

ROBUST HIGH-FIDELITY DATASET-BASED SHAPE OPTIMIZATION FOR AIRCRAFT DESIGN

M. Maier^{*†}, K. Sørensen-Libik[†], C. Breitsamter^{*}

^{*} Technical University of Munich, Chair of Aerodynamics and Fluid Mechanics,
Boltzmannstr. 15, 85748 Garching, Germany

[†] Airbus Defence and Space GmbH, Rechliner Str., 85077 Manching, Germany

Abstract

Reliable and robust aircraft design requires accurate understanding of aerodynamic behavior across the full flight envelope to predict mission performance, as well as stability and control characteristics. This work demonstrates the potential of automated, high-fidelity dataset-based shape optimization of aircraft for mission-level performance evaluation, using a simple yet realistic test case: the Optimization Test Interceptor with Fan (OTIFAN) configuration. To ensure a viable aircraft design, lower-fidelity methods are incorporated for the disciplines of mass properties, flight mechanics, and structural analysis. Three optimization strategies have been implemented into an automated framework: Grid search (GS) for framework setup and validation, gradient-based optimization (GO) for efficient local optimization and Bayesian optimization (BO) for global, gradient-free optimization. The strategies are applied across two objective functions, illustrating the applicability of the framework on geometric and mission profile optimizations.

Keywords

High-Fidelity Optimization; Dataset-Based Optimization; Aerodynamic Shape Design; Aircraft Design

NOMENCLATURE

Symbols

α	Angle of attack	[°]
b	Span	[m]
cr	Constraint rating	[%]
D	Diameter	[-]
δ	Control deflection angle	[°]
f	Factor	[-]
fr	FOM rating	[pts]
H	Altitude	[m]
J	Impulse	[N·s]
m	Mass	[kg]
Ma	Mach number	[-]
mr	Mission rating	[pts]
mw	Mission weight	[-]
n	Load factor	[g]
obj	Objective	[pts]
ROC	Rate of climb	[m/s]
s	Sensitivity	[pts/m]
SEP	Specific excess power	[m/s]
t	Time	[s]
u	Utility	[-]
V	Velocity	[m/s]

w	Weight	[-]
$\vec{x}$	Design parameter vector	[m]

Subscripts

0	Initial value
$climb$	Climb condition
C	Control
eps	Perturbed value
F	Fuselage
inc	Increment
inv	Inverted
$mean$	Mean
pu	Pull-up condition
S	Stall
tgt	Target
W	Wing
xb	X-direction in body-fixed coordinates
xg	X-direction in earth-fixed coordinates
zb	Z-direction in body-fixed coordinates
zg	Z-direction in earth-fixed coordinates

Acronyms

BO	Bayesian optimization
CAD	Computer-aided design
CFD	Computational fluid dynamics
CPU	Central processing unit
EDF	Electric ducted fan
EI	Expected improvement
FOM	Figure of merit
GO	Gradient-based optimization
GS	Grid search
HPC	High-performance computing
LCB	Lower confidence bound
MPI	Maximum probability of improvement
OTIFAN	Optimization Test Interceptor with Fan
PLA	Polylactic acid, plastic material
RANS	Reynolds-averaged Navier-Stokes
UAV	Unmanned aerial vehicle

1. INTRODUCTION

Aircraft design is an inherently multidisciplinary process, requiring accurate treatment of all relevant disciplines to achieve viable configurations [1]. This work focuses on high-fidelity treatment of aerodynamics while introducing low-fidelity treatment of mass properties, flight performance, flight mechanics and structural analysis.

In addition to experimental methods for assessing aerodynamic characteristics, various computational fluid dynamic (CFD) methods have emerged, with a wide span of fidelity and associated computational cost. Reynolds-averaged Navier-Stokes (RANS) CFD simulations are currently considered the state of the art high-fidelity method for aerodynamic problems. The application to complex geometries such as full aircraft configurations often provides sufficiently accurate data even in flow fields featuring non-linear phenomena such as separated flow [2, 3].

State-of-the-art computational hardware enables the execution of quasi-steady RANS simulations for complete aircraft configurations within feasible industrial timeframes. Such multi-point analyses typically involve selecting a set of flow conditions — such as Mach number Ma , angle of attack α , and control settings including control surface deflections and engine parameters — thus covering only a subset of the flight envelope. However, a priori assumptions about the flight conditions of interest can introduce significant uncertainties in a design process as many aerodynamic characteristics are difficult to predict [4].

Truly robust aircraft design requires knowledge of the aerodynamic behavior across the full range of operating conditions, to ensure improved flight performance whilst maintaining adequate stability and control characteristics. Achieving this coverage typically necessitates a large number of simulations including all relevant flow conditions and control settings, from which the quasi-steady aerodynamic

integral loads are extracted and compiled into an aerodynamic dataset.

Combining high-fidelity aerodynamic datasets with information from other disciplines — such as mass properties and actuators — enables realistic mission-level performance evaluations. This integrated approach supports more accurate assessments of aircraft configurations and reduces the risk in the design process.

The generation of RANS-based aerodynamic datasets has recently become feasible in industrial settings [5]. The computational process typically requires several weeks of wall-clock time on currently available high-performance computing (HPC) systems, a marked improvement compared to traditional dataset generation processes based primarily on experimentally generated data.

These recent developments in mass production of numerical data have made the introduction of robust automatic optimization strategies possible. The computational cost of optimization is primarily driven by the dimensionality of the design space and the modality of the objective function [6]. The dimension in a shape optimization problem is equivalent to the number of design parameters, while the modality is characterized by the number of the minima in the objective function. For automated dataset-based aircraft shape optimization, it is estimated that in the order of 3-20 parameters will suffice for the major design drivers, with only a few dominant modes. This is in agreement with other recent works [7]. Automated optimization is expected to yield higher-performing designs while significantly reducing manual design effort compared to state-of-the-art aircraft design processes.

Continuous advancement of HPC hardware and numerical algorithms render the routine application of high-fidelity dataset-based aircraft design optimization a feasible goal for the coming decade. However, currently very few aircraft design tools fully exploit high-fidelity for the optimization of new aircraft configurations [8].

This work demonstrates the potential of leveraging high-fidelity CFD simulations in an efficient and robust framework to enable high-fidelity dataset-based aerodynamic shape optimization for future aircraft configurations.

2. METHODOLOGY

To develop and demonstrate the proposed optimization framework, a simplified yet representative test case is selected. By staying within the subsonic regime and a limited angle-of-attack range the test case enables fast and efficient development, whilst maintaining physical relevance across the complete flight envelope. The configuration includes key components such as wings, control surfaces and a propulsion system, capturing all essential aerodynamic interactions.

The shape design parameters cover the wing position, the wing span, the wing root chord and the tail span. The fuselage geometry is held constant, reflecting a common industrial scenario where the propulsion system is pre-selected prior to aerodynamic optimization. Further details on the test case are presented in Chapter 3.1.

A Python-based framework has been developed to optimize parametric computer-aided design (CAD) models by evaluating high-fidelity aerodynamic datasets, using one of the

three strategies implemented: Grid search (GS), gradient-based optimization (GO) and Bayesian optimization (BO). Figure 1 provides an overview of the steps involved in one optimization iteration of the framework.

Each iteration begins by retrieving templates (step 0) for subsequent processes, including the CAD model, mesh generation scripts, and scripts required for the dataset pipeline setup and processing. Following successful geometry adaptation (step 1) a new mesh is generated (step 2) and passed to the dataset pipeline (step 3), limiting geometrical changes only by the allowable parameter range of the parametric CAD model. The pipeline produces the high-fidelity dataset, which is then evaluated by the objective function (step 4). Finally the parameter set for the next iteration is selected according to the chosen optimization strategy (step 5). This iterative process continues until the predefined number of iterations is reached, pausing for manual intervention if errors occur during any step.




FIG 1. Iteration steps of the high-fidelity dataset-based shape optimization framework.

To validate and verify the efficiency of the selected optimization strategies, a linear surrogate model was constructed using 20 design points obtained from a grid search over three shape parameters and their corresponding objective function values. This surrogate facilitates rapid evaluation, enabling multiple optimization runs under various settings. Such repeated testing aids in refining the optimization methods and mitigating risks in subsequent computationally expensive productive optimization iterations.

Two optimization strategies are examined for high-fidelity dataset-based shape optimization for aircraft design: Gradient-based optimization (GO) and Bayesian optimization (BO).

GO is a local optimization approach that can be straightforwardly implemented using finite differencing to approximate the gradients. While computationally efficient, GO can be sensitive to initial parameter values and prone to convergence to local minima. These risks can be mitigated by executing multiple GO runs in parallel, though at the cost of reduced overall efficiency.

In contrast, BO is a gradient-free global optimization technique that employs probabilistic surrogate models, typically Gaussian Processes, to efficiently explore the design space. The Efficient Global Optimization algorithm presented by Jones in 1998 [9] as one of the first implementations advanced BO and stated its applicability for expensive objective functions with moderate dimensionality, as often encountered in aerodynamic shape optimization [10]. However, Bayesian optimization is expected to require

significantly more computational resources than a more localized gradient-based optimization with a good starting point, rendering its expected introduction to industrially sized dataset-based optimizations to a later stage.

Both GO and BO have been applied to the validated test case and associated objective functions to demonstrate the feasibility of robust high-fidelity dataset-based optimization for aircraft design. The optimization procedures yield improved shape parameter sets for the test configuration tailored to the design missions, derived from high-fidelity datasets. These optimizations demonstrate the feasibility of the high-fidelity dataset-based approach currently under development.

3. TEST CASE

A simplified test case is used to demonstrate the dataset-based optimization approach. Whilst requiring considerably fewer resources than typical combat aircraft, most features of the methodology can be analyzed in an efficient way using this configuration.

3.1. Aircraft Configuration

Demonstration of the dataset-based high-fidelity optimization framework is performed on the OTIFAN configuration (Optimization Test Interceptor with **F**an). It is an expendable, subsonic unmanned aerial vehicle (UAV), designed for catapult launch and fast intercept missions against low-flying subsonic targets.

The configuration features an axisymmetric fuselage that houses a sensor head, batteries, control systems, and an internally mounted electric ducted fan (EDF) propulsion system. The wing has 5° incidence, while the all-movable, rectangular V-tail is inclined at 45° dihedral. Wing and tail use symmetric NACA 64-010 airfoils, which have been modified to a trailing edge thickness of 1%. Key specifications and components of the OTIFAN configuration are listed in Table 1.

A parametric CAD model of the configuration was constructed in CATIA V5, allowing for large variations in design parameters for the wing and empennage as shown in Figure 2. The full design space is parameterized using four shape variables: wing position x_W , wingspan b_W , wing root chord $c_{r,W}$, and V-tail span b_C .

Shape adaptation results in changed mass properties, which are estimated via lower-fidelity methods. The mass change of the components affects the center of gravity, influencing the trimming as well as stability and control characteristics of the configuration. The estimated mass breakdown for the minimum span configuration — with minimum wingspan and V-tail span — is provided in Table 2.

3.2. Numerical Setup

The numerical setup used to evaluate the configuration's performance consists of a meshing routine, the CFD solver configuration, the dataset definition and the objective computation. Each of the components is detailed in the following.

To solve the RANS equations, firstly the fluid domain has to be discretized into cells. A mesh generator creates unstructured, hybrid, vertex-centered dual cell grids by extruding prism layers from a tetrahedral surface mesh, which are




FIG 2. Parameterized CAD model of the OTIFAN configuration highlighting the design parameters and their respective bounds.

TAB 1. OTIFAN configuration specifications.

Specification	Value	Unit
Altitude range	0 ... 2000	[m]
Mach range	0.05 ... 0.25	[—]
Take-off mass	$\sim 3 \dots \sim 5$	[kg]
Propulsion type	Internal 90 mm EDF	
Batteries	4x Li-Po 6S, 1.6 Ah @ 80C	
Gross thrust	38 (static, uninstalled)	[N]
Flight time	118	[s]
Fuselage length	0.545	[m]
Fuselage max dia	0.180	[m]
Wing long. position	0.2 ... 0.3	[m]
Wing ref. area	0.05 ... 0.20	[m ²]
Wing root chord	0.060 ... 0.200	[m]
Wing tip chord	0.060	[m]
Wing span	0.5 ... 2.0	[m]
V-tail span	0.175 ... 0.275	[m]

TAB 2. Estimated mass breakdown of the OTIFAN configuration at minimum wing and tail span.

Component	Mass	Percentage
Fuselage	0.700 kg	23.3%
Sensors	0.280 kg	9.3%
Propulsion	0.370 kg	12.3%
Batteries	0.888 kg	29.6%
Servos	0.048 kg	1.6%
Wing	0.100 kg	3.3%
Tail	0.120 kg	4.0%
Payload	0.495 kg	16.5%
Total	3.000 kg	100%

then merged into a surrounding tetrahedral volume mesh. Control surface deflections are handled efficiently via the

Chimera approach [11] by modeling the surfaces as sub-meshes that can be displaced and connected using interpolation at runtime. High-curvature geometric features such as leading and trailing edges are locally refined.

A mesh convergence study has been performed across four mesh levels for a full dataset on three HPC nodes with 48 central processing units (CPUs) each. Starting from level 0 with approximately $3.5 \cdot 10^6$ nodes, the sources of the other levels were scaled so that the next higher and lower level corresponds to double and half the nodes, respectively. The overall mesh resolution distribution was kept constant, resulting in a cost doubling with each level increment. The first prism layer height was chosen to ensure $y^+ \leq 1$ and a growth rate of 1.33 was applied to generate up to 30 prism layers for accurate boundary layer capturing.

Mesh level -1 exhibited significant sensitivity in the control power derivatives (see Figure 3), indicating insufficient grid resolution. Mesh level 0 (see Figure 4), was selected for the optimization runs due to its favorable balance between computational cost and low sensitivity in integral coefficients and control derivatives (see Figure 3), suggesting acceptable grid independence.




FIG 3. Mesh convergence study showing mesh sensitivity of pitch control power for the OTIFAN configuration.




FIG 4. Mesh used for the OTIFAN optimization showing the farfield a) and details of the wing b), the fuselage c), the wing tip d) and the V-tail e).

Once the fluid volume is discretized with adequate resolution, the dataset pipeline calls the DLR TAU-Code [12] to solve the compressible RANS equations. The aerodynamic simulations are performed as transient RANS computations with the dual time-stepping approach with an implicit Backward-Euler scheme. Figure 5 shows a representative flow solution. The symmetry plane and the walls are colored by the mean Mach number and the mean pressure coefficient respectively.

The stagnation point at the sensor pod, as well as the wake of the fuselage are clearly visible in the symmetry plane.




FIG 5. Example flow field for the OTIFAN configuration: Contour plot of mean Mach and surface pressure coefficient on symmetry and viscous walls, respectively.

On the upper wing surface alongside the adjacent fuselage hull the typical suction area in the forward half of the wing chord is present. The tip of the control surface indicates also a suction region as expected from the tip vortex. In contrast, the aftbody and the engine exhaust indicate areas of higher pressure due to the boat-tailing and boundary condition respectively.

The aerodynamic dataset used to characterize the OTIFAN configuration is built from the flow conditions and control settings listed in Table 3. Reynolds and compressibility effects are neglected due to low altitude and subsonic operation. One complete dataset comprises of 30 RANS simulations and linear interpolation is employed for fast and robust evaluation.

Parameter	Symbol	Values
Altitude	H	0 m
Mach	Ma	0.25
Alpha	α	$[-2^\circ, 0^\circ, 2^\circ, 4^\circ, 6^\circ, 8^\circ]$
Delta	δ	$[-15^\circ, -5^\circ, 0^\circ, 5^\circ, 15^\circ]$

TAB 3. Dataset flow conditions and control settings.

In industrial applications the engine forces and moments are usually introduced by employing an engine deck. In addition, appropriate boundary conditions have to be set at the engine interfaces in the mesh, while the internal aerodynamics are omitted from the CFD simulation. In the test case presented here, a similar approach is used, where the EDF propulsion is modeled by an engine deck and boundary conditions generated by explicit simulations of the rotating fan in the fuselage of the design as shown in Figure 6. Several computations with varying onflow Mach numbers were performed applying the Chimera overset mesh approach. It is presently assumed that the optimizations are conducted at maximum power, sea level. Given the limited angle of attack range of the OTIFAN configuration, the effect of the angle of attack on propulsion performance was neglected.




FIG 6. Cut view of transient RANS Chimera simulation with rotating EDF used in the engine deck generation, showing pressure coefficient on viscous walls.

4. OBJECTIVE FUNCTIONS

A critical part of any optimization is the definition of the objective function. The availability of a dataset covering the entire flight envelope allows for the accurate evaluation of objectives of direct interest in the design process of an aircraft, including stability and control evaluation as well as the simulation of entire design missions.

4.1. Objective I

The objective for the configuration is composed of two design missions weighted 4:1. The primary mission is a direct intercept at target height H with initial launch from a catapult as illustrated in Figure 7. The secondary mission represents a contingency scenario: an inverted intercept following a half-loop maneuver after an unsuccessful primary intercept as shown in Figure 8.




FIG 7. Mission profile of the primary intercept mission.

The mission performance is rated based on figures of merit (FOMs) and constraint penalties, which are scaled and weighted, similarly to the approach by [13]. The FOM values are scaled across three reference points corresponding to 0, 100 and 200 mission points (pts), while the constraint values are scaled across three reference points corresponding to 105 %, 95 % and 0 %. Table 4 and 5 detail the FOMs, the constraints, the respective scaling values and the weights for the primary and secondary mission, respectively.

The flight time for the primary mission is set to 80 % of the total flight time and includes the launch from a catapult with




FIG 8. Mission profile of the secondary intercept mission.

a rail length of 8 m. To add a 5% margin on stall speed, the launch velocity is set to $1.05 V_S$ which is achieved by constant acceleration starting at $V_0 = 0 \text{ m/s}$, while the stall speed V_S is evaluated from the dataset at trimmed level flight with the lowest feasible velocity. To ensure structural integrity during the catapult launch segment the axial load factor is constrained to $n_{xb} < 15 \text{ g}$.

The climb segment is fixed to 80% of the primary mission time, where the climb is performed at the maximum rate of climb ROC to maximize the altitude H at the end of the primary mission as specified in Equation 1.

$$(1) \quad H = ROC \cdot 0.8^2 \cdot t_{max}$$

The optimum ROC is determined by evaluating the dataset for the highest value of the specific excess power SEP in trimmed level flight, which is equivalent to the maximum ROC . To compute the SEP in trimmed level flight, first the trimmed dataset is evaluated for the SEP as specified in Equation 2, in a second stage the results are then filtered for level flight condition where $n_{zg} = 1$ and $n_{xg} > 0$.

$$(2) \quad SEP = V \cdot n_{xg} = ROC$$

The last 20% of the primary mission time are allocated for acceleration from the climb velocity to the terminal velocity V_{max} , which is evaluated from the dataset at trimmed level flight, with the aim of maximizing the terminal impulse J .

TAB 4. FOMs and constraint of the primary mission including scaling and weighting.

FOMs	Scaling [0, 100, 200] pts	Weight
H	[500, 1000, 2000] m	3
J	$3 \text{ kg} \cdot [0.1, 0.2, 0.3] \cdot 340 \text{ m/s}$	1
Constraints	Scaling [105, 95, 0] %	Weight
n_{xb}	[0, 10, 15] g	4

The secondary mission utilizes the remaining 20% of the total flight time for a half-loop maneuver, enabling a second attempt at intercepting the target in inverted flight.

The minimization of the time to inverted is realized by finding the highest normal load factor at trimmed pitching moment and applying the corresponding flight conditions for the time needed to complete the half-loop as specified in Equation 3. The factor of two is applied to account for the quasi-steady modeling. The maximum velocity of the target V_{tgt} is computed at the maximum velocity at trimmed level flight, which is assumed to be close to the maximum velocity at inverted

level flight as specified in Equation 4. To ensure structural integrity during the pull-up or half-loop maneuver the normal load factor is constrained to $n_{zb} < 9 \text{ g}$.

$$(3) \quad t_{inv} = \frac{\pi \cdot V_{xb, pu}}{g \cdot (n_{zb, pu} - 1)} \cdot 2$$

$$(4) \quad V_{max, tgt} = \frac{V_{climb}}{1 + \frac{t_{inv}}{0.2 \cdot t_{max}}}$$

TAB 5. FOMs and constraint of the secondary mission including scaling and weighting.

FOMs	Scaling [0, 100, 200] pts	Weight
t_{inv}	$t_{inv}(n_{zb} = [1.5, 2.5, 9.0] \text{ g})$	1
V_{tgt}	[50, 150, 200] km/h	3
Constraints	Scaling [105, 95, 0] %	Weight
n_{zb}	[0, 6, 9] g	4

The mission rating is computed as the weighted sum of the rated FOMs, multiplied with the weighted product of the rated constraints acting as dampers as stated in Equation 6. Before computing the mission rating, the FOMs and constraints are normalized per mission as given in Equation 5. Ultimately, the objective is computed as weighted sum of the mission ratings mr_i as stated in Equation 7, where the mission weighting is 4:1 between the primary and the secondary mission.

$$(5) \quad \bar{w}_i = \frac{w_i}{\sum w}$$

$$(6) \quad mr(\vec{x}) = \left(\sum fr_i(\vec{x}) \cdot \bar{w}_i \right) \cdot \left(\prod cr_j(\vec{x}) \cdot \bar{w}_j \right)$$

$$(7) \quad obj(\vec{x}) = \sum mr_i(\vec{x}) \cdot \bar{w}_i$$

To account for mass changes due to geometric variations a mass-scaling model as given in Equation 8 is applied to the wing and the tail. It relates the mass scaling to a span scaling, corrected for the local fuselage diameter:

$$(8) \quad \left(\frac{m}{m_{min}} \right) = \left(\frac{b - D_F}{b_{min} - D_F} \right)^k \quad \text{where } k = 1.4$$

The component mass at the current and minimum spans b and b_{min} are noted as m and m_{min} , respectively. The parameter D_F denotes the fuselage diameter at the root of the corresponding lifting surface. The exponent $k = 1.4$ was chosen as a moderate deviation from the square-cube law, reflecting the structural mitigation effects of a fixed taper ratio of 0.6 and internal support structures such as hinge rods. These design features reduce the bending loads near the root, leading to a less-than-cubic growth in mass with span.

4.2. Objective II

To demonstrate the added possibilities made available by the introduction of dataset-based optimization, an objective was formulated in which a complete design mission is flown and where the flight path of the design mission is optimized for each geometric evolution to obtain a truly optimal design based on the high-level objective of interest. In this case, the maximum obtainable altitude for a 100 second flight was se-

lected as objective. The aircraft is launched at a 45° climb angle at a speed of 30 m/s and elevation above ground of 5 m . The mission integration is performed with a temporal resolution of 1 s and a quasi steady-state is assumed. The mission is divided into five linear segments in which the commanded climb angle is allowed to vary in a linear manner. A simplified aircraft flight control system attempts to follow the commanded path as closely as possible, whilst staying within the allowable flight envelope. In the search for the optimal climb strategy, an iterative approach is used, in which the optimizer adjusts the flight path incrementally in the direction for which the objective improves. This is performed by examining the combinatorial possibilities of each of the six flight path node points, by either increasing, decreasing or keeping the climb angle of each node. For each outer iteration, all $3^6 = 729$ possibilities of changing the flight path are computed and the best is selected as basis for the next iteration. The increments with which the flight path is varied is decreased as the optimization progresses. Flight envelope exceedance or ground collisions are detected and penalized to ensure realistic results.

For increased fidelity treatment of the structural wing weight, a localized structural analysis based on beam theory was applied. This was performed by slicing the wing into 100 segments, where the outer skin thickness of each segment was determined assuming a 3D-Printed Polylactic Acid (PLA) structure and a safety factor of around 3. A minimum skin thickness 0.4 mm was imposed. In addition to the outer skin, mass of a fixed pattern of ribs and spars was added. The present optimizations did not incorporate the local pressure distributions available from the CFD-dataset, but rather assumed a constant force distribution on the wing. Each wing was sized to withstand a lifting force of 90 N .

5. OPTIMIZATION STRATEGIES

The three optimization strategies demonstrated in this work are described in the following subsections.

5.1. Grid Search (GS)

Grid search, or brute-force optimization, samples the design space using a predefined grid, where the sample points in the design space are determined solely by the grid spacing in each dimension and are completely independent of the objective function.

This method is particularly useful for setting up the optimization problem in a robust manner and allows for the evaluation of the objective function without any dependence of an optimization strategy. When the grid resolution is sufficiently fine to capture the main modes of the objective function, the results generated via grid search can also serve to verify various optimization strategies, without the need for additional HPC resources.

The computational cost of grid search scales exponentially with the number of design parameters, as it is proportional to the product of grid points per dimension. Consequently, fine grid resolutions quickly become infeasible at higher-dimensional design spaces.

Nevertheless, the independence of each simulation allows for full parallelization, enabling all grid points to be evaluated simultaneously, potentially requiring the same wall-clock time as a single CFD simulation of a control sweep.

5.2. Gradient-Based Optimization (GO)

Gradient-based optimization (GO) is a local, iterative strategy that uses sensitivity information to evolve the design. In this framework, finite differencing is used to approximate gradients (sensitivities), as shown in Equation 9:

$$(9) \quad \frac{\delta obj}{\delta x_i} \approx \frac{\Delta obj}{\Delta x_i} = \frac{obj_{eps,i} - obj_{0,i}}{x_{eps,i} - x_{0,i}}$$

The quantities subscripted with *eps* denote the perturbed geometry values, while the 0 subscript denotes the values of the baseline geometry.

This requires generating a high-fidelity dataset for each design parameter at every iteration. The parameter updates are computed by scaling an initial step size according to the normalized sensitivity, using the mean of previously observed sensitivities while assuming linear dependence between parameters:

$$(10) \quad \Delta x_i = \Delta x_0 \cdot \frac{\left(\frac{\delta obj}{\delta x}\right)_i}{\left(\frac{\delta obj}{\delta x}\right)_{mean,i}}$$

To dampen overshoots, the parameter increment is reduced to 25 % after a sign change is detected in the sensitivity history.

The geometric increments used in the finite differencing are small, resulting in similar flow fields compared to the baseline geometry. This feature can be exploited to reduce the computational cost of the procedure by employing restart strategies.

It is worth noting that finite-difference gradients are sensitive to numerical noise, particularly in high-fidelity CFD environments where solver convergence criteria and mesh adaptation may introduce fluctuations, especially for separated flow. Proper selection of perturbation magnitudes and convergence tolerances as well as application of temporal averaging of the CFD computation is employed to reduce the noise.

While adjoint-based methods are more efficient if the number of optimization parameters is large [14, 15], in this work, finite-differencing was deployed due to its robustness and overall good efficiency in combination with the dataset generation procedure.

5.3. Bayesian Optimization (BO)

Bayesian Optimization (BO) is a global, gradient-free optimization strategy well suited for problems involving expensive or noisy objective functions. BO builds a surrogate model of the objective function and uses an acquisition function to identify promising evaluation points, balancing exploration and exploitation.

The implementation used in this study is based on the GPyOpt Python module [16]. For the demonstration on the OTIFAN test case, Gaussian Processes were employed to construct the surrogate, and the Expected Improvement (EI) acquisition function was selected. EI favors solutions with the potential to outperform the current best observed value obj' , as expressed in Equation 11:

$$(11) \quad u(\vec{x}) = \max(0, obj' - obj(\vec{x}))$$

At each iteration, the design parameters with the highest expected improvement, or maximal expected utility, are selected by trading exploitation and exploration of the objective function. This strategy offers a sample-efficient search mechanism, particularly advantageous in aerospace applications where each CFD run requires significant resources.

The scalability of BO is limited in high-dimensional design spaces due to the increasing cost of surrogate model updates and acquisition function optimization [17]. Since the implementation is based on the GPyOpt module, other supported acquisition functions such as lower confidence bound (LCB) or maximum probability of improvement (MPI) can easily be deployed.

6. RESULTS

The following sections present the results of various optimization runs, applying different strategies to distinct objectives.

6.1. Objective I: Grid search (GS)

The first optimization strategy applied to the OTIFAN test case was a coarse grid search over the three shape parameters of wing position x_W , wing span b_W and tail span b_C , covering the bounds of the parameter space. Additionally, three equally spaced points were sampled for the wingspan parameter, yielding a total of $2 \times 5 \times 2 = 20$ aircraft configurations. All simulations were executed in parallel at a total computational cost of $20 \cdot 6 \cdot 1536 = 184 \cdot 10^3$ CPU-hours, completed in under four hours of wall-clock time.

The resulting data were stored in an interpolator, which serves as surrogate model to efficiently benchmark the GO and BO strategies, as shown in Figure 9 and 10, respectively. Red circles indicate optima within 2 pts of the true maximum optimum value, while the dashed lines mark the respective objective and parameter values for the true optimum found through the grid search.

Both GO and BO were initialized at the midpoint of the parameter bounds. Within ten iterations the objective, as specified in Chapter 4.1, was increased from the initial value of $obj_0 = 72$ pts to the true optimum value of $obj_{max} = 90$ pts. The corresponding optimal parameter set was identified with negligible deviation. The wing position x_W reached its lower bound, while the tail span b_C had minimal influence on the objective. The optimum range for the wing span b_W appears to be in the lower third of the parameter bounds. The GO method found the optimum at the sixth iteration. Assuming that each sensitivity dataset requires one-third of the resources of a baseline dataset, the effective cost was $6 \times (1 + 3 \times \frac{1}{3}) = 12$ full baseline datasets. This corresponds to 60 % of the cost, but takes six times the wall-clock time compared to the grid search. The BO strategy also found the optimum in six iterations, corresponding to six full datasets. This achieves the optimum at 30 % of the cost and also requires six times the wall-clock time comparison to the grid search.

Due to its non-deterministic nature, BO's repeatability was assessed in ten independent runs. As shown in Figure 11, the method consistently converged to within 5 % of the global optimum.

These results confirm the validity of all three optimization strategies for the OTIFAN test case, while simultaneously verifying the correct implementation of both the GO and BO algorithms.

6.2. Objective I: Gradient-based optimization (GO)

Figure 12 presents the results of a GO run focused on optimizing the wingspan only, as it was identified as the most influential parameter in the GS study. The wing position and the tail span were fixed to their optimum values from




FIG 9. Gradient-based optimization (GO) based on grid search interpolation for Objective I.




FIG 10. Bayesian optimization (BO) based on grid search interpolation for Objective I.




FIG 11. Relative error for ten repeats of the BO based on grid search interpolation for Objective I.

the GS.

Starting from the parameter midpoint, the optimizer progressively adjusted the wingspan toward a local optimum, halting when the sensitivity approached zero.

The objective peaked at $obj_{max} = 95$ pts, surpassing the grid search result in iteration five. The final design was obtained within one day of wall-clock time and required $5 \times (1 + 1) = 10$ full baseline datasets.

The optimal wingspan was found to be $b_{W,opt} = 0.721$ m, yielding mission ratings of 106 pts and 61 pts for the primary and secondary missions, respectively.




FIG 12. Gradient-based optimization results on parameter wingspan only for Objective I.

6.3. Objective I: Bayesian optimization (BO)

Figure 13 shows the BO results for all three shape parameters over twelve iterations. The optimization history shows the five initial iterations where a latin hypercube sampling is applied. From iteration six onward, the objective remained above the grid search maximum, delivering a sufficiently optimal design for the three parameters within one full day of wall-clock time at the cost of six datasets.

The optimal shape parameters were found to be $x_{W,opt} = \min x_W$, $b_{W,opt} = 0.581$ m and $b_{C,opt} = \min b_C$.

The corresponding mission ratings were $mr_1 = 110$ pts and $mr_2 = 55$ pts, resulting in a maximum objective value of $obj_{max} = 97$ pts - the highest across all tested strategies.




FIG 13. Bayesian optimization results on all parameters

6.4. Comparison of GS, GO and BO

An overview of the optimization results is provided in Table 6. The lower portion of the table lists the FOMs and constraints.

All optimization strategies improved on the initial design, primarily by reducing the wingspan and increasing the objective by over 20 pts - approximately a 30 % improvement versus the configuration at the midpoint of the parameter

bounds. The optimized designs can intercept targets at altitudes of $H \approx 1000$ m and velocities of $V_{tgt} \approx 100$ km/h within the total mission time of ≈ 2 minutes, while the constraints on the load factors during catapult launch and pull-up maneuver are satisfied.

TAB 6. OTIFAN optimization results overview for Objective I.

Property	Midpoint	GS	GO	BO	Unit
obj	69	90	95	97	[pts]
x_W	0.218	0.208	0.208	0.208	[m]
b_W	1.250	0.500	0.721	0.581	[m]
b_C	0.225	0.275	0.175	0.175	[m]
mr_1	68	104	106	112	[pts]
mr_2	70	49	62	52	[pts]
cr_1	100	95	97	96	[%]
cr_2	97	100	99	100	[%]
$mass$	3.78	3.25	3.16	3.06	[kg]
H	733	1015	1058	1138	[m]
J	238	238	226	224	[N · s]
n_{xb}	5.2	9.9	8.1	9.0	[g]
t_{inv}	10.5	21.3	15.2	17.7	[s]
V_{tgt}	102	90	103	98	[km/h]
n_{zb}	4.7	3.1	3.9	3.6	[g]

The BO approach achieved the highest objective value by further compromising on the secondary mission and reducing the mass through decreasing wing and tail spans, resulting in the highest intercept altitude (the most heavily weighted performance metric).

Figure 14 visually compares the optimized BO geometry to the initial design at the mid-bound parameter values. However, as mentioned above, the optimized geometry presented is solely a demonstration of the shape optimization framework.




FIG 14. Comparison of the BO shape optimum (green) versus the initial shape (grey), including the fixed fuselage geometry (blue) for Objective I.

6.5. Objective II: Gradient-based optimization (GO)

In Figures 15, 16 and 17 the results of the climb mission optimization are depicted. The wingspan and root chord of the wing were selected as parameters in the optimization. It can be seen that a relatively complex climbing strategy results, terminated by a zoom climb to obtain the maximum controllable final altitude. For this objective function, the optimal wingspan and exposed root chord was found to be 0.636 m and 0.098 m respectively, with a wing mass of 0.147 kg, resulting in a maximum climb altitude of 1353 m. The optimized geometry was reached after nine optimization steps.

The optimization procedure thus involved the generation of 27 wing meshes and the solution of 810 RANS computations. This optimization demonstrates the possibility of relaxing the constraints of design missions, allowing the tool to identify the optimal combination of geometry and mission profile to achieve higher level mission goals. This is considered a very attractive capability in the aircraft design process.




FIG 15. Flight altitude over flight time for GO optimization result on Objective II.




FIG 16. Flight path angle, angle of attack and control surface deflection angle over flight time for the optimization result on Objective II.




FIG 17. Top view of the GO-optimized geometry for Objective II at $Ma = 0.25$, $\alpha = 2^\circ$ and $\delta = 0^\circ$.

7. CONCLUSION AND OUTLOOK

This study has demonstrated the feasibility of high-fidelity RANS dataset-based aerodynamic shape optimization within an automated framework. The software was applied to the Optimization Test Interceptor with Fan (OTIFAN) - a fixed-wing unmanned aerial vehicle configuration featuring

an internal electric ducted fan (EDF) propulsion system. Two objective functions and several geometrical parameters were considered.

The multi-mission objective function evaluated multiple flight phases, including launch, climb, and intercept, for a primary mission. A secondary mission, involving a follow-up intercept after a failed initial attempt, executed via a half loop and an inverted intercept maneuver. A combined wing geometry and flight path optimization was also shown, further demonstrating the possibilities available when the entire dataset is present for the objective evaluation.

The influence of non-aerodynamic disciplines such as mass properties, performance, flight dynamics, and structural loads was modeled using lower-fidelity methods to enhance the validity of the obtained design solutions.

Three optimization strategies were deployed, each with distinctive advantages and limitations:

- **Grid Search:** Useful with coarse spacing for initial setup and objective function validation or manual tuning, but computationally too expensive for automated optimization.
- **Gradient-Based Optimization:** Enabled efficient local convergence, or global via multi-start at reduced efficiency.
- **Bayesian Optimization:** Provided global search capabilities without requiring sensitivity computation.

All optimization strategies exhibited stable convergence behavior and were successful in identifying local or global optima, demonstrating the potential of high-fidelity dataset-based optimization in aircraft design.

ACKNOWLEDGMENTS

The funding of the activities by Airbus Defence & Space GmbH (www.airbus.com) and the scientific guidance by the Chair of Aerodynamics and Fluid Mechanics at the Technical University of Munich (AER, www.epc.ed.tum.de/aer) is gratefully acknowledged. The authors want to thank the German Aerospace Center (DLR, www.dlr.de/en) for providing the computational fluid dynamics solver TAU, Airbus Defence & Space GmbH for providing additional software frameworks as well as the Gauss Centre for Supercomputing e.V. (GCS, www.gausscentre.eu) for funding this project by providing computing time on the Linux Cluster and the SuperMUC NG at Leibniz Supercomputing Center (LRZ, www.lrz.de).

Contact address:

markus.v.maier@tum.de

References

- [1] Jun-xue Leng, Zhen-guo Wang, Wei Huang, Yang Shen, and Kai An. Multidisciplinary design optimization processes for efficiency improvement of aircraft: State-of-the-art review. *International Journal of Aeronautical and Space Sciences*, 26(4):2020–2042, Jul 2025. ISSN: 2093-2480. DOI: [10.1007/s42405-024-00811-8](https://doi.org/10.1007/s42405-024-00811-8).
- [2] Arthur Rizzi and James M. Luckring. Historical development and use of cfd for separated flow simulations relevant to military aircraft. *Aerospace Science and Technology*, 117:106940, 2021. ISSN: 1270-9638. DOI: <https://doi.org/10.1016/j.ast.2021.106940>.

- [3] Matteo Moioli, Christian Breitsamter, and Kaare Sørensen. *Turbulence Modeling for Leading-Edge Vortices: an Enhancement based on Experimental Data*. DOI: [10.2514/6.2020-1032](https://doi.org/10.2514/6.2020-1032).
- [4] Mengmeng Zhang, Jing Gong, Lilit Axner, and Michaela Barth. Automation of high-fidelity cfd analysis for aircraft design and optimization aided by hpc. In *2020 28th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*, pages 395–399, 2020. DOI: [10.1109/PDP50117.2020.00067](https://doi.org/10.1109/PDP50117.2020.00067).
- [5] Kaare Sørensen-Libik. *Overview and Results of Sub-project VitAMIn-ABC*, pages 123–127. Springer International Publishing, Cham, 2025. ISBN: 978-3-031-69425-7. DOI: [10.1007/978-3-031-69425-7_7](https://doi.org/10.1007/978-3-031-69425-7_7).
- [6] Robin Grapin, Youssef Diouane, Joseph Morlier, Nathalie Bartoli, Thierry Lefebvre, Paul Saves, and Jasper H. Bussemaker. *Regularized Infill Criteria for Multi-objective Bayesian Optimization with Application to Aircraft Design*. DOI: [10.2514/6.2022-4053](https://doi.org/10.2514/6.2022-4053).
- [7] Remy Priem, Hugo Gagnon, Ian Chittick, Stephane Dufresne, Youssef Diouane, and Nathalie Bartoli. *An efficient application of Bayesian optimization to an industrial MDO framework for aircraft design*. DOI: [10.2514/6.2020-3152](https://doi.org/10.2514/6.2020-3152).
- [8] Yaolong Liu, Siyu Zhang, Jiechao Zhang, Kexuan Yao, Mingqiang Luo, and Yantao Liu. Towards the design of future aircraft: A critical review on the tools and methodologies. *Aerospace Research Communications*, Volume 2 - 2024, 2024. ISSN: 2813-6209. DOI: [10.3389/arc.2024.13096](https://doi.org/10.3389/arc.2024.13096).
- [9] Donald Jones, Matthias Schonlau, and William Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 12 1998. DOI: [10.1023/A:1008306431147](https://doi.org/10.1023/A:1008306431147).
- [10] Andreas Kümmel and Christian Breitsamter. *Efficient Aerodynamic and Aeroacoustic Optimization of Propeller Sections Using Bayesian Methods*, pages 365–375. 07 2021. ISBN: 978-3-030-79560-3. DOI: [10.1007/978-3-030-79561-0_35](https://doi.org/10.1007/978-3-030-79561-0_35).
- [11] Kaare Sørensen, Udo Tremel, Herbert Rieger, and Stephan Hitzel. *Simulation of Manoeuvring Aircraft with the Unstructured Chimera Approach*. DOI: [10.2514/6.2007-123](https://doi.org/10.2514/6.2007-123).
- [12] Dieter Schwamborn, Thomas Gerhold, and Ralf Heinrich. The dlr tau-code: Recent applications in research and industry. 01 2006.
- [13] Clifton Davies, Marc Stelmack, Paul Scott Zink, Antonio De La Garza, and Peter Flick. *High Fidelity MDO Process Development and Application to Fighter Strike Conceptual Design*. DOI: [10.2514/6.2012-5490](https://doi.org/10.2514/6.2012-5490).
- [14] Michael Giles and Niles Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65, 04 2000. DOI: [10.1023/A:1011430410075](https://doi.org/10.1023/A:1011430410075).
- [15] Elssel Kolja, Sørensen Kaare, and Petersson Ögmundur. *Application of Adjoint Based Optimization on a MALE Platform*, pages 119–133. Springer International Publishing, Cham, 2018. ISBN: 978-3-319-72019-7. DOI: [10.1007/978-3-319-72020-3_8](https://doi.org/10.1007/978-3-319-72020-3_8).
- [16] The GPyOpt authors. Gpyopt: A bayesian optimization framework in python. <http://github.com/SheffieldML/GPyOpt>, 2016.
- [17] Alexander Forrester, Andras Sobester, and Andy Keane. *Engineering Design Via Surrogate Modelling: A Practical Guide*. 07 2008. ISBN: 978-0-470-06068-1. DOI: [10.1002/9780470770801](https://doi.org/10.1002/9780470770801).