Aerospace Science and Technology 178 (2026) 113022

Contents lists available at ScienceDirect

Aerospace Science and Technology

journal homepage: www.elsevier.com/locate/aescte

Original article

A locally validated surrogate-assisted design strategy for a hypersonic waverider under coupled aerodynamic and aerothermal constraints

Murat Metehan Türkoğlu [a,*], Emirhan Dönmez [a], Ibrahim Özkol [b]

[a] İstanbul Gelisim University, Aeronautical Engineering, Avcılar, İstanbul, 34310, Turkey
[b] Istanbul Technical University, Aeronautical Engineering, Maslak, İstanbul, 34475, Turkey

ARTICLE INFO

Communicated by Dr Mehdi Ghoreyshi

Keywords:
Hypersonic waverider
Surrogate-assisted optimization
Aerothermal constraints
Local design-space refinement
Waverider design
Flight-envelope assessment

ABSTRACT

This study aims to develop and assess a physically interpretable, locally validated surrogate-assisted design strategy for a hypersonic waverider under coupled aerodynamic and aerothermal constraints. Hypersonic waverider design requires aerodynamic efficiency and preliminary aerothermal feasibility to be resolved within a coupled and physically interpretable framework. This study presents a bounded local design investigation for a hypersonic waverider at Mach 10, 30 km altitude, and $\alpha = 4°$, using a physics-based evaluation chain, local design-space refinement, target-specific surrogate modeling, and post-optimization full-physics reevaluation. The local problem is defined by four geometric variables: shock angle β, width-to-length ratio W/L, leading-edge radius r_{le}, and shaping exponent n_{power}. The results show that β and W/L act primarily as aerodynamic performance drivers, whereas r_{le} acts mainly as a thermal-feasibility variable with only weak influence on L/D over the examined local range. The best feasible solution is boundary-controlled rather than interior, occurring near the upper bounds of β, W/L, and n_{power}, and near the smallest thermally admissible value of r_{le}. Two final reference designs are retained: a peak feasible point and a recommended nominal point. The nominal design preserves nearly the same aerodynamic efficiency as the peak feasible design while increasing the thermal margin from 0.19 to 4.04 W/cm^2. A local surrogate framework trained only within the trusted four-dimensional design box reproduces the optimizer-relevant structure of the problem under random hold-out, blocked hold-out, and post-opt reevaluation. Both retained reference points are shown to remain within the nominal admissible hypersonic flight corridor, while a heat-flux-bias sensitivity check indicates that the recommended nominal point has a larger screening-level thermal reserve. These results establish a physically interpretable and locally validated surrogate-assisted design strategy for hypersonic waverider development under coupled aerodynamic and aerothermal constraints.

1. Introduction

Hypersonic waveriders continue to attract sustained interest because they provide a physically structured route to high lift-to-drag performance through compression-lift generation and shock-attached or near-attached flow organization [1–3]. Since the classical viscous-optimized studies, the waverider concept has evolved into a broader design family encompassing parametric geometry generation, osculating-flow-based constructions, inverse design, and airframe/inlet integration [1,4–6]. Recent review and methodology studies have confirmed that waveriders remain central to hypersonic vehicle design because they naturally couple external shape, compression topology, aerodynamic efficiency, and propulsion-relevant forebody flow organization [2,3]. At the same time, these same advantages make the design problem inherently constrained, since the geometric features that enhance aerodynamic performance also interact strongly with viscous effects, thermal loading, and practical configuration limits. Accordingly, contemporary waverider design must be treated as a coupled aerothermodynamic problem rather than as a purely aerodynamic optimization task.

This coupled character becomes especially important when off-design behavior, leading-edge bluntness, and local geometric variation are considered simultaneously. Off-design analyses have shown that waverider lift-to-drag behavior is sensitive to operating condition and local geometry [7], while numerical and experimental studies on blunt leading edges have demonstrated that thermal relief is typically achieved through changes that may also alter aerodynamic response [8]. More recent optimization studies have reinforced this point by treating bluntness as an active design variable rather than a purely geometric correction, particularly when aerodynamic efficiency and heating must be balanced within a constrained design space [9]. In parallel, reduced-order

* Corresponding author.
E-mail addresses: mmturkoglu@gelisim.edu.tr (M.M. Türkoğlu), 200414421@ogr.gelisim.edu.tr (E. Dönmez), ozkol@itu.edu.tr (I. Özkol).

https://doi.org/10.1016/j.ast.2026.113022
Received 11 April 2026; Received in revised form 16 June 2026; Accepted 22 June 2026
Available online 26 June 2026

and surrogate-based frameworks have shown clear potential for accelerating hypersonic design analysis [5,10]. However, these developments do not remove the underlying engineering requirement to interpret how specific local design variables govern aerodynamic gain and thermal admissibility within a physically interpretable and explicitly bounded evaluation chain.

Recent research has expanded hypersonic design toward direct inverse-design methods, wide-speed-range aerodynamic optimization, surrogate-assisted shape–trajectory co-design, and dynamic surrogate-based vehicle/engine integration [5,11–14]. These studies demonstrate the growing importance of surrogate and data-driven methods for reducing computational cost and enlarging the accessible design space. Nevertheless, comparatively less attention has been given to tightly bounded local waverider studies in which the roles of the active geometric variables are first interpreted physically, the trusted feasible neighborhood is then defined explicitly, and the resulting surrogate is assessed not only by nominal regression accuracy but also by blocked hold-out testing, optimizer-level backchecking, and post-optimization true-physics confirmation. A further issue is that the geometric scope of many studies is not always stated sharply enough. In the present work, the geometry-generation framework is treated explicitly as OCM-inspired and parametric rather than as a full classical streamline-tracing implementation of the Osculating Cone Method. Under this more careful framing, a clear opportunity remains for a locally validated surrogate-assisted design strategy that is physically interpretable, computationally efficient, and methodologically transparent. Recent studies have also continued to broaden hypersonic-vehicle research toward trajectory-planning, guidance/control, and optimization-oriented methodologies, further confirming the need for timely reduced-order, surrogate-assisted, and physics-interpretable design frameworks for constrained hypersonic configurations [15,16].

More specifically, the present study addresses six closely related gaps in the existing waverider-design literature. First, the local roles of active geometric variables are not always separated explicitly into aerodynamic-performance and thermal-feasibility effects. Second, the trusted local design neighborhood is often not defined sharply before surrogate construction. Third, surrogate models are frequently judged mainly by nominal regression accuracy, without complementary blocked hold-out testing and optimizer-level backchecking. Fourth, surrogate-selected candidates are not always confirmed through post-optimization true-physics reevaluation. Fifth, the distinction between a mathematical peak feasible point and a more conservative nominal engineering reference is often left implicit. Sixth, the fidelity level of OCM-inspired parametric geometries is not always separated clearly from full classical streamline-traced Osculating Cone Method implementations.

Relative to prior waverider optimization and surrogate-assisted hypersonic design studies, the present work is positioned not by expanding the global design space, but by tightening the local validation logic inside a bounded and interpretable trust region. Its main methodological distinction is the sequential use of physics-based local refinement, target-specific surrogate validation, optimizer-level backchecking, and mandatory post-optimization true-physics reevaluation. This combination allows the retained surrogate to be assessed not only as a regression model, but also as a local decision-support tool that preserves variable-role interpretation, candidate ranking, and feasibility structure inside the stated design box.

In this context, the present study investigates a hypersonic waverider within a bounded local design regime at Mach 10, 30 km altitude, and $\alpha = 4°$. The bounded scope is not treated as a limitation of convenience, but as a deliberate methodological choice. By freezing the operating condition and restricting the geometry to a trusted four-dimensional local design box, the aerodynamic, thermal, geometric, and surrogate-modeling effects can be interpreted without making unsupported global extrapolation claims.

Accordingly, the specific objectives of the study are defined as follows:

1. To define and justify a bounded local waverider design box at the frozen Mach–altitude–angle-of-attack condition.
2. To identify the physical roles of the active geometric variables, especially the aerodynamic influence of β and W/L and the thermal-feasibility role of r_{le}.
3. To construct a target-specific local surrogate framework for the retained design box and validate it using random hold-out testing, blocked hold-out validation, optimizer-level backchecking, and post-optimization true-physics reevaluation.
4. To distinguish the peak feasible design from a recommended nominal reference design with improved thermal reserve.
5. To place the retained designs within a broader hypersonic flight-corridor context while explicitly limiting the conclusions to the stated local trust region.

These objectives are directly linked to the limitations identified above. Objective 1 responds to the need for a sharply bounded and reproducible trust region. Objective 2 addresses the lack of explicit local variable-role interpretation. Objective 3 responds to the limitation of relying only on nominal surrogate regression accuracy. Objective 4 addresses the practical distinction between local mathematical optimality and engineering usefulness near active constraints. Objective 5 prevents over-generalization by separating local design insight from global Mach–altitude–angle-of-attack prediction.

The expected outcome of this formulation is therefore not a universal optimum, but a locally credible design interpretation. The analysis is expected to reveal which variables control aerodynamic performance, which variables control thermal accessibility, whether the feasible optimum is interior or boundary-controlled, and whether a local surrogate can preserve the true-physics candidate ranking inside the trusted design box. Under this scope, the contribution of the paper is a physically interpretable and locally validated surrogate-assisted optimization strategy for waverider design under coupled aerodynamic and aerothermal constraints, rather than a universal hypersonic surrogate or a global waverider optimum.

The remainder of this manuscript is organized as follows. Section 2 defines the fixed operating regime, the retained local design variables, the physics-based evaluation chain, and the surrogate validation workflow. Section 3 reports the governing-module verification, literature-based aerodynamic-efficiency comparison, and numerical-consistency checks used to bound the credibility of the reduced-order framework. Section 4 presents the retained local design-space behavior, surrogate-guided candidate refinement, true-physics confirmation, drag-share interpretation, and flight-envelope embedding. Section 5 discusses the physical interpretation of the retained optimum family, the thermal-margin versus peak-performance trade-off, and the trust-region limitations of the study. Section 6 summarizes the main findings and outlines the required future work, while the appendices provide supplementary implementation details, constrained-sweep slices, single-radius corridor views, and reproducibility information.

2. Methodology

This section describes the methodological framework used in the bounded local waverider study. It first defines the fixed operating regime, the retained local design box, and the feasibility constraints adopted for the final analysis. It then summarizes the coupled physics-based evaluation chain, the local refinement logic, and the surrogate construction and validation workflow used to identify and confirm the retained reference designs. For clarity, Algorithm 1 summarizes the full retained procedure.

2.1. Baseline regime, design variables, and constraints

The present study is formulated as a bounded local design investigation rather than as a global hypersonic optimization campaign over

Table 1
Fixed operating condition, retained local design box, and feasibility constraints used in the final $\alpha = 4°$ study.

Item	Symbol	Value / Range
Design Mach number	M_∞	10
Altitude	h	30 km
Angle of attack	α	4°
Reference length	L_{ref}	10 m
Shock angle	β	[23.0, 25.0]°
Width-to-length ratio	W/L	[0.72, 0.80]
Leading-edge radius	r_{le}	[23.25, 27.00] mm
Shaping exponent	n_{power}	[0.85, 1.00]
Heat-flux constraint	$q_{\max}$	≤ 245 W/cm^2
Volume constraint	V	≥ 3.0 m^3
Lift constraint	C_L	≥ 0.01

multiple flight regimes. All frozen core analyses were conducted at a fixed operating condition of Mach 10, 30 km altitude, and $\alpha = 4°$, with a reference vehicle length of 10 m. This regime was retained because it provided a numerically stable and physically interpretable basis for the final local refinement campaign, surrogate construction, and post-optimization true-physics confirmation. Accordingly, the claims of the manuscript are restricted to this local regime and its associated bounded design neighborhood rather than to a universal Mach–altitude–angle-of-attack surrogate.

Within this frozen regime, the final local study is defined in terms of four geometric design variables: shock angle β, width-to-length ratio W/L, leading-edge radius r_{le}, and shaping exponent n_{power}. In the frozen local design workflow at $\alpha = 4$, the trusted local design box is $\beta \in [23.0, 25.0]°$, $W/L \in [0.72, 0.80]$, $r_{le} \in [23.25, 27.00]$ mm, and $n_{power} \in [0.85, 1.00]$. These bounds were not chosen arbitrarily; rather, they were retained after local refinement and post-opt validation had identified this neighborhood as the region in which aerodynamic behavior, thermal feasibility, and surrogate interpretation remained physically consistent and numerically robust. Here, n_{power} is not a conical-flow state variable. It is the exponent of the retained leading-edge seed-curve parametrization and acts only on the spanwise growth and planform curvature of the conical-shock-surface leading-edge seed. Smaller values produce a more pointed seed curve, whereas values closer to unity produce a fuller and locally blunter planform within the same retained geometric family.

The frozen constraint set used for the final local study was $q_{\max} \leq 245$ W/cm^2, $V \geq 3.0$ m^3, and $C_L \geq 0.01$. Under this formulation, the methodology is directed toward identifying and interpreting a trusted feasible neighborhood, rather than toward claiming a universal optimum over the broader hypersonic operating space.

Table 1 summarizes the frozen baseline regime, active design variables, and local feasibility constraints adopted in the present study.

2.2. Physics-based evaluation framework

Each candidate design was evaluated through a coupled physics-based framework linking atmospheric conditions, conical-flow solution, geometry generation, aerodynamic prediction, and aerothermal analysis. This evaluation chain served as the common reference model throughout the local refinement, surrogate training, and post-optimization confirmation stages.

Atmospheric properties were first computed at the prescribed altitude using a comprehensive atmosphere model based on the US Standard Atmosphere in the lower and middle atmosphere together with an extended high-altitude approximation above 86 km. The framework returned temperature, pressure, density, speed of sound, viscosity, and mean free path, which were then used to determine whether high-temperature real-gas support or rarefied-flow support should be activated. Rarefaction was monitored through the Knudsen number,

$$Kn = \frac{\lambda}{L_{\text{ref}}}, \tag{1}$$

where λ is the mean free path and L_{ref} is the characteristic vehicle length. In parallel, the stagnation-temperature level was used to assess whether calorically perfect gas assumptions remained adequate.

The external flow field was then characterized through a conical-flow solution obtained by integrating the Taylor–Maccoll equations for the prescribed shock angle. In normalized conical coordinates, the governing system is written as

$$\frac{dV_r}{d\theta} = V_\theta, \tag{2}$$

$$\frac{dV_\theta}{d\theta} = \frac{V_\theta^2 V_r - \frac{\gamma-1}{2}\left(1 - V_r^2 - V_\theta^2\right)\left(2V_r + V_\theta \cot\theta\right)}{\frac{\gamma-1}{2}\left(1 - V_r^2 - V_\theta^2\right) - V_\theta^2}, \tag{3}$$

where V_r and V_θ are the nondimensional radial and polar velocity components, respectively. The resulting conical-flow quantities were then used to construct the vehicle geometry through an OCM-inspired parametric quasi-waverider procedure. In the present study, this geometry-generation framework is treated explicitly as OCM-inspired and parametric rather than as a full classical streamline-tracing implementation of the Osculating Cone Method.

The leading-edge definition was revised to avoid any planar interpretation of the conical shock. In conical-flow theory, the shock associated with the generating cone is a conical surface rather than a plane. Therefore, the retained implementation does not define. Instead, the leading-edge seed is treated as a three-dimensional spatial curve constrained to lie on the prescribed conical shock surface. For a cone with its axis aligned with the streamwise x direction, the revised leading-edge seed is written as

$$\begin{aligned}
R_\perp &= \sqrt{y^2 + z^2} = x \tan\beta, \\
y_{le} &= \frac{W}{2}\eta^{n_{\text{power}}}, \\
\phi_{le}(\eta) &= \phi_0 + (\phi_1 - \phi_0)\eta, \\
r_{\perp,le} &= \frac{|y_{le}|}{\sin\phi_{le}}, \\
x_{le} &= \frac{r_{\perp,le}}{\tan\beta}, \\
z_{le} &= r_{\perp,le}\cos\phi_{le}.
\end{aligned} \tag{4}$$

This construction gives $\sqrt{y_{le}^2 + z_{le}^2} = x_{le}\tan\beta$, so that the leading-edge seed lies on the conical shock surface rather than in a planar leading-edge representation. Here, W is the vehicle width, $\eta \in [0, 1]$ is the nondimensional spanwise parameter, β is the prescribed shock angle, n_{power} governs the spanwise growth of the seed curve, and ϕ_{le} defines the azimuthal placement of the seed curve on the conical shock surface. Smaller n_{power} values produce a more pointed leading-edge seed, whereas values closer to unity produce a fuller and locally blunter spanwise growth.

It should be emphasized that Eq. (4) is still part of an OCM-inspired parametric quasi-waverider geometry generator. It is not claimed to be a full classical flow-capture-tube/intersection-curve construction or a complete streamline-traced osculating-cone waverider method. In a classical cone-derived or osculating-cone waverider, the leading edge is obtained as a three-dimensional intersection curve between the flow capture tube and the conical or osculating shock surface. The retained method instead uses a conical-shock-surface seed curve and a downstream conical-flow-informed surface mapping to generate a bounded parametric quasi-waverider family for local design-space screening. The lower surface is generated from this seed by conical-flow-informed downstream mapping, whereas the upper surface is introduced through a volume-preserving freestream-aligned closure with prescribed thickness and leading-edge blunting. The geometry module returns the full surface coordinates together with planform area, wetted area, base area, enclosed volume, and a closed-surface diagnostic mesh. Under this construction, the retained shapes should be interpreted as OCM-inspired parametric quasi-waveriders rather than as full classical streamline-traced osculating-cone waveriders.

For terminology, in a classical cone-derived or osculating-cone waverider construction, the inlet capture curve (ICC) defines the prescribed capture-boundary curve used to construct the osculating conical or shock geometry, while the flow-capture tube (FCT) denotes the three-dimensional streamtube or captured-flow boundary associated with that prescribed capture curve and the selected conical or osculating shock surface. In such a classical construction, the leading edge is obtained as a three-dimensional intersection curve associated with the FCT and the conical or osculating shock surface. The present study does not prescribe an independent ICC and FCT as full classical design inputs. Instead, it uses a conical-shock-surface leading-edge seed curve and a reduced-order downstream surface mapping to generate an OCM-inspired parametric quasi-waverider family for bounded local screening. This distinction is made explicit to avoid interpreting the present geometry generator as a complete FCT/ICC streamline-traced OCM implementation.

For implementation-level completeness, supplementary details on the retained lower-surface mapping, freestream-aligned upper-surface closure, and leading-edge blunting are provided in Appendix A; these supplementary forms make explicit that the retained geometry is an OCM-inspired parametric quasi-waverider construction rather than a full classical streamline-traced osculating-cone realization.

Aerodynamic coefficients were evaluated using a panel-integrated engineering model combining pressure, viscous, and base-drag contributions. Inviscid pressure loads were obtained from a modified Newtonian formulation,

$$C_p = C_{p,\max} \sin^2 \theta, \tag{5}$$

where $C_{p,\max}$ is the stagnation-point pressure coefficient and θ is the local surface inclination angle. Pressure forces were integrated panel-wise over the upper and lower surfaces to obtain the inviscid force components. Skin-friction drag was then estimated through a reference-temperature-based treatment, and base drag was added separately in order to preserve physical interpretability of the final drag structure. Accordingly, the total drag coefficient was written as

$$C_D = C_{D,\text{inv}} + C_{D,f} + C_{D,b}, \tag{6}$$

where $C_{D,\text{inv}}$, $C_{D,f}$, and $C_{D,b}$ denote the inviscid, friction, and base-drag components, respectively. This decomposition is important because the final design interpretation depends not only on total L/D, but also on the stability of the drag shares across the retained candidate set.

Aerothermal analysis was carried out using a combined engineering heating framework. Stagnation-point heating was assessed using Fay–Riddell and Sutton–Graves correlations, while surface heat fluxes were estimated using a reference-enthalpy/reference-temperature treatment over the body surface. A representative local surface heat-flux relation may be written as

$$q = St\, \rho_\infty V_\infty c_p \left(T_{aw} - T_w\right), \tag{7}$$

where St is the local Stanton number, ρ_∞ and V_∞ are the freestream density and velocity, c_p is the specific heat at constant pressure, T_{aw} is the adiabatic wall temperature, and T_w is the wall temperature. A radiative-equilibrium wall-temperature estimate was also included. When stagnation conditions exceeded the calorically perfect range, a real-gas support module was invoked to provide variable specific heats, variable γ, and dissociation-sensitive thermodynamic properties. The framework also included a bridging-function rarefied-gas correction module for slip, transition, and free-molecular regimes; however, the retained local design point remained within the admissible corridor of the continuum-oriented study. The aerothermal model is therefore used as an engineering screening model rather than as a fully coupled boundary-layer/thermal-protection-system solver. In particular, the retained implementation does not solve the viscous boundary-layer equations in a coupled manner with the external inviscid field, does not include finite-rate surface chemistry, and does not distinguish between fully catalytic, partially catalytic, and non-catalytic wall boundary conditions. It also does not model ablation, pyrolysis, recession, material response, or wall-catalyzed recombination heating. Consequently, the reported heat-flux values are used for relative local feasibility screening, radius-sensitivity interpretation, and nominal-versus-peak comparison inside the stated trust region, rather than as certified TPS design loads or high-fidelity surface-temperature predictions.

Under this formulation, the complete evaluation chain provided the reference outputs used to define feasibility, interpret local variable roles, train the local surrogate layer, and confirm the retained candidate set through post-optimization full-physics reevaluation.

2.3. Local refinement and optimization workflow

The broader implementation supports single-objective optimization, multi-objective optimization, and parametric sweep modes over the four active geometric variables β, W/L, r_{le}, and n_{power}. In that generic setting, aerodynamic efficiency, maximum heating, and volumetric performance can be combined within a constrained aerothermodynamic design framework. The retained manuscript workflow, however, does not report this broader capability as a global design campaign. Instead, it focuses on a bounded local refinement strategy in which the same physics-based evaluation chain is applied repeatedly at a fixed operating condition in order to identify a trusted feasible neighborhood and its optimizer-relevant structure.

Within the fixed local regime, the retained design problem may be written in compact form as

$$\max_{\mathbf{x}} \frac{L}{D}(\mathbf{x}), \qquad \mathbf{x} = \begin{bmatrix}\beta & W/L & r_{le} & n_{power}\end{bmatrix}^T \tag{8}$$

subject to

$$q_{max}(\mathbf{x}) \le 245\ \text{W/cm}^2, \qquad V(\mathbf{x}) \ge 3.0\ \text{m}^3, \qquad C_L(\mathbf{x}) \ge 0.01. \tag{9}$$

Under this formulation, the objective of the local campaign is not to claim a universal optimum over the wider hypersonic operating space, but to resolve how the active variables govern feasibility and performance inside a physically interpretable neighborhood.

The first local refinement axis examined the effect of leading-edge radius at fixed geometric settings and multiple angle-of-attack values. This stage showed that r_{le} acts primarily as a thermal-feasibility variable: increasing bluntness reduced peak heating in a systematic manner, while the corresponding variation in L/D remained weak over the examined local range. This behavior established the local thermal boundary and motivated the later distinction between the smallest thermally admissible design and a slightly more conservative nominal alternative.

The second local refinement stage examined the $(\beta, W/L)$ plane at fixed $\alpha = 4°$, with the remaining variables held at locally relevant values. This sweep showed that the dominant aerodynamic gain within the retained neighborhood is driven mainly by these two variables. In particular, the best feasible solution was found on the upper boundary of the local box rather than at an interior stationary point, indicating a boundary-controlled local optimum. This result is important because it shows that the local problem should be interpreted in terms of constrained edge behavior rather than in terms of an unconstrained interior maximum.

For reproducibility, the retained local search was implemented as a finite local candidate-generation and screening workflow rather than as an unconstrained global stochastic optimization campaign. Candidate feasibility was enforced by a hard constraint filter using $q_{\max} \le 245$ W/cm^2, $V \ge 3.0\,\text{m}^3$, and $C_L \ge 0.01$. No infeasible candidate was retained by penalty relaxation. The surrogate layer was used only to accelerate local candidate discovery inside the trusted four-dimensional design box, and all retained decision candidates were subsequently reevaluated through the frozen physics chain. The retained local search and confirmation settings are summarized in Table 2.

Table 2
Reproducibility settings for the retained local search and post-optimization confirmation workflow.

Item	Retained setting
Operating condition	$M_\infty = 10$, $h = 30$ km, $\alpha = 4°$
Design variables	β, W/L, r_{le}, n_{power}
Design box	Table 1
Search type	Finite local candidate screening with surrogate-assisted refinement
Objective	Maximize L/D among feasible candidates
Constraint handling	Hard feasibility filter; no penalty-relaxed infeasible retention
Heat-flux constraint	$q_{max} \leq 245$ W/cm^2
Volume constraint	$V \geq 3.0$ m^3
Lift constraint	$C_L \geq 0.01$
Final confirmation	Full frozen-chain reevaluation of retained candidates
Randomness	Used only for surrogate data splitting and validation audits
Random seed	Fixed and reported in the shared reproduction package
Stopping criterion	Exhaustion of retained local candidate set and no higher-L/D feasible confirmed candidate

Based on these refinement stages, two final reference designs were retained for the remainder of the study: a peak feasible point and a recommended nominal point. The peak feasible point corresponds to the highest-performing solution that still satisfies the retained local constraints, whereas the nominal point preserves nearly the same aerodynamic quality while introducing additional thermal margin. This retained-design distinction is carried forward into the surrogate-validation and post-optimization confirmation stages, where the local optimum structure is assessed not only in terms of nominal performance, but also in terms of engineering robustness and interpretability.

2.4. Local surrogate construction and validation strategy

A local surrogate framework was constructed only for the retained four-dimensional design box defined by β, W/L, r_{le}, and n_{power}, while Mach number, altitude, angle of attack, and reference length were held fixed at their baseline values. This restriction was deliberate: the surrogate was intended as a local engineering accelerator for the Mach 10, 30 km, $\alpha = 4°$ study, rather than as a global Mach–altitude–angle-of-attack meta-model. Accordingly, all surrogate-related claims in the manuscript are restricted to the retained local neighborhood identified through the refinement workflow.

The surrogate dataset was generated by repeated evaluations of the same physics-based reference chain over the retained local design box, so that all surrogate targets were derived directly from the model used in the refinement and post-optimization confirmation stages. The retained surrogate targets were the principal aerodynamic, thermal, and geometric outputs required by the local design problem, namely L/D, $q_{\max}$, volume, C_L, and C_D. For each target, a separate surrogate was trained according to the general mapping

$$\hat{y}_t = f_t(\beta, W/L, r_{le}, n_{power}), \tag{10}$$

where t denotes the response quantity of interest. This target-specific strategy was adopted deliberately rather than using a single global regressor, because the local smoothness, scale, and optimizer relevance of the retained responses were not identical across aerodynamic, thermal, and geometric outputs.

Within the retained implementation, several target-specific surrogate families were evaluated, including ridge-regularized polynomial models, radial-basis-function models, and Gaussian-process regression models with an automatic-relevance-determination squared-exponential kernel. The final retained surrogate selected by the validation procedure was the ARD squared-exponential Gaussian-process model for all five target responses. This choice was adopted because it provided the lowest hold-out prediction errors while preserving smooth local interpolation behavior inside the trusted four-dimensional design box. The surrogate layer is therefore interpreted as a local interpolation and candidate-discovery tool for the retained physics-evaluation dataset, not as an extrapolative global model. A supplementary implementation-level representation of the retained polynomial and RBF-type surrogate families is provided in Appendix A.

To make the surrogate layer reproducible, the dataset-generation logic, split definition, model-family settings, and validation protocol are summarized in Table 3. The complete split files, target-specific hyperparameter metadata, and trained-model reproduction scripts are provided in the shared Google Drive reproducibility package. The final retained surrogate was not selected from nominal training accuracy alone; it was selected after random hold-out testing, blocked hold-out testing, optimizer-level backchecking, and post-optimization true-physics reevaluation.

The present study does not involve population-level statistical inference, experimental group comparison, or hypothesis testing. Therefore, statistical quantities such as p-values, confidence intervals, and effect sizes are not applicable to the retained analysis. The term “sampling” in this manuscript refers to deterministic and fixed-split computational sampling of the retained four-dimensional design box, not to biological, clinical, or survey sampling. Quantitative reliability is therefore assessed through surrogate-validation and numerical-consistency metrics, including RMSE, MAE, R^2, MAPE, blocked hold-out validation, post-optimization true-physics reevaluation, grid-convergence auditing, and deterministic heat-flux and volume-bias checks.

For the polynomial family, model fitting followed a ridge-regularized least-squares formulation of the form

$$\min_{\mathbf{w}_t} \sum_{i=1}^{N} \left(y_{t,i} - \hat{y}_{t,i}\right)^2 + \lambda_t \|\mathbf{w}_t\|_2^2, \tag{11}$$

where $\mathbf{w}_t$ is the coefficient vector for target t and λ_t is the corresponding regularization parameter. The same target-specific philosophy was retained for the RBF family through separate kernel and regularization choices for each response. In this way, the surrogate layer remained subordinate to the retained local physics problem rather than becoming an unconstrained global regression exercise.

The local surrogate package was not accepted on the basis of nominal regression accuracy alone. Instead, four complementary checks were applied. First, random train/validation/test hold-out was used to assess interpolation behavior within the retained local dataset. Second, blocked hold-out validation was used to remove entire local regions from training so that out-of-region robustness could be examined explicitly. Third, an optimizer-level backcheck was performed by comparing the surrogate-optimized solution against the best feasible design found directly in the local dataset. Finally, in the post-opt confirmation stage, the surrogate optimum, the dataset-best feasible point, and the retained nominal design were all reevaluated through the full physics-based chain. Under this workflow, surrogate fidelity was interpreted not only in terms of pointwise prediction error, but also in terms of preserving the optimizer-relevant ranking and feasibility structure of the retained local design problem.

Table 3
Surrogate dataset, model-family, and validation settings used for reproducibility of the retained local surrogate layer.

Item	Setting
Input vector	$\mathbf{x} = [\beta, W/L, r_{le}, n_{power}]^T$
Output targets	L/D, $q_{\max}$, V, C_L, C_D
Operating condition	Fixed $M_\infty = 10$, $h = 30$ km, $\alpha = 4°$
Dataset source	Frozen-chain physics evaluations inside the retained local design box
Sampling strategy	Retained local sweep and refinement samples; full sample list in repository
Number of samples	Reported in `physics_evaluations_local_box.csv`
Random split	Fixed split files provided in the repository
Blocked hold-out definition	Six local-region blocked audits; block labels provided in repository metadata
Polynomial family	Ridge-regularized polynomial surrogate
Polynomial features	Constant, linear, quadratic, and interaction terms of the four input variables
Ridge coefficient	Target-specific λ_t values reported in repository metadata
RBF family	Target-specific radial-basis-function surrogate
RBF kernel	Gaussian RBF kernel, $\phi(r) = \exp[-(r/\sigma_t)^2]$
RBF width parameter	Target-specific σ_t values reported in repository metadata
Final retained model	ARD squared-exponential Gaussian-process surrogate for all five targets
Validation metrics	RMSE, MAE, R^2, and MAPE on excluded physics-evaluation points
Post-opt confirmation	Retained surrogate-selected candidates reevaluated through frozen physics chain

For the random validation/test hold-out and blocked hold-out assessments, surrogate accuracy was quantified using the coefficient of determination R^2, the root-mean-square error RMSE, the mean absolute error MAE, and the mean absolute percentage error MAPE. These metrics were computed only on physics-evaluation points excluded from surrogate fitting. For a target response y_t, the retained definitions were

$$\mathrm{RMSE}_t = \sqrt{\frac{1}{N_h}\sum_{i=1}^{N_h}\left(\hat{y}_{t,i} - y_{t,i}^{phys}\right)^2}, \tag{12}$$

$$\mathrm{MAE}_t = \frac{1}{N_h}\sum_{i=1}^{N_h}\left|\hat{y}_{t,i} - y_{t,i}^{phys}\right|, \tag{13}$$

$$R_t^2 = 1 - \frac{\sum_{i=1}^{N_h}\left(y_{t,i}^{phys} - \hat{y}_{t,i}\right)^2}{\sum_{i=1}^{N_h}\left(y_{t,i}^{phys} - \bar{y}_t^{phys}\right)^2}, \tag{14}$$

$$\mathrm{MAPE}_t = \frac{100}{N_h}\sum_{i=1}^{N_h}\left|\frac{\hat{y}_{t,i} - y_{t,i}^{phys}}{y_{t,i}^{phys}}\right|. \tag{15}$$

where N_h is the number of hold-out physics-evaluation points, $\hat{y}_{t,i}$ is the surrogate prediction, and $y_{t,i}^{phys}$ is the corresponding frozen-chain physics output. These quantitative measures complement the parity plots and directly assess prediction accuracy on data not used for fitting.

To summarize the retained local methodology in a single reproducible workflow, Algorithm 1 presents the OCM-inspired physics-based evaluation, target-specific surrogate construction, surrogate-level candidate search, and mandatory post-optimization true-physics confirmation adopted in the present study, consistent with established waverider-design and reduced-order design-acceleration approaches [1,3,5,10].

Algorithm 1 Retained local surrogate-assisted waverider design workflow.

Require: Frozen operating condition $(M_\infty, h, \alpha, L_{\mathrm{ref}})$, trusted local design box $\mathcal{D}_{\mathrm{tr}}$ in $(\beta, W/L, r_{le}, n_{power})$, and feasibility limits $(q_{\max}^{\lim}, V_{\min}, C_{L,\min})$
Ensure: Retained peak feasible design $\mathbf{x}_{\mathrm{peak}}$ and recommended nominal design $\mathbf{x}_{\mathrm{nom}}$

Fix the baseline regime at $M_\infty = 10$, $h = 30$ km, $\alpha = 4°$, and $L_{\mathrm{ref}} = 10$ m
Define the trusted local design box $\mathcal{D}_{\mathrm{tr}}$ for $(\beta, W/L, r_{le}, n_{power})$
for all sampled designs $\mathbf{x} \in \mathcal{D}_{\mathrm{tr}}$ **do**
    Compute atmospheric properties and regime indicators
    Solve the conical-flow field using the Taylor–Maccoll model
    Generate the OCM-inspired parametric waverider geometry
    Evaluate aerodynamic outputs $(C_L, C_D, L/D)$
    Evaluate aerothermal outputs, including $q_{\max}$
    Evaluate geometric outputs, including enclosed volume V
    Mark $\mathbf{x}$ as feasible if $q_{\max} \le q_{\max}^{\lim}$, $V \ge V_{\min}$, and $C_L \ge C_{L,\min}$
end for
Interpret local variable roles from the physics-based refinement study
Build target-specific local surrogates for L/D, $q_{\max}$, V, C_L, and C_D
Validate the surrogate set using random hold-out and blocked hold-out tests
Search the surrogate-defined feasible region to identify high-value local candidates
Re-evaluate the surrogate optimum and retained reference candidates through the full physics-based chain
Retain $\mathbf{x}_{\mathrm{peak}}$ as the highest-performing feasible design under true-physics reevaluation
Retain $\mathbf{x}_{\mathrm{nom}}$ as the nearby feasible design with nearly unchanged aerodynamic quality and improved thermal margin
Embed the retained designs in the Mach–altitude flight corridor as a final consistency check

3. Verification and numerical credibility

The numerical credibility of the frozen physics chain was assessed through a dedicated verification and validation suite designed to determine which parts of the workflow can be treated as quantitatively reliable, which parts are only trend-faithful within a limited neighborhood, and where the scope of the subsequent claims must remain explicitly bounded. After re-auditing the oblique-shock verification block for angle-convention consistency, the former oblique-shock maximum-error outlier was removed from the governing-module verification summary. Over the retained 65 individual diagnostic comparisons, the updated aggregate mean absolute error was 3.678%, the aggregate root-mean-square error was 8.877%, the maximum absolute error was 40.611%, and the median absolute error was 0.004%. Accordingly, the frozen-chain closure status remains REVIEW rather than unconditional acceptance, because the Taylor–Maccoll and literature-based absolute aerodynamic comparisons still impose clear bounds on the predictive scope of the workflow. The results in the following sections are therefore interpreted as outputs of a partially verified engineering workflow, suitable for constrained local design assessment, but not as globally validated high-fidelity predictions.

Table 4
Compact governing-module verification summary for the retained physics chain.

Module	Tol.	Status	MAE	RMSE	Max
Taylor–Maccoll	2%	F	13.20	15.50	31.53
Oblique shock audit	1%	P	< 0.01	< 0.01	< 0.01
Prandtl–Meyer	0.5%	P	0.01	0.02	0.04
USSA 1976	1%	P	0.00	0.01	0.03
Stag. heating	15%	P	5.08	5.08	5.08
FR–SG	Diff. < 10%	P	4.84	4.84	4.84

Note: P = pass, F = fail, FR = Fay–Riddell, SG = Sutton–Graves, and USSA = US Standard Atmosphere. The oblique-shock entry corresponds to the corrected shock-angle audit in which β is treated as the prescribed shock-wave angle. The Taylor–Maccoll failure remains the principal governing-module limitation; therefore, the retained aerodynamic conclusions are restricted to local comparative ranking and internally consistent trend interpretation within the stated trust region.

3.1. Governing-module verification

The strongest agreement within the governing-module verification campaign was obtained for the Prandtl–Meyer expansion and atmospheric property modules. The Prandtl–Meyer block, assessed against standard analytical reference values with a 0.5% tolerance, passed all six cases and produced a group mean absolute error of 0.0105%, an RMSE of 0.0172%, and a maximum error of 0.0402%. Similarly, the atmosphere module, checked against the US Standard Atmosphere 1976 with a 1% tolerance, passed all nine sampled altitudes and yielded a group mean absolute error of 0.004%, an RMSE of 0.008%, and a maximum error of 0.025%. These results indicate that the expansion and atmosphere blocks provide a stable and accurate backbone for the retained local design study. A compact summary of the governing-module verification results is provided in Table 4.

The aeroheating branch also showed acceptable engineering consistency within the retained reduced-order scope. The stagnation-point heating block, evaluated with a 15% tolerance, passed all four test cases and produced a group MAE and RMSE of 5.08%. In addition, the Fay–Riddell versus Sutton–Graves cross-check passed all eight cases, with a consistent difference of 4.836%, which is fully compatible with the expected systematic coefficient offset reported in the validation log. This level of agreement supports the use of the thermal branch for comparative constraint enforcement and trend interpretation. It should not, however, be interpreted as validation of a fully coupled boundary-layer, finite-rate chemistry, catalytic-wall, or TPS material-response model. The retained thermal results are therefore used to rank local candidates and assess relative thermal margin, not to certify absolute surface heating or material-level thermal protection performance.

By contrast, the conical-flow verification remained the main unresolved governing-module limitation in the frozen implementation. The Taylor–Maccoll verification, assessed with a 2% tolerance, failed all seven benchmark cases and yielded a group MAE of 13.200%, an RMSE of 15.504%, and a maximum error of 31.527%. This result prevents the present chain from being interpreted as a fully verified absolute conical-flow predictor and requires the downstream aerodynamic results to be read primarily in terms of local comparative ranking and internally consistent trend behavior.

The originally reported oblique-shock discrepancy was re-audited separately because the retained implementation uses β as the prescribed shock-wave angle. In the oblique-shock module, the normal upstream Mach number is evaluated as $M_{1n} = M_1 \sin\beta$, and the pressure and temperature jumps are computed from the classical normal-Mach oblique-shock relations. The audit showed that the previously reported maximum error originated from an angle-convention inconsistency in the hard-coded benchmark table rather than from a numerical instability of the oblique-shock routine. For example, for $M_1 = 6$ and $\beta = 25°$, the shock-angle formulation gives $p_2/p_1 = 7.335$, whereas the original benchmark entry used $p_2/p_1 = 3.898$, which alone produces the previously reported maximum error of approximately 88%. After correcting the benchmark convention, the oblique-shock angle audit passed the adopted 1% tolerance, with MAE = 0.000008%, RMSE = 0.000010%, and Max = 0.000021%.

Diagnostic oblique-shock angle-convention audit using β as the prescribed shock-wave angle.

M_1	β [deg]	θ [deg]	p_2/p_1	T_2/T_1	M_2
2.0	45.0	14.744	2.167	1.264	1.456
3.0	40.0	21.846	4.172	1.630	1.894
4.0	35.0	22.180	5.974	1.942	2.412
5.0	30.0	20.174	7.125	2.138	3.006
6.0	25.0	17.134	7.335	2.173	3.724

The corrected oblique-shock audit therefore removes the former benchmark-error outlier from the governing-module verification summary. Nevertheless, the use of β-dependent trends remains bounded by the overall reduced-order nature of the frozen chain and by the unresolved Taylor–Maccoll limitation. Consequently, shock-angle-dependent conclusions are interpreted as local comparative trends within the retained trust region, rather than as broad high-fidelity shock-controlled aerodynamic predictions. In the retained workflow, the oblique-shock module is not used to claim quantitatively validated tabulated shock-state accuracy. Instead, it provides a consistent engineering closure within the same reduced-order chain used for all candidate designs. Consequently, trends involving the prescribed shock-angle parameter β are interpreted as local, internally consistent design-space trends within the retained trust region, rather than as absolute shock-state predictions. This distinction is important because the present manuscript uses the frozen chain for comparative design ranking, feasibility discrimination, and surrogate backchecking, not as a globally validated high-fidelity shock solver.

The Taylor–Maccoll result should therefore not be interpreted as a fully verified absolute conical-flow prediction within the present frozen implementation. Its role in the retained workflow is more limited: it provides a consistent conical-flow-informed basis for constructing and comparing members of the same OCM-inspired parametric quasi-waverider family inside a narrowly bounded local design box. Consequently, the subsequent optimization results are interpreted in terms of local ranking, feasibility discrimination, and internally consistent design trends rather than as absolute high-fidelity aerodynamic predictions. This limitation is explicitly carried through the discussion and conclusion, where the retained framework is described as a partially verified engineering workflow requiring higher-fidelity confirmation before broad predictive use.

Taken together, the governing-module verification results establish a mixed but still usable credibility structure for the frozen chain. The atmosphere, Prandtl–Meyer, corrected oblique-shock angle audit, and heating modules are sufficiently reliable for the local constrained study reported here, whereas the Taylor–Maccoll discrepancy limits the extent to which the workflow can be presented as an absolute conical-flow predictive solver. For this reason, the remainder of the manuscript interprets the retained designs primarily through local comparisons, feasibility ranking, and cross-consistency between aerodynamic, thermal, and geometric trends. The group-level verification metrics reported above and in Table 4 were evaluated using Eq. (16)-(18).

$$\varepsilon_i = \left| \frac{\phi_i^{\mathrm{calc}} - \phi_i^{\mathrm{ref}}}{\phi_i^{\mathrm{ref}}} \right| \times 100 \tag{16}$$

$$\mathrm{MAE} = \frac{1}{N} \sum_{i=1}^{N} \varepsilon_i \tag{17}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} \varepsilon_i^2} \tag{18}$$

Table 5
Compact literature-based L/D validation summary using only explicitly cited reference studies.

M	L/D_{ref}	L/D_{calc}	Err. [%]	Stat.
6	6.50	3.86	40.6	F
10	5.00	6.88	37.6	F

Note: P = pass and F = fail under the adopted 15% tolerance. These literature comparisons are not one-to-one geometry-matched validation cases; the failed entries are retained deliberately to show that the frozen reduced-order chain is not used as an absolute L/D predictor over generic waverider configurations.

Here, ϕ_i^{calc} and ϕ_i^{ref} denote the calculated and reference values of the selected verification quantity, respectively, and N is the number of comparisons in the corresponding verification group.

3.2. Literature-based absolute L/D comparison

Beyond the governing-equation checks, the retained physics chain was also compared against representative literature-based aerodynamic-efficiency levels. This comparison is deliberately treated as an absolute L/D range check rather than as a strict one-to-one validation exercise, because the cited configurations are not geometrically identical to the present OCM-inspired parametric quasi-waverider and do not share the same viscous optimization level, closure assumptions, operating details, or drag-model implementation. In the present manuscript, the comparison is restricted to explicitly cited reference studies, namely the viscous-optimized Mach-6-class result reported by Bowcutt et al. and the Mach-10-class axisymmetric-flowfield result reported by Corda and Anderson [1,17]. The purpose of this comparison is therefore to expose the level of absolute aerodynamic-efficiency bias in the retained reduced-order chain, not to claim geometric equivalence with the published vehicles. The retained cited benchmark comparison is summarized in Table 5.

The resulting agreement is limited. For the Mach 6 comparison, the validation routine used a reference value of $L/D = 6.50$ and produced a calculated value of $L/D = 3.86$, corresponding to a relative error of 40.6% and a failed case under the adopted 15% tolerance. For the Mach 10 comparison, the reference value was $L/D = 5.00$ and the calculated value was $L/D = 6.88$, corresponding to a relative error of 37.6%, which also failed the tolerance criterion. Taken together, these retained cited comparisons indicate that the present workflow does not reproduce the absolute literature values with sufficient accuracy to support broad predictive claims. For this reason, Table 5 is interpreted as a negative absolute-performance benchmark rather than as a passed validation result. Instead, the retained chain is more appropriately interpreted as a reduced-order engineering workflow that preserves useful local ranking behavior while exhibiting noticeable bias in absolute aerodynamic-efficiency prediction.

This result does not negate the utility of the retained chain for the purposes of the present study. Rather, it constrains the admissible interpretation of the subsequent results. The retained local optimization and surrogate-assisted refinement stages are therefore discussed comparatively, with emphasis on internal consistency, local feasibility ranking, and constrained trend analysis rather than on exact agreement with legacy absolute L/D benchmarks. This framing is consistent with the broader credibility bounds established for the retained physics chain.

3.3. Numerical consistency and cross-checks

In addition to the governing-equation verification and literature-based comparison, the retained physics chain was subjected to internal numerical-consistency checks in order to assess discretization sensitivity and cross-model agreement within the local workflow. This part of the assessment addressed two questions: first, whether the principal aerodynamic and geometric outputs exhibited bounded mesh-level sensitivity under systematic refinement; and second, whether the retained thermal submodels remained mutually consistent over representative hypersonic conditions. A compact summary of the numerical-consistency results is provided in Table 7.

Table 6
Raw grid-level values used in the extended numerical-consistency audit.

Grid level	L/D	C_D	V [m^3]
15×25	6.051947	0.011091	3.220130
30×50	6.008061	0.011190	3.251668
60×100	5.951879	0.011310	3.273308
120×200	5.875884	0.011465	3.287901

Note: The raw values correspond to the frozen validation-chain grid audit. The extended GCI values in Table 7 were computed from the 30×50, 60×100, and 120×200 levels, whereas the 15×25 level is retained to document the original three-grid audit.

Table 7
Extended numerical-consistency summary for the frozen physics chain using the 30×50–60×100–120×200 grid sequence.

Qty.	p	$\mathrm{GCI}_{\mathrm{fine}}$ [%]	$\Delta_{m\to f}$ [%]	Status	f_{ext}
L/D	0.436	4.584	1.293	P	5.660386
C_D	0.373	5.748	1.357	F	0.011993
V	-0.568	0.444	0.444	P	3.287901

Note: p is the observed order of convergence, $\mathrm{GCI}_{\mathrm{fine}}$ is the fine-grid Grid Convergence Index, $\Delta_{m\to f}$ is the relative change from the medium to the fine grid, f_{ext} is the Richardson-extrapolated estimate, P = pass, and F = fail under the adopted criterion $\mathrm{GCI}_{\mathrm{fine}} <$ 5%. The original 15×25–30×50–60×100 sequence gave C_D $\mathrm{GCI}_{\mathrm{fine}} = 6.074\%$; the added 120×200 audit reduced this value to 5.748%, but not below the 5% threshold. Therefore, the failed C_D entry is retained deliberately. The low observed orders show that the sequence should be interpreted as a conservative numerical-consistency audit rather than as a formal asymptotic grid-convergence demonstration.

A grid-convergence audit was first performed using structured resolutions of 15×25, 30×50, and 60×100. In response to the drag-coefficient credibility concern, this audit was extended with an additional 120×200 fine-grid calculation. The original three-grid sequence produced C_D $\mathrm{GCI}_{\mathrm{fine}} = 6.074\%$, exceeding the adopted 5% numerical-credibility threshold. The extended sequence, using 30×50, 60×100, and 120×200 as the three GCI levels, reduced the corresponding C_D $\mathrm{GCI}_{\mathrm{fine}}$ to 5.748%, but did not bring it below the acceptance threshold. Therefore, the failed C_D grid-convergence status is retained deliberately in the revised manuscript.

$$p = \frac{\ln\left|\dfrac{f_3 - f_2}{f_2 - f_1}\right|}{\ln r}, \tag{19}$$

where f_1, f_2, and f_3 denote the coarse-, medium-, and fine-grid solutions, respectively, and r is the refinement ratio. The fine-grid Grid Convergence Index was then estimated as

$$\mathrm{GCI}_{\mathrm{fine}} = F_s \frac{|(f_3 - f_2)/f_3|}{r^p - 1} \times 100, \tag{20}$$

with a three-grid safety factor of $F_s = 1.25$. For completeness, the Richardson-extrapolated grid-independent estimate was written as

$$f_{\mathrm{ext}} = f_3 + \frac{f_3 - f_2}{r^p - 1}. \tag{21}$$

These relations were those used by the frozen validation routine for the grid-sensitivity block.

In addition to the summarized GCI metrics, the raw grid-level values used in the extended convergence audit are provided in Table 6.

These values allow the reported observed orders, fine-grid GCI values, and Richardson-extrapolated estimates to be independently checked. The complete raw grid table is also included in the shared reproduction package.

The extended-grid behavior was mixed but informative. For L/D, the extended 30×50–60×100–120×200 sequence gave $p = .436$, $\mathrm{GCI}_{\mathrm{fine}} = 4.584\%$, a medium-to-fine relative change of 1.293%, and $f_{\mathrm{ext}} = 5.660386$, so the aerodynamic-efficiency metric remained within the adopted 5% threshold. For total drag coefficient, the same extended sequence gave $p = .373$, $\mathrm{GCI}_{\mathrm{fine}} = 5.748\%$, a medium-to-fine relative change of 1.357%, and $f_{\mathrm{ext}} = 0.011993$. This remains a failed case because the GCI is still above the adopted 5% threshold. For enclosed volume, the extended sequence gave $p = -0.568$, $\mathrm{GCI}_{\mathrm{fine}} = 0.444\%$, a relative change of 0.444%, and $f_{\mathrm{ext}} = 3.287901$, which was accepted. The extended asymptotic-convergence ratios for L/D and C_D were 0.540 and 0.604, respectively, indicating that the refinement sequence is informative but not asymptotic.

The low observed orders of convergence indicate that the refinement sequence should not be interpreted as a formal asymptotic grid-convergence demonstration. This behavior is expected to some extent in the present reduced-order workflow because the computed aerodynamic outputs combine geometry reconstruction, panel-based pressure integration, skin-friction estimation, base-drag closure, and leading-edge blunting effects rather than a single smooth high-order discretization of the governing equations. Accordingly, the grid study is used here as a conservative numerical-consistency audit rather than as evidence of asymptotic discretization-error closure. The resulting credibility claim is therefore limited: L/D and volume show acceptable fine-grid consistency for local ranking and geometry-level interpretation, whereas C_D and drag-share quantities remain supporting diagnostics only.

These outcomes support a restrained interpretation of numerical consistency. The two quantities most important for the retained local design narrative, namely aerodynamic efficiency and volume preservation, remain within the adopted grid-acceptance threshold even after adding the 120×200 audit level. By contrast, the drag coefficient still does not satisfy the adopted numerical-credibility threshold: its extended fine-grid GCI is 5.748%, exceeding the 5% acceptance limit. Therefore, the present manuscript does not treat absolute C_D values as numerically closed predictions. Instead, C_D is used as a supporting diagnostic for interpreting local trends, drag-share structure, and surrogate consistency within the retained design box. Absolute drag-related quantities should therefore be read more cautiously than L/D, volume, thermal feasibility, and relative ranking outcomes.

A second internal consistency check was performed on the thermal side by directly comparing the Fay–Riddell and Sutton–Graves stagnation-heating correlations over eight representative test conditions spanning multiple Mach numbers, altitudes, and nose radii. All eight cases passed the adopted difference criterion, and the reported discrepancy remained uniformly 4.836%, matching the expected systematic coefficient offset noted in the validation output. This result is important because it shows that the thermal branch does not merely pass isolated point checks, but also preserves internal cross-correlation consistency across the operating envelope used in the retained study. This cross-check is an internal correlation-consistency test, not a substitute for coupled boundary-layer/aerothermochemical validation. Effects such as catalytic wall recombination, transition-sensitive boundary-layer growth, surface chemistry, ablation, and material response remain outside the retained model fidelity.

Taken together, the numerical-consistency results support the use of the retained chain as a bounded local engineering tool rather than as a fully converged universal predictor. Mesh refinement suggests acceptable fine-grid consistency for L/D and volume, but not asymptotic convergence or complete closure for C_D, while the thermal cross-check confirms strong internal consistency of the stagnation-heating branch. The combined implication is that the retained local design conclusions are numerically defensible when interpreted in terms of local feasibility, surrogate backchecking, and relative design ranking, whereas absolute drag-related claims should remain explicitly qualified. The observed convergence orders are therefore used as diagnostic evidence of the present mesh-sensitivity limit, not as proof of formal asymptotic mesh independence.

With these verification bounds and numerical qualifications established, the retained local design behavior and surrogate-assisted refinement results can now be interpreted on a controlled and physically transparent basis.

4. Results

This section presents the retained local design behavior, the role of the active geometric variables within the trusted Mach 10, 30 km, $\alpha = 4°$ regime, and the surrogate-assisted refinement outcomes. In accordance with the verification bounds established in Section 3, the results are interpreted comparatively and locally, with emphasis on physically consistent trend extraction, constrained feasible ranking, and optimizer-relevant structure rather than on universal global prediction.

4.1. Local design-space behavior

The retained local design study showed that the dominant aerodynamic variation within the trusted neighborhood is governed primarily by the shock angle β and the width-to-length ratio W/L. When the successful local samples are projected onto the $(\beta, W/L)$ plane and colored by aerodynamic efficiency, the highest-performing candidates accumulate toward the upper-right portion of the retained box. This pattern indicates that aerodynamic efficiency increases as both β and W/L increase within the examined local range, and that the local optimum is better interpreted as a smooth boundary-seeking trend than as a clearly isolated interior stationary point.

The contour-based local maps support the same interpretation, while also separating the roles of the principal aerodynamic responses. Within the retained $\alpha = 4°$ slice, C_L increases mainly with β, whereas total C_D increases with both β and W/L. At the same time, the enclosed volume rises monotonically across the retained local box, so the upper-right region is favored not only by aerodynamic efficiency but also by geometric capacity. Taken together, these trends explain why the best feasible solution is pushed toward the upper boundary of the retained local domain rather than toward the interior.

Over the retained feasible slice shown in Fig. 1–3, a compact local summary of the observed aerodynamic trends is

$$\frac{\partial(L/D)}{\partial\beta} > 0, \qquad \frac{\partial(L/D)}{\partial(W/L)} > 0, \qquad \frac{\partial C_L}{\partial\beta} > 0 \tag{22}$$

within the trusted local design box. Eq. (22) is not intended as a universal global law for hypersonic waverider design; rather, it summarizes the observed local behavior of the retained design space at the fixed Mach 10, 30 km, $\alpha = 4°$ condition.

A second important observation is that, at fixed leading-edge radius, the heating field in the $(\beta, W/L)$ plane remains comparatively flat relative to the aerodynamic fields. Thus, within the retained local neighborhood, the in-plane variation is controlled mainly by aerodynamic structure, while the thermal admissibility boundary is influenced more strongly by the leading-edge radius than by strong in-plane changes of β and W/L. This separation of roles is important because it makes the local problem physically interpretable: aerodynamic gain is obtained mainly through the $(\beta, W/L)$ subspace, whereas thermal accessibility is governed primarily by leading-edge bluntness (see Fig. 4).

Accordingly, the retained local design-space behavior may be interpreted as a two-variable aerodynamic escalation embedded within a wider four-variable feasible region. In this formulation, β and W/L act mainly as aerodynamic performance drivers, while the remaining

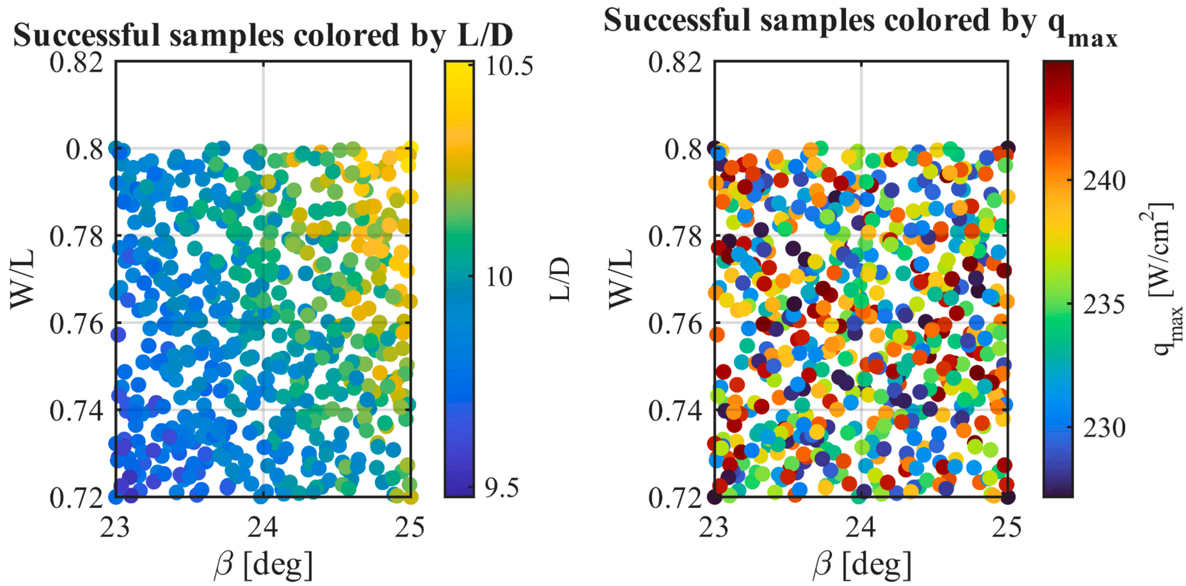

Fig. 1. Successful local design samples projected onto the retained $(\beta, W/L)$ plane: (a) colored by aerodynamic efficiency L/D, and (b) colored by maximum surface heating $q_{\max}$. Together, the two panels illustrate the local aerodynamic–thermal structure of the retained design region.

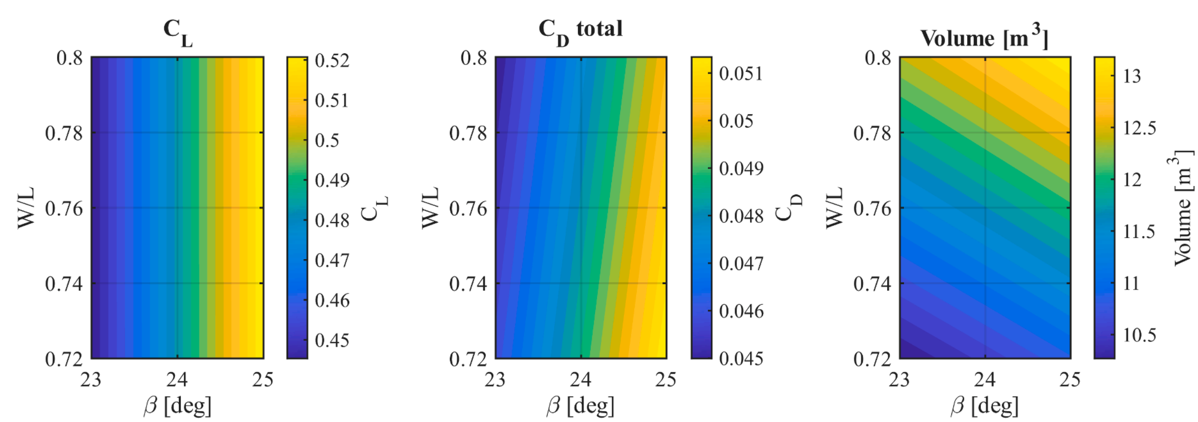

Fig. 2. Retained local response maps on the $(\beta, W/L)$ plane at the representative Mach 10, 30 km, $\alpha = 4°$ condition: (a) lift coefficient C_L, (b) total drag coefficient C_D, and (c) enclosed volume V. Together, the three panels summarize the local aerodynamic and geometric trends within the retained design region.

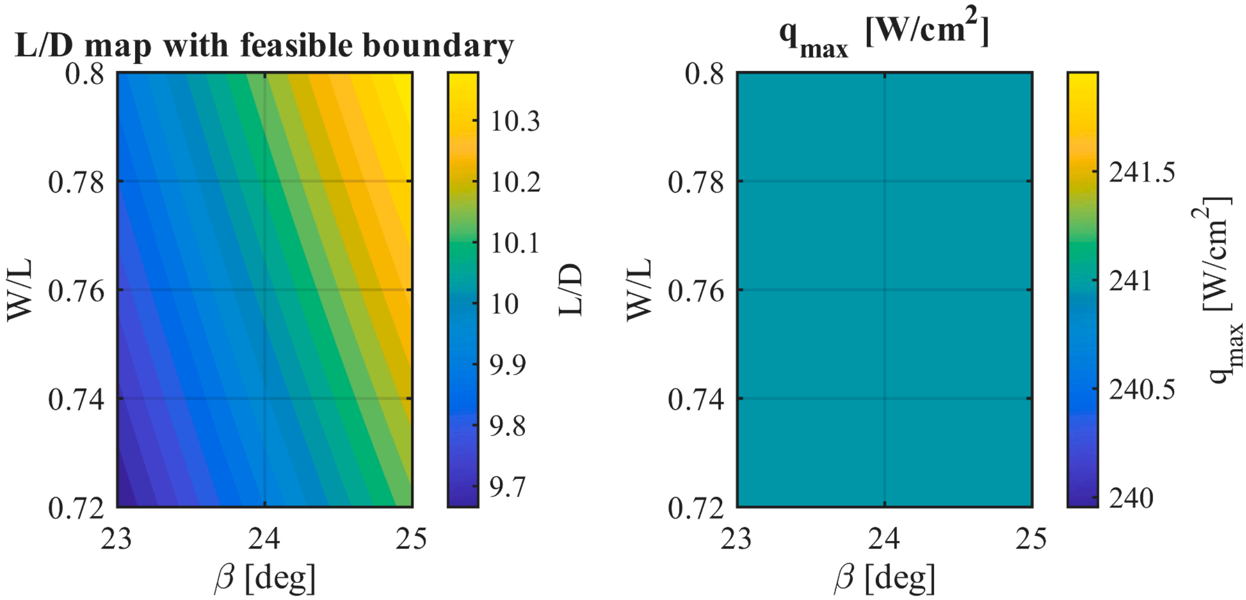

Fig. 3. Local aerodynamic-efficiency and heating maps in the retained $(\beta, W/L)$ plane.

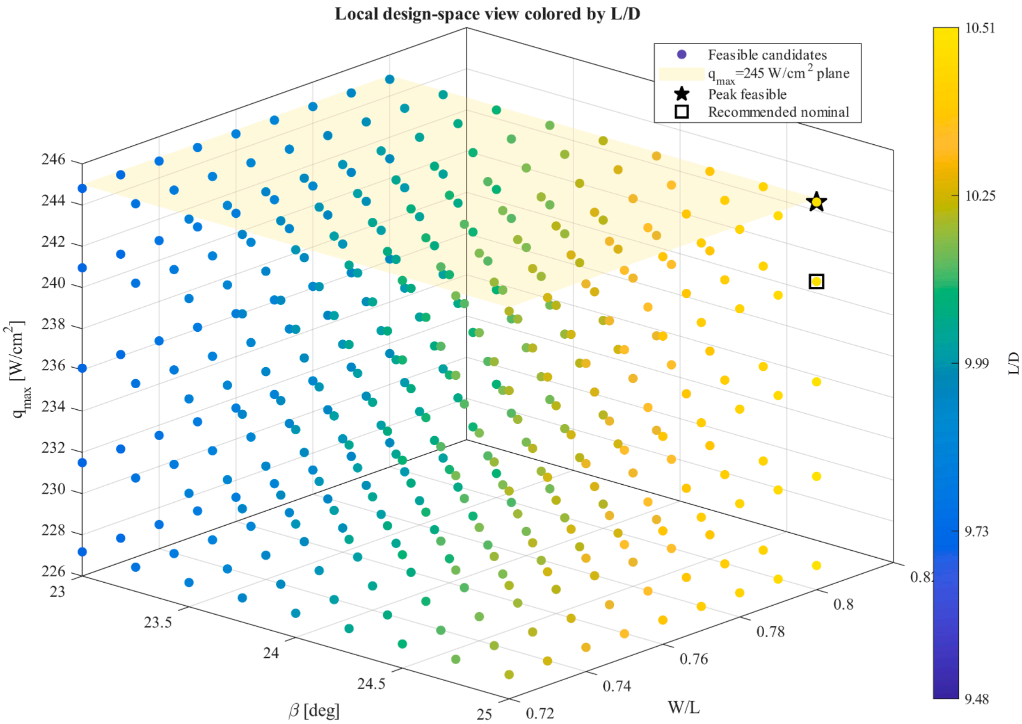

Fig. 4. Three-dimensional local design-space visualization in terms of shock angle β, width-to-length ratio W/L, and maximum surface heating $q_{\max}$, with marker color indicating aerodynamic efficiency L/D. The upper plane corresponds to the retained thermal feasibility limit $q_{\max} = 245$ W/cm^2. The figure shows that the retained local design family approaches the thermal boundary near the high-β, high-W/L aerodynamic-performance corner, while the recommended nominal design steps back from the heat-flux limit with nearly unchanged L/D.

variables determine whether the corresponding high-performance region remains admissible under the frozen thermal and geometric constraints. This interpretation provides the physical basis for the leading-edge-radius analysis in Section 4.2 and for the surrogate-guided refinement discussed in the subsequent subsections.

4.2. Thermal role of leading-edge radius

Across the retained local range, the dependence of maximum heating on leading-edge radius is strongly monotonic and decreasing, while the corresponding aerodynamic response remains weak in relative terms. In practical terms, increasing bluntness provides systematic thermal relief, but does not generate a commensurate first-order aerodynamic benefit over the retained radius range. Accordingly, Fig. 5 reports the collapsed representative heating response, whereas Fig. 6 reports the relative change in aerodynamic efficiency for representative local angle-of-attack slices, with each slice referenced to the smallest retained leading-edge radius in that slice. The local behavior may therefore be summarized as

$$\frac{\partial q_{\max}}{\partial r_{le}} < 0, \qquad \frac{\partial \Delta(L/D)_{\mathrm{rel}}}{\partial r_{le}} < 0 \tag{23}$$

within the retained Mach 10, 30 km local neighborhood, where the heating response is nearly slice-invariant across the representative $\alpha = 1°$, 4°, and 6° cases. Eq. (23) should therefore be interpreted as a local relative-trend statement rather than as a claim about universal absolute lift-to-drag levels for hypersonic waveriders.

The detailed constrained-sweep slices are provided in Appendix B. These supplementary maps show that increasing the leading-edge radius shifts the retained $(\beta, W/L)$ region toward lower maximum-heating levels while preserving the same overall in-plane topology. For the representative radii, the reported $q_{\max}$ bands are approximately 218.4–220.2 W/cm^2 at $r_{le} = 29$ mm, 199.0–200.8 W/cm^2 at $r_{le} = 36$ mm, and 173.8–175.6 W/cm^2 at $r_{le} = 50$ mm. Thus, the dominant effect of increasing bluntness is not to reshape the aerodynamic topology of the local box, but to move the same retained local region farther away from the thermal limit.

This interpretation can be stated more formally by defining the retained feasible in-plane set at a fixed leading-edge radius as

$$\mathcal{F}(r_{le}) = \left\{ (\beta, W/L) \;\middle|\; \begin{array}{c} q_{\max}(\beta, W/L, r_{le}) \le 245 \text{ W/cm}^2, \\ V(\beta, W/L, r_{le}) \ge 3.0 \text{ m}^3, \\ C_L(\beta, W/L, r_{le}) \ge 0.01 \end{array} \right\}. \tag{24}$$

Under the retained formulation, increasing r_{le} expands $\mathcal{F}(r_{le})$ primarily through thermal relief rather than through a major change in aerodynamic efficiency. In other words, r_{le} determines whether the locally favorable aerodynamic region is thermally accessible, whereas β and W/L determine how favorable that region is once it becomes admissible.

This same trade-off appears in the retained nominal-versus-peak comparison. In the corridor-based embedding developed later in the paper, the recommended nominal point is associated with $r_{le} = 24.00$ mm and $q_{\max} = 240.96$ W/cm^2, while the peak feasible point occurs at $r_{le} = 23.25$ mm with $q_{\max} = 244.81$ W/cm^2. The aerodynamic distinction between these two retained points is intentionally small, but the thermal margin differs in a meaningful way. This confirms that the recommended nominal design is not a fundamentally different aerodynamic solution; rather, it is a slightly blunter member of the same retained local family, selected to preserve nearly the same performance while stepping back from the thermal boundary. Because the aerothermal model is correlation-based and does not include catalytic-wall or coupled boundary-layer effects, this margin is interpreted as a relative engineering screening margin within the retained model rather than as an absolute TPS certification margin.

Accordingly, the role of leading-edge radius in the present study is best interpreted as that of a thermal-access variable. The retained peak feasible design is obtained near the smallest thermally admissible radius, whereas the recommended nominal design is obtained by moving

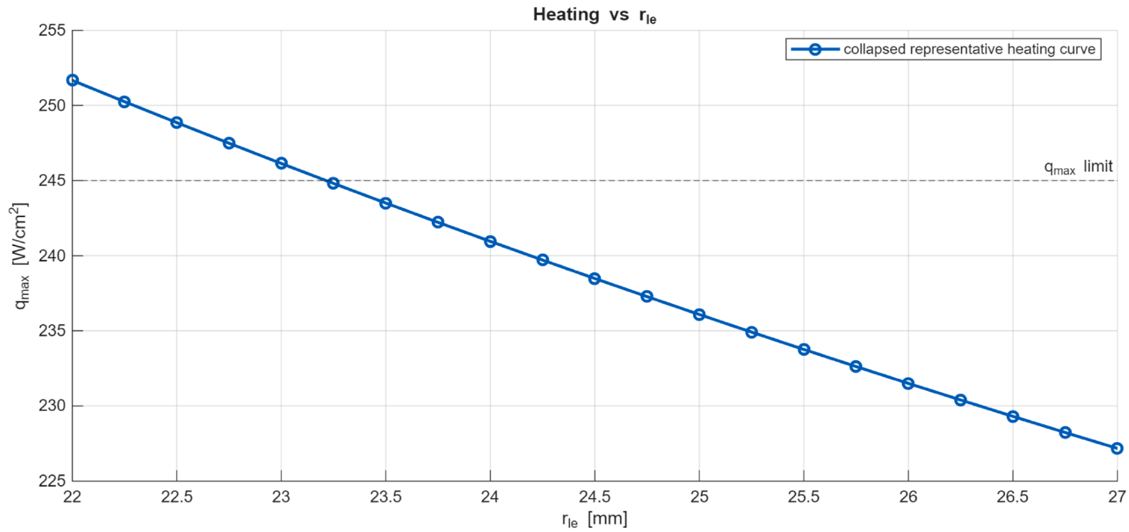

Fig. 5. Variation of maximum surface heating with leading-edge radius in the retained local refinement. The representative $\alpha = 1°$, $4°$, and $6°$ slices collapse onto a nearly identical heating response, indicating that the dominant sensitivity in this sweep is to leading-edge radius rather than angle of attack.

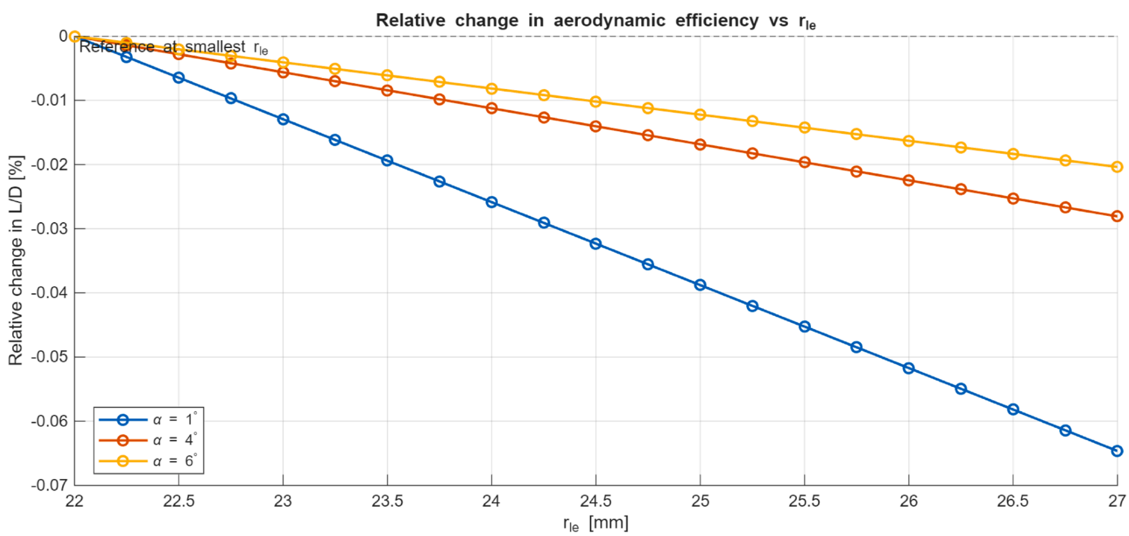

Fig. 6. Relative change in aerodynamic efficiency with leading-edge radius for representative local angle-of-attack slices. Each slice is referenced to its value at the smallest retained leading-edge radius in that slice. The figure is intended to emphasize relative local sensitivity to bluntness rather than absolute lift-to-drag levels.

to a slightly larger bluntness level in order to gain thermal margin at minimal aerodynamic cost. This interpretation is central to the logic of the retained local design study, because it explains why the final design decision is not governed solely by peak L/D, but by the balance between local performance and proximity to the thermal feasibility boundary.

4.3. Surrogate-guided refined candidates

After the local variable roles had been identified through the physics-based refinement study, the retained four-dimensional neighborhood was explored using target-specific local surrogates in order to expose the optimizer-relevant structure of the admissible design space more efficiently. The purpose of this stage was not to replace the retained physics chain with a universal predictive model, but to use the surrogate layer as a local accelerator for identifying a compact set of refined candidates worthy of subsequent full-physics reevaluation.

Within the retained local box, the surrogate-guided search reproduced the same qualitative design logic already visible in the physics-based maps. The highest-performing predicted candidates accumulated near the upper portion of the retained β and W/L bounds, while remaining close to the smallest thermally admissible values of r_{le}. Thus, the surrogate did not introduce a qualitatively different optimum structure; instead, it recovered the same boundary-seeking local pattern in a smoother and computationally more efficient form. This behavior is illustrated directly by the seed-search distribution and retained candidate clustering shown in Figs. 7 and 8. The retained surrogate optimum may therefore be written in compact form as

$$\mathbf{x}^*_{\text{surr}} = \arg\max_{\mathbf{x}\in\hat{\mathcal{F}}} \left(\widehat{L/D}\right)(\mathbf{x}), \qquad \mathbf{x} = \left[\beta,\ W/L,\ r_{le},\ n_{\text{power}}\right]^T, \tag{25}$$

where $\hat{\mathcal{F}}$ denotes the surrogate-defined feasible set under the retained local constraints. In this formulation, the surrogate search is not used to redefine feasibility, but to identify the most promising subset of locally admissible candidates for later physics-based confirmation.

A second important result is that the refined surrogate candidates occupy a compact and physically interpretable region rather than a scattered cloud. This behavior is significant because it indicates that the retained local problem is not dominated by surrogate-induced noise or spurious optimizer artifacts. Instead, the local design space contains a coherent high-performance neighborhood, and the surrogate-guided refinement acts mainly to resolve its internal ordering and proximity to the thermal boundary. In engineering terms, this is precisely the situation

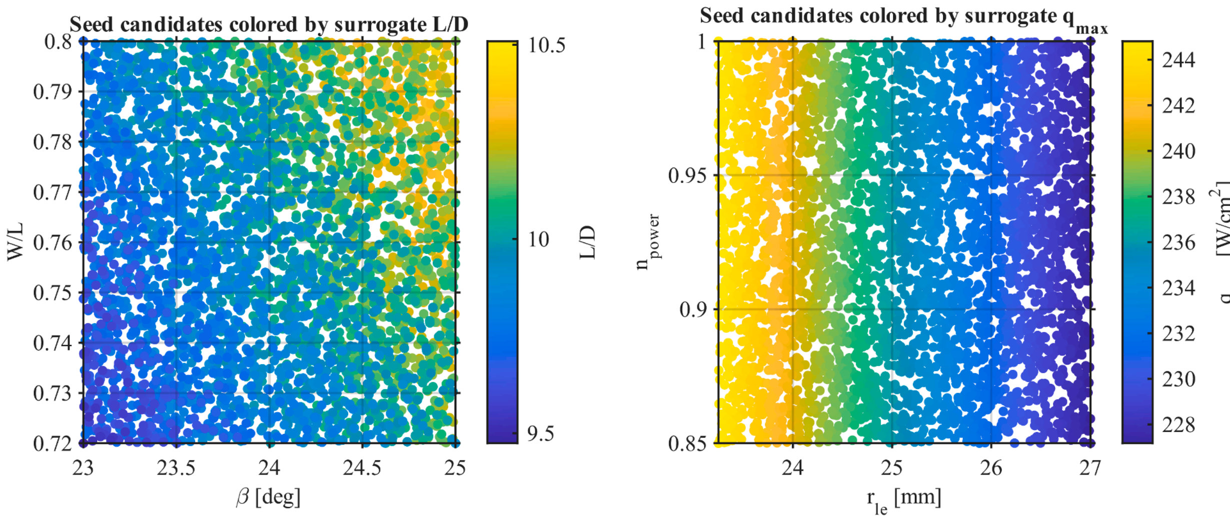

Fig. 7. Surrogate-guided seed search within the retained four-dimensional design box.

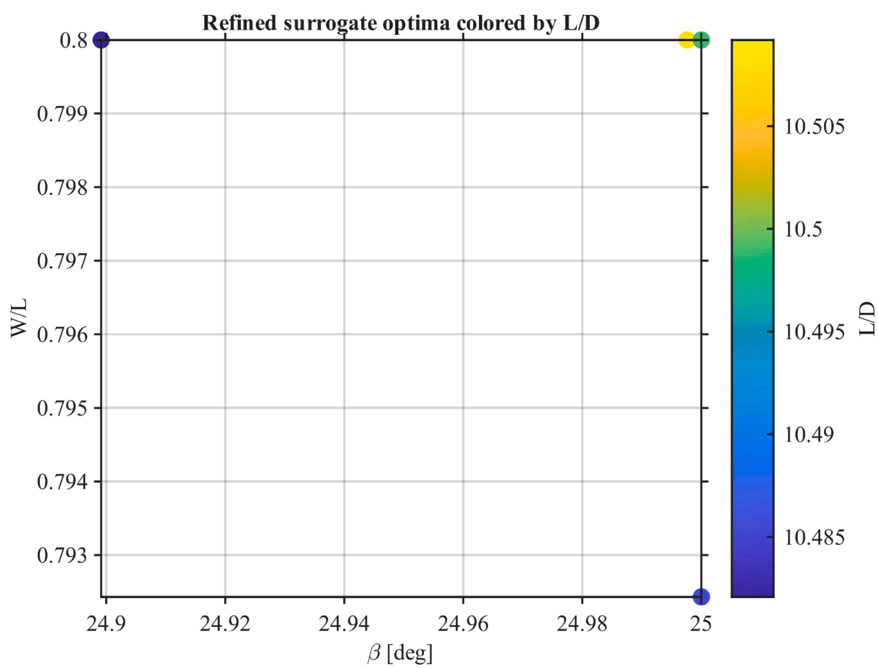

Fig. 8. Refined surrogate candidate set in the retained local design space.

in which a local surrogate is most useful: it does not replace physical interpretation, but sharpens it.

The parity plots shown in Figs. 9 and 10, together with the quantitative random test and worst blocked hold-out metrics in Table 8, reinforce this conclusion. Across the retained local dataset, the predicted and physics-based values of C_D, C_L, L/D, $q_{\max}$, and V remain closely concentrated around the one-to-one line. The random test and blocked hold-out metrics further show that the surrogate errors remain small on physics-evaluation points excluded from fitting, including blocked cases where local regions of the design box were removed from the training set. Thus, the surrogate assessment is not based only on visual parity, but also on explicit R^2, RMSE, MAE, and MAPE measures. To express this consistency compactly, the local surrogate deviation for a target quantity y_t may be written as

$$\Delta y_t(\mathbf{x}) = \hat{y}_t(\mathbf{x}) - y_t^{\mathrm{phys}}(\mathbf{x}), \tag{26}$$

where $\hat{y}_t$ is the surrogate prediction and y_t^{phys} is the corresponding output of the frozen physics chain. In the present study, the key requirement is not exact pointwise identity at a global level, but sufficiently small local deviation such that candidate ranking, feasibility discrimination, and boundary-seeking behavior are preserved inside the trusted neighborhood. Under that criterion, the retained parity behavior supports the use of the surrogate layer as a reliable local decision-support tool.

This point is especially important for the thermal and aerodynamic responses taken together. If the surrogate distorted $q_{\max}$ more strongly than L/D, or vice versa, the retained optimum structure could shift artificially. The observed parity behavior suggests that such distortion is not dominant within the retained local regime. Therefore, the refined surrogate candidates may be interpreted as physically meaningful approximations of the same local optimum family identified by the direct refinement campaign, rather than as artifacts created by regression bias.

Taken together, these results show that the retained surrogate stage succeeds in its intended role: it compresses the local design search into a smaller and more interpretable set of high-value candidates while preserving the same boundary-controlled local optimum structure already indicated by the direct physics-based analysis. This provides the necessary bridge between the local response interpretation developed in

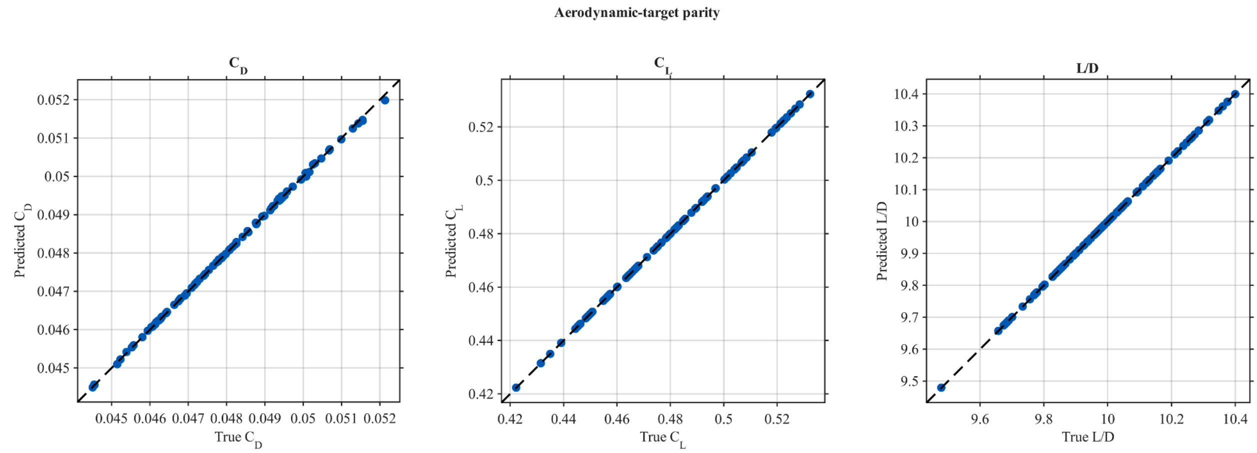

Fig. 9. Aerodynamic-target parity between surrogate predictions and frozen-chain physics outputs for the retained local dataset: C_D, C_L, and L/D. Quantitative random test and worst blocked hold-out metrics are reported in Table 8.

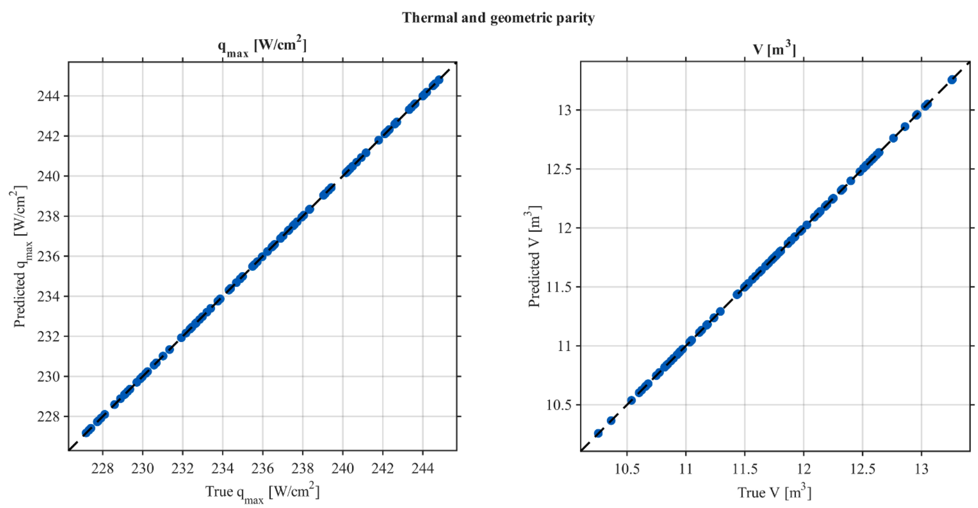

Fig. 10. Thermal and geometric parity between surrogate predictions and frozen-chain physics outputs for the retained local dataset: $q_{\max}$ and V. Quantitative random test and worst blocked hold-out metrics are reported in Table 8.

Sections 4.1 and 4.2 and the true-physics retained-design confirmation presented in the following subsection.

Although the present study does not claim a high-fidelity CFD-level cost reduction, the retained surrogate layer provides a practical computational advantage within the reduced-order screening workflow. A full frozen-chain evaluation requires sequential atmosphere, conical-flow, geometry-generation, aerodynamic, aerothermal, and geometric-volume calculations for each candidate, whereas the trained surrogate evaluates the same target responses through direct target-specific regression calls. Therefore, the surrogate is used as a low-cost local candidate-screening layer, while only the retained high-value candidates are passed back to the full physics chain for confirmation. This workflow reduces the number of full-chain evaluations required during local refinement and preserves physical credibility because the final retained candidates are still judged only after post-optimization true-physics reevaluation.

4.4. True-physics confirmation of retained designs

The surrogate-guided refinement stage produced a compact set of retained candidates, but the final design interpretation must ultimately be based on reevaluation through the frozen physics chain rather than on surrogate prediction alone. For this reason, the dataset-best feasible point, the recommended nominal design, and a representative dataset-center point were re-run through the full analysis chain. The objective of this step was not to search for a new optimum, but to verify whether the retained candidates preserve their local ordering, feasibility status, and engineering interpretation once the surrogate layer is removed.

The post-optimization reevaluation confirmed that the retained candidates belong to the same local optimum family. In particular, the surrogate optimum and the dataset-best feasible point remain nearly indistinguishable at the level of optimizer-relevant performance, while the recommended nominal design preserves essentially the same aerodynamic quality with slightly increased thermal margin. By contrast, the dataset-center point remains feasible but clearly less attractive in aerodynamic terms, confirming that the retained high-performance region is not a broad flat plateau across the full local box, but a compact neighborhood concentrated near the boundary-seeking optimum structure identified in the preceding subsections.

A compact way to express the post-opt confirmation requirement is

$$\Delta J_k(\mathbf{x}) = \hat{J}_k(\mathbf{x}) - J_k^{\text{phys}}(\mathbf{x}),$$
$$\varepsilon_{J_k}(\mathbf{x}) = 100\left|\frac{\hat{J}_k(\mathbf{x}) - J_k^{\text{phys}}(\mathbf{x})}{J_k^{\text{phys}}(\mathbf{x})}\right|, \quad J_k \in \{L/D,\ q_{\max}\}. \tag{27}$$

Here, $\hat{J}_k$ denotes the surrogate prediction and J_k^{phys} denotes the corresponding frozen-chain true-physics reevaluation. The purpose of the post-opt check is not to enforce exact pointwise identity, but to verify

Table 8
Quantitative surrogate accuracy on random test and worst blocked hold-out physics-evaluation points in the retained local design box.

Target	Split	Block	RMSE	MAE	R^2	MAPE [%]
L/D	Test	–	1.74×10^{-4}	1.26×10^{-4}	0.99999917	0.00127
L/D	Worst blocked	n_{power} high	1.36×10^{-3}	9.86×10^{-4}	0.99994363	0.00990
q_{max}	Test	–	1.40×10^{-3}	1.18×10^{-3}	0.99999993	0.00050
q_{max}	Worst blocked	r_{le} high	1.64×10^{-3}	1.38×10^{-3}	0.99999752	0.00061
V	Test	–	5.94×10^{-4}	4.47×10^{-4}	0.99999934	0.00382
V	Worst blocked	volume corner	1.53×10^{-3}	1.38×10^{-3}	0.99988368	0.01049
C_L	Test	–	2.59×10^{-5}	2.06×10^{-5}	0.99999898	0.00429
C_L	Worst blocked	β low	1.57×10^{-5}	1.31×10^{-5}	0.99999824	0.00294
C_D	Test	–	3.07×10^{-5}	2.06×10^{-5}	0.99969797	0.04230
C_D	Worst blocked	r_{le} high	1.28×10^{-5}	8.16×10^{-6}	0.99994591	0.01714

Note: Test metrics correspond to random hold-out physics-evaluation points excluded from fitting. Worst blocked metrics correspond to the lowest-R^2 case among six blocked hold-out audits for each target, where a local region of the retained design box was excluded from training. The C_D metrics quantify surrogate reproduction of the frozen-chain output only; they do not remove the separate grid-convergence limitation of absolute drag prediction.

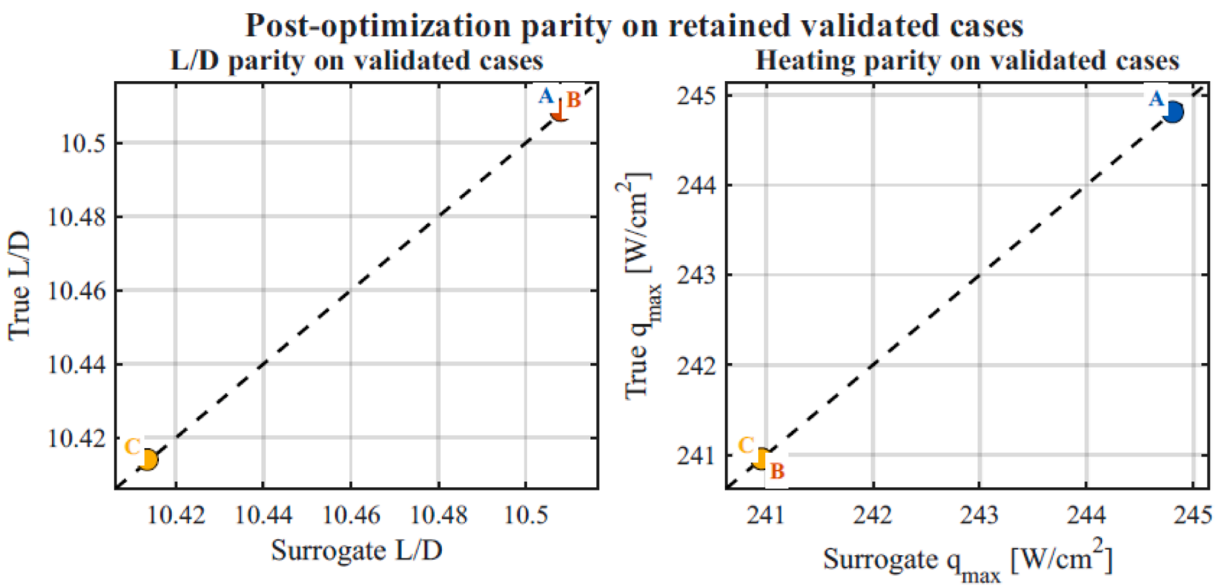

Fig. 11. Post-optimization parity on retained validated cases for L/D and q_{max}. The corresponding numerical surrogate-to-physics discrepancies, $\Delta J = \hat{J} - J_{phys}$, are reported in Table 10.

that ΔJ_k and ε_{J_k} remain sufficiently small so that the retained ranking and feasibility logic are preserved after reevaluation through the full physics chain.

The parity comparison on the validated post-opt cases shown in Fig. 11, together with the explicit post-optimization discrepancy values in Table 10, supports this interpretation. For the retained candidates, the surrogate and true-physics values of L/D and q_{max} remain tightly clustered about the one-to-one line, indicating that the surrogate layer preserves both aerodynamic ranking and thermal-feasibility discrimination at the point where the final design decision is actually made. This result is more important than nominal dataset-wide regression quality alone, because the engineering credibility of the retained optimum depends specifically on whether the post-opt candidate set survives true-physics reevaluation without a qualitative change in local ordering.

The full-physics reevaluation shown in Fig. 12 also clarifies the retained design roles. The surrogate optimum and the dataset-best feasible point correspond to the sharpest locally admissible designs and remain closest to the thermal boundary. The recommended nominal point, by contrast, occupies nearly the same local performance level while shifting slightly away from that boundary. In the corridor-based comparison, this distinction is expressed directly through the retained radii and heating levels: the recommended nominal point is associated with $r_{le} = 24.00$ mm and $q_{max} = 240.96$ W/cm^2, whereas the peak feasible point occurs at $r_{le} = 23.25$ mm and $q_{max} = 244.81$ W/cm^2. The significance of this difference is not that the nominal point achieves a different aerodynamic regime, but that it preserves the same local design logic while introducing a more comfortable thermal separation from the imposed limit (see Table 9).

Viewed together, the post-opt true-physics results validate the retained-candidate strategy adopted in the present study. The surrogate layer successfully identifies the relevant local optimum family, and the full-physics reevaluation confirms that the final design conclusion is not a regression artifact. Instead, the retained design set consists of a peak feasible solution, a slightly more conservative nominal alternative, and lower-priority neighboring candidates whose reevaluated behavior remains fully consistent with the local physical interpretation established earlier. This closes the surrogate loop: the surrogate is used to accelerate local candidate discovery, but the final retained-design judgment remains anchored to the frozen physics chain.

Because the enclosed volume is obtained from the retained parametric upper–lower surface construction, an additional geometric-volume bias sensitivity was evaluated for the volume constraint. For a retained candidate, the surrogate-to- physics volume discrepancy and the negative volume bias required to reach the imposed lower bound were de-

Table 9
Retained design parameter and true-physics performance summary.

Candidate	β [deg]	W/L	r_{le} [mm]	n_{power}	L/D	$q_{\max}$ [W/cm^2]	Δq [W/cm^2]	V [m^3]
Peak feasible	25.00	0.80	23.25	1.00	10.509195	244.810557	0.189443	13.3269
Recommended nominal	25.00	0.80	24.00	1.00	10.508754	240.955032	4.044968	13.3291

Note: $\Delta q = 245 - q_{\max}$ is the retained screening-level thermal margin. The table reports frozen-chain true-physics values for the retained peak feasible and recommended nominal designs.

Table 10
Post-optimization surrogate-to-physics discrepancy for the retained validated candidates.

Candidate	$\widehat{L/D}$	$(L/D)_{\text{phys}}$	$\Delta(L/D)$	Err. [%]	$\hat{q}_{\max}$	$q_{\max,\text{phys}}$	$\Delta q_{\max}$	Err. [%]	Feas.
Dataset-best feasible	10.508415	10.509195	-7.80×10^{-4}	0.00742	244.807040	244.810557	-3.52×10^{-3}	0.00144	Y/Y
Recommended nominal	10.508017	10.508754	-7.38×10^{-4}	0.00702	240.956558	240.955032	$+1.53 \times 10^{-3}$	0.00063	Y/Y
Dataset-center point	10.413473	10.414084	-6.11×10^{-4}	0.00587	240.956558	240.955032	$+1.53 \times 10^{-3}$	0.00063	Y/Y

Note: $\Delta J = \hat{J} - J_{\text{phys}}$ is reported for the two optimizer-relevant post-opt responses, L/D and $q_{\max}$. Here, $\hat{J}$ denotes the surrogate prediction and J_{phys} denotes the frozen-chain true-physics reevaluation. Feas. = Y/Y indicates that feasibility was predicted by the surrogate and preserved after true-physics reevaluation. The table quantifies the post-optimization discrepancy at the retained decision points and does not represent an additional global surrogate-validation claim outside the retained local design box.

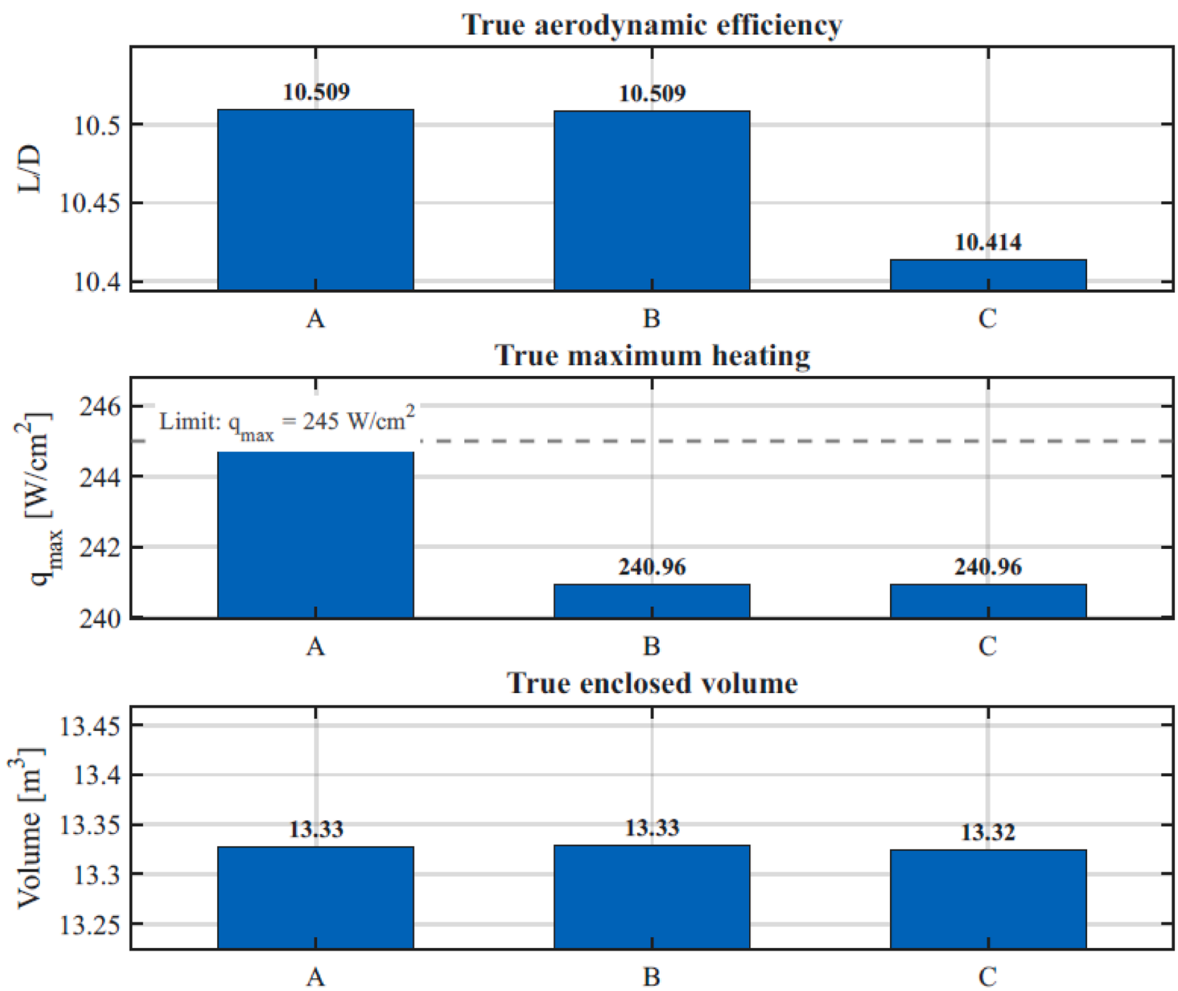

Fig. 12. True-physics comparison of the retained post-opt candidate set.

fined as

$$\Delta V = V_{\text{surr}} - V_{\text{phys}}, \tag{28}$$

$$\delta_{V,\text{lim}} = 100\frac{V_{\text{phys}} - V_{\min}}{V_{\text{phys}}}, \tag{29}$$

where $V_{\min} = 3.0$ m^3. This sensitivity does not convert the geometric envelope into a CAD-certified usable payload volume. It only quantifies whether plausible deterministic bias in the retained parametric volume estimate could change the enforced feasibility decision.

The results in Table 11 show that the volume constraint is not active for the retained candidates. The largest surrogate-to-physics volume discrepancy is only 0.0193%, whereas all retained true-physics volumes exceed the imposed lower bound by approximately 10.32–10.33 m^3. A negative geometric-volume bias of about 77.5% would be required to move any retained candidate onto the $V_{\min} = 3.0$ m^3 boundary. Even

Table 11
Enclosed-volume constraint sensitivity to deterministic geometric-volume bias for the retained post-opt candidates.

Candidate	V_{surr}	V_{phys}	ΔV	Err. [%]	$V_{\text{phys}} - V_{\min}$	Bias to fail [%]	$V_{\text{phys}}^{-10\%}$
Dataset-best feasible	13.3251	13.3269	-1.79×10^{-3}	0.0134	10.3269	77.489	11.9942
Recommended nominal	13.3272	13.3291	-1.91×10^{-3}	0.0143	10.3291	77.493	11.9962
Dataset-center point	13.3218	13.3243	-2.57×10^{-3}	0.0193	10.3243	77.485	11.9919

Note: $\Delta V = V_{\text{surr}} - V_{\text{phys}}$ and $V_{\min} = 3.0\,\text{m}^3$. Bias to fail is computed as $100(V_{\text{phys}} - V_{\min})/V_{\text{phys}}$ and denotes the negative geometric-volume bias required to move a retained candidate onto the imposed volume boundary. $V_{\text{phys}}^{-10\%} = 0.90 V_{\text{phys}}$. All three retained candidates remain feasible under a deterministic -10% volume-bias perturbation.

under a deterministic -10% volume perturbation, all retained candidates remain above $11.99\,\text{m}^3$ and therefore remain feasible. Thus, the enclosed-volume constraint is not sensitive to small geometric-parametrization fidelity errors in the retained post-opt candidate set.

4.5. Drag-share and performance interpretation

The retained local design trends become more physically interpretable when the total drag coefficient is decomposed into its constituent parts. In the present workflow,

$$C_D = C_{D,\text{inv}} + C_{D,f} + C_{D,b}, \tag{30}$$

where $C_{D,\text{inv}}$, $C_{D,f}$, and $C_{D,b}$ denote the inviscid, skin-friction, and base-drag contributions, respectively. Since the final retained designs differ only within a bounded local neighborhood, the key question is not merely which candidate yields the highest L/D, but also which drag component is primarily responsible for the remaining performance differences.

The retained drag-breakdown maps shown in Fig. 13 show that the inviscid component remains the dominant contributor throughout the trusted $(\beta, W/L)$ slice, while the friction component provides a smaller but still non-negligible secondary contribution. In contrast, the diagnostic map for the pressure-coefficient refresh indicates no effective activation within the retained local slice, which implies that the observed local drag structure is not being reshaped by a separate off-design pressure refresh mechanism in this neighborhood. This is an important consistency result, because it means that the local performance ordering is governed mainly by the smooth variation of the native inviscid and viscous terms already embedded in the frozen physics chain.

A convenient way to express the retained drag structure is through the component shares

$$s_k = \frac{C_{D,k}}{C_D} \times 100, \qquad k \in \{\text{inv}, f, b\}, \tag{31}$$

with

$$s_{\text{inv}} + s_f + s_b = 100. \tag{32}$$

Under this formulation, the retained candidate set can be compared not only through total L/D, but also through the extent to which each design is dominated by inviscid or viscous drag.

The post-opt true-physics drag-share comparison shown in Fig. 14 shows that the retained candidates exhibit a remarkably stable drag partition. Across the dataset-best feasible point, the recommended nominal design, and the representative dataset-center point, the inviscid share remains the largest component, the friction share remains the second largest, and the base-drag share stays very small. This stability is significant because it indicates that the local optimum family is not generated by a qualitatively different drag mechanism at one isolated point. Instead, the retained candidates belong to the same aerodynamic regime, and the final design selection is governed by small but meaningful differences in total drag magnitude and thermal proximity rather than by a structural transition in drag composition.

This interpretation also clarifies why the nominal design remains attractive despite not being the sharpest thermally admissible point. Since the drag-share structure changes only weakly across the retained candidate set, moving from the peak feasible point to the recommended nominal point does not imply a major loss of aerodynamic character. Rather, the nominal design preserves essentially the same inviscid-dominated drag architecture while stepping back from the thermal boundary through a modest increase in leading-edge radius. In this sense, the nominal-versus-peak distinction is not a switch between two different design philosophies, but a controlled movement within the same locally coherent drag regime.

Accordingly, the retained performance differences should be interpreted through two coupled facts. First, aerodynamic efficiency remains primarily controlled by the inviscid component and its smooth response to the local geometric variables. Second, the final design decision is not determined solely by maximizing L/D, because the relative stability of the drag shares means that a small aerodynamic sacrifice can be accepted when it produces a more robust thermal margin. This drag-share interpretation therefore provides an additional physical explanation for why the recommended nominal design remains a credible retained reference alongside the peak feasible solution.

4.6. Flight-envelope embedding of the retained designs

The local retained-design comparison becomes more meaningful when it is embedded into a broader hypersonic flight-envelope context. The local optimization problem was solved at a fixed design point of Mach 10 and 30 km altitude, but that point does not exist in isolation: its engineering relevance depends on whether it remains inside an admissible corridor once dynamic-pressure, thermal, rarefaction, and stagnation-temperature considerations are viewed together in the wider (M, h) space. In the present study, this broader admissibility context is not used as a second optimizer, but as a consistency check that connects the retained local design conclusions to a physically interpretable hypersonic operating corridor.

A compact way to express this interpretation is to define the corridor-compatible set at a given nose-radius basis R_n as

$$C(R_n) = \left\{ (M, h) \;\middle|\; \begin{array}{c} q_{\text{stag}}(M, h; R_n) \le 245\ \text{W/cm}^2, \\ Kn(M, h) \le 0.001, \\ (M, h) \in \Omega_{q_\infty} \end{array} \right\}, \tag{33}$$

where Ω_{q_∞} denotes the region compatible with the retained dynamic-pressure envelope. Under this formulation, the local retained designs are not judged only by their behavior inside the four-dimensional geometric box, but also by whether the fixed Mach-10, 30-km operating point remains embedded within the admissible corridor once the broader thermodynamic and rarefaction limits are overlaid.

Because the corridor boundaries are generated from reduced-order aerodynamic and aerothermal models, a screening-level model-error sensitivity was also evaluated for the heat-flux constraint. The positively biased heat-flux estimate was defined as

$$q_{\max}^{+} = q_{\max}(1 + \delta_q), \tag{34}$$

where δ_q is a prescribed positive heat-flux model-bias factor. The corresponding biased thermal margin is

$$\Delta q_{\text{margin}}^{+} = q_{\text{lim}} - q_{\max}(1 + \delta_q). \tag{35}$$

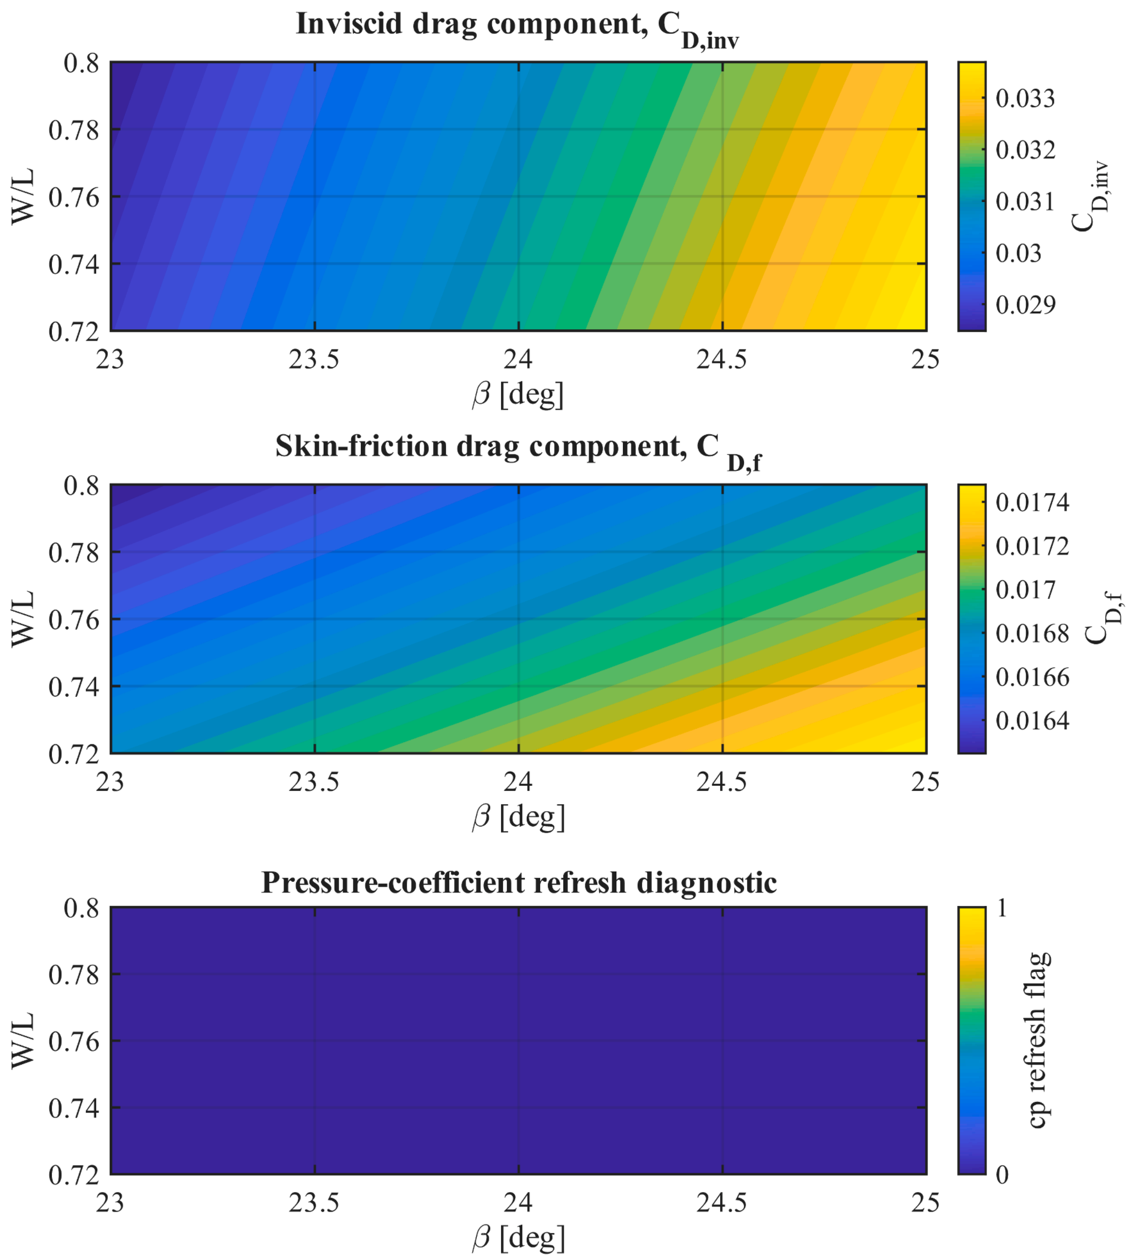

Fig. 13. Local drag-breakdown maps in the retained $(\beta, W/L)$ slice.

The heat-flux bias required to move a retained point exactly onto the thermal boundary is therefore

$$\delta_{q,\mathrm{lim}} = 100\frac{q_{\mathrm{lim}} - q_{\max}}{q_{\max}}. \tag{36}$$

This calculation is not intended as a full probabilistic uncertainty quantification over the Mach–altitude map. It is a deterministic model-error sensitivity check used to determine how much positive heat-flux bias can be tolerated before the retained design crosses the imposed thermal boundary.

The nominal-versus-peak flight-envelope embedding in Fig. 15 confirms that both retained local reference points remain inside the nominal admissible corridor at the selected design condition. However, they do not occupy the corridor in the same way. The recommended nominal point, associated with $r_{le} = 24.00$ mm, corresponds to a reported maximum heating of $q_{\max} = 240.96$ W/cm^2, whereas the peak feasible point, associated with $r_{le} = 23.25$ mm, reaches $q_{\max} = 244.81$ W/cm^2. Thus, both points remain admissible, but the peak feasible design operates much closer to the imposed thermal boundary. This confirms the interpretation developed in the preceding subsections: the recommended nominal design is not selected because it defines a different aerodynamic regime, but because it preserves nearly the same local design logic while moving the final retained reference point slightly farther away from the thermal limit.

The sensitivity results in Table 12 show that the two retained designs are not equally robust to heat-flux model bias. The peak feasible point has only 0.189 W/cm^2 of nominal thermal margin and would cross the imposed thermal limit under a positive heat-flux bias of only 0.077%. By contrast, the recommended nominal point has a 4.045 W/cm^2 nominal margin, tolerates approximately 1.679% positive heat-flux bias before reaching the limit, and remains feasible under a +1% heat-flux perturbation. Therefore, the flight-envelope result is interpreted as nominal corridor compatibility supplemented by a deterministic model-error sensitivity check, rather than as probabilistic flight-envelope robustness.

Detailed single-radius corridor views for the two retained radius bases are provided in Appendix C (Fig. 20). These supplementary views

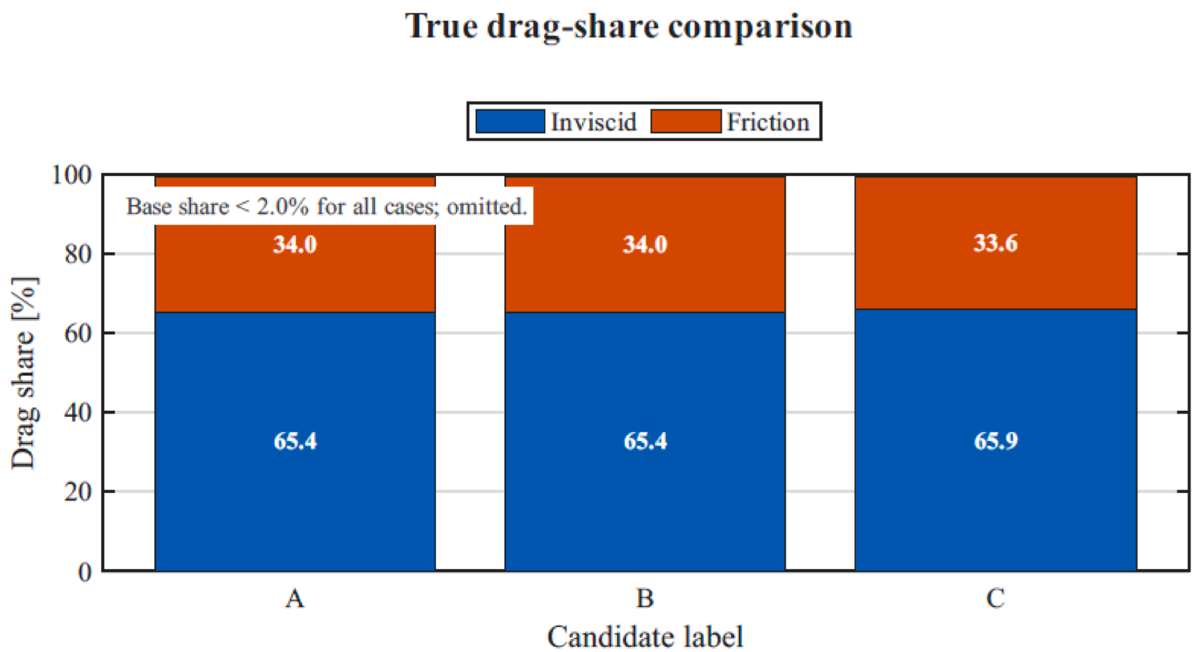

Fig. 14. Supporting drag-share diagnostic for the retained post-opt candidate set. Because the extended fine-grid GCI for C_D remains above the adopted 5% numerical-credibility threshold, the plotted drag shares are interpreted qualitatively as decomposition trends rather than as grid-independent absolute drag predictions.

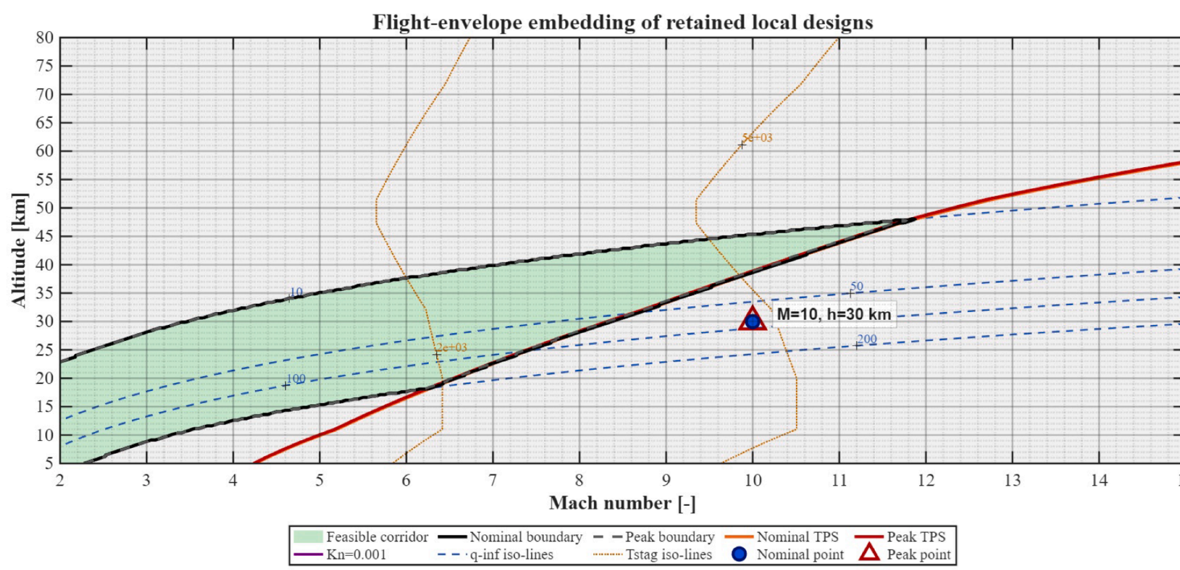

Fig. 15. Flight-envelope embedding of the retained nominal and peak local designs at the Mach-10, 30-km operating point. The retained nominal case uses $r_{le} = 24.00$ mm with $q_{\max} = 240.96$ W/cm^2, whereas the peak feasible case uses $r_{le} = 23.25$ mm with $q_{\max} = 244.81$ W/cm^2.

Table 12
Screening-level heat-flux model-error sensitivity for the retained flight-envelope design point.

Case	$q_{\max}$ [W/cm^2]	Margin [W/cm^2]	Bias to limit [%]	Margin after +1% bias [W/cm^2]	Interpretation
Peak feasible	244.811	0.189	0.077	−2.259	Not robust to positive heating bias
Recommended nominal	240.955	4.045	1.679	+1.635	Larger screening-level reserve

Note: Bias to limit is computed as $100(q_{\lim} - q_{\max})/q_{\max}$ using $q_{\lim} = 245$ W/cm^2. The +1% column reports the remaining thermal margin after applying $q^{+}_{\max} = 1.01 q_{\max}$. This table is a deterministic heat-flux model-error sensitivity check, not a full probabilistic uncertainty propagation over the Mach–altitude map.

support the same conclusion: the smaller retained radius remains admissible but operates closer to the heating-controlled boundary, whereas the slightly larger nominal radius preserves corridor compatibility while increasing separation from the thermal limit.

This embedding result is important for the interpretation of the retained local study. A purely local optimum can always be identified inside a trusted design box, but such a point is of limited engineering value if it lies too close to a wider system-level constraint boundary. By contrast, the present retained-design pair shows that the local aerodynamic optimum family remains corridor-compatible, and that the final choice between the peak feasible and recommended nominal points is governed primarily by robustness to the thermal boundary rather than by a qualitative change in flight-envelope admissibility. This strengthens the practical significance of the retained nominal design, because it links the local surrogate-assisted refinement result to a broader vehicle-level operating context.

Accordingly, the flight-envelope embedding closes the interpretation loop of the present study. Sections 4.1 and 4.2 established that β and W/L drive local aerodynamic gain while r_{le} governs thermal accessibility; Sections 4.3 and 4.4 then showed that this structure is preserved under surrogate-guided refinement and full-physics reevaluation. The corridor analysis now demonstrates that the retained nominal and peak designs are not merely local mathematical artifacts, but physically admissible members of a broader hypersonic operating envelope.

5. Discussion

This section interprets the retained local optimum family from a physical and engineering perspective, with emphasis on why the final design choice is governed not only by peak aerodynamic efficiency, but also by thermal proximity, local robustness, and the bounded credibility of the frozen evaluation chain.

5.1. Physical interpretation of the retained local optimum

The retained local optimum should not be interpreted as a single isolated point with universal significance, but rather as a compact family of closely related designs that occupy the same local aerodynamic regime. Within the trusted Mach 10, 30 km, $\alpha = 4°$ neighborhood, the best-performing candidates are consistently driven toward higher values of β and W/L, together with the smallest thermally admissible leading-edge radius. The resulting retained optimum family is therefore boundary-controlled in a physically meaningful sense: aerodynamic gain continues to increase toward the upper retained geometric bounds, while thermal admissibility prevents indefinite continuation in the same direction.

This interpretation is reinforced by the revised conical-shock-surface geometry-audit views shown in Figs. 16 and 17. These views should be read as consistency views of the revised OCM-inspired parametric quasi-waverider construction and as evidence that the retained leading-edge seed is no longer confined to a planar $x = 0$ construction. They are not intended to imply a full classical FCT/ICC streamline-tracing implementation.

To address the geometric-closure concern, the revised geometry generator also returns a closed diagnostic surface mesh for each retained reference case. The diagnostic mesh connects the lower and upper surfaces through leading-edge, trailing-edge, and side-cap faces, rather than leaving the parametric surface as an open upper–lower shell. For both retained reference cases, the generated closed diagnostic mesh contains 38,962 vertices and 77,920 triangular faces, and the internal closure flag is returned as `is_closed_surface_topology`. This check is not presented as a substitute for CFD validation; rather, it demonstrates that the revised geometry can be exported as a closed, mesh-ready diagnostic surface for subsequent CFD preprocessing, while the aerodynamic results in the present manuscript remain interpreted within the reduced-order local screening framework.

The peak feasible and recommended nominal designs share the same overall planform class, centerline growth behavior, and cross-sectional architecture. In other words, the nominal design is not a qualitatively different waverider concept, but a slightly modified member of the same local family. The aerodynamic skeleton of the design remains intact: the compression-lift-oriented lower-surface character, the same basic upper-surface closure strategy, and the same planform curvature logic are preserved across the retained pair. What changes is not the geometric identity of the vehicle, but the degree to which that same geometry is sharpened toward the thermal limit.

From this perspective, the role of the retained nominal design becomes clearer. The peak feasible point corresponds to the sharpest locally admissible member of the retained family, and therefore occupies the smallest thermal margin relative to the imposed heating constraint. The nominal point, by contrast, introduces a modest increase in bluntness while remaining within the same locally favorable aerodynamic architecture. As a result, the final retained pair should be read not as "best design" versus "alternative design," but as two neighboring realizations of the same optimum family: one biased toward maximum local performance, the other biased toward a more comfortable engineering margin.

The same conclusion follows from the drag-share interpretation developed earlier. Because the retained candidates remain in the same inviscid-dominated drag regime, the movement from the peak feasible design to the nominal design does not reflect a transition to a new aerodynamic mechanism. Instead, it represents a controlled trade within a stable local architecture. This matters physically, because it shows that the final design decision is not being made between two unrelated shapes, but between two nearby members of the same coherent local response manifold.

A useful qualitative summary of the retained optimum family is

$$
\begin{aligned}
\mathbf{x}_{\text{peak}} &\sim (\beta\uparrow,\ W/L\uparrow,\ r_{le}\downarrow_{\text{adm}},\ n_{\text{power}}\uparrow),\\
\mathbf{x}_{\text{nom}} &\sim (\beta\uparrow,\ W/L\uparrow,\ r_{le}\uparrow_{\text{slight}},\ n_{\text{power}}\uparrow)
\end{aligned}
\tag{37}
$$

where the arrows denote the observed local direction of the retained design tendency rather than exact global monotonic laws. Eq. (37) summarizes the essential engineering meaning of the retained pair: both points belong to the same locally preferred aerodynamic corner of the design box, but the nominal design steps slightly away from the thermal boundary while preserving the same overall geometric and aerodynamic identity.

This interpretation is important for the practical reading of the present study. In a tightly bounded local optimization problem, the mathematically highest feasible point is not always the most useful engineering reference. A retained design is more valuable when it preserves the same physical design logic while avoiding excessive sensitivity to the active constraint boundary. The present results show that the recommended nominal point satisfies exactly this role. It remains part of the same optimum family as the peak feasible point, but expresses that family in a slightly more robust form.

5.2. Representative CFD feasibility check of the revised closed geometry

A representative CFD feasibility check was added for the revised closed quasi-waverider geometry in order to verify that the corrected configuration can be imported, meshed, and advanced in a flow solver. This check was not intended to replace the reduced-order optimization framework or to provide a grid-converged high-fidelity aerodynamic validation. Instead, it was introduced as a meshability and flow-field plausibility assessment for one representative member of the retained design family.

The recommended nominal configuration was selected for this check because it provides a larger thermal margin than the peak-feasible candidate while remaining in the same retained local design family. A CFD-safe, volume-closed surface derived from this configuration was embedded in a rectangular external-flow farfield domain. The calculation was performed in ANSYS Fluent using a density-based, steady, inviscid Euler formulation with ideal-gas air. The freestream condition was set to $M_\infty = 10$, $h = 30$ km, and $\alpha = 4°$, consistent with the retained local design condition. Pressure-far-field boundary conditions were assigned to the farfield boundaries, and the body surfaces were treated as inviscid walls. The preliminary solution was advanced for 250 iterations without fatal solver failure.

The CFD check produced a finite flow solution, a residual history, mid-plane Mach-number and static-pressure fields, and wall-surface pressure distributions. Integrated wall-force reporting over the two body-wall zones gave a drag force of $D = 3.09 \times 10^5$ N in the freestream direction and a lift force of $L = 3.20 \times 10^5$ N in the corresponding lift direction. Using the Mach-10, 30-km freestream properties and $S_{\text{ref}} = 80$

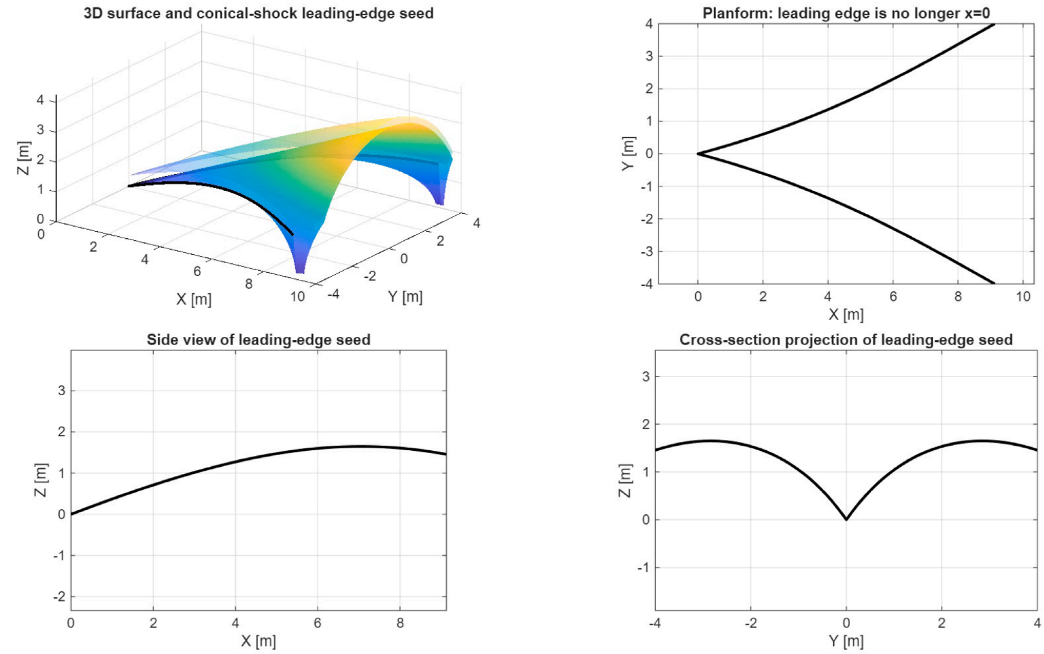

Fig. 16. Revised conical-shock-surface leading-edge seed and volume-closed diagnostic geometry for the retained peak-feasible reference case. The panels show the 3D surface with the conical-shock leading-edge seed, the planform projection demonstrating that the leading edge is not confined to $x = 0$, the side-view projection of the seed, and the cross-section projection. The maximum conical-surface residual of the leading-edge seed is 8.88×10^{-16} m. The corresponding diagnostic mesh includes upper-surface, lower-surface, leading-edge-cap, trailing-edge-cap, and side-cap faces for closed-surface preprocessing.

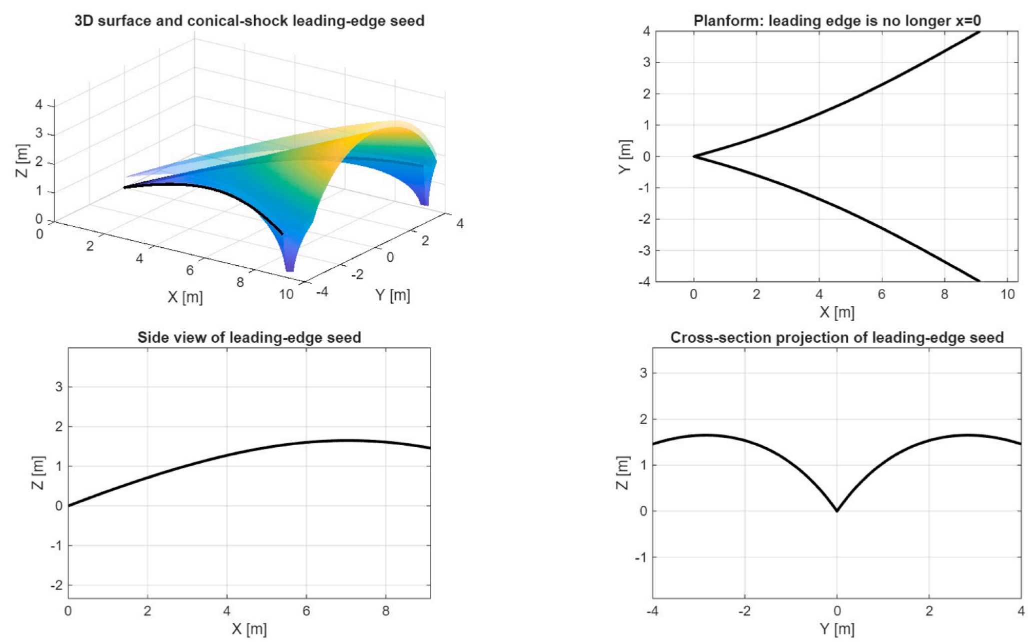

Fig. 17. Revised conical-shock-surface leading-edge seed and volume-closed diagnostic geometry for the recommended nominal reference case. The panels show the 3D surface with the conical-shock leading-edge seed, the planform projection demonstrating that the leading edge is not confined to $x = 0$, the side-view projection of the seed, and the cross-section projection. The maximum conical-surface residual of the leading-edge seed is 8.88×10^{-16} m. The corresponding diagnostic mesh includes upper-surface, lower-surface, leading-edge-cap, trailing-edge-cap, and side-cap faces for closed-surface preprocessing.

m^2, these values correspond to approximately

$$C_D \approx 0.0461, \qquad C_L \approx 0.0477, \qquad L/D \approx 1.03. \tag{38}$$

Because this was a preliminary Euler calculation on a CFD-safe trimmed geometry and a coarse feasibility mesh, these coefficients should not be interpreted as replacement values for the reduced-order retained-design tables. Rather, they demonstrate that the revised geometry is not merely a schematic surface but can be embedded in an external-flow domain, meshed, solved, and post-processed to obtain flow-field and integrated-force quantities (see Fig. 18).

5.3. Thermal margin versus peak performance

The most important engineering distinction between the two retained reference designs is not a qualitative change in aerodynamic character, but their relative distance from the active thermal constraint. As

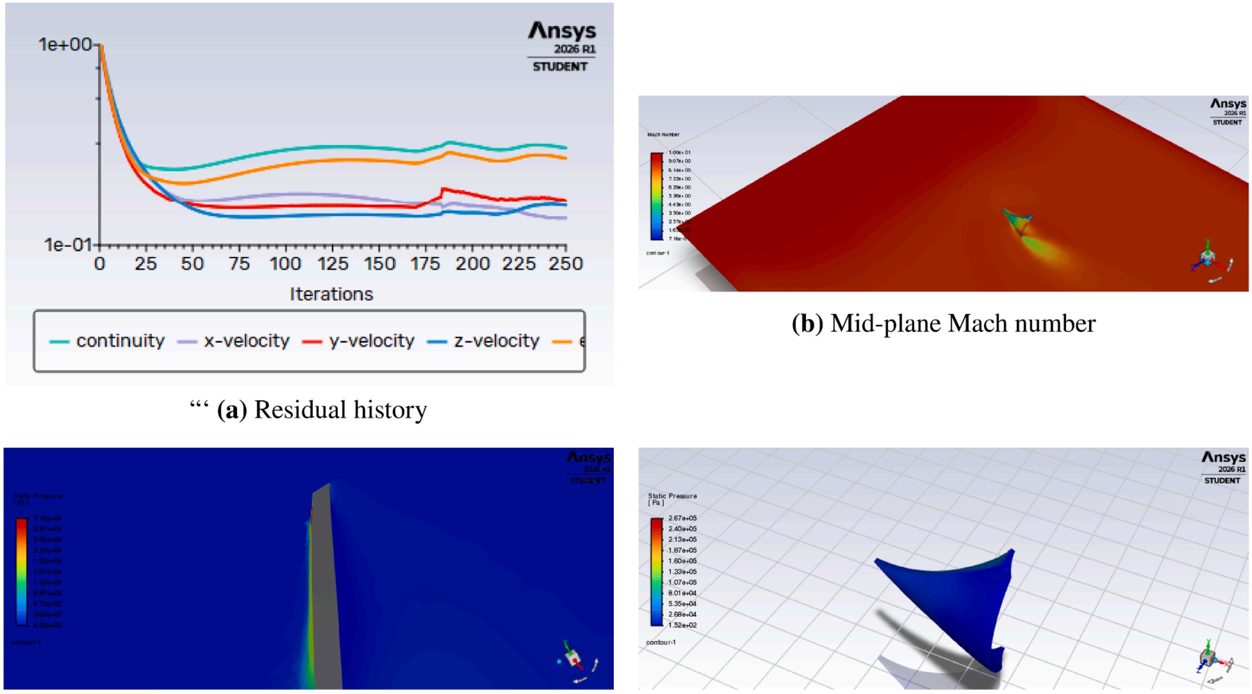

‘‘‘ **(a)** Residual history

(b) Mid-plane Mach number

(c) Mid-plane static pressure

(d) Wall-surface static pressure

Fig. 18. Representative CFD feasibility check for the revised volume-closed quasi-waverider configuration at $M_\infty = 10$, $h = 30$ km, and $\alpha = 4°$. The panels show the residual history, mid-plane Mach-number contour, mid-plane static-pressure contour, and wall-surface static-pressure distribution. The calculation was performed as a preliminary density-based Euler check and is used here to demonstrate meshability and flow-field plausibility rather than grid-converged aerodynamic validation.

established in the results section, both the peak feasible point and the recommended nominal point belong to the same retained local optimum family and both remain inside the admissible hypersonic corridor. However, they do not occupy that corridor with the same level of thermal reserve. The peak feasible point is the sharpest locally admissible member of the retained family, whereas the nominal design is a slightly blunter realization chosen to preserve nearly the same aerodynamic logic while stepping away from the heating boundary.

A convenient way to express this distinction is through the thermal margin

$$\Delta q_{\text{margin}} = q_{\text{lim}} - q_{\text{max}}, \tag{39}$$

where q_{lim} is the imposed heat-flux limit. For the retained study, $q_{\text{lim}} = 245\ \text{W/cm}^2$. Using the retained nominal and peak values,

$$\begin{aligned} \Delta q_{\text{margin,peak}} &= 245 - 244.81 = 0.19\ \text{W/cm}^2, \\ \Delta q_{\text{margin,nom}} &= 245 - 240.96 = 4.04\ \text{W/cm}^2. \end{aligned} \tag{40}$$

Thus, the recommended nominal point provides an additional absolute thermal reserve of $3.85\ \text{W/cm}^2$ relative to the peak feasible point, and its retained margin is more than one order of magnitude larger.

This result is important because it clarifies the meaning of the final design choice. If the nominal design had achieved its additional thermal margin only by moving into a clearly inferior aerodynamic regime, the retained trade would have been weak. That is not the case here. The nominal point remains embedded in the same locally favorable design family and preserves the same basic geometric and drag-share structure, while reducing sensitivity to the active heating constraint. In that sense, the nominal design should not be interpreted as a compromise imposed by poor surrogate accuracy or by an arbitrary conservative choice; rather, it is a deliberately selected engineering reference that converts a boundary-hugging local optimum into a more usable retained configuration.

The corridor views reinforce the same interpretation. In the combined nominal-versus-peak flight-envelope plot, both retained points lie within the admissible Mach–altitude corridor, but the peak feasible point lies visibly closer to the heating-controlled edge, whereas the nominal point remains inside the same corridor with a clearer thermal buffer. The single-radius corridor views tell the same story from another angle: reducing the effective nose-radius basis keeps the design point admissible but pushes it nearer the thermal boundary, while increasing the basis slightly preserves admissibility and improves separation from that limit. Accordingly, the difference between the two retained references is not corridor compatibility versus incompatibility, but narrow-margin admissibility versus more comfortable admissibility within the same operating family.

This distinction has direct engineering significance. In tightly constrained hypersonic design, a mathematically maximal feasible point is not always the most useful retained reference if it is excessively close to an active limit. A design with nearly unchanged aerodynamic quality but visibly improved distance from the constraint boundary is often more valuable for subsequent vehicle-level interpretation, design maturation, and uncertainty-aware decision-making. The present retained pair illustrates exactly this situation: the peak feasible point is useful as the upper local performance marker, while the recommended nominal point is the more balanced engineering reference because it preserves the same local optimum logic with a materially improved thermal reserve.

For this reason, the recommended nominal design is best interpreted as the practically preferred representative of the retained optimum family. The peak feasible point remains important because it defines the performance edge of the trusted local box, but the nominal point provides a more robust expression of the same underlying design logic. The difference between them is therefore not merely numerical; it captures the central engineering lesson of the present bounded local study, namely that local peak performance and local design usefulness are not always identical once thermal proximity is treated as an explicit part of the decision.

5.4. Trust region, model limits, and scope of claims

The final point that must be made explicit is that the present study is intentionally local in both geometric and interpretive scope. The retained surrogate framework, the post-opt confirmation stage, and the nominal-versus-peak comparison are all valid only within the trusted neighborhood defined by the frozen Mach 10, 30 km, $\alpha = 4°$ regime and the retained four-dimensional design box. Accordingly, the conclusions of this paper should not be read as statements about a universal waverider optimum, a global Mach–altitude–angle-of-attack surrogate, or a fully validated absolute-performance predictor over a broad hypersonic design space.

A compact way to express the retained trust region is

$$\mathcal{D}_{\mathrm{tr}} = \left\{ \mathbf{x} \;\middle|\; \begin{aligned} 23.0° \le \beta \le 25.0°, \\ 0.72 \le W/L \le 0.80, \\ 23.25\ \mathrm{mm} \le r_{le} \le 27.00\ \mathrm{mm}, \\ 0.85 \le n_{\mathrm{power}} \le 1.00 \end{aligned} \right\} \tag{41}$$

Within this formulation, the surrogate, refinement, and retained-design logic are intended to be used only for $\mathbf{x} \in \mathcal{D}_{\mathrm{tr}}$ under the fixed operating condition adopted in the manuscript. Outside this region, the smoothness, ranking consistency, and optimizer-relevant fidelity demonstrated in the retained local study should not be assumed without further verification.

The atmosphere, Prandtl–Meyer, corrected oblique-shock angle audit, and thermal cross-check blocks showed strong or acceptable agreement, whereas the Taylor–Maccoll verification, literature-based absolute L/D comparisons, and the fine-grid C_D GCI exhibited non-negligible discrepancies. In particular, the added 120×200 audit reduced the C_D $\mathrm{GCI}_{\mathrm{fine}}$ from 6.074% to 5.748%, but this value still exceeds the adopted 5% numerical-credibility threshold. For this reason, the revised manuscript treats absolute drag-coefficient values and drag-share decompositions as supporting diagnostics rather than as grid-independent quantitative predictions. This limitation does not remove the local ranking value of the retained workflow, but it narrows the interpretation of drag-sensitive conclusions. The oblique-shock correction does not change the bounded interpretation of the study: β-dependent trends are still used for local comparative design ranking inside a common reduced-order chain, while broad absolute aerodynamic prediction remains outside the claim scope of the present manuscript. The grid-sensitivity results further showed that the retained workflow is sufficiently stable for local L/D-level ranking and volume-preserving feasibility interpretation, but not yet fully closed for absolute drag prediction. The same audit also confirmed that the observed orders of convergence remain low for both L/D and C_D; therefore, the numerical evidence is not presented as an asymptotic mesh-convergence proof. Instead, it defines the level at which the frozen chain can be interpreted: reliable enough for bounded local ranking and trend analysis, but insufficient for high-precision absolute drag prediction. Therefore, the present manuscript supports local comparative inference, trend interpretation, and retained-candidate confirmation more strongly than it supports broad absolute-performance extrapolation.

A second limitation concerns geometric interpretation. The geometry-generation framework used here is explicitly OCM-inspired and parametric rather than a full classical streamline-tracing implementation of the Osculating Cone Method. This distinction matters because it defines the level at which geometric fidelity should be interpreted. The present workflow is well suited for constructing a physically interpretable local family of quasi-waverider geometries and for tracking how the active variables influence performance and feasibility inside that family. However, the results should not be over-interpreted as direct substitutes for higher-fidelity streamline-traced or CFD-resolved shape optimization carried out over a broader configuration class.

A third limitation concerns regime scope. The retained local point lies within the admissible continuum-oriented corridor used in the present study, and the codebase includes real-gas and rarefied-gas support logic. Nevertheless, those modules are not used here to claim validated predictive capability across strongly rarefied, strongly dissociated, or fully off-design flight conditions. Instead, they function as part of the consistency architecture that helps define where the retained design point sits relative to wider hypersonic constraints. In this sense, the present study should be read as a bounded engineering investigation of a locally credible design neighborhood, not as a regime-universal flow solver.

A fourth limitation concerns aerothermal fidelity. The retained Fay–Riddell/Sutton–Graves and reference-enthalpy heating treatment does not include coupled viscous boundary-layer solution, transition-resolved heat-transfer prediction, catalytic or partially catalytic wall chemistry, finite-rate gas–surface reactions, ablation, or material-response feedback. Therefore, the heat-flux constraint and the reported thermal margins are used as comparative engineering screening quantities. They are suitable for identifying whether a candidate moves toward or away from the retained thermal boundary inside the local design box, but they should not be interpreted as final TPS sizing loads or material-certification predictions. Higher-fidelity CFD/aerothermal analysis with appropriate wall-catalysis and material-response modeling is required before vehicle-level thermal design decisions can be made.

A fifth limitation concerns uncertainty propagation in the flight-envelope maps. The Mach–altitude corridor shown in the present work is a deterministic screening map based on the retained reduced-order atmosphere, dynamic-pressure, heating, Knudsen-number, and stagnation-temperature models. It does not constitute a full probabilistic uncertainty quantification over model-form error, atmospheric variability, wall-catalysis uncertainty, transition uncertainty, or TPS material-response uncertainty. To address this limitation at the screening level, the revised manuscript reports a heat-flux model-error sensitivity at the retained design point. This check shows that the peak feasible point is very close to the thermal boundary, while the recommended nominal point has a larger but still model-dependent thermal reserve. Therefore, the flight-envelope result is interpreted as nominal corridor compatibility with a deterministic model-error sensitivity check, not as probabilistic flight-envelope robustness.

A sixth limitation concerns the interpretation of enclosed volume. The enclosed-volume quantity should be interpreted as a parametric geometric-envelope volume obtained from the retained upper–lower surface construction, not as a CAD-certified usable payload volume after structural layout, thermal-protection thickness, tanks, ducts, or subsystems are inserted. The added volume-bias sensitivity shows that the retained feasibility decision is insensitive to deterministic geometric-volume bias at the present constraint level, but detailed usable-volume certification remains a later CAD/packaging task.

A seventh limitation concerns vehicle-level integration and objective completeness. The present local optimization prioritizes aerodynamic efficiency, thermal admissibility, lift feasibility, and a parametric enclosed-volume constraint. It does not yet include structural mass, TPS thickness, internal packaging, center-of-gravity constraints, propulsion integration, controllability, or mission-level trajectory coupling as simultaneous design objectives. A broader vehicle-level multi-objective formulation, including structural weight and usable packaged volume, is therefore required before the retained nominal geometry can be interpreted as a complete integrated vehicle design.

These limitations do not weaken the contribution of the work; rather, they define its proper strength. The manuscript does not claim that the surrogate can replace the frozen physics chain everywhere, nor that the retained nominal design is the best possible waverider under all conditions. What it does show is that, once the active variable roles are interpreted physically and a trusted local neighborhood is identified, a target-specific surrogate layer can reproduce the optimizer-relevant structure of that neighborhood with sufficient fidelity to support retained-design discovery, post-opt confirmation, and engineering decision-making.

Accordingly, the scope of the present claims may be stated succinctly as follows: the study establishes a locally validated and physically interpretable surrogate-assisted design strategy for a bounded Mach 10, 30 km, $\alpha = 4°$ waverider problem, in which β and W/L act mainly as aerodynamic drivers, r_{le} acts mainly as a thermal-access variable, and the final design choice is governed by the balance between local peak performance and thermal margin inside a trusted local box. Any extension beyond this trust region, beyond this flight condition, or beyond this modeling fidelity should be treated as future work rather than as an implicit conclusion of the present manuscript.

6. Conclusion

This study presented a bounded local surrogate-assisted design investigation for a hypersonic waverider at Mach 10, 30 km altitude, and $\alpha = 4°$, using a frozen physics-based evaluation chain, local refinement, target-specific surrogate modeling, optimizer-level backchecking, and post-optimization true-physics reevaluation. The contribution of the work is therefore not a universal hypersonic surrogate or a global optimum claim, but a physically interpretable and locally validated design strategy for a trusted four-dimensional neighborhood defined by β, W/L, r_{le}, and n_{power}.

Within that retained local regime, the results showed that β and W/L act primarily as aerodynamic performance drivers, whereas r_{le} acts mainly as a thermal-access variable with comparatively weak influence on L/D over the examined range. The best-performing feasible candidates were found near the upper retained bounds of β and W/L and near the smallest thermally admissible values of r_{le}, confirming that the local optimum is boundary-controlled rather than interior within the examined design box. This boundary-controlled behavior should not be interpreted as a global waverider optimum structure; it is specific to the constrained Mach-10, 30-km, $\alpha = 4°$ trust region and should be re-evaluated in expanded geometric and flight-envelope studies.

The retained peak feasible and recommended nominal designs quantify the main aerodynamic–thermal trade-off. Moving from the peak feasible design to the nominal design changes L/D only from 10.509195 to 10.508754, whereas the maximum heat flux decreases from 244.810557 to 240.955032 W/cm^2 and the screening-level thermal margin increases from 0.189443 to 4.044968 W/cm^2. Thus, the nominal design provides a substantially larger thermal reserve with a negligible aerodynamic-efficiency penalty inside the retained local model. This reserve should be interpreted as improved screening-level tolerance to heat-flux model bias and thermal-boundary proximity, not as final TPS material certification.

The surrogate layer reproduced the optimizer-relevant structure of the retained local box with sufficient fidelity to identify a compact family of high-value candidates, and the subsequent true-physics reevaluation confirmed that the final retained ranking was not a regression artifact. In this sense, the surrogate did not replace the frozen physics chain; instead, it acted as a local accelerator that sharpened candidate discovery while leaving the final retained-design judgment anchored to full-physics reevaluation. The resulting design interpretation is therefore physically coherent: the peak feasible point defines the local performance edge, whereas the recommended nominal point represents a slightly more conservative member of the same optimum family, selected because it preserves essentially the same local design logic while moving away from the active thermal boundary.

The verification results also define the limits of the present claims. The atmosphere, Prandtl–Meyer, corrected oblique-shock angle audit, and thermal cross-check blocks showed strong or acceptable agreement, whereas the Taylor–Maccoll verification, literature-based absolute aerodynamic comparisons, and the fine-grid C_D GCI exhibited non-negligible discrepancies. The added 120×200 audit reduced the C_D GCI$_{\text{fine}}$ from 6.074% to 5.748%, but this remained above the adopted 5% threshold. In addition, the observed convergence orders remained low, confirming that the grid study should be regarded as a conservative numerical-consistency audit rather than a formal asymptotic mesh-convergence demonstration. Therefore, absolute drag-coefficient and drag-share results were treated as supporting diagnostics rather than as grid-independent quantitative predictions. The oblique-shock benchmark audit clarified that the previously reported large error was caused by an angle-convention inconsistency in the benchmark table; nevertheless, the retained β-dependent conclusions remain deliberately restricted to local comparative behavior within the stated design box. Accordingly, the present manuscript supports local comparative inference, retained-design ranking, and bounded engineering interpretation more strongly than broad absolute-performance extrapolation over a wide hypersonic regime. In addition, because the retained geometry generator is OCM-inspired and parametric rather than a full classical streamline-traced osculating-cone implementation, the geometric conclusions should be interpreted as trends within the retained quasi-waverider family and not as claims of full OCM geometric fidelity.

The flight-envelope maps should also be extended with formal uncertainty propagation over model-form error, atmospheric variability, transition/catalysis assumptions, and TPS material response. In the present work, the added heat-flux-bias sensitivity provides only a screening-level uncertainty check for the retained Mach–altitude embedding. A deterministic enclosed-volume bias check was also added, showing that the retained candidates remain far from the $V_{\min} = 3.0$ m^3 boundary even under a -10% geometric-volume perturbation; future work should replace this geometric-envelope volume with CAD-level usable-volume and packaging assessments.

The retained results also close the specific methodological gaps identified in the Introduction. The study separates the local roles of the active variables, defines an explicit four-dimensional trust region before surrogate construction, validates the surrogate through random and blocked hold-out checks, confirms retained candidates through post-optimization true-physics reevaluation, distinguishes a peak feasible point from a more conservative nominal engineering reference, and states the fidelity limits of the OCM-inspired parametric geometry. These points define the contribution as a locally validated design strategy rather than as a global hypersonic optimization result. Under these bounds, the main conclusion is that a trusted local waverider design neighborhood can be explored effectively through a target-specific surrogate-assisted workflow, provided that the active variable roles are interpreted physically, the trust region is stated explicitly, and the final retained candidates are revalidated through the underlying physics chain. Future work should therefore focus on extending the retained framework beyond the current trust region through strengthened shock/conical-flow verification, broader off-design assessment, and higher-fidelity aerodynamic and aerothermal confirmation over enlarged Mach–altitude–angle-of-attack domains. In particular, coupled boundary-layer/aerothermochemical modeling, catalytic-wall sensitivity, transition effects, and TPS material-response analysis should be incorporated before the retained thermal margins are used for final vehicle-level thermal-protection design.

CRediT authorship contribution statement

Murat Metehan Türkoğlu: Writing – review & editing, Writing – original draft, Visualization, Supervision, Software, Methodology, Data curation, Conceptualization; **Emirhan Dönmez:** Writing – original draft, Resources, Investigation, Formal analysis, Conceptualization; **Ibrahim Özkol:** Writing – review & editing, Supervision, Conceptualization.

Data and code availability

The data and source-code package required to reproduce the retained local study is provided through a public Google Drive reproducibility folder. The shared package includes the MATLAB source files, retained local design inputs, physics-based evaluation data, surrogate-validation tables, random and blocked hold-out results, post-optimization true-physics reevaluation cases, grid-convergence raw data, thermal cross-validation data, figure-generation outputs, and a step-by-step README file.

Public reproducibility folder: https://drive.google.com/drive/folders/1dIcgDgZGN3ui69d_o758ex7_12kinx0v

The shared package is intended to reproduce the bounded Mach 10, 30 km, $\alpha = 4°$ local design study reported in this manuscript. It is not intended to provide a globally validated hypersonic design code outside the stated trust region.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Appendix A. Supplementary implementation details for geometry construction and surrogate representation

A.1. Lower-surface mapping

TThe main text states that the retained lower surface is generated through conical-flow-informed downstream mapping from a leading-edge seed that lies on the conical shock surface. In the revised construction, each spanwise station starts from its own leading-edge point (x_{le}, y_{le}, z_{le}) rather than from a common $x = 0$ plane. The local streamwise coordinate is therefore defined as

$$s = \frac{x - x_{le}}{L - x_{le}}, \qquad 0 \le s \le 1, \tag{42}$$

where L is the retained vehicle length.

At each spanwise station, the local conical-flow angle is obtained from the Taylor–Maccoll solution through a monotone interpolation of the conical-flow field, written here in compact form as

$$\theta_{\mathrm{loc}}(s) = \mathcal{I}_\theta(s; \beta, \theta_c), \tag{43}$$

where β is the prescribed shock angle and θ_c is the cone angle associated with the retained conical solution. When a fallback linear form is used for numerical robustness, the same quantity may be written as

$$\theta_{\mathrm{loc}}(s) = (1 - s)\beta + s\theta_c. \tag{44}$$

For a spanwise station starting from (x_{le}, y_{le}, z_{le}), the downstream spanwise contraction is represented by

$$\chi(s) = \frac{\sin\theta_{\mathrm{loc}}(s)}{\sin\beta}, \tag{45}$$

so that the retained lower-surface spanwise coordinate becomes

$$Y_\ell(x) = y_{le}\left[1 - s(1 - \chi(s))\right]. \tag{46}$$

The corresponding lower-surface height is obtained from the local conical-surface relation,

$$Z_\ell(x) = \sqrt{\max\left(\left[x\tan\theta_{\mathrm{loc}}(s)\right]^2 - Y_\ell^2(x), 0\right)}. \tag{47}$$

At $s = 0$, this formulation recovers the leading-edge seed point on the conical shock surface, $(X_\ell, Y_\ell, Z_\ell) = (x_{le}, y_{le}, z_{le})$, and therefore avoids the earlier planar leading-edge interpretation. This mapping remains a reduced-order, conical-flow-informed parametric construction and should not be interpreted as a full FCT/ICC streamline-traced osculating-cone method.

A.2. Upper-surface closure and blunting

As stated in the main text, the retained upper surface was introduced as a freestream-aligned and volume-preserving closure surface rather than through a separate inverse-design procedure. For a given spanwise station, let z_{le} denote the local leading-edge height associated with the conical-shock-surface seed and let

$$s = \frac{x - x_{le}}{L - x_{le}}$$

be the local nondimensional streamwise coordinate. The retained upper-surface closure is written as

$$X_u(x, y) = X_\ell(x, y), \qquad Y_u(x, y) = Y_\ell(x, y),$$

$$Z_u(x, y) = z_{le} + t_{\min}(1 - s),$$

where

$$t_{\min} = 0.005\,L$$

is the minimum structural thickness scale used in the retained implementation. In the revised volume-closed diagnostic implementation, the upper–lower clearance is allowed to scale upward when required to preserve the retained enclosed-volume constraint while maintaining the same conical-shock-surface leading-edge seed.

To prevent nonphysical upper/lower surface intersection, the implementation enforces the inequality

$$Z_u(x, y) \ge Z_\ell(x, y) + 0.1\,t_{\min}.$$

Leading-edge bluntness is then applied over the first n_{blend} streamwise stations, where

$$n_{\mathrm{blend}} = \min(5, n_{\mathrm{stations}}).$$

Defining the discrete blend factor

$$b_j = \frac{n_{\mathrm{blend}} - j + 1}{n_{\mathrm{blend}}}, \qquad j = 1, \ldots, n_{\mathrm{blend}},$$

the retained blunting update is written as

$$Z_\ell^b(i, j) = Z_\ell(i, j) - r_{le}\, b_j \left[1 - \cos\left(\frac{\pi}{2}\frac{j}{n_{\mathrm{blend}}}\right)\right],$$

$$Z_u^b(i, j) = Z_u(i, j) + r_{le}\, b_j \left[1 - \cos\left(\frac{\pi}{2}\frac{j}{n_{\mathrm{blend}}}\right)\right],$$

where r_{le} is the retained leading-edge radius.

This construction makes the geometric role of r_{le} explicit. Within the retained local study, bluntness modifies the near-leading-edge thickness distribution while preserving the same overall local planform family and the same broader OCM-inspired geometric logic.

A.3. Supplementary surrogate representation

The main text already defines the target-specific local surrogate strategy through the mapping

$$\hat{y}_t = f_t(\beta, W/L, r_{le}, n_{\mathrm{power}})$$

and gives the ridge-regularized polynomial fitting statement. For completeness, let the retained local design vector be

$$\mathbf{x} = \begin{bmatrix}\beta & W/L & r_{le} & n_{\mathrm{power}}\end{bmatrix}^{\mathsf{T}}.$$

For the polynomial family, the surrogate may be written in basis form as

$$\hat{y}_t^{(\mathrm{poly})}(\mathbf{x}) = \boldsymbol{\phi}(\mathbf{x})^{\mathsf{T}}\mathbf{w}_t,$$

where $\boldsymbol{\phi}(\mathbf{x})$ is the retained polynomial feature vector and $\mathbf{w}_t$ is obtained from the ridge-regularized problem reported in the main text.

For the RBF-type family, a generic local representation may be written as

$$\hat{y}_t^{(\mathrm{RBF})}(\mathbf{x}) = c_{0,t} + \mathbf{c}_t^{\mathsf{T}}\mathbf{x} + \sum_{i=1}^{N} a_{i,t}\,\varphi_t\!\left(\frac{\|\mathbf{x}-\mathbf{x}_i\|}{\sigma_t}\right),$$

where N is the number of retained local training samples, $\mathbf{x}_i$ are the retained sample locations, σ_t is a target-specific width parameter, $\varphi_t(\cdot)$ is the selected radial-basis kernel, and $a_{i,t}$, $c_{0,t}$, and $\mathbf{c}_t$ are fitted coefficients.

Under this supplementary form, the surrogate layer remains a target-specific local accelerator for candidate discovery inside the trusted neighborhood, while the final retained-design judgment remains anchored to reevaluation through the frozen physics chain.

Appendix B. Detailed constrained sweep slices

See Fig. 19.

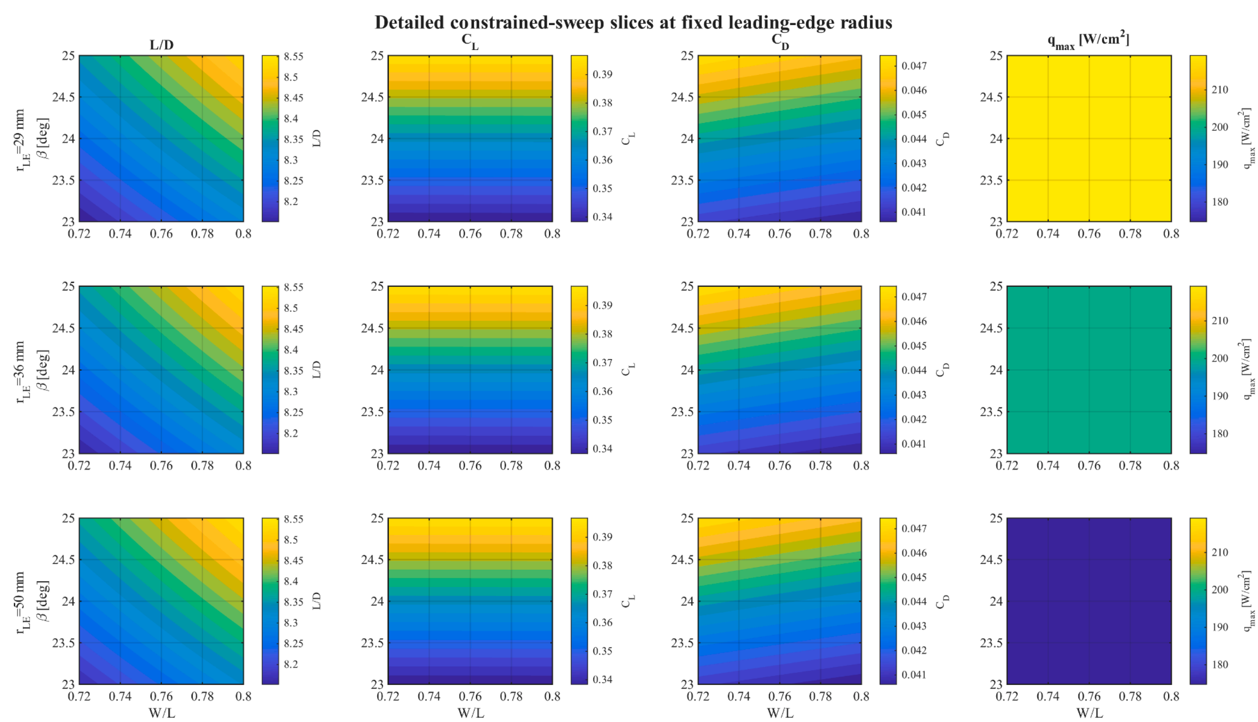

Fig. 19. Detailed constrained-sweep slices at fixed leading-edge radii, showing L/D, C_L, C_D, and $q_{\max}$ over the retained $(\beta, W/L)$ plane for r_{le} = 29, 36, and 50 mm. These maps provide the supplementary constrained-sweep support for the thermal-access interpretation discussed in Section 4.2.

Appendix C. Single-radius corridor views

See Fig. 20.

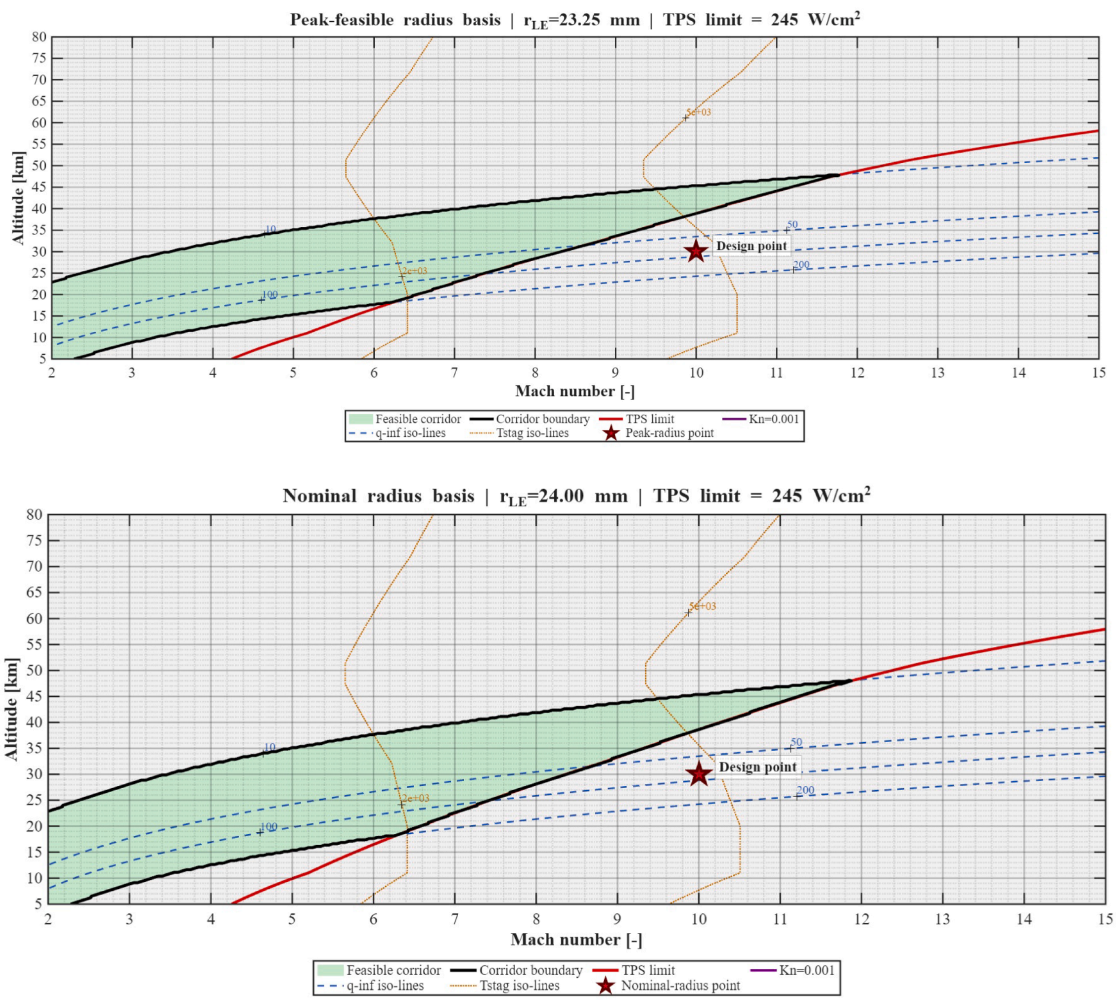

Fig. 20. Single-radius corridor views of the Mach-10, 30-km design point for two retained radius bases. The upper panel corresponds to the peak-feasible radius basis, r_{le} = 23.25 mm, whereas the lower panel corresponds to the retained nominal radius basis, r_{le} = 24.00 mm.

Appendix D. Reproducibility package and execution sequence

A supplementary reproducibility package is provided through the public Google Drive folder reported in the Data and Code Availability section. The package contains the retained MATLAB source files, local physics-evaluation dataset, surrogate-validation tables, blocked hold-out definitions, post-optimization true-physics reevaluation cases, raw grid-convergence data, thermal cross-validation data, figure-generation outputs, and a step-by-step README file.

The reproduction workflow is organized as follows:

1. Fix the operating condition at $M_\infty = 10$, $h = 30$ km, $\alpha = 4°$, and $L_{ref} = 10$ m.
2. Generate candidate geometries inside the retained four-dimensional design box in $(\beta, W/L, r_{le}, n_{power})$.
3. Evaluate each candidate through the frozen atmosphere, conical-flow, geometry, aerodynamic, aerothermal, and geometric-volume modules.
4. Apply the hard feasibility filter using $q_{\max} \leq 245$ W/cm^2, $V \geq 3.0$ m^3, and $C_L \geq 0.01$.
5. Train the target-specific local surrogate models for L/D, $q_{\max}$, V, C_L, and C_D using only points inside the trusted design box.
6. Perform random hold-out and blocked hold-out validation using the fixed split definitions provided in the repository.
7. Conduct surrogate-assisted local candidate screening and optimizer-level backchecking.
8. Reevaluate the retained peak feasible, recommended nominal, and comparison candidates through the frozen physics chain.
9. Reproduce the grid-convergence audit, Fay–Riddell/Sutton–Graves thermal cross-validation, heat-flux-bias check, and volume-bias check.
10. Regenerate the manuscript figures and tables using the plotting scripts and retained output files.

The accompanying README file specifies the folder structure, software environment, execution order, expected input/output files, and interpretation limits of the shared package. The package is intended to reproduce the bounded local study reported here and should not be interpreted as a globally validated hypersonic design environment outside the stated trust region.

References

[1] K.G. Bowcutt, J.D. Anderson, D. Capriotti, Viscous optimized hypersonic waveriders, in: 25Th Aerospace Sciences Meeting, American Institute of Aeronautics and Astronautics, 1987. https://doi.org/10.2514/6.1987-272

[2] F. Ding, J. Liu, C.-B. Shen, Z. Liu, S.-H. Chen, X. Fu, An overview of research on waverider design methodology, Acta Astronaut. 140 (2017) 190–205. https://doi.org/10.1016/j.actaastro.2017.08.027

[3] F. Ding, J. Liu, C.-B. Shen, W. Huang, Z. Liu, S.-H. Chen, An overview of waverider design concept in airframe/inlet integration methodology for air-breathing hypersonic vehicles, Acta Astronaut. 152 (2018) 639–656. https://doi.org/10.1016/j.actaastro.2018.09.002

[4] L.-L. Chen, X.-L. Deng, Z. Guo, Z.-X. Hou, W.-K. Wang, A novel approach for design and analysis of volume-improved osculating-cone waveriders, Acta Astronaut. 161 (2019) 430–445. https://doi.org/10.1016/j.actaastro.2019.02.033

[5] J. Son, C. Son, K. Yee, A novel direct optimization framework for hypersonic waverider inverse design methods, Aerospace 9 (7) (2022) 348. https://doi.org/10.3390/aerospace9070348

[6] C. Liu, P. Bai, Waverider design using osculating method, AIAA J. 62 (5) (2024) 1643–1661. https://doi.org/10.2514/1.J063558

[7] R.P. Starkey, M.J. Lewis, Analytical off-Design lift-to-Drag-Ratio analysis for hypersonic waveriders, J. Spacecr. Rockets 37 (5) (2000) 684–691. https://doi.org/10.2514/2.3618

[8] J.-X. Liu, Z.-X. Hou, G.-h. Ding, X.-Q. Chen, X.-Q. Chen, Numerical and experimental study on waverider with blunt leading edge, Comput. Fluids 84 (2013) 203–217. https://doi.org/10.1016/j.compfluid.2013.06.005

[9] S.-Q. Guo, W. Liu, C.-A. Zhang, Y. Liu, F.-M. Wang, Aerodynamic optimization of hypersonic blunted waveriders based on symbolic regression, Aerosp. Sci. Technol. 144 (2024) 108801. https://doi.org/10.1016/j.ast.2023.108801

[10] X. Chen, L. Liu, T. Long, Z. Yue, A reduced order aerothermodynamic modeling framework for hypersonic vehicles based on surrogate and POD, Chin. J. Aeronaut. 28 (2015) 1328–1342. https://doi.org/10.1016/j.cja.2015.06.024

[11] F. Qu, T. Wang, C. Liu, J. Fu, J. Bai, Aerodynamic shape optimization of the vortex-shock integrated waverider over a wide speed range, Aerosp. Sci. Technol. 143 (2023) 108696. https://doi.org/10.1016/j.ast.2023.108696

[12] J. Carvalho, E. Mooij, Integrated shape and trajectory optimisation of hypersonic waveriders, in: AIAA SCITECH 2024 Forum, American Institute of Aeronautics and Astronautics, 2024. https://doi.org/10.2514/6.2024-0156

[13] Z. Zhao, Y. Ma, Y. Tian, Z. Ding, H. Zhang, S. Tong, Research on integrated design method of wide-range hypersonic vehicle/engine based on dynamic multi-objective optimization, Aerosp. Sci. Technol. 160 (2025) 110029. https://doi.org/10.1016/j.ast.2025.110029

[14] Y. Chen, J. Li, J. Cai, Aerodynamic shape optimization of hypersonic aircraft using data-driven generative nonlinear parameterization, Chin. J. Aeronaut. 39 (3) (2026) 103924. https://doi.org/10.1016/j.cja.2025.103924

[15] Q. Yang, S. Liu, X. Wu, T. Han, H.-S. Shin, R. Tekin, A review of trajectory planning for high speed vehicles: current advancements and future prospects, Aerosp. Sci. Technol. 176 (2026) 112667. https://doi.org/10.1016/j.ast.2026.112667

[16] K. An, W. Huang, S. Liu, Fixed-time robust attitude tracking control for high-speed aircraft: a precise funnel-guided approach, J. Zhejiang Univ.-SCI. A 27 (2026) 348–364. https://doi.org/10.1631/jzus.A2500172

[17] S. Corda, J.D. Anderson, J., Viscous optimized hypersonic waveriders designed from axisymmetric flow fields, in: 26Th Aerospace Sciences Meeting, Reno, NV, 1988. https://doi.org/10.2514/6.1988-369

Murat Metehan Türkoglu received his M.Sc. in Physics from Ulupınar Observatory and Astrophysics Research Center (2010) and his Ph.D. in Physics from Istanbul Technical University (2020), with a focus on magnetized neutron stars and disk interactions. He is currently pursuing a second Ph.D. in Defense Technologies at ITU and is also an undergraduate student in the Department of Aerospace Engineering at Middle East Technical University (ODTÜ). His research encompasses neutron star physics, alternative gravity theories, reinforcement learning, and aerospace systems, emphasizing multidisciplinary and cross-domain approaches. Dr. Türkoglu has authored, translated, and edited numerous scientific publications and currently serves as a faculty member at Istanbul Gelişim University, Department of Aeronautical Engineering.

Prof. Dr. İbrahim Özkol (born 1962, İzmir, Türkiye) is a Professor of Aeronautical Engineering at Istanbul Technical University (ITU). He received his BSc and MSc in Aeronautical Engineering and his PhD in Aeronautical and Astronautical Engineering from ITU in 1985, 1988, and 1992, respectively, and has been on the ITU faculty since 1987 (Associate Professor in 1996, Full Professor since 2005). In 1992–1993 he was seconded by the Defence Industry Undersecretariat to Spain, where he worked on the CASA-3000 and CN-235 aircraft production programs. He has since led or contributed to numerous aerospace and defence projects supported by national agencies, industry, and international partners, including the Turkish defence sector, the EU, Boeing and TÜBİTAK. Prof. Özkol is the founder and coordinator of the ITU–Boeing–Turkish Airlines Air Transport Management MBA and the ITU–TUSAŞ Air Vehicle Structures and Materials MBA. He is the founding Director of the ITU Institute of Aviation (since 2019) and of the ITU Center for Aviation and Space Technologies (since 2016), which contributed to ITU receiving the 2018 YÖK International Collaboration Award. He served as ITU Vice Rector between 2012 and 2018. Externally, he has held various board and advisory roles and has served as an independent Board Member at ASELSAN since 2017, including a term as Chairman of the Board in 2017–2018.