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, 1:24 PM GMT+5:30

File Name

Thesis for plag report.docx

File Size

5.3 MB

57 Pages

19,630 Words

109,102 Characters

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)

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)

Top Sources

The sources with the highest number of matches within the submission. Overlapping sources will not be displayed.

1	Student papers	University of New South Wales on 2013-11-01	<1%
2	Internet	sportdocbox.com	<1%
3	Internet	www.docstoc.com	<1%
4	Student papers	Clarkson University on 2023-03-10	<1%
5	Student papers	Emirates Aviation College, Aerospace & Academic Studies on 2026-05-08	<1%
6	Internet	archive.org	<1%
7	Internet	c.coek.info	<1%
8	Internet	www.cfd-online.com	<1%
9	Publication	Giorgio Besagni, Nicolò Cristiani, Lorenzo Croci, Gaël Raymond Guédon, Fabio Inz...	<1%
10	Student papers	University of New South Wales on 2013-05-31	<1%

11	Publication	Yong-Jie Liu, Tsung-Wei Zeng. "Evaluation of the effects of synthesis conditions o...	<1%
12	Internet	jurnal.sttkd.ac.id	<1%
13	Internet	theses.gla.ac.uk	<1%
14	Internet	core.ac.uk	<1%
15	Internet	mdpi-res.com	<1%
16	Internet	vdoc.pub	<1%
17	Publication	Aaron M. Kirk, Joaquin I. Gargoloff, Othon K. Rediniotis, Paul G.A. Cizmas. "Numer...	<1%
18	Publication	Aircraft Engineering and Aerospace Technology, Volume 88, Issue 2 (2016)	<1%
19	Publication	I.K. Gavich. "Hydrogeodynamics", A.A. Balkema, Rotterdam, 2020	<1%
20	Student papers	Monash University on 2015-10-23	<1%
21	Student papers	Technical University Delft on 2010-03-22	<1%
22	Internet	eprints.gla.ac.uk	<1%
23	Internet	qspace.library.queensu.ca	<1%
24	Internet	drum.lib.umd.edu	<1%

25	Internet	
e-archivo.uc3m.es		<1%
26	Internet	
hdl.handle.net		<1%
27	Internet	
pubs.aip.org		<1%
28	Student papers	
North West University on 2015-10-27		<1%
29	Student papers	
School of Engineering, The University of Tokyo on 2010-06-03		<1%
30	Publication	
Vanky, Patricia. "Numerical Simulations of the Urban Microclimate", Chalmers Te...		<1%
31	Student papers	
Higher Education Commission Pakistan on 2012-03-21		<1%
32	Publication	
Rui-Yong Liu, Joseph C. Tsai, Chia-Cheng Liu, Agusta Wicaksono, Shih-Ming Chan...		<1%

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 = 3×10^6, 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 = 3×10^6–9×10^6 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 $Re = 2\times10^5$–2×10^6, 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–$Re\theta$), 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/m*^{3}*,* $\mu = 1.7894\times10^{-5}$ *kg/(m·s)) across* $Re = 1$–12×10^6*,* $\alpha = -4°$ *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 (1–12×10^6), 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 $\alpha L0 \approx -2.08°$; clean-condition Cl,max ≈ 1.60–1.65 at $Re = 6\times10^6$; minimum drag coefficient Cd,min ≈ 0.0065 at $Re = 6\times10^6$ 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 $Re = 9\times10^6$, 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 ($Ma \leq 0.258$ at $Re = 12\times10^6$) 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 R̃eθt, enabling prediction of laminar-to-turbulent transition through an empirical correlation linking the freestream turbulence intensity Tu to the transition onset Reynolds number Reθt. 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.

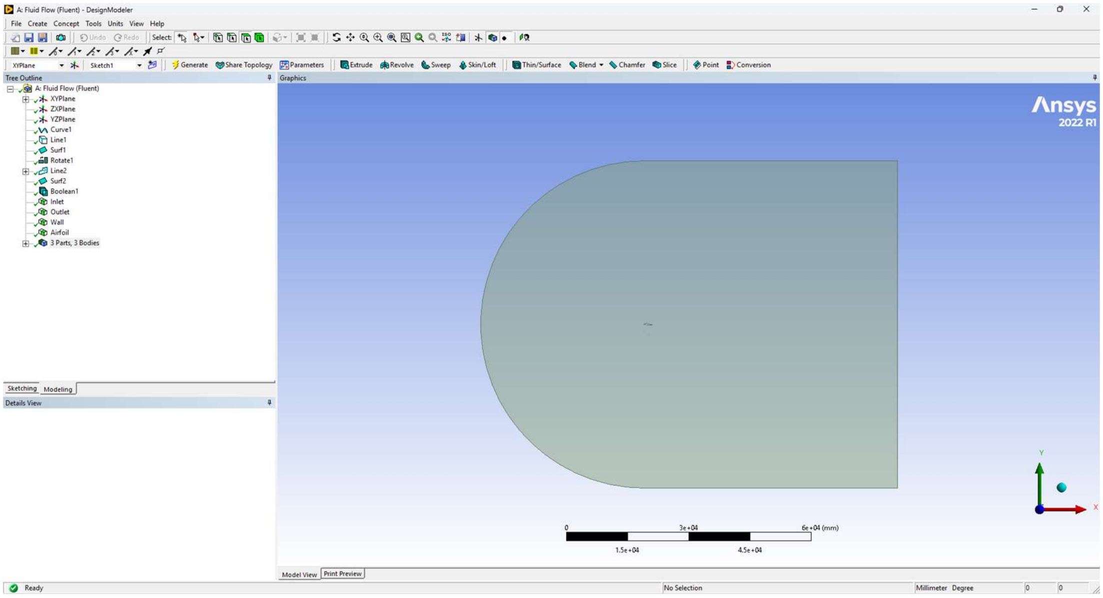

Figure 2.1: C-type computational domain for the NACA 2415 aerofoil, constructed in ANSYS DesignModeler (XY-plane, Z = 0 face shown). The semicircular inlet has radius R = 20c = 40 m; the rectangular downstream section extends L = 20c = 40 m from the inlet axis. At this stand-off distance, far-field velocity 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**.

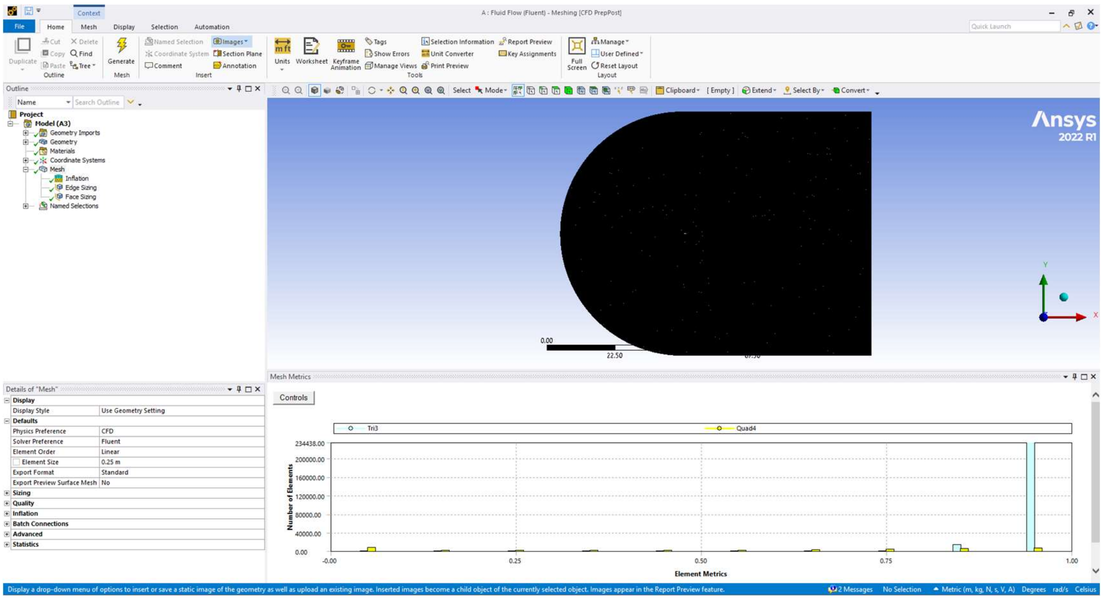

Details of "Mesh"

Display
Display Style: Use Geometry Setting
Defaults
Physics Preference: CFD
Solver Preference: Fluent
Element Order: Linear
☐ Element Size: 0.25 m
Export Format: Standard
Export Preview Surface Mesh: No
Sizing
Quality
Inflation
Batch Connections
Advanced
Statistics

Display a drop-down menu of options to insert or save a static image of the geometry as well as upload an existing image. Inserted images become a child object of the currently selected object. Images appear in the Report Preview feature.

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.

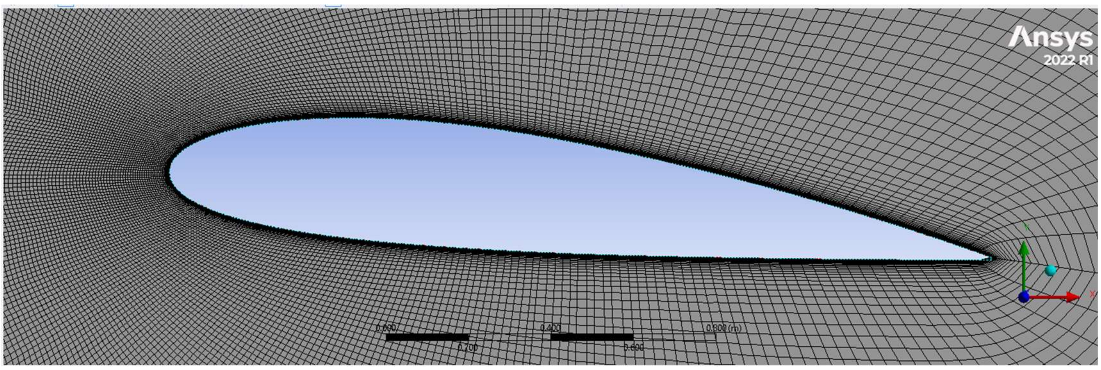

Figure 2.3: Near-wall structured mesh around the NACA 2415 profile (chord c = 2 m, production Mesh M4). The inflation layer comprises 25–30 prismatic cell layers with first-cell height $y_1 = 9.2$ µm — computed from the flat-plate friction velocity at Re = 12×10^6 to enforce $y^+ \leq 1$ across all surface 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 CD = 0.00816 at Re = 9×10^6, CL ≈ 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**.

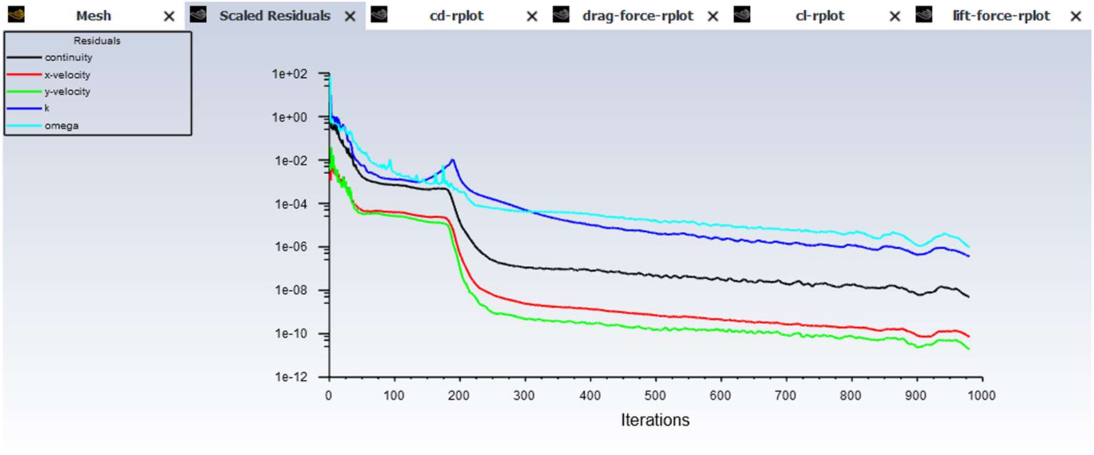

Figure 2.4: Scaled residual convergence history for a representative NACA 2415 k-ω SST simulation (ANSYS Fluent 2024 R1, production Mesh M4, Re = 6×10^6, α = 6°, TI = 0.10%). All five transported quantities — continuity (black), x-velocity (red), y-velocity (green), turbulent kinetic energy k (blue), 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.

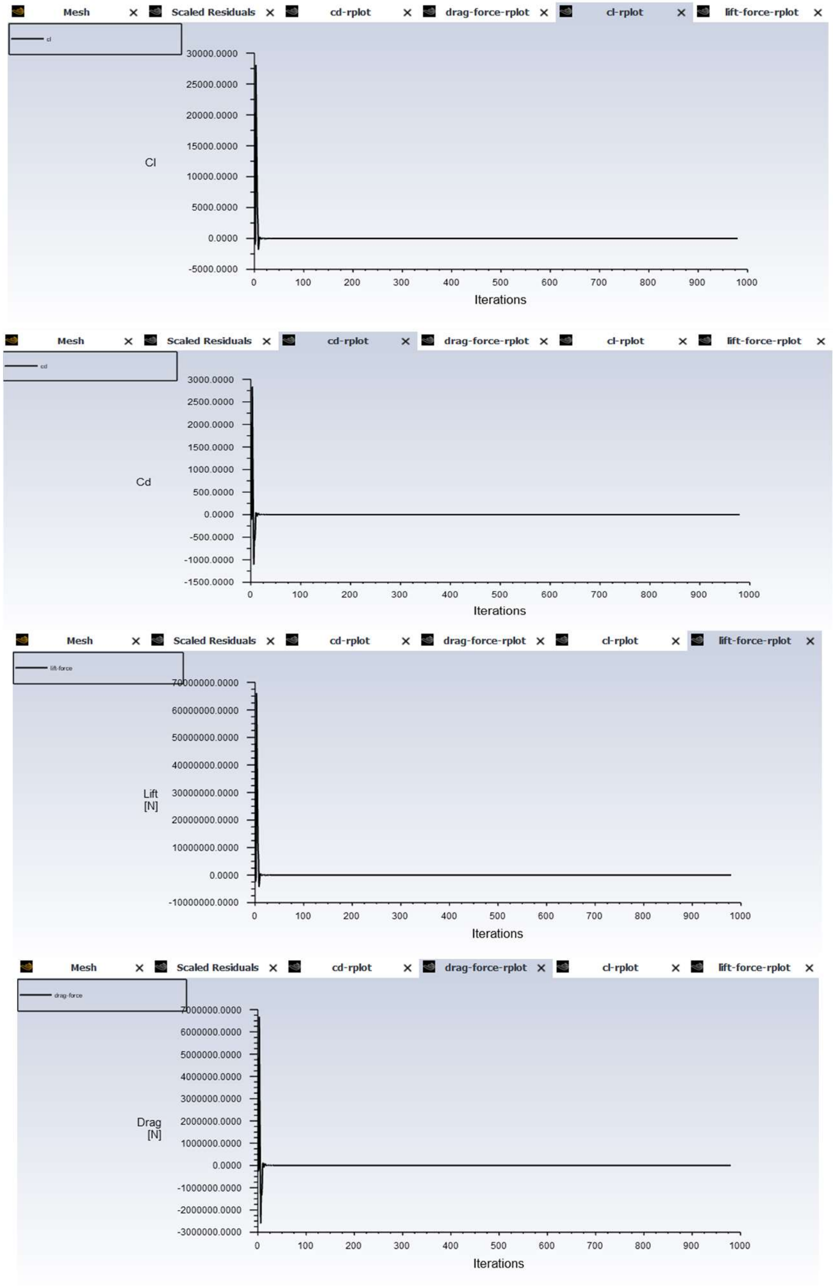

Figure 2.5: Force monitor convergence histories for the same representative k-ω SST simulation shown in Figure 2.4 (Re = 6×10^6, α = 6°, TI = 0.10%): (a) lift coefficient CL, (b) drag coefficient CD, (c) dimensional lift force [N], and (d) dimensional drag force [N], all plotted versus solver iteration 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 ($\times10^6$)	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
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

Continued on page no.17

Table 3.1 continued

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 ($Re = 1\times10^6$) to 87.644 m/s ($Re = 12\times10^6$). 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
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

Continued on page no.18

Table 3.2 Continued

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.

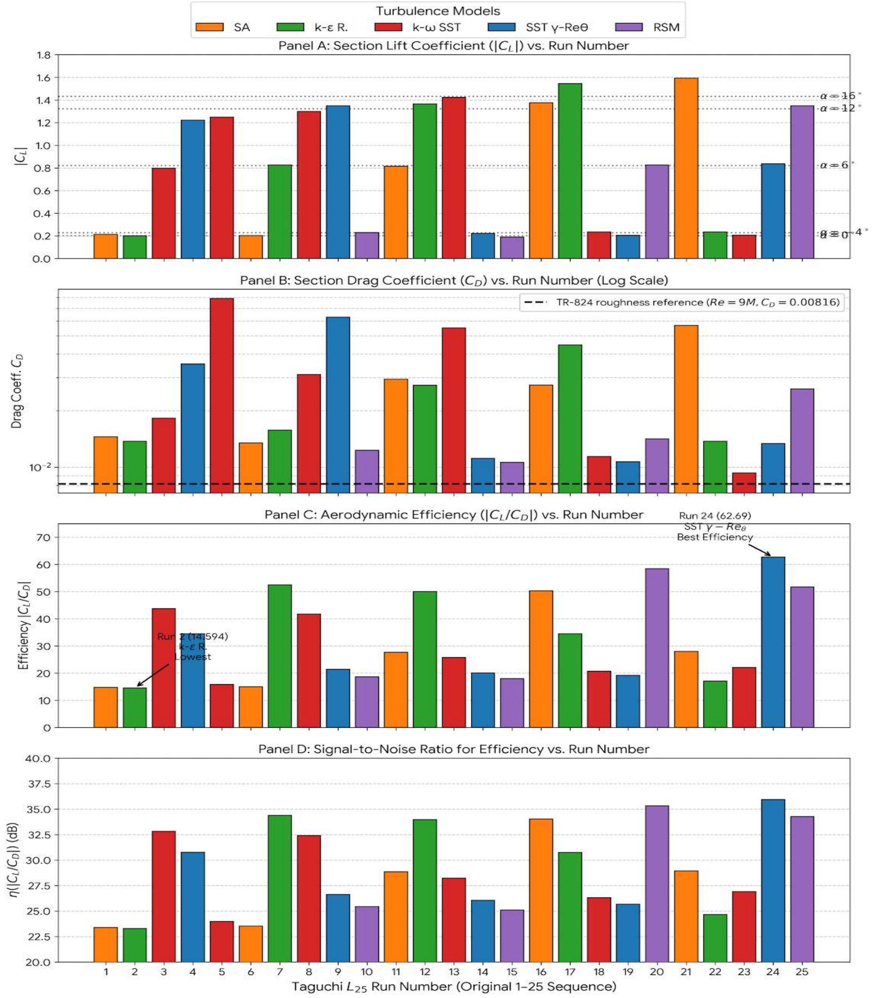

Figure 3.1: Complete Phase 1 L25 aerodynamic dataset visualised across all 25 runs in their original experimental sequence (Runs 1–25, x-axis). Bars are coloured by turbulence model (Factor B): SA — orange; k-ε Realizable — green; k-ω SST — red; SST γ–Reθ — blue; RSM — purple. **Panel A:** Section lift 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 = 9×10^6, 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 $\alpha = -4°$, CD falls from 0.01453 (Re = 1M, Run 1) to 0.01114 (Re = 6M, Run 14), consistent with $Cf \propto Re^{-0.2}$. Lift at $\alpha = 6°$ 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 $Re = 9\times10^6$, 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 $Re = 12\times10^6$ (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).

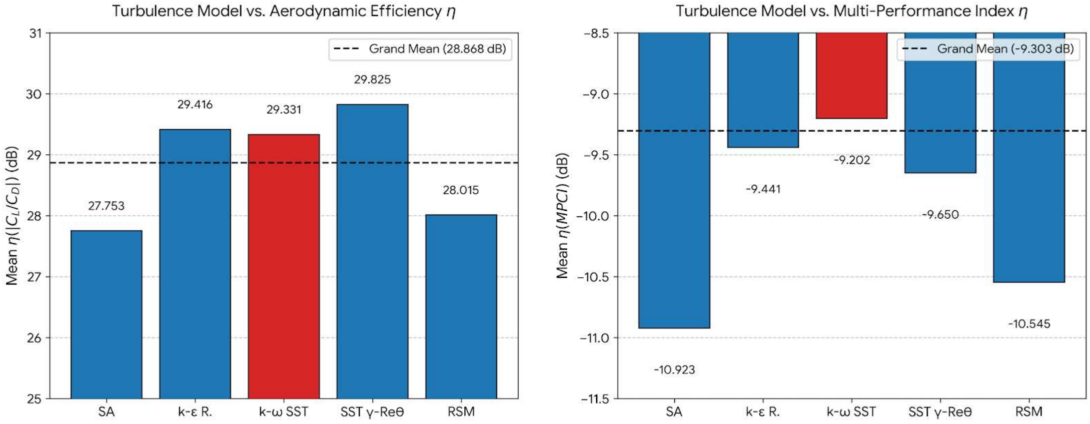

Figure 3.2: Phase 1 turbulence model comparison — five-run level means of both quality metrics for all five RANS closures, coloured by Phase 1 recommendation status (red = recommended k-ω SST; blue = all others). Horizontal dashed lines mark the grand means across all 25 runs. **Left panel:** Mean η(|CL/CD|) 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, α = −4°), where the optimiser correctly drives α toward the zero-lift angle αL0 = −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.

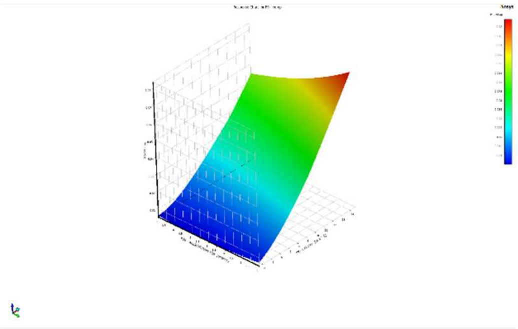

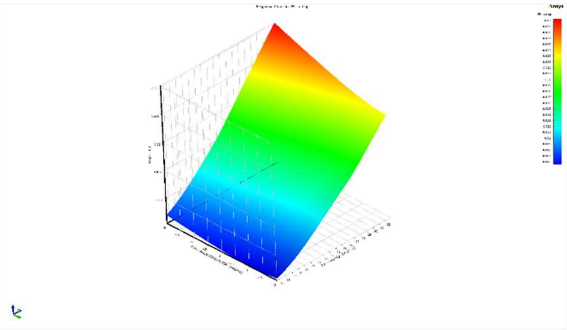

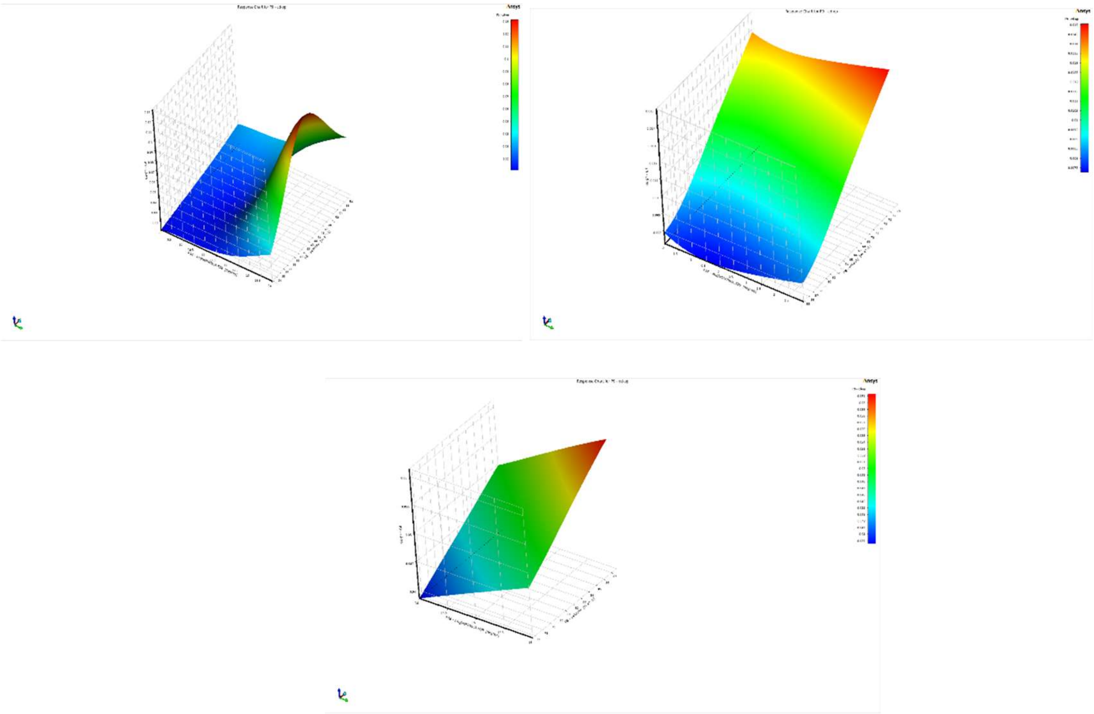

Figure 3.3: Representative NSGA-II CD response surface for Run L12 (k-ε Realizable, Re = 6M, α = 9°–14°, TI = 0.02%–0.07%), generated from the 15-point central composite DoE evaluated in ANSYS Fluent and fitted in ANSYS DesignXplorer. The surface colour maps CD from minimum (blue, ≈0.008) to maximum (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.

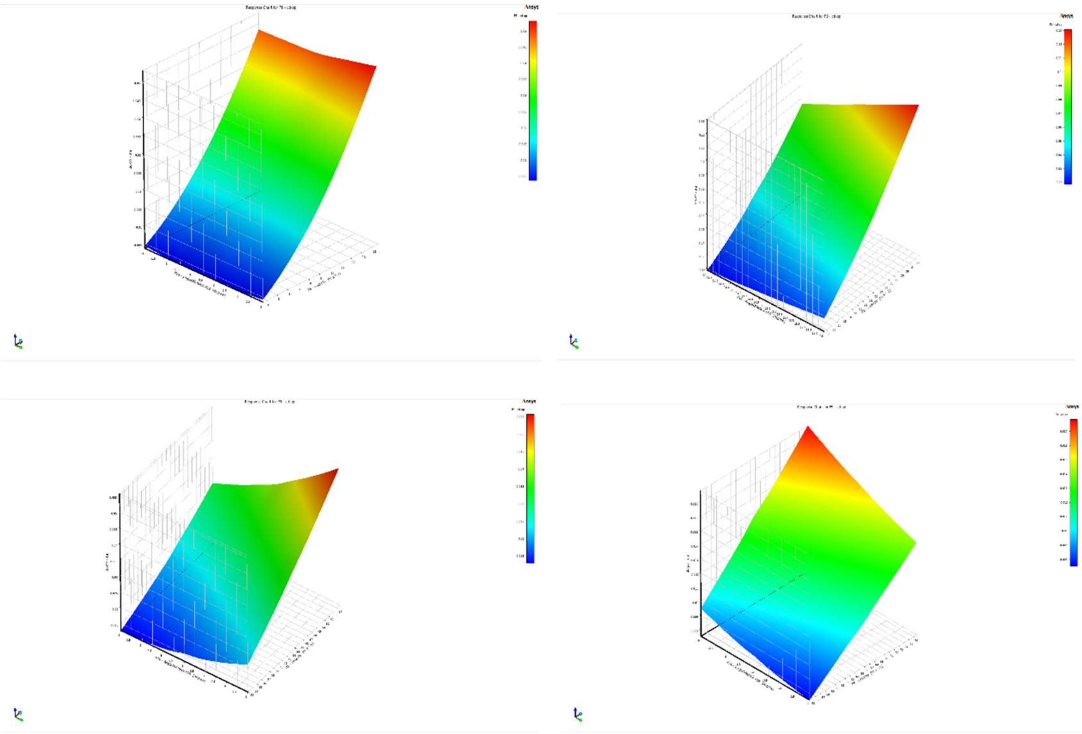

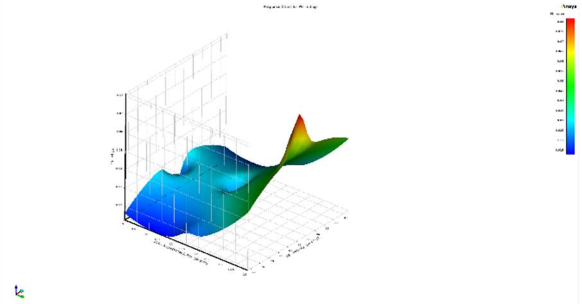

Figure 3.4: RSM-Kriging CD response surface for Run L25 (RSM, Re = 12M, α = 9°–14°, TI = 0.22%–0.71%), illustrating the catastrophic Kriging extrapolation failure documented in Section 3.7. The Kriging surrogate, constrained to exact interpolation through all 15 central-composite DoE training points, 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 α = −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, α = −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.

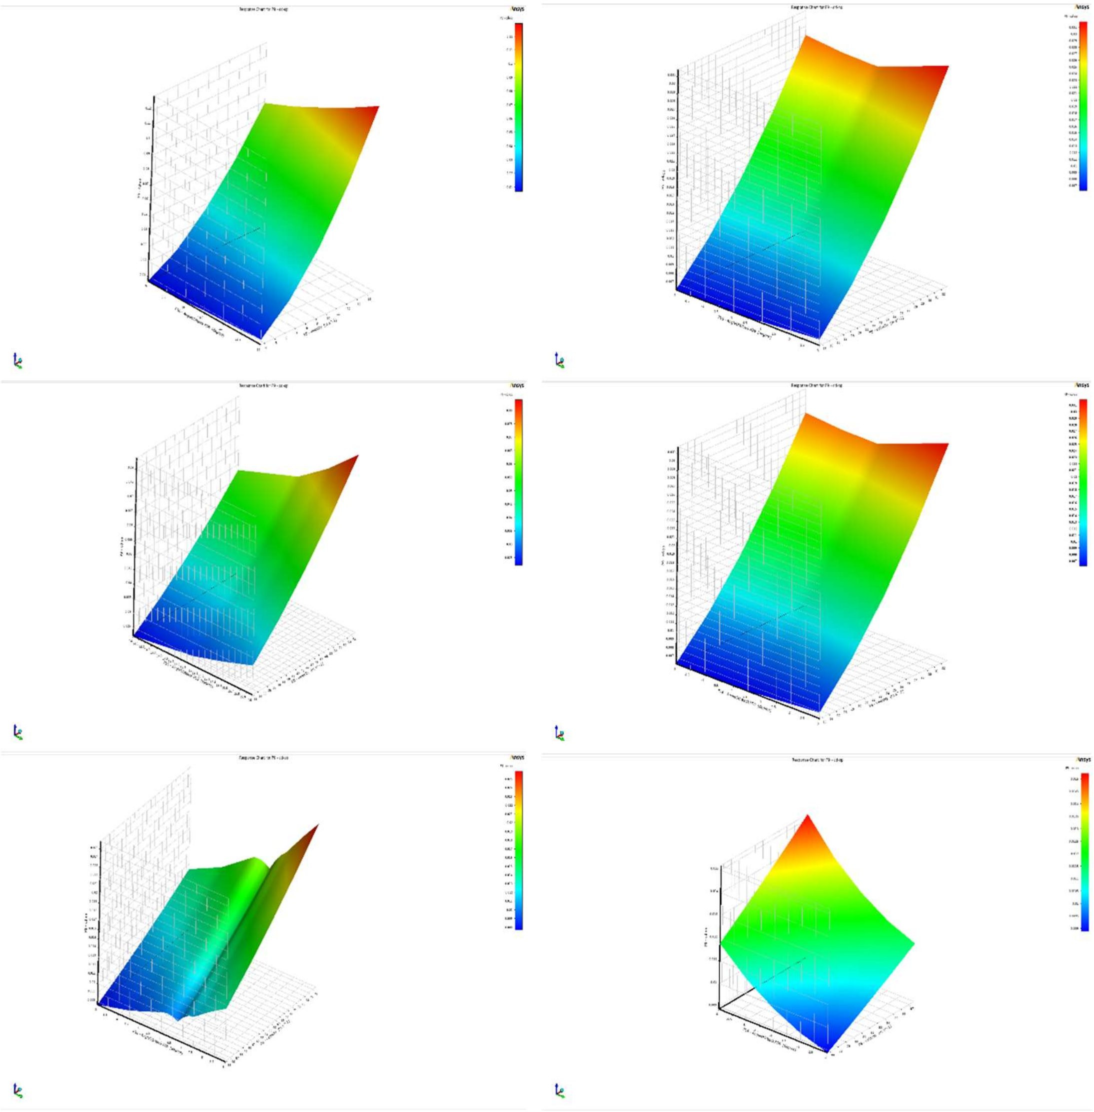

Figure 3.5: Sparse RSM CD response surface for Run L6 (SA model, Re = 3M, α = −2° to +2.5°, V = 15–32 m s^{-1}), generated from only 7 axial training points (1 centre + 2 axial per variable) fitted to a main-effects-plus-quadratic polynomial in ANSYS DesignXplorer. The surface colour maps CD from minimum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

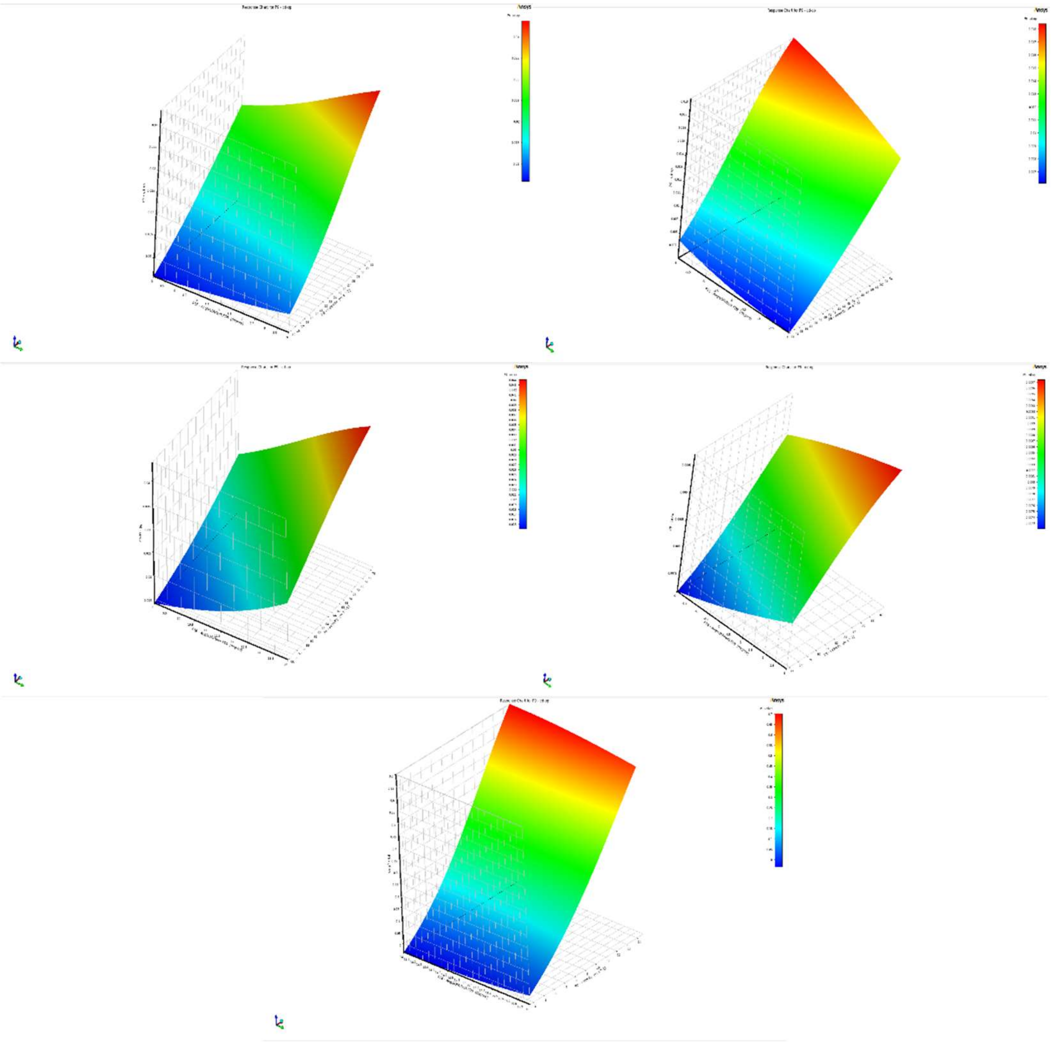

Figure 3.6: Neural Network Screening CD response surface for Run L5 (RSM model, Re = 1M, α = 14°–16°, V = 4–14 m s^{-1}), illustrating the four-stage cascading failure documented in Section 3.9. The ANN was trained on 15 central-composite DoE evaluations contaminated by RSM post-stall divergence, which 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 $\alpha = 6°$ or $\alpha = 12°$ 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 $\alpha = \min(0.653, 0.760, 0.544) = 0.544$ 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.

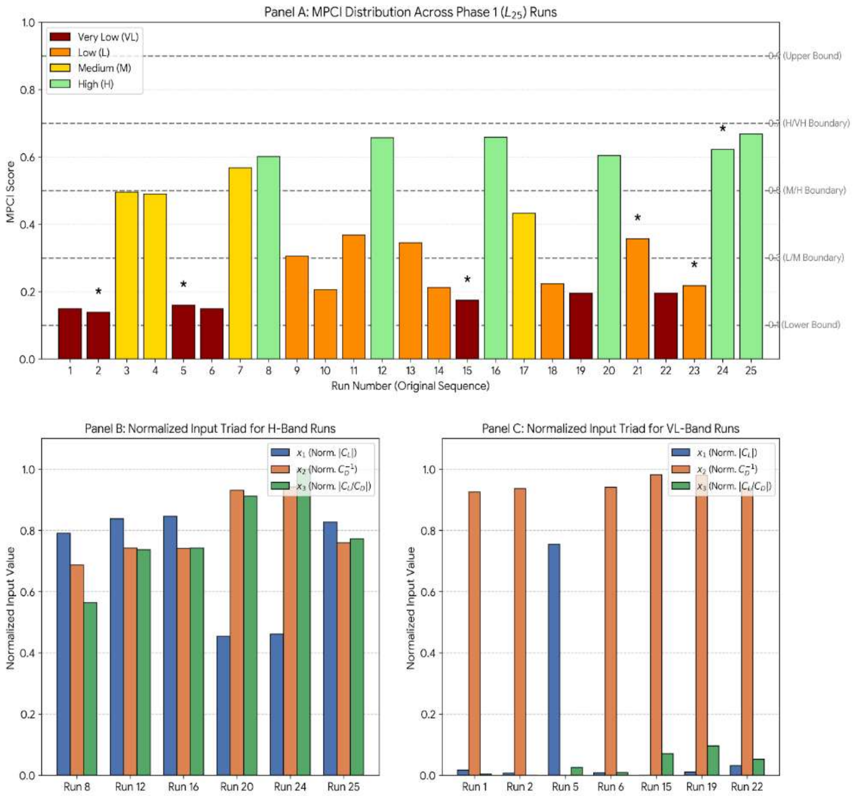

Figure 3.7: Phase 1 MPCI distribution and physical stratification analysis — three-panel figure. **Panel A:** MPCI scores for all 25 L25 runs in original experimental sequence, coloured by quality band: Very Low (dark red, MPCI < 0.20), Low (orange, 0.20–0.30), Medium (yellow, 0.30–0.50), High (green, 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 $\alpha = -4°$ or $\alpha = 0°$) — 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 $(\times10^6)$	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

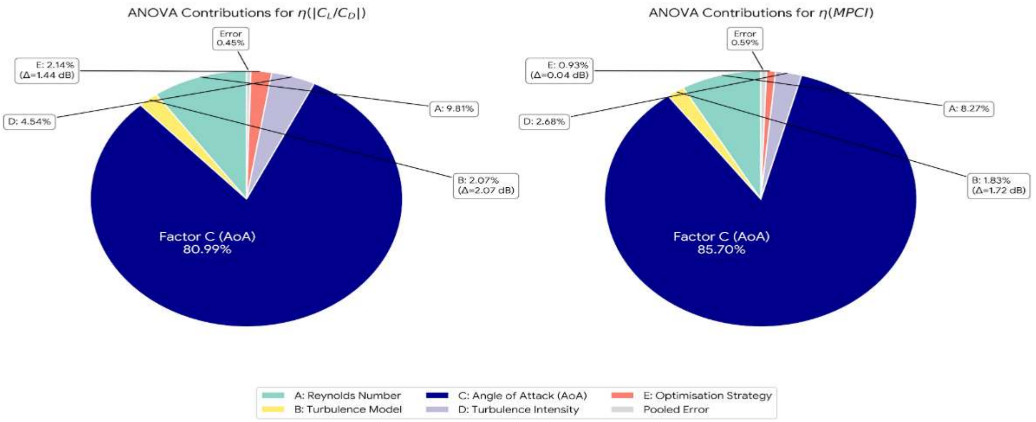

Figure 3.8: Phase 1 one-way ANOVA contribution pie charts for both quality metrics. **Left:** η(|CL/CD|) variance decomposition — Factor C (angle of attack, dark blue) dominates at 80.99%, followed by Factor A (Reynolds number, teal, 9.81%), Factor D (turbulence intensity, yellow, 4.54%), Factor E (optimisation 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 α ≈ −0.75°, cross-confirmed at α ≈ −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 = 6–12×10^6, 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

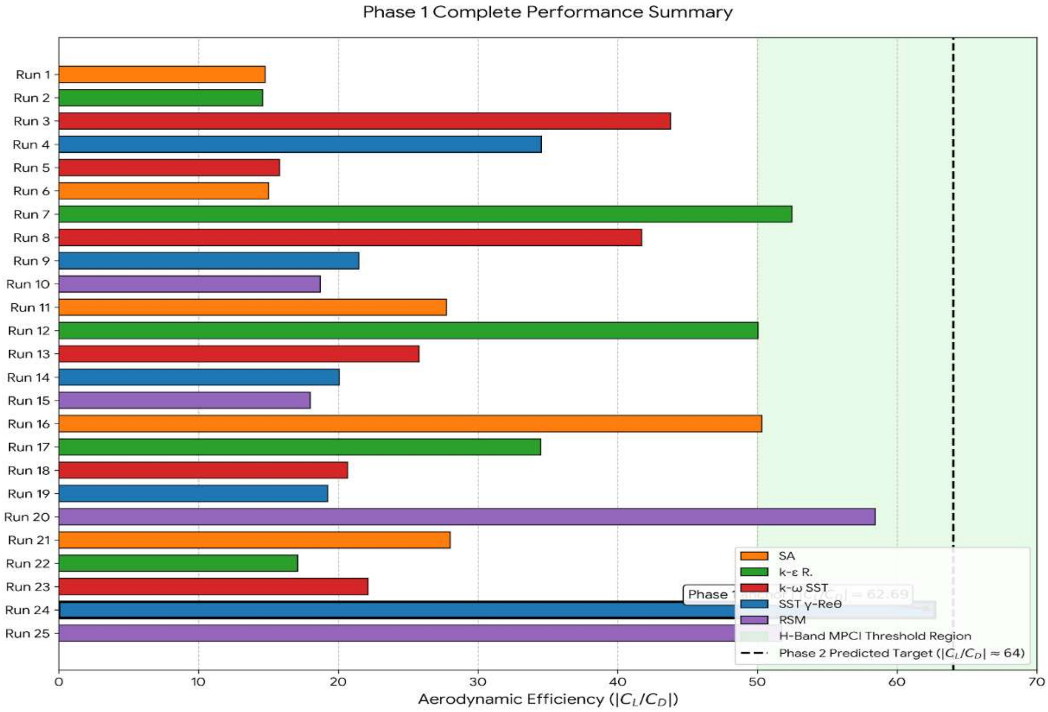

Figure 3.9: Phase 1 complete aerodynamic efficiency summary — horizontal bar chart of |CL/CD| for all 25 L25 runs in their original experimental sequence (Run 1–25, y-axis), coloured by turbulence model: SA (orange), k-ε Realizable (green), k-ω SST (red), SST γ–Reθ (blue), RSM (purple). The green shaded 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 = 6–12×10^6 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 = 6–12×10^6): 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 = 6×10^6 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 $(i,j) \in \{1,2,3\}^2$ 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%.

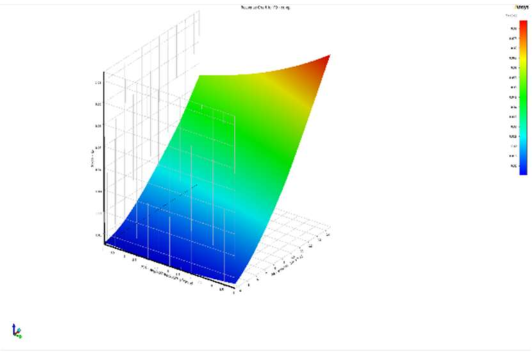

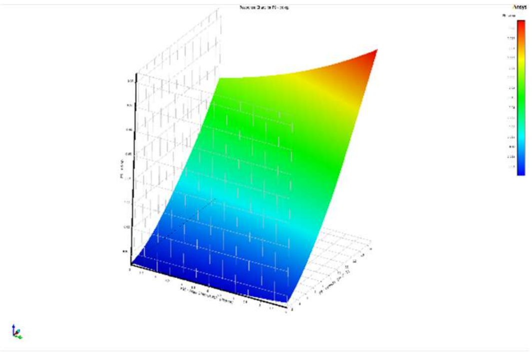

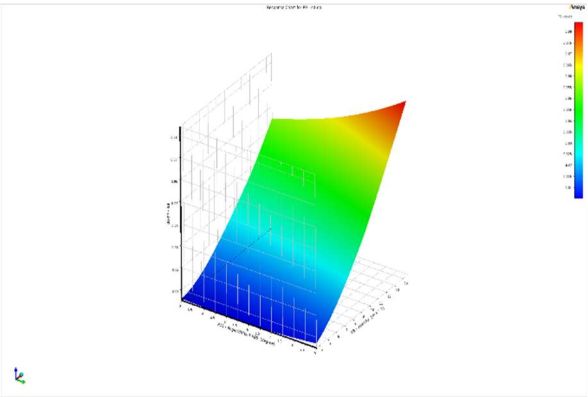

Figure 4.1: Phase 2 Sparse RSM CD response surfaces for the three Reynolds number cells, generated from 7-point axial designs within each cell's local (α, V) bounds. All three surfaces show CD as a function of angle of attack α (front axis, 3°–9°) and velocity V (right axis), with the colour scale mapping 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 ($\times 10^6$)	35.063 (6M)	35.532 (9M)	35.847 (12M)	0.784	2	**L3 = 12×10^6**
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 = 12×10^6, 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 = $6–12\times10^6$. 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 dB^2 — 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:

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

Confirmation run results. A direct ANSYS Fluent k-ω SST simulation at Re = 12×10^6 (V∞ = 87.644 m/s), α = 8°, TI = 0.10°, V_ref = 87.644 m/s, A_ref = 2.0 m^2 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| = |η̂_opt − η_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%**
Additive model error	—	0.022 dB (0.26%)	—	Additivity validated
Convergence	< 10^{-5} predicted	< 10^{-5} confirmed ✓	< 10^{-5} ✓	Attached flow confirmed

Continued on page no. 42

Table 4.5 Continued

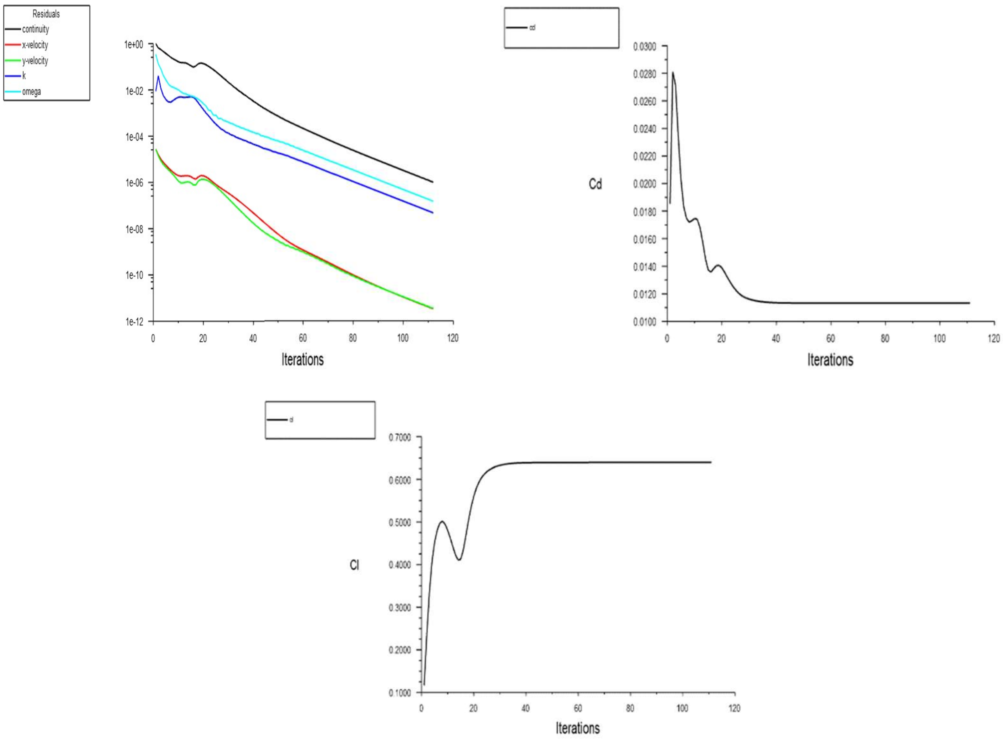

Figure 4.2: Phase 2 k-ω SST simulation convergence histories (representative base run, Re = 12×10^6, α = 4°, TI = 0.05%, ANSYS Fluent 2024 R1), demonstrating the convergence quality achieved uniformly across all nine Phase 2 L9 simulations. **Top:** Scaled residual convergence history for all five transported 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 = 12×10^6, 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.

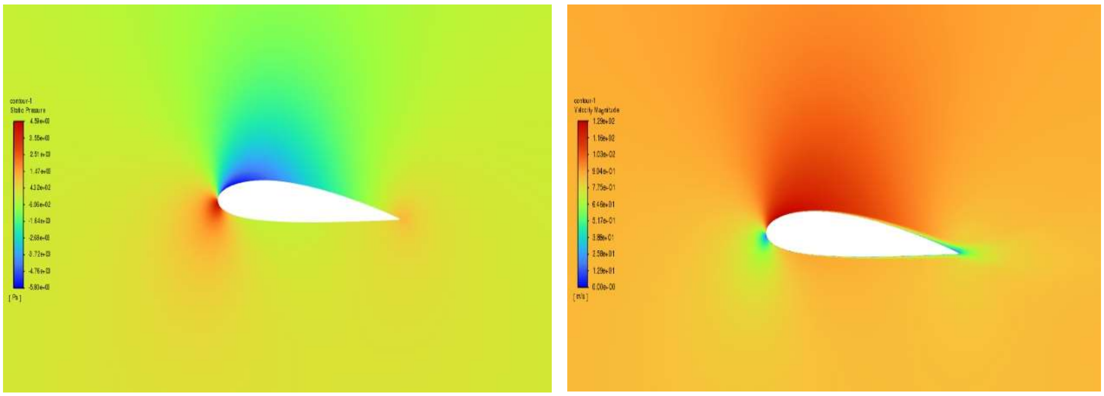

Figure 4.3: Phase 2 confirmed optimum flow field — NACA 2415 at $Re = 12\times10^6$, α = 8°, TI = 0.10%, k-ω SST (ANSYS Fluent 2024 R1, production Mesh M4).

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∞ = 129/87.644 ≈ 1.47, Cp,min ≈ −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 α = 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 $Re = 12\times10^6$, $\alpha = 8°$, 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 $Re = 12\times10^6$ 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 m^2, 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 × b_computational = 2.0 × 1.0 = 2.0 m^2 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 × b = 2.0 × 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 = 12×10^6, α = 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.

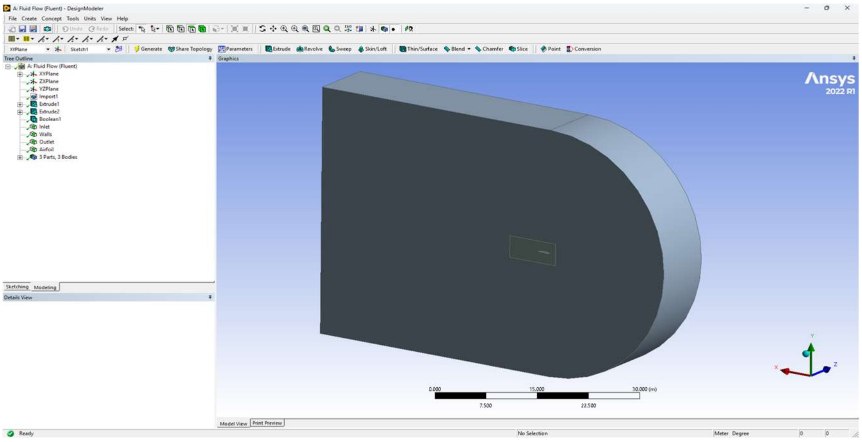

Figure 5.1: Phase 3 three-dimensional computational domain for the NACA 2415 finite wing (b = 0.5 m, c = 2.0 m, AR = 0.25), constructed in ANSYS DesignModeler 2022 R1. The isometric view shows the C-type domain geometry: the semicircular velocity inlet face (radius R = 20c = 40 m, curved front), the 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.

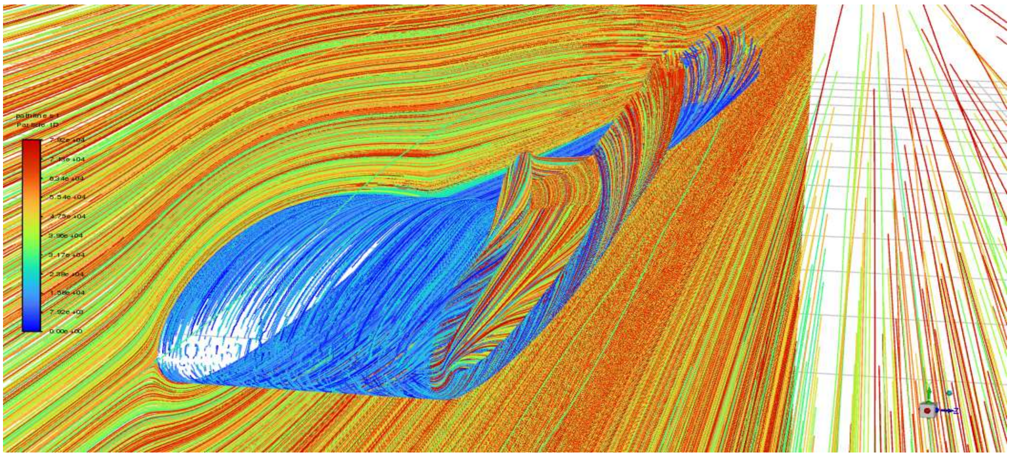

Figure 5.2: Velocity pathline distribution for the Phase 3 NACA 2415 finite wing (b = 0.5 m, c = 2.0 m, AR = 0.25, Re = 12×10^6, α = 8°, TI = 0.10°, k-ω SST), coloured by Particle ID — blue (low ID) = near-surface particles seeded close to the aerofoil; orange/red (high ID) = outer flow and far-field 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 = CL/(π × AR) × V∞ = 0.12921/(π × 0.25) × 87.644 ≈ **14.38 m/s**

This corresponds to ε = arctan(14.38/87.644) ≈ **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.

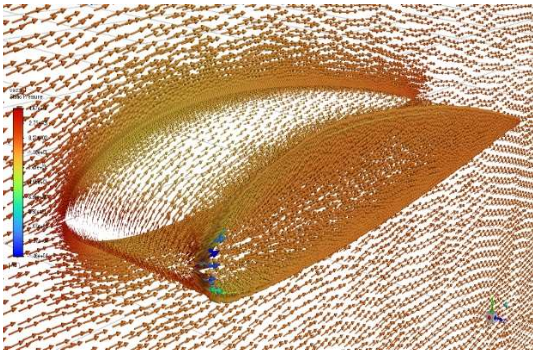

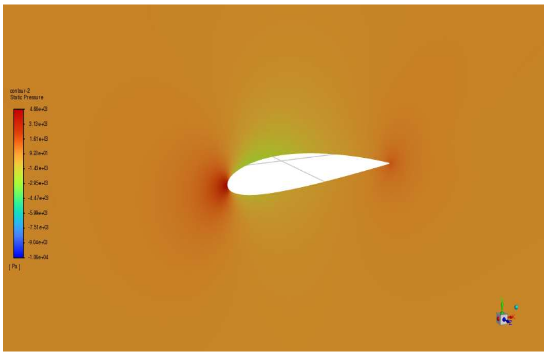

Figure 5.3: Static pressure field for the Phase 3 NACA 2415 finite wing (b = 0.5 m, c = 2.0 m, AR = 0.25, Re = 12×10^6, α = 8°, k-ω SST), providing the thermodynamic proof of the tip vortex formation mechanism. **Panel A (top):** Three-dimensional static pressure vector field — velocity direction vectors 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 ΔP = P_lower − P_upper ≈ (+1,500) − (−7,000) ≈ **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 (Cp,min ≈ −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 Cp,lower ≈ +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 Cp,tip ≈ −14,600/4,704.9 ≈ −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:

ΔP_tip = P_lower − P_upper ≈ (+1,500) − (−7,000) ≈ **8,500 Pa** (≈ 1.81 × q∞)

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.

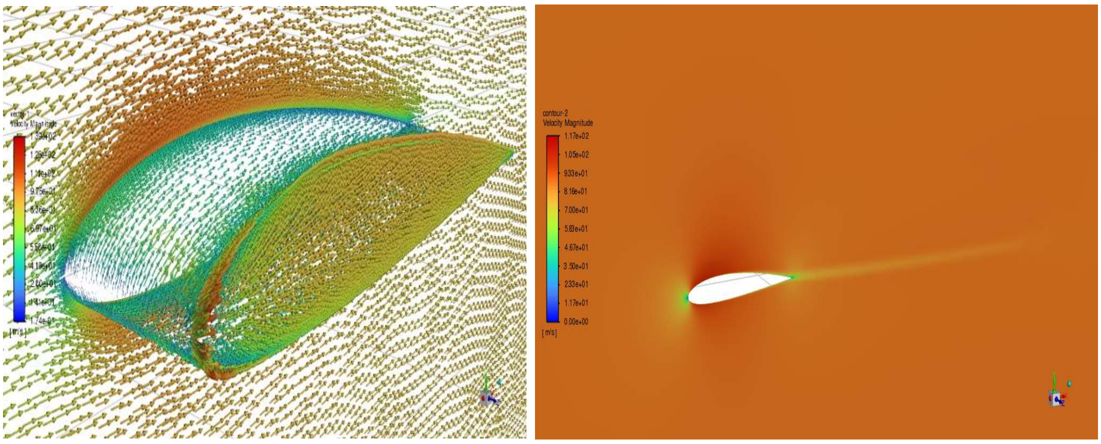

Figure 5.4: Velocity magnitude field for the Phase 3 NACA 2415 finite wing (b = 0.5 m, c = 2.0 m, AR = 0.25, Re = 12×10^6, α = 8°, k-ω SST), providing the kinematic proof that completes the three-visualisation physical argument for the 86.9% efficiency collapse. **Panel A (top):** Three-dimensional velocity 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∞ cos 8°, V_y = V∞ 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∞ = 139/87.644 ≈ 1.587, Cp,min ≈ −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 ≈ 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 ≈ 43.9 N (profile drag at effective α ≈ −1.43°) plus D_induced ≈ 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:

$$V\_peak/V_\infty = 139/87.644 \approx \mathbf{1.586}$$

$$Cp,min = 1 - (V\_peak/V_\infty)^2 = 1 - (1.586)^2 \approx \mathbf{-1.515}$$

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 $CD\_profile \approx 0.00933$ (Run 23: $\alpha = 0°$, $Re = 12M$, $CD = 0.009332$), 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:

$\varepsilon = \arctan(w\_induced/V\infty) = \arctan(14.38/87.644) \approx \mathbf{9.33°}$

The effective angle of attack is then:

$\alpha\_eff = \alpha\_geometric - \varepsilon = 8.00° - 9.33° \approx \mathbf{-1.33°}$ (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:

$CL{,}3D \approx (dCl/d\alpha)\_2D \times (\alpha\_eff - \alpha L0) \times [AR/(AR+2)] = 0.1097 \times (-1.33 - (-2.08)) \times [0.25/2.25] = 0.1097 \times 0.75 \times 0.1111 \approx \mathbf{0.0914}$

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 α_eff ≈ −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∞ = 87.644 m/s
Drag (N)	—	70.096	—	Physical drag force; D_induced ≈ 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°	—	ε = CL,3D/(π × AR) × (180/π); 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

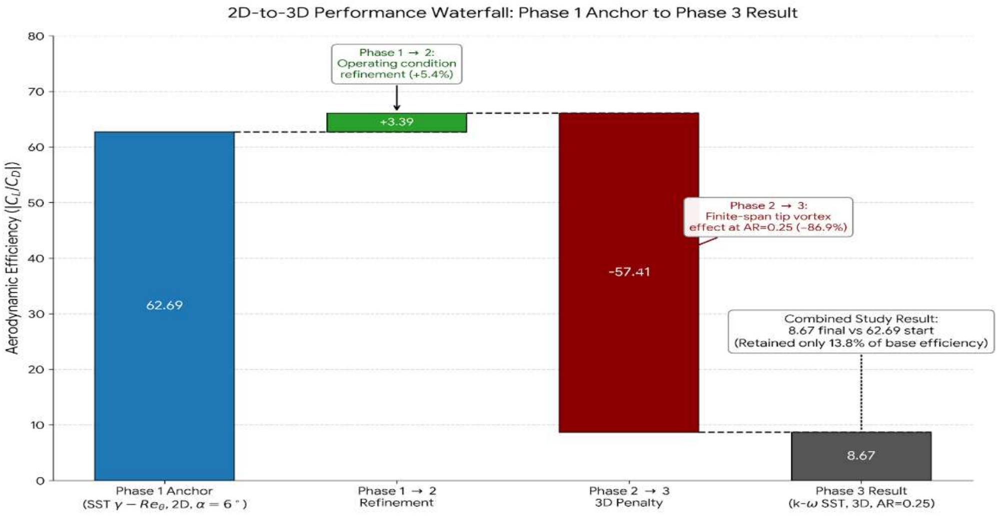

Figure 5.5: Three-phase aerodynamic efficiency waterfall chart — the complete study performance narrative from Phase 1 anchor to Phase 3 finite-wing result. **Blue bar (Phase 1 Anchor):** |CL/CD| = 62.69, achieved by SST γ–Reθ at Re = 12M, α = 6°, TI = 0.10°, two-

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.

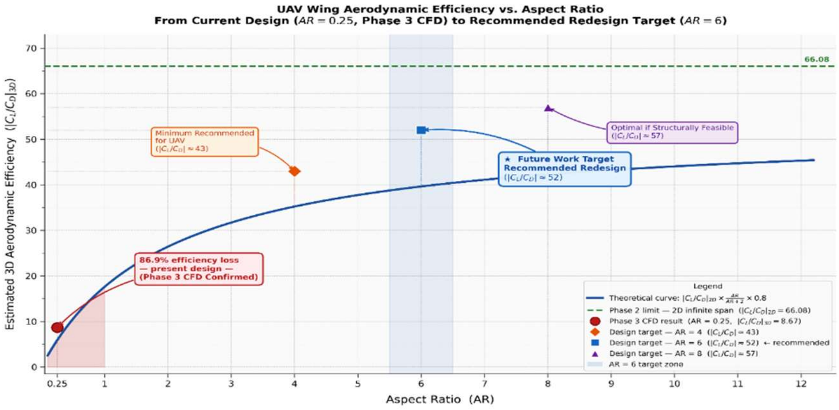

Note: smooth curve uses simplified formula $|C_L/C_D|_{3D} = 66.08 \times AR/(AR + 2) \times 0.8$. Diamond/square/triangle markers show Chapter 6 estimates from a more complete analysis.

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 $Re = 9\times10^6$ 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).

27

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 **$Re = 12\times10^6$, $\alpha = 8°$, $TI = 0.10°$**, 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 **$\rho\_C \approx 0.00\%$ with $F = 0.04$** — definitively confirming that TI specification within 0.05%–0.50% is aerodynamically irrelevant for k-ω SST at $Re = 6$–12×10^6. 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 $\varepsilon \approx 9.33°$ 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 $\alpha \approx -0.75°$ 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 m^2) 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 = 1×10^6 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 = 6–12×10^6 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 10^5–4×10^5) 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 = 3–6×10^6 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 ≈ 10^5–4×10^5, 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 = 10^5–5×10^5, 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 = 10^5–4×10^5 (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.

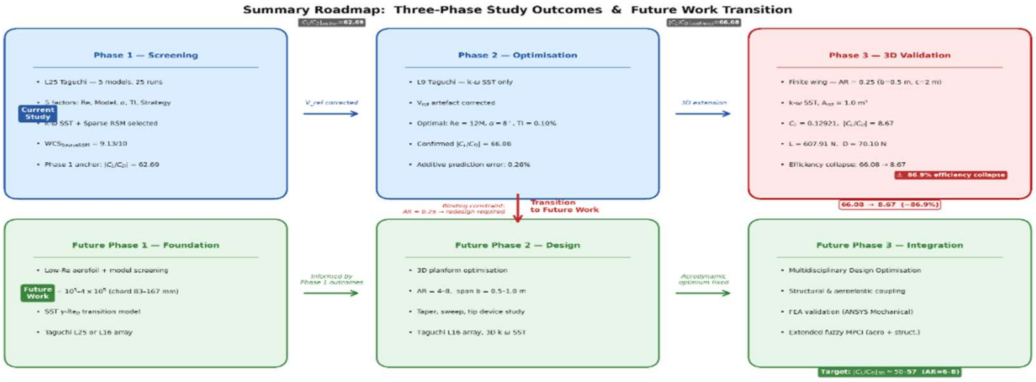

Figure 6.1: Three-phase study outcomes and future work transition roadmap. **Top row (Current Study, blue/red boxes):** Phase 1 Screening — Taguchi L25(5^5), 25 RANS simulations across five turbulence models, k-ω SST and Sparse RSM selected (WCS = 9.13/10), Phase 1 anchor |CL/CD| = 62.69. Phase 2 Optimisation — 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