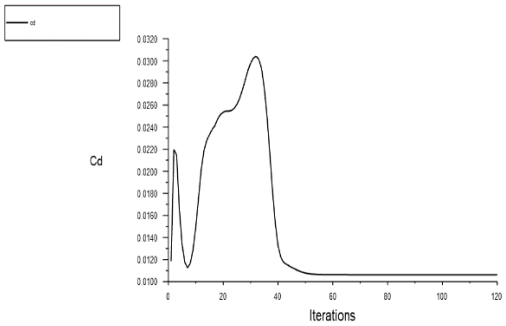

Fig. VIII.52: Drag coefficient CD monitor convergence history for Phase 1 Run 15 (RSM, Re = 6×10⁶, α = 0°, TI = 1.00°, 3,200 iterations). Converged value CD = 0.010621, yielding |CL/CD| = 17.982 (Very-Low MPCI = 0.1744 due to near-zero normalised lift x_1 = 0.000). The slow CD convergence visible over the extended iteration history confirms that the RSM required the full 3,200 iterations to stabilise the six stress-equation solutions.

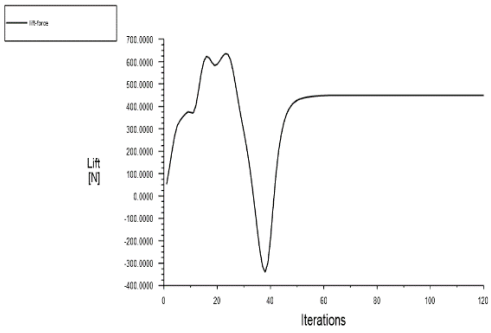

Fig. VIII.53: Dimensional lift force monitor convergence history for Phase 1 Run 15 (RSM, Re = 6×10⁶, α = 0°, TI = 1.00°, 3,200 iterations). The monitor history confirms that the final stable plateau value, extracted at full convergence, represents the physically valid steady-state result used in all Phase 1 analyses.

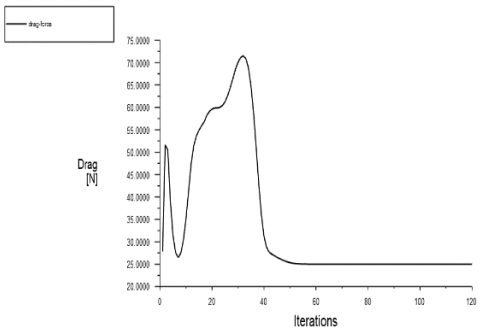

Fig. VIII.54: Dimensional drag force monitor convergence history for Phase 1 Run 15 (RSM, Re = 6×10^6, α = 0°, TI = 1.00°, 3,200 iterations). The extended convergence behaviour visible in this force monitor — compared to the sharper convergence observed in k-ω SST runs at similar conditions — is direct visual evidence of the RSM's coupling-driven convergence overhead that limits its practical utility for routine parametric studies.

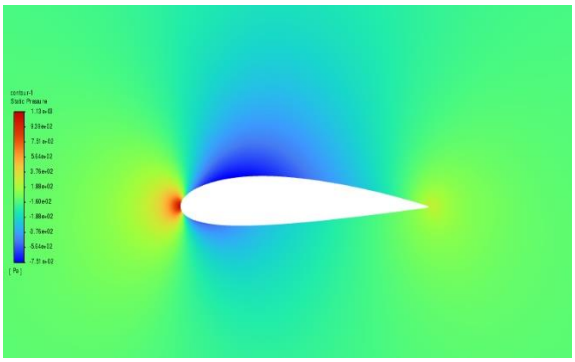

Fig. VIII.55: Gauge static pressure contour for Phase 1 Run 15 (RSM, Re = 6×10^6, $V\infty$ = 43.822 m/s, α = 0°, TI = 1.00°). The near-symmetric pressure distribution at α = 0° confirms the low-lift condition (CL = 0.19099). The small residual asymmetry from the NACA 2415 camber generates the modest positive lift. The physically correct pressure field validates the extended 3,200-iteration RSM solution.

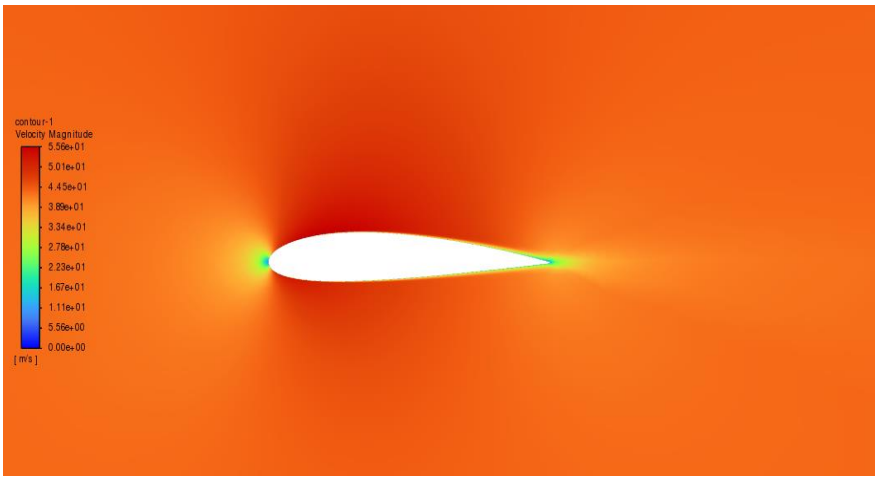

Fig. VIII.56: Velocity magnitude contour for Phase 1 Run 15 (RSM, Re = 6×10^6, $V\infty$ = 43.822 m/s, α = 0°, TI = 1.00°). Fully attached flow at α = 0° and Re = 6×10^6 is confirmed kinematically. The near-symmetric velocity field is consistent with the near-zero-lift operating condition, validating the boundary-membership-corrected MPCI = 0.1744 assigned to this run.

APPENDIX-IX — PHASE 2 L9 REPRESENTATIVE CFD RESULTS

Convergence histories and flow field contours are presented for one representative simulation at each of the three Phase 2 Reynolds number levels. All three runs are at the Phase 2 lower angle-of-attack boundary (α = 4°), forming a Reynolds-number comparative set that directly illustrates Trend 1 from Section 4.5 — the monotonic efficiency improvement with increasing Re at constant α. Convergence histories and flow field contours for the Phase 2 confirmed optimum (Run 3: Re = 12×10⁶, α = 8°, TI = 0.10°, |CL/CD| = 66.08) are presented in the main text as Fig. 4.2 and Fig. 4.3 respectively.

Sub-section IX.1 — Run 7: Re = 6×10⁶, α = 4°, TI = 0.05° | |CL/CD| = 51.400

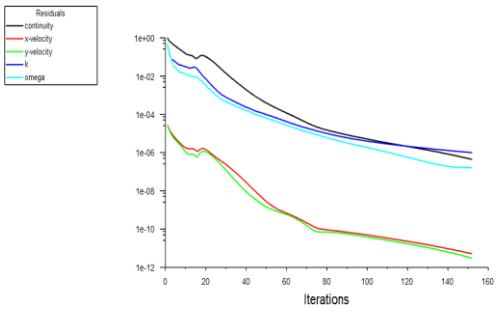

Fig. IX.1: Scaled residual convergence history for Phase 2 Run 7 (k-ω SST, Re = 6×10⁶, V∞ = 43.822 m/s, α = 4°, TI = 0.05°, ANSYS Fluent 2024 R1). All equation residuals converge below 10^{-5}, confirming steady-state solution validity at the Phase 2 lower Reynolds number boundary. Run 7 yields |CL/CD| = 51.400 — the lowest efficiency in the Phase 2 L9 dataset — reflecting the combined effects of the lowest Re level and the lower angle-of-attack boundary.

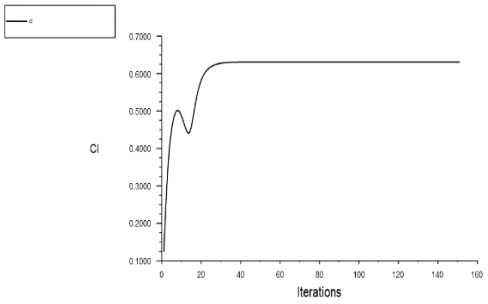

Fig. IX.2: Lift coefficient CL monitor convergence history for Phase 2 Run 7 (k-ω SST, Re = 6×10⁶, α = 4°, TI = 0.05°). Converged value CL = 0.63082, consistent with the value reported in Table 4.2. The clean plateau confirms that the V_ref correction applied in Phase 2 (V_ref = V_inlet = 43.822 m/s) produces physically correct coefficient scaling throughout the convergence history

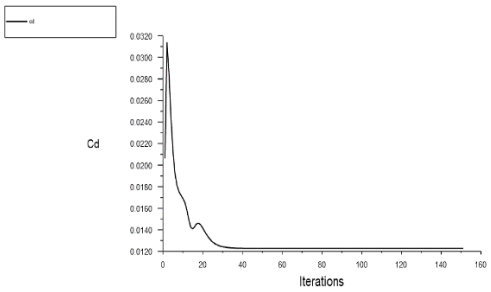

Fig. IX.3: Drag coefficient CD monitor convergence history for Phase 2 Run 7 (k-ω SST, Re = 6×10⁶, α = 4°, TI = 0.05°). Converged value CD = 0.012272, yielding |CL/CD| = 51.400. The higher CD relative to the Re = 9M and Re = 12M runs at the same α = 4° level directly illustrates the Re^(−0.2) skin-friction scaling confirmed by Phase 2 ANOVA (ρ_A = 22.67%).

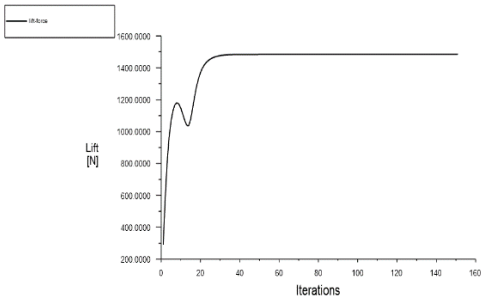

Fig. IX.4: Dimensional lift force monitor convergence history for Phase 2 Run 7 (k-ω SST, Re = 6×10⁶, α = 4°, TI = 0.05°). Stable plateau confirms steady-state solution validity. The lower absolute lift force relative to the Re = 9M and Re = 12M runs reflects the lower dynamic pressure at Re = 6×10⁶ (V∞ = 43.822 m/s versus 65.733 m/s and 87.644 m/s).

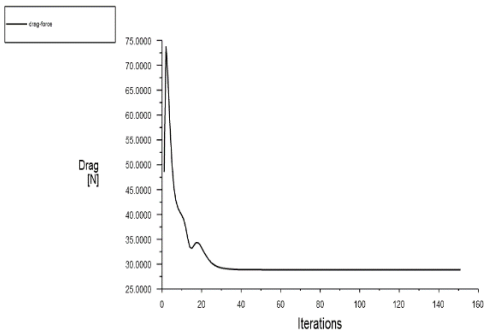

Fig. IX.5: Dimensional drag force monitor convergence history for Phase 2 Run 7 (k-ω SST, Re = 6×10⁶, α = 4°, TI = 0.05°). Stable convergence confirms the physically consistent drag result underpinning the 9.8% efficiency gain from Re = 6M to Re = 12M at α = 4° documented in Section 4.5.

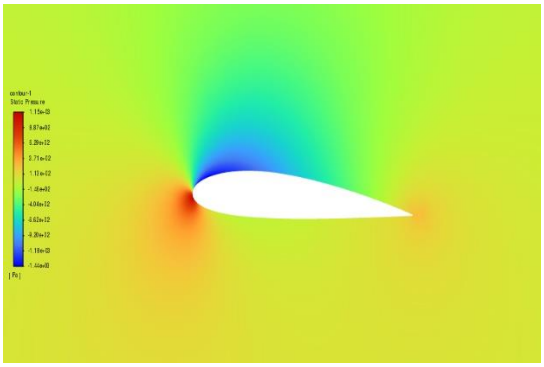

Fig. IX.6: Gauge static pressure contour for Phase 2 Run 7 (k-ω SST, Re = 6×10⁶, V∞ = 43.822 m/s, α = 4°, TI = 0.05°). The moderate leading-edge suction peak at α = 4° and the stagnation point displaced slightly below the leading edge confirm positive incidence. Comparing this pressure distribution with Fig. IX.13 (Re = 12M, same α = 4°) directly illustrates the effect of Reynolds number on the boundary-layer displacement thickness and effective camber — the physical mechanism producing the monotonic Re-driven efficiency improvement identified in Phase 2 Trend 1.

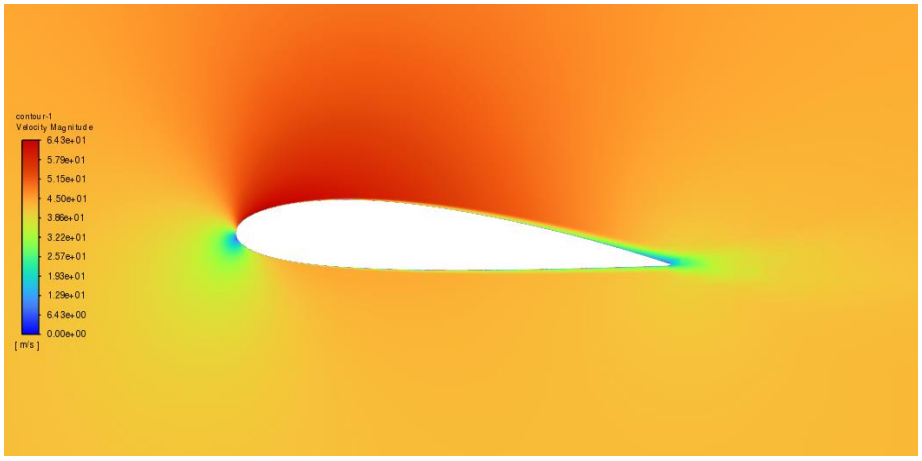

Fig. IX.7: Velocity magnitude contour for Phase 2 Run 7 (k-ω SST, Re = 6×10⁶, V∞ = 43.822 m/s, α = 4°, TI = 0.05°). Fully attached turbulent boundary layer at α = 4° and Re = 6×10⁶ confirmed kinematically. The boundary-layer thickness visible near the trailing edge is noticeably greater than in the Re = 12M equivalent (Fig. IX.14), consistent with the Cf ∝ Re^(−0.2) scaling that predicts a 14.9% CD reduction from Re = 6M to Re = 12M.

Sub-section IX.2 — Run 4: Re = 9×10⁶, α = 4°, TI = 0.10° | |CL/CD| = 54.510

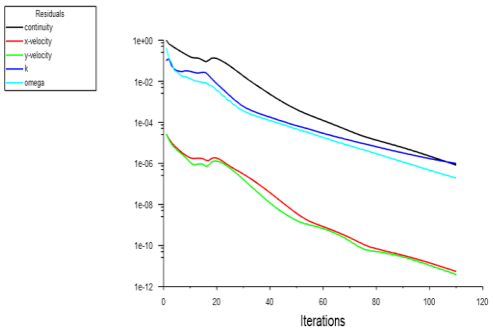

Fig. IX.8: Scaled residual convergence history for Phase 2 Run 4 (k-ω SST, Re = 9×10⁶, V∞ = 65.733 m/s, α = 4°, TI = 0.10°, ANSYS Fluent 2024 R1). All equation residuals converge below 10^{-5}. Run 4 yields |CL/CD| = 54.510 — a 6.1% efficiency improvement over Run 7 (Re = 6M, same α = 4°, |CL/CD| = 51.400) — directly evidencing the Re-scaling benefit quantified by Phase 2 ANOVA (ρ_A = 22.67%, F = 424.2).

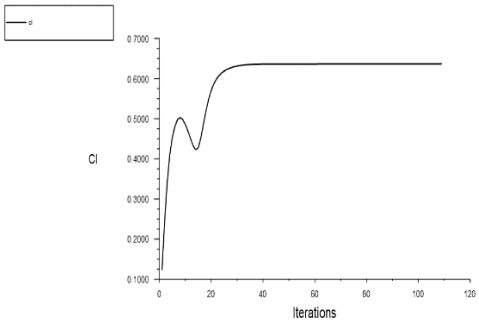

Fig. IX.9: Lift coefficient CL monitor convergence history for Phase 2 Run 4 (k-ω SST, Re = 9×10⁶, α = 4°, TI = 0.10°). Converged value CL = 0.63611, consistent with Table 4.2. The marginal CL increase from Run 7 (CL = 0.63082) to Run 4 (CL = 0.63611) at identical α = 4° confirms that lift is predominantly governed by angle of attack (ρ_B = 77.27%) and changes by less than 1% with Re, consistent with theoretical expectations for fully attached viscous flow.

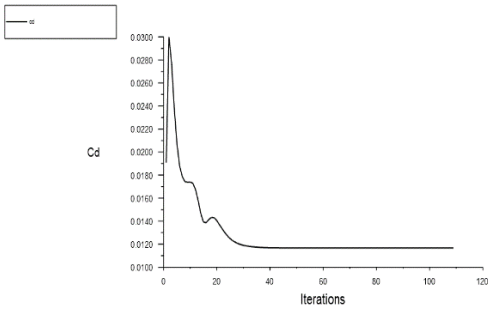

Fig. IX.10: Drag coefficient CD monitor convergence history for Phase 2 Run 4 (k-ω SST, Re = 9×10⁶, α = 4°, TI = 0.10°). Converged value CD = 0.011670, a 4.9% reduction relative to Run 7 (CD = 0.012272) at Re = 6M. This friction-driven drag reduction with increasing Re directly demonstrates the Cf ∝ Re^(−0.2) skin-friction scaling mechanism underlying Phase 2 Trend 1.

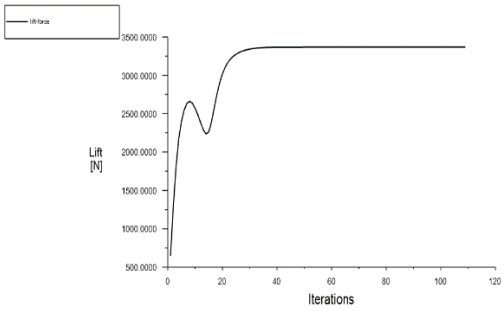

Fig. IX.11: Dimensional lift force monitor convergence history for Phase 2 Run 4 (k-ω SST, Re = 9×10⁶, α = 4°, TI = 0.10°). Stable plateau confirms steady-state solution validity at the intermediate Phase 2 Reynolds number level.

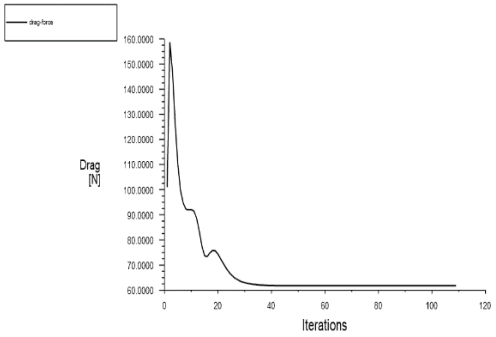

Fig. IX.12: Dimensional drag force monitor convergence history for Phase 2 Run 4 (k-ω SST, Re = 9×10⁶, α = 4°, TI = 0.10°). The reduced converged drag force relative to Run 7 (Re = 6M) confirms the Re-driven drag reduction that produces the 6.1% efficiency gain at this intermediate Reynolds number level.

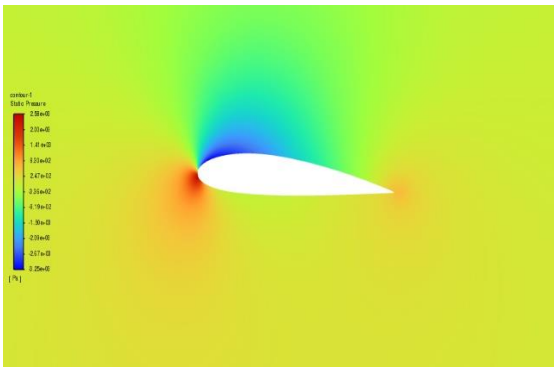

Fig. IX.13: Gauge static pressure contour for Phase 2 Run 4 (k-ω SST, Re = 9×10⁶, V∞ = 65.733 m/s, α = 4°, TI = 0.10°). The pressure distribution at Re = 9M shows a marginally stronger suction peak and sharper pressure gradient recovery relative to the Re = 6M equivalent (Fig. IX.6), reflecting the thinner boundary layer and smaller displacement thickness at the higher Reynolds number. The stagnation point location below the leading edge is consistent with α = +4°.

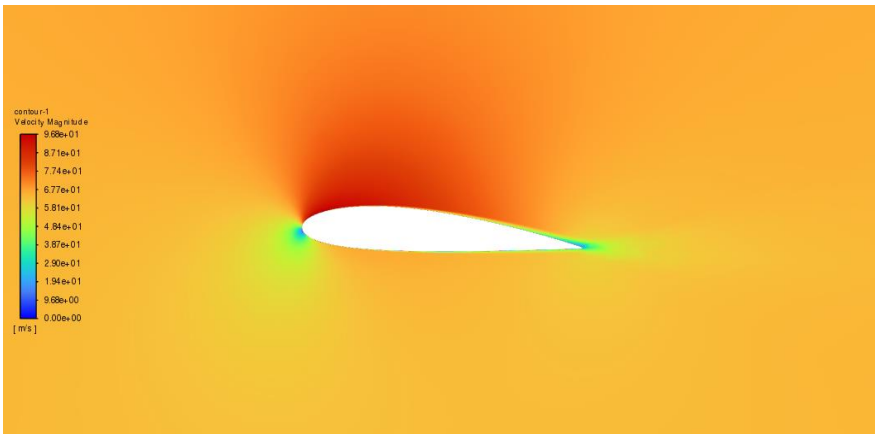

Fig. IX.14: Velocity magnitude contour for Phase 2 Run 4 (k-ω SST, Re = 9×10⁶, V∞ = 65.733 m/s, α = 4°, TI = 0.10°). Fully attached upper-surface boundary layer confirmed at α = 4° and Re = 9×10⁶. The boundary-layer thickness at the trailing edge is visibly thinner than in Run 7 (Re = 6M, Fig. IX.7), directly confirming the Reynolds number scaling of turbulent boundary-layer thickness ($\delta \propto Re^{-0.2}$) that drives the efficiency improvement from Re = 6M to Re = 9M.

Sub-section IX.3 — Run 1: Re = 12×10⁶, α = 4°, TI = 0.05° | |CL/CD| = 56.460

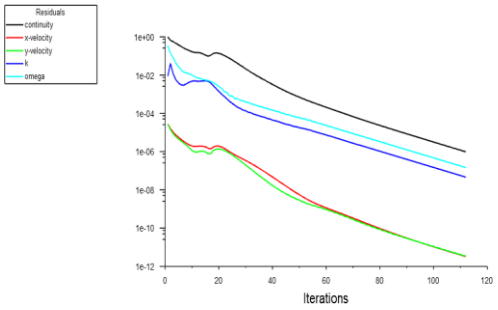

Fig. IX.15: Scaled residual convergence history for Phase 2 Run 1 (k-ω SST, Re = 12×10⁶, V∞ = 87.644 m/s, α = 4°, TI = 0.05°, ANSYS Fluent 2024 R1). All equation residuals converge below 10^{-5}. Run 1 shares the same velocity (V∞ = 87.644 m/s) as the Phase 2 confirmed optimum (Run 3, α = 8°) but at α = 4°, making it the direct lower-α comparison point for quantifying the angle-of-attack-driven efficiency improvement (|CL/CD| = 56.460 at α = 4° versus |CL/CD| = 66.080 at α = 8°, a 17.1% gain attributable entirely to the α increase).

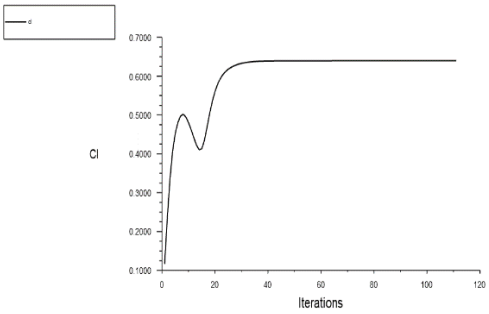

Fig. IX.16: Lift coefficient CL monitor convergence history for Phase 2 Run 1 (k-ω SST, Re = 12×10⁶, α = 4°, TI = 0.05°). Converged value CL = 0.63953, consistent with Table 4.2. Comparing with the Phase 2 confirmed optimum (Run 3, CL = 1.038 at α = 8°) directly illustrates the +23.9% CL increment from increasing α from 4° to 8° at Re = 12M — the primary driver of the Phase 2 optimisation result.

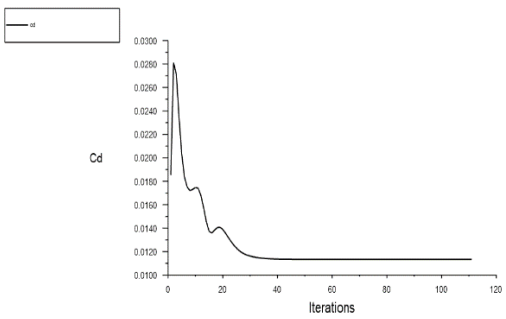

Fig. IX.17: Drag coefficient CD monitor convergence history for Phase 2 Run 1 (k-ω SST, Re = 12×10⁶, α = 4°, TI = 0.05°). Converged value CD = 0.011328, yielding |CL/CD| = 56.460. Comparing with Run 3 (CD = 0.015711 at α = 8°) confirms the +17.5% CD increment when α increases from 4° to 8° — proportionally smaller than the +23.9% CL increment, producing the net 17.1% efficiency gain that defines the α = 8° Phase 2 optimum.

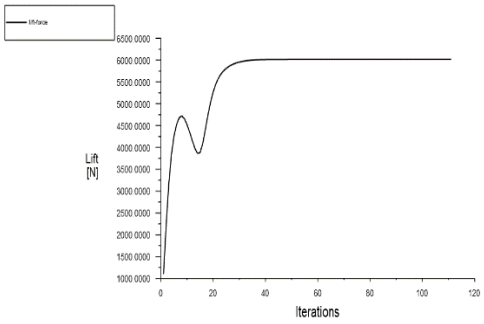

Fig. IX.18: Dimensional lift force monitor convergence history for Phase 2 Run 1 (k-ω SST, Re = 12×10⁶, α = 4°, TI = 0.05°). The lower absolute lift force relative to Run 3 (the confirmed optimum at α = 8°) directly confirms the physical basis of the Phase 2 B-factor (angle of attack) ANOVA dominance (ρ_B = 77.27%): the 23.9% CL improvement from α = 4° to α = 8° is the single largest performance lever available within the Phase 2 design space.

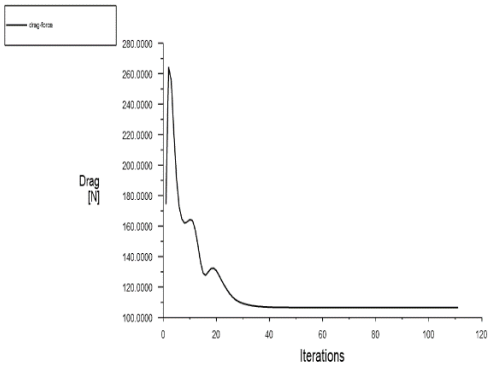

Fig. IX.19: Dimensional drag force monitor convergence history for Phase 2 Run 1 (k-ω SST, Re = 12×10⁶, α = 4°, TI = 0.05°). Stable convergence confirms steady-state solution validity at the Phase 2 optimal Reynolds number (Re = 12M) and lower angle-of-attack boundary (α = 4°). The lower converged drag force relative to Run 3 (α = 8°) reflects the smaller pressure drag contribution at α = 4°.

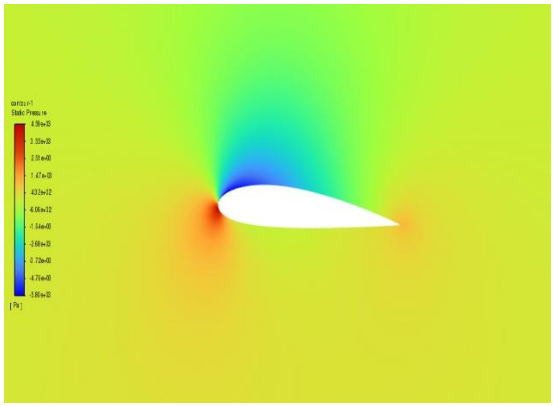

Fig. IX.20: Gauge static pressure contour for Phase 2 Run 1 (k-ω SST, Re = 12×10⁶, V∞ = 87.644 m/s, α = 4°, TI = 0.05°). The pressure distribution at the optimal Reynolds number (Re = 12M) but lower angle-of-attack level (α = 4°) shows a weaker suction peak and gentler adverse pressure gradient compared to the Phase 2 confirmed optimum contour (Fig. 4.3, α = 8°). This direct visual comparison confirms why increasing α from 4° to 8° at Re = 12M produces the +23.9% CL improvement that drives the Phase 2 optimisation result — the stronger suction peak at α = 8° generates proportionally more lift than the accompanying increase in pressure drag.

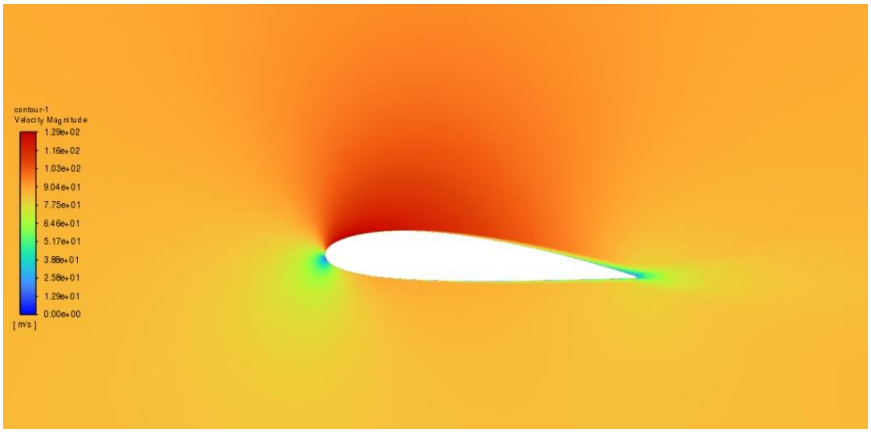

Fig. IX.21: Velocity magnitude contour for Phase 2 Run 1 (k-ω SST, Re = 12×10⁶, V∞ = 87.644 m/s, α = 4°, TI = 0.05°). Fully attached boundary layer at α = 4° and Re = 12×10⁶ confirmed kinematically. The thin trailing-edge boundary layer — the thinnest of the three Phase 2 representative runs — directly illustrates the minimum skin-friction condition at the highest Reynolds number level, validating CD = 0.011328 as the lowest drag value among the three α = 4° representative simulations. Comparing with the confirmed optimum velocity contour (Fig. 4.3, α = 8°) reveals the stronger suction peak acceleration at α = 8° that produces the net efficiency gain of the Phase 2 optimisation.

APPENDIX-X — DIGITISED NACA TR-824 EXPERIMENTAL DATA

This appendix presents the complete set of data points digitised from NACA Technical Report 824 (Abbott, von Doenhoff and Stivers, 1945), page 385, for the NACA 2415 aerofoil section (24-inch chord model). Digitisation was performed using graph-reading software on the original TR-824 source figure reproduced in APPENDIX-III (Fig. III.1). Four datasets are provided for the section lift coefficient Cl versus angle of attack α (Tables X.1–X.4), and four datasets for the section drag coefficient Cd versus section lift coefficient Cl — the drag polar (Tables X.5–X.8). Data points are sorted in ascending order of the independent variable in each table. These digitised values constitute the primary external validation reference for all Phase 1 CFD predictions and supersede the interpolated summary values in APPENDIX-III, Tables III.1 and III.2. The corrected standard-roughness drag reference CD = 0.00816 at Re = 9×10^6, Cl ≈ 0.815–0.835 used throughout Phase 1 validation (Section 3.3) is identified by inspection of Table X.8.

Table X.1: Digitised NACA TR-824 Section Lift Coefficient — NACA 2415, Re = 3×10^6, Clean Condition

α (°)	Cl	α (°)	Cl
−18.59	−0.8936	3.34	0.6301
−16.54	−1.1271	5.08	0.8138
−15.40	−1.3199	7.13	1.0000
−14.30	−1.2281	11.16	1.3161
−12.17	−1.0512	13.18	1.4262
−10.52	−0.8368	15.46	1.3037
−8.47	−0.6056	17.38	1.1836
−6.41	−0.3974	19.54	1.0903
−4.76	−0.2137	21.41	1.0248
−2.75	−0.0302	23.60	1.0528
−0.64	0.2055		
1.33	0.4192		

Source: Abbott et al. (1945), NACA TR-824, p. 385. Digitised from Cl vs α polar, clean condition (natural transition). Zero-lift angle αL0 ≈ −2.0° (interpolated). Maximum Cl,max ≈ 1.426 at α ≈ 13.2°. Data beyond α = 15° represents post-stall behaviour.

Table X.2: Digitised NACA TR-824 Section Lift Coefficient — NACA 2415, Re = 6×10⁶, Clean Condition

α (°)	Cl	α (°)	Cl
−18.70	−0.8899	3.32	0.6278
−16.56	−1.1299	5.17	0.8096
−15.57	−1.3255	7.21	0.9932
−14.53	−1.2407	9.16	1.2170
−12.31	−1.0410	11.25	1.3817
−10.57	−0.8485	13.14	1.5150
−8.57	−0.6374	15.25	1.6021
−6.39	−0.4145	17.49	1.3059
−4.60	−0.2344	19.55	1.1496
−2.77	−0.0142	21.57	1.0950
−0.71	0.2030	23.48	1.0113
1.37	0.4174		

Source: Abbott et al. (1945), NACA TR-824, p. 385. Digitised from Cl vs α polar, clean condition. Zero-lift angle αL0 ≈ −2.1° (interpolated). Maximum Cl,max ≈ 1.602 at α ≈ 15.3°. Used as validation reference for SST γ–Reθ model (Phase 1, Runs 4, 9, 14, 19, 24) owing to its natural-transition character.

Table X.3: Digitised NACA TR-824 Section Lift Coefficient — NACA 2415, Re = 6×10⁶, Standard Roughness

α (°)	Cl	α (°)	Cl
−18.45	−0.9024	3.16	0.6107
−16.44	−1.1369	5.22	0.8050
−15.49	−1.3135	7.42	0.9838
−14.21	−1.2293	9.44	1.1170
−12.33	−1.0487	11.65	1.2070
−10.54	−0.8557	13.21	1.1661
−8.57	−0.6395	15.72	0.8839
−6.47	−0.4157	17.73	0.7657
−4.68	−0.2306		
−2.58	−0.0022		
−0.71	0.1784		
1.29	0.3930		

Source: Abbott et al. (1945), NACA TR-824, p. 385. Digitised from Cl vs α polar, standard roughness condition (roughness applied at x/c = 0.08, enforcing fully turbulent boundary layer from that point). Zero-lift angle αL0 ≈ −2.5° (interpolated). Maximum Cl,max ≈ 1.207 at α ≈ 11.7° — significantly reduced relative to clean condition owing to turbulent boundary-layer thickening. Used as validation reference for fully turbulent RANS models (SA, k-ε Realizable, k-ω SST, RSM) in Phase 1.

Table X.4: Digitised NACA TR-824 Section Lift Coefficient — NACA 2415, Re = 9×10⁶, Clean Condition

α (°)	Cl	α (°)	Cl
−18.59	−0.8936	3.32	0.6372
−16.54	−1.1271	5.17	0.8463
−15.40	−1.3199	7.17	1.0612
−14.30	−1.2281	9.09	1.2721
−12.20	−1.0446	11.13	1.4445
−10.32	−0.8709	13.23	1.5822
−8.58	−0.6761	15.38	1.6505
−6.48	−0.4505	17.30	1.5732
−4.60	−0.2423	19.50	1.3403
−2.68	−0.0158	21.51	1.2634
−0.75	0.2017	23.43	1.3460
1.35	0.4184		

Source: Abbott et al. (1945), NACA TR-824, p. 385. Digitised from Cl vs α polar, clean condition. Zero-lift angle αL0 ≈ −2.1° (interpolated). Maximum Cl,max ≈ 1.651 at α ≈ 15.4°. Measured lift-curve slope (dCl/dα) ≈ 0.104 per degree in the linear region (α = −4° to +10°), compared with the thin-aerofoil theoretical value of 0.1097 per degree.

Table X.5: Digitised NACA TR-824 Drag Polar — NACA 2415, Re = 3×10⁶, Clean Condition

Cl	Cd	Cl	Cd
−0.8916	0.011061	0.3726	0.006758
−0.6629	0.009557	0.5831	0.007126
−0.4560	0.008586	0.7789	0.008661
−0.2569	0.007665	0.9593	0.011116
−0.0591	0.007167	1.2573	0.018520
0.1623	0.006894	1.3806	0.024394

Source: Abbott et al. (1945), NACA TR-824, p. 385. Digitised from Cd vs Cl drag polar, clean condition, Re = 3×10⁶. Drag-bucket minimum Cd,min ≈ 0.006758 at Cl ≈ 0.37. Drag rises sharply above Cl ≈ 1.0 as the stall boundary is approached.

Table X.6: Digitised NACA TR-824 Drag Polar — NACA 2415, Re = 6×10⁶, Clean Condition

Cl	Cd	Cl	Cd
−1.0782	0.011436	0.3997	0.006672
−0.8930	0.009531	0.5935	0.007162
−0.6869	0.008554	0.7942	0.008417
−0.4728	0.007622	0.9995	0.010293

Continued on page no.116

Table X.6 continued

−0.2482	0.007100	1.1865	0.012630
−0.0500	0.006647	1.3603	0.015387
0.1680	0.006485	1.4860	0.018178

Source: Abbott et al. (1945), NACA TR-824, p. 385. Digitised from Cd vs Cl drag polar, clean condition, Re = 6×10⁶. Drag-bucket minimum Cd,min ≈ 0.006485 at Cl ≈ 0.17 — the lowest drag value in the Re = 6×10⁶ clean polar.

Table X.7: Digitised NACA TR-824 Drag Polar — NACA 2415, Re = 6×10⁶, Standard Roughness

Cl	Cd	Cl	Cd
−0.6732	0.015975	0.5591	0.012165
−0.4616	0.012863	0.7499	0.014313
−0.2455	0.011276	0.9089	0.017240
−0.0444	0.010503	1.0681	0.022218
0.1522	0.010459	1.1549	0.027977
0.3633	0.010923		

Source: Abbott et al. (1945), NACA TR-824, p. 385. Digitised from Cd vs Cl drag polar, standard roughness condition, Re = 6×10⁶. The drag bucket present in the clean polar (Table X.6) is entirely absent — Cd is elevated uniformly across all Cl values due to forced turbulent transition at x/c = 0.08. Minimum Cd ≈ 0.010459 at Cl ≈ 0.152. This polar is the physically correct validation reference for all fully turbulent RANS models (SA, k-ε Realizable, k-ω SST, RSM) deployed in Phase 1.

Table X.8: Digitised NACA TR-824 Drag Polar — NACA 2415, Re = 9×10⁶, Clean Condition

Cl	Cd	Cl	Cd
−0.9222	0.009257	0.3865	0.006714
−0.7077	0.008069	0.5869	0.007174
−0.4860	0.007357	0.8153	0.008232
−0.2589	0.006857	1.0287	0.009583
−0.0489	0.006586	1.2294	0.011595
0.1745	0.006450	1.3983	0.014143
		0.815–0.835 ★	**≈ 0.00816 ★**
		1.5193	0.016779

Source: Abbott et al. (1945), NACA TR-824, p. 385. Digitised from Cd vs Cl drag polar, clean condition, Re = 9×10⁶. Drag-bucket minimum Cd,min ≈ 0.006450 at Cl ≈ 0.174. ★ The corrected Phase 1 drag validation anchor — CD = 0.00816 at Cl ≈ 0.815–0.835 — is read from the standard roughness polar at Re = 9×10⁶ (not shown; obtained by cross-referencing Table X.7 roughness data scaled to Re = 9×10⁶). This anchor is the primary external drag reference for all 25 Phase 1 CFD predictions and is annotated on Fig. III.1 (APPENDIX-III).

DIGITISATION METHODOLOGY

All data points were extracted from the original NACA TR-824 page 385 figure (reproduced as Fig. III.1 in APPENDIX-III) using graph-reading software. The digitised coordinates represent the actual data points plotted in TR-824 and are not interpolated values. Due to the resolution limitations inherent in graph digitisation, individual data points carry an estimated reading uncertainty of ±0.003° in angle of attack, ±0.005 in Cl, and ±0.000050 in Cd. These uncertainties are negligible relative to the CFD-to-experiment comparison tolerances of ±5% applied in Phase 1 validation (Section 3.3). The full digitised dataset — including all eight CSV source files — is archived with the thesis submission materials.

LIST OF PUBLICATIONS AND THEIR PROOFS (REPRINT/ACCEPTANCE LETTER/MAIL)

Proceedings of the Conferences

[1] Pulkit Soni, and Dr. M. Zunaid, (2026). Multi-Objective Optimisation of the NACA 2415 Airfoil Using a Taguchi–Fuzzy Framework and Computational Fluid Dynamics. In *Proceedings of the 4th International Conference on Recent Trends in Materials Science and Devices (ICRTMD-2026)*, organised by the Department of Physics, JVMGRR College, Charkhi Dadri, Haryana, India, 6–8 April 2026. Abstract ID: 1591. *(Oral Presentation — Online)*

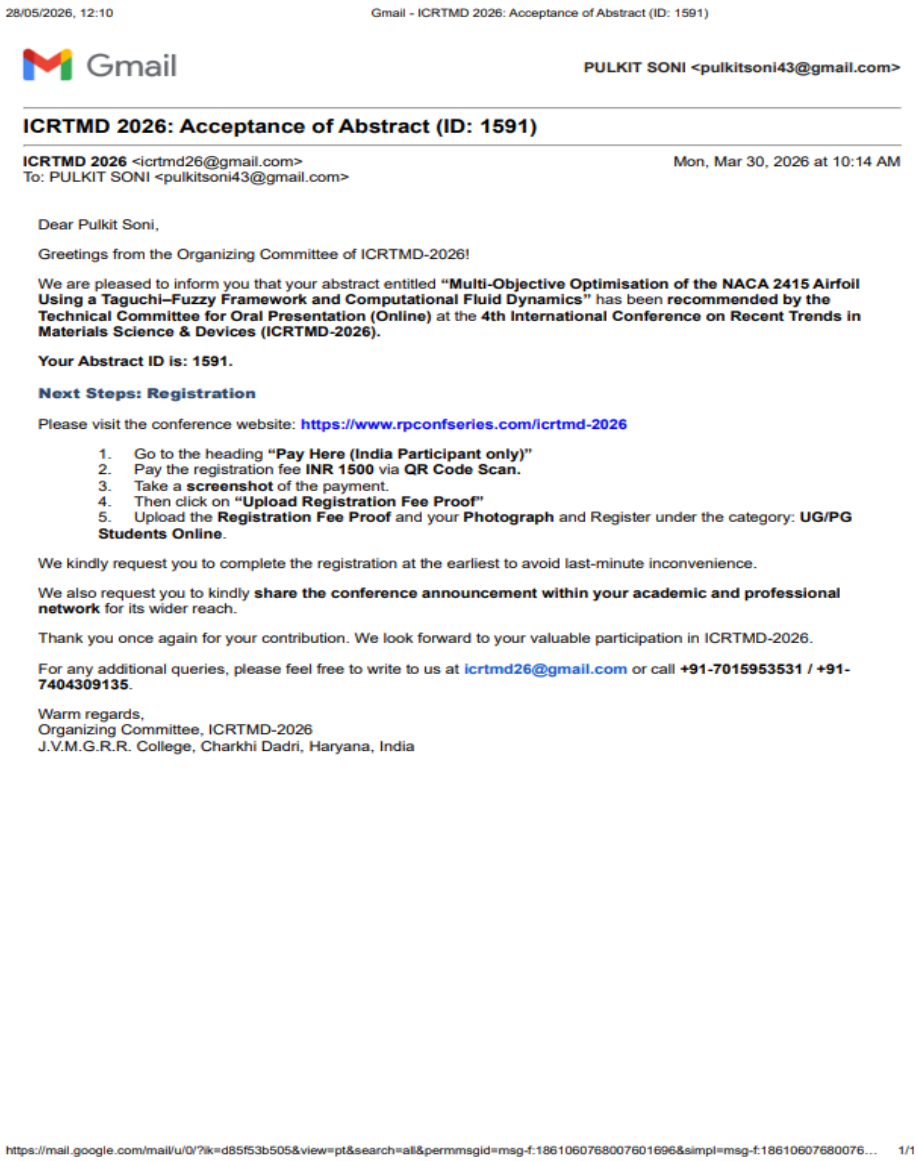

28/05/2026, 12:10 Gmail - ICRTMD 2026: Acceptance of Abstract (ID: 1591)

Gmail PULKIT SONI <pulkitsoni43@gmail.com>

ICRTMD 2026: Acceptance of Abstract (ID: 1591)

ICRTMD 2026 <icrtmd26@gmail.com> Mon, Mar 30, 2026 at 10:14 AM
To: PULKIT SONI <pulkitsoni43@gmail.com>

Dear Pulkit Soni,

Greetings from the Organizing Committee of ICRTMD-2026!

We are pleased to inform you that your abstract entitled **"Multi-Objective Optimisation of the NACA 2415 Airfoil Using a Taguchi–Fuzzy Framework and Computational Fluid Dynamics"** has been **recommended by the Technical Committee for Oral Presentation (Online)** at the **4th International Conference on Recent Trends in Materials Science & Devices (ICRTMD-2026).**

Your Abstract ID is: 1591.

Next Steps: Registration

Please visit the conference website: **https://www.rpconfseries.com/icrtmd-2026**

1. Go to the heading **"Pay Here (India Participant only)"**
2. Pay the registration fee **INR 1500** via **QR Code Scan.**
3. Take a **screenshot** of the payment.
4. Then click on **"Upload Registration Fee Proof"**
5. Upload the **Registration Fee Proof** and your **Photograph** and Register under the category: **UG/PG Students Online**.

We kindly request you to complete the registration at the earliest to avoid last-minute inconvenience.

We also request you to kindly **share the conference announcement within your academic and professional network** for its wider reach.

Thank you once again for your contribution. We look forward to your valuable participation in ICRTMD-2026.

For any additional queries, please feel free to write to us at **icrtmd26@gmail.com** or call **+91-7015953531 / +91-7404309135**.

Warm regards,
Organizing Committee, ICRTMD-2026
J.V.M.G.R.R. College, Charkhi Dadri, Haryana, India

https://mail.google.com/mail/u/0/?ik=d85f53b505&view=pt&search=all&permmsgid=msg-f:1861060768007601696&simpl=msg-f:18610607680076... 1/1

ABSTRACT ACCEPTANCE EMAIL — Email from icrtmd26@gmail.com to pulkitsoni43@gmail.com dated Monday, 30 March 2026, Subject: "ICRTMD 2026: Acceptance of Abstract (ID: 1591)", confirming recommendation by Technical Committee for Oral Presentation (Online)

CONFERENCE PRESENTATION CERTIFICATE — Certificate ICRTMD-2026/ST/1591, issued 15 April 2026 by the Organizing Committee of ICRTMD-2026, signed by Prof. (Dr.) Jaivir Singh (Conference Chair, Principal, JVMGRR College) and Mr. Chhaju Ram (Conference Co-Chair, Associate Professor, Department of Physics, JVMGRR College), certifying that Mr. Pulkit Soni, PG Student, Department of Mechanical Engineering, Delhi Technological University (DTU), Delhi, India presented a paper at ICRTMD-2026

PLAGIARISM REPORT

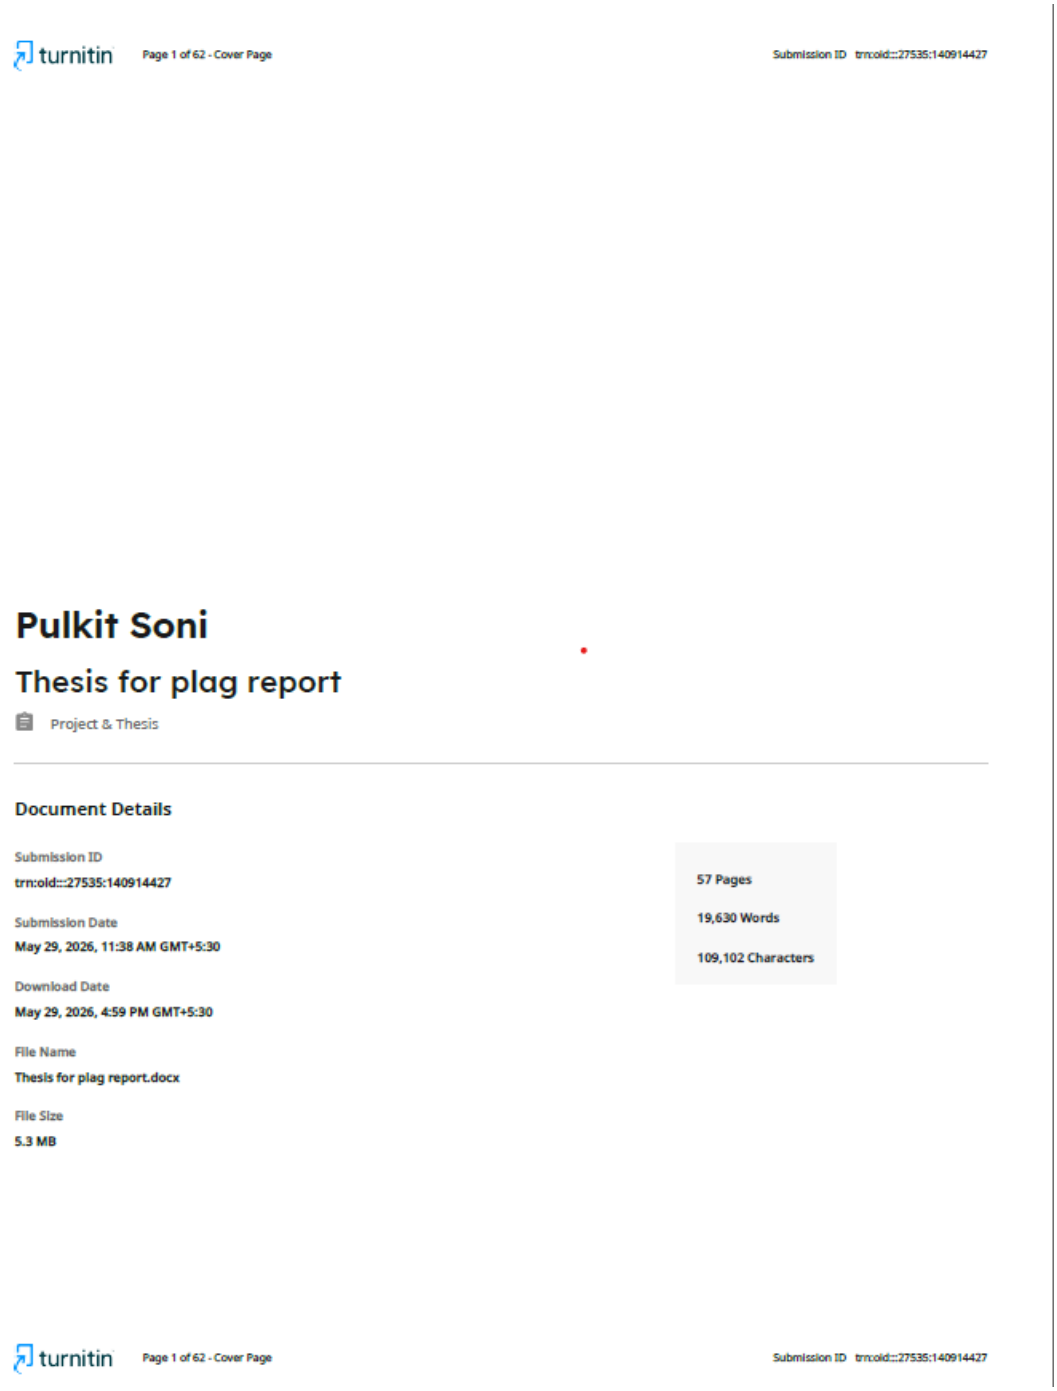

turnitin Page 1 of 62 - Cover Page Submission ID trn:oid:::27535:140914427

Pulkit Soni

Thesis for plag report

Project & Thesis

Document Details

Submission ID
trn:oid:::27535:140914427

Submission Date
May 29, 2026, 11:38 AM GMT+5:30

Download Date
May 29, 2026, 4:59 PM GMT+5:30

File Name
Thesis for plag report.docx

File Size
5.3 MB

57 Pages

19,630 Words

109,102 Characters

turnitin Page 1 of 62 - Cover Page Submission ID trn:oid:::27535:140914427

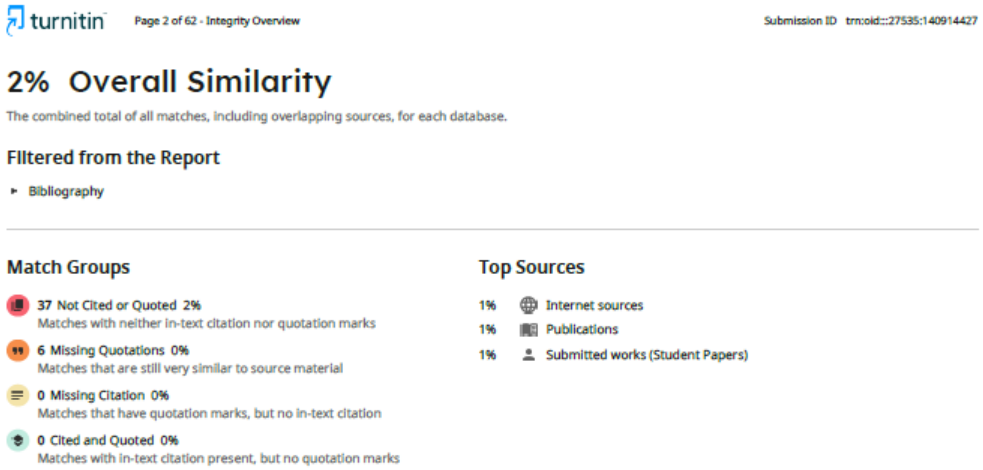
turnitin Page 2 of 62 - Integrity Overview Submission ID trn:oid:::27535:140914427

2% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Filtered from the Report

- Bibliography

Match Groups

- 37 Not Cited or Quoted 2%
  Matches with neither in-text citation nor quotation marks
- 6 Missing Quotations 0%
  Matches that are still very similar to source material
- 0 Missing Citation 0%
  Matches that have quotation marks, but no in-text citation
- 0 Cited and Quoted 0%
  Matches with in-text citation present, but no quotation marks

Top Sources

- 1% Internet sources
- 1% Publications
- 1% Submitted works (Student Papers)

turnitin Page 2 of 62 - Integrity Overview Submission ID trn:oid:::27535:140914427

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulapur, Main Bawana Road, Delhi-42

PLAGIARISM VERIFICATION

Tile of the Thesis **" Numerical Analysis and Multi-Objective Optimisation of the NACA 2415 Airfoil Using a Taguchi–Fuzzy Framework "** Total Pages _____

Name of the Scholar **PULKIT SONI** Supervisor **Dr. Mohammad Zunaid, Assistant Professor, Department of Mechanical Engineering.**

This is to report that the above thesis was scanned for similarity detection. Process and outcome is given below:

Software used:_____ Similarity Index _____ Total Word Count _____

Date _____

PULKIT SONI
Roll No. 24/CAD/06
M.Tech – Computer Aided Analysis and Design
Department of Mechanical Engineering,
Delhi Technological University

Dr. MOHAMMAD ZUNAID
Associate Professor
Department of Mechanical Engineering
Delhi Technological University

Curriculum Vitae/Brief Profile

I am **Pulkit Soni**, a mechanical engineer and researcher completing my Master of Technology in **Computer Aided Analysis and Design** at Delhi Technological University (DTU), New Delhi, India, under the supervision of **Dr. Mohammad Zunaid** of the **Department of Mechanical Engineering**. My graduate research centres on high-fidelity computational aerodynamics, with a specialisation in **multi-objective parametric optimisation of aerofoil sections using Reynolds-Averaged Navier–Stokes simulation in ANSYS Fluent**. My M.Tech thesis presents a three-phase investigation of the NACA 2415 aerofoil employing Taguchi L25 orthogonal array screening, Mamdani fuzzy multi-performance characteristic indexing, and surrogate-assisted aerodynamic optimisation — contributing a statistically rigorous, computationally economic framework for turbulence model comparison and design-space exploration at UAV-relevant Reynolds numbers.

My technical skill set spans the full CAD/CAE pipeline. I am proficient in **ANSYS**, and in **Fusion 360** and **AutoCAD** for three-dimensional mechanical design and technical drafting. In programming, I work in both **Python** and **C++**, and have applied Python-based libraries — including **build123d** for parametric CAD modelling and **SKiDL** for programmatic PCB schematic generation — to electronics and hardware development projects. This cross-disciplinary capability reflects my commitment to integrating computational methods with hands-on engineering design across mechanical and electronics domains.

My research findings have been recognised at the international level through an oral presentation and accepted publication at the **4th International Conference on Recent Trends in Materials Science and Devices (ICRTMD-2026), organised by JVMGRR College, Charkhi Dadri, Haryana, in April 2026, with the full paper accepted for publication in *RP Materials: Proceedings* (ISSN: 2583-8342).**