Rapid Aerodynamic Optimization Using Digital DATCOM and SLSQP

Emre CAMUSOĞLU[1] and Samet Çaka ÇAKMAKÇIOĞLU[2]
Turkish Aerospace Industries, Ankara, Turkey

H.Burak KURT[3]
Turkish Aerospace Industries, Ankara, Turkey

The present study introduces a rapid aerodynamic optimization framework developed for the preliminary phase of aircraft conceptual design. By coupling Digital DATCOM, a semi-empirical aerodynamics prediction tool, with a Sequential Least Squares Quadratic Programming optimization method, a practical and computationally efficient framework is established to refine aircraft geometry. The methodology is intentionally lightweight, leveraging Digital DATCOM to enable rapid design iterations prior to high-fidelity Computational Fluid Dynamics analyses. The optimization prioritizes a balance between aerodynamic efficiency, static stability, and geometric feasibility constraints, providing a robust framework for preliminary design under limited time and computational budgets. Furthermore, the framework is integrated into a graphical user interface that streamlines input processing, optimization execution, and result visualization. This architecture is highly reconfigurable for different aircraft types and extendable to various mission-profile optimization tasks.

I. Nomenclature

C_L = Lift coefficient
C_D = Drag coefficient
C_m = Pitching moment coefficient
M = Pitching moment
S_{ref} = Reference wing area specified in the original DATCOM input file
S_{geom} = Wing area computed from the geometric definition in the DATCOM output
ρ = Air density
c = Chord
$\bar{c}$ = Mean aerodynamic chord
V = Velocity
α = Angle of attack
W = Aircraft weight
ZV = Z-Coordinate of vertical tail root
C_L^{Req} = Lift coefficient required to balance the weight
VT = Vertical tail
HT = Horizontal tail
CFD = Computational fluid dynamics
UAV = Unmanned aerial vehicle

[1] Flight Control Algorithm Design Engineer
[2] Flight Mechanics and Performance Chief Engineer
[3] Flight Control Algorithm Design Engineer

II. Introduction

In the preliminary phases of aircraft development, rapid assessment of various configurations is crucial for guiding conceptual design decisions. High-fidelity methods such as Computational Fluid Dynamics (CFD) and wind tunnel experiments, while indispensable for later stages, impose prohibitive computational and logistical costs when applied during early-stage design. Instead, simplified models like Digital DATCOM, based on semi-empirical correlations, offer a trade-off by providing reasonable accuracy with significantly reduced computational effort. This study aims to exploit the speed of Digital DATCOM and the rigor of nonlinear constrained optimization to deliver a fast, autonomous, and effective design improvement tool.

Aircraft conceptual design often involves navigating trade-offs between competing performance metrics, stability and control, and geometric constraints. Manually iterating design parameters to identify an optimal configuration is not only time consuming but increasingly labor intensive as complexity grows. As the number of design parameters increases, this approach becomes computationally intractable. Therefore, an optimization driven approach, coupled with intelligent selection of design variables and constraints, is essential for exploring the design space systematically.

In the conceptual phase of aircraft design, the high computational cost of high-fidelity simulations has necessitated the development of fast and efficient optimization tools. Accordingly, several studies have proposed integrating semi-empirical models with modern optimization algorithms to achieve a balance between accuracy and computational efficiency [1] [2].

In study [1], the Sequential Least Squares Quadratic Programming (SLSQP) algorithm is used to optimize the geometric parameters of a two-dimensional airfoil, validating the results using both Digital DATCOM and CFD simulations. However, the scope of that study was limited to isolated airfoil profiles and a limited set of design parameters (such as angle of attack and thickness). Conversely, the present study expands the design space to cover full aircraft geometry and employs 30 physically meaningful design parameters (see Appendix A) in a comprehensive composite objective (weighted sum) optimization process. Furthermore, in addition to improving aerodynamic coefficients, the proposed methodology incorporates static stability constraints. While structural efficiency is not explicitly modeled, the cost function can be adapted to include additional metrics based on the aircraft type, offering a flexible and extensible design approach.

In study [2], a fast optimization framework was proposed specifically for flying wing configurations using a simplified aerodynamic model. While the method enables rapid iterations and emphasizes flexibility, it remains restricted to unconventional platforms and does not incorporate stability metrics or geometric constraints. In contrast, the proposed framework is applicable to both conventional and unconventional aircraft configurations and directly enforces geometric limitations and static stability requirements within the optimization loop. Although explicit structural analysis is not performed, constraints on wing area are included to indirectly address size and structural feasibility. As a result, the proposed framework distinguishes itself not only in terms of computational efficiency but also in its alignment with practical preliminary design considerations.

In conclusion, while existing studies demonstrate the effectiveness and speed of optimization algorithms in aerodynamic design, this work presents a comprehensive framework for full-aircraft geometry optimization based on Digital DATCOM and a SLSQP algorithm. A multi-start optimization strategy is employed to enhance robustness against local optima and ensure convergence toward a global solution. The methodology is implemented within a graphical user interface (GUI), enabling automated input handling, optimization execution, and result visualization. The proposed tool is versatile enough to support both civilian transport aircraft, where aerodynamic efficiency is prioritized, and high-performance aircraft, where maneuverability and stability considerations are dominant. Consequently, this study presents an adaptable optimization framework that leverages the efficiency of Digital DATCOM and the rigor of the SLSQP algorithm to facilitate rapid aerodynamic refinement across diverse aircraft architectures, effectively streamlining the transition from initial conceptualization to a flight-stable configuration.

III. Problem Description

The aerodynamic optimization problem addressed in this study aims to balance conflicting design objectives during the preliminary phase of aircraft design. Specifically, the optimization simultaneously seeks to maximize aerodynamic efficiency (e.g., the lift-to-drag ratio), maintain an adequate static margin, and constrain the geometric expansion of the wing and tail surfaces, effectively enforcing realistic wing loading $\left(W/S_{geom}\right)$ limits to indirectly account for structural weight and feasibility. These objectives inherently conflict with one another; for instance, increasing lift often leads to a rise in induced drag, while minimizing drag can negatively impact longitudinal stability.

This problem is fundamentally complex due to the nonlinear and interdependent nature of the design variables. Wing sweep affects the critical Mach number and longitudinal stability, while span influences both induced drag and lift distribution. Furthermore, dihedral angle contributes to lateral stability, and twist distribution can delay stall and

balance structural loading. Each of these variables affects multiple aspects of aerodynamic and flight performance simultaneously, creating a highly coupled design space.

Moreover, the feasible design space is strictly bounded by fundamental physical constraints, specifically concerning aerodynamic performance, structural integrity, and geometric feasibility. Stability and control criteria impose rigid geometric boundaries to guarantee safe flight dynamics, while parallel sizing constraints ensure the airframe remains structurally robust. Given the moderately high dimensionality of the problem and the interdependent nature of the design variables, manual parameter tuning is impractical. Especially in the early conceptual design phase, when rapid evaluation of the design space is essential, a systematic and constraint-aware optimization approach becomes critical. The objective is not only to improve performance metrics but also to identify viable configurations that rigorously satisfy aerodynamic, stability, and geometric criteria.

Consequently, because different aircraft types are governed by distinct mission profiles, an effective optimization framework must possess the flexibility to integrate variable design objectives and constraint boundaries. Although the proposed framework can be adapted to a broad range of configurations, from transport aircraft to high-performance aircraft, the semi-empirical limitations and validity boundaries of Digital DATCOM must be carefully considered when extending the analysis to transonic or supersonic flight regimes.

IV. Optimization Strategy and Methodology

The optimization objective is defined according to aerodynamic performance, stability and control, and geometrical metrics. When an aerodynamic performance metric is targeted, such as maximizing lift or the lift-to-drag ratio, the objective is formulated as a negative expression such as $-C_L$ or $(-L/D)$ to ensure compatibility with standard minimization solvers. To solve this optimization problem, the SLSQP algorithm is employed. SLSQP reformulates the nonlinear optimization problem into a sequence of linearly constrained quadratic subproblems, solving each to obtain a descent direction before updating the solution iteratively. A key motivation for selecting SLSQP lies in the numerical characteristics of Digital DATCOM. Because DATCOM evaluates aerodynamic coefficients through semi-empirical formulas and discrete lookup tables, gradient estimation via finite differences can yield noisy or locally discontinuous sensitivities at interpolation boundaries. Such gradient noise poses a significant challenge for many gradient-based solvers. SLSQP handles this effectively through its line search strategy and approximate Hessian updates, offering improved robustness over standard SQP implementations for this class of problem. Alternative optimization methods, such as genetic algorithms or derivative-free approaches, were excluded owing to their excessive computational burden. To ensure controlled but flexible exploration of the design space, the GUI initially assigns default lower and upper bounds to each design variable based on a ±5% variation around the baseline configuration. However, these bounds are fully user-configurable, allowing the optimization range to be adjusted according to the aircraft type, design objective, and engineering judgment. [3] [4]

Because SLSQP is a gradient-based solver, it is susceptible to local minima. To improve robustness and reduce sensitivity to initial conditions, a multi-start strategy is employed. The optimization is initialized from multiple starting points distributed across the design space, increasing the probability of converging to a global, or near-global, optimum while maintaining computational efficiency.

Digital DATCOM automatically computes a reference wing area to nondimensionalize aerodynamic parameters when the user does not explicitly specify one. If the user provides a fixed reference wing area as an input, Digital DATCOM uses that value for nondimensionalization instead [5].This behavior introduces a significant challenge when Digital DATCOM is integrated into an optimization loop. During the optimization process, Digital DATCOM is executed iteratively; however, it computes aerodynamic parameters based on the fixed reference wing area supplied at the outset, while the optimization algorithm simultaneously modifies the wing's geometric parameters. As a result, changes to the wing geometry are not reflected in the reference area used for nondimensionalization, causing the optimization procedure to produce incorrect results.

To resolve this inconsistency, a normalization procedure is introduced. Rather than relying on the user-defined fixed reference area (S_{ref}), the geometric wing area (S_{geom}) extracted directly from the Digital DATCOM output at each iteration is used to normalize the aerodynamic parameters. This ensures that the aerodynamic outputs remain consistent with the evolving wing geometry throughout the optimization process. The corrected aerodynamic coefficients are redefined as:

$$C_L^{Corr} = C_L^{DATCOM} \frac{S_{ref}}{S_{geom}}, \quad C_D^{Corr} = C_D^{DATCOM} \frac{S_{ref}}{S_{geom}}, \quad C_m^{Corr} = C_m^{DATCOM} \frac{S_{ref}}{S_{geom}} \tag{1}$$

This transformation ensures that all aerodynamic coefficients are evaluated consistently with respect to the evolving wing geometry. In parallel, the required lift coefficient is recomputed at each iteration using the updated geometric area:

$$C_L^{Req,Corr} = C_L^{Req,DATCOM} \frac{S_{ref}}{S_{geom}} = \frac{W}{\frac{1}{2}\rho V^2 S_{ref}} * \frac{S_{ref}}{S_{geom}} \tag{2}$$

where W is the aircraft weight, ρ is the air density, and V is the flight velocity. This formulation computes the lift coefficient required to support the aircraft weight at the given flight condition, ensuring that the lift constraint remains physically coupled with the evolving wing geometry throughout the optimization. The aircraft weight is held constant across all optimization iterations.

Finally, it is essential to recognize that the standard nondimensionalization of the pitching moment coefficient (C_m) which is given in Eq. (3), inherently depends on both the reference wing area and the mean aerodynamic chord (MAC). When the user explicitly provides a fixed MAC value as input, Digital DATCOM adopts that value for all pitching moment calculations. If no MAC value is provided, Digital DATCOM automatically computes the theoretical MAC from the wing planform geometry at each run. This distinction has critical implications within an optimization context. When a fixed MAC is user-specified, it remains constant throughout the entire optimization regardless of any changes to the wing geometry imposed by the optimizer. Consequently, as the wing geometry evolves across iterations, the pitching moment coefficients continue to be nondimensionalized using the original fixed reference chord rather than the MAC corresponding to the current wing planform. This mismatch between the evolving wing geometry and the frozen reference chord renders the resulting pitching moment values dimensionally inconsistent and unreliable for pitching moment calculation. To prevent this, the optimization framework actively checks whether the user has specified a fixed MAC value and, if so, raises an error and terminates the optimization. This safeguard ensures that the MAC is always derived from the current wing geometry at each iteration, maintaining physical consistency between the aerodynamic outputs and the evolving wing geometry throughout the optimization process.

$$C_m = \frac{M}{\frac{1}{2}\rho V^2 S_{ref} \bar{c}} \tag{3}$$

V. Optimization Framework

The proposed optimization framework is structured as an automated, modular pipeline that integrates design parameter definition, aerodynamic analysis, and optimization within a unified environment. The framework is designed to facilitate the rapid evaluation of aircraft configurations while maintaining physical consistency and numerical robustness.

The architecture comprises three primary phases: initialization, aerodynamic analysis and design parameter optimization. During initialization phase, design variables and constraints are defined, in subsequent aerodynamic analysis phase, Digital DATCOM is utilized to predict static stability and control derivatives. Finally, the optimization phase iteratively updates the design using a SLSQP algorithm.

These components are tightly coupled through an automated workflow, ensuring seamless data exchange between geometry generation, aerodynamic analysis, and cost function evaluation. To enhance usability and minimize manual intervention, the entire framework is implemented within a GUI. This integration allows users to define optimization settings, manage variables, execute tasks, and visualize results within a single environment, making the framework particularly suitable for rapid conceptual design studies. The workflow layout is illustrated in Fig. 1.

A. Initialization Phase

In the initialization phase, the selection of design parameters is based on their aerodynamic significance and relevance to early-stage aircraft design. Thirty design variables are defined, including sweep angles, spanwise dimensions, dihedral angles, twist distributions and chord lengths. Variables are grouped by aerodynamic surfaces (wing, horizontal tail, and vertical tail) and are bounded within physically realistic limits to ensure feasibility and manufacturability. The full list of parameters is provided in Appendix A.

Design variable bounds are defined relative to the baseline configuration, allowing the optimization to explore variations around a physically meaningful starting point. In the implemented framework, these bounds are automatically generated from the initial geometry within a constrained variation range to maintain realistic

configurations. This approach ensures that the optimization remains within a feasible design space while allowing sufficient flexibility for performance improvement.

Unlike traditional script-based approaches, the parameter definition and constraint handling are integrated into the GUI, where users can select active design variables and specify corresponding bounds interactively. This structure enables flexible configuration of the optimization problem without requiring manual modification of input files or source code.

The initialization phase outputs the initial design vector and corresponding constraint bounds, which are passed to the optimization module. This structured setup ensures that all subsequent optimization steps operate within a consistent and physically valid design space.

B. Aerodynamic Analysis Phase

Digital DATCOM is a semi-empirical aerodynamic analysis tool developed by the United States Air Force to estimate stability and control characteristics based on regression models derived from wind tunnel and flight test data. Although it relies on simplifying assumptions such as attached flow and linear aerodynamic behavior, it provides rapid and sufficiently accurate predictions, making it highly suitable for conceptual and preliminary aircraft design applications [5].

Within the proposed optimization framework, Digital DATCOM functions as the primary aerodynamic analysis module in a fully automated loop. For each iteration, the design variables generated by the optimizer are written into a temporary DATCOM input file, preserving the original baseline configuration. This input file defines the updated aircraft geometry, including parameters such as sweep angles, span, dihedral angles, twist distributions, chord lengths, and tail surface positions. The DATCOM executable is then invoked, producing an output file containing the corresponding aerodynamic coefficients.

The resulting output provides key aerodynamic quantities such as the lift coefficient (C_L), drag coefficient (C_D), and pitching moment coefficient (C_m), evaluated under the specified flight conditions. These coefficients are automatically extracted and passed back to the optimization module for cost function evaluation.

The aerodynamic analysis phase includes the normalization of aerodynamic coefficients within Digital DATCOM. If reference wing area (S_{ref}) is specified in the input file, all aerodynamic coefficients are nondimensionalized using a reference wing area (S_{ref}), which remains fixed throughout the optimization. However, because the actual aircraft geometry is modified at each iteration, particularly the wing geometry parameters, Digital DATCOM computes a corresponding geometric wing area (S_{geom}) based on the updated configuration. This creates a discrepancy between the fixed reference area used for nondimensionalization and the actual geometric area of the wing.

To ensure consistency in aerodynamic evaluation, the optimization framework incorporates a correction mechanism that rescales the aerodynamic coefficients with respect to the geometric wing area. Specifically, the coefficients obtained from DATCOM (C_L, C_D, C_m), are systematically normalized using S_{geom} rather than S_{ref}. This adjustment ensures that comparisons between different configurations remain physically meaningful despite geometric changes, as all coefficients are consistently referenced to the actual wing area of each configuration.

The detailed mathematical formulation of this normalization procedure is presented in the methodology section. Following this correction, the adjusted aerodynamic data is used in cost function evaluation and constraint satisfaction. This closed-loop interaction (geometry update, aerodynamic analysis, and optimization) enables iterative convergence toward improved aircraft configurations while maintaining physical and numerical consistency.

C. Optimization Phase

The main optimization loop employs the SLSQP method, which is well suited for constrained nonlinear optimization problems. Each iteration of the loop follows these steps:

1) A subset of design parameters is selected by the user through the GUI, where active variables and their corresponding bounds are defined.
2) Based on the selected active parameters, a candidate design vector is generated within the prescribed bounds.
3) These parameters are written to a temporary DATCOM input file representing the updated aircraft geometry.
4) Digital DATCOM is executed to produce aerodynamic outputs (C_L, C_D, C_m).
5) The obtained aerodynamic coefficients are subsequently corrected to ensure consistency with the geometric wing area. If reference wing area is defined by user.
6) A user-defined cost function is then evaluated using the corrected aerodynamic coefficients and any additional terms specified by the user. In this way, the optimization objective is not fixed a priori, but can be adapted to different aircraft categories, mission priorities, and design preferences.

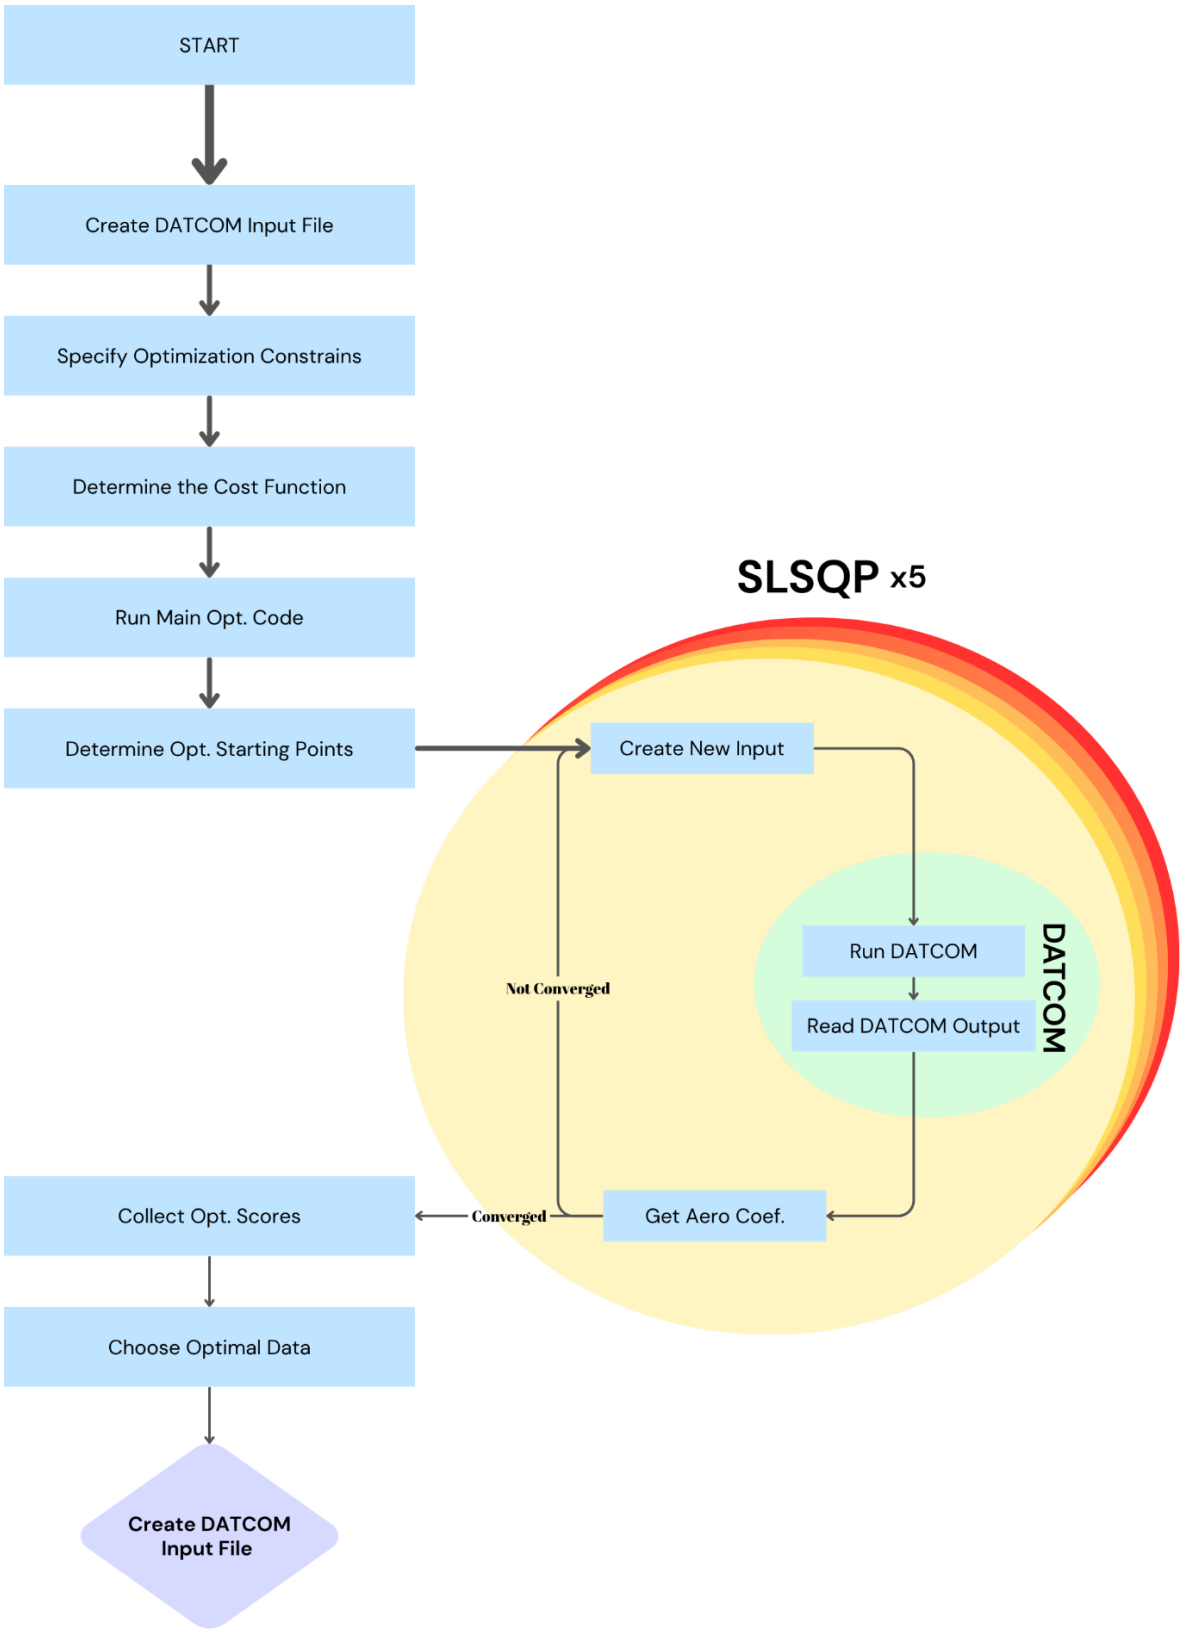

Fig. 1 Optimization framework schematic.

The optimization framework evaluates aerodynamic coefficients at a single user-specified constant angle of attack. Digital DATCOM computes aerodynamic coefficients across the full angle of attack sweep defined in the input file, and the framework extracts the coefficient value at the prescribed operating angle of attack for use in the cost function and constraint evaluations. Upon completion of the optimization, the aerodynamic coefficients across the full angle of attack sweep are recomputed for the optimized design parameters and presented to the user for evaluation.

The cost function used in the optimization process is defined through the GUI, allowing both flexibility and ease of use in shaping the optimization objective. The framework supports two modes of operation: users may either construct fully custom cost expressions manually or select from a set of predefined objective formulations provided within the interface. These predefined options are designed to represent common aerodynamic design goals, such as lift-to-drag ratio maximization, drag minimization, or stability-oriented optimization, enabling rapid setup without requiring detailed formulation. At the same time, advanced users retain full control by defining custom cost functions based on aerodynamic coefficients, stability characteristics, geometric properties, or any combination thereof.

This framework allows the optimization to be tailored to different aircraft types and mission requirements. For example, a user may prioritize lift-to-drag ratio for cruise efficiency, enforce stronger penalties on pitching moment behavior for enhanced stability, or emphasize lift generation for low-speed performance. Similarly, geometric terms such as wing area or span can be incorporated to indirectly account for structural or manufacturability considerations.

D. Multi-Start Strategy for Robustness

To address the initial condition dependency of gradient based optimization and improve convergence robustness, the present study employs a multi-start optimization strategy integrated within the GUI. Specifically, the allowable range (bounds) for each user-selected design variable is partitioned into multiple intervals of equal length. From each interval, a distinct initial design vector is constructed, ensuring that starting points span different regions of the feasible design space defined through the GUI.

The SLSQP optimization algorithm is then executed independently for each of these starting points. The optimization framework starts with five initial design points as a default. The choice of five initial conditions represents a practical balance between exploration capability and computational efficiency, enabling improved robustness against local minima without significantly increasing computational cost.

This method serves two primary objectives:

1) *Exploration of the design space:* By starting from diverse initial conditions, the optimizer is exposed to different regions of the design space. This increases the chance of discovering more promising basins of attraction and avoids over-committing to local minima that may be suboptimal on a global scale.
2) *Mitigation of sensitivity to initialization:* Variability in initial guesses often leads to widely different convergence outcomes. A single, arbitrarily chosen starting point might miss the optimal region altogether. The multi-start approach compensates for this risk by offering a form of ensemble exploration.

Upon completing all five optimization runs, each yielding a final design and corresponding cost value, the framework selects the solution with the lowest final cost as the global candidate. This not only enhances solution quality but also adds a layer of confidence in the robustness of the result, especially in the context of early-stage design where uncertainty is high and decision timelines are compressed.

It is worth emphasizing that this approach provides a pragmatic compromise between global and local optimization techniques. While true global methods such as genetic algorithms or particle swarm optimization are capable of comprehensively exploring the design space and identifying true global optima, they typically involve substantial computational expense and longer convergence times. In contrast, the multi-start SLSQP strategy retains the computational efficiency of gradient-based local search methods, while significantly mitigating their dependence on initial conditions and susceptibility to local minima. This makes the approach particularly well-suited for rapid, automated design evaluations in early-stage conceptual workflows, especially when paired with fast aerodynamic prediction tools like Digital DATCOM. [6]

VI. Case Studies: Application to Representative Aircraft Platforms

A. F-16 Fighting Falcon Case Study

To demonstrate the practical capability of the proposed optimization framework, the F-16 configuration [7] was selected as a well-documented benchmark platform and subjected to the evaluated through the complete optimization pipeline. This case is intended as a methodology demonstration rather than a mission-specific design study: the objective is to show how the framework navigates aerodynamic trade-offs under a user-defined cost function, and how resulting geometric changes manifest as coherent improvements in lift, drag, stability, and trim characteristics. The F-16 geometry was chosen for its aerodynamic complexity and its established role in the literature as a reference

configuration, making it a suitable platform for exercising the full range of design variables supported by the framework.

The optimization was performed at Mach 0.6 and 10,000 ft altitude with 10° angle of attack, representing a subsonic condition suitable for demonstrating the framework within the practical validity range of Digital DATCOM. This flight condition was selected to exercise the optimizer under a moderate-speed, moderate-altitude scenario in which lift generation, drag growth, and stability-related aerodynamic trends are all relevant. The purpose of this case study is not to propose a mission-specific redesign of the F-16, but rather to demonstrate how the framework responds to a user-defined engineering objective and how the resulting geometric modifications affect the aerodynamic behavior over a wider angle-of-attack range. For this demonstration, the GUI's default design-variable bounds were retained, corresponding to a ±5% variation around the baseline values. An engineering-judgment-based composite cost function was then defined as follows:

$$f_{cost} = -C_L + 2C_D + C_L^{req} + \frac{S_{geom}}{200} \tag{4}$$

In this formulation, the negative lift-coefficient term promotes lift generation at the selected design condition, while the drag term penalizes excessive aerodynamic drag. The required lift coefficient term penalizes configurations that exhibit increased lift demand at the prescribed operating angle of attack. The geometric wing area term discourages unnecessary growth of the lifting surface, preventing the optimizer from reducing C_L^{req} simply by excessively increasing wing area. For this case study, all design variable bounds were kept at the default ±5% range around their baseline values. The framework converged in approximately 5-10 minutes across all five multi-start runs.

The comparison between the baseline and optimized configurations in terms of lift and drag characteristics is presented in Fig. 2. The optimized configuration exhibits higher C_L values over the evaluated angle-of-attack range, particularly in the positive angle of attack region. This behavior is consistent with the $-C_L$ term in the selected cost function, which encourages increased lift generation at the 10° angle of attack. The improvement becomes more pronounced at higher angles of attack, indicating that the optimized geometry provides stronger lift capability at the high angle of attack region.

The drag coefficient (C_D) also increases slightly for the optimized configuration across most of the evaluated range. This increase is expected because the cost function prioritizes lift generation while applying only a moderate drag penalty through the $2C_D$ term. However, the drag rise remains limited and follows the same smooth trend as the baseline configuration, suggesting that the optimizer did not produce an aerodynamically unrealistic or numerically degenerate solution. The result therefore reflects a controlled lift and drag trade-off consistent with the selected engineering objective.

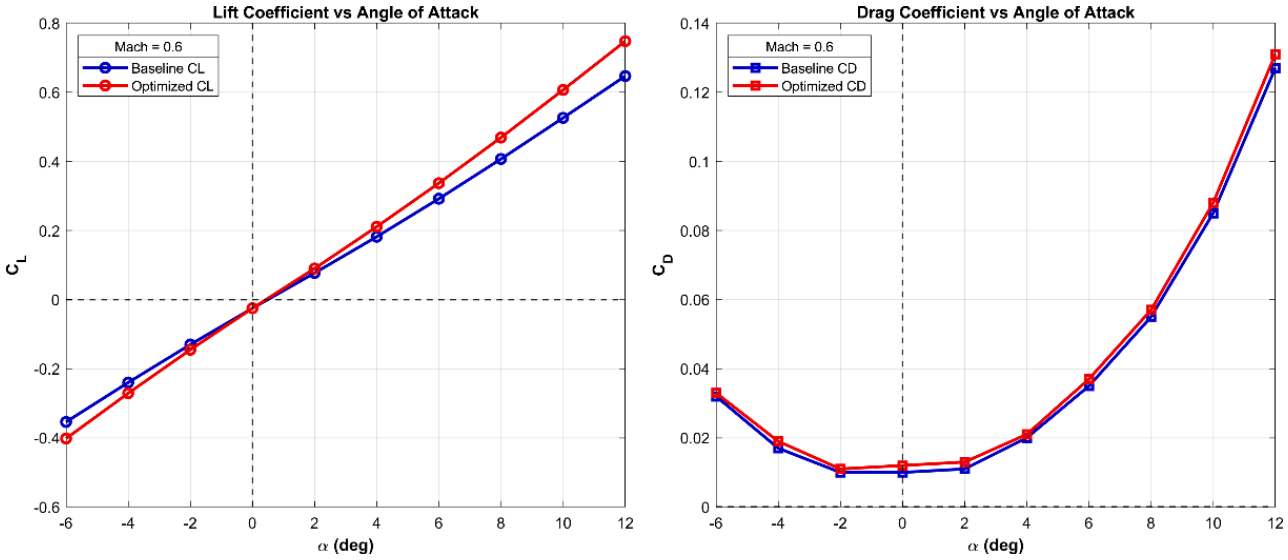

Fig. 2 Comparison of baseline and optimized aerodynamic performance (F-16).

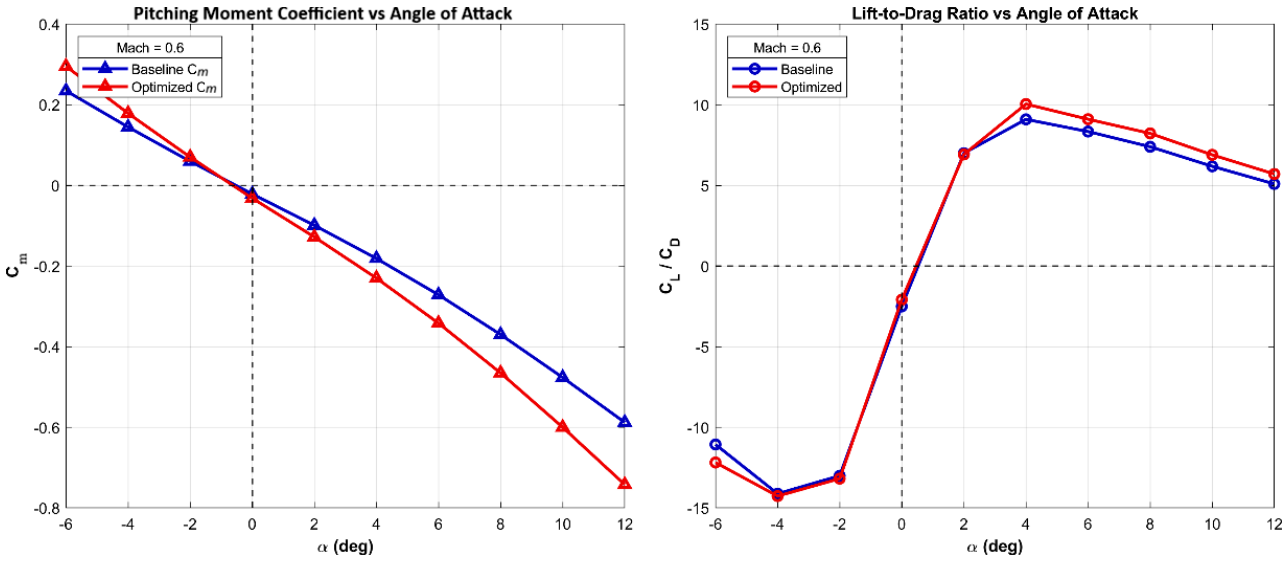

Fig. 3 Comparison of baseline and optimized aerodynamic stability and efficiency characteristics (F-16).

Fig. 3 presents the pitching-moment coefficient and lift-to-drag ratio trends for the baseline and optimized configurations. The lift-to-drag ratio improves over a broad portion of the positive angle-of-attack range, despite the slight increase in drag. This improvement occurs because the increase in lift outweighs the corresponding increase in drag. Although the cost function does not directly maximize the lift-to-drag ratio(C_L/C_D), the observed improvement emerges as a coupled outcome of the lift promoting and drag penalty terms.

The pitching moment coefficient (C_m) curve shows that the optimized configuration has a more negative C_m trend at positive angles of attack. Since C_m was not explicitly included in the selected cost function, this behavior should not be interpreted as a directly minimized objective. Instead, it represents an aerodynamic consequence of the geometric modifications produced by the optimizer. The C_m curve is therefore used here to assess the pitching moment.

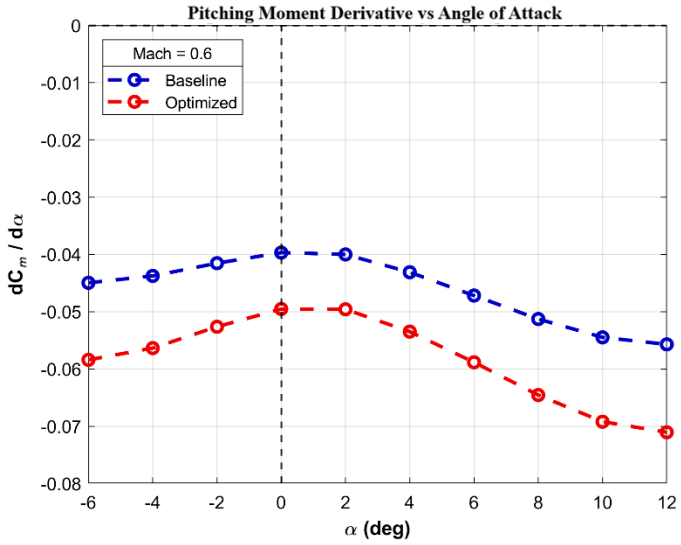

Fig. 4 Pitching moment derivative versus angle of attack for baseline and optimized configurations (F-16).

The variation of the pitching moment derivative with respect to angel of attack ($dC_m/d\alpha$) is shown in Fig. 4 The optimized configuration exhibits a consistently more negative $dC_m/d\alpha$ value than the baseline throughout the evaluated range. The more negative $dC_m/d\alpha$ trend at higher angles of attack suggests that the optimized wing geometry produces a stronger nose down restoring moment response to angle of attack perturbations, which is

generally a desirable stability characteristic. However, it should also be noted that excessively negative $dC_m/d\alpha$ values increase trim drag and control surface trim deflection requirements.

It should be noted that $dC_m/d\alpha$ was not explicitly included in the cost function. The improved slope behavior is therefore a result of the coupled geometric changes rather than a directly imposed optimization objective. The smooth variation of the optimized slope curve further indicates that the optimization did not introduce abrupt nonlinearities or irregular aerodynamic behavior across the evaluated angle of attack range.

Taken together, the F-16 demonstration shows that the proposed framework can generate aerodynamically coherent design changes under a user defined engineering objective. The optimized configuration improves lift generation and lift to drag behavior in the positive angle of attack range while maintaining smooth drag and pitching moment trends. Although the $dC_m/d\alpha$ was not directly included in the cost function, the resulting configuration also exhibits a more static longitudinal stability trend. These results demonstrate that the framework can balance lift enhancement, drag limitation, geometric feasibility, and stability related aerodynamic behavior within a rapid preliminary design workflow.

B. Cessna Citation 500 Case Study

As an additional case study, the Cessna Citation 500 [8] was selected as a representative civil aviation platform for a practical optimization case. The Citation 500 is a subsonic business jet characterized by a conventional low-wing configuration, making it a well-suited candidate for semi-empirical aerodynamic analysis using Digital DATCOM.

The primary objective of this case was to improve the lift-to-drag ratio while simultaneously minimizing the absolute pitching moment coefficient of aircraft configuration at the desired envelope point. To address this objective, the cost function was defined as:

$$f_{cost} = -\frac{C_L}{C_D} + 0.1 * abs(C_m) \tag{5}$$

The first term maximizes the lift-to-drag ratio, driving the optimizer toward aerodynamically efficient configurations at the desired envelope point. The second term imposes a penalty proportional to the absolute magnitude of the pitching moment coefficient, $abs(C_m)$. This penalty discourages designs that would gain efficiency through unbalanced moment characteristics. The combined objective seeks the optimal balance between aerodynamic efficiency and moment equilibrium, preferring configurations where the aircraft achieves high lift-to-drag ratio while yielding a near zero basic pitching moment coefficient ($C_m \approx 0$) at the design point, thereby minimizing the elevator trim requirement and its associated trim drag.

The weighting factor of 0.1 was selected to balance the relative contributions of the two terms. All design variable bounds were set to ±5% of their respective baseline values, consistent with the default framework configuration.

The optimization was performed at Mach 0.6 and 30,000 ft altitude with 2° angle of attack, representing a typical flight condition for this class of aircraft. The framework completed all five multi-start runs and converged in under 10 minutes, demonstrating the computational efficiency of the proposed methodology under realistic operating conditions.

The geometric modifications produced by the optimizer are summarized in Table 1. The design changes remain within the user-defined bounds and are distributed across all three aerodynamic surfaces. As shown in the table, certain design parameters reach their prescribed limits during the optimization. This suggests that further objective function improvements might exist beyond the current constraints; however, as the bounds represent the permissible design space, the resulting cost function is considered the optimum for the defined region. Additionally, the percentage-based bound strategy is not applicable when the initial value of a design parameter is zero, since a percentage variation around zero yields no meaningful search interval. Therefore, an additive bound strategy is employed for such parameters to provide a finite and physically controlled design range.

Notable adjustments include a slight increase in wing span (0.58%), a reduction in wing root chord (−5%), an increase in wing tip chord (5%), a reduction in wing sweep (−5%), and a reduction in wing twist (−5%). These changes collectively improve the spanwise lift distribution and shift the wing's aerodynamic center, which directly aids the optimizer in satisfying the pitching moment objective. The optimizer also introduced an increase in wing inner dihedral (3%). Although this parameter primarily affects lateral-directional stability characteristics, it also contributes marginally to spanwise lift redistribution by shifting loads toward the tip. Therefore, its emergence in the optimized solution likely reflects coupled geometric effects rather than a primary design driver for moment control.

The reduction in horizontal tail sweep (−5%) increases the aerodynamic effectiveness of the tail surface by enhancing normal force production, thereby improving its lift generation capability and moment generating capacity about the center of gravity. While this mechanism contributes to pitching moment balance, it represents a minor effect

compared to primary moment control mechanisms. The introduction of horizontal tail twist (2.0°) modifies the aircraft pitching moment by altering the spanwise lift distribution of the tail, thereby changing the total tail lift and its moment contribution about the center of gravity.

The vertical tail dihedral increased marginally (0.08°) from zero baseline. Given the small magnitude of this change relative to the design resolution, this parameter modification likely represents either a numerical convergence artifact or minimal aerodynamic significance. Its physical justification is not clear in the context of conventional vertical tail design.

The vertical tail twist parameter was optimized to 2.0°, representing a departure from the conventional baseline design (zero twist). Applying geometric twist to a single vertical tail generates asymmetric lateral-directional forces and residual yawing moments, even in symmetric cruise conditions. The emergence of this parameter suggests a secondary aerodynamic coupling effect or a numerical artifact within the optimization framework rather than a primary design driver, though its exact physical role requires further sensitivity analysis.

A significant limitation of the present optimization framework is that the cost function employs only longitudinal aerodynamic coefficients, while lateral-directional characteristics are not explicitly constrained. Consequently, design variables that influence lateral-directional behavior such as vertical tail dihedral and twist may be modified by the optimizer without regard for their lateral-directional implications. To prevent unintended coupling effects, these parameters should either be fixed to baseline values, excluded from the active optimization set, or bounded with conservative constraints when longitudinal performance is the primary objective. Future work should extend the optimization framework to include lateral-directional aerodynamic constraints.

Table 1 Optimized geometric parameter changes for the Cessna Citation 500 configuration.

Surface	Parameter	Baseline	Optimized	Change(%)
Wing	Inner Sweep (deg)	1.300	1.235	−5
Wing	Span (m)	25.850	26.000	0.58
Wing	Inner Dihedral (deg)	3.600	3.708	3
Wing	Twist (deg)	−3.000	−2.850	−5
Wing	Root Chord (m)	9.400	8.930	−5
Wing	Tip Chord (m)	3.010	3.161	5.00
Vertical Tail	Inner Sweep (deg)	32.300	33.260	2.97
Vertical Tail	Span (m)	9.420	9.390	−0.31
Vertical Tail	Inner Dihedral (deg)	0.000	0.080	NA
Vertical Tail	Twist (deg)	0.000	2.000	NA
Vertical Tail	Root Chord (m)	8.300	8.549	3
Vertical Tail	Tip Chord (m)	3.630	3.739	3
Horizontal Tail	Inner Sweep (deg)	5.320	5.054	−5
Horizontal Tail	Span (m)	9.420	9.500	0.84
Horizontal Tail	Inner Dihedral (deg)	9.200	9.470	2.93
Horizontal Tail	Twist (deg)	0.000	2.000	NA
Horizontal Tail	Root Chord (m)	4.990	5.240	5
Horizontal Tail	Tip Chord (m)	2.480	2.600	4.83

The aerodynamic results for the baseline and optimized configurations are presented in Fig. 5-7. The optimization of the cost function was evaluated at an angle of attack of 2°. As shown in the Fig. 5, the optimized configuration yields a marginal increase in the lift coefficient at this angle, while the drag coefficient remains virtually unchanged. Consequently, the lift-to-drag ratio is successfully increased at the given angle of attack. Because the primary objective of the cost function is to maximize lift-to-drag ratio while simultaneously minimizing the absolute pitching moment coefficient, the framework achieved the desired balance. As illustrated in Fig. 6, the absolute value of the pitching moment is effectively reduced at the 2° angle of attack.

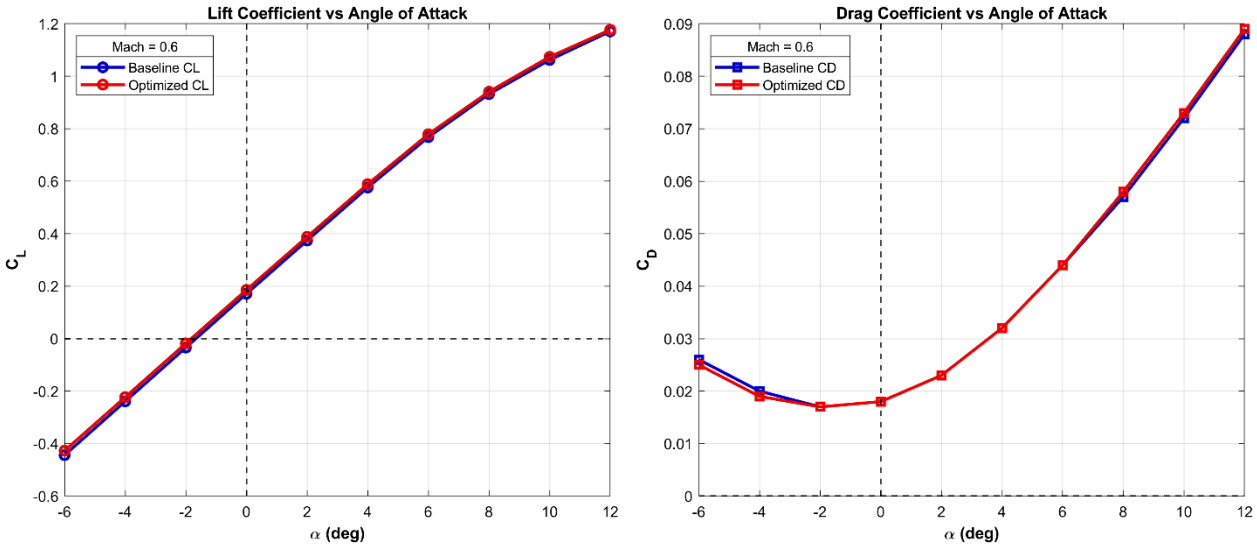

Fig. 5 Comparison of baseline and optimized aerodynamic performance (Citation 500).

An analysis of the aerodynamic trends across the angle of attack sweep reveals that the optimized configuration consistently improves upon the baseline design. The lift coefficient exhibits a standard linear increase, with the optimized geometry maintaining a marginal but consistent lift advantage over the entire positive angle of attack range. Crucially, this improvement is achieved without a structural drag penalty, as the optimized drag coefficient follows a classic parabolic profile that remains virtually indistinguishable from the baseline. Consequently, the aerodynamic efficiency represented by the lift-to-drag ratio sees a net improvement across the broad spectrum of positive angles of attack.

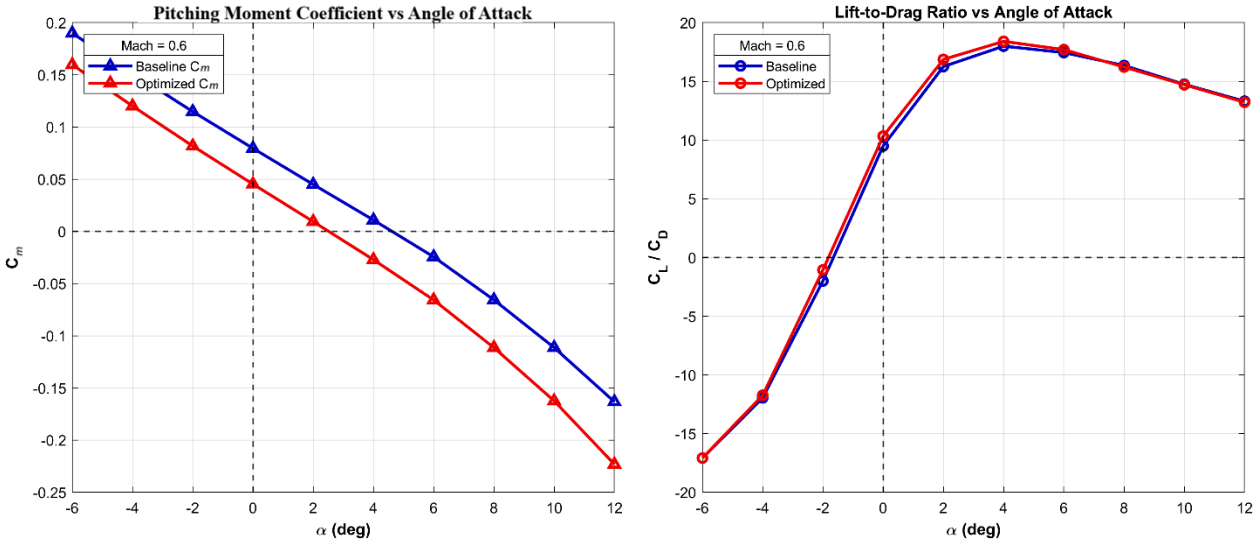

Fig. 6 Comparison of baseline and optimized aerodynamic stability and efficiency characteristics (Citation 500).

The optimization successfully shifted the pitching moment curve to reduce its absolute magnitude near the optimized angle of attack. The optimization process yielded a reduction in the pitching moment coefficient at zero angle of attack from approximately 0.08 to 0.045, which correlates with a downward shift of the entire curve. This shift effectively lowers the natural trim angle of attack identified by the zero-crossing of the pitching moment curve from approximately 4.3° to 2.3° in the optimized configuration.

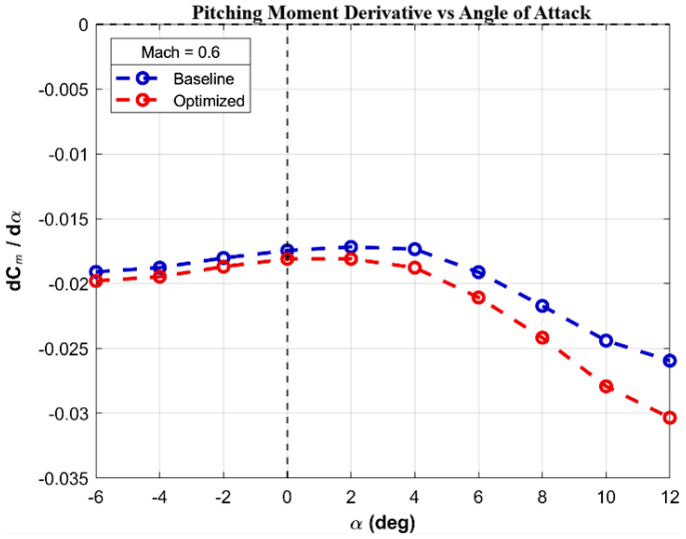

Fig. 7 Pitching moment derivative versus angle of attack for baseline and optimized configurations (Citation 500).

The pitching moment derivative with respect to angle of attack distribution further corroborates these findings. As shown in Fig. 7, the optimized configuration exhibits consistently more negative $dC_m/d\alpha$ values across the full angle of attack range, with the divergence between the baseline and optimized curves becoming more pronounced at higher angles. This behavior reflects increased pitch stiffness and improved disturbance rejection capability, suggesting that the optimized design generates a stronger initial restoring moment to oppose pitch perturbations.

Taken together, the results obtained for the Cessna Citation 500 validate the proposed framework's capability to deliver physically meaningful aerodynamic improvements. The simultaneous enhancement of lift-to-drag efficiency and reducing the absolute value of pitching moment coefficient, achieved within a tightly bounded design space and in under 10 minutes of computation time, demonstrates that the methodology is directly applicable to practical preliminary design workflows and suggests applicability to civil aviation platforms beyond the initial demonstration case.

VII. Results and Discussion

The developed optimization framework demonstrates a reliable, automated, and computationally efficient approach for aircraft aerodynamic sizing and refinement during the early stages of conceptual design. It demonstrates the ability to handle moderately high dimensional design spaces, enforce coupled aerodynamic and stability constraints, and produce physically consistent geometric modifications within sub-10-minute computation times.

A. Optimization Convergence and Computational Performance

The framework consistently achieves convergence within approximately 8 to 10 minutes, even when handling dozens of coupled design variables and performing Digital DATCOM aerodynamic evaluations at each function call. This performance is maintained across all multi-start runs, indicating that the overall optimization pipeline from geometry generation to aerodynamic evaluation and cost function computation is stable and computationally efficient.

By initiating multiple optimization runs from distinct regions of the design space, as defined through the user-specified bounds in the GUI, the framework reduces sensitivity to initial conditions and mitigates the risk of convergence to local minima. Since each multi-start run is initialized independently from a distinct region of the design space, the strategy is inherently amenable to parallel execution. Incorporating parallel processing represents a natural future extension that could significantly reduce total computation time, particularly as the number of design variables or multi-start runs is scaled up.

B. Geometry Modification Logic and Physical Realism

Throughout the optimization process, geometric variables such as wing sweep, span, twist, and chord distribution are constrained within user-defined bounds that reflect aerodynamic feasibility and practical manufacturing

constraints. At each iteration, the framework updates the DATCOM input file and evaluates the resulting configuration under consistent flight conditions. A key aspect of the framework is the consistent treatment of aerodynamic coefficients with respect to the geometric wing area, ensuring that geometric modifications are reflected accurately in the aerodynamic evaluation. This prevents misleading trends that could arise from normalization inconsistencies and supports physically meaningful comparisons between configurations.

C. Cost Function Behavior and Constraint Enforcement

The optimization framework employs a fully user-defined cost function, configured through the GUI, supporting both predefined objective templates and user-defined mathematical expressions. This flexibility enables users to tailor the optimization objective to specific aircraft types, mission profiles, or design priorities.

Depending on the selected formulation, the cost function incorporates aerodynamic performance metrics such as lift and drag, stability characteristics such as pitching moment behavior. The required lift coefficient (C_L^{Req}) can be included in the custom cost function when a lift-demand objective is desired; however, the lift requirement itself is enforced independently through an embedded feasibility penalty.

D. Adaptability and Scalability

The modular structure of the framework enables straightforward adaptation to a broad range of aircraft design applications. The system supports different aircraft categories within the modeling assumptions of Digital DATCOM, including transport aircraft, unmanned aerial vehicles, and high-performance configurations, without requiring modifications to the core optimization logic. Additionally, the cost function formulation can be reformulated to reflect different design priorities, such as efficiency, maneuverability, or stability. The framework can also be extended to include additional aerodynamic constraints, with further extension to structural or control considerations possible through integration with complementary analysis tools, making it applicable to both academic research and preliminary industrial design assessments.

Appendix A: List of Optimization Design Parameters

The optimization design parameters are detailed in Table 2. The geometric definitions and calculation methods for the symbols listed are provided in the Digital DATCOM User Manual [5]

Table 2 List of optimization design parameters.

Parameter	Symbol	Surface	Description	Unit*
Wing Inner Sweep Angle	SAVSI	Wing	Inner wing leading edge sweep	deg
Wing Outer Sweep Angle	SAVSO	Wing	Outer wing leading edge sweep	deg
Wing Span	SSPN	Wing	Wing total span	m/ft
Wing Breakpoint Span	SSPNOP	Wing	Span at wing breakpoint	m/ft
Wing Inner Dihedral	DHDADI	Wing	Inner wing dihedral angle	deg
Wing Outer Dihedral	DHDADO	Wing	Outer wing dihedral angle	deg
Wing Twist	TWISTA	Wing	Linear twist along wing span	deg
Wing Root Chord	CHRDR	Wing	Chord at wing root	m/ft
Wing Breakpoint Chord	CHRDBP	Wing	Chord at wing breakpoint	m/ft
Wing Tip Chord	CHRDTP	Wing	Chord at wing tip	m/ft
VT Inner Sweep Angle	SAVSI	Vertical Tail	Inner vertical tail sweep	deg
VT Outer Sweep Angle	SAVSO	Vertical Tail	Outer vertical tail sweep	deg
VT Span	SSPN	Vertical Tail	Vertical tail span	m/ft
VT Breakpoint Span	SSPNOP	Vertical Tail	Span at VT breakpoint	m/ft
VT Inner Dihedral	DHDADI	Vertical Tail	VT inner dihedral	deg
VT Outer Dihedral	DHDADO	Vertical Tail	VT outer dihedral	deg
VT Twist	TWISTA	Vertical Tail	VT twist angle	deg
VT Root Chord	CHRDR	Vertical Tail	Chord at VT root	m/ft
VT Breakpoint Chord	CHRDBP	Vertical Tail	Chord at VT breakpoint	m/ft
VT Tip Chord	CHRDTP	Vertical Tail	Chord at VT tip	m/ft
HT Inner Sweep Angle	SAVSI	Horizontal Tail	Inner HT sweep angle	deg
HT Outer Sweep Angle	SAVSO	Horizontal Tail	Outer HT sweep angle	deg
HT Span	SSPN	Horizontal Tail	HT total span	m/ft
HT Breakpoint Span	SSPNOP	Horizontal Tail	Span at HT breakpoint	m/ft
HT Inner Dihedral	DHDADI	Horizontal Tail	HT inner dihedral angle	deg
HT Outer Dihedral	DHDADO	Horizontal Tail	HT outer dihedral angle	deg
HT Twist	TWISTA	Horizontal Tail	HT twist distribution	deg
HT Root Chord	CHRDR	Horizontal Tail	Chord at HT root	m/ft
HT Breakpoint Chord	CHRDBP	Horizontal Tail	Chord at HT breakpoint	m/ft
HT Tip Chord	CHRDTP	Horizontal Tail	Chord at HT tip	m/ft

Note: All dimensional parameters follow the unit system defined in the Digital DATCOM input (DIM M or DIM FT)

Acknowledgments

The authors thank Turkish Aerospace Industries for supporting this research.

References

[1] M. R. I. H. R. K. P. Srinath R., "Aerofoil optimization using SLSQP and validation using numerical and analytical methods," Vietnam Journal of Science and Technology, 2024.

[2] J. H. S. L. Wei Chenhao, "Study on a Rapid Aerodynamic Optimization Method of Flying Wing Aircraft for Conceptual Design," International Journal of Aerospace Engineeringi WILEY, 2022.

[3] D. Kraft, "A software package for sequential quadratic programming," Deutsches Zentrum für Luft- und Raumfahrt, Institut für Dynamik der Flugsysteme, Cologne, 1988.

[4] P. Virtanen, R. Gommers and T. E. e. a. Oliphant, "SciPy 1.0: Fundamental algorithms for scientific computing in Python," Nature Methods, 2020.

[5] M. D. A. C. –. S. L. Division, "The USAF Stability and Control Digital DATCOM: Volume I, User's Manual," McDonnell Douglas Astronautics Company, St. Louis, MO, 1979.

[6] K. Deb, A. Pratap, S. Agarwal and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," IEEE Transactions on Evolutionary Computation, 2002.

[7] M. C. Fox and D. K. Forrest, "Supersonic aerodynamic characteristics of an advanced F-16 derivative aircraft configuration," NASA Dryden Flight Research Center, 1994.

[8] J. Goppert. [Online]. Available: https://github.com/arktools/pydatcom/blob/master/test/data/Citation.out.