A GENERALIZED APPROACH TO MULTIPHYSICAL AND MISSION-ADAPTIVE AEROSTRUCTURAL DESIGN WITH ROTORCRAFT APPLICATIONS

A Dissertation

by

ALLEN MILLER DAVIS

Submitted to the Graduate and Professional School of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Darren Hartl
Committee Members,	Douglas Allaire
	Moble Benedict
	John Valasek
Head of Department,	Ivett Leyva

December 2022

Major Subject: Aerospace Engineering

ABSTRACT

Mission-adaptive aerostructural design considers the alteration of structural geometries to improve multi-objective performance in multiple aerodynamic environments. Geometric structural alterations can be tailored to flight conditions derived from specific mission profiles, but adaptive structures design requires evaluating the aerostructural responses for each geometry to determine the optimal configuration for each mission stage. However, the addition of adaptive structures increases the design difficulty for several reasons. First, the aerostructural response must be evaluated over a range of feasible geometries, and morphing feasibility between each possible geometric configuration must be determined considering the proposed actuation method. Next, the effects of altering geometry must be related to vehicle performance. From a mission-driven perspective, altering the geometry may lead to changes in vehicle control inputs necessary to maintain specified velocity, attitude, etc., and the vehicle must be properly trimmed to match mission requirements. Finally, adaptive structures can increase design complexity, with configurations for each objective.

This work develops a mission-driven design framework combining aerodynamic, structural, mission, and optimization computational tools to design and optimize adaptive aerostructures. Aerodynamic conditions and vehicle properties are defined by mission requirements, where a mission is defined as a sequence of specified flight phases with a unique set of mission performance objectives. Aerodynamic and structural analysis tools iteratively trim the vehicle for each mission stage while considering geometric changes, viable actuation methods, and actuator sizing required for a complete structural description. Preferred geometries for each mission stage are then determined via optimization to improve performance metrics defined by mission objectives and requirements. The computational framework searches algorithmically for structural configurations that improve mission-driven objectives based on trim flight for each mission stage using a novel algorithm to consider both adaptive and fixed design variable selection to effectively solve this complex design problem.

Framework effectiveness will be evaluated by considering multiple mission-driven adaptive de-

sign and optimization problems. Specifically, the framework will evaluate mission-driven morphing rotorcraft designs, which consider dynamic multi-physical responses and cyclic control inputs to trim the vehicle. First, the computational framework will determine the placement of adaptive systems in the rotor blade based on multi-objective optimizations to improve individual mission objectives to determine preferred actuator locations and sets of adaptive geometric configurations for each mission phase.

A wide range of adaptive technologies can be explored for mission-driven configurations early in the design process with effective aircraft modeling. However, many existing design and optimization methods, including genetic optimization techniques, are not specifically designed for optimizations in which some design variables are adaptive while others are not. A common technique to utilize these tools is to consider each morphable design variable as a different design variable for each state or objective, which can increase the optimization problem complexity.

A novel optimization technique is developed and evaluated in this work for mission-driven, multi-objective design of adaptive structures. By considering design variables that can and cannot adapt differently, preferred configurations can be evaluated based on adaptive multi-objective performance and feasibility across the design space. It will be shown that this design space decomposition can used both as a post-processing technique and as a selection technique during multi-objective genetic optimizations to determine sets of realizable adaptive designs.

DEDICATION

For my dad, who taught me solving problems is fun.

For my mom, who taught me patience.

For Morgan, for her patience.

CONTRIBUTORS AND FUNDING SOURCES

If I have seen further, it is by standing on the shoulders of giants.

- Isaac Newton

Contributors

This work utilizes the efforts of numerous collaborators to develop a computational framework to efficiently design and optimize adaptive rotorcraft. None of this would have been possible without the contributions of a large number of collaborators and supporters from industry, academia, and defense.

I would like to thank Dr. Bochan Lee and Dr. Moble Benedict for developing and providing TRAC, a comprehensive rotorcraft analysis code that determines trim flight conditions efficiently and was essential for the forward flight analysis required for a rotorcraft design and optimization framework.

I would like to thank several collaborators from Boeing who contributed to this work under the Biomimetic Adaptive Airframe Technologies (BAAT) project. David Lazzara generated a large set of adaptive computational fluid dynamics data that was integrated into the framework as interpolation tables for efficient high-fidelity aerodynamics analysis. Michael Bass and I collaborated to turn a 2-D trailing edge camber morphing concept into a detailed design, which he then led fabrication efforts from Boeing R&T. Michael Bass and James Mabe then assisted me in rotor blade assembly and testing at Texas A&M University.

I would like to thank collaborators who contributed simulation tools from the MAESTRO Lab at Texas A&M University. The internal topology generation tool SPIDRS was created by Brent Bielefeldt, and its application to internal airfoils as an adaptive structures design tool was developed by Madalyn Mikkelsen. The fully-coupled aerostructural model requires relating local element-wise pressure fields to a structural rotor blade model, which was implemented using surrogate modeling and a finite element rotor blade model created by Trent White.

I would like to thank my dissertation committee for their guidance and feedback. This work was supported by a dissertation committee consisting of Professor Darren Hartl and Professors Moble Benedict and John Valasek of the Department of Aerospace Engineering and Douglas Allaire of the Department of Mechanical Engineering.

Funding Sources

This work was supported by Boeing R&T and U.S. Army Contracting Command Redstone Arsenal, Fort Eustis, VA 23604-5577. Under Technology Investment Agreement (TIA) No. W911W6-18-2-0004, titled Biomimetic Adaptive Airframe Technology (BAAT).

NOMENCLATURE

BAAT	Biomimetic Adaptive Airframe Technology
BET	Blade Element Theory
BEMT	Blade Element Momentum Theory
BVI	Blade Vortex Interaction
c	chord
c_∞	Speed of sound
C_l	Coefficient of Lift
C_P	Coefficient of Power
C_T	Coefficient of Thrust
CFD	Computational Fluid Dynamics
CST	Class/Shape Transformation
DEAP	Distributed Evolutionary Algorithms in Python
D_i	Single design with similar fixed variables with respect to a specified location in the design space $\underline{x}$
DOE	Design of Experiment
$\mathbb{D}_i$	Set of all designs with similar fixed variables with respect to a specified location in the design space $\underline{x}^f$
$\underline{f}()$	Multi-objective function
FEA	Finite Element Analysis
FM	Figure of Merit
GA	Genetic Algorithms

GJ_{blade}	Rotor Blade Torsional Stiffness
HHC	Higher Harmonic Control
J^*	Pareto-optimal solution
$\hat{J}^*(\underline{x}^f)$	Pareto-optimal solution considering a set of fixed variables $\underline{x}^f$
$\hat{\hat{J}}^*(\underline{x}^f)$	Approximate Pareto-optimal solution considering a set of designs with similar fixed variables
L	Lift
LHS	Latin Hypercube Sampling
$\frac{L}{D}$	Lift to Drag Ratio
M	Mach number
MAAB	Mission-Agnostic Adaptive Blade
MT	Momentum Theory
N_b	Number of rotor blades
NiTiHf	Nickel Titanium Hafnium
NSGA-II	Non-dominated Sorting Genetic Algorithm II
NSGA-AD	Non-dominated Sorting Genetic Algorithm for Adaptive Design
OML	Outer Mold Line
PSO	Particle Swarm Optimization
r	Rotor span-wise Location
R	Rotor radius
Re_x	Reynolds Number
SMA	Shape Memory Alloy
SPEA	Strength Pareto Evolutionary Algorithm
T	Thrust

TAA	Terminal Arrival Area
TRAC	TAMU Rotorcraft Analysis Code
$\underline{x}$	Variables for a single design
x^f	Fixed design variables
$\underline{x}^{f*}$	"Preferred" set of fixed variables for a single design
W	Vehicle Weight
α	Angle of Attack
ϵ	Similarity parameter for grouping approximate Pareto frontiers
λ	Inflow Ratio
λ_i	Induced Inflow Ratio
ψ	Rotor Blade Rotation Angle
ω	Rotor blade angular velocity
σ	Rotor Solidity
θ	Pitch Angle
θ_{blade}	Rotor Blade Manufactured Twist Angle
θ_c	Collective Pitch Angle
$\theta_{deflection}$	Structural Deflection Angle

TABLE OF CONTENTS

Page

LIST OF FIGURES

FIGURE Page

LIST OF TABLES

TABLE Page

1. INTRODUCTION AND LITERATURE REVIEW

1.1 Introduction

Mission-adaptive aerostructural design considers the altering of structural geometries to improve multi-objective performance in multiple aerodynamic environments. Geometric structural alterations can be tailored to flight conditions derived from specific mission profiles, but adaptive structures design requires evaluating the aerostructural responses for each geometry to determine the optimal configuration for each mission stage. However, the addition of adaptive structures increases the design difficulty for several reasons. First, the aerostructural response must be evaluated over a range of feasible geometries, and morphing feasibility between each possible geometric configuration must be determined considering the proposed actuation method. Next, the effects of altering geometry must be related to vehicle performance. From a mission-driven perspective, altering the geometry may lead to changes in vehicle controls to maintain the same velocity, attitude, etc. The vehicle must be properly trimmed to match mission requirements. Finally, the design space is dramatically increased for adaptive structures since morphable parameters can change for each configuration, often with a range of intermediate geometries.

In the aerospace engineering design process, it is advantageous to assess the maximum number of concepts and configurations in the earliest possible design stages. However, comprehensive vehicle modeling and simulation during preliminary design increases computational costs. Trade-offs between fidelity, computational requirements, and time must then be carefully considered.

Competing objectives and the resulting trade-offs are common in aerospace engineering design, especially in the design of rotorcraft. Some competing objectives are highly distinct (power vs. weight, strength vs. flexibility, etc.) while others represent identical performance metrics as measured at different times (power at takeoff vs. power in steady flight, etc.). With respect to both competing objective classes, the potential to replace traditional permanent design variables (i.e., set during design and built into the vehicle) with non-permanent alternatives (i.e., changeable

during flight) offers potential multi-objective performance improvements when prioritizing competing objectives at different times. Vehicle adaptivity, especially of the structural variety, enables this change in design approach. In this work, we will refer to the former static design variables as fixed while the reconfigurable variables that form the basis of morphing structures design will be referred to as adaptive.

Optimization methods navigate a design space by altering input variable values (typically engineering design variables) of individuals (designs) over generations (sets of designs) to determine preferred variables given specified objectives. The population evolves based on the performance of previous generations, prioritizing designs yielding better performance across multiple optimization selection criteria and constraints. When multiple objectives are considered in the design process, they must either be i) combined into a single objective (i.e., score), which allows down-selection of a single design later in the process or ii) considered independently, which allows effective design exploration early in the process. A common multi-objective optimization determines and prioritizes non-dominated designs, from which, given n objectives, "it is impossible to depart without making at least one of the n objectives worse [146]." Maximizing performance over multiple objectives requires considering trade-offs between each objective; final design requires design decisions to be made, variables to be set at a value, and a single down-selected design configuration to be propagated along the development process.

Restrictive trade-offs between competing objectives become possible and strict down-selection becomes unnecessary when *in situ* adaptivity is introduced. Structural geometric morphing is a particularly impactful, if highly challenging, option for such adaptivity. Altering vehicle geometry based on changes in forward flight speed, loading requirements, pilot maneuvers, and even environmental conditions can potentially eliminate the need for permanent compromises between performance metrics [5]. In other words, a single performance value is replaced by a performance range via adaptive changes in formerly fixed variables. Rotorcraft in particular experience large changes in flight conditions due to differences between hover and forward flight, as will be demonstrated later.

This novel work focuses on the adaptive aspect of aerostructural design to increase vehicle performance when considering multiple objectives and will introduce a completely new formulation of the design optimization problem to do so. In multi-stage missions, such as air assault or medical evacuation (MEDEVAC) missions, rotor morphing will be shown to improve performance via geometric alterations between mission stages such as hover or high-speed flight. However, many non-morphable design decisions, such as the fixed location and size of adaptive structural components, must also be determined in the engineering design problem. Since structural reconfiguration generally focuses on a limited number of concepts (e.g., twisting, camber changes, etc.), not all design variables can be adaptive, and it is important to recognize which structural variables should be targeted for geometric adaptivity.

Relating design variables to changes in objective space can highlight potential performance increases via adaptivity. Variables that, when altered, improve some competing objectives (Design Variable 2 from Fig. 1.1) without requiring any significant alteration of other variables offer adaptive advantages such as increased performance across multiple missions. Thus, the selection of adaptive design variables is dependent on both the performance benefits and the feasibility of morphing the design variable in question. A new mathematical problem formulation and associated computational framework has been developed to determine optimal morphing design parameters over multi-stage missions considering both fixed and adaptive variables. Rotorcraft performance analysis tools have been developed to relate design variables to both single-objective and more general multi-objective mission performance goals. It will be shown how genetic optimization algorithms can drive a targeted search exploring the multi-objective design space to determine both traditional fixed design variables (i.e., optimal actuator placement and sizing) and adaptive variables (e.g. active material stimulus level), arriving at a truly realizable adaptive rotorcraft design with the performance advantages of adaptivity. These optimization methods are powerful tools for optimizing adaptive structural designs while considering both local morphing subsystem performance and overall vehicle performance, including the effect of fixed design variables on objectives.

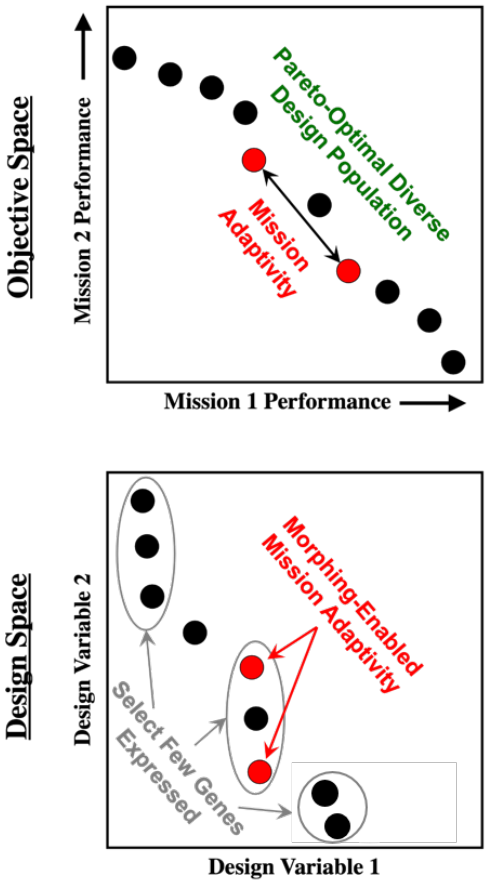

Figure 1.1: Adaptivity offers the potential to realize multiple positions in the objective space (mission performance) via changes in a limited number of variables in the design space. Here, changes in variable 2 effect performance even when variable 2 is unchanged, indicating that variable 2 may be a good candidate for adaptation.

1.2 Adaptive Aircraft

Adaptive aerostructures in the most general sense have been developed for over 100 years. In the late 1800s Otto Lilienthal demonstrated George Cayley's proposed concept originally from the late 1700s of breaking aircraft functions into lift, propulsion, and control; the concept has remained ubiquitous in aircraft design to this day [45, 114]. Controls (logic, electronics, and physical surfaces) by definition adapt an aerospace system to its environment or updated functional goals, and thus all aircraft are adaptive in some manner. Since control elements exist throughout

all aircraft, the specific term "adaptive" has come to refer to changes in geometry that go beyond the conventionally accepted and implemented control surfaces required for flight maneuvers and/or stability. It has specifically come to be associated with conformal alterations of structural shape, and such conformal changes represent the focus of this work.

Aircraft morphing beyond controls has also existed for a long time, though less often in mass production. The addition of such adaptive systems must prove advantageous when compared to drawbacks such as additional weight, controls, and complexity. One example of adaptive advantages warranting mass production were the Mission Adaptive Wing (MAW) and the variable sweep wing, used on the F-111 and F-14 [94, 125]. These aircraft could change the trailing edge camber and level of sweep while keeping the actuation systems inside the fuselage. The effects of different wing sweep values at various supersonic velocities on shock waves and maneuverability could be leveraged to select the best wing sweep for a given flight condition.

Additional work was performed to explore the effects of trailing edge camber on the variable sweep wing for the F-111 after its development [58, 145, 20]. While this work showed advantages and applicable actuation mechanisms, the technology was only demonstrated and was not transitioned to mass production. One limitation on the adoption of adaptive novel aerospace technologies as applied to existing aircraft concepts is the difficulty of changing existing technologies and designs. It would be far more attractive for a team to consider adaptive technologies targeted towards novel aircraft concepts relatively early in design. For example, the rapid expansion of unmanned aerial vehicles (UAVs) [23, 18] research and development has provided an attractive opportunity to explore new adaptive structures solutions from the very beginning of their design.

Several aircraft wing morphing methods have been evaluated to match optimal configurations with specific mission operations [77], as shown in Fig. 1.2. Often designers have looked to nature for inspiration, such as studying bat and bird wings [25, 159], where the shape of lifting surfaces may change substantially as the animal seeks to match changing flight conditions or objectives. Thus, it may be advantageous to apply similar principles to designing aircraft for sets of specific aerospace missions that may be strikingly diverse and may change over the course of operating

a given vehicle. This is particularly true of high-performance military aircraft. Of these, rotorcraft missions are particularly diverse, considering both the aerodynamic environment and flight requirements. For this reason, rotorcraft design was considered an attractive demonstration case for the new fixed-adaptive design approach.

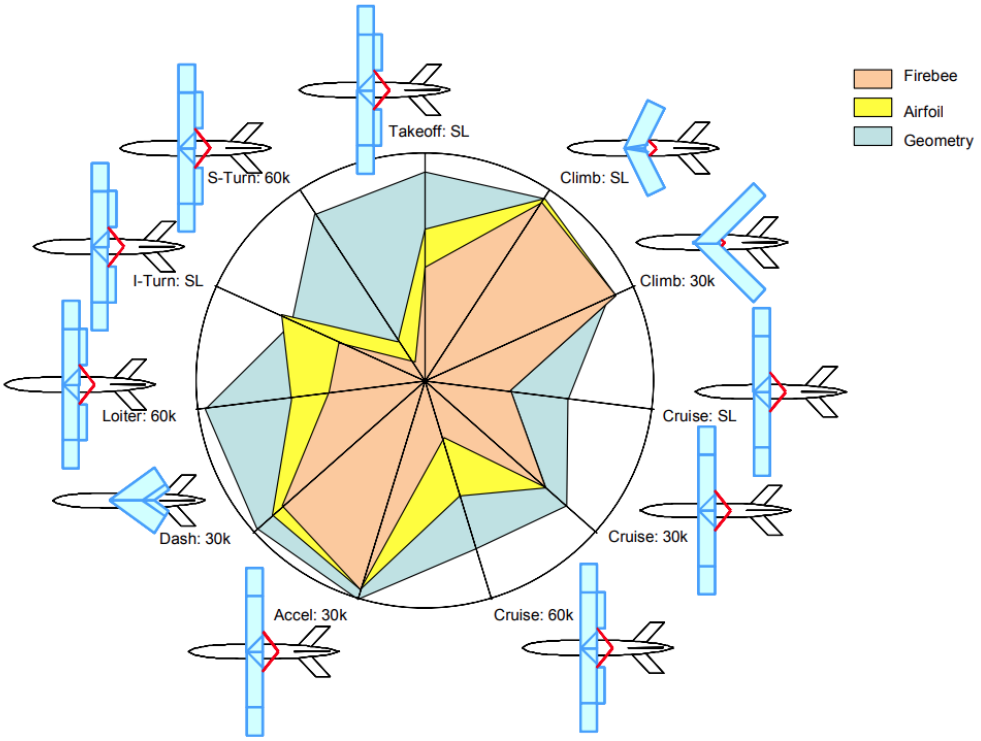

Figure 1.2: Different maneuvers and flight conditions will lead to different optimal designs.

1.3 Adaptive Rotorcraft

Adaptive rotorcraft technologies must consider a wide range of dynamic effects not seen in fixed wing aircraft due in general to the unique multi-modality of vertical lift and in particular to the challenges of rotor rotation. These added dynamic, and often harmonic, responses may increase the difficulty quantifying adaptive advantages over fixed rotor designs. For instance, the mass distribution of the rotor is critical due to inertial forces in rotation, and those effects change

with rotor angular velocity and disk loading. Adding actuation systems into a rotor blade alters rotor mass distribution, leading to increased modeling complexity. Rotors are typically carefully designed to minimize unwanted drag, moments, and the creation of and sensitivity to vibrations. The addition of adaptive elements on or in the rotor blade must both consider these effects and others while justifying added complexity.

In general, given a specific objective for an aerodynamic surface (i.e., minimize drag) under a specific aerodynamic condition, there exists an optimal set of design variables (chord, taper, twist, etc.) [55, 56]. Thus, the number of optimal shapes for a given mission will in general equal the number of different flight aerodynamic conditions, which are driven by both environmental conditions and mission requirements such as a specified forward flight speed. Changes in these conditions inform the optimal design, particularly for rotorcraft, where inflow conditions can vary significantly as a function of velocity through the rotor disk. Since there exists for a given rotor blade an optimal shape schedule corresponding to the best performance, a rotor blade that matches such an optimal shape considering the aerodynamic profile during a mission stage such as cruise, is preferred. The local angle of attack is a function of rotor blade attitude, twist, incoming airflow due to forward flight (if applicable), and inflow induced by the rotor blade [101, 74, 149], which will affect the rotor blade performance. High performance static rotor blades are typically designed with a twist and geometry schedule to equalize the angle of attack and inflow such that high lift-to-drag ratios can be achieved along the rotor blade across a broad range of flight conditions. However, such designs are inherently a compromise between the best designs for each given aerodynamic environment.

Changes in flight condition and objectives may impede static, manufactured designs from meeting all the challenges of a mission effectively; as an example, the local trim flight pitching angle for a high lift objective may be very different from the local pitch required for a low lift objective. In the case of competing objectives and changing flight conditions, adaptive structures have been studied in rotorcraft with varying levels of success to morph between preferred geometries. One study [175] listed the active subsystems considered as including 1) full blade feathering, 2) con-

trollable twist, 3) propulsive jet flap, 4) non-propulsive jet flap, 5) circulation control, 6) X-wing, 7) advanced center blades.

Tests were performed in the early 1990's on a number of actuation devices and active structures at the subsystem level (local regions of the rotor blade) to explore practicality [151, 103]. Piezoelectric methods were explored for both internal actuators and external flaps [150, 152]. These piezoelectric systems were popular actuation methods due to their fast response times and smaller volume when compared to hydraulic systems. Work was performed to develop actuators that could cycle significant strains over large numbers of cycles [150, 119, 152], then those actuators cycled trailing edge flaps to explore a number of potential adaptive advantages [169, 170, 73].

Other actuation mechanisms include integrated shape memory alloys (SMAs), known for their high power density and large recoverable strains [91]. SMAs have slower actuation frequencies than piezoelectric actuators, but can generate higher strains with high power density. Embedded shape memory alloy (SMA) torque tubes have been thoroughly investigated, with studies demonstrating their feasibility for morphing trailing edges [154]. Numerous morphing rotor blade actuation systems have been developed and tested for various morphing rotor mechanisms, to be described herein. One such conclusion stated in [151] was: "it is seen that imbedded actuator concepts, i.e. pitch, twist, and camber control, are not practical at this time. Servoflap control, using hinged control surfaces driven by discrete actuators emerges as the most suitable candidate for smart material actuation." This would notably not discourage other work on internal, embedded actuation, but encouraged new research exploring active edge flaps on the trailing edge of rotor blades.

1.3.1 Adaptive Rotorcraft Advantages

Several benefits of active structures in rotor blades have been identified, including noise reduction. Interactions between the main rotor blades and the vortices generated by those same rotors, called blade-vortex interactions (BVI) [123, 175, 83], generate a significant level of noise. Fixed wing aircraft rarely encounters similar phenomena because vortices generated by the wing trail behind the vehicle in forward flight. A rotorcraft in hover or low speed forward flight may encounter

vortices generated by a previous rotation or another rotor blade several times a second. A number of methods have been investigated to reduce the noise generated by BVI. One such method slightly adjusts the swash plate mechanism to change the blade pitch and twist slightly to limit BVI without requiring high-frequency active actuations systems. This method is limited in that the generally swash plate responses are necessary for full vehicle control. Another method to reduce BVI is to disrupt the trailing edge aerodynamic environment via active trailing edge tabs, which will be discussed later.

Reducing power required for trim flight is another common design objective. Minimizing main rotor power P, coefficient of power C_P and coefficient of thrust C_T typically result in reduced fuel burn and emissions [74]. Lift can generally be increased by increasing the planform, while increased camber is generally correlated with increased lift and drag [55, 56]. Thus, it can be advantageous at different times under different loading requirements to change the relevant geometric parameters best suited to minimize drag and the associated power to overcome said drag. Reducing trim power for a given lifting surface can generally be posed as reducing drag given a certain lift and rotor disk plane requirement, thus maximizing lift over drag (L/D). For fixed wing aircraft this L/D performance metric may be sufficient for design selection, especially if the aerodynamic environment stays largely within the linear region of the airfoil. This local performance metric enables easier analysis by focusing on the local airfoil region under a steady state flight condition. However, often a more robust vehicle evaluation may be required to determine trim power. For instance, changing incoming angles of attack and velocities on the rotor blade due to rotation around the hub requires additional analysis.

Rotor blades are subject to a particularly dynamic aerodynamic environment when compared with fixed wings. Increased rotor performance analysis to determine required power follows from increased complexity due to dynamic lift generation, which may be difficult to treat as a steady analysis. One reason for this increased difficulty is determining the overall power required to rotate the rotor, which may change depending on the angle of rotation around the hub. Another reason is the changing aerodynamic environment may lead to aerodynamic conditions requiring

analysis beyond steady laminar flow, such as rotor wakes in forward flight. When flow is laminar or simply does not separate from the airfoil, panel methods and other low-fidelity analysis tools may be employed with reasonable accuracy. As viscous effects increase, more complex analysis tools may be required. A more thorough review of the aerodynamic analysis tools will be further discussed and evaluated as a trade-off between fidelity and computational costs.

Vibration reduction is another objective that varies significantly between rotorcraft and fixed-wing aircraft. Rotor blade rotation generates dynamic effects such as vibrations which can lead to harmonic effects impacting performance and stability. Responsive actuation systems as the rotor rotates around the hub would be valuable for balancing unsteady forces or neutralizing unwanted high-speed dynamic and vibratory effects. Piezoelectric actuators have been proposed as an efficient, lightweight method for high-frequency adaptive vibration control, with actuation cycles over 4 times per revolution.

Rotor blade morphing takes many forms with a particular distinction between reconfigurable and high frequency active morphing. Reconfigurable camber morphing considers different aerodynamic environments for the vehicle, while high frequency active morphing considers different aerodynamic environments for the local rotor blade regions. For example, a reconfigurable rotor may alter camber between sea level and 10,000 feet to account for changes in air density or different forward flight speeds between mission stages such as hover and cruise. Meanwhile, high frequency active camber morphing may change camber between advancing and retreating sides of the rotorcraft to reduce power, BVI, vibrations, etc.

1.3.2 Adaptation I: Active Rotor Twist

A wide range of adaptive rotor blade technologies altering twist, airfoil shape, and dynamic responses have been explored for the main rotor [60]. The first branch of adaptive technologies discussed in this work is active twist, in which the rotor blade is altered from its manufactured twist schedule during flight, either as a function of rotation around the rotor blade or as a function of the environment. The advantages and limitations of altering the rotor blade twist schedule during flight have been explored using multiple actuation mechanisms. Piezoelectric actuation has been devel-

oped for active twist rotor blades for several purposes, include reducing vibrations generated due to the dynamic effects of the rotor blade. Fiber-reinforced piezoelectric actuators were developed to actively dampen the rotor blade and thus reduce vibrations [169, 126]. These designs have been tested in a full-scale wind tunnel, generating a maximum active twist change of 1.1-1.4 degrees and reducing vibration-generated loading up to 90% [169, 165]. In another work a Mach-scale piezo-induced active twist rotor blade achieves an amplitude of 0.78 degrees, increasing thrust authority [14]. Such active twist rotor blades have also been optimized to maximize active rotor twist deflection while maintaining mass and structural constraints [29].

Other piezoelectric active twist systems have been developed for lower voltage requirements and increased actuation power densities [168]. Extensive work has also been conducted towards controlling and modeling adaptive piezoelectric adaptive systems [28, 119, 31]. As piezoelectric technology has become more used and standardized, the technological readiness and applications have steadily increased. Active rotor blade twist has also been shown to increase required power, particularly small harmonic active twist in higher-speed forward flight [63].

Not all active rotor twist systems are high frequency. There also exist advantages for reconfigurable twist rotor blades, enabling different levels of twist for different missions, flight conditions, mission stages, etc. One common objective-driven need for adaptive rotor blade twist is due to a change in required lift, either due to increased vehicle cargo or changes in air density. Such changes do not require actuation systems with actuation frequencies comparable to rotor blade rotation.

One such adaptive design method utilizes shape memory alloys as compact actuators. Shape memory alloys couple strain with stress, temperature, and even electromagnetic state and can generate large recoverable strains. Shape memory alloys (SMAs) have been developed for high strength and extended life cycles [8, 26].

Shape memory alloys also have a wide range of uses as active actuators, including in active twist rotor blades [104, 103, 151, 70]. Shape memory alloy torque tubes have been developed as twisting mechanisms with high forces, ideal for twisting high-stiffness rotor blades and have

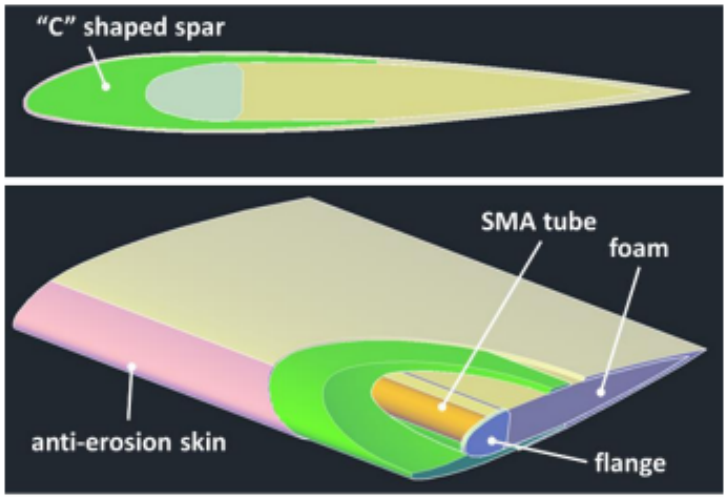

Figure 1.3: SMA torque tubes were embedded in the rotor blade near the spar to alter the twist schedule [34].

been. SMA wires have also been implemented as tensegrity mechanism components [61] rotor rotor blade twisting. Other morphing SMA active twist wire exploration includes skewed SMA wires around an internal cylinder, but this study deemed simply using an SMA torque tubes as more promising [34] and is shown in Fig. 1.3. SMA actuators are often actuated thermally [137], leading to much more compact designs than other high power-dense actuation mechanisms such as hydraulic systems. These SMA tubes are often heated with small heaters, which can be actuated very quickly with little excess weight for controls [108]. Camber morphing is another method of adaptive rotors which has generated significant interest recently, as actuation response has increased and actuator size has decreased. Often camber morphing devices on rotorcraft are referred to as active trailing edge flaps. In this work, the change of trailing edge camber due to active trailing edge flaps in which the chord is not actively altered will be considered alongside work done specifically changing the camber of the rotor blade.

1.3.3 Adaptation II: Active Camber Morphing

Camber morphing alters the shape of the local airfoil, often around the trailing edge. These changes in the airfoil outer mold lines (OMLs) alter the local lift, drag, and moment profiles with

respect to incoming velocity and angle of attack. This adaptive morphing category can be used for a number of purposes; the dynamic environment limits the advantages of simply maximizing a single lift-to-drag ratio, and camber morphing can also be actuated at higher velocities as demonstrated using piezoelectric actuation. As with active twist, cyclic changes have been investigated for camber morphing to match the dynamic local aerodynamic environment of the rotor blade, particularly noting the differences between the advancing and retreating sides of the rotor disk plane. Most work considering cyclic active camber morphing considers the outer regions of the rotor blade; those regions encounter the most substantial local rotor blade velocity changes in forward flight.

Several applications for adaptive camber morphing have been identified for rotorcraft, as have different methods of altering the overall airfoil shape. Geometric changes near the rotor blade trailing edge can significantly alter the local aerodynamics. These changes to the trailing edge can notably alter the rotor blade aerodynamic response and are often easier to implement compared with geometric changes to other OML regions due to the highest pressures on the airfoil usually being located on the leading edge. However, this has not eliminated effects to morph the airfoil leading edge, as active leading-edge morphing has been investigated to reduce dynamic stall with an active dropping leading edge [54]. Integrated shape memory alloy actuators have also been developed for active deicing via leading-edge rotor blade morphing [155]. Active trailing edge flaps driven by piezoelectric actuators have been integrated into the rotor blade for vibration control [116, 92]. Another camber morphing system changes the trailing edge shape via an internal piezo-electric active linkage structure near the center of the rotor blade [51]. These active trailing edge flaps have also been shown to offer power reductions [106] due to vibration reduction [135] (shown in Fig. 1.4) and have also been used to reduce rotor noise due to blade vortex interactions [175, 84, 83]. Pitch link load reduction has also been achieved by trailing edge flaps [86, 85]. Numerous local aerostructural advantages have been demonstrated via active camber morphing. Active frequency inputs have enabled active camber morphing to reduce vibrations, prevent dynamic stall, and reduce required power. The effects vary depending on the specific controls and in-

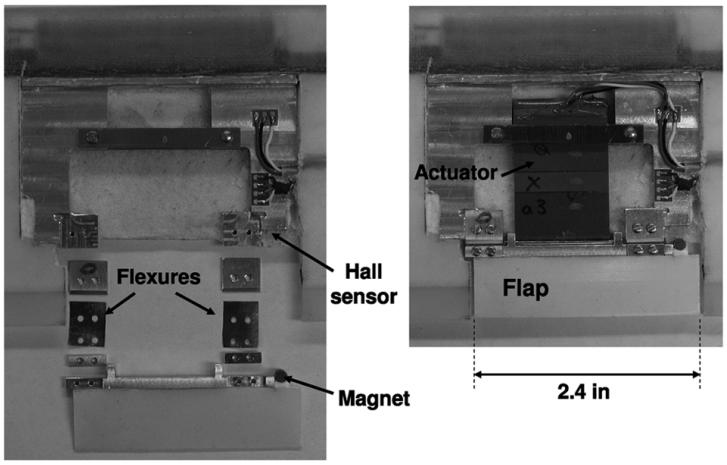

Figure 1.4: Trailing edge flap mechanisms have been developed and tested for vibration reduction [135].

puts, as well aerodynamic environment. Torsionally softer blades have been introduced to facilitate easier actuation with lower loading and stresses [134, 133, 172]. Such active systems require careful controls and dynamic analysis. Comprehensive analysis tools have been developed to account for the dynamic effects of active trailing edge camber morphing added into a rotor blade [113].

Reconfigurable camber morphing using shape memory alloys offers increased ranges of actuator motion when compared with piezoelectric actuation at the cost of frequency. These actuators have also been integrated into rotor blade design as passively adaptive systems with geometries correlated with specific temperatures, morphing the rotor based on the external environment. For instance, at higher temperatures air density decreases, so passively increasing camber can increase the coefficient of lift to maintain lift generation under changing temperatures [41], as shown in Fig. 1.5. Work has also been performed to change the rotor blade airfoil camber using SMA wires [153]. Computational frameworks have been developed to consider camber morphing which does not require high frequency actuation [118], some considering optimal mission-based performance [109]. Piezo-electric actuators have also been used for reconfigurable trailing edge rotor blade design to increase L/D over a fixed rotor blade when fight requirements change [176].

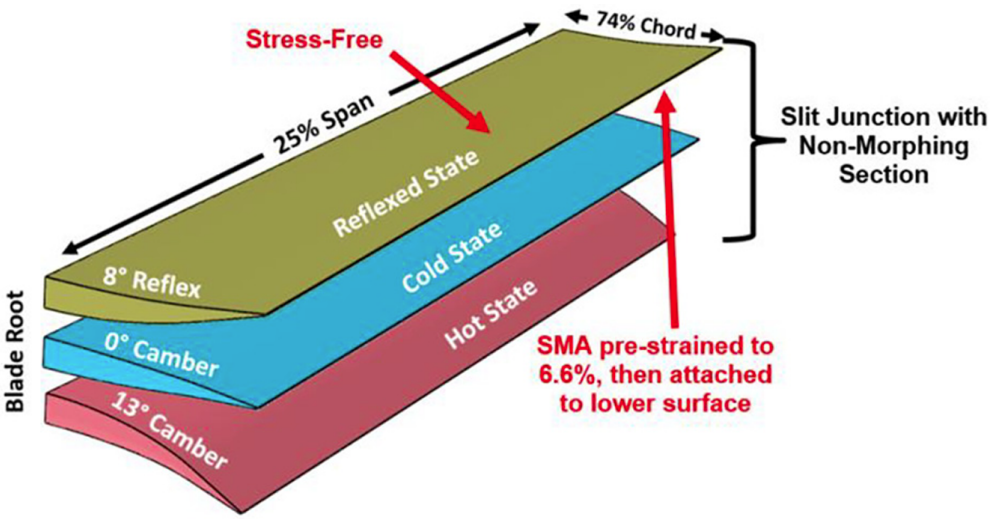

Figure 1.5: In this work by DiPalma and Gandhi, a region of the rotor blade span passively changes camber as a function of temperature using shape memory alloys [41].

1.3.4 Adaptation III: Active Chord and Span

The rotor blade morphing previously described considers changing twist and shape, altering the local aerodynamic profile and angle of attack. Additional morphing rotor blade methods have been developed to explore the benefits of an active extendable trailing edge tab to reduce blade vortex interactions. The addition of active tabs has been shown to reduce rotor noise and power when properly applied [121].

The use of internal extendable honeycomb structures in the rotor blade has enabled adaptive rotors in which the rotor chord can be extended for different mission objectives [11]. This extendable chord research has included NiTiNOL bistable arches as actuators and von-Mises trusses with variable length links [10], along with a biased, spring-based bistable systems [129]. Other research has investigated optimal chord changes based on mission-driven objectives in larger design optimization frameworks. Linearly-variable chord extensions are still being investigated and developed, including additional work under SABRE [13, 110, 109]. Extendable trailing edges to increase the chord have been shown to provide different power benefits at different forward flight speeds, suggesting benefits for change the extendable chord based on advance ratio [132]. Such

extendable trailing edges were predicted to increase power at low speeds and loading conditions, but decrease required power under higher loads and altitudes, resulting in faster maximum speeds and loads [81]. Active trailing tabs have also been modeled to help balance slightly unstable rotor blade motion for the main rotors in the UH-60, reducing in-plane forces at hover and low speed flight and all forces at higher forward velocities [87]. Required control models were also developed through this research to simulate the necessary active controls [88]. One simulation considered a combination of 2-5 trailing edge flap control inputs per revolution, at best reducing power by 1.9% and vibratory loads by 70% [133].

The bistable actuation mechanisms developed to enable rotor blade chord extension have also been developed for rotor blade span extension [117]. The mechanism holding the rotor extension in its shorter configuration serves as a biased load, such that when the rotor rpm is increased beyond a critical threshold the system will overcome the biased mechanism and the rotor will extend and increase span due to the increased rotor inertial forces at higher rpm. Another work considering reconfigurable rotor morphing evaluates variable span, rpm, and camber using a quasi-static analysis [118].

1.3.5 Non-Rotor Adaptations

Active systems do not have to be located directly on the rotor blade to improve rotor performance. One active method typically used for vibration and noise reduction is higher harmonic control (HHC) [140, 80]. Higher harmonic control systems have been developed and tested [141, 124] to adjust swash plate oscillations with hydraulic actuators to reduce BVI [69]. HHC reduced BVI on an XV-15 main rotor tested in the large NASA Ames wind tunnel up to 12 dB in peak noise reduction [123].

These actuation systems often replace connections between the swash plate and rotor blade to create active pilot inputs beyond a typical rotor blade swash plate for specific objectives. Such HHC inputs have been shown to dramatically lower vertical vibratory loads, particularly at higher forward speeds and advance ratios [49]. Piezoelectric actuators have also been combined with hydraulic systems to combine the larger extensions of hydraulic systems with the faster piezoelectric

actuation [143].

1.4 Aerodynamic Computational Modeling

1.4.1 Local Subsystem Aerodynamic Characterization

Adaptive geometries driven by structural morphing must be related to the rotorcraft aerodynamic response. Vehicle aerodynamic response can be related to parameterized morphing by first relating local subsystem aerodynamic responses to local geometries based on parameterized shapes whether they are morphed or not. Total vehicle response can be determined only after all such local responses are calculated and integrated together.

To date, the bulk of literature published on morphing rotorcraft design focuses on the relationship between local (subsystem/subcomponents) shape [150, 135, 28] and vehicle performance [90, 79]. Work on morphing rotorcraft, and adaptive structures more generally, has been performed to relate local morphing to changes in vehicle performance at the design stage. Morphing design and analysis typically begins with the subsystem level (i.e. individual local components such as the main rotor, tail rotor, etc.). The aerodynamic effects of altering geometry and structure must be quantified locally before evaluating vehicle-wide performance.

Significant existing work, including the work described herein, on subsystem morphing has demonstrated the effects of local alterations for increased lift [41], lower loading, and other changes to structural and aerodynamic response. Demonstrating adaptivity may be valuable in its own regard, since adaptive technologies can be used in situations tailored for these specific advantages when competing objectives are determined early in the design process.

Existing experimental data provides a resource-efficient but often limited aerodynamic characterization. This data may have limited applicability beyond the geometries and flight conditions tested. Existing aerodynamic data can provide an efficient though limited set of known designs for optimization when parameterized then interpolated. While using existing interpolation tables requires capturing the entire range of aerodynamic responses to be evaluated, such tables can be generated with ranges informed by knowledge of the expected aerodynamic environment for a

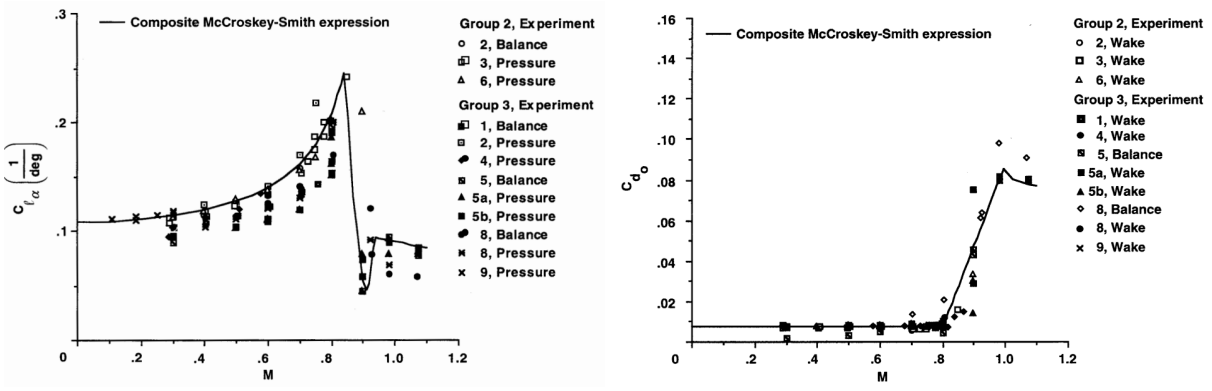

Figure 1.6: Experimental data can be used to relate aerodynamic environments with local airfoil lift and drag properties.

given subsystem. An example of experimental data to characterize the SC-1095 airfoil used on the UH-60 main rotor at various Mach numbers and is shown in Fig. 1.6 [156]. Numerous tools have been demonstrated for effective aerodynamic characterization at the subsystem level. In analyzing the overall effects of local rotor blade changes, 2-D methods are often employed, neglecting or simplifying 3-D effects to increase computational efficiency. Thus, 3-D effects are often approximated as a series of continuous 2-D local effects, though XFLR5 is also used to consider 3D effects similarly to XFOIL [2]. These 2-D aerodynamic tools used for rotor blades include JBlade [120] and XFOIL [43], as well as empirical solutions for existing airfoils.

Low-fidelity panel methods such as XFOIL allow for wide, generalized geometric ranges, but with less accuracy and more convergence difficulties for the some unconventional geometric shapes and aerodynamic environments [89]. Other tools have been developed to define airfoil shapes using parameterized functions such as class/shape transformations (CSTs) [33] then determine the aerodynamic characteristics using low-fidelity tools such as XFOIL [95, 96, 97]. Improved aerodynamic properties such as increased lift-to-drag ratio were determined for different adaptive NACA airfoils depending on the required angle of attack and airspeed [166]. These tools enable relatively cost efficient geometric parameterized aerodynamic structural optimization.

More work has been developed relating adaptive geometric shapes with higher-quality aero-

dynamic responses, often via computational fluid dynamics (CFD) [63, 46]. Other parameterized interpolation functions such as Bezier curve parameterization and Non-Uniform Rational B-splines (NURBS) have been used with higher fidelity CFD to determine optimal airfoil shapes for given flight conditions [174]. CFD typically requires much more time and computational cost to converge than lower-fidelity panel methods, but can more accurately determine pressure fields and the corresponding lift, drag, and moments on an airfoil [147]. These higher fidelity computational methods are particularly valuable at high Reynolds numbers and angles of attack beyond the stall condition by considering flow separation and viscous effects.

Different techniques have been developed for local aerodynamic analysis of varying computational cost and fidelity [162]. These tools have different strengths and drawbacks; tool selection depends on the specific application. Through computational aerodynamic modeling, most shapes and surfaces can be accurately characterized at the local level, which is invaluable for airfoil and rotor design. This local characterization is necessary but insufficient to determine overall vehicle performance.

1.4.2 Rotorcraft Aerodynamic Modeling

Traditional morphing research and development often follows this pattern: establish local morphing response, evaluate the adaptive subsystem under a specified configuration as part of the vehicle, then exhibit when and where morphing is advantageous. Sometimes only subsystem analysis is required; increasing lift to drag ratio over a certain angle of attack range for fixed wing vehicles may all but guarantee increased performance. However, rotorcraft morphing analysis may require more comprehensive study due to additional dynamic responses not typically encountered by fixed lifting surfaces. Inflow may vary dramatically depending on the rotation angle due to rotor blade rotation, flapping, and damping. Fig. 1.7 shows the various different higher-order effects not seen for fixed lifting surfaces. The contour represents local incoming velocity for the rotor blade imposed on the rotor disk. Two important areas are highlighted: regions of reverse and transonic flow; these two areas are important to account for accurately. Reverse flow occurs in high speed forward flight on the retreating (seen in Fig 1.7 as the left side of the rotor disk) region of the

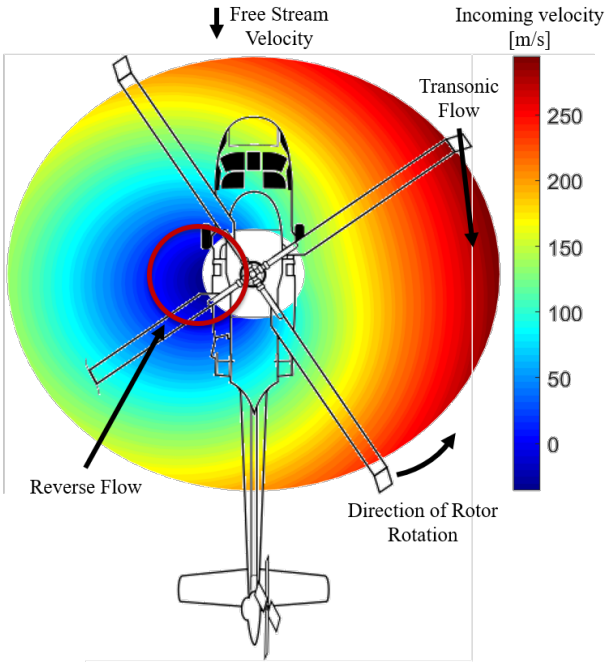

Figure 1.7: Local blade incoming velocity in forward flight.

rotor blade because the rotor is moving the same direction as the incoming velocity. This aerodynamic response contrasts fixed wing aircraft and most other rotorcraft flight conditions, leading to large angles of attack well beyond the stall flight condition. Another region of note is where the local airflow, due to rotor rotation, reaches much higher velocities than the rotorcraft. These high velocity regions on the outer edge of the advancing region of the rotor blade will experience transonic flow, where significant viscous effects appear. These effects may be further compounded when considering the changes in pilots inputs to maintain trim, or steady, flight. Thus, a subsystem morphing-enabled parameter change such as local lift requires analysis of the entire vehicle to evaluate benefits and trade-offs in trim flight.

1.4.3 Rotorcraft Vehicle Performance Modeling

A comprehensive analysis tool is required to determine the trim (level and steady) flight condition given a set of design variables and requirements. Such a tool considers the entire vehicle geometry and aerodynamic environment. Parameters relating adaptive rotor blades with overall

Table 1.1: Preliminary list of vehicle parameters.

Category	**Parameter**
Rotor	Chord
	Airfoil Section
	Twist
	Span
	Camber
Mission	Fixed Speed
	Fuel Load
	Altitude
	Stage Time
	Payload
	Optimal Speed
Actuation	Method
	Rate
	Location
Vehicle	Turbine Power
	Configuration
	Sizing

mission performance include rotor geometry, mission definitions, actuation geometry, and vehicle geometric parameters as seen in Table 1.1. Comprehensive rotorcraft analysis tools resolve local aerodynamic effects to trim the vehicle. As computational power has increased, so have rotorcraft analysis efforts and their aerodynamic tool fidelity. Significant existing work as been conducted towards developing and validating the comprehensive rotorcraft tools used today, including CAMRAD II, CHARM, UMARC and NDARC [133, 35, 21, 75, 76]. These tools all considered, among other rotorcraft, the UH-60, which is also used in this work due to the public availability of existing experimental test data and high-fidelity computational results. While the previously mentioned tools focus on (but are not limited to) power, other computational frameworks have been developed for noise reduction [138, 111]. One such framework utilizes CAMRAD-II for the rotor trim modeling combined with ANSYS to noise modeling to determine the noise profiles [83].

In mission-driven rotorcraft design the effects of each subsystem are interdependent. For exam-

ple, changes in main rotor lift and drag may lead to changes in the pilot tail rotor controls. Steady level (trim) flight requires balancing the loads and moments on the vehicle based on the pilot inputs. Thus, comprehensive rotorcraft analysis tools will be utilized to determine trim flight under specified mission definitions. RCAS [19], for example, considers a finite element structural model, an array of airloads models, induced velocity models, and controls models together. This computational framework can consider trim analysis (either static equilibrium or periodic steady-state), maneuver analysis, and a number of stability analysis (such as linearization and model reduction). Another comprehensive rotorcraft analysis code, TRAC, can perform trim analysis, extract linear models, and simulate flights as a complete helicopter model [99]. A more general comparison between different aerodynamic rotorcraft modeling methods can be seen in Table 1.2. These computational resources may not necessarily be selected over one another, but rather in conjunction depending on the specific modeling needs.

1.5 Design and Optimization

Accurate rotorcraft modeling enables comprehensive numerical tools for both analysis and design, but these should be deployed in an efficient manner. In particular, determining the optimal design variables based on a given objective or set of objectives requires substantially more computational resources than a single analysis. There exist several methods of relating variables at the design stage to objectives determined by the vehicle modeling efforts.

A common optimization formulation is defined by an objective function, constraints, and bounds [24]. A single-objective design optimization problem is defined as:

$$\min_{\underline{x}} f(\underline{x}), \tag{1.1}$$

subject to inequality and equality constraints:

$$\begin{aligned} g(\underline{x}) &\leq 0, \\ h(\underline{x}) &= 0. \end{aligned} \tag{1.2}$$

Table 1.2: Comparison of Various Rotorcraft Aerodynamic Solvers

	Aerodynamic Solvers	**Comprehensive Analysis Tools**	**Physics Solvers**
Examples	• MachUp • XFLR5 • XFOIL	• TRAC • UMARC • RCAS	• CFD • COMSOL
Pros	• Fast runtimes (< 1 second) • Relatively accurate at low Reynolds numbers • Solves for propellers, wings, and simple geometries	• Can have relatively fast runtimes (with low-fidelity inputs < 5 min) • Solves for trim condition • Evaluates aircraft geometry beyond lifting surfaces • Can be coupled with structural analysis tools	• High fidelity • Robust • Can solve for a large number of geometries • Can solve static or dynamic states • Can be coupled with structural FEA
Cons	• Limited geometries • Requires airfoil-specific constants. • These constants may not be available, especially when morphing the outer mold line • Limited variables to adjust beyond lifting line parameters	• May ignore effects between separate systems (interference, etc.) • Solves as a steady state problem instead of dynamic • Limited to known or pre-computed rotor geometries	• Slow runtimes • Longer to set up • May require prior analysis to determine boundary conditions

Here $\underline{x}$ is a vector $\{x_1...x_n\}$ denoting the *optimization variable* for the *objective function* $f(\underline{x})$. $g(\underline{x})$ and $h(\underline{x})$ are inequality and equality functions used to define the constraints in equation 1.2. In a single-objective optimization (i.e., $f(\underline{x})$ is a single function or evaluation of the optimization variable $\underline{x}$ resulting in a single output per evaluation) the final solution to equation 1.1 is a single $\underline{x}$, in the context of this work typically a set of design variables. The best design, or solution to this problem is defined as J^*:

$$J^* = \min_{\underline{x}} f(\underline{x}). \tag{1.3}$$

This single-objective optimization can be extended to multiple objectives whereby a single score is defined as some weighted sum of individual objectives. Trade-offs between the multiple objectives could be further explored by altering the individual weights for each objective. However, multiple objectives can be considered during the optimization based on trade-offs and relative dominance in the objective space, as will be discussed next.

1.5.1 Multi-Objective Design and Optimization Methods

In multi-objective analysis, the best set of non-dominated designs represents trade-offs between competing objectives. In traditional multi-objective analysis, the existence of trade-offs between competing objectives means near-optimal designs cannot be fully realized. Often an optimization is performed to determine the non-dominated designs, designated Pareto-optimal. Non-dominated in this respect refers to designs wherein, for n objectives, "it is one from which it is impossible to depart without making at least one of the n objectives worse (i.e., larger)[146]." This set of non-dominated designs is referred to as the Pareto frontier, and shows trade-offs between competing objectives. An example Pareto frontier for minimizing two objectives is shown in Fig. 1.8; for every design on the frontier no other design has a lower Objective 1 *and* Objective 2.

Population-based optimization methods explore the design space by altering input variables (typically engineering design variables) of individuals (designs) over iterations or time-like increments (sets of designs) to determine the preferred variables given specified objectives. The entire set of designs in a given iteration or increment, designated the population, changes based on the

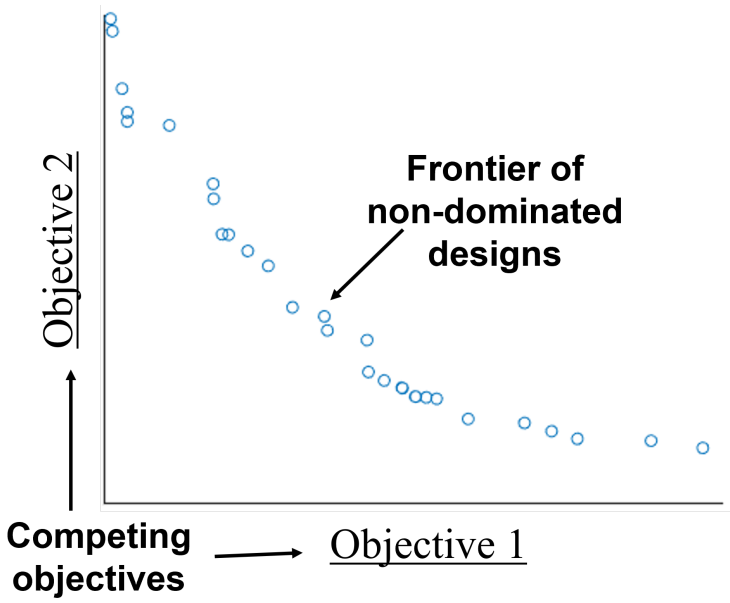

Figure 1.8: The non-dominated designs in the objective space represent the best known trade-offs between competing objectives. No design in such a set is worse than another across every objective.

response of the previous generations.

Genetic algorithms (GAs) are commonly applied to complex engineering design problems, proposed in the early 1970's as computer programs which mimic natural evolutionary processes [148, 67, 47]. These optimization methods were developed based on encoding individual properties (parameters) as genes, with individual designs part of a larger population. This larger population is evaluated based on objective values and 'evolved' over generations, favoring individuals exhibiting preferred objectives. This evolution process crosses individual genes and introduces mutations to increase genetic diversity.

NSGA-II [37], a genetic algorithm, explores these trade-offs by encouraging designs spanning the non-dominated objective space. Other population-based optimization methods used for rotorcraft design exploration include particle swarm optimizations (PSOs) [161], which consider the individual designs in the design space as locations, updating positions and velocities based on the best performance of each individual and the entire population. Particle swarm optimizations have

been used in design and control optimization for rotorcraft missions [52].

More efficient trade-offs between competing objectives can be accomplished by morphing certain vehicle components. In an aerospace engineering context, changes in shape allow for different objectives to be reached at different times and under different flight conditions. Altering vehicle geometry to match changes in forward flight speed, loading requirements, pilot maneuvers, and even environmental conditions can reduce trade-offs between performance metrics [5]. Rotorcraft in particular see large changes in flight conditions due to differences between hover and forward flight, as will be demonstrated later.

Optimization methods such as genetic algorithms require careful parameterization and tuning. Work has been performed to characterize the proper methods for selecting 'good' designs and design variables in genetic optimization frameworks. One popular genetic optimization method is to encourage both genetic diversity and elitism when selecting designs. For multi-objective design optimization, the Non-dominated Sorting Genetic Algorithm-II (NSGA-II) [37] weighs both the best sets of non-dominated designs, but also the diversity of the objective space. This work was extended for larger objective spaces as NSGA-III, using reference points for comparison to reduce computational costs [39]. When considering multiple objectives, these methods rank designs based on frontiers of non-dominated designs. The first rank of non-dominated designs are identified as 'Rank 1', then the next frontier designated 'Rank 2', and so forth, as can be seen in Fig. 1.9.

NSGA-II prioritizes designs with better rank to improve sets of trade-offs between objectives, then selecting designs with larger empty objective spacing to promote diversity. Since its original implementation, NSGA-II and other genetic algorithms have been combined with adaptive mutation mechanisms [27, 177] and other selectors [102, 30, 158] to improve convergence and diversity for specific applications.

Other multi-objective population-based selection criteria include the Strength Pareto Evolutionary Algorithm (SPEA) [178, 180] and its updated version SPEA-2 [179], which maintain an archive of non-dominated designs and then truncate sets of equal rank based on the distance between individuals. Both of these algorithms have been applied to numerous multi-objective

test [38] and design [57, 78] problems. It is worth noting that computational design tools require

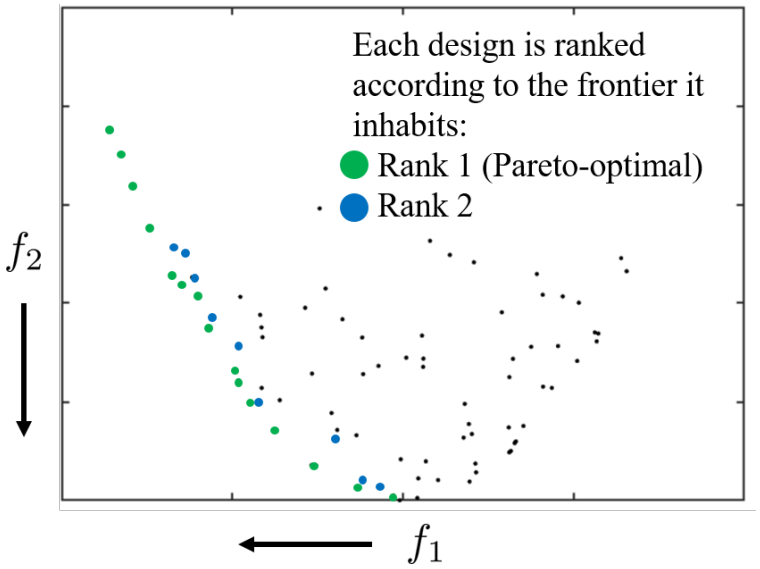

Figure 1.9: Designs can be ranked by non-dominated frontiers in the objective space, signifying the best trade-offs in the multi-objective space.

optimization parameters such as selection criteria, population size, convergence criteria, etc. The selection of such parameters is highly problem-dependent, and work has been performed to better optimize engineering design problems using these tools [71, 42, 72]. Often the selection of parameters which do not necessarily need to be minimized is of notable importance, and recent work have been performed to expand the engineering design problem to parameterized optimization [50, 166].

With increased computational power , numerical rotorcraft design and optimization frameworks have been developed for a number of purposes. Mission-driven objectives such as fuel burn, noise signature, and emissions have been incorporated into comprehensive analysis and design tools [59, 146]. Other trade-offs such as resiliency vs efficiency have been explored using ecosystems as a natural inspiration [32]. Some recent computational frameworks include adaptive technologies, which increases optimization with the addition of different design variables for different mission stages or objectives [5]. Much of this work is recent due to the high computational

cost of modeling and evaluating complex aerostructural responses. Mission-based optimization has also extended to flight planning [52, 164, 171], with specific requirements for the mission but flexibility in the path selection.

1.5.2 Mission-Driven Design and Optimization

Mission-driven design extends beyond a traditional engineering multi-objective design formulation, in that the objectives in different mission stages are not necessarily independent of each other. Missions also have unique qualities present in many, but not all multi-objective engineering design problems. An example mission can be seen in Fig. 1.10 [5]. For instance, mission stages may occur over a range of different flight conditions. In the aerospace engineering design context, changes in external environmental conditions lead to changes in air density, temperature, altitude, pressure, etc. Some mission stages may contain similar flight conditions, but other parameters such as weight may change due to fuel burn, ablation, etc. even if mission stages, objectives, and requirements remain constant. Sometimes these changes may be considered insignificant compared to other changes between mission stages, but such effects are generally initially considered, even if then dismissed.

A given mission may have specific requirements between mission stages, or require trim (steady, constant attitude) flight. A wide range of requirements and objectives can define a mission, such as distance covered or trim flight at a given forward velocity or climb rate. As previously mentioned, different mission stages may compete depending on the mission objective(s). If an objective is to minimize fuel burn, the best designs at different stages may compete in an optimization framework, depending on design problem definition. Also, a mission stage can be defined as a given time or distance. This in turn may lead to different preferred designs depending on if the mission objective is speed or range. Other objectives such as fastest climb will likewise compete with efficiency requirements.

Mission definitions are traditionally sequential; a series of stages in pre-defined order. By this formulation, the required pilot inputs and controls can be determined for each stage, reducing the design problem complexity from continuous to discrete. However, this does not consider transi-

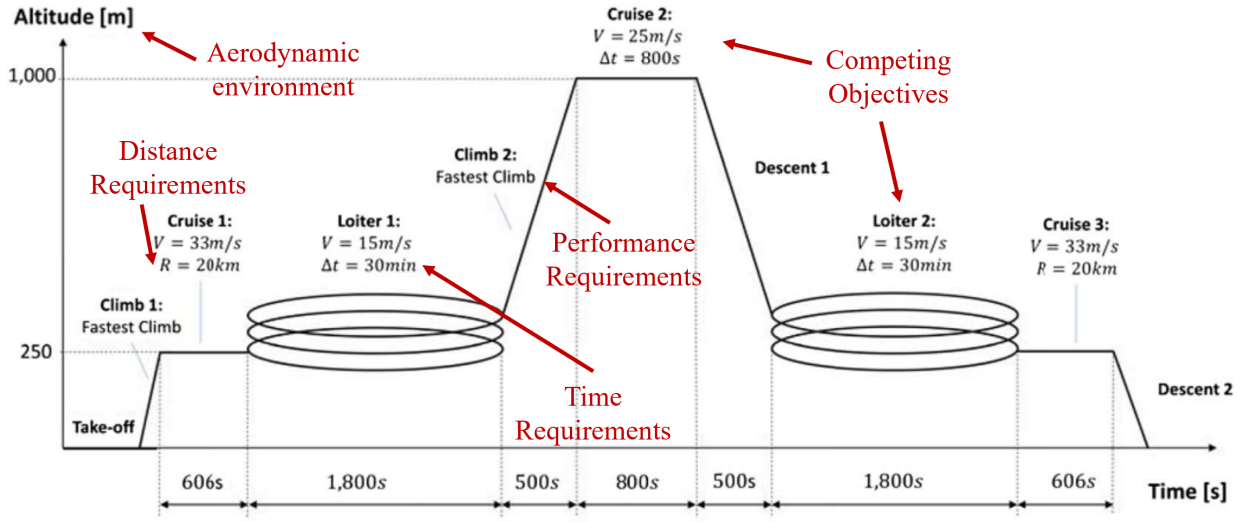

Figure 1.10: A complete mission definition considers changes in aerodynamic environments, requirements, and objectives.

tions between mission stages, which require increased computational costs [7]. Such transients may also be considered, but at a considerable cost when compared with a set of discrete calculations equal to, at most, the total number of mission stages. Often transients are considered in the engineering design problem after initial structures have been determined, with a focus on optimal control given a final aerostructural design. Even when considering optimal control airspeeds for a given structural design, often the optimization problem is confined to discrete mission stages and yields valuable solutions [3]. Work has also been performed to determine optimal control across multiple missions based on pilot inputs, though in [62] these missions were limited to single, averaged, and discrete stages with a constant prescribed vehicle structural definition.

The transient 'optimal control' problem for a given mission increases the complexity beyond some traditional mission definitions, for structural design but particularly adaptive structures design. Much of the work with transient design occurs after initial structural design, with a focus on control system design after initial geometries are selected [9], while others may consider other active or adaptive systems and some other design variables that could be considered in a limited multi-disciplinary design optimization such as engine sizing [127, 6].

1.6 Chapter 1 Summary

In this chapter, mission-driven adaptive aircraft (and more specifically adaptive rotorcraft) design and optimization is introduced and the existing research explored. A large number of adaptive rotorcraft techniques have been developed via modeling, fabrication, and experimental testing. Concerning rotorcraft modeling, a number of numerical methods have been developed over a range of accuracies and computational costs. Trade-offs between accuracy and resources must be considered, especially in early design stages. Both the local rotor blade aerodynamics and the entire vehicle performance for trim, or constant attitude, flight must be determined. Additionally, large scale optimizations are often used to determine optimal input parameters, both those that can and cannot morph, for overall mission performance. Missions have been modeled both as a sequential set of stages or as a set of multiple objectives during different flight conditions.

This work builds on existing design and optimization techniques to consider the differences between adaptive and fixed design variables to reduce the design space and determine optimal adaptive solutions using the rank and dominance methods discussed in this chapter. These optimization techniques are applied to mission-based adaptive rotorcraft design and optimization, which requires numerous rotorcraft modeling techniques relating design inputs to vehicle performance over the course of a mission. Several levels of fidelity are considered while modeling a wide range of mission stage types to evaluate a wide range of adaptive rotorcraft technologies.

A comprehensive mission-driven adaptive rotorcraft framework was developed in this work encapsulating modeling, mission analysis, and multi-objective design and optimization. The framework is modular and general, but applied specifically to the UH-60 Blackhawk due to existing experimental and numerical data provided by the NASA UH-60 airloads program. Many of the missions considered in this work are selected from Future Vertical Lift proposed missions as to be realistic and of future interest [66].

A wide range of adaptive technologies can be explored for mission-driven configurations early in the design process with effective aircraft modeling. However, many existing design and optimization methods, including genetic optimization techniques, are not specifically designed for op-

timizations in which some design variables are adaptive while others are not. A common technique to utilize these tools is to consider each morphable design variable as a different design variable for each state or objective, which can increase the optimization problem complexity. A novel optimization technique is developed and evaluated in this work for mission-driven, multi-objective design of adaptive structures. By considering design variables that can and cannot adapt differently in the optimization problem, preferred adaptive configurations can be determined in the objective space and evaluated based on multi-objective performance and feasibility across adaptive designs. It will be shown that this design space decomposition can be used both as a post-processing and a selection technique during multi-objective genetic optimizations.

1.7 Dissertation Summary

In summary, this work is organized as follows:

- Chapter 2 provides an overview of the numerical and computational methods used to model the rotorcraft aerodynamics considering adaptive rotor blades. This chapter describes the implementation of structural and aerodynamic computational tools of various fidelities and computational costs to capture the physics from local rotor blade sections to larger vehicle-wide analysis under trim flight, including validation between existing computational and experimental data for the UH-60 Blackhawk.

- Chapter 3 describes the most novel aspect of this dissertation: a decomposition of the engineering design space into fixed and adaptive design variables for multi-objective optimization. Methods for both optimizing over such a combined design space and the development of adaptivity as an evolutionary selection criteria are developed, evaluated, and discussed here.

- Chapter 4 applies the low-fidelity methods from Chapter 2 toward a comprehensive mission-driven adaptive rotorcraft computational framework, which is applied to several adaptive rotorcraft technologies from Chapter 1. Then the design space composition outlined in Chapter 3 is applied to multi-objective optimization results to determine feasible adaptive

designs for improved mission performance. An alternative mission-driven adaptive design and optimization approach via sequential overall mission scoring is introduced and evaluated here.

- Chapter 5 applies the computational framework developed in Chapter 4 to higher-fidelity methods from Chapter 2 and extends the mission-driven framework beyond single missions towards more general adaptivity. Multiple missions are considered for the same rotor blade developed either via mission scoring or multi-objective optimization. Comparisons between the design space decomposition approach from Chapter 3 and existing multi-objective adaptive optimization approaches are evaluated for mission-driven adaptive rotorcraft design. Aerostructural coupling will also be introduced and evaluated.

- Chapter 6 details the development of a camber morphing optimized rotor blade created from the results of aerodynamic and structural optimizations to develop a 2-D prototype. Internal structural topologies are determined via genetic optimizations and then sizing optimizations determine a morphing rotor blade design to match mission-driven preferred geometries. The rotor blade is then tested to evaluate the structural morphing mechanisms developed to match the mission-optimized blade. The outer mold line shapes are compared between the aerodynamic optimization, structural optimization, and actuated prototype.

- Chapter 7 summarizes the entire dissertation. General conclusions, contributions, and future work are discussed.

2. ADAPTIVE ROTORCRAFT MODELING

Aerodynamic engineering design requires capturing and predicting both the structural and aerodynamic vehicle responses. Mission-driven design requires these models accommodate a range of requirements and environments. Further, adaptivity introduces different geometric configurations at different times during a mission. Thus, efficient modeling is required to consider relating a range of adaptive geometries to mission-wide performance.

In this work a range of modeling tools are utilized and developed to determine local aerodynamic effects, trim flight solutions under aerodynamic environments defined by mission stages, changes in aerodynamic properties due to adaptive structures, and coupled aero-elastic effects. Trade-offs between model fidelity and computational costs are considered; a range of aerodynamic tools of varying fidelities and computational costs are evaluated.

The design and optimization of adaptive structures can require significant evaluations to determine optimal design parameters, which in turn requires efficient rotorcraft performance assessment. These adaptive rotorcraft modeling tools are all required to relate adaptive structures at the rotor blade level to overall mission performance in a computationally efficient manner. The work relates rotorcraft subsystems, currently focusing on main rotor blade contributions, to overall vehicle performance under trim flight during mission stages.

Chapter 2 Outline

This chapter is divided into 4 sections:

- In section 2.1, local airfoil characterization methods are analyzed relating rotor blade geometries to aerodynamic properties. This work considers several models to determine lift, drag, and moments based on local rotor geometries using tools ranging from linear aerodynamic properties and panel methods to computational fluid dynamics.
- In section 2.2, methods to characterize changes in aerodynamic properties via rotor blade morphing are developed considering adaptive parameterized trailing edge camber, twist, and outer mold lines.
- In section 2.3, two methods of vehicle aerodynamic trim flight analysis are considered, a hover BEMT code developed in this work and an existing comprehensive rotorcraft analysis code for forward flight.
- In section 2.4, an efficient method for fully coupling aerodynamic and structural responses is developed based on the relative speeds of convergence for both structural and aerodynamic analysis.

2.1 Rotor Blade Aerodynamic Characterization

Rotor blade aerodynamic properties must be specified or determined before any full-vehicle aerodynamic analysis is possible. Aerodynamic properties such as lift and drag can be defined as functions of local geometries and environmental parameters such as position, attitude, velocity, and air properties. Changes in geometries such as twist, chord, and camber must also be considered for local airfoil characterization. The methods of rotor blade aerodynamic characterization discussed in section 1.4.1 are applied in this work to consider multiple techniques of various fidelity and applicability while considering geometries to be altered during flight.

2.1.1 Linear Aerodynamic Properties

One of the simplest relations between lift and drag is a linear relationship between angle of attack and lift, and a low drag coefficient. This lift and drag profile can be seen in Fig. 2.1. This method also includes a Prandtl-Glauert correction factor for 2D airfoil sections to account for compressibility [64, 101, 157]:

$$C_P = \frac{c_{p_0}}{\beta}, \tag{2.1}$$

$$C_l = \frac{c_{l_0}}{\beta}, \tag{2.2}$$

$$C_m = \frac{c_{m_0}}{\beta}, \tag{2.3}$$

where:

$$\beta \equiv \sqrt{1 - M^2}. \tag{2.4}$$

This method yields the fastest computations but loses accuracy beyond the linear regions of lift, such as at high angles of attack beyond stall, and through transonic flow ($M > 0.7$) where the Prandtl-Glauert correction loses accuracy, where:

$$as\ M \longrightarrow 1,\ \beta \longrightarrow 0 \text{ and } C_l \longrightarrow \infty. \tag{2.5}$$

This aerodynamic characterization method is useful for low angles of attack and low Reynolds numbers, but is limited outside of these regions.

Only considering the changes in a lift curve slope and constant drag also limits applicability to capturing the effects of changing the rotor blade airfoil shape. Since altering the discretized twist schedule of the rotor blade does not necessarily change the 2-D shape, only the local angle of attack, twist morphing would still be feasible using this method for determining local lift and drag within limited angle of attack bounds. It will be shown that this method is insufficient for many missions and morphing applications, over-predicting performance for higher loading conditions.

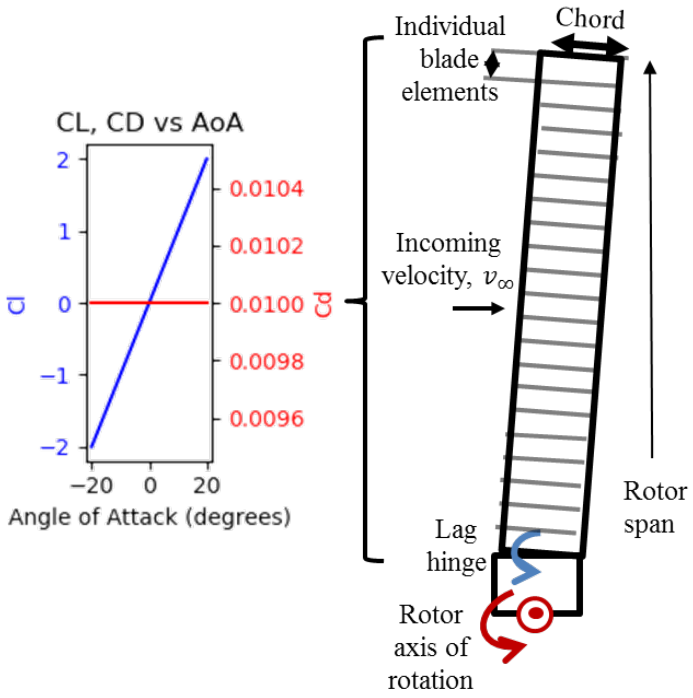

Figure 2.1: The simplest but most limited method of airfoil aerodynamic characterization: a linear lift curve slope with a constant drag. Each rotor blade element lift and drag is determined via two constants specific to that blade element.

2.1.2 Low-Fidelity Panel Methods

A more robust approach was developed to consider a wider range of adaptive rotor blade geometries and an increased range of aerodynamic effects. This approach utilized XFOIL [43] to generate large interpolation tables of lift, drag, and moments as a function of Mach number and angle of attack. A single, large interpolation table was generated for a given shape. When given the local angle of attack, air properties, and incoming velocity, these interpolation tables determine the the lift, drag, and moments for each blade element. An example of the lift, and drag profiles generated for the SC-1095 airfoil used on the UH-60 can be seen in Fig. 2.2. These tables were generated in Python [160] via an XFOIL interface developed in the Python library AeroPy [97] and applied to fixed-wing airfoil design and optimization [96]. Evaluating the airfoil shapes individually during an optimization increases the runtime, but not significantly. This method also

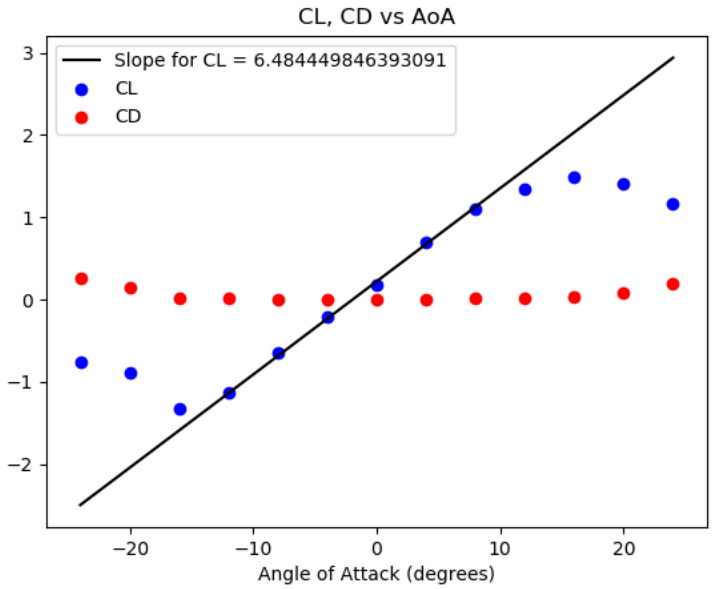

Figure 2.2: Interpolation tables are generated across a range of angles of attack and Mach numbers.

allows for evaluating any airfoil shape as a closed set of points.

XFOIL tables increase the accuracy of local lift, drag, and moments for complex airfoil profiles, particularly near the stall condition, but were limited to angles of attack and Reynolds numbers with converged solutions, which is not guaranteed. XFOIL can consider boundary layer effects and separation, but not in the detail of a Navier-Stokes based method. Beyond stall, XFOIL may lose accuracy or have trouble converging, limiting interpolation table ranges. In this implementation, local blade elements are considered to generate zero lift and the maximum drag beyond stall and XFOIL convergence ranges. Higher fidelity tools are needed to generate tables that encapsulate regions well beyond stall, as well as trans-sonic effects near the rotor blade tip in forward flight. Higher fidelity tools such as computational fluid dynamics (CFD) typically require more input variables than XFOIL and take much longer to converge due to solving full-field equations. In the case of a static rotor blade OML, i.e., no adaptive geometries, CFD tables could be created separately for the current rotor blade shapes, and then incorporated into the lower fidelity tools as a replacement for XFOIL. However, considering adaptive geometries in CFD may be cost-

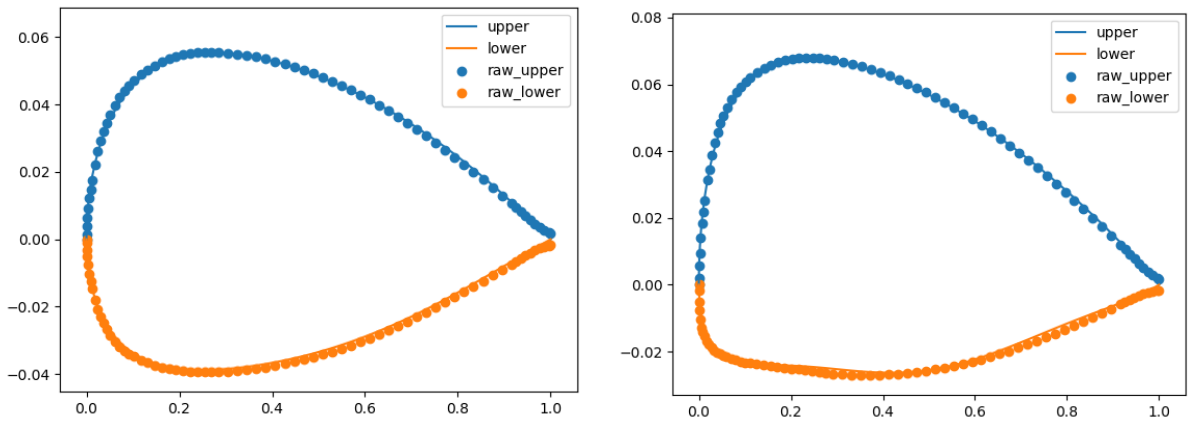

Figure 2.3: The 4-parameter CST equations were able to accurately represent both the SC-1095 (left) and SC-1094 r8 (right) airfoils.

prohibitive; XFOIL could generate interpolation tables for each geometry in minutes while it may takes days, weeks, or months to generate the same data using CFD.

Since XFOIL can consider any closed set of points, any parameterized design optimization must relate a limited set of shape parameters into a full airfoil outer mold line (OML). Class/shape transformation equations (CSTs) were used to parameterize the airfoil OML, as they have been shown to accurately represent realistic airfoil geometries with a minimum number of variables [89, 166].

A four-parameter CST equation was selected in an effort to both limit the number of parameters necessary to capture conventional airfoil shapes and to accurately model existing airfoils. A minimization procedure developed in [97] was used to minimize the mean squared error between the existing UH-60 airfoils (SC-1095 and SC-1094 r8) by changing the CST parameters. These CST parameters then defined the reference, unmorphed rotor blade configuration and a discretized point-based OML could be defined and used in XFOIL to generate aerodynamic properties. A comparison between the airfoils and 4-parameter CST approximations can be seen in Fig. 2.3. A report on the aerodynamic characteristics for the two rotor blades pictured provides values for the point-wise shape profiles [22].

2.2 Adaptive Rotor Blade Modeling

The rotorcraft modeling methods developed and applied in this work can determine the trim flight condition for a rotorcraft in hover and forward flight when given the aerodynamic properties for the main rotor blade. Other contributions such as tail rotor contributions fuselage drag, and even additional lifting surfaces (side wings, etc.) can be considered as additions to the trim flight optimization considering the additional respective pilot inputs.

However, each of these contributions require other specified properties to model these geometric and aerodynamic effects accurately while morphing. Functions such as $C_l = f(\alpha, M)$ (equation 2.15) require the methods described in this section 2.1 applied to changing design variables in a computationally efficient manner.

2.2.1 Parameterized Adaptive Camber Morphing

Parameterized adaptive rotor blade methods can take many forms, two examples being bistable actuation mechanisms [117, 10, 11] and trailing edge flaps [135, 116]. Piezo-electric elements have been used to actively twist the rotor blade [28], while shape memory alloy (SMA) torque tubes have been used to deploy noise reduction tabs [108]. These adaptive technologies offer a large potential to improve vehicle performance if properly applied.

One such geometric rotor blade morphing method is the smooth alteration of the outer mold line of the rotorcraft via internal actuators. Proper parameterization of such airfoil shape morphing (e.g., by limiting deflections to particular forms and regions, such as altering trailing edge camber) imposes limits on more realizable geometries, which can reduce the design space and computational costs in multi-objective optimizations. Changes in rotor blade shape alter lift, drag, and moments generated given a specified local angle of attack and Mach number. These changes in aerodynamic properties alter the induced inflow generated by the rotor blade, particularly at hover. Thus, the effects of morphing the rotor blade by changing the local camber on overall vehicle performance must be determined iteratively until convergence with the complete vehicle analysis.

Computational fluid dynamics (CFD) (i.e., FUN3D [15]) tools were used to evaluate steady-

state solutions for the UH-60 main rotor airfoil SC-1095 over a wide range of angles of attack, Reynold's numbers, and levels of parameterized trailing edge camber morphing. The converged pressure fields were integrated over the airfoils to determine the local coefficients of lift, drag, and moment. The ranges were informed by previous analysis of trim flight conditions for the UH-60, and the expected angle of attack and Reynolds number ranges observed at different locations down the rotor span. Span-wise local lift, drag, and moment coefficients from the CFD results were saved as interpolation tables, which could be quickly accessed without further aerodynamic analysis.

The CFD tables were geometrically defined by altering the SC-1095 airfoil camber. The trailing edge camber was altered parametrically, considering morphing to both a higher camber and, conversely, a reflex shape. The trailing edge alterations generate several effects on the lift, drag, and moment profiles of the airfoil, most significantly shifting the relationship between local angle of attack, lift, and drag.

The dynamic state and rotorcraft attitude determine a local angle of attack and Mach number for discretized 2D blade elements, then local aerodynamic properties are determined via the CFD interpolation tables, as shown in Fig. 2.4. The ranges were informed by previous analysis of trim flight conditions for the UH-60, and the expected ranges of angle of attack and Mach number observed at different locations down the rotor span. The trailing edge camber of the SC-1095 airfoil was altered parametrically, considering morphing to both a higher and to a lower camber (reflex) shape. The trailing edge alterations significantly shift the relationship between local angle of attack, lift, and drag. An example of the lift profile of camber morphing data used in this work can be seen in Fig. 2.5. Morphing the trailing edge of the airfoil down shifts the lift profile such that the zero-lift angle of attack c_{l_0} (the angle of attack where lift is zero) decreases and the maximum lift coefficient $c_{l_{max}}$ increases.

It is worth noting the camber morphing considered here is quasi-static; morphing is considered between different mission stages or competing objectives but not on a per-revolution basis. High-frequency morphing technologies have been considered elsewhere for adaptive rotor blades on a per-revolution basis and cyclic inputs could be related to changing camber at different angles of

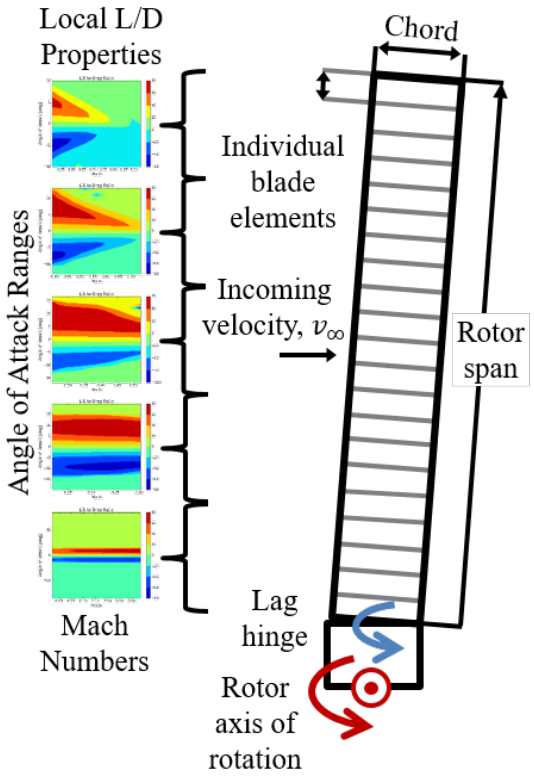

Figure 2.4: CFD interpolation tables associated with the aerodynamic conditions for each local blade element

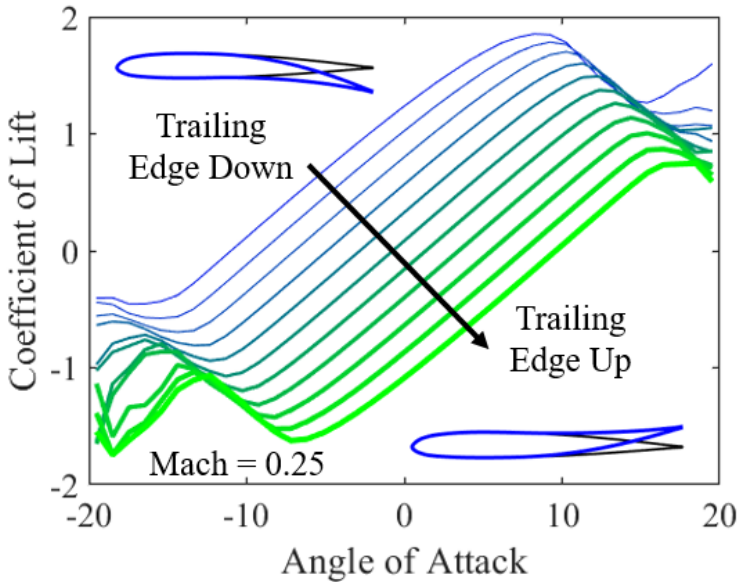

Figure 2.5: The effects of changing rotor camber localized near the trailing edge on the entire lift profile. Morphing the rotor blade trailing edge down results in lower stall angles and higher c_{l_0}.

rotation ψ, but that is beyond the scope of this work.

Minimizing total drag integrated over the rotor disk for trim flight given a specific aerodynamic environment reduces required power, a common performance objective. However, it is difficult to determine how local morphing in one rotor blade region will effect the overall integrated rotor disk considering the changes in pilot inputs required to trim the vehicle.

Level, steady flight requires pilot inputs adjusting main rotor pitch (cyclic and collective) to balance vehicle forces and moments appropriately. Determining the effect of local camber morphing on lift, drag, and moment locally requires solving for the pilot controls necessary for trim flight under aerodynamic conditions specified by mission stages. Thus, after the local morphing configuration is selected, the changes in aerodynamic properties are conferred to their respective local rotor blade regions as described in this section. Next, the vehicle aerodynamic modeling must incorporate the effects of morphing in calculating the trim flight condition.

2.2.2 Adaptive Rotor Blade Twist Morphing

Adaptive rotor morphing mechanisms have been developed and tested to alter the rotor blade twist schedule down the span (see section 1.3.2). A static twist schedule is typically designed to maintain the same or a similar inflow angle to maximize the lift-to-drag ratio down the rotor blade span. Such a manufactured twist schedule (and the overall schedule, which is also dependent on collective and cyclic pitching inputs as described in equation 2.19) is dependent on the flight condition. An adaptive rotor blade twisting mechanism seeks to change the twist schedule based on changing mission stages.

A simple adaptive twist analytical model is used in this work to based on specified torsional deflection. While useful for preliminary design and optimization, a full analysis could be performed with a complete finite element analysis method in later stages of detailed design and sizing. The torsional stiffness profile for the UH-60 rotor blade was taken from the RCAS implementation compared in [40]. This model considers the torsional stiffness of the UH-60 rotor blade but not the internal structure that ultimately creates the torsionally rigidity. Internal SMA actuators are installed inside the rotor blade and provide a torque that is resisted by the rotor blade. The SMA material has stress and strain limits, and is sized such that the inner radius is fixed at 0.6 the outer radius.

The relationship between twist (θ), torsional stiffness (GJ), torque (T), and the length down the rotor span (L) is given by:

$$\theta = \frac{T\ L}{GJ_{blade}}, \tag{2.6}$$

which can be rearranged to calculate the applied torque necessary to generate a specified change in twist via:

$$T = \frac{\theta\ GJ_{blade}}{L} \tag{2.7}$$

for a given region of the rotor blade. Then the internal SMA torque tubes are sized based on the material properties of NiTiHf. Since the ratio between torque tube inner and outer diameter is

specified, the minimum outer diameter required is:

$$OD_{SMA} = (\frac{16\,T}{\sigma^{y}{}_{SMA}\,\pi(1-0.6^4)})^{\frac{1}{3}}. \tag{2.8}$$

A constraint is applied limiting the outer diameter to a specified fraction of the total clearance inside the rotor blade at the spar location. Alternatively, the maximum strain, and thus the maximum twist, can be calculated in advance and used to set the bounds of any adaptive twist optimization studies. Given a maximum SMA stress, the maximum twist rate can be written:

$$\frac{\theta}{L}_{max} = \frac{\sigma^{y}{}_{SMA}\,{r_{clearance}}^3\,(1-0.6^4)\,\pi}{2\,GJ_{blade}}, \tag{2.9}$$

where $r_{clearance}$ is the maximum torque tube radius satisfying internal rotor blade clearance constraints. Setting the solution to equation 2.9 as the bounds for an adaptive twist design optimization problem ensures all solutions meet SMA yield stress and internal clearance requirements. The maximum stress and strain limits for the NiTiHf are shown in Table 2.1.

Table 2.1: NiTiHf Material Limits

Property	Value	Units
Maximum Yield Stress	275	MPa
Maximum Transformation Strain	0.02	

Multiple torque tubes could be introduced into the rotor blade, each changing the twist schedule for not just the installed region, but also all outboard regions. For a rotor blade consisting of 2D discrete elements, changes in twist schedule are easily implemented into rotorcraft trim analysis tools described next in section 2.3, since the only addition necessary to include adaptive twist is to update θ_{blade} from equation 2.36 to account for changes in manufactured twist due to adaptive structures such as SMA torque tubes. In this analytical structural model, the attachment of internal

SMA torque tubes is not considered, but the mass is added into the local rotor blade elements based on the sizing described in equation 2.9 or 2.8.

2.2.3 Adaptive Rotor Outer Mold Lines

Morphing the main rotor blade outer mold line (OML) while considering a single trailing edge parameter is valuable when generating high-fidelity but also computationally expensive aerodynamic data, but limits the range of explorable geometries. Wider ranges of OMLs can be defined and evaluated using an increased number of parameters, as described in section 1.4.1.

With class/shape transformations (CSTs), airfoil shapes can be altered beyond simply parameterized changes in trailing edge camber. Further, they can be altered from a parent (unmorphed or manufactured) shape to a child (morphed) shape as outlined in works by Leal and Hartl [95]. These transformation functions consider internal structural constraints applied as geometric restrictions, a method of considering internal structures such as spars between top and bottom surfaces. This formulation also ensures the distance between specified rotor blade geometries can be maintained through morphing.

Combined with a faster aerodynamic analysis tool, such as a panel method solver (for example, XFOIL or XFLR5), parameterized morphed OMLs can quickly be generated, then evaluated over angle of attack and Mach number ranges to generate 2-D interpolation tables for lift, drag, and moments in a manner not unlike those presented for CFD results. The key differences are two-fold: i) CFD solutions require orders of magnitude more computational costs compared to XFOIL evaluations and likely must be pre-computed instead of generated inside any optimization loop for feasibility of computational costs in a large design and optimization framework, and ii) the XFOIL interpolation tables are 2-D for a given CST shape (over angle of attack and Mach number for a specific morphed/unmorphed geometry) while the CFD interpolation tables are 3-D interpolations (across angle of attack, Mach number, *and* level of camber morphing).

2.3 Rotorcraft Modeling Methods

The objectives of a mission-driven optimization can range from minimizing fuel burn and power to maximizing range and endurance, and rotorcraft trim analysis is necessary to relate design variable inputs with these objectives. For a given mission stage pilot inputs must be determined to trim the rotorcraft. The effects of morphing are first quantified at the subsystem level before vehicle performance can be considered, though properties such as local lift and drag will typically change with pilot trim inputs, altering the local aerodynamic environment.

The main focus in this and much other adaptive rotorcraft design research is the main rotor blade. An efficient method for determining the main rotor contributions under specified pilot inputs to the overall vehicle performance is required and developed herein.

2.3.1 Blade Element Momentum Theory

One method to determine the main rotor aerodynamic response, particularly induced inflow at hover, is blade element momentum theory (BEMT). The BEMT code combines Blade Element Theory (BET) and Momentum Theory (MT) for a discrete set of local rotor blade elements [101, 74].

In hover a BEMT code coupled with parameterized pilot inputs is utilized to determine the trim flight condition. The pilot determines collective pitch θ_c, which defines the span-wise local section pitch θ when combined with the twist schedule down the rotor blade. The implemented BEMT code must determine the correct pilot collective pitch input such that there are no net forces. The correct trim flight collective pitch is determined via the minimization function:

$$\min_{\theta_c} f(\theta_c), \text{ where } f(\theta_c) = |T - W| \tag{2.10}$$

where T is the time-averaged vertical thrust component generated by the rotor disk and W is the vehicle weight. Once collective pitch is determined, the total power and figure of merit (FM) can be determined for trim, steady flight.

From BET, the local thrust dC_T generated by a single blade element is a function of the lift

coefficient C_l, the blade solidity σ, and the rotor span-wise position r as measured from the root:

$$dC_T = \frac{\sigma}{2} C_l r^2 dr. \tag{2.11}$$

From MT, local thrust dC_T is a function of total inflow (λ), induced inflow (λ_i), and rotor span location (r):

$$dC_T = 4\lambda\lambda_i r dr. \tag{2.12}$$

In hover, induced inflow and total inflow are equal. For both theories (BET and MT) to yield the same local thrust coefficient, the local inflow is given as:

$$\lambda^2 = \frac{\sigma C_l r}{8}. \tag{2.13}$$

σ is the local rotor blade solidity, given by:

$$\sigma = \frac{N_b c}{\pi R}, \tag{2.14}$$

where N_b is the number of blades and c is the chord length of the local blade element. The lift coefficient (C_l) is a function of angle of attack (α), which is in turn a function of local inflow. The lift coefficient is also a function of the Mach Number (M), or Reynolds number (Re_x) for a rotor blade with a constant chord under constant air properties:

$$C_l = f(\alpha, M). \tag{2.15}$$

The local Reynolds number Re for each element is given by:

$$Re = \frac{r\omega\rho c}{\mu}, \tag{2.16}$$

where ω, ρ, c, and μ represent the rotor angular velocity, air density, rotor chord, and dynamic

viscosity of air. The Mach number is given locally by:

$$M = \frac{\omega r}{c_\infty}, \tag{2.17}$$

where c_∞ is the speed of sound. It is worth noting the speed of sound is often referred to as c in literature, but since chord is also referred to as such, the ∞ subscript has been added to the speed of sound to avoid confusion when both are referenced together.

Due to the coupling between angle of attack, lift coefficient, and inflow, α, C_l, and λ must be iteratively updated until convergence. The local angle of attack is the difference between the local section pitch angle (θ and the local induced inflow angle:

$$\alpha = \theta - \arctan\left(\frac{\lambda}{r}\right), \tag{2.18}$$

where in hover the pitching angle θ is defined as

$$\theta = \theta_c + \theta_{blade}, \text{ where} \tag{2.19}$$

θ_{blade} is the manufactured rotor blade twist schedule. Once the inflow is determined, the local angle of attack and coefficient of thrust are determined via equations 2.18 and 2.12.

The local power coefficient (dC_P) is the sum of induced (dC_{p_i}) and profile (dC_{p_0}) power of the blade element:

$$dCp_i = dCt\lambda. \tag{2.20}$$

The local profile power coefficient dC_{P0} is given by:

$$dC_{P_0} = \frac{\sigma C_d r^3}{2} dr, \tag{2.21}$$

and the total coefficient of power sums the profile and induced power such that:

$$dC_P = dC_{P_i} + dC_{P_0}. \tag{2.22}$$

The local coefficient of drag (dC_d) is a function of angle of attack (α) and Mach number (M) similarly to the coefficient of lift, either through experimental or computational methods.

The individual elements are summed to determine the coefficients of thrust and power for the entire rotor blade:

$$C_T = \sum_{i=1}^{N} dC_T(i), \tag{2.23}$$

$$C_{P_i} = \sum_{i=1}^{N} dC_{P_i}(i), \tag{2.24}$$

$$C_{P_0} = \sum_{i=1}^{N} dC_{P_0}(i), \tag{2.25}$$

$$C_P = \sum_{i=1}^{N} dC_P(i). \tag{2.26}$$

The total trim power in hover, both a common objective and an input used to determine other common objectives (fuel consumption, endurance, and range), is then calculated as:

$$P_h = C_P \rho A {V_{tip}}^3 = C_P \rho \pi R^2 \left(\omega R\right)^3. \tag{2.27}$$

A comparison between BEMT code predictions, experimental testing results, and results from the well known Comprehensive Hierarchical Aeromechanics Rotorcraft Model (CHARM) [131, 35, 138] can be seen in Fig. 2.6. The BEMT code was run using aerodynamic properties both from XFOIL and from the parameterized CFD tables described in this chapter. It can be seen that the use of XFOIL tables leads to over-predicting performance while the CFD tables result in much better correlation with existing experimental and computational results. The figure of merit (FM)

is defined as:

$$FM = \frac{C_{P_{ideal}}}{C_{P_{actual}}} = \frac{P_{ideal}}{P_{actual}}, \tag{2.28}$$

where the ideal coefficient of power $C_{P_{ideal}}$ is given by:

$$\frac{{C_T}^{\frac{3}{2}}}{\sqrt{2}} \tag{2.29}$$

The BEMT code determines pilot collective pitch input for trim flight in hover by determining the

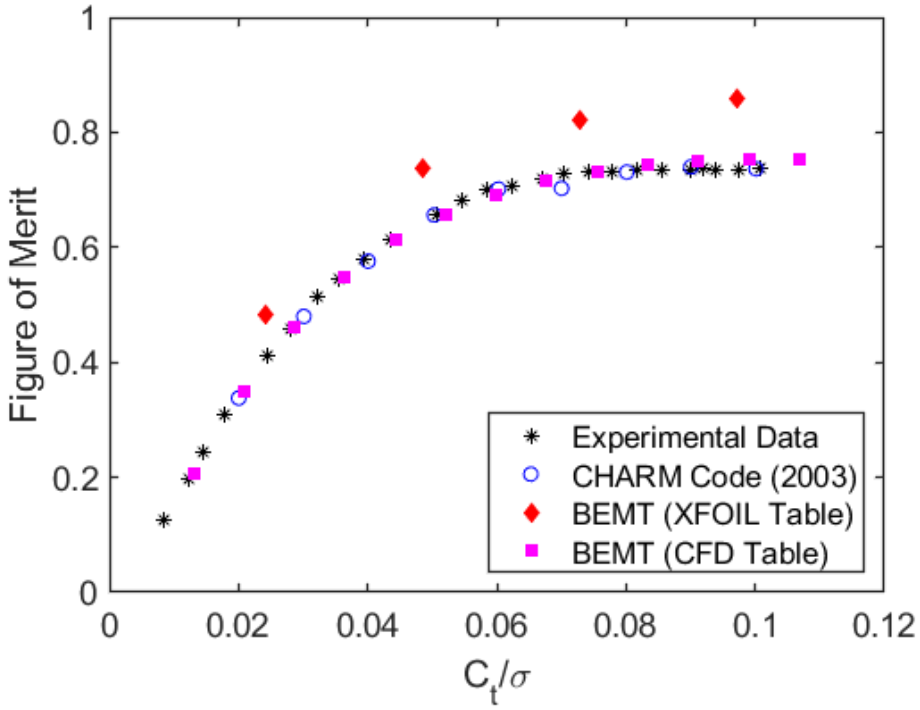

Figure 2.6: Comparison of UH-60 main rotor performance from experimental data and computational models in hover as a function of coefficient of thrust normalized by rotor blade solidity ($\frac{C_t}{\sigma}$)

lift, drag, and moment profiles integrated over the entire rotor disk considering local morphed states for each blade element then calculating the force equilibrium. The trim hover condition can be extended to climb based on hover-climb empirical relations. The hover rotorcraft trim code focuses

on the main rotor and determining the trim condition for an axisymmetric rotor disk dependent only on the pilot's collective pitch input. However, this code does not consider additional forward flight aerodynamic responses, which will be discussed in the following section.

The local lift coefficient is determined from interpolation tables over a wide range of angles of attack and Mach numbers: $C_l(\alpha, M)$. The BEMT code determines the trim flight condition in hover given the required lift, drag, and moment profiles for each local blade element. These aerodynamic profiles can be calculated via additional aerodynamic tools or precomputed in advance over a range of rotor blade geometries as previously described for morphing CFD interpolation tables. This versatility allows for a large range of parameterized designs to be explored using the same BEMT code.

The BEMT code determines trim flight for hover and climb, based on hover-climb empirical relations. For vertical climb, first the total power for hover is calculated via equation 2.27. Then the power ratio for upward climb is given by:

$$\frac{P}{P_h} = \frac{V_c}{2v_h} + \sqrt{\frac{V_c}{2v_h}^2 + 1}, \tag{2.30}$$

which is valid for:

$$\frac{v_c}{v_h} \geq 0, \tag{2.31}$$

signifying a positive climb rate.

In a vortex ring state ($-2 \leq \frac{v_c}{v_h} \leq 0$), the assumptions required for momentum theory to hold (a steady, quasi 1-D problem in which air flows in one direction through the rotor blade without recirculation) break down as the rotor can experience unsteady, recirculating flow. In this work, the power required for hover (steady, level flight) is used for the vortex ring state as a fairly conservative estimate, simply due to the physical limitations of calculating a true solution for adaptive rotor blades without experimental data or high-fidelity solutions: in the case of low speed descent, the preferred rotor blade in this work will be the best rotor blade for hover under the exact same flight conditions. In higher speed descent (a 'windmill brake' state), the required

power to maintain constant velocity descent is given by:

$$\frac{P}{P_h} = \frac{V_c}{2v_h} - \sqrt{\frac{V_c}{2v_h}^2 - 1}, \tag{2.32}$$

which is valid for:

$$\frac{v_c}{v_h} \leq -2. \tag{2.33}$$

A more detailed description of the momentum conservation analysis behind BEMT and its assumptions can be found in [101].

Additionally, the main rotor generates a moment around the main hub. One common method to alleviate this moment for rotorcraft is the addition of another rotor; coaxial rotors generate an opposing moment along the same rotational axis, while tail rotors generate thrust to generate an opposing moment arm farther away from the main rotor rotational axis. This work considers the UH-60, which utilizes a tail rotor, so the tail rotor contributions were determined based on the required torque necessary to balance out the main rotor contributions.

The tail rotor generates torque to counteract the main rotor. Main rotor torque can be calculated from the total coefficient of power C_P from equation 2.26:

$$\text{Torque} = C_P A \rho \omega^2 R^3. \tag{2.34}$$

With the torque defined, the tail rotor center a distance l must generate thrust T_{TR} such that:

$$T_{TR} = \frac{\text{Torque}}{l}. \tag{2.35}$$

With tail rotor thrust defined, the required power is then determined using standard relations between thrust, power, and efficiency using momentum theory (with a figure of merit of 0.7) as given in equation 2.28. The total required power is then determined from the tail rotor power C_P similarly to that of the main rotor as defined in equation 2.27, now using the angular velocity and geometry of the tail rotor. It is worth noting that for active tail rotor morphing a more comprehen-

sive BEMT implementation would be necessary to determine the true figure of merit.

The hover rotorcraft trim code focuses on the main rotor and determining the trim condition for an axisymmetric rotor disk dependent only on the pilot's collective pitch input. However, this code does not consider additional forward flight aerodynamic responses and pilot cyclic inputs, which will be discussed next.

2.3.2 Comprehensive Rotorcraft Analysis Tools

In forward flight additional aerodynamic effects must be considered beyond solely the induced inflow active in hover. Parasitic power is required to propel the vehicle forward and is proportional to the forward flight speed cubed, becoming dominant at high speeds. The local rotor blade inflow and incoming velocity contributions to the section and thus blade lift, drag, and moment in forward flight are heavily dependent on the rotation angle of a given blade around the hub, ψ. Such responses require a more robust, comprehensive analysis tool that incorporates aerodynamic effects not seen in hover.

TAMU Rotorcraft Analysis Code (TRAC) was developed by Advanced Vertical Flight Laboratory (AVFL) and written in MATLAB [112, 65] to integrate these aerodynamic responses with pilot controls [100, 98]. Main rotor dynamics are developed according to BEMT, including dynamic forward flight effects such as rotor blade lead/lag and flapping with the Pitt-Peters linear inflow model. TRAC determines attitude, power, and control inputs at designated forward velocities by resolving all forces and moments around the vehicle center of gravity, considering full rotation of all main rotor blades [99].

Different rotorcraft can be modeled in TRAC by setting the relative locations between aerodynamic surfaces and the vehicle center of gravity. TRAC resolves forces and moments of a main rotor with cyclic controls, flapping, and damping. Prediction from the BEMT code and TRAC were compared with experimental results on the UH-60 from NASA Ames [142] as shown in Fig. 2.7, along with the different regions of inflow effects. TRAC is used in flight conditions where linear effects dominate the inflow (advance ratio μ over 0.15 [101]); BEMT nonlinear inflow conditions where considered, primarily in hover.

The previously described computational tools (BEMT and TRAC) determine trim flight for hover and forward flight based on the inflow conditions, as significant induced inflow nonlinearities can exist in hover and low speed forward flight [101]. The BEMT code determines the collective pitch for trim flight in hover, while TRAC also considers pilot cyclic inputs in forward flight. The trim rotorcraft solutions are used to determine the performance metrics of interest in this design optimization study, most often a measure of power in a given phase of flight. The power required for steady-state, trim flight is the sum of the profile, parasitic, and induced power, as shown in Fig. 2.8. Parasitic power, as previously described, is required to propel the vehicle forward and overcome drag forces due to the incoming free-stream velocity on the airframe. Profile power is required to rotate the rotor blade, overcoming the rotor blade aerodynamic drag forces. Induced power produces lift, overcoming the rotor drag created by the induced inflow. The overall power

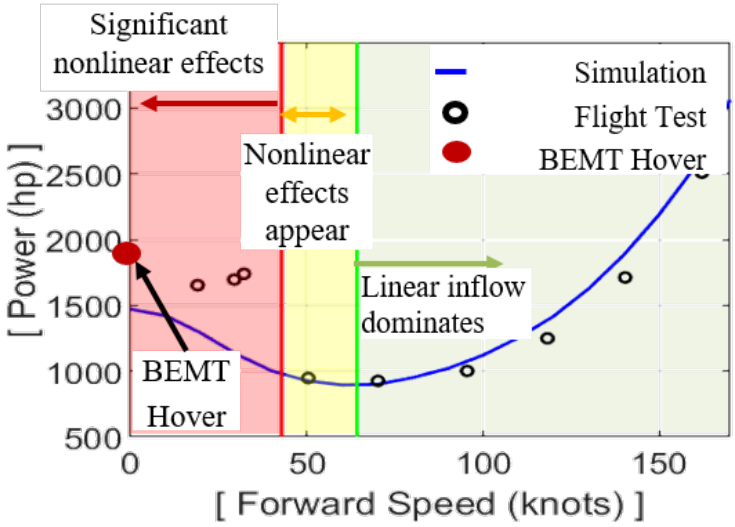

Figure 2.7: Comparison of UH-60 power profile from TRAC and from experimental test data in forward flight. BEMT is used for hover, the low speed forward flight region is not considered, and TRAC is used where there are insignificant nonlinear effects.

varies with vehicle velocity, with different aerodynamic contributions dominating for different objectives. In general, maximum loiter time (or maximum endurance) occurs in forward flight due to

the decrease in induced drag, while the maximum range occurs at the lowest power-to-forward velocity ratio. Cruise can be defined as either the maximum vehicle velocity given a power constraint or as forward flight at a specified velocity. The combination of BEMT and TRAC allows for any of these flight conditions to be accurately modeled as individual stages in a larger mission which is the primary goal of this work as we consider adaptivity to maximize mission performance. The

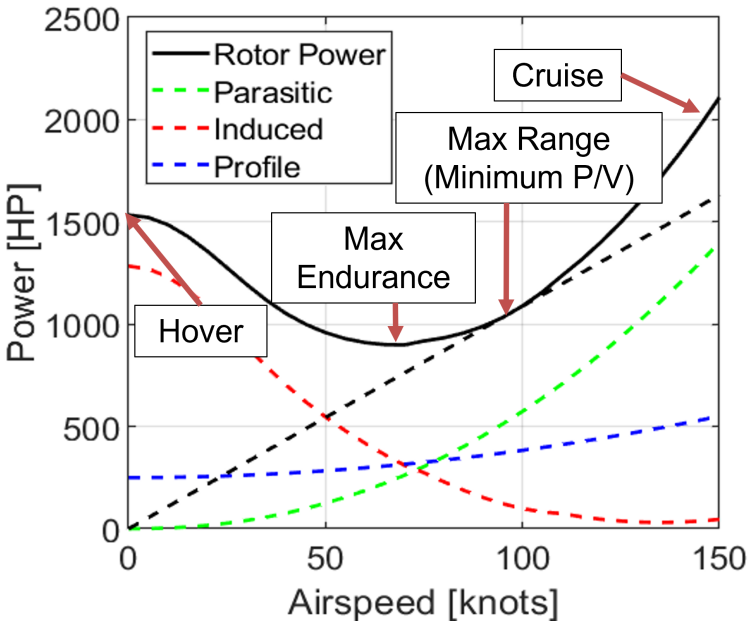

Figure 2.8: Different airspeeds in the rotorcraft power profile correspond to different mission-driven performance metrics.

general aerodynamic tools and trim solvers are applied here to the UH-60 Blackhawk helicopter due to the generous existing experimental and computational results for comparison, partly due to NASA's UH-60 airloads program [1]. However, this computational framework can be applied to other vehicles, as all the vehicle parameters can be altered to match aerostructural definition. Some of those aerostructural parameters can be seen in Table 2.2.

Table 2.2: Vehicle Parameters

Parameter Name	Number of Values
Center of Gravity Location	3
Vehicle Moment of Inertia	3
Main Rotor Hub Location	3
Tail Rotor Hub Location	3
Horizontal Tail Location	2
Vertical Tail Location	2
Horizontal Tail Planform	1
Vertical Tail Planform	1

2.4 Fully Coupled Aerostructural Analysis

The previous analysis codes, BEMT and TRAC, consider the aerodynamic response for a defined rotor blade shape. However, the rotor blade shape can change due to inertial and aerodynamic loading during operation. Changes in these loads will affect the rotor blade geometry, while these changes will alter local aerodynamic properties. This interdependency between aerodynamic and structural responses has led to the development of a field known as aero-elasticity. To fully capture the aerostructural response of the rotor blade and full rotorcraft performance, the coupling of aerodynamic and structural solvers must converge to a single solution.

Typically, aero-elastic coupling is achieved by passing values shared by both the structural analysis and the aerodynamic analysis back and forth until convergence. In a fully-coupled analysis, an aerodynamic evaluation could pass the pressure field to the structural analysis, which would then pass the deflection back to the aerodynamic analysis as shown in Fig. 2.9. This process is repeated until both the deflection and pressure field are converged.

However, such an aerodynamic process of calculating the pressure fields using CFD for every rotor blade location while still solving for a trim flight condition is computationally impractical during large design and optimization procedures without extensive time and resources. Large numbers of aerostructural evaluations are required to determine the pilot inputs necessary for trim flight, while large numbers of iterations are required for each fully coupled solution, each requiring

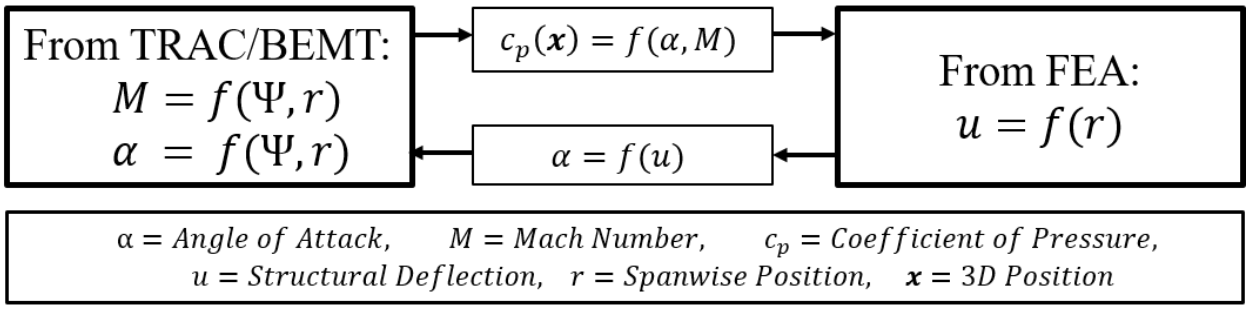

Figure 2.9: At a high level, pressure information is passed from the aerodynamic analysis to the structural analysis, while the deflection is passed from the structural analysis to the aerodynamic analysis until both solutions are converged.

structural and aerodynamic evaluations. Considering the aerodynamic tools previously described, a typical trim flight solution requires 60-120 evaluations in hover using the BEMT code, and 300-400 evaluations for forward flight.

The addition of another iterative loop for each evaluation to represent a converged fully-coupled aerostructural solution increases the number of evaluations by a multiplicative power while adding additional structural evaluation time. Conducting structural analysis for each rotor blade aerodynamic condition adds additional feasibility problems. In hover, assuming radial symmetry allows for one structural evaluation for every aerodynamic evaluation: the aerostructural state for each rotation angle ψ is the same for a rotor blade with no cyclic inputs, no incoming velocity v_{inf}, and a constant collective pitch θ_0. In reality this radially symmetric rotor disk may be slightly askew due to the tail rotor positioning and contributions.

In contrast to the hover flight condition, the rotor disk in forward flight experiences a wide range of aerodynamic conditions as a function of rotation angle ψ. Thus, additional structural analysis is needed to capture the full aerodynamic profile in forward flight. The aerodynamic analysis used in TRAC evaluates the four UH-60 Blackhawk rotor blades over 72 rotation angles using a specified inflow angle and pilot controls. Thus, 288 structural evaluations are necessary to include the structural deflection under a given aerodynamic environment. As previously mentioned, those 288 structural responses would then be iteratively updated along with the trim flight minimization

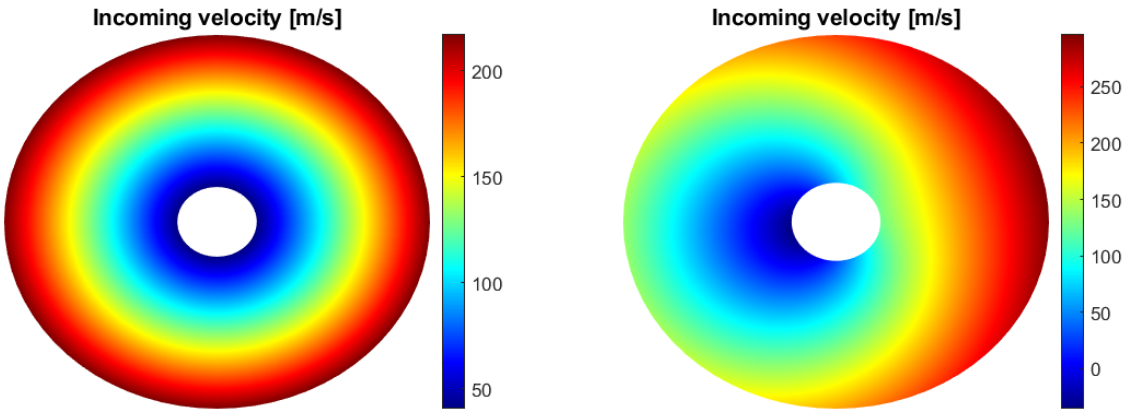

Figure 2.10: The rotor disk aerodynamic environment is modeled as radially symmetric in hover (left), but additional aerodynamic effects in forward flight (right) require considering different radial locations.

to determine the true fully-coupled aerostructural solution. Since the hover BEMT code would only require one structural evaluation per aerodynamic iteration, as opposed to 288 evaluations, the fully coupled method was first developed for the hover flight condition. A comparison of rotor disks between hover and forward flight can be seen in Fig. 2.10, highlighting the differences in aerodynamic environments.

High fidelity solutions for structural and aerodynamic models are associated with high computational costs. CFD is very expensive, as previously mentioned, and the camber morphing interpolation tables described previously were again employed to reduce the runtime while preserving accuracy. Given the pre-computed CFD results, a given angle of attack and Mach number could also determine the pressure field around the CFD results. This pressure field is then applied to the rotor blade in the structural analysis.

2.4.1 Structural Finite Element Model

A structural finite element analysis (FEA) model was developed in ABAQUS [144] to simulate the UH-60 Blackhawk rotor blade under inertial and external pressure loading. The pressure loading was applied as a field onto the surface of the rotor blade. The model specifications are

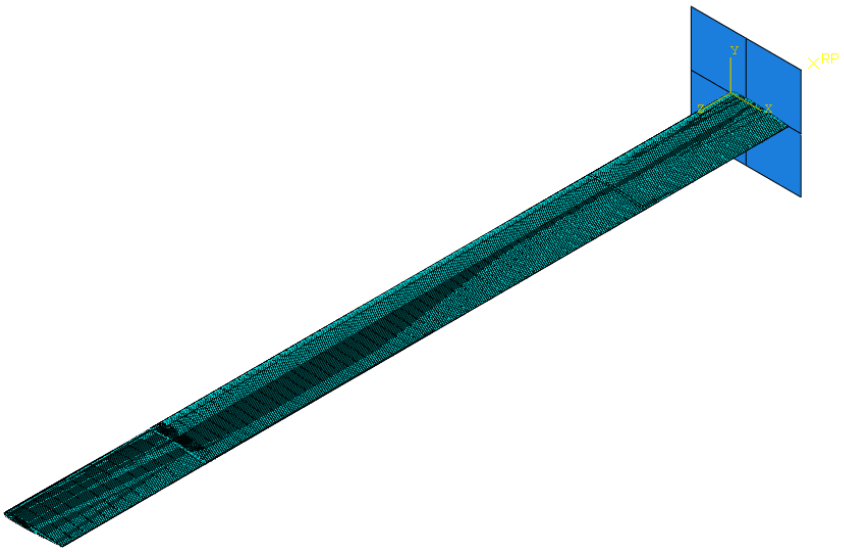

Figure 2.11: Meshed FEA rotor blade model.

described in more detail in previous rotorcraft modeling and design works [90, 92] as well as an application for uncoupled static aeroelastic analysis (USAA) [128]. This model was first developed by Trent White and Francis Phillips for the USAA analysis then generously shared with me in a collaboration in which the BEMT code, TRAC, and the FUN3D CFD results provided the aerodynamic analysis to be coupled with the FEA model. The contribution in this work is to take the existing FEA rotor blade model and couple it with aerodynamic analysis to determine converged aero-elastic solutions under mission-driven trim flight conditions. The meshed rotor blade can be seen in Fig. 2.11.

The FEA model first applies an inertial load to the rotor blade, then applying the pressure fields based on the local angle of attack values and Mach numbers for every rotor blade element determined by the aerodynamic analysis, along with the starting pitching angle for the rotor blade. An example rotor blade under loading can be seen in Fig. 2.12. The unloaded rotor blade twist is then compared with the aerodynamically and inertially loaded rotor blade to determine the structural deflection. This deflection is then discretized into local rotor blade elements, passing the deflection each element center to the aerodynamic solver. The deflection is then integrated into the overall

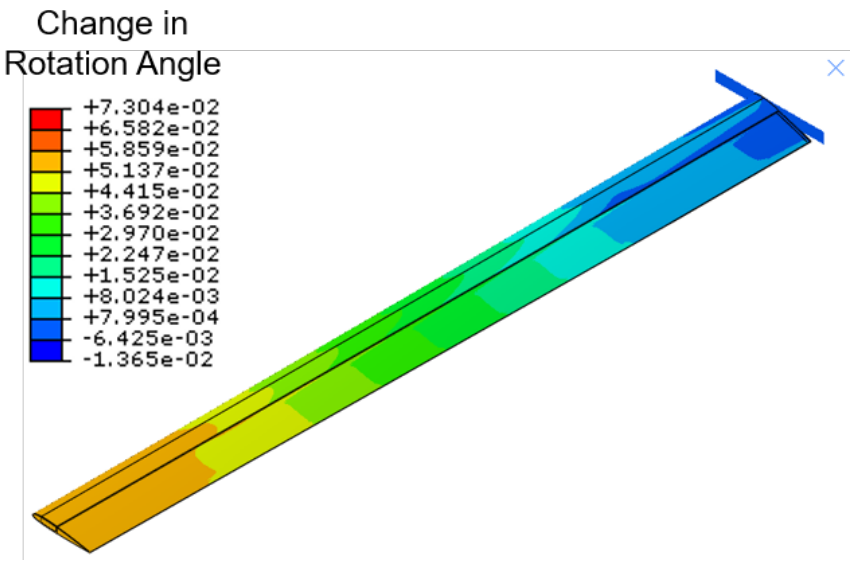

Figure 2.12: Rotor blade under aerodynamic loading results in changes in local rotation angle (in radians).

rotor blade twist schedule as:

$$\theta_{twist} = \theta_{deflection} + \theta_{root} + \theta_{blade}, \tag{2.36}$$

where $\theta_{deflection}$ is a function of the structural model displacements u, θ_{root} is the rotor blade root initial angle, a function of collective and cyclic pitching inputs, and θ_{blade} is the manufactured undeflected rotor blade twist schedule. Convergence is approached as the changes in deflection, and thus the changes in aerodynamic conditions, approach zero.

2.4.2 Aeroelastic Coupling

An aero-elastic code was developed to minimize the number of FEA evaluations required to determine a converged, trim, fully-coupled solution for hover rotorcraft mission-driven simulations, as shown in Fig. 2.13. After some initial observations iterating both deflection and aerodynamic evaluations, it was clear the FEA deflection converged after a few iterations, with much longer evaluation times. The fully-coupled simulation leverages this by only evaluating the FEA solution

for trim aerodynamic solutions instead of every aerodynamic evaluation.

Vehicle and mission parameters are initialized at the start, including an initial aerodynamic condition based on vehicle loading. Next, the rotor blade deflection is calculated under aerodynamic loads and updated. Then the trim flight aerodynamic solution is calculated considering the previous structural deflection. Each iteration of the trim flight minimization assumes the rotor blade has the previous rotor blade deflection, which is incorporated into the aerodynamic analysis via equation 2.36.

The aerodynamic analysis calculates the lift, thrust, and power for a given collective pitch using the BEMT code, which iteratively determines the local thrust, angle of attack, and inflow for each local blade element. The BEMT code originally used a fixed-point iteration method to determine the inflow, but now utilizes a gradient-based solver in Matlab: fsolve. This change was introduced due to some convergence difficulties with fixed point iteration due to non-smooth interpolation tables, leading to unconverged local minima; no solutions are accepted if any local rotor blade elements are unable to converge. The fsolve function provides a more robust method to find the inflow solution than the fixed point method. Other methods such as genetic algorithms and particles swarm optimizations were explored, but required orders of magnitude more evaluations to determine trim flight solutions.

Once the trim collective pitch is determined, the new aerodynamic environment is applied to the rotor blade in the FEA model and the deflection is updated. The updated FEA deflection is evaluated for trim, a relatively quick aerodynamic calculation. If the new solution is still a trim flight solution, the total power is calculated and aerodynamic with structural results passed forward in the simulation. If the new solution is no longer a trim flight solution due to the updated structural deflections, the aerodynamic analysis once again iteratively solves for the collective pitch necessary for trim flight. Thus, the collective pitch for trim flight and the structural deflection are determined iteratively. Finally, the tail rotor power contributions necessary to balance the torque generated by the main rotor are determined. These tools will be demonstrated toward mission-driven adaptive design later in section 5.3.

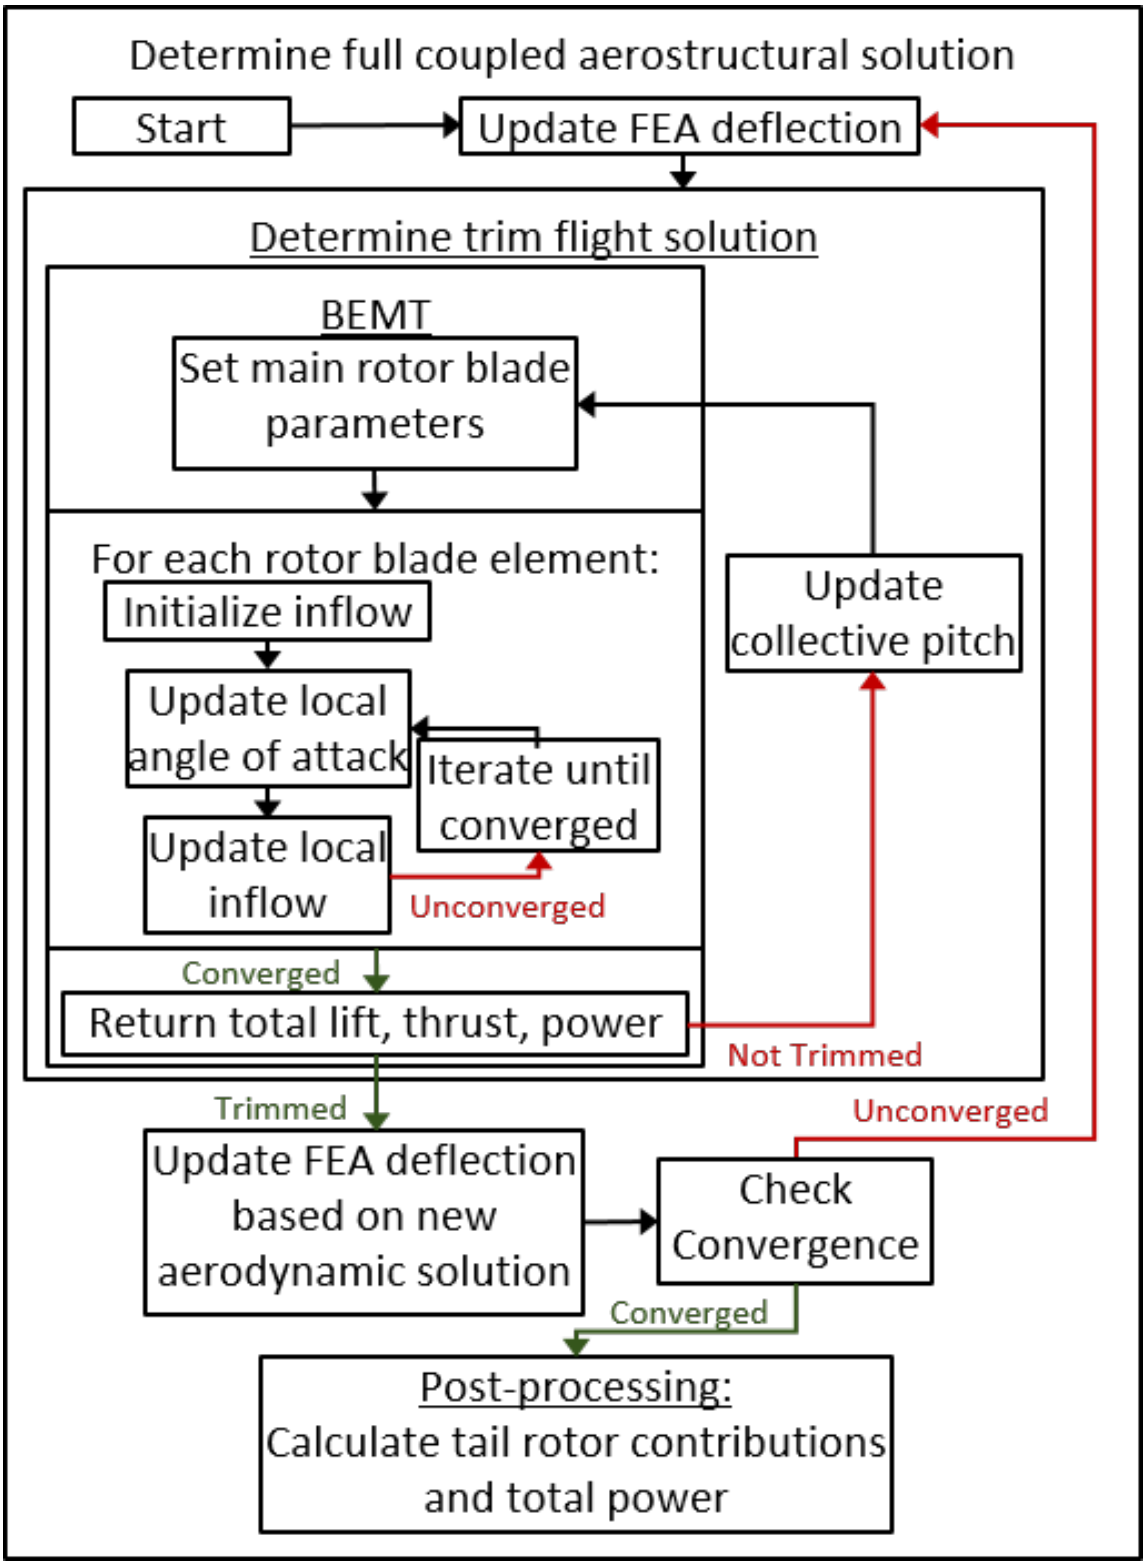

Figure 2.13: The process to determine a converged coupled aeroelastic solution, with a converged trim flight solution, using converged BEMT solutions.

2.5 Chapter 2 Summary

In this chapter, a wide range of computational tools are developed to design and optimize adaptive rotorcraft. At the local rotor blade element level, a variety of numerical tools are implemented, from linear properties to blade element momentum theory to computational fluid dynamics and interpolation tables. Additionally, each method of morphing the rotorcraft must be considered along with impact on vehicle performance. Active twist, chord, and camber (or outer mold line) are considered in this work, and changes in these adaptive variables are related to lift, drag, and moments based on the aerodynamic environment, which in turn is related to pilot inputs from trim flight.

The rotorcraft performance must be modeled to determine specific trim flight conditions necessary for a given mission stage. Two numerical methods are employed in this work to determine pilot inputs required for trim flight, their selection depending on inflow state linearity. An efficient blade element momentum theory code combined with multiple optimization loops is developed in this work to determine the hover flight condition, while an existing comprehensive vehicle analysis code (TRAC) determines trim forward flight.

These tools are general and can determine trim flight over a wide range of velocities and other mission stage parameters with much lower computational costs than many existing rotorcraft analysis tools, which often require computational fluid dynamics solutions solved during the analysis. These relatively efficient aerodynamic calculations can then determine preferred structural configurations based on mission objectives over a wide range of design variables.

3. A THEORY FOR ADAPTIVE DESIGN SPACE DECOMPOSITION

Often in adaptive structural design and optimization the location and placement of adaptive structures is pre-determined based on the designer's intuition about feasibility, morphing effectiveness, mission objectives, etc. Optimization tools are then used to determine the optimal adaptive design variables which can morph between preferred configurations. A multi-objective adaptive structures optimization considers multiple designs or configurations and evaluates them based on some objectives $\mathbf{f}$. In this case, the design problem does not include design variables that do not morph as unknown values to be determined; only the morphable design variables are varied to minimize the objective functions $\mathbf{f}$. By considering design variables that do not morph, the engineering problem design space becomes the union of fixed design variables, referred to in this work as $\mathbf{x}^f$, and the adaptive design variables $\boldsymbol{\alpha}$.

One common alternative to simply optimizing for adaptive structures considers every adaptive design variable for each mission state or objective. For example, in a five-stage rotorcraft mission with two adaptive twist parameters and two location/placement variables that cannot morph during a mission, the designer would optimize over 12 design variables: the two location design variables and twist values for every objective (two active twist rates over five mission stages), resulting in a larger design space that scales with increasing mission stages and their requisite objectives. Mission-driven rotorcraft design and analysis can be computationally expensive, even using the modeling methods described in Chapter 2 to avoid running the highest fidelity tools inside an optimization loop. As such, reducing the design space may prove valuable for determining the best sets of both fixed and adaptive variables.

While many design processes focus entirely on the most advantageous selection of $\boldsymbol{\alpha}$, it is also important to select fixed design variables $\mathbf{x}^f$ that enable the best morphing advantages to minimize $\mathbf{f}$. Another approach is to consider the adaptive design variables for each morphed state as different design variables. An example would be considering the same adaptive twist parameter for hover and cruise as two separate design variables in the design space. However, this ap-

proach increases the design space, particularly when considering larger numbers of objectives. The methodology presented in this section follows a novel approach to determining preferred fixed and adaptive design variables for multi-objective design optimizations. A selection procedure is developed to determine optimal morphing configurations given a multi-objective adaptive structures design problem.

Chapter 3 Outline

This chapter is divided into 3 sections:

- In section 3.1 a general conventional multi-objective optimization and the decomposition of the design space into subsets of design variables based on adaptivity is described and developed. Then methods of quantifying potential adaptive designs are formulated considering both similarity between non-adaptive variables in the design space and adaptive advantages in the objective space.
- In section 3.2 techniques are introduced and developed considering adaptivity for optimization of both fixed and adaptive design variables. Adaptivity can be considered via i) multiple sets of optimizations, ii) post-processing to determine sets af adaptive designs after an optimization, or iii) *during* an optimization as *in-situ* selection criteria.
- In section 3.3 the methods developed in this chapter are applied to simple mathematical and engineering design problems with analytical solutions. Multiple optimization techniques are used to determine preferred adaptive designs over a design space containing design variables that cannot and cannot be altered. The use of adaptivity as a selection algorithm is also evaluated and compared to existing selection algorithms.

In engineering design, the ideal configuration maximizes the full span of performance objectives considered by the designer. However, this is rarely the case, and compromises between objectives must be weighed even in adaptive systems design, as some elements are fixed after early

design stages (e.g., actuator placement, structural member thickness, number of rotor blades, etc.). For the vast majority of adaptive structural design and optimization studies in the literature, location and placement of actuating components is pre-determined based on designer intuition regarding subsystem feasibility, morphing effectiveness, mission objectives, past successes, etc. Given this *a priori* input, optimization tools are then limited to determining preferred geometric configurations both of these active components and of the more conventional structure that surrounds them.

If the optimization study is multi-objective, the designer will determine a final adaptive structures configuration based on necessary trade-offs between engineering goals. In the end, one arrives at a final set of parameters that define the geometry of all structural members, active and passive. Before formally decomposing the full design space (i.e., of input variables), let us first consider the conventional multi-objective optimization problem in more detail.

3.1 Mathematical Introduction

In conventional systems a multi-variable multi-objective design optimization problem is based on the determination of $\mathbf{J}^*$ where:

$$\mathbf{J}^* = \min_{\mathbf{X}} \mathbf{f}(\mathbf{X}) = \mathbf{f}(\mathbf{X}^*) \tag{3.1}$$

subject to inequality and equality constraints:

$$\mathbf{g}(\mathbf{X}) \leq \mathbf{0}, \tag{3.2}$$

$$\mathbf{h}(\mathbf{X}) = \mathbf{0}, \tag{3.3}$$

such that the set $\mathbf{J}^*$ of optimal designs across multiple objectives $\mathbf{f}()$ form the non-dominated Pareto frontier. Note that $\mathbf{X}$ represents the full set of design variables, and the solution $\mathbf{X}$ then represents the set of all Pareto-optimal design configurations.

In the conventional case that adaptive structures are not considered, the designer will eventually be forced to choose a single fixed design configuration $\mathbf{X_{final}} \in \mathbf{X}^*$ such that this final configuration represents a compromise between competing objectives. The final system performance is then $\mathbf{f}(\mathbf{X_{final}})$, as seen in Fig. 3.1. It is now valuable to imagine an idealized case in which possible

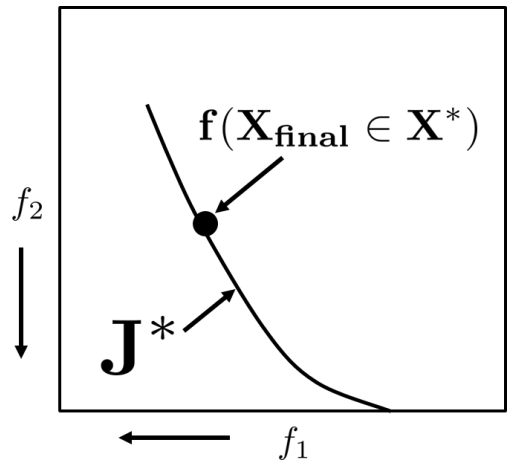

Figure 3.1: $\mathbf{J}^*$ represents the Pareto frontier within an objective space considering the full set of design variables $\mathbf{X}$.

configurations of a system were all fully adaptive and *every* design variable could be altered within operational bounds at any time during operation. This would allow an *in situ* preference for the maximization of one performance metric over any others. Such adaptivity would be represented by full mobility along the frontier $\mathbf{J}^*$ such that a designer was no longer required to accept compromises between competing objectives. This ability to adjust to changing operational conditions or functional objectives is the ultimate goal of adaptive design, and such set of designs will be referred to as the *idealized Pareto frontier*. However, fully adaptive systems are unrealistic and we must recognize that final values for many design variables must be determined and fabricated/coded into a system. *Partially* adaptive systems, on the other hand, have been demonstrated, can be advantageous, [10, 11, 116, 108, 86, 41, 118, 109], and represent the central topic of this work. We will represent the current configuration of such partially adaptive systems as being represented

by the combination of i) a set of conventional fixed design variables $\mathbf{x}^f$ and ii) a set of adaptive variables $\boldsymbol{\alpha}$ that represent the capability, where applicable, for in situ system adjustments. That is, $\mathbf{X} = \mathbf{x}^f \cup \boldsymbol{\alpha}$. In such systems, the performance can be quantified by the set of objectives

$$\mathbf{f}(\mathbf{X}) = \mathbf{f}(\mathbf{x}^\mathbf{f}, \boldsymbol{\alpha}). \tag{3.4}$$

Mobility along $\mathbf{J}^*$ is not realistic given the fixed nature of the variables in $\mathbf{x}^f$. However, mobility along local frontiers $\hat{\mathbf{J}}^*(\mathbf{x}^f)$ is possible due to the *in situ* adaptivity of $\boldsymbol{\alpha}$, where in particular we define these local frontiers via:

$$\hat{\mathbf{J}}^*(\mathbf{x}^f) = \min_{\boldsymbol{\alpha}} \ \mathbf{f}(\mathbf{x}^f, \boldsymbol{\alpha}) = \mathbf{f}(\mathbf{x}^\mathbf{f}, \boldsymbol{\alpha}^*(\mathbf{x}^\mathbf{f})). \tag{3.5}$$

Given the idealized Pareto frontier represented by $\mathbf{J}^*(\mathbf{X})$ and the decomposition of $\mathbf{X}$, the essential challenge is then finding the single fixed set $\mathbf{x}^{\mathbf{f}^*}$ that results in a local frontier $\hat{\mathbf{J}}^*(\mathbf{x}^{\mathbf{f}^*})$ which is minimally offset from $\mathbf{J}^*$. We will rigorously define this optimally adaptive design as the solution to the following:

$$\hat{\hat{\mathbf{J}}}^* = \min_{\mathbf{x}^f} \ (\mathrm{Dist}\{\hat{\mathbf{J}}^*(\mathbf{x}^f), \mathbf{J}^*\}) = \mathbf{f}(\mathbf{x}^{\mathbf{f}^*}, \alpha^*), \tag{3.6}$$

where $\mathrm{Dist}\{\bullet\}$ is a measure of mutual proximity and similarity between two sets and will be specifically defined in following subsection. These concepts are illustrated in Fig. 3.2. In other words, we define the preferred adaptive system to be that which provides the closest possible performance to an idealized fully adaptable design while being limited by the physical reality that only a portion of the configurational variables can be altered in-situ.

In an engineering sense, $\mathbf{x}^f$ contains the values permanently manufactured into a structure (e.g., fixed structure thicknesses and locations, active component placement and reference size), while $\boldsymbol{\alpha}^*$ are the full range of values for the adaptive variables which can and should be altered during operation (e.g., component length, twist angle, etc. associated with actuation levels).

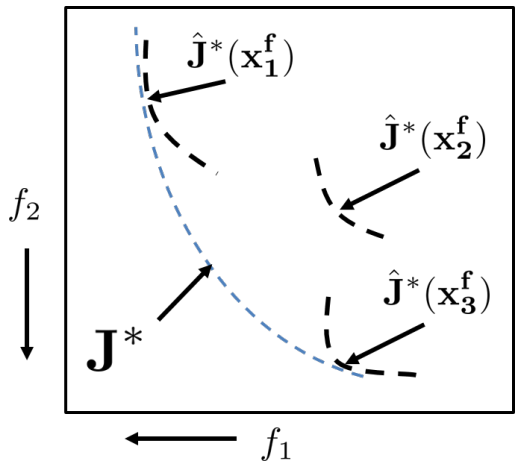

Figure 3.2: $\mathbf{J}^*$ represents the idealized Pareto frontier and considers best performing designs across the full range of $\mathbf{x}$ considered.

3.1.1 Definition of Distance Metric

Essential to the implementation of equation 3.6 is the mathematical definition of the distance between $\mathbf{J}^*$ and $\hat{\mathbf{J}}$. Note that, due to the manner in which these surfaces are commonly generated (e.g., genetic algorithms, design of experiments, etc.), each is composed of a discrete set of points. Therefore, the function $\mathrm{Dist}\{\bullet\}$ must consider the relationship between such sets of points. During development of this approach, it was found that two distinct classes of distance-like metrics were well-suited to the design problems at hand. The first considers points throughout the surfaces while the second addresses only the limits of $\mathbf{J}^*$ and $\hat{\mathbf{J}}$.

A distance metric for continuous adaptivity requires measuring the distance between approximated surfaces defined multi-dimensional datasets with different numbers of points. The *Hausdorff* distance is the maximum distance of a set to the nearest point in another set, used in this work to define the distance between the approximate ideal Pareto frontier and local adaptive frontiers [12, 53]. A traditional *Hausdorff* distance can be expressed as:

$$\mathrm{Dist}_{Hausdorff}\{\mathbf{A},\mathbf{B}\} = \tag{3.7}$$

$$\max\left\{\max_{a\in\mathbf{A}} d(a,\mathbf{B}),\ \max_{\mathbf{b}\in\mathbf{B}} \mathbf{d}(\mathbf{b},\mathbf{A})\right\}, \tag{3.8}$$

where $d(a, \mathbf{B})$ is the distance between point a and the nearest point in set $\mathbf{B}$. Specifically, the use of the *Hausdorff* distance encourages local frontiers to remain near all point in the approximate Pareto frontier, but there is no such requirement for a design in the approximate Pareto frontier to be relatively close to all the adaptive designs in a local frontier. The goal of continuous adaptivity is for the adaptive designs to match the approximate Pareto frontier; the inverse is unnecessary. Several distance and error metrics are required for comparing designs in both the design and objective space in this work. Further details can be found in Appendix E.

As an alternative to computing a distance which considers all points, we introduce a much less computationally intensive measure based on the *utopia points* of the surfaces involved. This is illustrated in Fig. 3.3. The ideal utopia point is defined by the best performance across each objective and is given by (cf. 3.1):

$$\mathbf{u}^* = \{\min({f_1}^*),\ \min({f_2}^*)...\ \min({f_{N_f}}^*)\}.$$

A similar quantity can be defined for any $\hat{\mathbf{J}}^*$ calculated assuming $\mathbf{x}_i^f$, which will be known as the adaptive utopia point for each set of fixed design variables (cf. equation 3.5):

$$\begin{aligned}\mathbf{u}_i^*(\mathbf{x}_i^f) = \{ \quad & \min(f_1(\mathbf{x}_i^f, \boldsymbol{\alpha}^*(\mathbf{x}_i^f))), \\ & \min(f_2(\mathbf{x}_i^f, \boldsymbol{\alpha}^*(\mathbf{x}_i^f))),\ ... \\ & \min(f_n(\mathbf{x}_i^f, \boldsymbol{\alpha}^*(\mathbf{x}_i^f)))\}\end{aligned}$$

If we define $\mathbf{d_i} = \hat{\mathbf{u}}_\mathbf{i}^*(\mathbf{x}_\mathbf{i}^\mathbf{f}) - \mathbf{u}^*$, then Dist can be defined as:

$$\mathrm{Dist}\{\hat{\mathbf{J}}^*(\mathbf{x}_i^f), \mathbf{J}^*\} = ||\mathbf{d}_i||, \tag{3.9}$$

and equation 3.6 can be calculated using much less computation than measures such as the Hausdorff distance. In addition to increased simplicity, this option is also highly faithful to the common adaptive structures design goal of maximizing each performance metric independent of the others

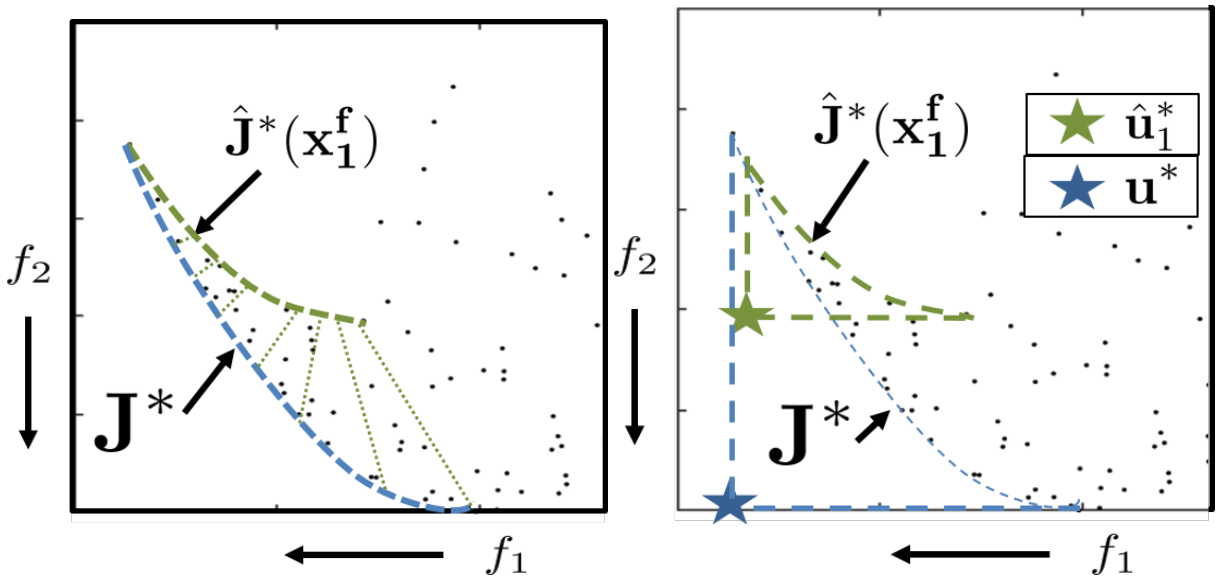

Figure 3.3: One metric for evaluating adaptive sets of designs is to measure the distance between the ideal Pareto frontier $\hat{\mathbf{J}}$ and the local frontiers $\hat{\mathbf{J}}^*$ for sets of adaptive designs. Another metric, the utopia point, considers the distance between the best values for local frontiers $\hat{\mathbf{u}}_i^*$ and those on the approximate ideal Pareto frontier $\mathbf{u}^*$.

when desired (i.e., reaching the 'tails' of the Pareto frontier).

3.2 Optimization Methods

Regardless of which distance metric is chosen, equation 3.6 must be solved. This itself requires both the solution of equation 3.1 and a sufficiently diverse set of $\hat{\mathbf{J}}^*$ frontiers associated with some number of fixed designs $\mathbf{x}_i^f$. Given that an accurate calculation of $\mathbf{J}^*$ is essential to the success of this approach, here we will describe two general methods for determining $\hat{\hat{\mathbf{J}}}^*$.

3.2.1 Option 1: Rigorously Calculated Local Frontiers

The generation of $\mathbf{J}^*$ requires a set of designs with functional evaluations, often generated via a multi-objective optimization and in this work considering all design variables $\mathbf{X}$. Determining the best adaptive set requires comparing the distance, however defined, between $\mathbf{J}^*$ and local frontiers $\hat{\mathbf{J}}^*$. Thus, the generation of different sets of local frontiers $\hat{\mathbf{J}}^*$ is essential. For each design and set of fixed design variables $\mathbf{x}_i^f$, there exists a range of performance that may be unknown to the designer without additional functional evaluations. Note that these functional evaluations do not

require a full optimization to explore the design space; a design of experiment (DOE) of sufficient size could generate the initial population in the same manner, typically in some objective-agnostic space-filling manner. An optimization is first proposed because prioritizing better objectives during the initial design space exploration is likely to result in a more accurate determination of the approximate Pareto frontier $\mathbf{J}^*$. For example, the genetic optimization selection algorithm NSGA-II selects the best designs for each generation based on Pareto-optimality and spacing along frontiers.

The first approach to determining each $\hat{\mathbf{J}}^*$ and the best local frontier $\hat{\hat{\mathbf{J}}}$ is via additional functional evaluations. In this manner, N sets of fixed design variables $\mathbf{x}_i$ are evaluated and the adaptive design variables $\boldsymbol{\alpha}$ are optimized to determine the range of performance for each $\mathbf{x}_i^f$. By this approach the result of every optimization is a local frontier surface $\hat{\mathbf{J}}^*$ representing a fully-realizable set of designs. Each optimization result is then compared using the distance metrics from section 3.1.1 to determine the best adaptive design $\hat{\hat{\mathbf{J}}}^*$.

After solving equation 3.1, via a full conventional multi-objective optimization process (Stage 1), one would then proceed with a second stage where, for each fixed design variable $\mathbf{x}_i^f$, we solve equation 3.9 such that:

$$\hat{\mathbf{J}}^*(\mathbf{x}_i^f) = \min_{\alpha} \ \mathbf{f}(\mathbf{x}_\mathbf{i}^\mathbf{f}, \boldsymbol{\alpha}) \ \forall \ \mathbf{i} = \mathbf{1}...\mathbf{N}. \tag{3.10}$$

This results in N localized Pareto frontiers sharing at least one point with $\mathbf{J}^*$, illustrated in Fig. 3.4. Note that no frontier $\hat{\mathbf{J}}^*(\mathbf{x}_i^*)$ can outperform $\mathbf{J}^*$ as per the definition of the latter in equation 3.1. After the rigorous calculation of equation 3.10 for N points, these results can be combined with the original $\mathbf{J}^*$ solution and used in equation 3.9. The solution to this procedure would then represent the preferred adaptive systems design as per the guiding conjecture.

The obvious drawback to this approach is computational cost, as an additional optimization is required to determine every local Pareto frontier. The number of fixed configurations, and thus the number of local frontiers, can down-selected based on previous performance. One approach would be to only consider fixed design values where previous functional analysis yielded Pareto-optimal multi-objective performance. However, this may exclude frontiers that have a large coverage of

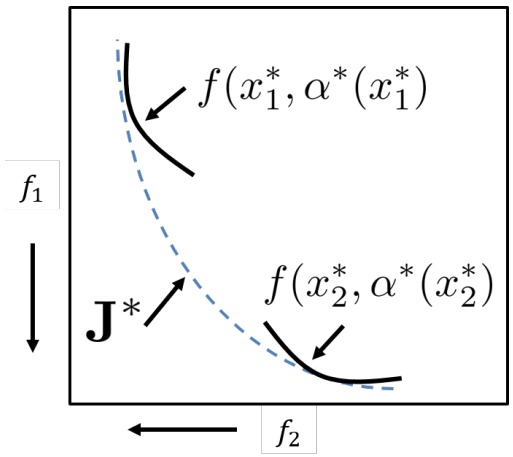

Figure 3.4: Each point on the idealized frontier $\mathbf{J}^*$ corresponds to a local non-idealized frontier representing the range of objective values attainable via adaptivity.

the objective space near but not on the Pareto frontier.

3.2.2 Option 2: Approximation of Local Frontiers using Existing Data

A second option approximates the solution to equation 3.9 by directly exploiting all discrete data generated toward the original solution of $\mathbf{J}^*$ such that additional functional evaluations are averted. Rather, this becomes a simple post-processing problem if one introduces an algorithmic approximation to the adaptive Pareto frontiers. In generating $\mathbf{J}^*$, individual designs $D_i = \{\mathbf{x}_i^f, \alpha_\mathbf{i}\}, \forall = \mathbf{1}...\mathbf{N_{eval}}$ have already been evaluated, many of them near or on the idealized Pareto frontier. Instead of requiring that each sub-frontier $\mathbf{f}(\mathbf{x}^{\mathbf{f}^*}, \boldsymbol{\alpha}^*(\mathbf{x}^*))$ (see Fig. 3.4) is strictly associated with a single $\mathbf{x}_i^*$, we alternatively consider a local neighborhood around each configuration $\mathbf{x}_i^f$ as being representative of designs we should consider.

Thus, for each design D_i, we create sets of similar designs:

$$\mathbb{D}_i = \{D_j | \parallel \mathbf{x}_j^f - \mathbf{x}_i^f \parallel < \epsilon\}, \tag{3.11}$$

where ϵ is based on data density (e.g. $\epsilon \approx 5\%$ range$(\mathbf{x}^f)$), as shown in Fig. 3.5. Thus, $\mathbb{D}_i$ is the set of all previously evaluated designs near $\mathbf{x}_i^f$. Because the objectives required to determine

$\mathbb{D}_i,\ i = 1...N$ were generated in finding $\mathbf{J}^*$, no additional functional evaluations are required. By

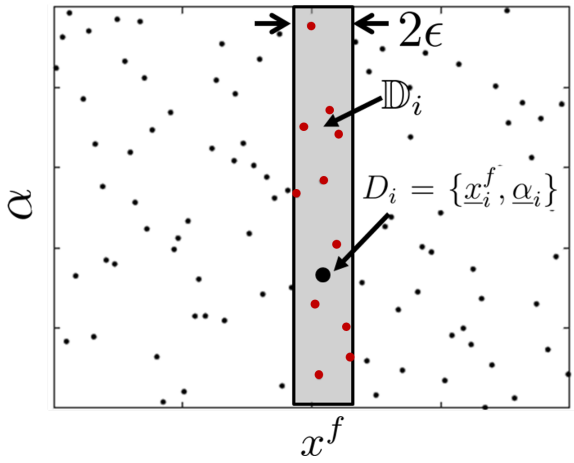

Figure 3.5: $\mathbb{D}_i$ is the set of designs having design variable values similar to $\mathbf{x}_i^f$ to within some range ϵ. This is illustrated for a case of $\text{len}(\mathbf{x}^f) = \text{len}(\boldsymbol{\alpha}) = 1$.

this formulation, each local frontier is determined via:

$$\hat{\mathbf{J}}^*(\mathbf{x}_i^f) = \min_{\boldsymbol{\alpha}} \ \mathbf{f}(\mathbf{x}_\mathbf{i}^\mathbf{f}, \boldsymbol{\alpha}) \approx \hat{\mathbf{J}}^*(\mathbb{D}_\mathbf{i}) = \min_{\mathbb{D}_\mathbf{i}} \ \mathbf{f}(\mathbb{D}_\mathbf{i}). \tag{3.12}$$

To clarify, the right side of equation 3.12 represents an approximation to the left side in two ways:

1. the set of points $\{\mathbf{x_j}\ \mathbf{such\ that}\ \parallel \mathbf{x_i} - \mathbf{x_j} \parallel < \epsilon\}$ only imprecisely represents $\mathbf{x}_i$, and

2. the surfaces equated in equation 3.12 are only approximately continuous, defined by sets of discrete points.

This assumption from equation 3.12 is invalid if i) there are not enough data points to construct approximate frontiers $\hat{\mathbf{J}}^*(\mathbf{x}_i^f)$ (which occurs if ϵ is too small or the generation of $\mathbf{J}^*$ is too sparse) or ii) ϵ is too large and designs considered 'similar' yield unrealizable objectives due to the large range of fixed design variables in the sets $\mathbb{D}_i$. The final configuration is a set of approximately

similar designs $\mathbb{D}_i$ with respect to fixed design variables $\mathbf{x}^f$ following from equation 3.6 given by:

$$\hat{\hat{\mathbf{J}}}^*(\mathbb{D}_i) = \min_{\mathbb{D}_i} \mathbf{f}(\mathbb{D}_\mathbf{i}), \tag{3.13}$$

which produces an approximation to the true Pareto-optimal adaptivity performance of any fixed design $\mathbf{x}_i$. Note that, since the values of $\mathbf{x}_i^f$ used to define each member of the set $\mathbb{D}_i$ only differ within some range to ϵ and not are identical, these are only approximations to the exact surfaces shown in Fig. 3.6. The fidelity of $\hat{\hat{\mathbf{J}}}^*$ to $\mathbf{J}^*$ is critically dependent on the density and diversity of

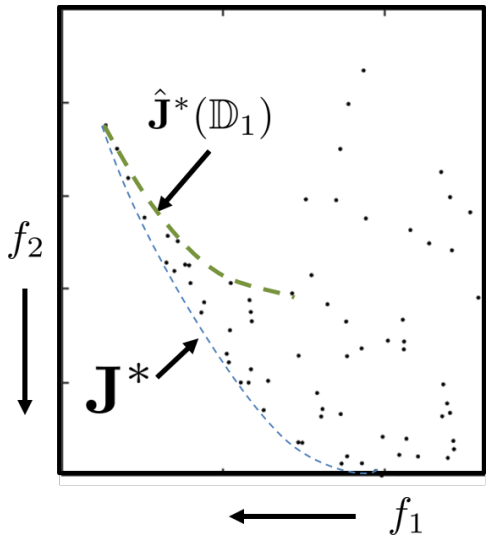

Figure 3.6: Each set $\mathbb{D}_i$ creates a Pareto frontier $\hat{\mathbf{J}}^*(\mathbf{x}_i)$ based on adaptive parameters $\boldsymbol{\alpha}$ while keeping the fixed parameters $\mathbf{x}$ approximately constant. The best $\hat{\mathbf{J}}^*$ is given as $\hat{\hat{\mathbf{J}}}^*$ and determined via section 3.1.1.

previously generated data (existing design points), which also drives the useful value of ϵ. The final selection determines the 'optimal' set $\mathbb{D}_i$ (from equation 3.6) via:

$$\hat{\hat{\mathbf{J}}}^*(\boldsymbol{\alpha}) = \min_{\mathbb{D}_i} \left(\text{Dist}\{\hat{\mathbf{J}}(\mathbb{D}_i), \mathbf{J}^*\}\right) \tag{3.14}$$

which provides a final approximate best adaptive design $\hat{\mathbf{x}^f} \rightarrow \mathbf{f}(\hat{\mathbf{x}^\mathbf{f}}, /\boldsymbol{\alpha}_\mathbf{f}{}^*)$. Alternatively,

$$\hat{\hat{\mathbf{J}}}^* = \{\hat{\mathbf{f}}^*(\mathbb{D}_i) \mid \min_{\mathbb{D}_i} \parallel \mathbf{d} \parallel\} \tag{3.15}$$

can be used to solve the adaptive design problem as defined in section 3.1.1.

The approximation of local frontiers approach from section 3.2.2 introduces a selection procedure to determine sets of similar designs $\mathbb{D}_i$ based on fixed design variables $\mathbf{x}_i^f$ to determine the best adaptive Pareto frontier $\hat{\hat{\mathbf{J}}}^*(\mathbb{D}_i^*)$. This selection procedure functions as a post-processing search algorithm which can be applied to any dataset consisting of fixed and adaptive design variables. The proposed algorithm could select the best approximate adaptive design or inform a future optimization with preferred fixed design variables. An optimization to determine the best adaptive design variables, conducted after the selection procedure determines such preferred fixed design variables, would ensure that the advantages of morphing are maintained as the data density criteria approaches zero ($\epsilon \longrightarrow 0$).

3.2.3 Option 3: *In Situ* Adaptive Optimization

The proposed design space decomposition into adaptive and fixed design variables as a post-processing selection procedure described in section 3.2.2 occurs after an optimization or other design space exploration has been completed. To continue, we now consider that the previous formulation could also serve as an evolutionary criteria during optimizations: designs exhibiting strong adaptive potential based on of similar designs $\hat{\mathbf{J}}^*$ are preferred for future generations in evolutionary design and optimization. In a genetic optimization context, implementing these two options occurs at different points as shown in Fig. 3.7. For example, NSGA-II utilizes a selection criteria for multi-objective optimization in which, before genetic crossovers and mutations, members used to generate the next generation are selected based on two criteria to be described: rank and crowding distance. To determine *rank*, a given population is sorted into frontiers, starting with the non-dominated Pareto frontier (rank 1) and continuing to sort additional frontiers until all designs are members of a ranked frontier. When selecting members of the next generation, NSGA-II

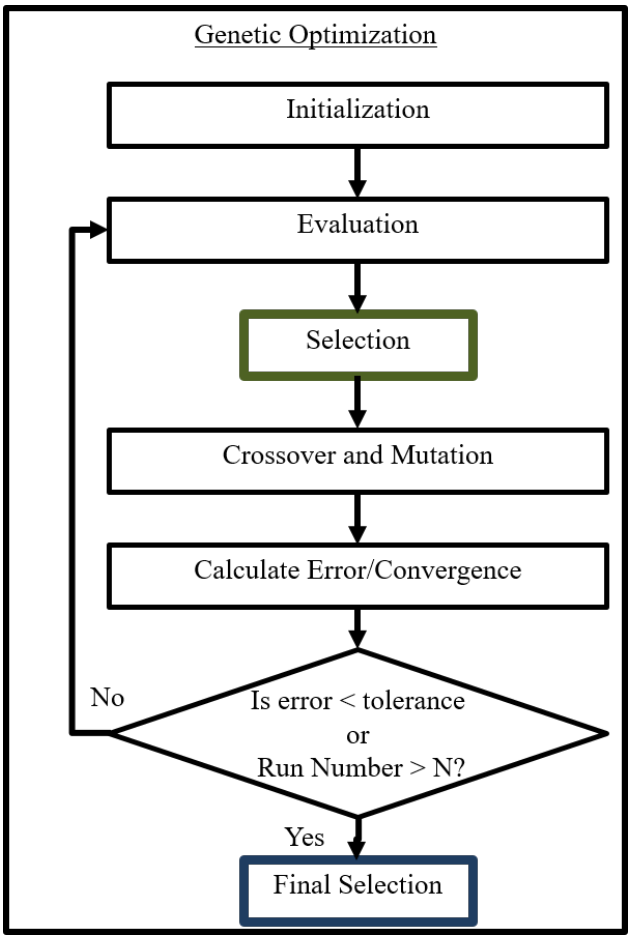

Figure 3.7: Genetic algorithms consider crossover and mutation of genes (in this case design variables), and then selection of new designs for the next generation. The adaptive design methodology laid out in section 3.1.1 can be applied to select optimal designs either *during* or *after* an optimization.

first selects members based on rank. However, typically the algorithm must select between members of the same rank to fill the final positions for the next generation; if four individuals must be selected from a set of 7 designs of equal rank, a additional distinction is required. For this purpose, crowding distance differentiates between designs of the same rank in the selection process. The *crowding distance* is defined by the spacing in the objective space between the objectives of an individual and the nearest other individual of the same rank.

In a *in-situ* adaptive optimization, a new selection criteria could also (or alternatively) include

the proposed adaptive values from equation 3.6 and quantified using the distance metrics from section 3.1.1 and use these methods to select between designs of equal rank based on adaptivity. An extension of NSGA-II, titled Non-dominated Sorting Genetic Algorithm for Adaptive Design (NSGA-AD) has been developed in the course of this work and is then explored considering the best adaptive designs based on the best adaptive sets $\mathbb{D}_i$ in each generation as previously described. The morphing-based evolutionary application applies the adaptive performance metrics from section 3.1.1, and may also consider the original crowding distance from NSGA-II. The selection method, shown in 3.8b, begins with the current parent and offspring generations P_t and Q_t, respectively, selecting designs for the next generation with the best ranks (e.g., F_1, F_2, and a portion of F_3). Next, when designs of similar rank must be distinguished for selection, the best approximate adaptive designs are determined and given a distance from the approximate ideal Pareto frontier $\mathbf{J}^*$ for the entire generation. Then designs of equal rank are selected based on the distance between approximate local frontier $\hat{\mathbf{J}}^*$ and the approximate Pareto frontier $\mathbf{J}^*$ as in equation 3.14. Thus, the evolutionary optimization prioritizes well-performing designs across multiple objectives that also posses high adaptive potential based on sets of similar fixed design variables as described in section 3.2.2. Since the algorithm still prioritizes non-dominated designs via the rank metric, selection still encourages elitism across the approximate Pareto frontier before considering adaptivity, but the effect of such a criteria on diversity still must be considered.

3.2.4 Relationship to Parameterized Optimization

In the works by Malak [50, 32, 166], the parameterized optimization problem is stated as:

$$\mathbf{J}^*(\boldsymbol{\theta}) = \min_{\mathbf{x}} \ (\mathbf{f}(\mathbf{x}, \boldsymbol{\theta})), \tag{3.16}$$

which appears mathematically equivalent to the fixed design adaptive optimization statement from equation 3.5 if $\boldsymbol{\theta}$ is strictly taken as an input variable of $\mathbf{f}$ (not necessary in the method of Malak, as output parameters are considered). Casting this in the current notation:

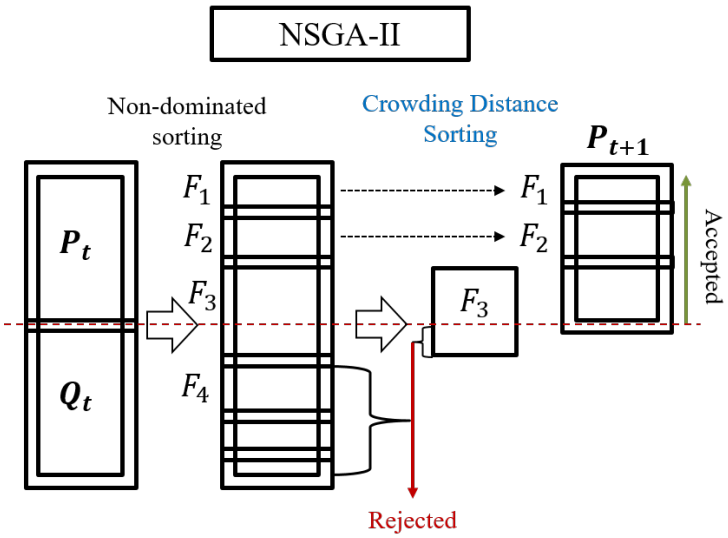

(a) The common existing genetic selection algorithm NSGA-II selects designs from the parent (P_t) and offspring (Q_t) designs by first ranking the non-dominated frontiers (F_i), selecting the best ranks, then differentiating between designs of the same rank by sorting based on crowding distance in the objective space.

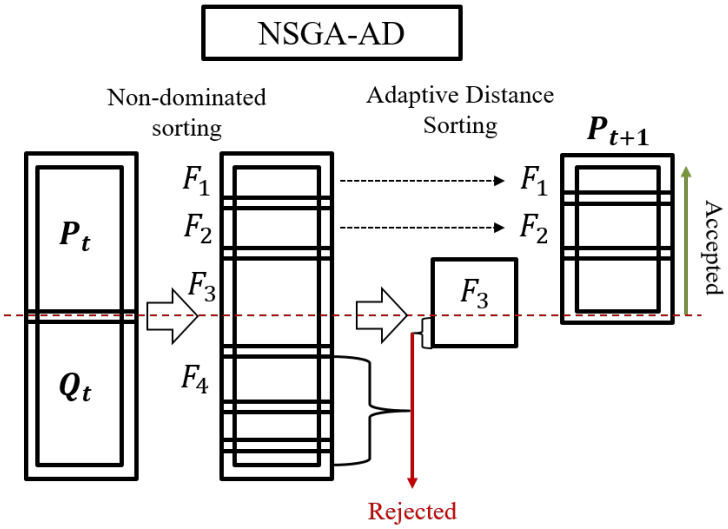

(b) The proposed selection criteria for an *in situ* adaptive optimization considers frontier rank, then adaptivity distance as described in section 3.1.1.

$$\hat{\mathbf{J}}^*(\mathbf{x}^f) = \min_{\boldsymbol{\alpha}} \ \mathbf{f}(\mathbf{x}^{\mathbf{f}}, \boldsymbol{\alpha}). \tag{3.17}$$

Finding the preferred single adaptive system design then follows the same form as equation 3.6 from section 3.1:

$$\hat{\hat{\mathbf{J}}}^* = \min_{\mathbf{x}^f} \ (\mathrm{Dist}\{\hat{\mathbf{J}}^*(\mathbf{x}^f), \mathbf{J}^*\}). \tag{3.18}$$

This approach does not provide an alternative to the local adaptive frontiers approach developed in this work, since the algorithms previously proposed by Malak et. al. for finding $\hat{\mathbf{J}}^*(\boldsymbol{\theta})$ do not consider preferred parameters, while the optimization approaches presented in this chapter search for the 'best' fixed design variables throughout the optimization. Similarly, the work described in this chapter does not provide an alternative to the parameterized optimization problem because the local adaptive frontiers approach does not explore fixed design variables ($\mathbf{x}^f$, $\boldsymbol{\theta}$) evenly or exhaustively, prioritizing 'good' fixed design variables or parameters.

3.3 General Demonstration of Adaptive Design and Optimization

The design space decomposition, adaptive design performance metrics, and optimization methods described in this chapter are broadly applicable to any multi-objective design and optimization problem in which the design space consists of variables that can and cannot change based on the different objectives. This is particularly applicable to mission-driven designs with separate stages, but can also be applied to analytical problems with competing objectives and a design space composed of inputs that both can and cannot be altered between objectives. In this section, these methods are applied to general problems to demonstrate applicability and initial implementation.

A simple physics-based adaptive structures design problem was created to demonstrate the utility of the selection criteria and optimization methods laid out in sections 3.1.1 and 3.2. The utopia point method determines the optimal adaptive configuration considering both fixed and adaptive variables in a multi-objective design context.

3.3.1 Euler-Bernoulli Beam-SMA Wire Adaptive Design Demonstration

Consider a 100 mm long cantilevered Euler-Bernoulli beam for which the designer prefers two unique shapes depending on the loading condition at the end of the beam. A shape memory alloy wire is attached to the wall 5 mm below the beam, its other end being attached to location L on the cantilevered beam that is set when manufactured. The SMA wire is given initial recoverable strain that can be recovered by heating the wire, transforming the wire from martensite to austenite and in doing so, will contract by up to 5% of its length [93]. The amount of strain recovery can be altered as needed to meet the objectives. In this case the strain ε is adaptive (α) while the wire-beam connection location L is fixed ($\mathbf{x}^f$). The objective shapes and loading conditions can be seen in Fig. 3.9.

The beam geometric and material properties are given in Table 3.1. Assuming an initially straight beam and using small angle approximations ($\sin(\theta) \approx \theta,\ \cos(\theta) \approx 1$), this problem is reduced to a 1-D thin beam solution in which local curvature ρ is defined by the local moment M such that:

$$\frac{\partial^2 y}{\partial x^2} = \frac{1}{\rho} = \frac{M}{EI}, \tag{3.19}$$

where E and I are the elastic modulus of the material and the cross-sectional moment of inertia, respectively.

In this example, the goal shapes are known to be contradictory and not attainable using a common wire attachment point. Since the objective shapes were derived using different connection locations L, and thus there is no combination of a single ε and L to exactly match both shapes. In fact, condition 1 can only be matched when $L = 100$ mm, while condition 2 can only be matched when $L = 60$ mm. Ideally, a single design with fixed variable L and two strain states ε allows for morphing from the unmorphed, initially straight configuration to shape 1 and shape 2 almost exactly. Thus, an optimization method is needed to determine the multi-objective set of adaptive variables ε and fixed variable L which minimize the shape errors across a range of values.

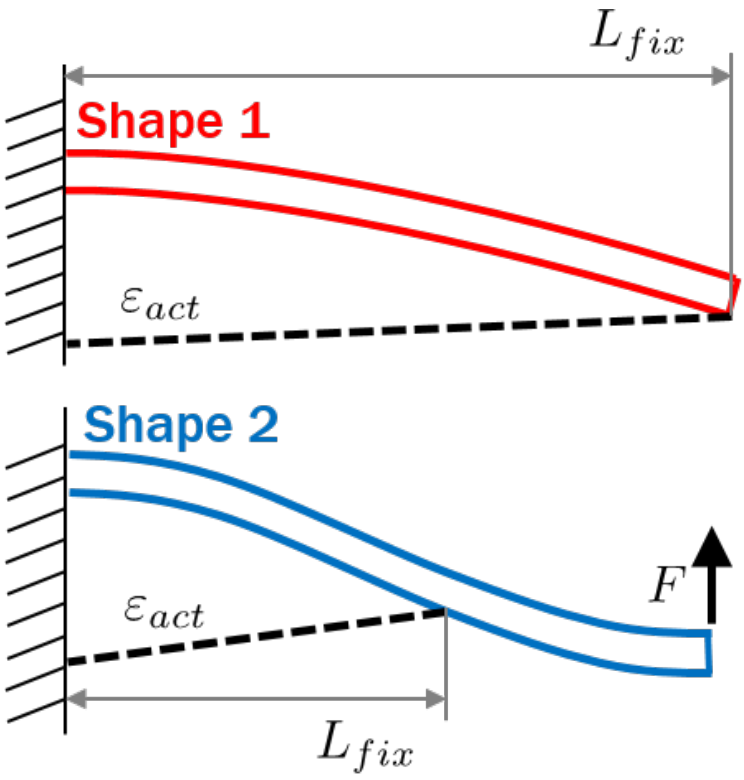

Figure 3.9: Two unique objective shapes for an Euler-Bernoulli beam defined by two loading conditions.

Table 3.1: Euler-Bernoulli Beam Parameters

PROPERTY	VALUE	DESCRIPTION
w	5 mm	Beam out of plane width
h	5 mm	Beam height
E	90 GPa	Material elastic modulus
L_0	100 mm	Beam length
P	200 N	Condition 2 End Load

The design space was initially explored via Latin Hypercube Sampling (LHS) over all fixed and adaptive variables. For each design, a single fixed and a single adaptive parameter were defined

by the LHS and return an error evaluation for each of the two shape objectives based on the mean squared error between the target and computed beam shapes.

A clear trade-off between the best designs to match each shape can be seen in Fig. 3.10, which shows the approximate Pareto frontier $\mathbf{J}^*$. An ideal design fully spanning $\mathbf{J}^*$ would attain the approximately best shape for each objective, located at the utopia point.

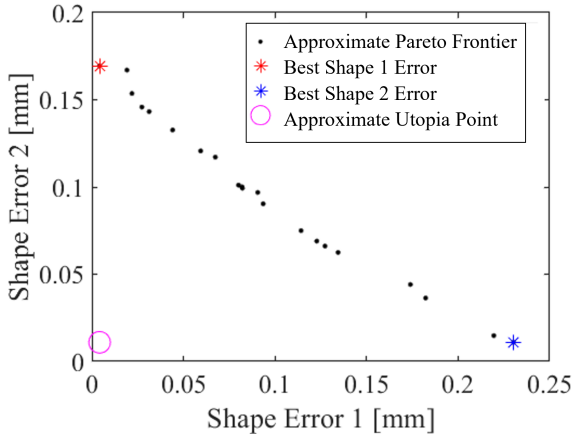

Figure 3.10: An ideal design would morph from the best design for Shape 1 to the best design for Shape 2, with objective values at the utopia point.

For each set of fixed design variables $\mathbf{x}_i^f$ evaluated to generate $\mathbf{J}^*$, a corresponding set of 'similar' designs $\mathbb{D}_i$ is determined. The set of designs within range ε represents a group of designs assumed feasible via adaptivity, and the best fixed design set $\mathbb{D}_i$ was determined via the utopia point method described in section 3.1.1. The set $\mathbb{D}_i$ with the shortest distance between local frontier utopia point $\hat{\mathbf{u}}_i$ and the approximate utopia point $\hat{\mathbf{u}}$ was selected, with the best fixed and adaptive design variables shown in Fig. 3.11a. A range of similar fixed parameters determined the adaptive design nearest to the utopia point via morphing. The distance between utopia point and best adaptive design only considers the extrema of the full approximate Pareto and local frontiers. It can be seen in Fig. 3.12 while the Pareto frontiers differ, the best adaptive design is near-optimal. The addition of an adaptive design variable allowed for the best realizable design to minimize the

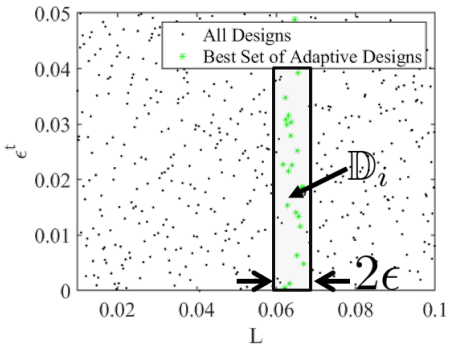

(a) The best set $\mathbb{D}_i$ consists of designs with a full range of adaptive design variables $\boldsymbol{\alpha}$ but a small range of fixed design variables $\mathbf{x}^f$ assumed to be similar enough to facilitate morphing.

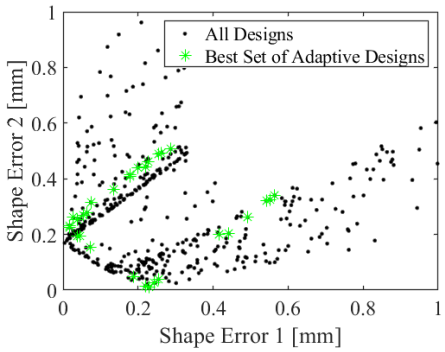

(b) The design and objective space, with emphasis on the best set of designs with a narrow range of L.

competing objective functions significantly better than a single fixed design, using the selection methodology for similar sets $\mathbb{D}_i$ from equation 3.11 to minimize the distance to the utopia point as described in section 3.1.1.

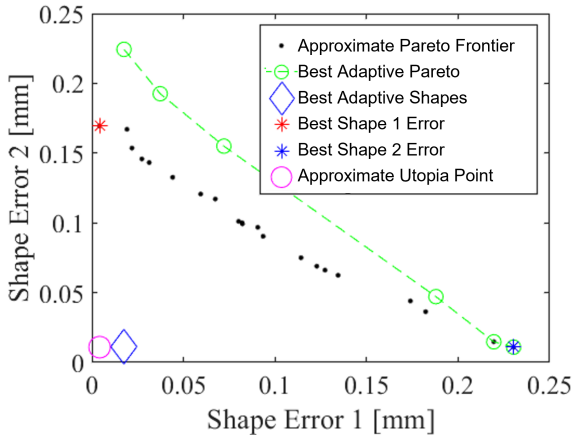

Figure 3.12: Adaptive and ideal Pareto frontiers, showing a design capable of morphing between two competing error minimization objectives.

After selecting the preferred fixed design variable L, a new optimization was performed to

determine if the assumption that "similar" designs (i.e., the designs in $\mathbb{D}_i$) were actually feasible. A final optimization determined the optimal adaptive variables, considering the fixed variable L given by the best $\mathbb{D}_i$.. This ensures the final designs are realizable, since the previous results had slightly different fixed variables L, within ε. The final morphed and target shapes can be seen in Fig. 3.13, with the final objective values in Table 3.2. In this example, the novel selection approach described in section 3.2.2 determines an adaptive design minimizing the distance between an approximate ideal design and an adaptive, realizable design as described in section 3.1.1. Next, this selection criteria is applied to a much more complex mission-driven adaptive rotorcraft optimization.

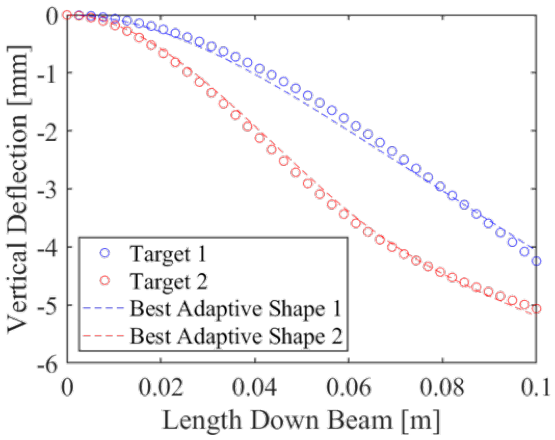

Figure 3.13: Final target and realizable shapes.

3.3.2 *In Situ* Optimization Examples

The *In Situ* optimization approach from section 3.2.3 is first demonstrated using mathematical functions as a method for selecting preferred adaptive designs *in-situ* during the genetic optimization. The procedure follows from NSGA-II in that non-dominated designs are important for improving multi-objective populations. After the initial generation, a parent generation (P_t) generates a following generation (Q_t) via genetic crossover and mutation. The designs selected to continue to the next generation are chosen based first on rank to promote elitism. However, gener-

Table 3.2: Final Shape Objectives

DESIGN	SHAPE 1 Error [mm]	SHAPE 2 Error [mm]
Best Shape 1	0.004	0.169
Best Shape 2	0.230	0.011
Approx. Utopia Pt.	0.004	0.011
Adaptive Design	0.018	0.011
Final Opt. Design	0.017	0.011

ally selections must be made between designs of the same rank. When selection must differentiate between members of the same rank, NSGA-II selects designs based on the distance between similar designs with respect to the objective space. NSGA-AD differentiates designs with the same rank based on adaptivity, either using the Utopia Point or Hausdorff distance metrics as described in sections 3.1.1 and E.

3.3.2.1 Three-Variable Adaptive Optimization

A set of periodic equations is used to explore and demonstrate the concept of an *In Situ* optimization where adaptivity is considered as a *selection* criteria and not a post-processing criteria. An optimization was defined in which three design variables, two fixed and one adaptive, are optimized to minimize two objective functions. As such, the two fixed design variables were inputs to competing periodic functions for both objectives, with some adaptive design variable coupling:

$$\begin{aligned} f_1(\mathbf{x}) &= \sin(\frac{\pi}{2}x_0)\alpha + \cos(5x_1) - \frac{e^{x_0}}{5}, \\ f_2(\mathbf{x}) &= \sin(\frac{\pi}{2}x_1)\alpha + \cos(7x_0) - x_1\alpha. \end{aligned} \tag{3.20}$$

The bounds of $\mathbf{x}$ are set as $\mathbf{x} \in \{0, 1\}$, and the properties of the design variables can be seen

in Table 3.3. A genetic algorithm was run with a population of 20 individuals with $\epsilon = 0.025$ with various selection algorithms.

Table 3.3: *In Situ* Custom Optimization Example Design Variables

Design Variable	Lower Bound	Upper Bound	Adaptive
x_0	0	1	No
x_1	0	1	No
x_2	0	1	Yes

Two reference genetic selection algorithms were compared alongside the methods developed in this work: NSGA-II and SPEA-II, both Pareto-optimally based well established multi-objective selection algorithms. Since both these criteria base selections toward designs spanning the Pareto frontier, both should explore the objective space effectively and provide an effective benchmark for this current work. In a case where the Pareto frontier is composed entirely of adaptive designs, NSGA-II should effectively search the objective space, and the approximate local frontiers approach from section 3.2.2 could then be applied to the entire population in post-processing for final selection (via the blue box in Fig. 3.7). However, if the best design is not entirely along the Pareto frontier, such a search may be less effective.

Objective scaling can also be performed via normalization either based on the current approximate Pareto frontier bounds defined either before the optimization or calculated every generation, or set such that each objective is scaled equally. In this study, the equations were defined such that the objectives were similar in magnitude and no additional scaling was applied.

Each optimization was initialized randomly using DEAP (Distributed Evolutionary Algorithms in Python) [48, 82], a Python optimization toolbox that allows for custom mutation, crossover, and selection algorithms. Each optimization was run 50 times to account for inherent randomness of the initialization, mutation, and crossover processes. Customizing the selection algorithm was nec-

essary to consider adaptive design, but the crossover and mutation parameters were kept constant across all criteria for an equal comparison. The crossover probability of a selected gene was 0.8, while the mutation range for a given selected gene was 5% of the design space.

The best adaptive design was evaluated after each generation using the Utopia Point method with $\epsilon = 0.05$ to track the convergence of best adaptive designs, as well as the size of each 'family' of similar designs. The results for NSGA-II can be seen in Fig. 3.14, where the black lines are the averages across every optimization, and the shaded regions are, for each generation, the mean plus/minus the standard deviation.

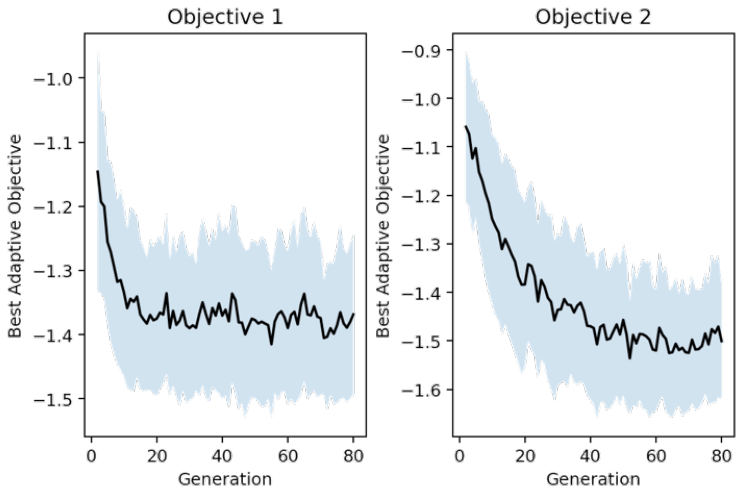

Figure 3.14: The NSGA-II algorithm searched the objective space, and the adaptive design each generation converged to $\mathbf{f}(\mathbf{x}) = \{\mathbf{-1.35}, \mathbf{-1.5}\}$.

Additionally, the size of the best adaptive family was recorded for each generation, as shown in Fig. 3.15. With an ϵ value of 0.025, the maximum fraction of the design space a single 'family' could consider would be 5% (or $2\ \epsilon$), or 1 in 20. If a single fixed design variable yielded the best Pareto-optimal multi-objectivity performance, it is likely the best adaptive design would contain a large 'family' of similar designs. In this case, however, the best adaptive design typically considered morphing between 2 or 3 designs.

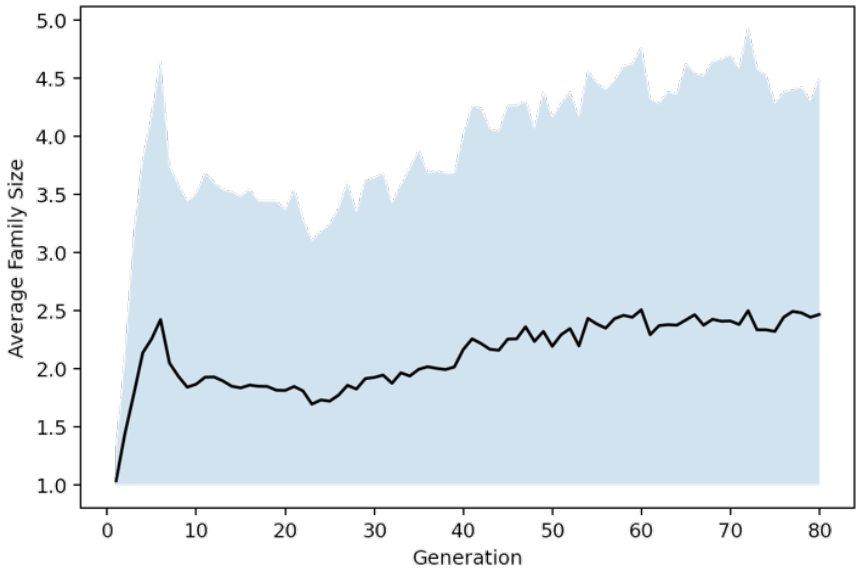

Figure 3.15: The size of the optimal adaptive families after each generation typically remained between 2 and 3.

Next, the multi-objective design problem was optimized using the same population, crossover, and mutation parameters with the NSGA-AD selection approach from Fig. 3.8b. In this procedure, the non-dominated sorting is applied the same as NSGA-II and described in section 1.5.1. The adaptive distance sorting can use either the Utopia Point method or the Local Pareto Frontier method. For both methods, the best designs in a given rank are selected by minimum distance between the approximate ideal and adaptive design. The Utopia Point method is used as the selection algorithm for each generation. Each selection prioritizes the rank of each design in the objective space, then the distance from the approximate ideal utopia point to the best adaptive utopia point.

While NSGA-AD still encourages elitism (non-dominated designs), preference to adaptive designs prioritizes adaptive over space-filling objectives. In an optimization where the best designs all reside near the Pareto frontier, simply searching the Pareto frontier is likely sufficient in post-processing to determine advantageous adaptive designs, as will be shown later in this work. However, an *In Situ* optimization approach considering adaptivity *during* the optimization may provide benefits beyond simply post-processing data generated for $\mathbf{J}^*$ using traditional multi-objective op-

timization techniques for these reasons:

- it cannot be assumed that all adaptive design problems will yield solutions on or near the Pareto frontier,
- the addition of a second stage in the design and optimization problem may prove difficult when data from intermediate stages may be less accessible in large scale engineering problems, and
- a generational population selected with a preference towards good adaptive designs may lead to larger 'families' of similar designs in the objective space, encouraging more feasible designs.

The best adaptive designs and 'family' sizes for each generation (averaged over 50 optimizations) can be seen in Fig. 3.16. This adaptive selection algorithm using the NSGA-AD selection algorithm and Utopia point distance metric approached an adaptive objective function $\mathbf{f}(\mathbf{x}) = \{-\mathbf{1.5},\ -\mathbf{1.6}\}$ during *each* generation, where the best adaptive design consisted of a 'family' of 6-7 similar designs with respect to the two fixed design variables x_0 and x_1. These objectives are lower in value than the adaptive designs found using NSGA-II and generated larger adaptive families of similar designs.

Next, NSGA-AD was implemented using the Hausdorff distance metric to select between designs of the same rank. The Hausdorff distance approach performed similarly to the Utopia Point metric, (as seen in Fig. 3.17. even though the adaptive objectives were evaluated as a single point, not the entire distance between frontiers used in this selection process.

For a given number of generations N, the first half would select designs using NSGA-II, with the final half using NSGA-AD. Thus, the optimization for equation 3.20 was performed using first a space-filling selection algorithm followed by an adaptive selection using both the Utopia Point and Hausdorff distance metrics.

The results can be seen in Fig. 3.18, and behave similarly to the previous approaches during their respective generation used. From these results, there appears to be little benefit to starting

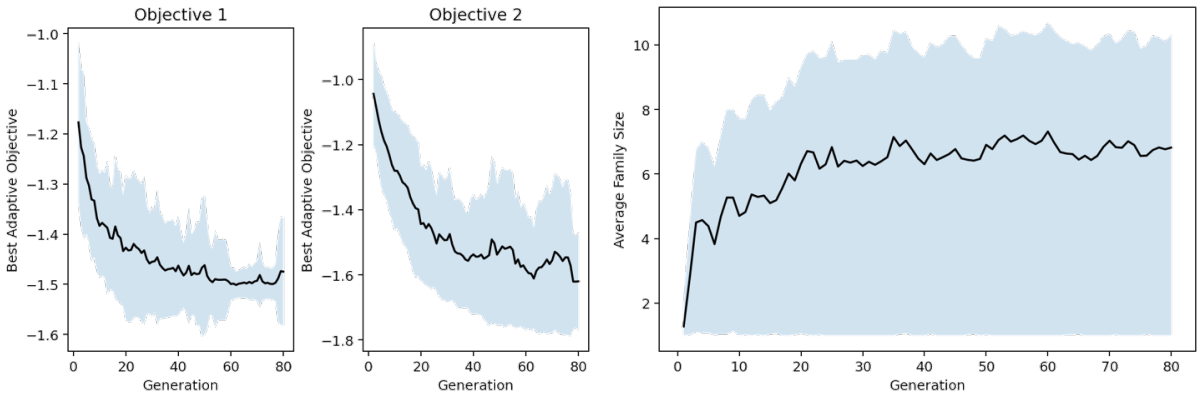

Figure 3.16: The best adaptive designs selected using the Utopia Point NSGA-AD during the optimization process converged toward a design with an adaptive objective function $\mathbf{f}(\mathbf{x}) = \{-\mathbf{1.5},\ -\mathbf{1.6}\}$ with an adaptive family size in each generation of 6-7 designs.

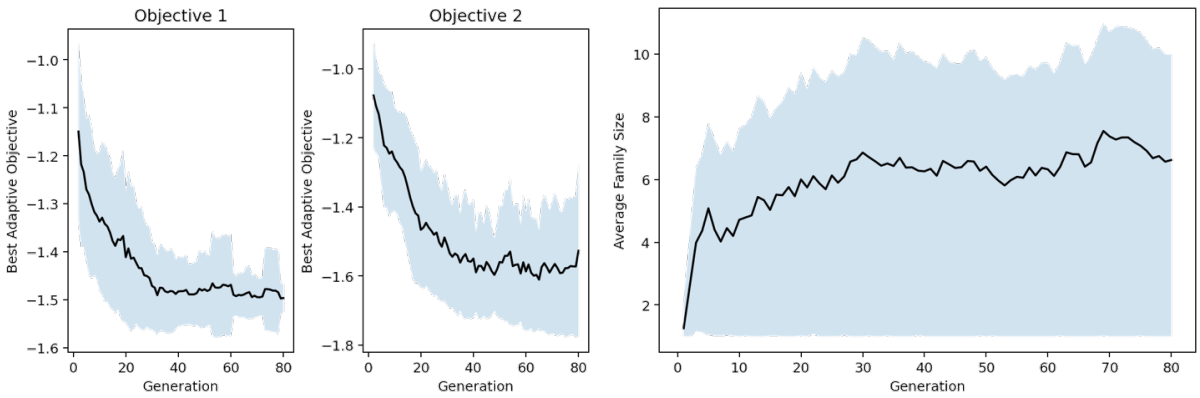

Figure 3.17: NSGA-AD using the Hausdorff distance metric for design selection performed similarly to the Utopia Point method, with adaptive designs near $\mathbf{f}(\mathbf{x}) = \{-\mathbf{1.5},\ -\mathbf{1.6}\}$, composed of adaptive families of 6-7 designs in each generation.

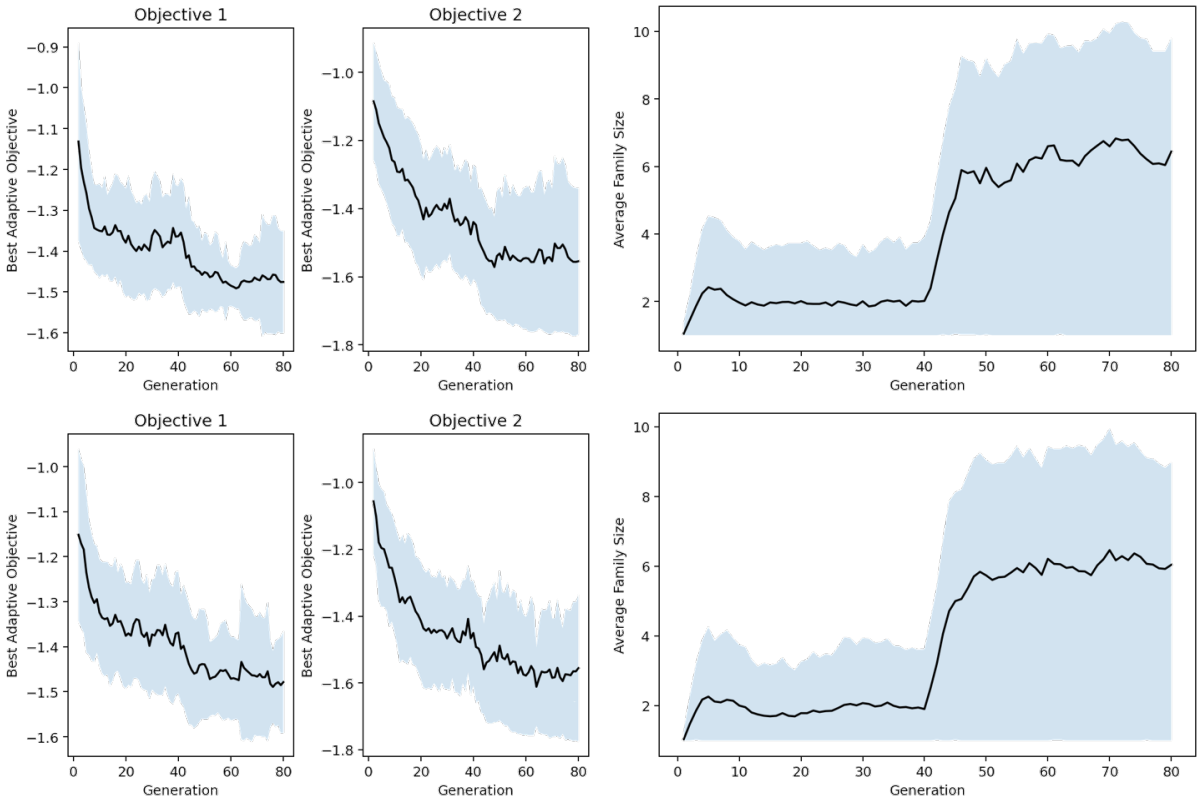

Figure 3.18: The NSGA-II-then-NSGA-AD hybrid approach resulted in a similar final result using both the Utopia Point (top) and Hausdorff (bottom) approaches for the second half of the optimization.

with NSGA-II for this specific problem before using NSGA-AD. In higher dimension problems such objective space exploration may be necessary, but NSGA-AD was sufficient to explore the objective space in this two-objective problem. The best adaptive design converged slower using NSGA-II then converged toward similar results as the previous NSGA-AD approaches over the second half of the optimization, suggesting the initial NSGA-II objective space search did not improve the final results.

Finally, the previous five optimizations, along with another common multi-objective optimization selection algorithm SPEA-II, were compared for convergence in Fig. 3.19. It can be seen that NSGA-AD converges to lower values across both objectives for this adaptive optimization whether using the Utopia Point or Hausdorff selection. As expected, the hybrid approaches follow similarly

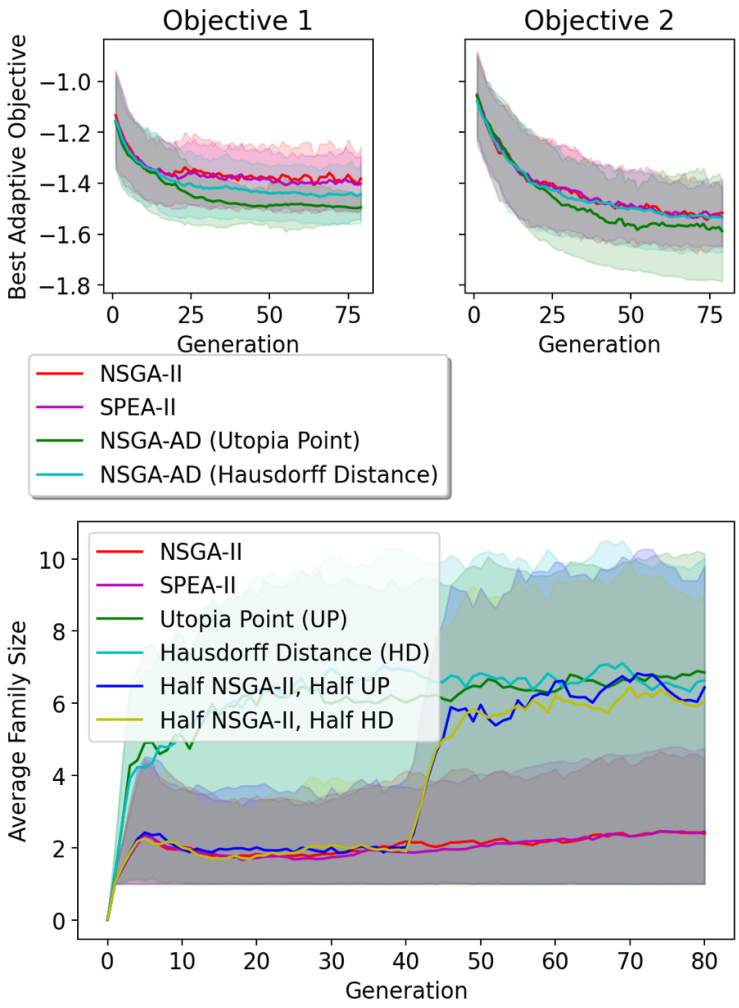

Figure 3.19: The adaptive designs after each generation using six selection criteria approaches for the *In Situ* optimization approach.

to the NSGA-II selection for the first 40 generations, then converge to similar results as the optimizations that utilized NSGA-AD methods for the entire optimization. The best adaptive family sizes were compared similarly, and it can be clearly seen that optimizations prioritizing adaptivity in the selection criteria resulted in larger sets of adaptive designs.

The distribution of adaptive designs for the final generation can be seen in Fig. 3.20. The NSGA-AD methods resulted in, for the each final generation, lower final Objective 1 results, with the Utopia point distance metric outperforming the other selection algorithms for Objective 2.

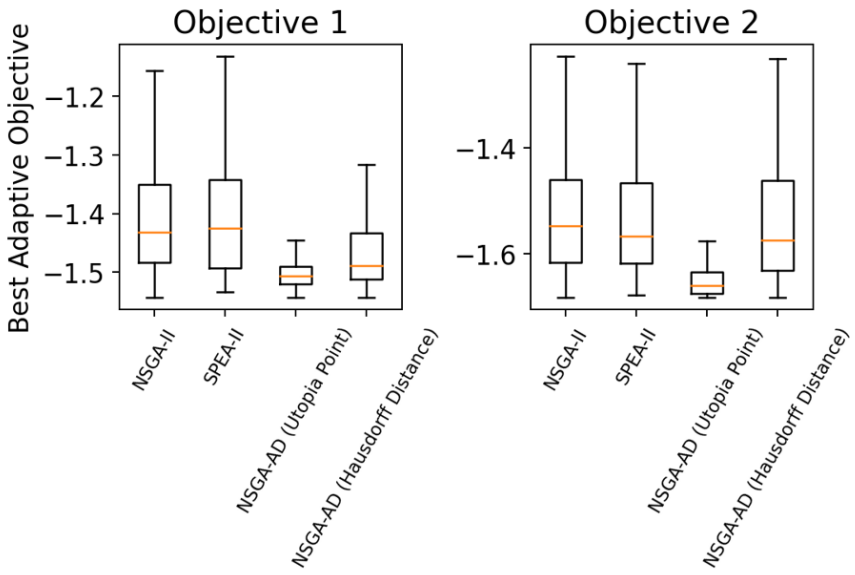

Figure 3.20: The distribution of final adaptive objectives using each selection criteria, where UP and HD represent the utopia point method and the Hausdorff distance method, respectively.

In this example larger families of adaptive designs allowed for improved objectives inside each generation using the *In Situ* optimization approach titled in this work NSGA-AD. Both the utopia point and Hausdorff distance methods applied to the prioritization of adaptive designs yielded similar results and improvements over more well-established traditional genetic optimization selection criteria. However, the number of designs in this optimization problem were small enough that most NSGA-II and SPEA-II generations resulted in a best adaptive design with a family of only two individual designs. Next another multi-objective design problem is taken from literature and run with a larger population.

3.3.2.2 Zitzler-Deb-Thiele Function 3

Another multi-objective function, the Zitzler-Deb-Thiele Function 3 (ZDT3), with a larger design space from existing literature [38] was tested as a multi-objective adaptive design problem by making some design variables adaptive, while others remained fixed. The function was selected

due to an interesting Pareto frontier, as shown in Fig. 3.21. The objectives to this problem are:

$$f_1(\mathbf{x}) = x_1, \tag{3.21}$$

$$f_2(\mathbf{x}) = g(\mathbf{x})h\left(f_1(\mathbf{x}), g(\mathbf{x})\right), \text{ where} \tag{3.22}$$

$$g(\mathbf{x}) = 1 + \frac{9}{29}\sum_{i=2}^{30} x_i \text{ and} \tag{3.23}$$

$$h(f_1(\mathbf{x}), g(\mathbf{x})) = 1 - \sqrt{\frac{f_1(\mathbf{x})}{g(\mathbf{x})}} - \left(\frac{f_1(\mathbf{x})}{g(\mathbf{x})}\right)\sin(10\pi f_1(\mathbf{x})). \tag{3.24}$$

An adaptive design and optimization problem was then specified where, for the 30 values in $\mathbf{x}$, each value was adaptive except for x_{28}, x_{29}, and x_{30}, meaning:

$$\mathbf{x}^f = \{x_1 ... x_i ... x_{27}\} \tag{3.25}$$

$$\boldsymbol{\alpha} = \{x_{28}, x_{29}, x_{30}\}. \tag{3.26}$$

The population was increased to 100 individuals and run for 100 generations. Optimizations using the selection algorithms NSGA-II, SPEA-II, NSGA-AD (using the Utopia Point metric) and NSGA-AD (using the Hausdorff distance metric) were all run 200 times. It was suspected that NSGA-II would perform better with a larger relative set of adaptive design variables, since when every design variable is adaptive, the idealized Pareto frontier can also be considered the best adaptive frontier.

The convergence of the best adaptive designs can be seen for each selection algorithm in Fig. 3.22. The mean and standard deviations were plotted for the best adaptive design for each method over every generation. While all three selection criteria converged almost identically for objective 2, NSGA-AD generated lower values for objective 1 using both the utopia point and Hausdorff distance methods. With a much larger number of adaptive design variables, the selection methods presented in this work found slightly better adaptive designs.

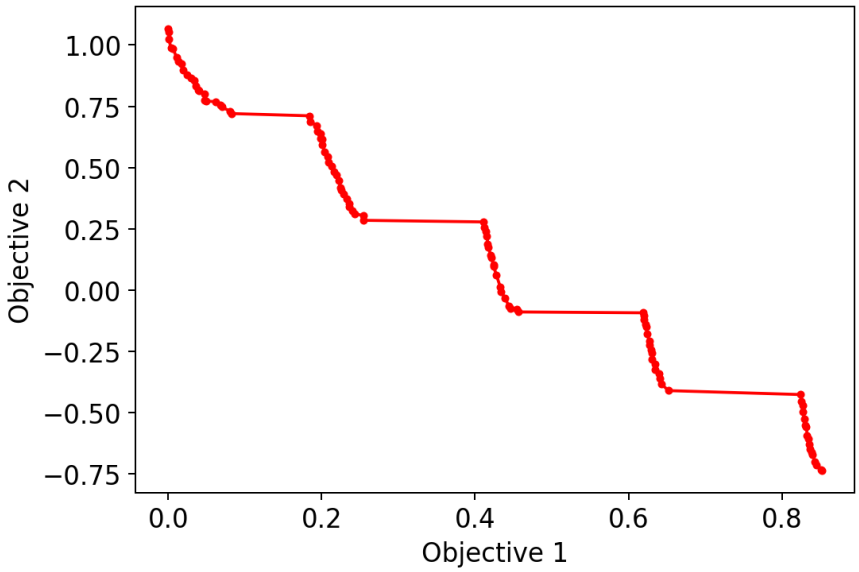

Figure 3.21: An example final population for the multi-objective Zitzler-Deb-Thiele Function 3, which features a stepped Pareto frontier.

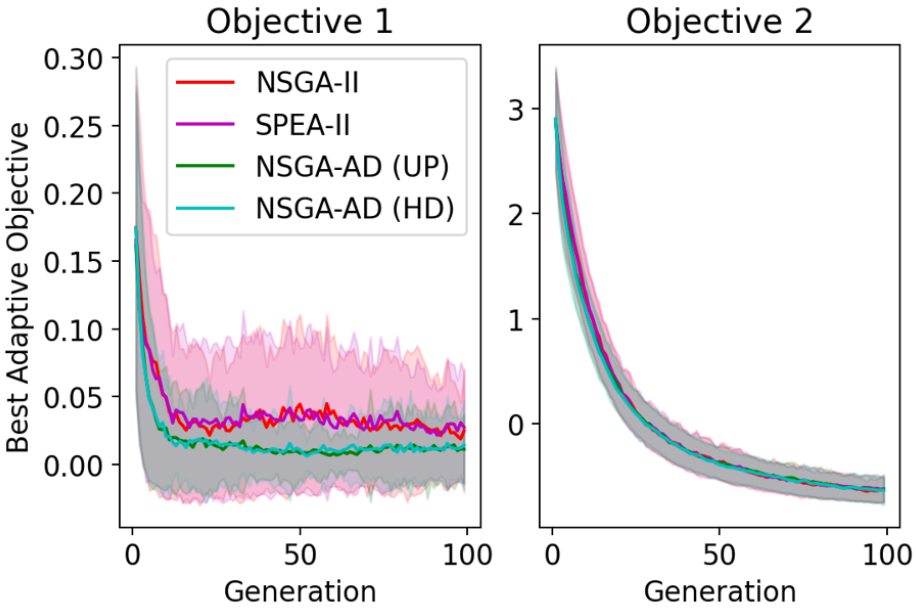

Figure 3.22: The NSGA-AD selection criteria generated populations with better inter-generational adaptive designs that converged faster than NSGA-II and SPEA-II.

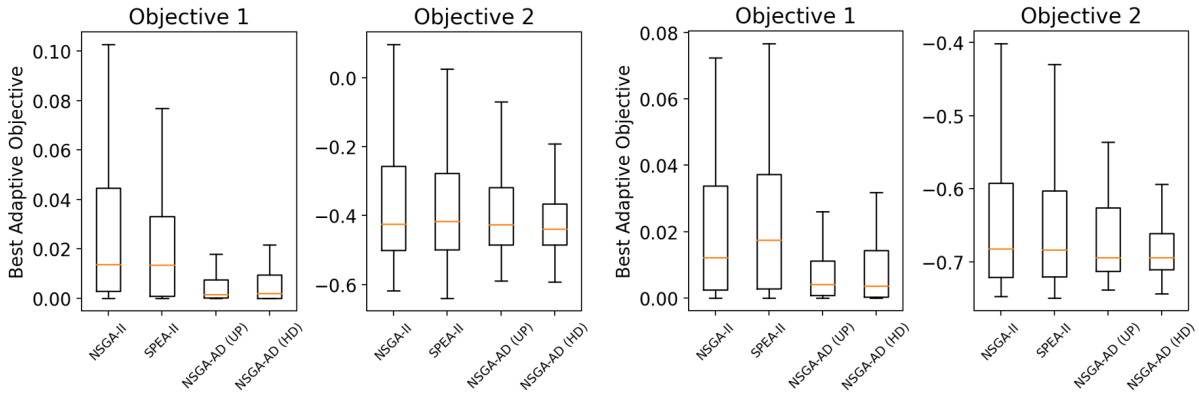

Figure 3.23: The NSGA-AD selection criteria resulted in lower adaptive designs across both objectives than NSGA-II and SPEA-II through 50 generations (left) and 100 generations (right).

The adaptive design objective distributions are shown in Fig. 3.23 for after 50 and 100 generations. The results are similar for objective 2 while using NSGA-AD with either the utopia point (UP) or Hausdorff distance (HD) results in lower values for objective 1 and a smaller variance across optimizations.

3.3.2.3 Two-Wire Analytical Euler-Bernoulli Beam

Next, the methodology and adaptive design tools in this chapter were applied to an analytical engineering adaptive structures problem using the *In Situ* optimization approach from section 3.2.3. In this example, a cantilevered Euler-Bernoulli beam has two SMA linear actuators attached along the span, anchored above and below the beam as shown in Fig. 3.24. The wires generate a strain when actuated and are attached at a distance of $h = \frac{L}{2}$ above and below the beam.

The beam parameters are shown in Table 3.4. The aluminum beam has a cross-section of 1 cm by 1 cm with a length of 1 m. A constraint is applied where the minimum distance between the two wire connections on the beam must be at least 0.05 m apart; the design problem is strain-driven and a singularity occurs when the connection points approach the same location.

The objectives in this design and optimization problem are two different shapes, as shown in Fig. 3.25. Both of these shapes are achievable via applied strain from the SMA wires, but the wires

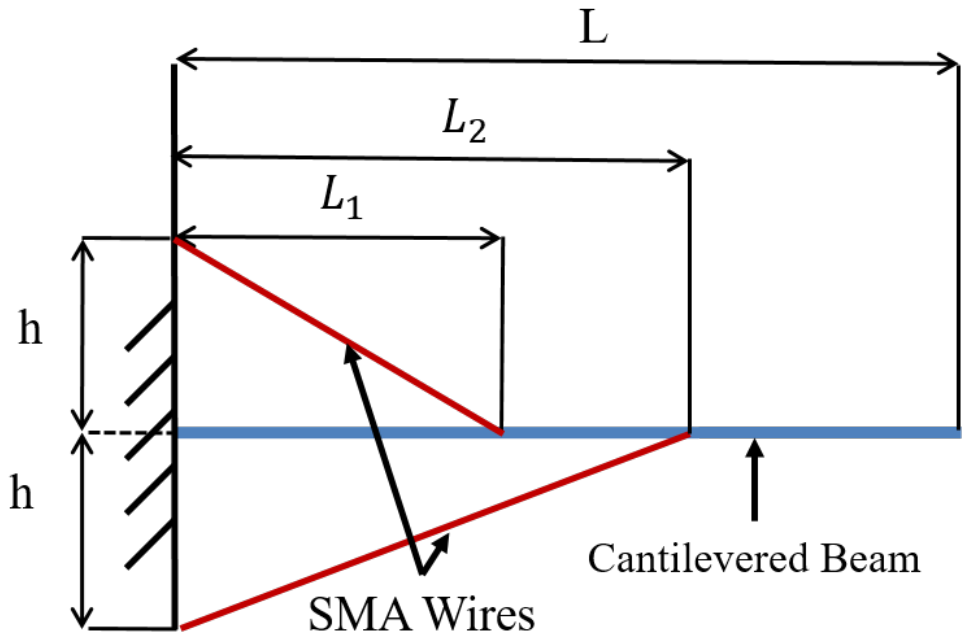

Figure 3.24: Consider a cantilevered beam with two SMA wires attached at two locations.

are attached to different locations on the beam for each shape as shown in Table 3.5 For the first shape the upper wire is attached 0.9 m down the beam with a total strain of 0.04, while the lower wire is attached at 0.3 m also with a total strain of 0.04. The second shape has upper and lower wire attachments at 0.3 and 0.8 m with strains of 0.03 and 0.01, respectively.

The analytical solution for the Euler-Bernoulli beam is determined via super-position. First, the strain from both wires is input and the vertical displacements at the connection points determined. Next, the vertical reaction forces at those two locations required to generate the vertical displacements are determined analytically; there are two equations (one for each displacement) and two unknowns (the forces on each wire connection). These forces can be determined analytically, then the entire beam profile is calculated for a set of evenly spaced points down the span.

First the vertical displacements ($\delta_1,\ \delta_2$) are determined as a function of the two strains (ϵ_1 and ϵ_2). Next the vertical displacements P_1 and P_2 are determined via the equations for each

Table 3.4: Two-Wire Beam Parameters

Parameter	Description	Units
Beam Height	0.01	m
Beam Width	0.01	m
Beam Elastic Modulus	70	GPa
Wire Attachment Height	0.5	m
Wire Radius	0.001	m

Table 3.5: Two-Wire Target Shape Definitions

Design Variable	Target Shape 1	Target Shape 2
L_1	0.9	0.5
L_2	0.3	0.3
ϵ_1	0.02	0.02
ϵ_2	0.02	0.00

displacement (where $L_1 \leq L_2$):

$$\delta_1 = \frac{P_1 L_1^3}{3EI} - \frac{P_2 L_1^2 (3L_2 - L_1)}{6EI}, \tag{3.27}$$

$$\delta_2 = \frac{-P_2 L_2^3}{3EI} + \frac{P_1 L_1^2 (3L_2 - L_1)}{6EI}. \tag{3.28}$$

By setting the constants to fixed values, the simplified equation becomes:

$$\delta_1 = P_1 c_1 + P_2 c_2, \tag{3.29}$$

$$\delta_2 = P_2 c_3 + P_1 c_4, \text{ where:} \tag{3.30}$$

$$c_1 = \frac{L_1^3}{3EI}, \tag{3.31}$$

$$c_2 = \frac{L_1^3 - 3L_1^2 L_2}{6EI}, \tag{3.32}$$

$$c_3 = \frac{-L_2^3}{3EI}, \tag{3.33}$$

$$c_4 = -c_2 = \frac{3L_1^2 L_2 - L_1^3}{6EI}. \tag{3.34}$$

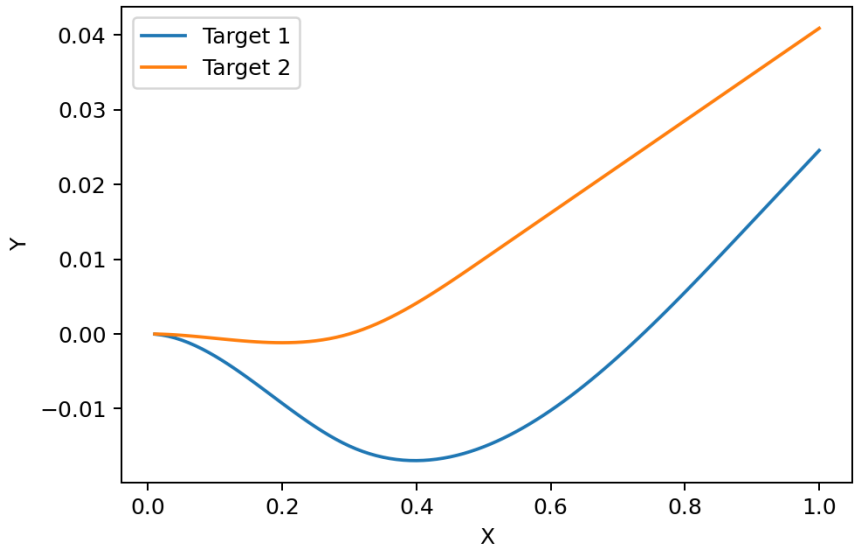

Figure 3.25: Two target shapes were defined by morphed beam shapes created via actuation, but with wires attached at different locations for each shape.

The forces (P_1 and P_2) are then calculated via:

$$P_2 = \frac{\delta_2 - \frac{\delta_1 c_4}{c_1}}{c_3 - \frac{c_2 c_4}{c_1}}, \tag{3.35}$$

$$P_1 = \frac{\delta_1 - P_2 c_2}{c_1}. \tag{3.36}$$

Once the forces are calculated, the vertical displacement can be calculated via the superposition of both forces applied at their respective points using the solution for both loads P applied at locations a:

$$\delta(x) = \frac{Px^2}{6EI}(3a - x) \qquad \text{for } 0 < x < a, \tag{3.37}$$

$$\delta(x) = \frac{Px^2}{6EI}(3x - a) \qquad \text{for } a < x < L, \tag{3.38}$$

where the total deflection is the sum of individual deflections from both P_1 and P_2.

A genetic optimization was performed with a population of 100 individuals over 100 generations to optimize 4 designs variables: two attachment locations for the SMA wires on the beam (L_1 and L_2) and 4 total strains in the SMA wire, one for each shape and wire. The strains in this design problem are adaptive, with a range from 0-6 percent, while the attachment locations are fixed and must be set in advance.

The objectives to be minimized are the error between the two target and actuated shapes by determining the best placement and actuation for the SMA wires. An additional constraint is applied to the SMA wire limited the maximum stress in the wire to 300 MPa, with a multiplicative penalty applied to both shape fitting objectives for violating the constraint. Each optimization was run 200 times using four different selection criteria (NSGA-II, SPEA-II, and NSGA-AD using both the Utopia point and Hausdorff distance metrics), and after each generation the best adaptive design was determined and recorded.

Additionally, another approach to this design problem is to consider each strain for each objective as a separate design variable. Instead of optimizing 4 design variables: 2 fixed and 2 adaptive, the adaptive variables for shape 1 and shape 2 are all considered as different variables in the design space. In this case, the design space considers 6 design variables: 2 fixed locations, 2 strains for shape 1, and 2 strains for shape 2. This approach considers all the design variables over all the objectives, and is referred to in this work as 'All DVs' for considering every adaptive design variable over every objective.

The convergence of adaptive designs in each generation using each method can be seen in Fig. 3.26, along with the distributions of the final adaptive designs using the post-processing described previously to determine the best final adaptive design. Using NSGA-II and considering all the adaptive design variables over every objective converged the fastest. This was due to the relatively small design space and population of 100 designs; this approach was sufficient to quickly explore the entire objective space and determine the best adaptive design despite the increased design space. It will be shown in the next study that this is not always the case, but the value of a exploring a larger design space is demonstrated for relatively large generational populations.

For the 4 design variable optimizations, NSGA-AD using the Utopia Point method resulted in the best adaptive shapes for both objectives, and a smaller variance for the final adaptive objectives. Using NSGA-AD with a Hausdorff distance metric resulted in a better objectives than exploring the objective space using NSGA-II and SPEA-II.

The size for each adaptive design at the end of each generation was also plotted for each selection criteria. As expected, the NSGA-AD methods resulted in larger sets of adaptive designs for each generation than the traditional selection criteria. Since NSGA-II and SPEA-II focus on exploring the objective space and encouraging diversity, it is worth comparing the entire populations generated during the optimization using the approximate local frontiers approach from section 3.2.2 as a post-processing selection over every generation.

Though the *In Situ* optimization approach outperformed the traditional multi-objective optimization techniques when considering only 4 design variables, the 6 design variable optimization quickly converged to the best adaptive shape fitting solution for two objectives. However, this approach scales with the number of objectives. Multi-objective designs (such as multi-stage missions) are rarely limited to only two objectives, so the same shape-fitting optimization was performed with two additional shape objectives. The target shapes can be seen in Fig. 3.27.

In this optimization, the design space for the optimization approaches in developed in this chapter remains in 4 dimensions, but considering every design variable for each objective now requires 10 design variables: 2 fixed wire locations and 2 strains for every objective. Both populations of 40 designs over 40 generations and populations of 100 designs over 100 generations were evaluated to compare the data density as well as the changes in objectives. The convergence of the best adaptive design for each generation for the 40 generation optimization can be seen in Fig. 3.28, along with the best final adaptive designs.

The increase in objectives reduced the effectiveness of considering every design variable for every objective, as both the convergence of adaptive designs and the final design considering every generation performed much worse considering four objectives as opposed to two. The NSGA-AD approaches performed the best over all four objectives as a whole, though NSGA-II found an

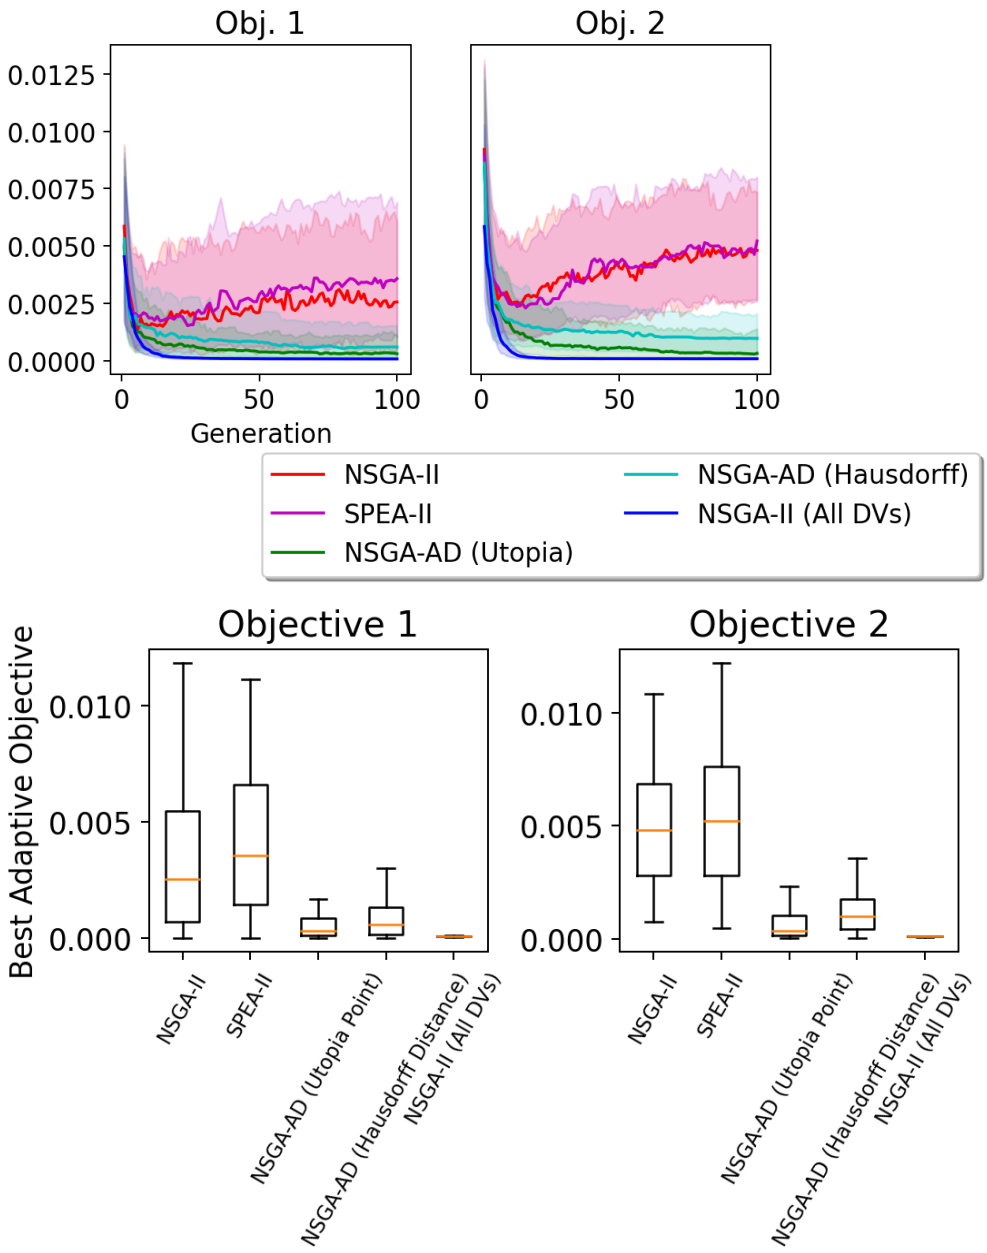

Figure 3.26: The selection algorithm NSGA-AD using the adaptive utopia point distance metric yielded better adaptive designs than the other selection criteria when considering only four design variables, but increasing the design space resulted in faster convergence and better adaptive designs by considering all design variables over all objectives separately.

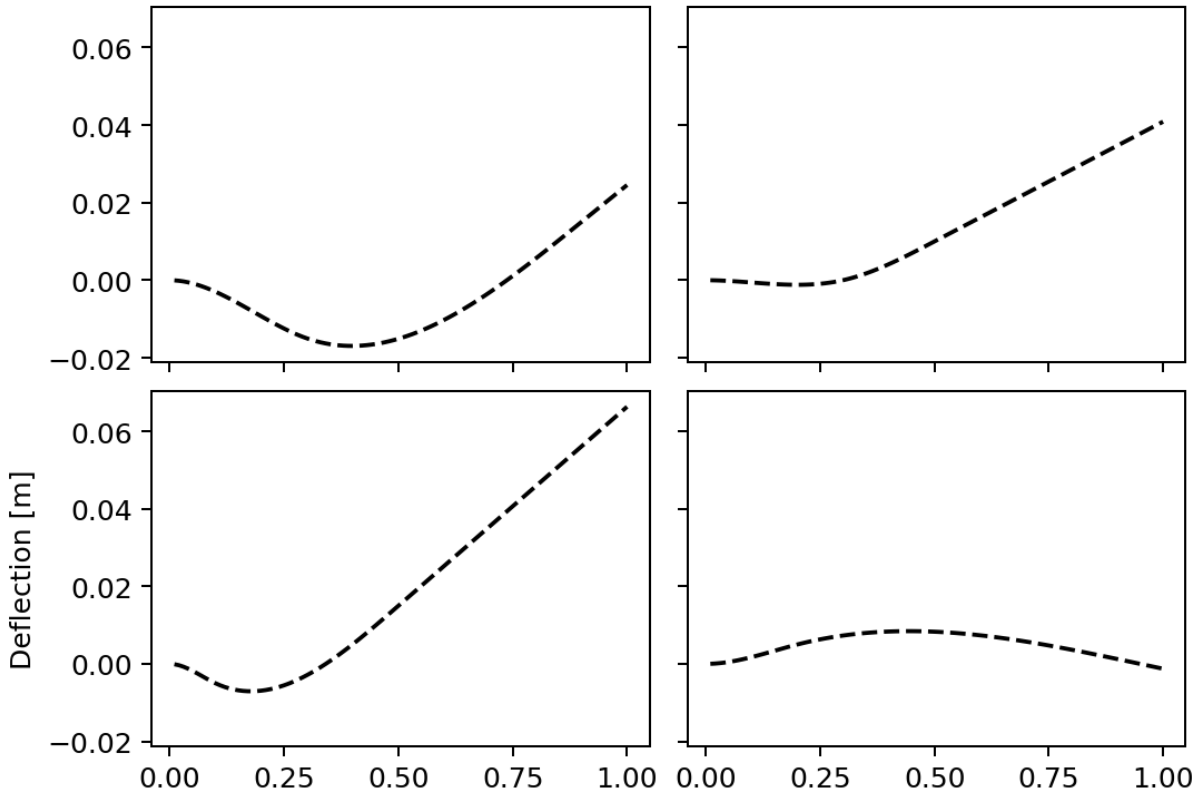

Figure 3.27: Four target shapes were defined for each set of optimizations to try and minimize the shape fitting errors using different selection criteria.

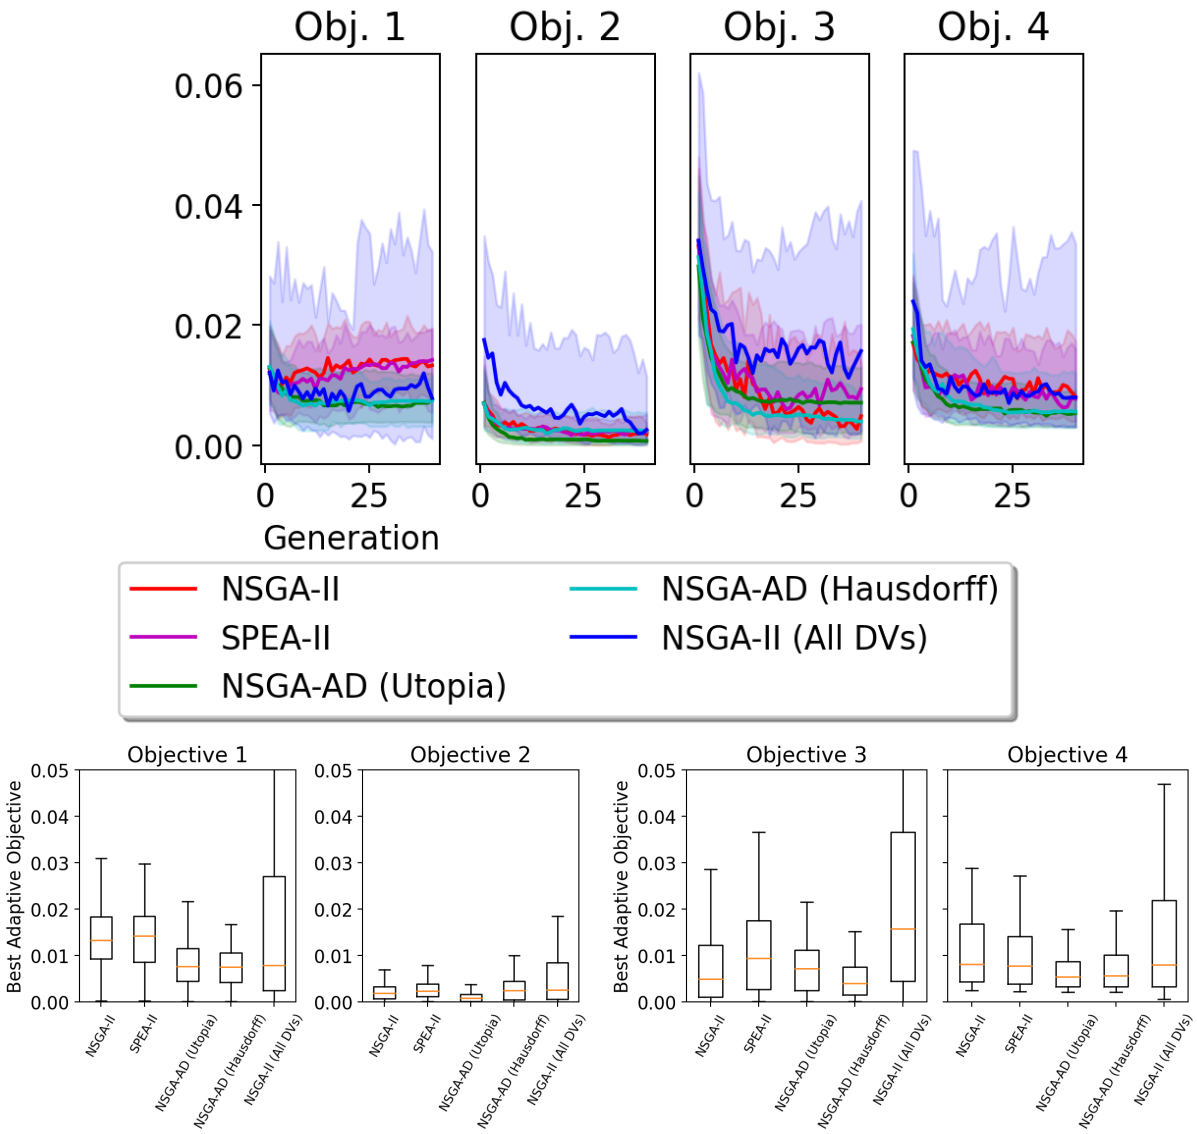

Figure 3.28: The best inter-generational adaptive designs converged the fastest over the 4 objectives using the NSGA-AD methods, though there were trade-offs between specific objectives. Considering every design variable for every objective led to much worse performance when compared with the previous two-shape optimization.

adaptive design with a comparable shape 3. The variance considering all DVs was much larger than the other approaches due to the larger design space. The results were similar for a 100 generation optimization of 100 individuals, as shown in Fig. 3.29, though the 10 design variable optimization performance improved when increasing the population from 40 to 100. Even with a population of 100 designs, the 'All DVs' approach had by far the higher upper distributions for all objectives except objective one, which performs similarly to NSGA-AD for that one objective.

Finally, the predicted final performance and actual final performance were compared as a function of ϵ for the four-shape Euler-Bernoulli beam final designs. The final objectives are scored based on the *Minkowski distance* where $p = 2$; in this case where the best theoretical objectives are zero the error metric is similar to a mean-squared error. As ϵ increases the predicted error should also decrease, but the assumption of similar fixed design variables decreases. A range of similarity parameters were evaluated using the approximate local frontiers approach to determine the best designs and predicted objectives. Then the selected designs were run with the best set of fixed design variables and best adaptive variables for each objective to compare the predicted and actual final performance.

The predicted and actual final performance for the best adaptive design sets as determined via the approximate local frontiers approach using the Utopia Point adaptive design metric are shown in Fig. 3.30. The lowest error metric, a *Minkowski distance* where $p = 2$, representing the best fully realizable designs can be found where $0.02 \leq \epsilon \leq 0.04$. Smaller ϵ-bands restrict the number of designs in a 'family', some of which are feasibly adaptive. Meanwhile, bands larger than 0.04 result in worse adaptive designs because the assumption of similar fixed design variables loses accuracy. Each decrease in predicted performance means increasing the ϵ-band range yielded a better adaptive design, while each decrease in actual performance due to increasing the ϵ-band range is the result of considering a new design thought to be fully realizable but was not. These trade-offs can be run fairly quickly, requiring only one evaluation for each value of ϵ to determine the best fully realizable design as determined via a single optimization (or other design space exploration technique as described in section 3.3.1). **Finally, if the best designs can be evaluated**

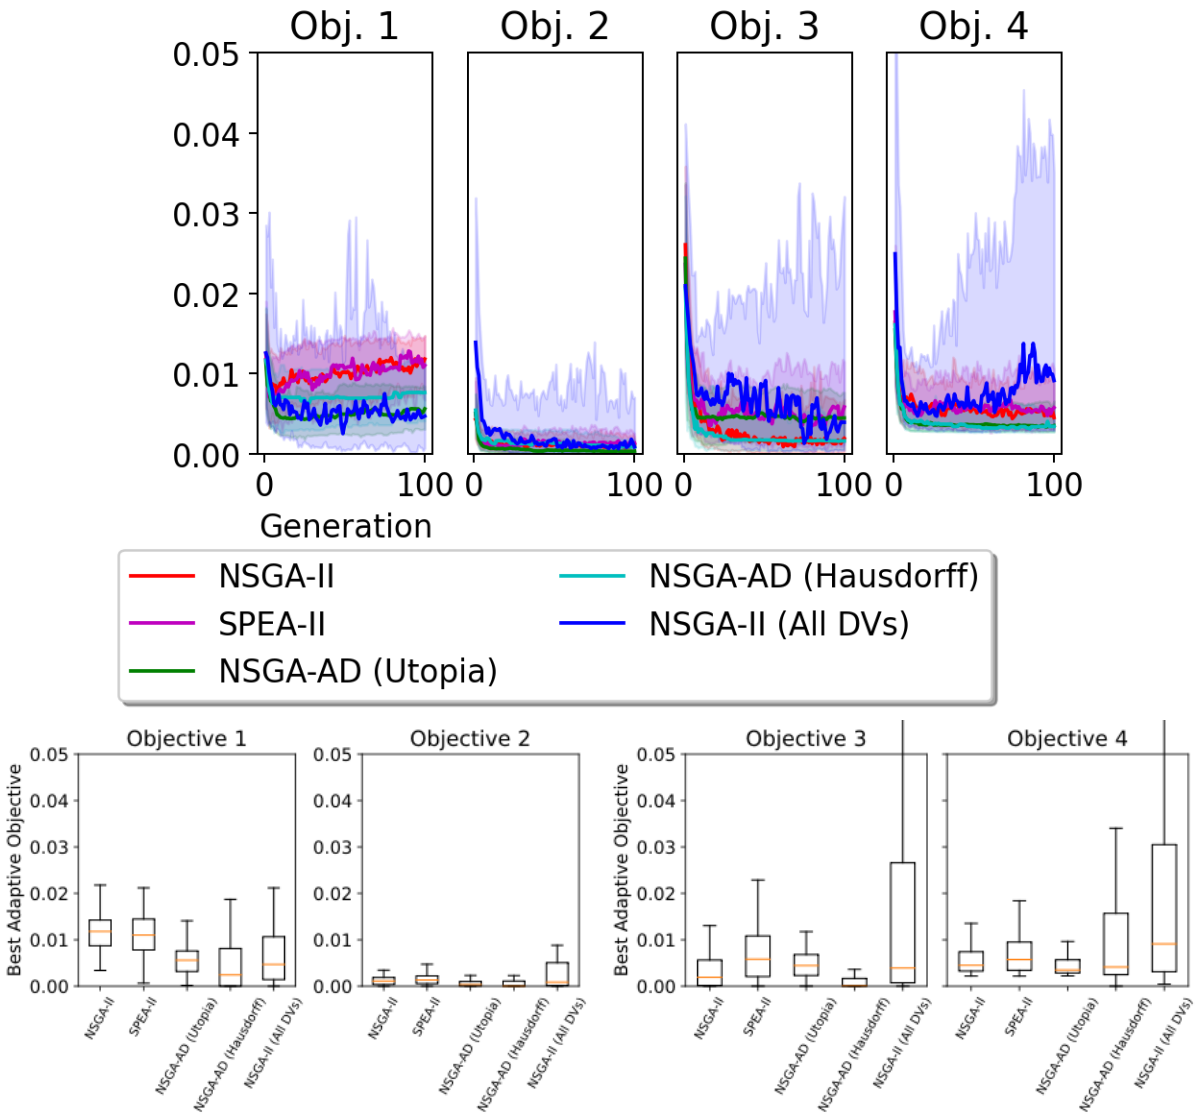

Figure 3.29: The trends with a larger population of 100 designs over 100 generations for 4 shape objectives were similar to the smaller optimization with 40 designs, though the larger population did allow the 10 design variable optimization better performance. Overall, the NSGA-AD selection criteria optimizations resulted in the lowest objectives.

over a range of ϵ values using the approximate local frontiers approach, there is no need to guess which similarity parameter yields the best results; the best design is directly selected from the final set of fully realizable designs.

3.4 Chapter 3 Summary

In this chapter, a methodology is developed to algorithmically determine good engineering designs for multi-objective performance when the design space is composed of fixed and adaptive variables. The design space is decomposed into variables that cannot be altered, labeled as 'fixed', and those variables that can be altered, denoted here as 'adaptive' or 'morphable.' A measure of design 'similarity' is developed to evaluate potential morphable designs, and two metrics for potential 'adaptive' designs' are formulated and tested.

Three optimization methods are laid out, with the two more practical approaches evaluated over mathematical and analytical engineering optimization designs problems formulated such that some design variables are fixed while other are adaptive. A post-processing procedure referred to in this work as the 'Approximate Local Frontiers' process was developed wherein after an initial optimization (or any design space exploration such as Latin Hypercube Sampling (LHS)) the best sets of similar designs is determined based on potential adaptive designs and the approximate Pareto-optimal objectives.

The methodology is further developed with the *In Situ* optimization approach in which adaptivity is consider *in-situ* during the selection process of a genetic algorithm. This selection criteria is compared with other well-known genetic algorithm selection processes NSGA-II and SPEA-II as applied to several mathematical and simple engineering design problems and performs well, particularly when applied using the Utopia Point method. This *In Situ* optimization approach can also be augmented with the final approximate local frontiers post-processing procedure to evaluate the best adaptive designs over every generation. These novel design techniques are compared to both conventional selection algorithms and adaptive structures problem formulations and show good performance.

There are clear trade-offs between the size of the design space, size of the population in an

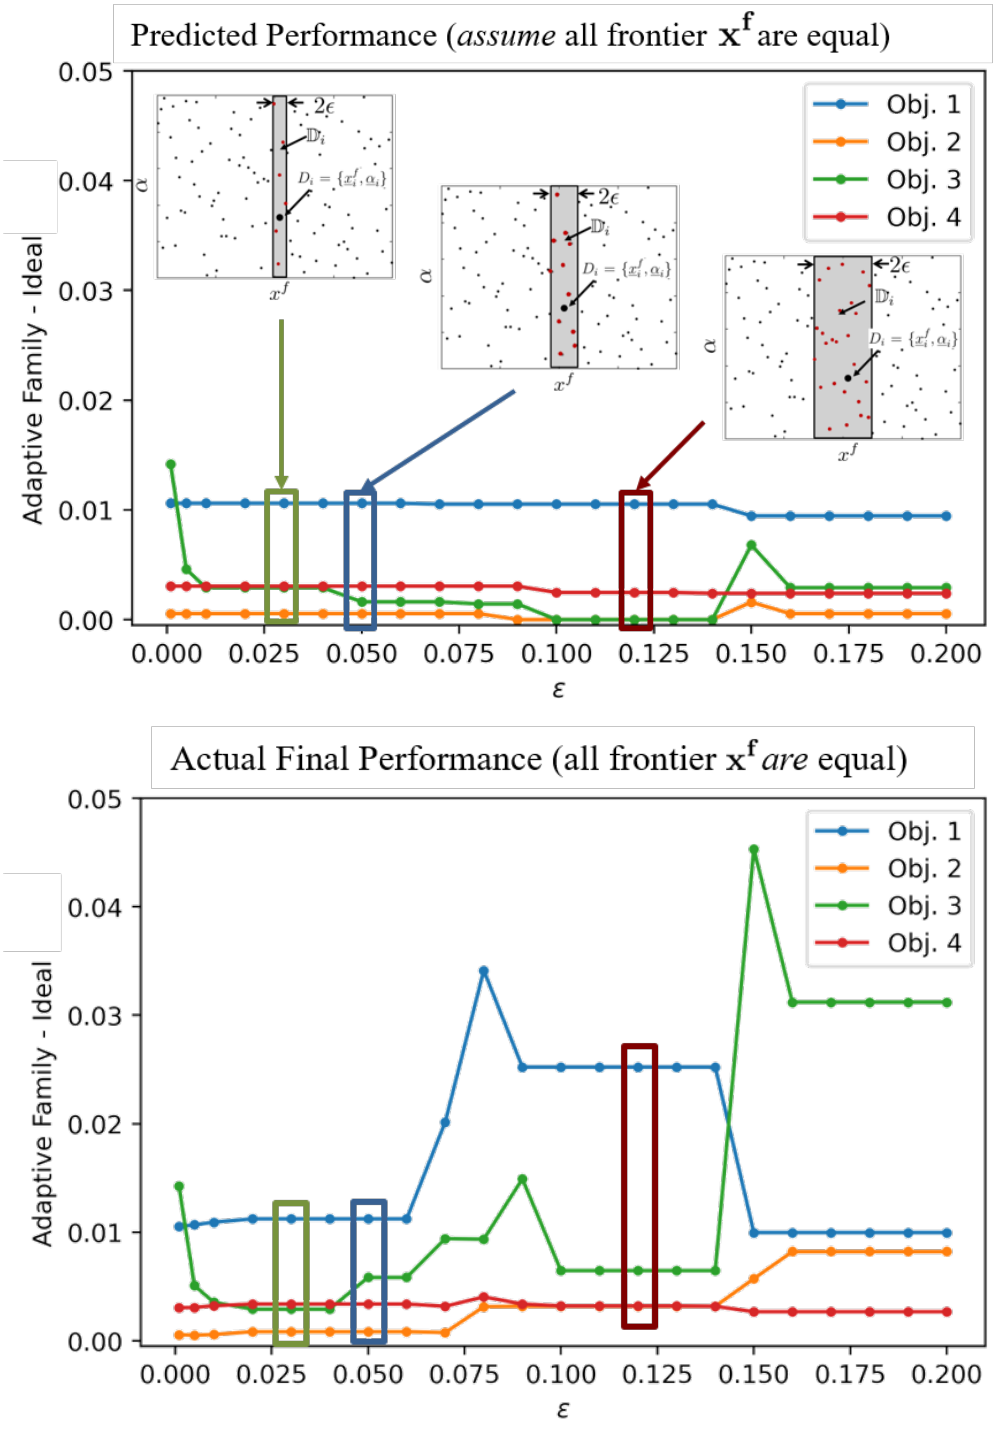

Figure 3.30: While the predicted best adaptive objectives decreased as per the Euclidean error metric with increasing similarity parameter ϵ, the actual final designs performed best in the $0.02 \leq \epsilon \leq 0.04$ range.

optimization, and benefits of the methods presented in this chapter. If an adaptive design problem is limited enough that exploring the design and objective spaces is easily done considering each adaptive design variable for each objective with the given population, then that is likely the most effective approach. However, simply increasing an Euler-Bernoulli beam shape fitting design problem from two to four shapes dramatically decreased the effectiveness of searching over an increased design space, demonstrating advantages and applications for the design space decomposition method when compared to conventional optimization techniques. The design space decomposition approaches also performed relatively better on the same four-shape design problem with smaller population sizes, which increases applications for optimizing designs requiring computationally expensive engineering analysis tools such as mission-driven adaptive rotorcraft designs. Finally, evaluating the best final designs over a range of ϵ values can determine the best fully realizable designs from an existing optimization or other design space exploration for a final adaptive configuration.

4. COMPUTATIONAL FRAMEWORK DEVELOPMENT AND MISSION-DRIVEN ADAPTIVE ROTORCRAFT DESIGN AND OPTIMIZATION

This chapter applies the adaptive rotorcraft computational methods described in Chapter 2, developing a framework to relate adaptive rotorcraft design variables such as chord, camber, and twist to mission-driven performance objectives such as required power for trim flight and fuel burn. Next the design space decomposition methodology developed in Chapter 3 will be applied to the designs generated via genetic optimizations to improve mission performance using the approximate local frontiers method to post-process all designs. In this chapter, a computational framework is developed to relate mission requirements and design variables such that preferred designs (i.e., designs resulting in improved mission performance) are determined algorithmically.

Trade-offs between the design space size, analysis fidelity, and costs such as time and computational resources are considered in the development of this framework with a focus on early stage design and optimization. Thus, lower fidelity tools offering similar results and trends are considered when doing so makes a larger design space feasible with respect to costs. The entire process from defining missions based on altitude, velocity, payload, passengers, etc, to calculating performance metrics from required trim power to total fuel burn and mission range is encapsulated in a modular design and optimization computational framework. As such, different design variables can be altered based on morphing mechanisms in self-contained modules. Aerodynamic properties due to changes in geometries can likewise be updated using range of tools described in section 2.1. First, missions are defined and adaptive geometries are optimized to improve either individual mission stage performance. Next coupled, sequentially evaluated missions are introduced and optimized to improve overall mission-wide performance metrics such as required fuel and final range.

This chapter begins with a high-level description of the computational framework. Next examples of mission stages are described in the context of calculating the trim flight condition. Then multi-objective optimizations are performed to demonstrate the framework capabilities to

determine preferred adaptive geometries for mission stages. Finally, a mission score sequentially-analyzed approach is introduced.

Chapter 4 Outline

This chapter is divided into 3 sections:

- In section 4.1 the high-level computational framework is described. Coupling between inputs, mission definitions, sizing tools, and rotorcraft trim performance analysis tools are detailed the large computational framework. Different mission stage types are considered in this work, with the methods for determining the trim flight and selecting the best forward flight velocities described in this section.
- In section 4.2, low-fidelity aerodynamic tools are utilized to explore the effects of geometric adaptive optimizations based on multi-objective missions consisting of competing stages. Chord, twist, and outer mold line shapes were all optimized to determine the best adaptive designs given specific mission stages. Additionally, the active twist optimization introduces the first design space decomposition application toward mission-driven rotorcraft design and optimization, which will be further explored with higher-fidelity tools in Chapter 5.
- In section 4.3, a sequential analysis method is developed as an alternative to the multi-objective approach from section 4.2. In this optimization approach, mission-wide objectives such as fuel burn provide an overall mission score such as, in this work, fuel burn and range.

4.1 Mission-Driven Design and Optimization Computational Framework

The computational framework described herein considers mission definition, analysis type, design variables, and parameterized data (weight, camber, chord, etc.), all of which will be addressed to some extent in the following. The full computational implementation was developed in

Table 4.1: Mission Definition Parameters

PARAMETER
Mission Stage Types
Stage Lengths (in either time or distance)
Passenger Weight
Cargo Weight
Fuselage Empty Weight
Air Properties

ModelCenter by Phoenix Integration [68], a program to manage simulation modeling work-flow. ModelCenter was selected for both it's trade-study and optimization tools and for its modularity. Additionally, ModelCenter interfaces directly with MATLAB, where TRAC and the BEMT code were developed. Optimizations in this work were run using the built-in ModelCenter suite of optimization tools.

4.1.1 Computational Workflow

The simulation work-flow, shown in Fig. 4.1, engages a set of sequential modules; any systems with iterative feedback (such as TRAC and hover BEMT) are self-contained modules. The names of specific module groups corresponding to Fig. 4.1 are listed in bold in this section. A more detailed description of the computational framework can be found in Appendix C.

First the specified mission is defined by flight conditions, fuel weight, cargo/passenger weight, mission stage types, and mission stage lengths (**Define Mission**). The full list of mission definitions for each stage, shown in Table 4.1, includes air density, summation of fuel, cargo, passengers, and vehicle empty weight. Missions are either defined as coupled (continuous) or uncoupled (separate mission stages). In an uncoupled mission each mission stage is evaluated independently, and mission stages under the same or similar flight conditions can be evaluated with a single trim flight minimization. In a coupled mission, each mission stage is run sequentially, with changes in fuel weight due to consumption included and updated during the analysis. For the coupled method a

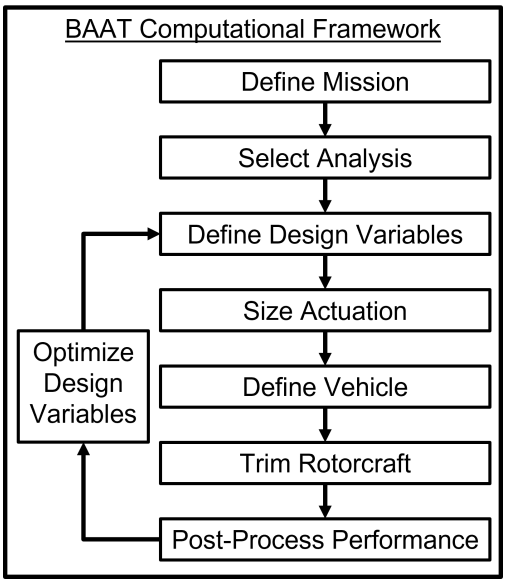

Figure 4.1: The computational framework. Design studies alter the design variables and evaluate the performance metrics after post-processing.

mission is treated as a series of connected stages without considering transitions between mission stages. This is due to the increased complexity and costs associated with transient effects and optimal control. Often surrogate models are developed to consider required power and fuel burn through the course of a mission, but this may be difficult in the early design stages, when many geometric and structural parameters are not yet selected [107], though work has been done to consider some geometric parameters at early stages with some success[139].

Next the analysis parameters are defined, such as whether aerodynamic properties are determined by XFOIL, CFD, precomputed interpolation tables, etc. as described in section 2.2 (**Select Analysis**). These analysis options are defined by a series of analysis flags. Some methods may be incompatible with each other, and checks ensure conflicting analysis parameters are resolved while warning the user. For example, an uncoupled analysis does not consider fuel burn, so the FuelBurn flag will be set to False if the Coupled Flag is set to False. Meanwhile, fuel burn can be considered for a fully coupled analysis, but does not have to be and can otherwise be set here.

Outputs of the previously mentioned modules are input into module defining vehicle design

variables (**Define Design Variables**). This module interfaces with ModelCenter's optimization and design tools such that, during an optimization, changes in objectives result only from changes in the Design Variables module. These parameters can include parametric adaptive design variables such as chord, camber, twist, and any respective locations and schedules down the length of the rotor blade. Modular models can be inserted between this module and the analysis tools to determine the local blade element properties based on parametric design variables.

The computational framework sizes morphing mechanisms then relates geometric inputs to global vehicle and local rotor blade properties (**Size Actuation** and **Define Vehicle**). For instance, if SMA torque tubes twist the rotor blade a specified angle over a specified regions, sizing and structural constraints are evaluated based on the torsional stiffness, internal clearance, and required twist as described in section 2.2.2. Changes in mass distributions due to actuation are determined and added to the existing rotor blade mass. Changes in geometric (such as local twist and shape) and aerodynamic (such as lift, drag, and moments) properties for each rotor blade element due to morphing are likewise determined in individual modules.

Next, rotorcraft aerodynamic tools determine the trim flight condition for each mission stage based on the mission parameters, flight conditions, and vehicle properties (**Trim Rotorcraft**). The pilot inputs for trim flight for each mission stage are determined, and required power (and velocity, if not specified) is calculated. For mission stages such as hover, climb, and cruise the forward velocity is specified, while stages such as loiter and endurance trim the rotorcraft over a range of forward flight speeds and the best forward flight speed selected. If fuel burn is considered for a coupled mission, the fuel is updated between each mission stage based on the previous trim flight solutions. If the final stage considers the final range of the vehicle, a common mission objective, then a coupled mission is evaluated with a fuel weight equal to half the remaining fuel.

Finally, the Post-Processing module relates trim flight outputs to mission performance in the form of fuel burn, power, range, etc., for each mission stage (**Post-Process Performance**). Other computational analysis tools could be added to calculate, for example, rotorcraft noise emissions under trim flight condition if such a tool were supplied to relate trim flight conditions with noise,

allowing for vehicle noise to be another objective. However, the objectives considered in this work are all related to power and fuel burn, not considering noise. The Post-Processing module returns all necessary outputs to ModelCenter for design studies and optimizations. During such trade studies, ModelCenter interfaces with the computational framework by altering the design variables and evaluating performance metrics from the Post-Processing module (**Optimize Design Variables**).

A computational, modular framework was developed to consider a wide range of adaptive rotor blade technologies in a multi-objective design and optimization framework. The framework considers specific rotorcraft missions and can consider objectives that any module produces; in this work the focus will be on trim power, fuel burn, and range. Next, the specific mission stage types are described.

4.1.2 Mission Stage Types

Mission stages consider environmental conditions (such as air density, temperature, altitude, humidity, etc.) and vehicle conditions (such forward flight speed, climb rate, number of passengers, additional weight, etc.) Calculating trim flight and mission objectives varies depending on the specific mission stage type considered. For each mission stage type, first trim flight is calculated, then required power, then (if necessary) fuel burn.

The hover trim flight condition is determined using the BEMT code to solve for the pilot's collective pitch, as outlined in section 2.3.1. For vertical climb, the power required is determined by first calculating the hover power, then applying hover-climb relations described in section 2.3.1. The best loiter (and endurance) condition is determined by calculating the trim flight condition over a range of forward flight velocities, then selecting the velocity yielding minimum required power. Similarly, the best range flight condition is also determined by calculating trim flight over a range of velocities, and the solution with the highest speed-to-power ratio is selected as the velocity for maximum range. Total range can be calculated for a given fuel input as:

$$\text{Range} = \frac{\text{Fuel} * \text{Velocity}}{\text{SFC} * \text{Power}}, \tag{4.1}$$

where SFC is the vehicle specific fuel consumption rate, and it can be seen that increasing the velocity and decreasing power results in increased range.

If fuel burn and remaining fuel are considered during a mission evaluation and the final stage is maximum range or endurance, the final mission stage is evaluated with half the remaining fuel. 'Cruise' in this work typically refers to high speed forward flight, but any specified forward flight speed is determined by specifying the incoming airspeed v_∞. In terms of computation speed, hover calculations are the fastest, while the range and loiter conditions require the most time since the preferred forward flight speed for both mission stage types is selected from a range of velocities.

Other mission stage types were considered but not part of any missions considered in this work. One mission stage type considered was the maximum forward flight speed given a maximum operating power. In this mission stage type, the forward flight speed is varied to determine the velocity where the required trim power equals the maximum allowable power. Another mission stage type with an additional iteration process developed was a maximum loading flight condition, where the maximum vehicle loading was determined iteratively based on the maximum operating power. A trim flight power curve where maximum loading was determined for the UH-60 considering a power limit of 3,000 HP can be seen in Fig. 4.2. For this mission stage type, an external optimization loop is added in which the total vehicle weight (W) is altered until the trim power equals the maximum allowable engine power (P_{max}):

$$J = \min_{W} f(W) \ = |\max(\text{Power(W)}) - P_{max}|. \tag{4.2}$$

In this example, the maximum gross weight before exceeding power constraints was approximately 85,000 lbs. Of course, calculating maximum power increases the number of required evaluations with the addition of another optimization loop around the trim flight solution.

In a mission-driven optimization, these mission stages can be evaluated either independently in a multi-objective optimization or sequentially as part of a single mission-wide score. Both approaches are considered in this work using the computational framework described here. First the we shall consider the multi-objective approach.

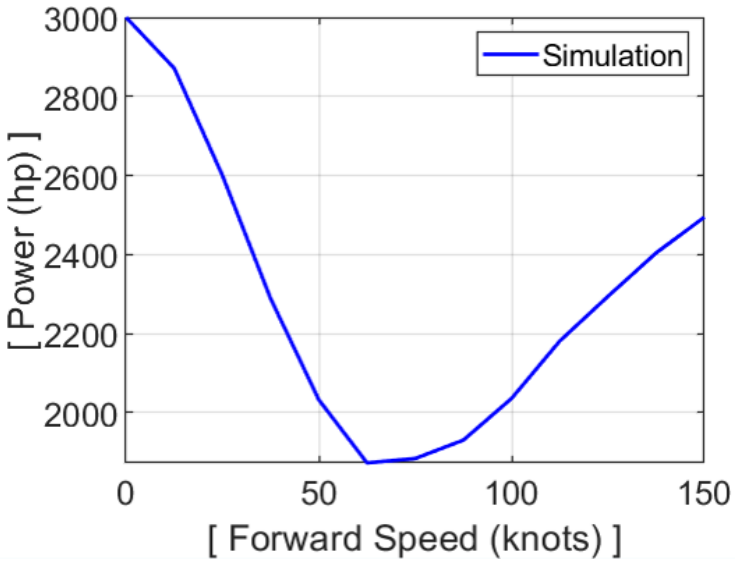

Figure 4.2: The maximum loading is determined iteratively such that the trim power matched the engine constraints of 3,000 HP.

4.2 Preliminary Geometric Design and Optimization Studies

With the computational framework and mission stage types described, geometric design and optimization studies will now exhibit the advantages of morphing, determined by the computational framework and design and optimization methods described previously. Parameterized adaptive design variables are optimized to minimize performance objectives across multiple mission stages to determine the best geometries that can be morphed via the technologies previously described in section 1.3.

4.2.1 Adaptive Rotor Blade Chord Aerodynamic Optimization

Consider an active chord system such as those described in section 1.3.4. Altering the chord between stages of a mission could yield improved performance based on required lift and forward flight speed, with a different preferred chord distribution for each stage. Now consider a mission consisting of four stage types: hover, loiter, high speed cruise, and maximum range. A parameterized optimization problem was defined in which a polynomial chord region is introduced into the

otherwise constant-chord rotor blade, as shown in Fig. 4.3

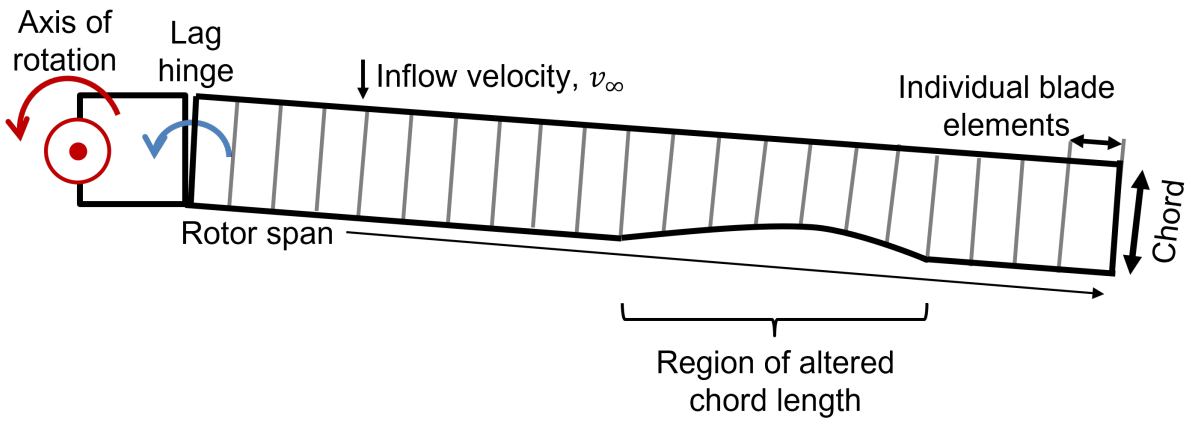

Figure 4.3: The rotor blade is altered by the addition of an adaptive region defined by a 2nd degree polynomial.

Over some region of the rotor blade, the chord is defined as:

$$c_{poly}(r) = c\left[1 + c_1\{r - r_1\} + c_2\{r - r_1\}^2\right], \; for \; r_1 < r < r_2, \tag{4.3}$$

where the design variables and constants are described in Table 4.2.

The rotor blade consists of 30 blade elements with the aerodynamic properties for C_l, C_d, and C_m taken from the SC-1095 airfoil. The changes in lift, drag, and moments are altered in the aerodynamic calculations as a function of chord (c), which determines local lift (and drag, and moment similarly) via:

$$dL = \frac{C_l * \rho v^2 c dr}{2}, \tag{4.4}$$

where v is the local incoming velocity.

A multi-objective optimization was run over the four objectives as given in Table 4.3 using the genetic algorithm NSGA-II to determine the best chord schedules across the objective space. The best trade-offs between objectives were determined, with a 'best' design assigned for each

Table 4.2: Adaptive Chord Variables and Parameters

Design Variable	Symbol	Description
Chord coefficients	c_1, c_2	Polynomial coefficients
Region start	r_1	Adaptive chord region inboard location
Region end	r_2	Adaptive chord region outboard location
Other Parameters		
Rotor span	r	Rotor span-wise location
Polynomial chord	$c_{poly}(r)$	Chord as a function of r, including polynomial Region
Chord	c	Rotor blade root chord

Table 4.3: Adaptive Chord Objectives

Objective	**Description**
Loiter	Minimum power required for trim flight at any velocity
Range	Maximum velocity/power ratio (proportional to maximum range)
Cruise	Minimum power required at 150 knots forward flight speed
Hover	Minimum power required at hover

of 4 objectives: i) minimum power at some loiter speed to be determined, 2) minimum power-to-velocity ratio, which correlates to maximum range as seen in equation 4.1, 3) minimum power at a 150 knots forward flight speed, and iv) minimum power at hover.

The optimization searched the design space ($\underline{x}$) from Table 4.2 to determine the best trade-offs between the objectives in Table 4.3:

$$J = \min_{\underline{x}} \underline{f}(\underline{x}). \tag{4.5}$$

The objective space has 4 dimensions, with can be difficult to completely visualize in 2D figures, so two 'slices' are shown displaying trade-offs between competing objectives in Fig. 4.4. The designs are all on the Pareto Frontier, which means the designs shown in Fig. 4.4 that do not appear Pareto-optimal for the objectives shown are Pareto-optimal across other dimensions, such as a design that

lies on the non-dominated frontier between Hover and Cruise power (or Hover, Cruise, and Loiter, for example).

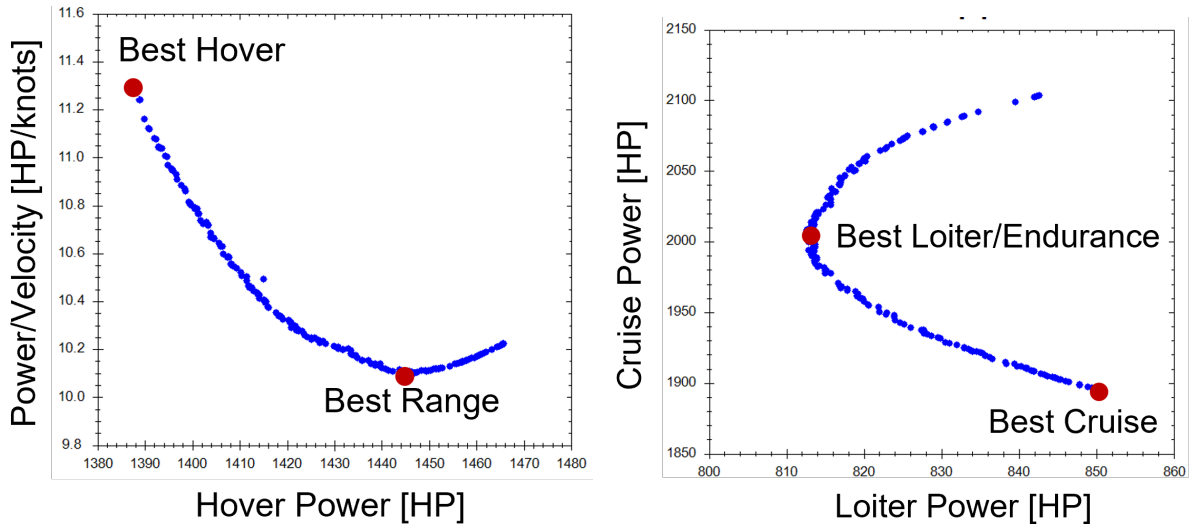

Figure 4.4: In a 4D optimization with competing objectives there exist a set of designs that represent the best trade-offs between each objective.

Each design yielded a different chord distribution, but with similar preferred geometric properties, generating designs similar to a trailing-edge taper down the outboard regions of the rotor blade span. The best design for each flight condition was a outer span taper; the best designs for mission phases with higher velocities in this study were less-tapered rotor blades. This result is because as the forward velocity is increased, the rotor disk must generate more forward thrust, which requires tilting the rotor disk forward, reducing the vertical thrust component. Thus, more thrust is required to both maintain trim flight vertically and forward flight, which is diminished by the fuselage drag. Trimming the rotorcraft requires matching the horizontal and vertical forces, and the parasitic power increases exponentially with airspeed (see 2.8 from section 2.3.2). Increasing the rotor blade planform via chord morphing increases the total lift, improving performance by lowering the required angle of attack for trim flight.

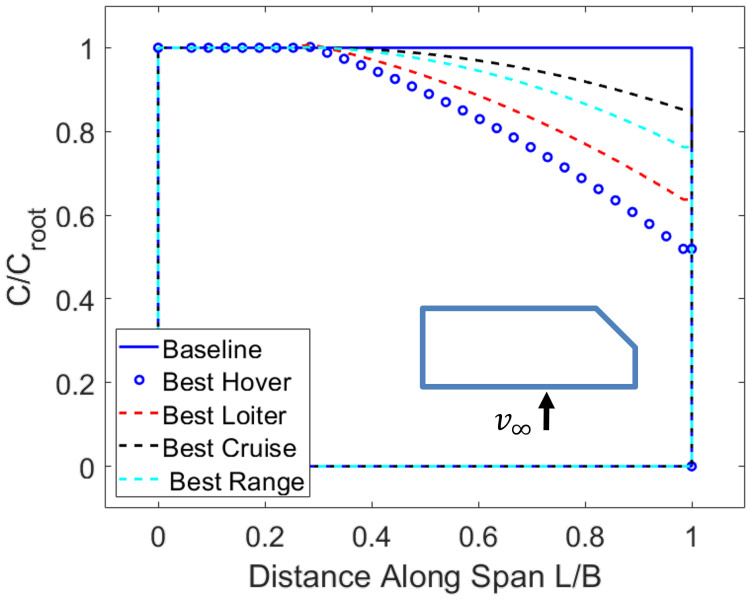

Figure 4.5: The best adaptive chord geometries were varying levels of taper down the outer rotor blade span. The best geometries have less tapered ends as forward velocity increases.

Of course, feasibility of adaptive rotor blades requires consideration of morphing mechanisms and the fidelity of aerodynamic data. Adaptive rotor blade chord methods have been explored and demonstrated in section 1.3.4, but a more detailed exploration of the morphing feasibility would be necessary beyond this purely aerodynamic optimization.

Additionally, in this optimization no distinction was made between fixed and adaptive design variables as defined in Chapter 3. Since the start and endpoints for the morphing region was almost identical for each of the preferred rotor blade geometries, these designs likely are all feasible via changing the level of chord contraction over the same rotor blade region. However, this type of solution (the best designs across all objectives having almost identical morphing regions) is unlikely to always be the case.

4.2.2 Effects of Variable Twist and Mission Stage on Power

Several methods of active twist have been explored and developed in the literature, as discussed in section 1.3.2. The computational framework described in this chapter was next applied toward

the design of an optimal adaptive twist rotor blade for a complex, multi-stage rotorcraft mission.

The design and optimization problem was defined as a series of embedded internal SMA torque tubes attached end-to-end inside the rotor blade. The stresses, strains, and torque tube sizing were calculated using the relations between material and geometric relations described in section 2.2.2.

The rotor blade was given a constant torsional stiffness down the rotor blade informed by the properties used in the RCAS implementation of the UH-60 [19]. The allowable internal clearance inside the rotor blade was also held constant down the span of the rotor blade. The maximum twist rate down the span of the rotor blade is a function of rotor blade torsional stiffness, SMA material properties, and internal clearance, which are all constant down the length of the rotor blade. Thus, the maximum SMA torque tube twist rate per unit length can be calculated before running any optimization or trade study, and was used to set the design space bounds such that all designs satisfied internal stress, strain, and clearance requirements.

4.2.3 SMA Torque Tube Twist Optimization

An optimization design problem was created where up to five SMA torque tubes could be installed in a rotor blade to alter the twist schedule. Changes in twist due to actuation are discretized to match the twist at the center of each local rotor blade element as shown in Fig. 4.6. Any torque tube length extending past the rotor blade tip is removed. The optimization design variables and objectives can be seen in Tables 4.4 and 4.5. The torque tube lengths begin at the end of the rotor blade cutout. The rotor blade is morphed via NiTiHf SMA torque tubes arranged within the rotor. The local blade segments maintain their sectional profiles while the twist distribution is altered by integrated SMA actuators. In subsequent analysis, these SMA actuators will be sized according to the torsional strength of the rotor blade, length of twist, and angle of twist. A multi-objective optimization considering both fixed and morphable design variables is formulated to minimize $\underline{f}(\underline{L}, \underline{\hat{\theta}})$, where the lengths of the SMA torque tube actuators $\underline{L}$ cannot change after the manufacturing stage, while the torque tube twist rates $\hat{\theta} = \frac{\theta}{R}$ can change between mission stages, altering the rotor twist schedule.

The optimization searched the design space $(\underline{L}, \underline{\hat{\theta}})$ from Table 4.4 to determine the best trade-

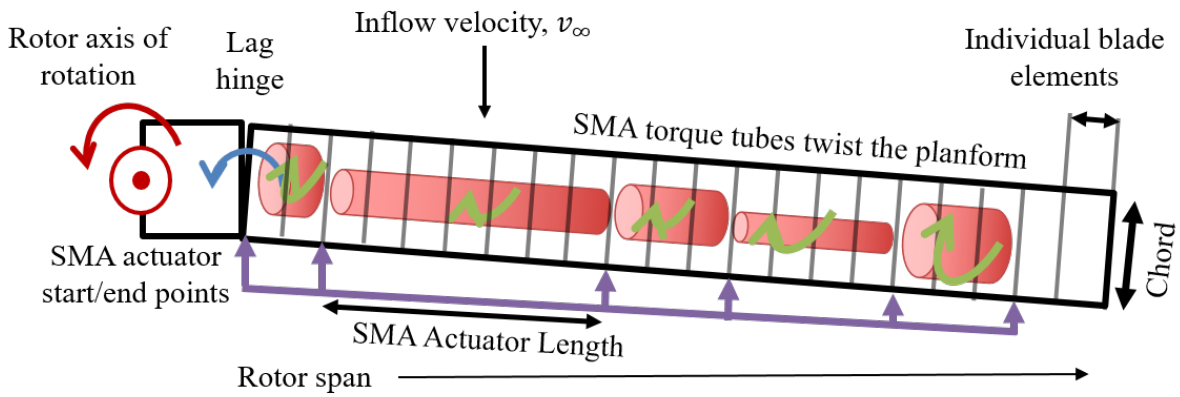

Figure 4.6: Up to five SMA actuators with varying twist rates are installed into a rotor blade to adapt the twist schedule based on the mission stage.

Table 4.4: Torque Tube Optimization Design Variables

Design Variable	Description	#	Variable Bounds
L_i	Normalized i^{th} SMA actuator length	5	$0 \leq L_i \leq 0.4$
$\hat{\theta}_i$	SMA actuator twist rate	5	$-2.0 \leq \hat{\theta}_i \leq 2.0 \ \left[\frac{\text{deg}}{\text{m}}\right]$

Table 4.5: Torque Tube Optimization Objectives

Objective	Description
P_{hover}	Power at Hover
P_{loiter}	Best Loiter Power
$\left[\frac{P}{V}\right]$	Best Range (minimum Power/Velocity)
$P_{v=140\ knots}$	Power required at Cruise

offs between the objectives ($\underline{f}(\underline{L}, \underline{\hat{\theta}})$) in Table 4.5:

$$J = \min_{\underline{L},\underline{\hat{\theta}}} \underline{f}\left(\underline{L}, \underline{\hat{\theta}}\right). \tag{4.6}$$

First, an optimization considering both the fixed and adaptive design variables is performed as part of the approximate frontier method described in section 3.2.2. In this initial study only the length and twist provided by each actuator are considered; all other blade parameters are held constant. Up to five SMA actuators are installed sequentially along the rotor blade space outboard of the lag hinge with no gaps between them. The length of each actuator is bounded from 0 to 40% of the rotor span and any actuator lengths extending beyond the rotor blade tip are removed.

The analysis determines trim flight performance from 0 to 140 knots in 20 knot increments to create the flight profile necessary to determine the best loiter and range velocities. To ensure simulation feasibility and validity, a constraint is implemented to ensure each evaluation converges to a trim flight condition. Unconverged solutions are removed from each flight profile whilst retaining converged solutions; a mission with a converged hover solution and unconverged range solution will still consider the hover objective for a multi-objective optimization. In this example a poor (or difficult to trim) range solution under a specified configuration does not preclude the hover solution because the pilot can morph to another configuration during the range mission stage. The minimum SMA actuator size is governed by SMA material properties, actuator length, actuator twist, and rotor blade torsional stiffness. The maximum actuator size was constrained such that the diameter of the actuator did not exceed the maximum internal clearance based on the rotor blade outer mold line.

The multi-objective optimization is performed using the NSGA-II evolutionary algorithm considering a population of 100 individuals over 25 generations to generate a non-dominated frontier of designs. Next, the design space is explored by taking these optimal designs and searching for designs of similar SMA actuator lengths. Actuator sizes cannot be modified during a mission, but the angles of twist may be altered mid-flight based on operator inputs and the range of feasible

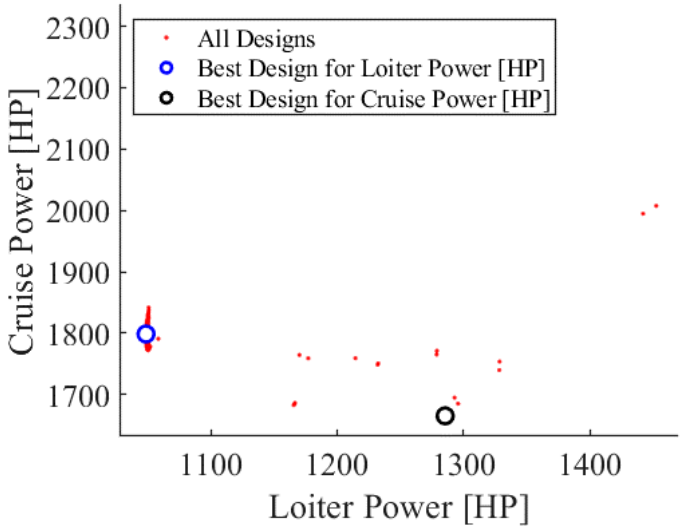

Figure 4.7: A clear trade-off between loiter and cruise can be seen; no single design has near-optimal performance for both objectives.

actuation.

The first optimization served to identify sets of similar actuator lengths which morph to optimal configurations for multiple objectives. The objective space, along with the best designs for each mission stage type, can be seen in Fig. 4.7. The best designs for each objective result in suboptimal performance for other, competing, objectives, suggesting advantages for adaptivity. These actuator lengths L_i are then fixed, and another optimization is run to optimize the angles of twist $\left[\hat{\theta}\right]_i$ given these fixed lengths. For each individual in the first optimization, all designs with actuator lengths within 10% of the design variable range were considered similar designs. Then an intuitive search was performed to find individual 'families' of similar designs that exhibited notable improvements with respect to different objectives. Such families of designs demonstrate the ability to fix the actuator lengths and morph (i.e., change the SMA torque tube twist rate) between different configurations depending on the current mission objective.

A final set of fixed design variables was selected based on both mission objectives satisfied in its base configuration and on improvements to mission objectives available via morphing. The

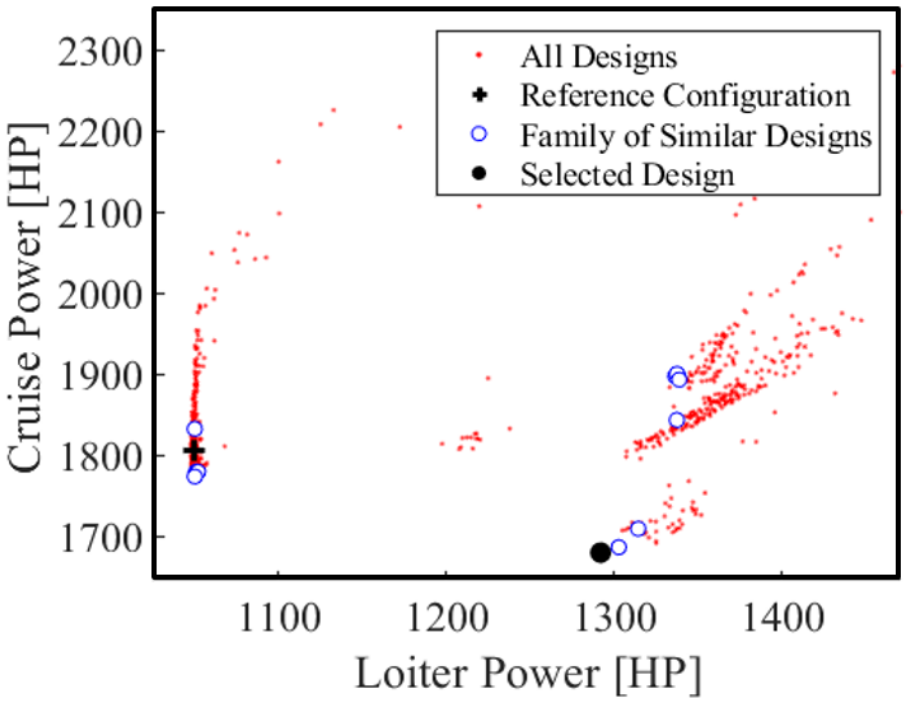

Figure 4.8: The best family of similar designs enable performance increases by morphing from higher-performing cruise designs to better loiter designs.

utopia point distance metric was used to selected the best 'family' of similar designs. The design chosen was the best design with respect to the cruise objective and featured a number of morphable configurations requiring lower loiter power but higher cruise power. The final population, selected design, and family of nearby designs for the first optimization run can be seen in Fig. 4.8. In the figure, the black circle represents the selected design that generates the family of similar designs. All the blue circles represent designs in which each fixed design variable (in this case, each SMA torque tube length) is within 5% of the selected design. To note, the cluster of designs within 5% (for a total band of 10% except where design variables are near the bounds) of a given design for every fixed design variable inhabits 0.001% of the total fixed design space.

With the best fixed design variables L_i selected, a final optimization was performed to determine the best operational design variables for each objective. The selected design variables for the final optimization are listed in Table 4.6.

Table 4.6: Torque Tube Optimization Selected Design

Variable name	Variable Type	Value	Units
L_1	Fixed	0.20	Non-dimensional
L_2	Fixed	0.37	...
L_3	Fixed	0.34	...
L_4	Fixed	0.28	...
L_5	Fixed	0.35	...
$\hat{\theta}_1$	Adaptive	0.75	$\frac{\text{deg}}{\text{m}}$
$\hat{\theta}_2$	Adaptive	-0.13	...
$\hat{\theta}_3$	Adaptive	0.65	...
$\hat{\theta}_4$	Adaptive	1.77	...
$\hat{\theta}_5$	Adaptive	0.75	...

With the actuator lengths fixed to those of the selected design from Table 4.6, the actuator twist rates were then optimized across the four objectives. This optimization was run using NSGA-II with a population of 100 individuals over 15 generations. The best designs for each objective were then selected from this morphable family of designs spanning the entire optimization. Thus, the multi-objective optimization first solved for the preliminary fixed design variables such as actuator lengths then determined the actuator twist rates.

The sum of the relative lengths of the first 4 actuators ($L_1 - L_4$), equals approximately 1.19 times the rotor blade length. Under the present problem formulation any actuator lengths beyond the rotor blade are ignored, as the relative length of the rotor blade is 1. The 5th actuator was omitted from the next optimization since the effects of each actuator are only considered between the lag hinge and the rotor blade; therefore, the 2nd optimization only considered the four inboard SMA actuators.

The final optimization generated a population of adaptive designs using the fixed actuator lengths from the first optimization determined via the Utopia point method. Unlike the objective space from the first optimization, all designs in this population can be achieved simply by

Table 4.7: Torque Tube Optimization Objectives

Objective	Units	Optimization 1 Selected Design	Optimization 2 Morphable Design	% Improvement
P_{hover}	[HP]	2103	2076	1.3
P_{loiter}	[HP]	1292	1049	18.8
$\left[\frac{V}{P}\right]$	$\frac{\text{Knots}}{\text{HP}}$	12.0	11.8	2.0
$P_{v=140\ knots}$	[HP]	1680	1665	0.9

morphing the SMA torque tubes to alter the twist rate between acceptable ranges. Thus, significant performance improvements can be realized throughout a single mission by morphing to the best design for each stage of the mission. The best performance for each objective was realizable via altering the SMA torque tube twist rates, and the final objectives can be seen in Table 4.7. Although the addition of SMA actuators and morphing during the mission offers only marginal improvements to hover, endurance, and cruise, these torque tubes enable a decrease in power required for loiter by over 18% via changes in twist rates.

The addition of torque tubes into a rotor blade for mission-driven performance was determined to improve performance across a mission comprised of four unique mission stages. The lengths of five torque tubes were determined first by an optimization in which every length was considered, then the design space was searched for similar designs with respect to fixed variables but also could span the objective space. Finally, a second optimization was performed over the adaptive design variables to determine i) if the assumption of similarity defined by ϵ was accurate, and ii) the final configuration and performance metrics. The second optimization generated a fully-realizable design with adaptive advantages that improved every objective, with the most noteworthy power performance improvement found by eliminating most of the trade-offs between Cruise and Loiter power.

4.2.4 Parameterized Adaptive Outer Mold Lines

Morphing the rotor blade outer mold line (OML) alters lift, drag and moment profiles. When changes in the rotor blade shape are parameterized, the OML can be optimized over certain regions of the rotor blade. Consider a rotor blade region in which internal actuation could morph the OML between mission stages. The level of morphing may change between missions, but the rotor blade location to install the actuators cannot changes after fabrication. Therefore, the regions on the rotor blade where morphing occurs must be fixed during manufacturing but the OML can still morph in those regions as long as the actuation system can morph between the preferred shapes. The location and OML shapes were optimized using NSGA-II to determine an optimal set of morphing shapes for the three mission stages. The base OML was defined using a four-parameter Class/Shape Transformation (CST) equation. Next, the first three coefficients were altered, with the final coefficient solved for analytically to ensure a closed contour.

An initial optimization was performed without any structural constraints. The results and procedure detailed are included here i) as a caution of purely aerodynamic optimizations that do not consider internal structure and ii) as a justification for structural constraints applied to maintain feasibility of airfoil morphing optimizations and final configurations in this work.

An optimization was run over the CST parameters to minimize four objectives: power at hover, power-to-velocity ratio for maximum range, loiter power, and power at 150 knots. The initial results yielded much better objectives than the reference design, as shown in Fig. 4.9 for the power-to velocity ratio for maximum range and hover power. The relative heights were calculated at the maximum vertical distance between the lower and upper airfoil shape, with the thinner airfoil shape as a reference. The optimization encouraged designs with thinner airfoils that may not be structurally feasible. Thus, structural constraints were added to ensure the leading edge spar shape was maintained. Spatial and rotation constraints were defined between the upper and lower surfaces to approximate internal spars and structures. These constraints were modeled as spars, restricting changes in the space between locations on the upper and lower surfaces as well as relative rotations. More details on these structural consistency equations can be found in works

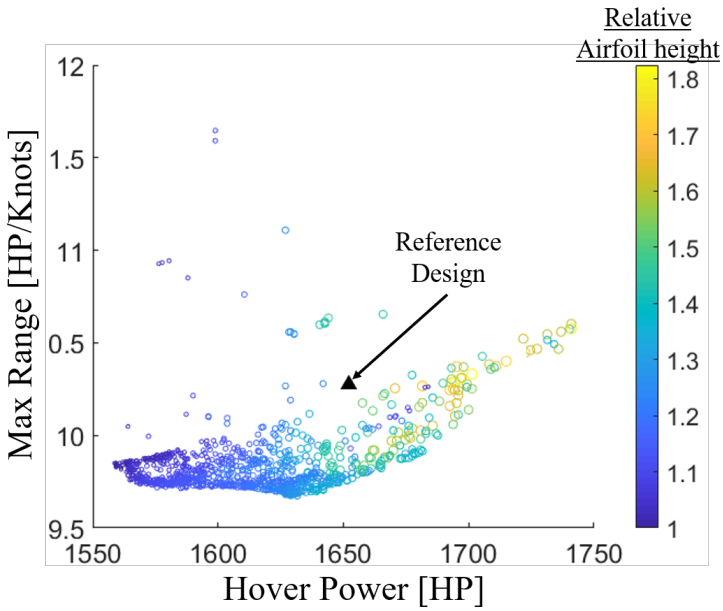

Figure 4.9: The shape optimization improved objectives by generating much thinner airfoil shapes that may not be structurally feasible.

by Leal and Hartl in [96, 95]. Constraints were applied to the Airfoil by placing spars at the 10, 20, and 30% locations down the airfoil chord to restrict the leading edge morphing that typically provides most structural strength in the rotor blade. An example airfoil can be seen in Fig. 4.10, as the leading edge spars restrict movement and rotations between the upper and lower surfaces near the rotor blade leading edge. Next, the airfoil shape was optimized for the unique mission stage types, resulting in optimal shapes for hover, loiter, range, and cruise. The objectives across the 4-D space is shown in Fig. 4.11, where there is a clear trade-off between good Cruise and Range designs. This makes some sense when considering that Cruise requires the highest power for trim flight, while Loiter and Range require the lowest power.

The best shapes for cruise and loiter can be seen in Fig. 4.12. The leading edge for both shapes is almost exactly the same due to the structural constraints, but the best trailing edge OML varies between objectives. Higher trailing edge camber in general increases lift and drag for a given angle of attack and can decrease the maximum angle of attack before stall ($C_{l_{max}}$). As higher forward

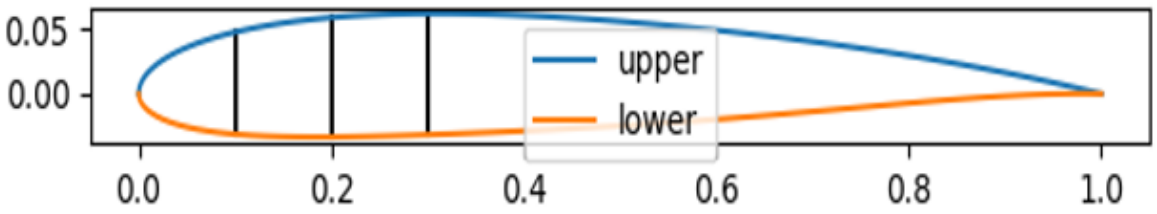

Figure 4.10: Spars were added to the OML definitions to maintain space near the trailing edge for internal structures.

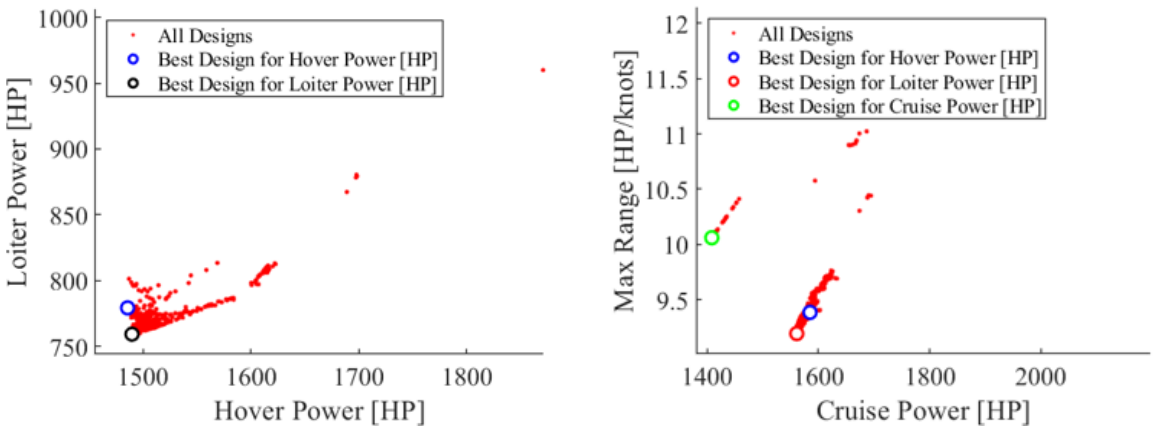

Figure 4.11: Trade-offs can be seen between good Cruise and Range designs.

flight speeds more lift is required, leading to higher collective pitches. Thus, less cambered airfoils can reduce the power requirements at higher velocities by keeping the local angle of attack down the rotor blade span below the stall angle.

4.3 Sequential Mission Analysis

Missions can be either analyzed in a multi-objective optimization over the unique mission stage types, or as a sequence of mission stages. While the multi-objective approach requires less evaluations and can be more easily parallelized, coupling between mission stages (due to fuel burn for combustion engines) is only considered when the mission is evaluated sequentially. Changes in weight due to fuel burn can be considered as the mission progresses for specific missions, which is

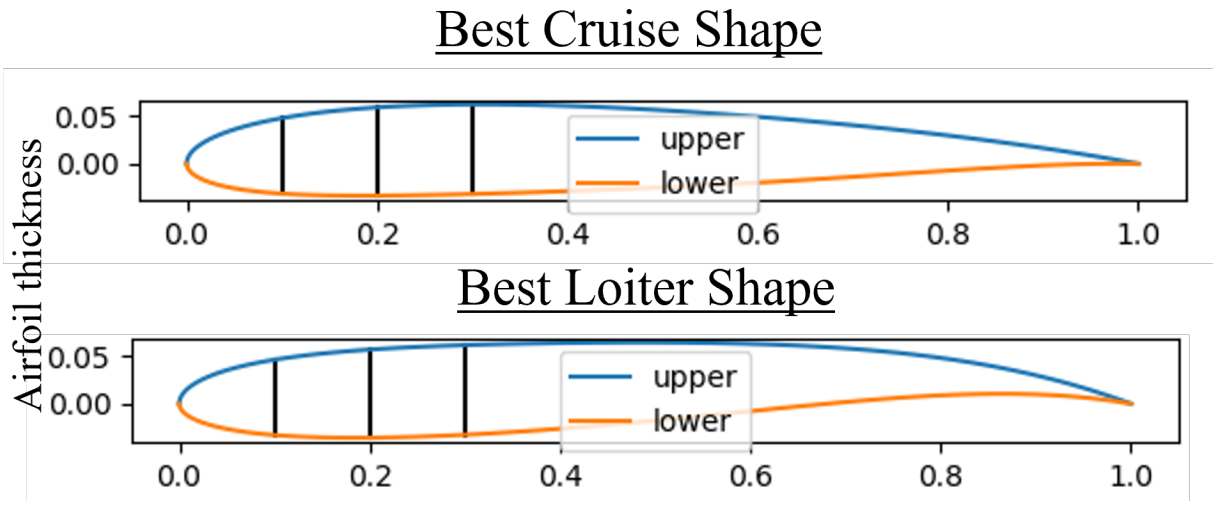

Figure 4.12: The best OMLs for Cruise and Range mission stage types.

not an insignificant portion of the fuel weight. For reference, at maximum take-off weight the full fuel tank accounts for approximately 10% of the total weight for the UH-60. However, this may limit adaptive structures developed in this manner to a single mission as opposed to more general objectives.

A comparison of the uncoupled vs sequential mission analysis can be seen in Fig. 4.13. In the top flowchart, each mission stage is calculated independently given a set of flight requirements. When considering a sequential analysis, the coupling between mission stages due to fuel burn must be calculated and updated after each mission stage. The fuel burned for each mission stage is given by the equation:

$$Fuel = Power * SFC * Time, \tag{4.7}$$

where SFC is the specific fuel consumption for the engines.

Consider a two-stage mission consisting of a loiter then a cruise stage to demonstrate the advantages of mission-driven morphing. First, the rotorcraft must loiter for 30 minutes, then cruise at 140 knots until the vehicle fuel is exhausted. The objective for this mission is to maximize the vehicle range given an initial fuel input of 1,000 lbs.

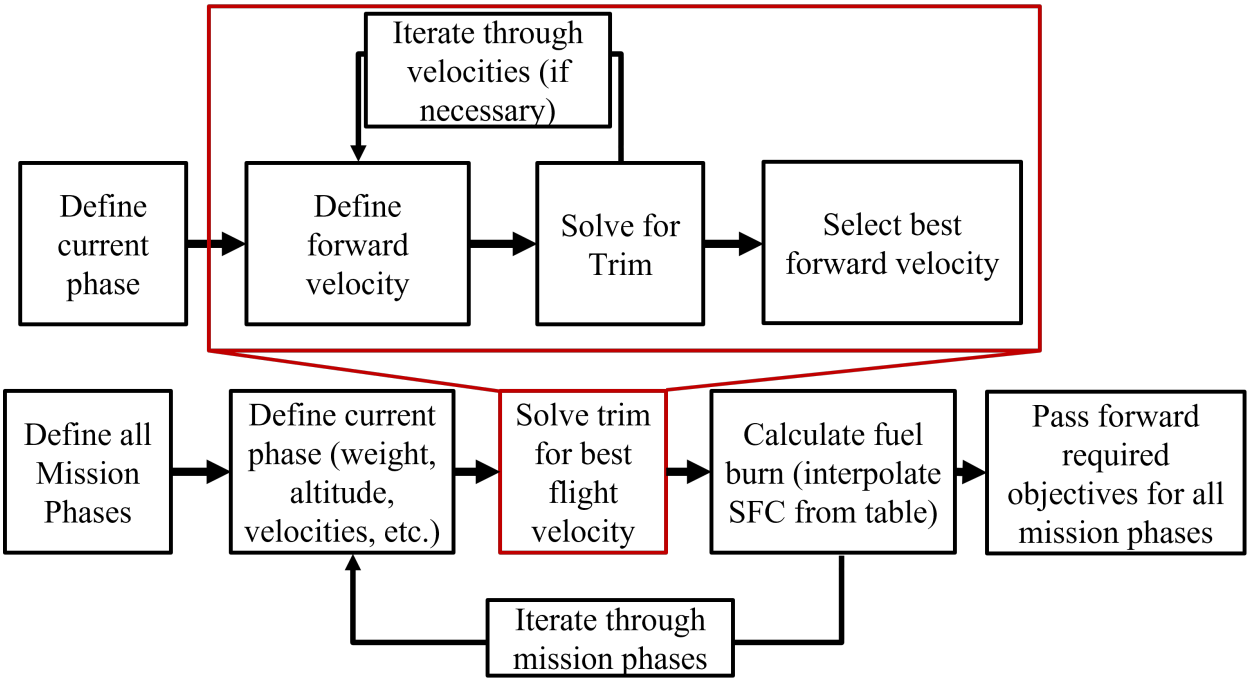

Figure 4.13: Mission stage objectives can either be determined independently (top) or as part of a sequential analysis (bottom).

This two-stage mission was run first using the best loiter and cruise airfoil shapes from Fig. 4.12 both mission stages, then using the best shape for each respective mission stage, morphing from the best loiter shape to the best cruise shape. The results with respect to range and fuel burn can be seen in Fig. 4.14. A comparison of the fuel burn for a given loiter and cruise time can be seen alongside the advantages of morphing the rotor blade between mission stages. Both maximum range and minimum fuel burn can define an overall mission score.

4.3.1 Mission Fuel Selection

The initial fuel for a mission can notably effect the required power and range for a given mission. Increasing the initial fuel increases the required lift and power, which in turn increase fuel consumption further, a coupling that can be seen in Fig. 4.15. However, enough fuel must be supplied to complete the mission. Often, as will be seen later in this work, there is a required fuel reserve at the end of a given mission. A mission in which the required fuel burn is larger than the

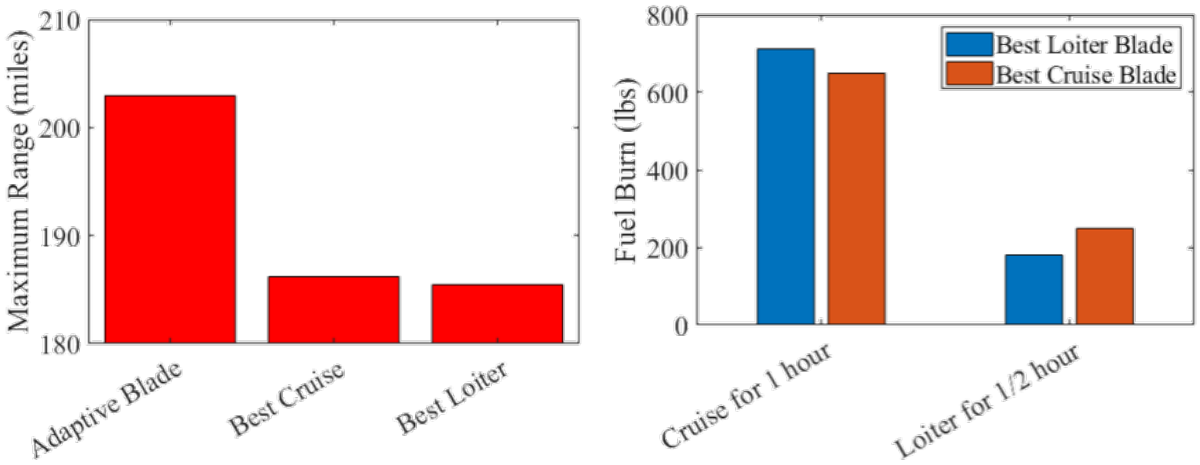

Figure 4.14: Morphing the rotor blade between mission stages increased the range for a two-stage mission by reducing the power, and thus fuel burn. Both of these can define overall mission-wide scores.

initial fuel input cannot be completed.

In this work, the minimum initial fuel for a specific mission is selected based on the fuel required to meet mission requirements for the unmorphed geometric configuration. The fuel required to complete a mission (considering any reserve fuel requirements) is determined for the reference design, and is thus the minimum fuel for a given optimization of the same mission. Less efficient designs than the reference blade can typically be discarded via constraints on reserve fuel, serving no added value for the specific mission.

4.3.2 Adaptive MEDEVAC Mission

Specific missions can be defined as a series of individual mission stages as described in section 4.1.2. Consider a medical evacuation (MEDEVAC) mission in which a rotorcraft and crew must go pick up a group on individuals and quickly transport them to a medical facility. This mission is fairly common in requirements for military capabilities and is considered one of the missions required for all four future vertical lift (VFL) mission capability sets for medium and heavy rotorcraft [66]. An optimization was performed to determine both the optimal positions and shapes for OML morphing during a multi-stage adaptive mission. In this MEDEVAC mission a

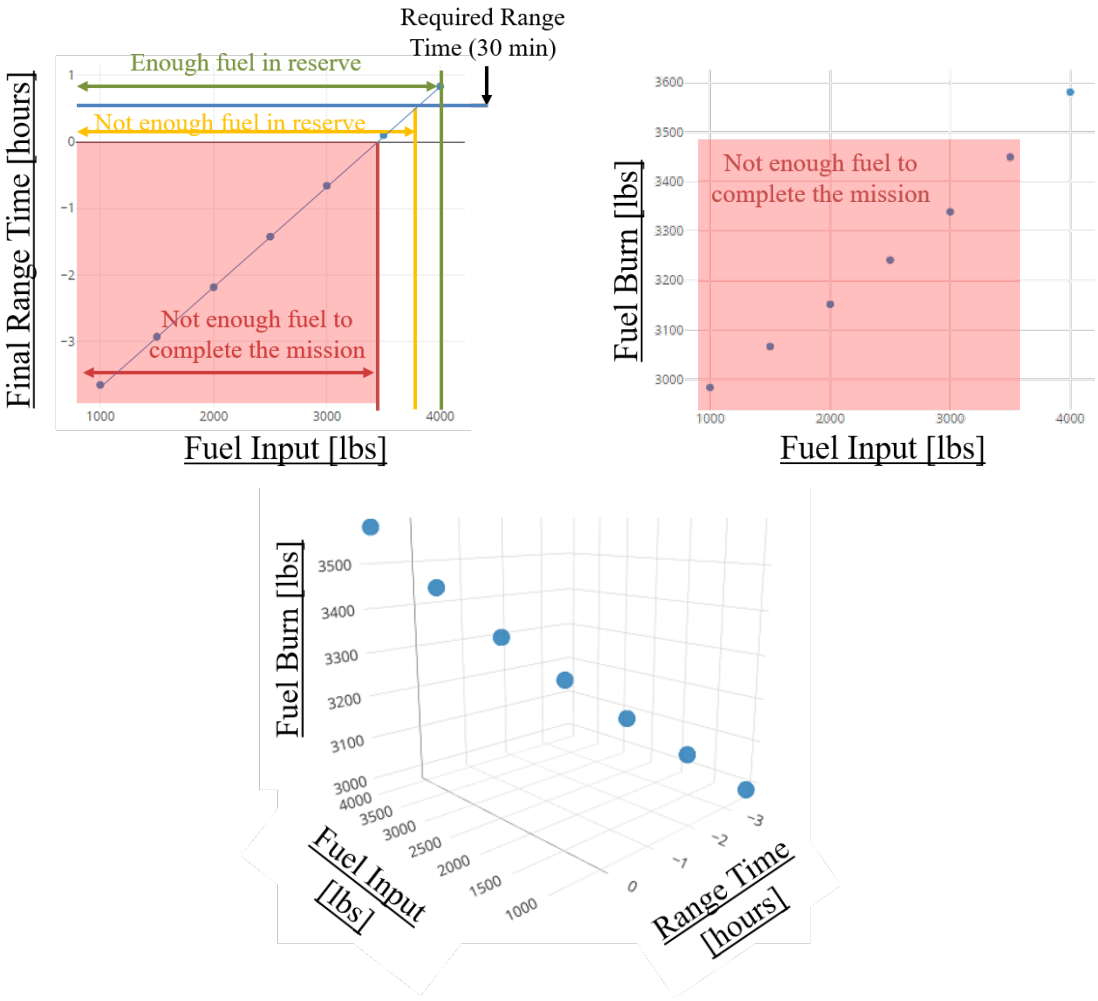

Figure 4.15: Fuel consumption for an example 14-stage mission. Increasing the fuel input increases total fuel burn, while too little initial fuel results in an incomplete mission.

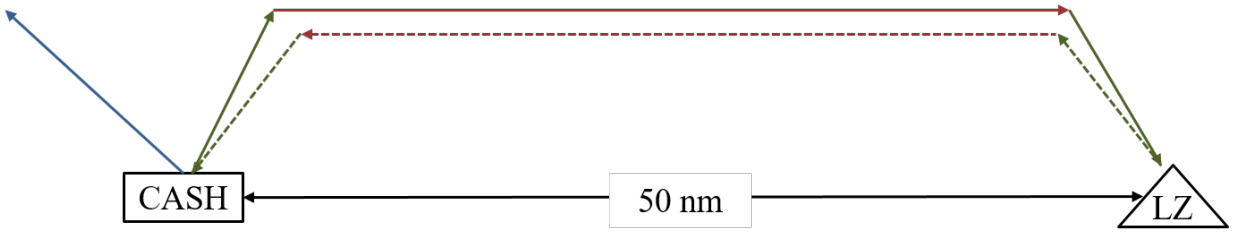

Figure 4.16: The MEDEVAC mission is 5-stage mission consisting of three different mission stage types.

Table 4.8: Five Stage MEDEVAC Mission Definition

Mission Stage Description	Flight Condition	Time/Distance
1) Hover over TAA at 6,000 ft. altitude on 95 deg. day	Hover	0.5 min
2) Cruise at high continuous speed (180 knots)	Cruise	50 nm
3) Descend to add passengers (6 passengers 365 lbs)	Hover	1 min
4) Cruise at high continuous speed (180 knots) to CASH	Cruise	50 nm
5) Determine max range based on fuel reserve	Max Range	Until out of fuel

rotorcraft starts in hover, quickly responds to an emergency 50 nm away, hovers, adds passengers, then returns at high speed to a combat army support hospital (CASH). Finally, the vehicle flies at best range speed. This final stage determines the range of missions the vehicle can still perform after the mission but before refueling. The MEDEVAC mission can be seen in Table 4.8 and Fig. 4.16.

Two constraints were applied to the mission, as shown in Table 4.9, the same as the air assault mission constraints . The maximum power during any mission stage cannot exceed a maximum power threshold based on the two engines used on the UH-60 Blackhawk. The other constraint is the common mission requirement of half an hour of fuel reserve at maximum range.

This MEDEVAC mission can be evaluated either as a single objective design and optimization problem with sequential set of mission stages and a final objective at the end of the mission or as a multi-objective design and optimization problem with separate objectives for each mission stage

Table 4.9: MEDEVAC Mission Constraints

Description	Value	Units
Final Range Time	30	Minutes
Maximum Power (for any mission stage)	3150	Horsepower

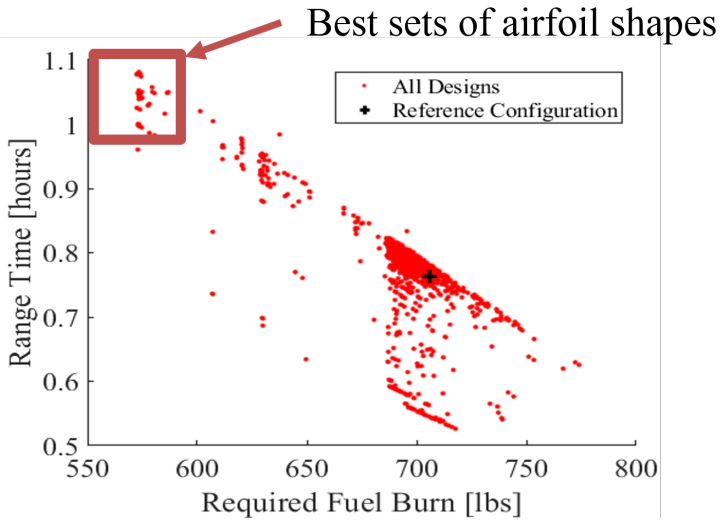

Figure 4.17: The sequential mission optimization evaluated designs based on overall mission performance metrics: fuel burn and final operational range.

type.

The OML is optimized to maximize the final mission range and reduce fuel burn over the course of the mission. The best design should minimize the fuel burn through the first 5 mission stages and maximize the final range time. The objective space can be seen in Fig. 4.17, while the best design (in the top left corner of the objective space) consists of three different airfoil shapes.

The optimization selected preferred OMLs for hover, range, and cruise as shown in Fig. 4.18. The results were similar to those from the multi-objective optimization performed in section 4.2.4, where lower required power led to higher preferred camber OMLs.

The MEDEVAC mission was optimized using the Mission Score approach, in which mission-

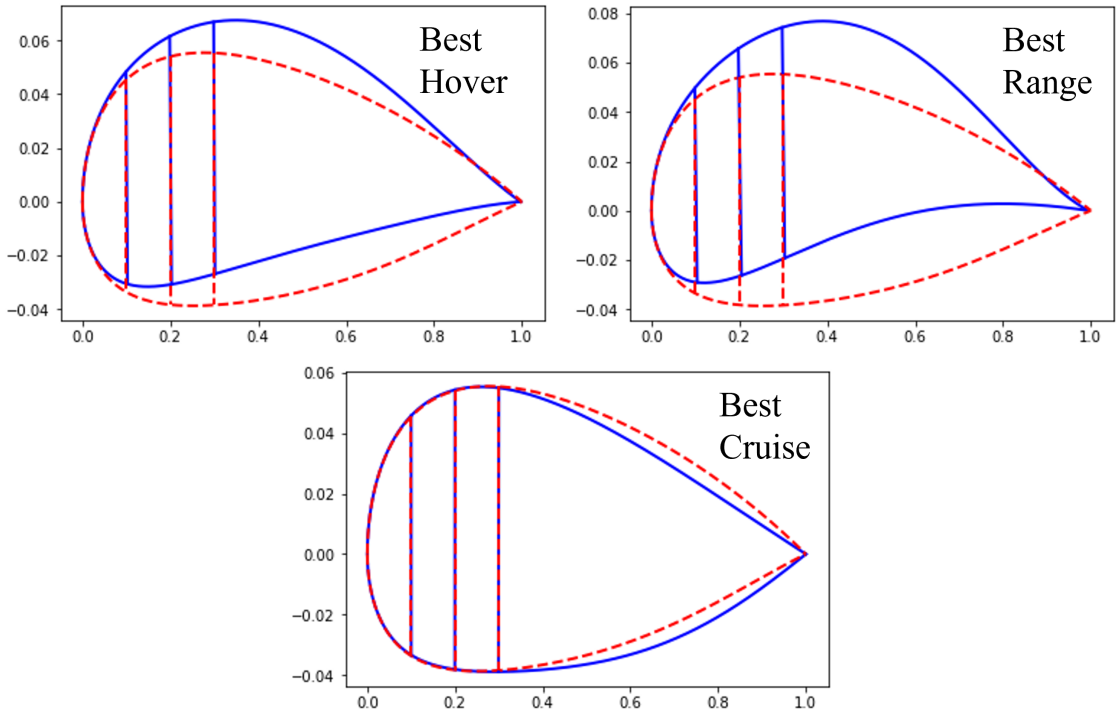

Figure 4.18: The best OMLs for each mission stage type varied largely in the change of trailing edge camber. Changes in camber are highlighted by the differences in axis scaling.

wide objectives are optimized and the adaptive design variables for each mission stage are considered separately. This optimization was performed for shapes associated with three mission stage types, resulting in 3 adaptive aerodynamic shapes to improve mission-wide performance. The mission-score approach provides an alternative to the multi-objective approach with a smaller objective space, but missions must be evaluated sequentially.

4.4 Chapter 4 Summary

In this chapter a computational framework is developed building on the adaptive rotorcraft technologies and methods described and developed in Chapter 2. Mission stage types were defined in the context of mission-driven optimization, where the complete mission can either be optimized as a set of independent stages in a multi-objective optimization or a sequential set of stages with an overall mission score.

Traditional adaptive rotorcraft morphing methods were optimized for multiple mission definitions, considering active chord, twist, and outer mold line shapes. The computational framework found different adaptive geometries necessary to maximize performance between competing mission stages. When considering active chord, the best designs for each mission stage (hover, loiter, range, and cruise) were a function of the forward flight speed. A more tapered outboard rotor blade improved vehicle performance for lower forward flight velocities, while at higher speeds more lift generation is required and the best rotor blade is less tapered to increase the planform.

An active twist rotor blade was optimized across the same four objectives as the chord/planform optimization, with five embedded torque tubes that alter twist rates between mission stages. In this example, a multi-objective optimization was performed first over all torque tube lengths and twist rates for 10 design variables. During this optimization each design could not be assumed to be morphable to another design, and each evaluation considered only one twist configuration; an alternative approach could consider fives twist rates for each objective, which would result in 25 adaptive design variables. Next, the approximate local frontiers adaptive design process is applied in which 'families' of similar designs with respect to fixed design variables (i.e., actuator lengths) across the entire optimization were evaluated to determine the best adaptive designs. Through this

process, a set of adaptive designs were determine resulting in improved objectives with respect to any fixed design. A second optimization was run to ensure the final objectives were feasible using the single set of fixed design variables determined after the initial optimization.

Adaptive outer mold lines were also investigated for morphing using XFOIL to determine the aerodynamic properties and class/shape transformation equations to define the geometries. Initial studies determined that any aerodynamic optimization required some constraints for structural feasibility, and the focus was shifted to trailing edge camber morphing to ensure the existing spar structure was not altered or thinned down. Geometric constraints on the leading edge shape ensured that the original spar sizing remained the same as the original UH-60 rotor blade.

Next a sequential mission analysis was performed in which a mission was optimized by an overall score as opposed to minimizing individual objectives for each mission. The benefits of camber morphing were demonstrated for a simple, two-stage mission a more complex five-stage mission. Preferred geometries for rotor blade outer mold lines were determined, where similar trade-offs between the best shapes and forward flight speed could be seen. The best outer mold line for the cruise mission stage was very similar to the base rotor blade, while at lower velocities performance was improved by morphing the trailing edge downward and increasing the camber.

5. HIGH-FIDELITY ADAPTIVE ROTORCRAFT OPTIMIZATION

In this chapter, the computational framework from Chapter 4 is extended using higher-fidelity inputs with more detailed adaptive rotorcraft design and optimization studies. First, adaptive computational fluids dynamics results are integrated into the aerodynamic tools, relating airfoil geometries, Mach numbers, and angles of attack to local lift, drag, and moments. These CFD tables consider a single parameterized morphing camber value, limiting any interpolations to three inputs without additional numerical analysis.

Then a series of optimizations are performed for adaptive camber morphing, and both the multi-objective and mission score approaches are compared and contrasted as mission-driven design and optimization techniques. The multi-objective approach is extended beyond single missions to any missions consisting over intersecting unique stage types.

Next, the design space decomposition optimization approach is compared with alternative methods for the Air-Assault mission with five unique mission stage types for multi-objective optimization of active twist, extending the work from section 4.2.2. The approximate frontiers technique is applied to optimizations using two different genetic algorithm selection methods: NSGA-II and NSGA-Adaptive Design (NSGA-AD). Additionally, these optimizations are compared with the alternative approach of considering every adaptive parameter for every objective as a different design variable in the optimization. The final designs are compared for each method, with the approaches outlined in this work determining improved performance for the mission-driven adaptive rotorcraft design and optimization problem.

Full aerostructural coupling is integrated into the computational framework in which changes in rotor blade deflection are considered during the aerodynamic trim flight analysis for hover. Currently this work is limited to the hover state due to computational costs, but the extension could be considered in full optimizations in the future via parallelization or dedicated run times over longer periods of time. Mission analysis with fully-coupled FEA and CFD interpolation tables provide high-fidelity solutions, but the computational costs are still limited for adaptive structures

optimization. Some adaptive elements in the rotor blade were also evaluated where the structural deflection due to embedded shape memory alloy wires was determined in the structural model instead of specified geometrically.

Chapter 5 Outline

This chapter is divided into 3 sections:

- In section 5.1, parameterized camber morphing tables from section 2.2.1 were integrated into the computational framework and compared over specific regions of the rotor blade at low and high loading. Then mission-driven optimizations are performed comparing the methodologies developed for adaptive rotorcraft optimization in Chapter 4.
- In section 5.2 an active rotor blade twist study compares the various multi-objective approaches similarly to the analytical models evaluated at the end of Chapter 3, this time applied to adaptive rotorcraft using the modeling from Chapter 2 and the computational framework developed in Chapter 4. **This study combines almost all the concepts developed in this work: adaptive rotorcraft modeling and analysis, mission-driven design, and engineering optimization techniques to determine optimal adaptive rotorcraft designs using the novel optimization techniques developed in Chapter 3**.
- In section 5.3, a fully-coupled aero-elastic model is integrated into the computational framework designed to reduced the computational costs to determine converged trim flight solutions under mission environments with structural deflections due to inertial and aerodynamic loading.

5.1 High Fidelity Parameterized Camber Morphing

The use of lower-fidelity aerodynamic tools is common in early stages of design. However, low-fidelity tools such as XFOIL and XFLR5 are limited in capturing aerodynamic effects consistent with higher Mach numbers, such as transonic effects and airflow separation. Therefore, it is preferred to consider higher fidelity tools such computational fluid dynamics (CFD) when relating angle of attack and Mach number to aerodynamic constants such as lift and drag.

However, higher-fidelity tools are much more computationally expensive. One solution is to pre-compute the CFD solutions over a wide range of aerodynamic conditions and morphable parameters [36]. However, each additional morphable parameter increases the dimensionality of the input space considered when both generating these tables and when extracting aerodynamic solutions during a system-level analysis. Therefore, the CFD inputs were limited to Mach number, angle of attack, and a single trailing edge camber morphing parameter. Mach number and angle of attack ranges were informed by previous low-fidelity analysis.

At the beginning of a mission analysis the camber morphing tables for each mission stage are generated via interpolating the pre-computed tables for every camber morphing parameter used during the mission as shown in Fig. 5.1. For example, if a five-stage mission considers five unique levels of camber morphing, the 3-D camber morphing tables for each element are used to generate five sets of 2-D aerodynamic tables spanning the appropriate Mach numbers and angles of attack for each camber value, along with aerodynamic constants. The effects of trailing edge parameterized camber morphing on the outer rotor blade span were evaluated under two different loading conditions: the first with a weight of 50 kN, and the second with a weight of 70 kN. The power profiles were generated over a range of trailing edge camber morphing, where -0.1 indicates trailing edge downward deflection, 0.0 indicates an unmorphed rotor, and 0.025 indicates a slight reflex shape. Example trailing edge camber morphing under this parameterization can be seen in Fig. 5.2. Note the figure is not to scale, but shows the ranges of trailing edge adaptivity considered in the CFD tables. The trailing edge camber in this example is morphed from 0.7-0.9 times the radius down the rotor blade span. In Fig. 5.3 it can be seen that for the high lift mission, the slight reflex

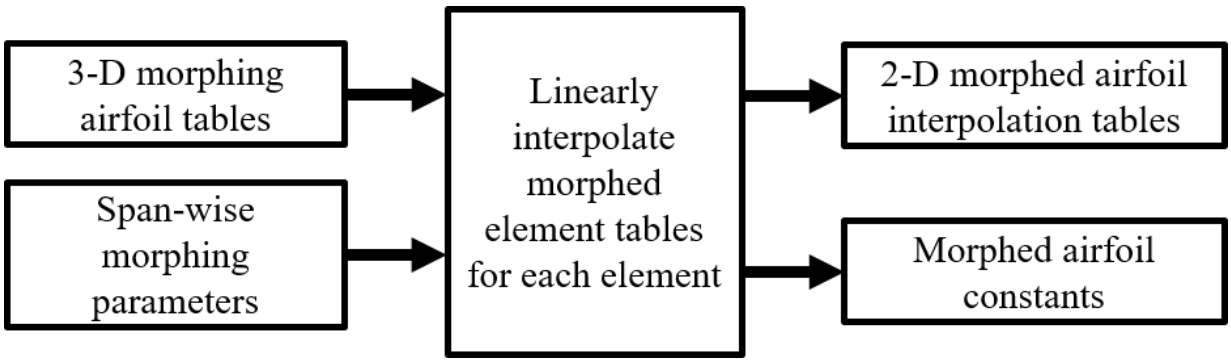

Figure 5.1: The 3-D morphing airfoil tables (spanning level of camber morphing, Mach number, and angle of attack) are interpolated into the appropriate 2-D camber morphing levels required for the current analysis to reduce the interpolation table dimensionality required when determining trim flight.

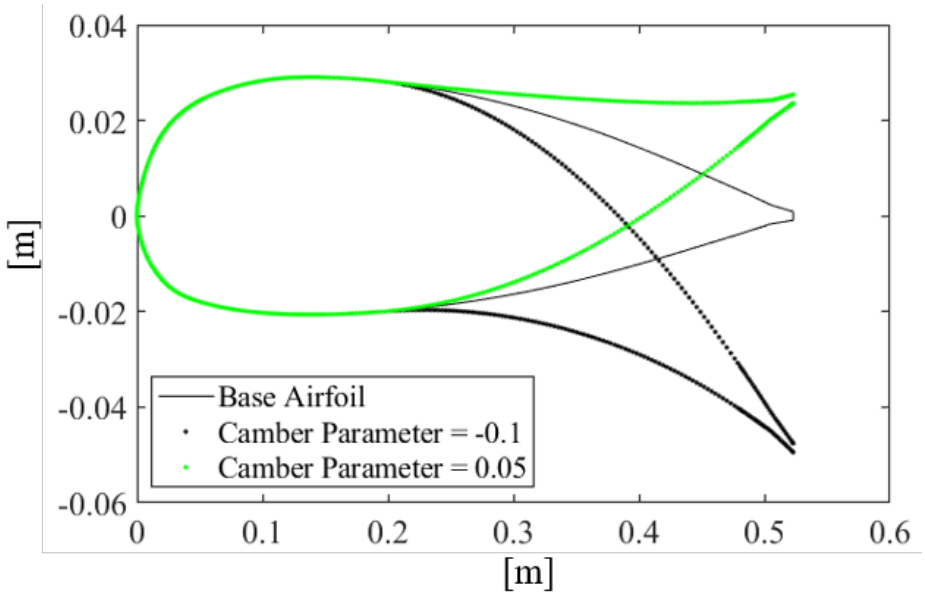

Figure 5.2: The ranges of adaptivity considered in the pre-computed parameterized CFD camber morphing tables generated for this work.

shape performs the worst, while the most cambered rotor blade performs the worst for the low lift mission. Unlike when comparing trim flight solutions between the NASA UH-60 airloads program and the computational results in this work for an unmorphed rotor blade, there does not exist

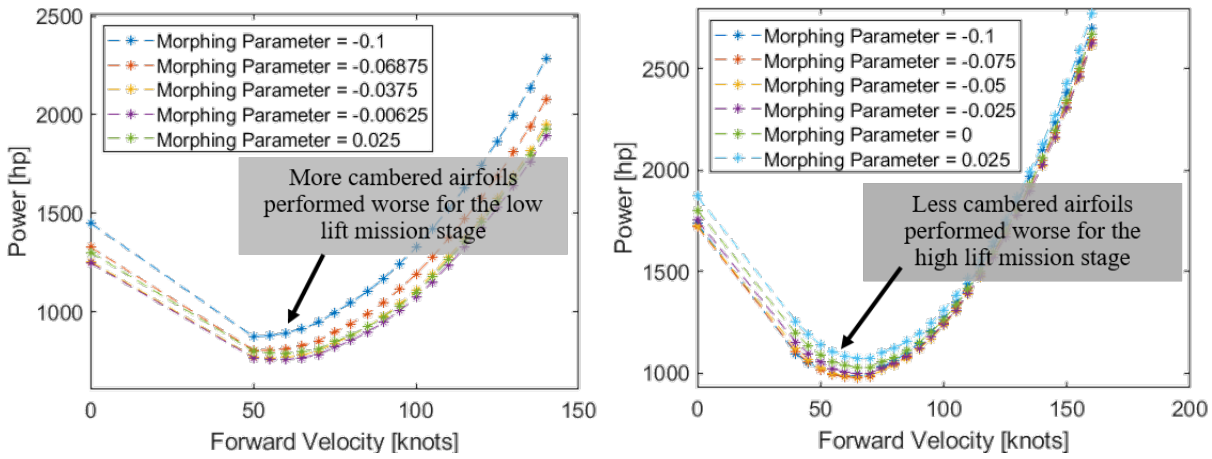

Figure 5.3: The effectiveness of camber morphing varied depending on the total vehicle loading, with more camber (trailing edge deflection down) leading to lower power required power under high lift missions and higher required power for low lift situations.

a UH-60 trailing edge camber morphing direct comparison set of experimental and computational data. However, these trends are similar to those found in works under SABRE [132] where modest trailing edge rotor blade camber morphing over the same region provided worse performance at lower coefficients of thrust but improvements at higher thrust levels.

5.1.1 14-Stage Air Assault Mission

Now consider a longer, more complex mission definition than the previous MEDEVAC mission: a 14-stage Air Assault mission. As mentioned in Chapter 4, this mission is considered an important mission across all four Future Vetical Lift (FVL) capability sets [66] for medium and heavy future rotorcraft. Two constraints were applied to the mission, as shown in Table 5.1. The maximum power during any mission stage cannot exceed a maximum power threshold based on the two engines used on the UH-60 Blackhawk. The other constraint is the common mission requirement of half an hour of fuel reserve at maximum range.

The multi-stage mission selected for the demonstration of genetic optimization in this particular work is a 14-stage Air Assault Mission comprised of five different mission phase types, each based

Table 5.1: Air Assault Mission Constraints

Description	Value	Units
Final Range Time	30	Minutes
Maximum Power (across all 14 stages)	3150	Horsepower

on payload and forward velocity. The mission given in Table 5.2 and Fig. 5.4, considers a UH-60 starting at the Terminal Arrival Area (TAA), climbing up 200 feet, then cruising at the most efficient range speed for 50 nautical miles (nm). The rotorcraft then descends to the Pickup Zone (PZ), where 12 fully equipped soldiers are loaded into the vehicle.

Next the rotorcraft hovers, climbs 80 feet, then cruises at 160 knots for 40 nm before descending to the Landing Zone (LZ). The rotorcraft climbs 80 feet, flies at the best range speed for 40 nm, climbs another 200 feet, then cruises at the best range speed for the remaining 50 nm before descending and unloading the soldiers back at the TAA. The final stage considers a maximum fuel reserve requirement of 30 minutes flying at the best range speed. The optimal morphing rotor blade is assured to be uniquely configured (e.g., a set of five trailing edge camber morphing settings, one for each flight condition).

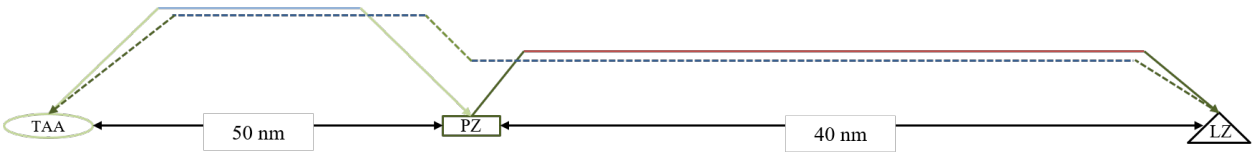

Figure 5.4: A proposed 14-stage mission with five mission stage types. The mission starts and ends at the Terminal Arrival Area (TAA).

The multi-stage mission is composed of different forward flight speeds under multiple loading conditions. As such, the best rotor blade for each mission stage type is assumed to be unique. The optimal morphing design will thus contain a set of different adaptive parameters corresponding to

Table 5.2: Proposed 14-stage mission

Mission Stage Description	Mission Stage Type	Time/Distance
1) Hover over TAA at 6,000 ft. altitude on 95 deg. day	Light Hover	0.5 min
2) Climb 200 ft.	Light Hover	2 min
3) Cruise at best range speed	Light Range	50 nm
4) Descend to land at PZ to pick up 12 soldiers	Light Hover	1 min
5) Hover before takeoff	Heavy Hover	0.5 min
6) Climb 80 ft.	Heavy Hover	1 min
7) Cruise at 160 knots	Heavy Cruise	40 nm
8) Descend to land at LZ (using hover blade shape)	Heavy Hover	1 min
9) Climb 80 ft. using hover	Heavy Hover	2 min
10) Cruise at best range speed	Heavy Range	40 nm
11) Climb 200 ft.	Heavy Hover	2 min
12) Cruise at best range speed to TAA	Heavy Range	50 nm
13) Descend to land at TAA unload soldiers	Heavy Hover	1 min
14) Determine max range based on fuel reserve	Light Range	

specific mission stages. It is necessary to define the parameterized adaptive optimization problem as applied specifically to this 14-stage mission.

5.1.2 Optimization Problem Statement

Toward optimizing the performance of a rotorcraft performing this diverse mission, consider a parameterized design problem in which camber could be altered elliptically over a single rotor blade span region, as shown in Fig. 5.5, where the level of maximum camber morphing in the center of such a region can be altered via shape memory alloys (SMA) or other internal actuation mechanisms. The local level of camber morphing for each blade element is thus based solely on this maximum level of camber morphing at the center of the span-wise region (δ). But the morphed region location (starting position x_s and region length l_c) is also described by two design variables, which cannot be altered between mission stages. The design variables defining camber morphing placement must be fixed during manufacturing of the blade, and are the fixed design variables in

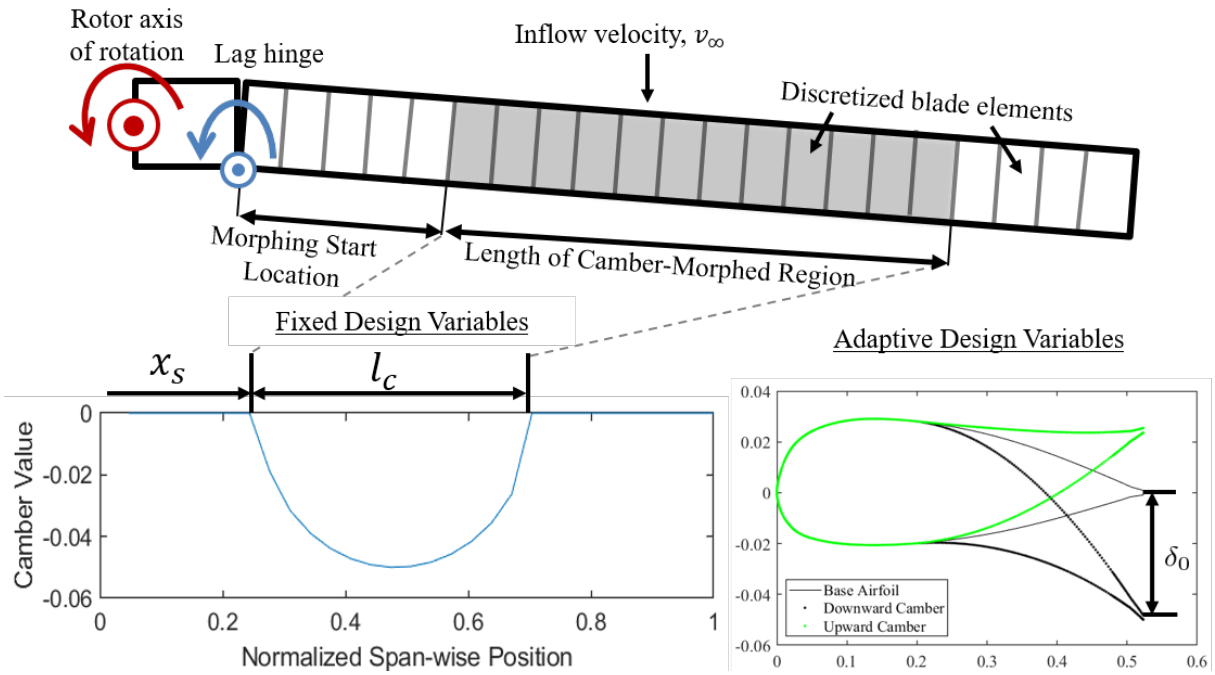

Figure 5.5: The camber of the rotor blade is morphed over some finite a region of the rotor span. The design then considers both fixed design variables (i.e., defining location of the adaptive camber region; $\mathbf{x}^f = \{x_s, l_c\}$) and adaptive variables (i.e., degree to which the trailing edge is deflected up or down in flight).

the design problem (i.e., $\mathbf{x}^f = \{x_s, l_c\}$). The design variable altered between mission phases in this example is the level of camber morphing, which is the adaptive design variable (i.e. $\alpha = \{\delta\}$).

The local camber of each element in the morphing region is defined by the following elliptical distribution function:

$$\delta(r) = \delta_0 \left[1 - \left(\frac{|r - (x_s + 0.5l_c)|}{l_c/2} \right)^2 \right], \tag{5.1}$$

where δ_0 is the adaptive design variable denoting maximum camber morphing, δ is the local camber morphing parameter for any blade element in the morphed region, and r is the span-wise location of the blade element as measured from the axis of rotation. The camber morphing fixed and adaptive variables can be seen in Table 5.3.

The design variable bounds are given in Table 5.4. The lower bounds of fixed design variables x_s and l_c were selected to ensure a minimum distance between the rotor lag hinge and adaptive ac-

Table 5.3: Camber Morphing Design Variables

Operational Variable	Camber
Fixed Variables	Morphing Start Location
	Morphing Region Length

Table 5.4: Adaptive ($\boldsymbol{\alpha} = \{\delta_0\}$) and fixed design variables ($\mathbf{x}^f = \{x_s, l_c\}$) considered in the current optimization problem

Variable Type	Design Variable		Lower Bound	Upper Bound	Units
Adaptive	Deflection	δ_0	-0.05	0.025	m
Fixed	Start Location	x_s	0.4	4	m
	Region Length	l_c	0.4	4	m

tuation mechanisms. The fixed design variables upper bounds were selected such that the morphing region did not extend beyond the rotor tip while also allowing for an adaptive region extending no more than half the rotor span. The limits for the adaptive design variables δ consider morphing the trailing edge from downward camber (negative δ) to a reflex shape (positive δ). The final design consists of a single set of fixed design variables $\mathbf{x}^f$ and requires a set of adaptive design variables $\boldsymbol{\alpha}$, one for each mission stage type. The best adaptive blade design for the mission described in Table 5.2 will then be defined by i) the fixed (i.e., manufactured) size and location of the camber morphing region for all mission stages ($\mathbf{x}^f = \{x_s, l_c\}$ and ii) a set of five adaptive cambers ($\boldsymbol{\alpha} = \{\delta_0\}$), each associated with an individual mission stage type, these being dynamically changed as the mission progresses. A large, genetic multi-objective optimization must consider all camber morphing regions and their respective scheduled cambers to capture the entire design space and determine optimal performance.

Next, two methods for mission-driven design and optimization are compared using the problem statement above. First, a mission-score optimization is performed for the mission in which the

location is fixed and the entire mission run sequentially. Then a multi-objective optimization is performed considering all locations over the course of a mission then selects the best region from families of similar designs as described in section 3.2.2. Each optimization is run in ModelCenter, with 4 evaluations run in parallel at a time.

5.1.3 Mission Score Optimization

Mission-based design optimizations can be performed either by determining the best performance for each independent mission stage type (see Table 5.2, referred to here as *multi-objective*, or by scoring mission performance via a single all-encompassing metric, referred to here as *mission score*. The mission score approach to improving performance simulates the entire mission through each stage, defining a metric at the end of the full mission to quantify overall performance. To accurately capture vehicle response throughout the mission, stages are evaluated sequentially, considering the change in weight due to fuel consumption during each stage. Such a simulation requires evaluating each mission stage while considering the effects of previous stages for each design.

The single score approach results in a final configuration for a specific mission under a single set of inputs (temperatures, initial fuel, number of passengers, etc.) but may generalize less well when compared with the multi-objective approach. Often aircraft are expected to perform a range of missions, particularly high-performance rotorcraft. There is particular interest in future rotorcraft design toward covering sets of missions [66], which will be discussed later in this work.

The mission score for this design problem was, for a fixed initial fuel input, to maximize the final range time at mission end. A design that maximizes final range time will also minimize fuel consumption before the final range stage by determining the best set of designs for the first 13 stages, followed by the best light range design for the final stage.

The common mission requirement of a 30 minute minimum fuel reserve at the final stage start was also applied. Increased initial fuel increases range and fuel consumption due to larger initial weight; the initial fuel mission parameter was selected such that the base rotor blade completed the proposed 14-stage mission while meeting the fuel reserve constraint. The selected initial fuel

ensured designs worse than the base, or reference, rotor blade were discarded, as there is little to no value in an adaptive rotor blade worse than the unmorphed rotor blade. A population of 24

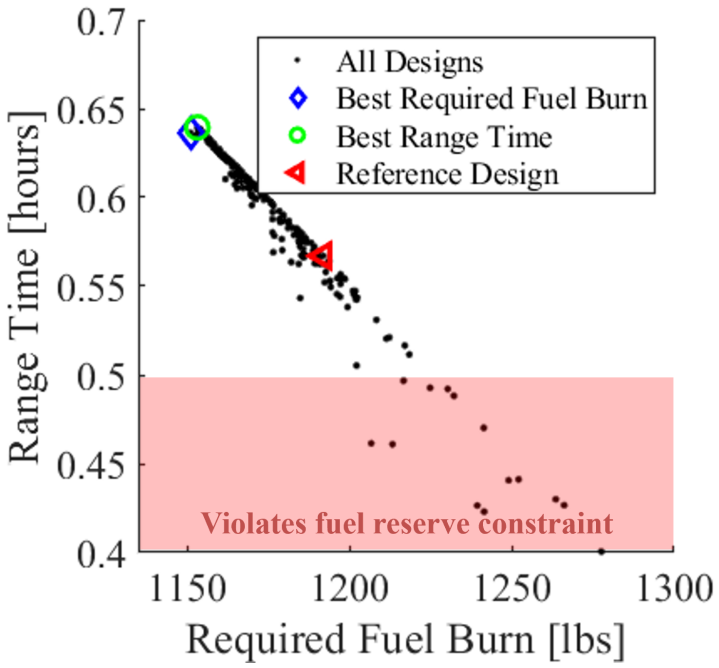

Figure 5.6: The best design considering a single mission score minimized fuel burn and maximized final range time via morphing between mission phases.

individuals was optimized over 15 generations to generate a final adaptive camber blade design (i.e., a blade with a fixed manufactured region $\mathbf{x}^f$ dynamically morphed to five different camber levels). The average evaluation time for each design was 4.6 minutes. The objective space can be seen in Fig. 5.6, where the best design minimizes fuel burn and maximizes the final range time. The final design will be compared to the multi-objective optimization blade in the following section in Table 5.6.

Since this fully-coupled optimization evaluates mission performance through all 14 stages sequentially, interactions between design variables arise due to the coupling between mission stages. For instance, poor initial stage performance leads to higher fuel burn, decreasing required power

at later stages due to a lower overall weight. These interactions may increase computational costs as it becomes more difficult to consider the individual effects of design variables through a single mission score.

5.1.4 Multi-Objective Design and Optimization

An alternative to the mission score approach could consider individual objectives separately. For instance, fuel burn is directly associated with required power for trim flight in all phases of flight. Thus, an adaptive design that minimizes power at each stage of the mission likewise reduces fuel burn.

It is worthwhile to consider the continuous nature of the fixed design variables (active camber region) versus the discrete blade element approaches being employed for analysis. In the previous study, designs with almost identical morphing regions were treated as distinct and thus incompatible, even if the differences in location were smaller than the width of discretized blade elements themselves. As previously demonstrated, there is a large benefit to considering sets of designs with similar fixed design variables $\mathbf{x}^f$ as compatible in a multi-objective optimization. Thus, the approximate local frontiers approach from section 3.2.2 provides an alternative to a mission-score approach while still considering every unique mission stage type.

To apply this new approach, the same 14-stage adaptive mission was then evaluated considering only the five specific mission stage types instead of the entire 14-stage sequential mission (from Table 5.2). This multi-objective study did not consider coupling between mission phases and was not evaluated using a single mission score, but instead using a multi-objective optimization across mission stage types.

The multi-objective optimization used the NSGA-II [37] evolutionary algorithm with a population of 24 individuals over 15 generations to generate a non-dominated frontier of designs. The genetic optimization generates an ever-changing population of designs in the search to minimize required power for each objective. The resulting population of all designs over every generation generated in the genetic optimization was explored to determine the best set of designs with approximately similar fixed design variables $\mathbb{D}_i$. For each individual in the optimization, all designs

with camber morphed regions defined as being within 2% of each other were considered similar. The utopia point method described in section 3.1.1 was used to select the best adaptive design represented by an approximate local frontier $\hat{\mathbf{J}}^*$.

As seen in Fig. 5.7, the selected adaptive design that performs the best across mission stages can morph approximately to any blue dots by changing only camber; the fixed design variables are approximately the same. It is worth noting that Fig. 5.7 shows a 2-D projection of the 5-D objective space, with a dimension for each objective. The singular descriptor of best performance across all possible morphing states is the 'Adaptive Utopia Point' from section 3.1.1, while any fixed design is limited to the trade-offs inherent in the Pareto frontier. A comparison in power reduction using both the best fixed and adaptive designs is shown in Table 5.5 alongside the unmorphed rotor blade. Performance can be increased across these specific mission stages via optimization toward a single fixed design, but further improved via the addition of active camber morphing. Optimizing a

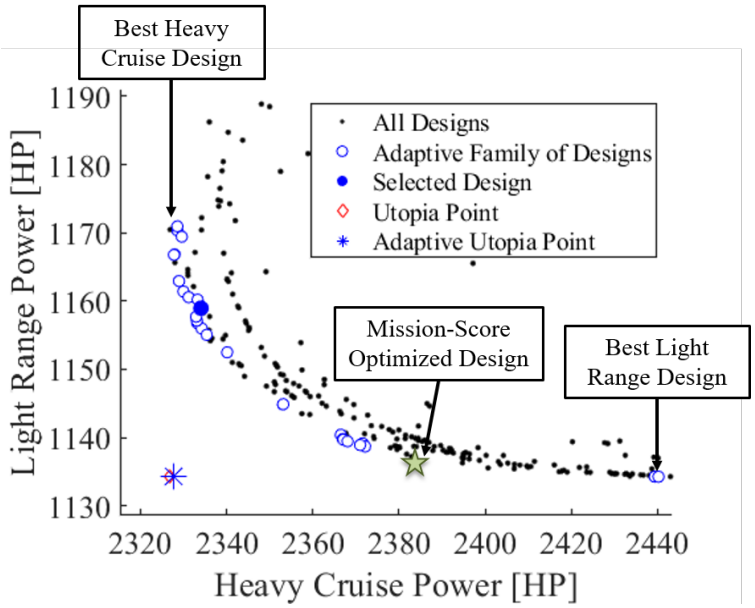

Figure 5.7: The multi-objective space, along with the best set of similar adaptive designs $\mathbb{D}_i$. Considering adaptive camber allows for an adaptive design with objectives near the utopia point (i.e., the rotor blade can be morphed such that minimum power is required for both heavy cruise and light range.)

Table 5.5: Mission Stage Trim Power [HP]

Mission Stage Type	Fixed Power Reduction	Adaptive Power Reduction	Base Power
Light Hover	14.27	14.87	1375
Heavy Hover	31.68	39.94	1982
Heavy Cruise	124.0	187.3	2515
Heavy Range	44.84	67.18	1493
Light Range	-1.67	1.91	1136

fixed rotor blade specifically for the 14-stage Air Assault Mission led to a morphing rotor blade design reducing required power when compared to the base rotor blade; the existing rotor blade cannot be optimized for every mission stage under every set of flight conditions. The base rotor blade performed almost identical to both the fixed and adaptive optimized rotor blades for heavy range but much worse for heavy cruise, a mission stage requiring nearly maximum UH-60 power. The adaptive rotor blade required the least power across all five objectives due to the addition of adaptive camber morphing parameters while using the same population of designs generated by the genetic optimization to determine the best fixed design.

A mission-specific morphing rotor blade was generated via a multi-objective optimization to determine the best local rotor camber changes to minimize trim power under five flight conditions. This optimization functioned as a targeted search to generate the best morphing regions and cambers. Next the population was post-processed to determine the best fixed region of camber morphing based on families of similar designs with respect to the fixed design variables (morphing location). By minimizing the distance between the best realizable design and the utopia point, the best morphable design was algorithmically determined, resulting the increased in power reduction only reachable via morphing.

Two methods of mission-based adaptive rotorcraft design were explored in this work: a mission score and a multi-objective approach. The mission score method defined mission-wide objectives

Table 5.6: Final Mission-Based Adaptive Design Variables

Design Variable	Mission Score Optimization	Multi-Objective Optimization
Camber Morphing Parameters	-0.48	-0.59
	-0.87	-0.84
	-0.73	-0.89
	-0.76	-0.89
	-0.63	-0.33
Start Location	0.058	0.059
Region Length	0.478	0.419

as a single score without considering individual mission stages in the optimization. The multi-objective method considered objectives for each mission stage type without considering an overall mission-wide objective. The objectives used in the framework are only limited by the modular modeling methods employed. BEMT and TRAC were used to determine trim power and fuel burn, but other objectives could be incorporated into the general computational framework if given the modular analysis tools, such as audio or structural objectives.

The final adaptive designs for both optimizations can be seen in Table 5.6. The adaptive regions, described by the Camber Start Location and Camber Region Length variables, found similar though not identical regions for optimal camber morphing. The camber morphing regions were inboard the rotor span, nearing the lower bound for the normalized start location of 0.05 along the normalized rotor span with a Camber Region Length of between 0.4 and 0.5. The optimal locations suggests, for this mission, downward camber morphing across an inboard region of the blade yields improved performance by increasing the local trailing edge camber various degrees throughout the mission. Both mission-based adaptive optimization methods generated a morphing rotor blade design with a unique adaptive variable for each mission stage type, but with different results and levels of computational cost. The average multi-objective evaluation runtime was 2 minutes, highlighting the relatively efficient computational cost to evaluate UH-60 trim flight considering

several subsystems while incorporating high-fidelity CFD data as well as the improved runtime per evaluation when compared to the fully coupled model. The multi-objective optimization only evaluates five mission stages while the mission score optimization requires evaluating all 14 mission stages, resulting in the faster runtime. Both studies were run on the same machine under the same optimization parameters with no other computationally expensive work run at the same time (the machine was running under a similar computational capacity for both studies with CPU and memory usage under 50%).

5.1.5 Multi-Mission Adaptive Rotorcraft Design

The multi-objective mission optimization described in section 3 considers the individual stages in a mission. The final rotor blade is optimized for all missions comprised of a combination of the five mission stage types in the 14-stage mission from Table 5.4. Thus, a multi-objective optimized rotor blade can be designed to span across multiple missions.

Now consider not one but three separate missions: i) the 14-stage Air Assault Mission, ii) a shorter Heavy Load mission, and iii) a MEDEVAC mission. These mission definitions can be found in Appendix D. One approach to designing a rotor blade for each mission is to perform a mission score optimization for each mission, resulting in 3 unique rotor blades. Another approach is to perform a multi-objective optimization across all unique mission stage types from all three missions. This multi-objective solution would yield a rotor blade which could morph toward preferred geometries across not only mission stages, but also missions. Such a rotor blade is termed a Mission-Agnostic Adaptive Blade (MAAB) in this work.

The three missions considered consist of combinations of six distinct mission stage types: hover, cruise, and range under loosely defined ‘light’ and ‘heavy’ loading conditions. The average ‘light’ and ‘heavy’ weights for each mission across all three missions were used in the analysis because the weights were slightly different for each mission stage, as shown in Table 5.7. Not every mission stage type is included in every mission.

An optimization was performed to determine a rotor blade with an adaptive camber morphing region with an elliptical morphing distribution as described previously in Fig. 5.5 across all three

Table 5.7: Weights across all Stage Types for Three Selected Missions

Mission Stage Type	Air Assault Weight [lbs]	MEDEVAC Weight [lbs]	High Lift Weight [lbs]	Average Weight [lbs]
Light Hover	13730	13530	13130	13463
Heavy Hover	17110	15720	17130	16987
Light Range	13730	13530	13130	13463
Heavy Range	18110	N/A	17130	17620
Light Cruise	N/A	13530	N/A	13530
Heavy Cruise	18110	15720	N/A	16915

missions. The required power for trim flight across each of the six unique mission stage types was minimized through a genetic optimization with a population of 24 designs. The optimizations were evaluated after 40 generations, but also run until convergence for comparison.

The mission-agnostic multi-objective optimization approach is applied to these six mission stages similarly to the previous single mission multi-objective optimization. The six unique camber shapes the MAAB morphs between are shown in Fig. 5.8.

The mission score optimization was also run for each mission, resulting in three different mission-specific rotor blades. The three mission-optimized rotor blades and the mission agnostic adaptive blade were then evaluated comparing the fuel burn required to complete each mission with the fuel burn required using the reference (unmorphed) rotor blade. A graphical comparison between the fuel savings for each mission using each adaptive rotor blade is shown in Fig. 5.9. When considering to overall fuel savings, each mission score optimized rotor blade performed the best for its respective mission, while the MAAB performed the best across three missions. The MAAB required many less evaluations and performed almost as well across each mission as the rotor blade optimized specifically to that mission. Thus, the MAAB is a more general solution for an adaptive rotor blade spanning multiple missions while also requiring less computational costs to develop.

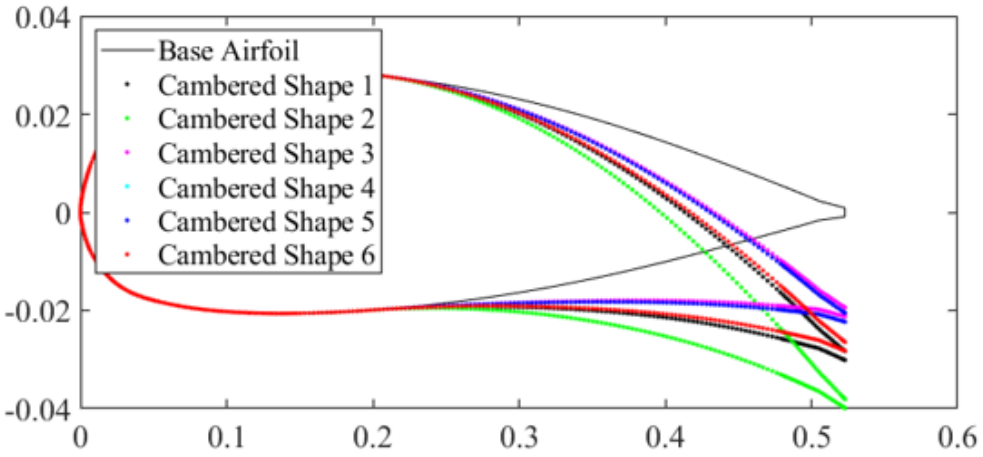

Figure 5.8: The Mission-Agnostic Adaptive Blade morphs between 6 uniques shapes, depending on the specific mission stage type.

5.2 Design Space Decomposition Comparison for Active Twist

It is also worth noting that the design space for the adaptive camber morphing optimization was fairly small, with one adaptive (over five objectives) and two fixed design variables. In contrast, consider the five torque tube design and optimization problem statement from Chapter 4. The multi-objective approach considered five fixed design variables (5 actuator lengths) and five adaptive design variables, the twist rates for each actuator. The design variables L_i and $\hat{\theta}$ and respective ranges are shown in Table 5.8 for reference.

Without the multi-objective approach and design space decomposition methodology developed in this work, two optimization strategies are typically applied to determine both the best lengths and twist rates for each torque tube over each mission stage. One solution to solve this design and optimization problem is the mission score approach: each mission has a set of fixed design variables and alters the adaptive design variables over each mission stage or mission stage type. This would result in either 25 adaptive design variables (5 actuator twist rates for each of five mission stage types) or 70 adaptive design variables (5 actuator twist rates for each of 14 mission stages). Similarly, a multi-objective approach could also optimize over five objectives, resulting in 25 adaptive design variables (five twist rates over five objectives), but with a faster evaluation

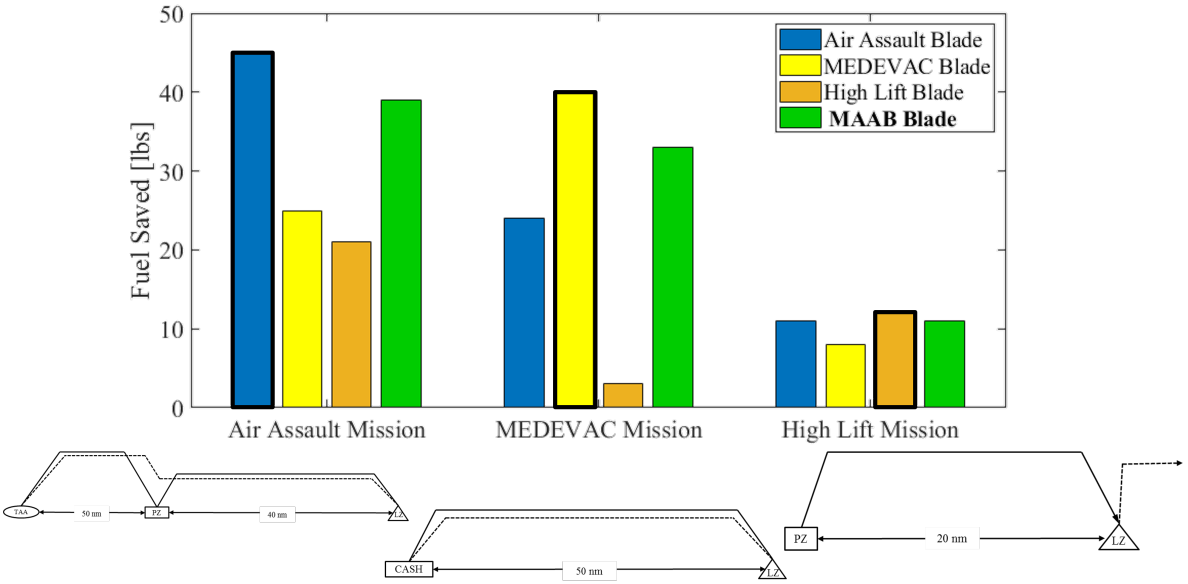

Figure 5.9: The fuel savings using each adaptive rotor blade for each mission. Each blade performed the best for its specific mission, while the MAAB performed the most consistently across all three missions.

Table 5.8: Torque Tube Optimization Design Variables

Design Variable	Description	#	Variable Bounds
L_i	Normalized i^{th} SMA actuator length	5	$0 \leq L_i \leq 0.4$
$\hat{\theta}_i$	SMA actuator twist rate	5	$-2.0 \leq \hat{\theta}_i \leq 2.0 \ \left[\frac{\text{deg}}{\text{m}}\right]$

time due to only solving for five trim flight conditions (one solutions for each mission stage type) instead of 14 (considering the entire mission sequentially).

5.2.1 Multi-Objective Optimizations

Two multi-objective optimizations were run to determine the best adaptive twist rotor blade for the 14-stage Air Assault mission with five SMA torque tubes embedded in each rotor blade. The first approach was to apply the local frontiers method after an initial 10-design variables optimization over the five torque tubes lengths and five twist rates, one for each mission stage. After the optimization, the best families of similar designs were determined to select a final adaptive design.

In the second approach each design has a single set of five torque tube lengths, but each objective is given a set of five torque tube twist rates. By this problem statement setup, each design is realizable and considers morphing between mission stages. However, this increases the design space as previously mentioned, and the runtime to determine trim flight for all five mission stages typically ranges from 10-15 minutes. Thus (without parallelization) a 20 generation optimization over a population of 24 individuals requires approximately 100 hours, which was selected as a trade-off between runtime and design space exploration. For comparison, a 100 generations of 100 designs optimization such as those explored in Chapter 3 for simple analytical functions would require

The optimization is performed with a similarity parameter of $\epsilon = 0.025$ to determine the bet adaptive family using the utopia point distance metric. Due to the relatively long runtimes, the design space explored was relatively sparse, and a 'family' of similar designs where $\epsilon = 0.025$ represents $3.125 * 10^{-5}$ percent of the design space. The other optimization considered all the design variables over all the optimization stages, ensuring feasible designs.

The final designs from both optimizations were compared, and the differences between trim power for each objective can be seen in Fig. 5.10. Both optimization methods determined adaptive configurations resulting in lower trim power, but optimization using the novel design space decomposition methodology outperformed the alternative method due to a smaller design space the the success in still finding adaptive 'families' of similar designs.

The largest gains in power reduction with respect to the unmorphed rotor blade were under the hover mission stages, while the reference blade performed almost identically for the light range mission stage. This suggests the twist schedule for this blade was design prioritizing light range and forward flight over hover, and the addition of active twist elements could reduce or eliminate such trade-offs via adaptivity.

5.2.2 Effects of ϵ on Multi-Objective Adaptive Designs

In section 3.2.2, a similarity parameter ϵ is introduced to approximate similar designs with respect to fixed design variables (x^f) to generate a set of designs (D_i). The set is assumed to have

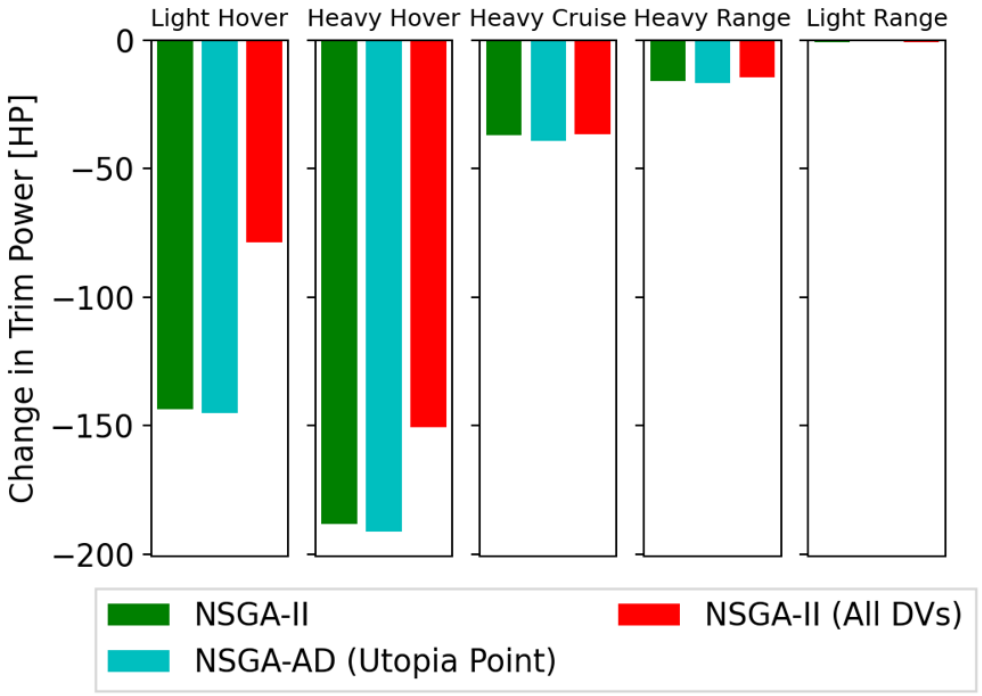

Figure 5.10: The addition of active twist reduced the required power when compared to the unmorphed rotor blade, with the design space decomposition methods more effectively searching the objective space for final adaptive configurations.

fixed design variables similar enough that adaptivity is possible between each design in the set. However, the selection of ϵ must be a trade-off between family size and potential adaptivity, while the family size is also a function of data density.

A large ϵ yields a larger set of $\epsilon - similar$ designs, but the designs may not be reachable via morphing if the fixed design variables are not 'similar'. Likewise, a small ϵ may exclude potential adaptive designs. At the bounds, when $\epsilon = 0$, each family D_i contains only a single design, while when $\epsilon >= 0.5$ a design in the 'center' of the design space (located halfway between the lower and upper bounds for each fixed design variable) would have a family of similar designs that spanned the entire population.

The effects of ϵ are explored specifically for multi-objective adaptive rotorcraft missions, using the previous five-objective mission and altering the similarity criteria. The final population from the previous rotor blade twist optimization is searched for adaptive designs using a range of similarity parameters $\epsilon \in \{0.001,\ 0.01,\ 0.02, 0.03,\ 0.04,\ 0.05\}$, and compares the trim hover power from the best 'adaptive' design to the approximate utopia point across the entire optimization for each objective. It is worth noting that as $\epsilon \longrightarrow 1$, the best adaptive designs will approach and eventually reach the approximate utopia point, but the assumption of similarity will clearly be violated. The convergence of the adaptive reachable objectives is shown in Fig. 5.11.

It can be seen that, while the low data density results in higher power requirements for $\epsilon = \{0, 1\}$, even an ϵ value of 2% of the design space bound results in similar performance to the approximate utopia point, while still limited adaptive 'families' to a cluster spanning $3.2x10^{-7}\%$ of the 5-D fixed design space. Since each adaptive design is a set of discrete points, it can be seen that in this example increasing ϵ from 0.01 to 0.02 improved all the objectives except light range, which actually increased in required power. This is because the design with the best adaptive light range for $\epsilon = 0.01$ had different fixed design variables than the adaptive family selected for $\epsilon = 0.01$, and the new best adaptive family was selected due to the improved performance across the other 4 objectives.

However, such an analysis assumes the overall performance is realizable even though the fixed

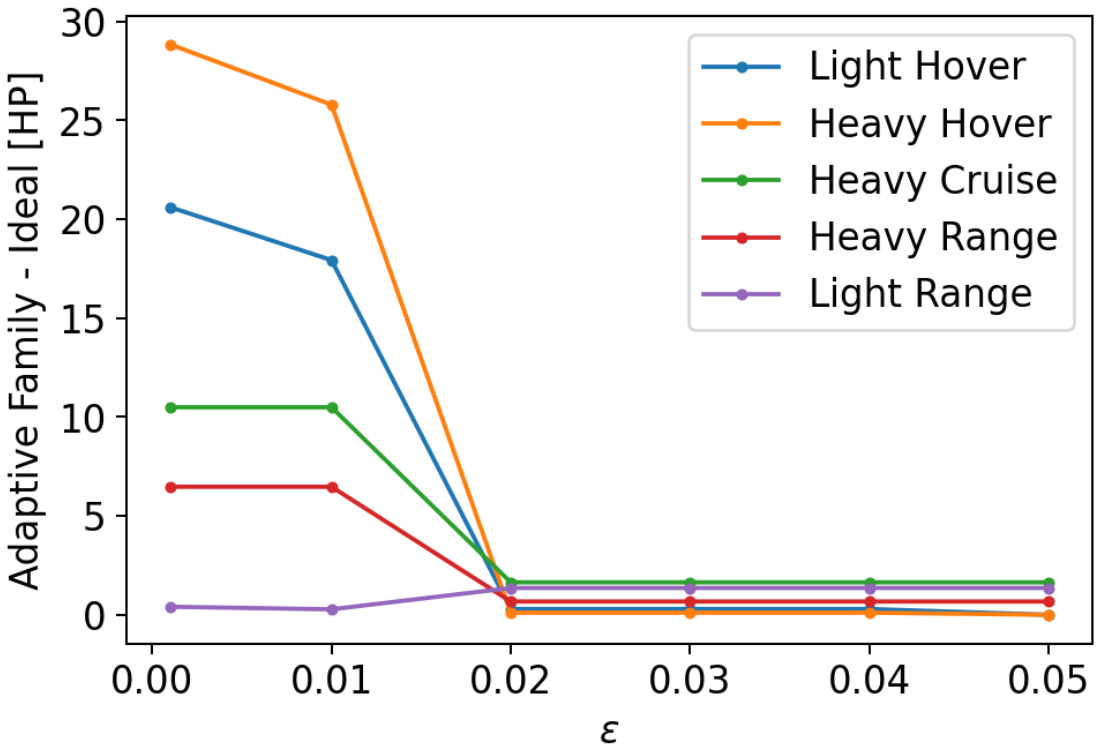

Figure 5.11: As ϵ increases, the farther apart designs can be while still considered 'similar,' with a limit at the approximate utopia (or ideal) point. However, these designs are 'approximately' similar and must be compared with the final results.

design variables are not identical. In some previous studies in this work a final optimization was performed to determine the final configuration after determining fixed design variables, and typically the adaptive variable bounds could be significantly reduced as well. The best similarity parameter may vary based on variable sensitivities, data density, and even the variability between identically-defined optimizations. Such variations were seen in the design studies from Chapter 3, where much faster analysis allowed for large numbers of optimizations to be evaluated as distributions instead of simply individual optimizations. However, determining the true objectives for each adaptive design only requires recalculating the objectives once for the best design. A range of similarity parameters were next used to determine the best predicted designs from Fig. 5.11 and compare the actual objectives (still compared with the same ideal approximate utopia point). In

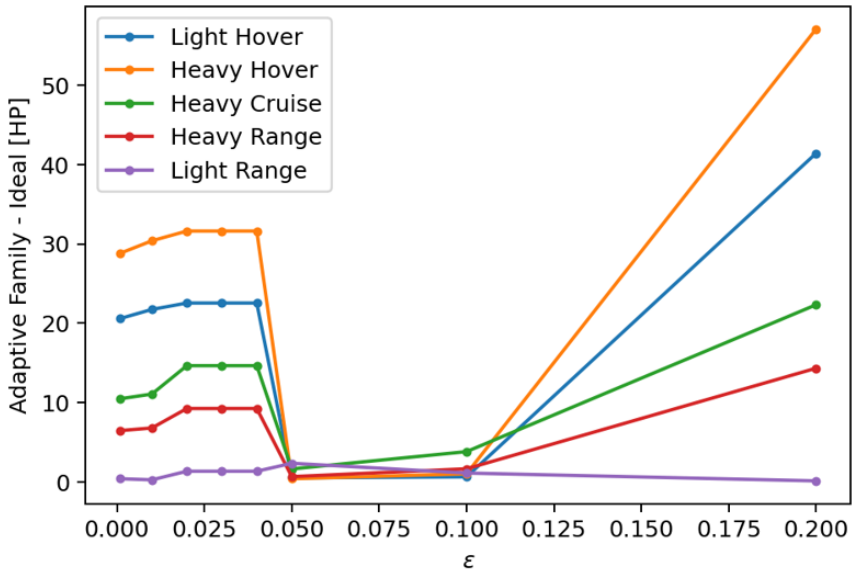

Figure 5.12: The objectives for the best families of adaptive designs when run using a final single set of fixed design variables.

this case, the best set of fixed design variables (meaning the fixed variables from the design that determined the center of the 'family') was run with the best adaptive design variables corresponding to each objective (which are selected based on the objectives from the entire 'family.')

In Fig. 5.12 the differences between objectives of these final evaluations and the 'ideal' design are shown over a range of similarity parameters. Additional values were evaluated for larger similarity parameters to show the divergence at larger ϵ parameters of $[0.1,\ 0.2]$. In this case, the best fully realizable adaptive design (as determined via this post-processing procedure) was determined to have a similarity parameter $\epsilon = 0.05$. Due to the relatively low data density and coupling between design variables (altering the length of the first torque tube will adjust the placement of the next torque tube, for instance), it was expected that too small and too high of ϵ values would provide worse objective when compared with some unknown value between the extremes.

5.2.3 Comparison Between Active Structures

Comparing the previous multi-objective results, it can be seen that different adaptive technologies offer various performance improvements depending on the mission stage type. For instance, the active camber morphing reduced trim power over 180 HP for heavy cruise while active twist only reduced heavy cruise power by less than 50 HP. Neither active twist nor camber offered any improvement for the light range mission stage, suggesting the static rotor blade was designed primarily for efficient flight with only a pilot and copilot. Meanwhile, the active twist rotor blade reduced trim power for hover by 150-180 HP depending on weight, while the active camber reductions for light and heavy hover were all under 50 HP. Both active systems resulted in modest heavy range power reductions of 40-80 HP. Overall, active twist primarily improves trim power for hover, while active camber primarily improves performance for heavy cruise, while both provide modest improvements for heavy range and the static rotor blade is optimized for light range.

5.3 Fully Coupled Aeroelastic Design

The majority of this work focuses on modeling, analyzing, and optimizing the geometry of a rotor blade that, beyond specified morphing, is rigid. The rotor blade begins with a manufactured geometry, then some internal actuation is added that alters the rotor blade profile. During previous aerodynamic optimizations, *in-situ* structural analysis was limited to low-fidelity sizing of necessary actuation. In some cases, aerodynamic results were input into a structural optimization to determine the placement of internal topologies.

However, the effects of inertial and aerodynamic loading were not considered in a fully-coupled high-fidelity environment due to the computational costs. Existing finite-element (FE) rotor blade models used in this work typically have a runtime of a few minutes. If the average finite element analysis (FEA) evaluation takes 3 minutes and a single forward-flight evaluation (using TRAC) requires 288 evaluations (four rotor blades at 72 increments around the hub), then one iteration of solving for trim flight requires 14.4 hours. The typical trim flight solution in forward flight requires >350 iterations using the fsolve function in Matlab, meaning a single forward flight mission stage

solution would require over 180 days. The hover BEMT code requires fewer evaluations due to radial symmetry, but would still require several hours for a single mission stage trim flight solution.

5.3.1 Computational Implementation

A fully-coupled aeroelastic procedure must therefore dramatically reduce the number of required structural FE evaluations to be useful as a design tool. In this section the fully coupled model is first limited to the hover BEMT code, but could be applied to forward flight (using TRAC) with associated increases in runtime. One approach to reduce the number of FEA evaluations as described in section 2.4 by only updating the rotor blade deflection after a trim solution was determined. The algorithm is implemented in Matlab and Python as described in section 2.4.

An 11-stage mission were defined with increasing loads in hover for each mission stage. First, each mission was run with no deflection and the collective pitch for each mission stage was determined for trim flight with a rigid rotor blade. The pitch for each mission stage, shown in Fig. 5.13, could then inform the initial guesses for each future mission stage. Each mission stage evaluation uses the solution to the rigid rotor trim solution due to convergence issues; the solution sometimes falls into local minima if the initial guess is too far away. Another method to improve robustness, also included in this work, is a restart function, in which a new initial guess is attempted if convergence is not reached after a set number of iterations. However, a restart function inherently increases the runtime and provides no guarantee for convergence. Since a single mission solution requires convergence of trim flight, structural deflection, and local rotor blade inflow, any procedures that improve convergence or reduce the number of iterations are valuable. The four-stage mission with pre-computed initial collective pitches quickly converged, requiring only 4 iterations for each mission stage. The residuals for all 4 mission stages can be seen in Fig. 5.14, converging logarithmically. The collective pitches did not change significantly while iterating through deflections for each mission stage since the initial guesses are the solutions for a rigid rotor blade.

The 11-stage mission was run with an initial collective pitch guess of 8 degrees and converged similarly to the four-stage mission, requiring 2two to five iterations depending on the mission stage. Mission stage six converged in only two stages, while the higher loaded mission stages 10 and 11

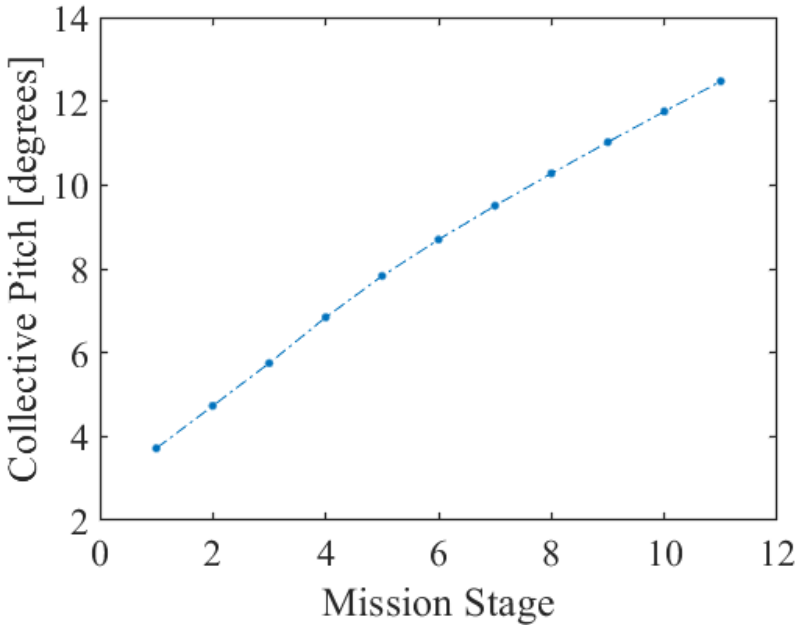

Figure 5.13: The trim flight collective pitch for each mission stage was first determined for a rigid rotor blade.

required five iterations each, as shown in Fig. 5.15. Thus, coupling only between trim solutions is shown to reduce the number of FEA calculations from between 60 and 120 (for Hover) to between two and five. The faster convergence of stage 6 was due to the initial guess of 8 degrees being close to the expected trim solution, as shown in Fig. 5.14.

5.3.2 Fully-Coupled Active Rotor Twist

Additionally, active rotor twist was introduced into the model when developed by Trent White. Angled SMA linear actuators in the rotor blade spar recover transformation strain, twisting the rotor blade. The internal spar configuration can be seen in Fig. 5.16. The SMA-driven twist was simulated as a thermal contraction for computational efficiency, though a more comprehensive material model could be substituted for a more detailed analysis. Another input, total SMA strain, was added into the BEMT code and passed to the FEA model. First, the model was run under a single load case considering 0, 1, and 2% strain. The change in local twist angle after aerodynamic loading and actuation changes as a function of span-wise location, as seen in Fig. 5.17. The 0%

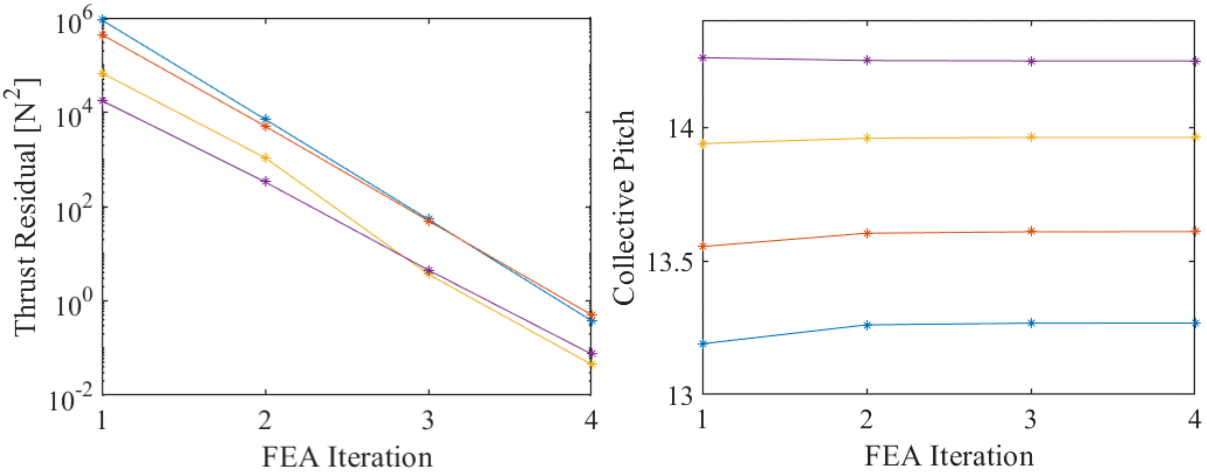

Figure 5.14: The fully-coupled 4-stage mission converged quickly when the rigid blade solutions were used to determine the initial collective pitches for each mission stage.

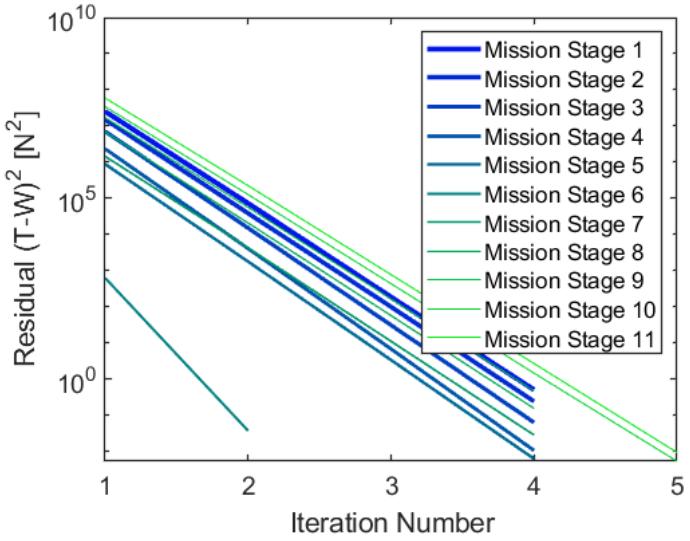

Figure 5.15: Each mission stage converged to a fully coupled aerostructural trim solution in between two and five FEA iterations.

actuation changes in twist are due solely to aerostructural deflection, while the addition of SMA actuation in this configuration leads to additional downward rotor blade rotation. Next, a series of mission stages under varying weights was run considering 0, 1, and 2% SMA actuation strain.

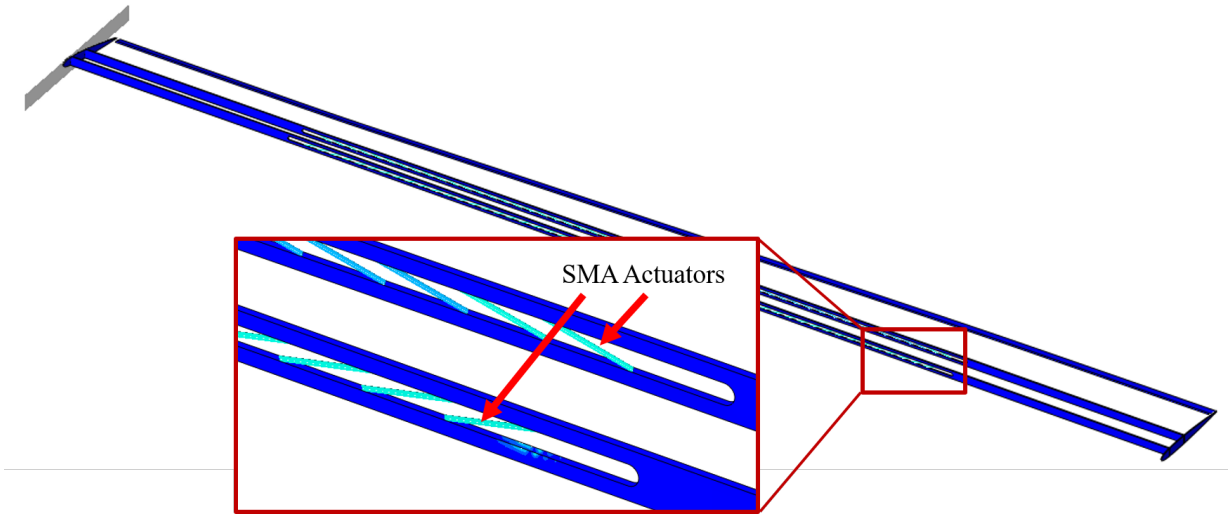

Figure 5.16: The structural model can consider active rotor blade twist via internal embedded SMA actuators.

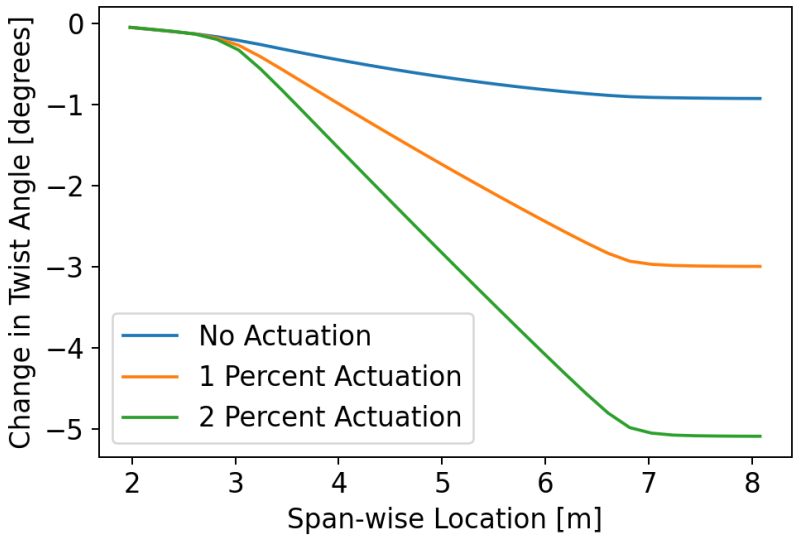

Figure 5.17: SMA actuation and aeroelastic deflection in the rotor blade alter the span-wise twist schedule.

At the 2% strain level, some higher loading fully coupled solutions were unable to converge, but the general trends, and the effects of morphing, can still be evaluated, as shown in Fig. 5.18. The overall trim power remained similar after considering rotor blade deflection due to structural

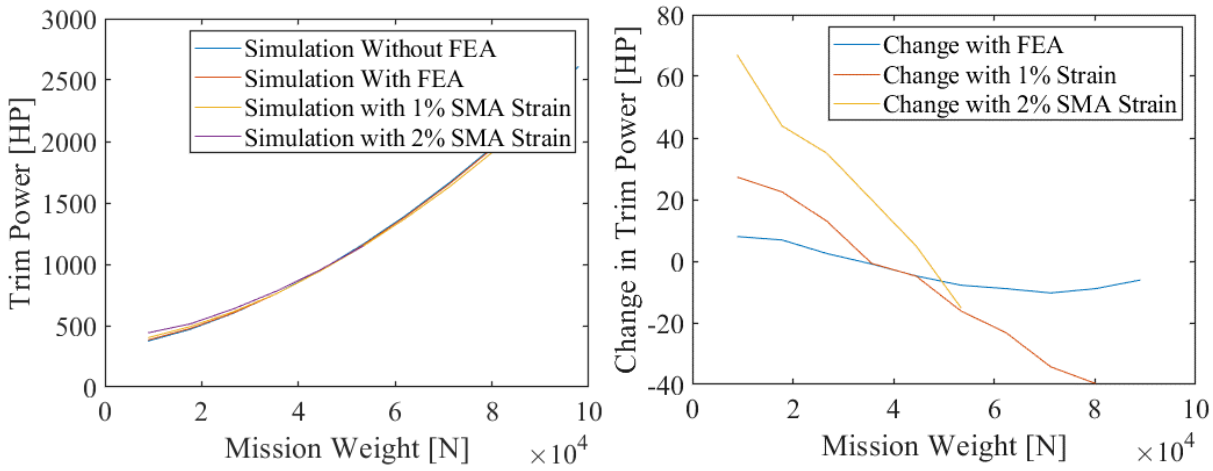

Figure 5.18: The changes in trim power across increasing mission weights were small, but noticeable.

loading. The addition of SMA-induced twist increased trim power at lower loading conditions, while decreasing trim power under higher loads.

5.4 Chapter 5 Summary

In this chapter high fidelity computational fluid dynamics parameterized camber morphing tables are integrated into the computational framework. A 14-stage air assault mission is introduced with five unique mission stages, and two optimization problem statements are compared: a mission score and multi-objective optimization for a single elliptical camber morphing region with a variable level of camber change. The starting and ending locations are treated as manufactured, or fixed, design variables, while the level of camber change can be morphed throughout the mission.

The multi-objective approach had less than 50% of the runtime required to determine trim flight for the mission due to less required evaluations. Both design and optimization methods led to improved performance metrics when compared to the reference rotor blade; the mission score approach increase final range time while reducing fuel burn, and the multi-objective optimization determined an adaptive design that reduced the required power for trim flight for each mission stage type.

Next, the multi-objective approach was extended to multiple missions with some shared mission stage types. Three mission were selected and adaptive rotor blade geometries were determined for each mission using the mission score optimization method, resulting in three mission-specific rotor blades. A mission-agnostic adaptive blade (MAAB) was then optimized using the multi-objective approach over all the uniques stages for all three missions. Then each adaptive rotor blade was run over each mission and the fuel savings compared. The mission score optimized rotor blades performed the best for their respective missions, and the mission-agnostic adaptive rotor blade performed nearly as well across all three missions.

An active twist rotor blade was next optimized for the 14 stage air assault mission to compare the design space decomposition method on a set of designs not guaranteed to be realizably adaptive and an optimization over all the adaptive designs variables for each objective, guaranteeing realizable final configurations. The presence of multiple adaptive elements meant the design space was much smaller using the design space decomposition. Additionally, both the NSGA-II and NSGA-AD selection criteria were utilized to selected the best generational designs. The best final designs were determined using the approximate local frontiers methodology from section 3.2.2, and the NSGA-AD selection algorithm slightly outperformed the NSGA-II selection algorithm. The chief benefits of active twist based on these multi-objective results suggest the reference rotor blade twist schedule was designed with a priority on forward flight, particularly flying at velocities with good maximum range. Meanwhile, the reference rotor blade performed worse in hover, with power required for trim flight. It stands to reason that, for the UH-60, forward flight and particularly light range was prioritized over hover, and most mission definitions required much more time in forward flight than in hover (particularly out of ground effect, which this work assumes).

The effects of the similarity parameter ϵ are also evaluated for this active twist optimization, with the difference between the best adaptive designs and the approximate utopia point for each objective compared to ϵ. As ϵ increases, the best adaptive design objectives approach the approximate utopia point, but the assumptions of similarity become less accurate. In this example over an ϵ value of 0.02, which represents of cluster of $3 * 10^{-7}\%$ of the fixed variable design space, the

objectives are largely converged.

Finally, aerostructural coupling is introduced into the computational framework. The nested aerodynamic and structural convergence iterative loops from section 2.4 are applied to hover stages to determine trim flight at varying loading conditions. Good convergence characteristics are observed before considering active twist, but the computational costs limited further exploration at this time. The changes due to aeroelastic deflections in rotor twist led to changes in trim power of under 10 HP, which is smaller than most of the changes due to morphing (changes in twist and camber reduced trim power by a maximum of approximately 180 HP). Additionally, active twist was introduced into the rotor blade via actuators embedded into the rotor blade spar, opening possibilities for simulating actuation mechanisms in finite element analysis as well, though any fully coupled optimization in this framework increases the runtime by nearly an order of magnitude when compared with the solely aerodynamic analysis.

6. FINAL ADAPTIVE ROTOR BLADE STRUCTURAL CONFIGURATION AND PROTOTYPING

Coupling aerodynamic and structural models requires significant computational costs. Structural and aerodynamic modeling at early stages in the design process are often performed separately to explore larger design spaces. For instance, if the changes in aerodynamic performance due to structural deflection are limited and much smaller than the effects of morphing and other design decisions, then a geometric configuration may be selected first based on the aerodynamic properties and then a more detailed structural configuration developed in later design stages.

Consider the five-stage MEDEVAC mission from section 4.3.2. An aerodynamic optimization was performed to optimize the rotor blade OML for three unique mission stage types. However, the method of morphing the OML under aerodynamic loading was not considered. Additional work is necessary to determine the structural configuration necessary to morphing between preferred geometries.

Chapter 6 Outline

This chapter is divided into three sections:

- In section 6.1, a structural shape-fitting optimization is performed to determine internal trailing-edge topologies for internal placement of actuators and structures to morph the rotor blade toward preferred shapes determined by this work.
- In section 6.2, a detailed design is optimized to match the high-fidelity camber morphing using materials and fabrication set out by Boeing for fabricating a prototype morphing rotor blade.
- In section 6.3, the morphing rotor blade 2-D section is fabricated and tested at Texas A&M to compare the results between aerodynamic optimization, structural optimization, and experimentally-tested prototype.

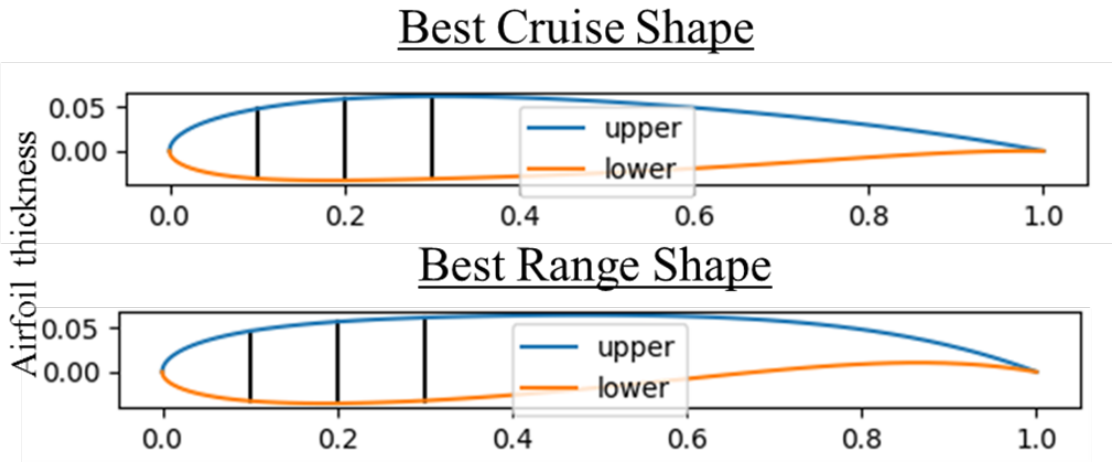

Figure 6.1: The least and most cambered shapes from the CST-driven OML optimization were selected for a structural optimization to morph between.

6.1 Adaptive Rotor Blade Modeling

A 2-D (or 2.5-D, as the model is 1 element wide) model was developed to determine the internal placement of actuators inside the rotor blade to morph between preferred shapes. The higher and lowest cambered results from the OML optimization were selected as the target shapes for morphing: the cruise and range configurations, as shown in Fig. 6.1. A shape optimization was defined where internal structure, including linear SMA actuators was added in the initial configuration. The initial configuration was defined as the Cruise OML shape, then the internal SMA actuators contract, altering the shape. Finally, the morphed shape's OML is compared with the target shape, the preferred shape for the range mission stage.

The internal topology was defined as a serious of construction commands based on the Lindenmayer system (L-system), a bio-inspired set of production rules originally used to describe the development of simple multi-cellular organisms [105, 16, 130]. This system has been implemented in engineering topological design in Python using the SPatial Interpreter for the Development of Reconfigurable Structures (SPIDRS). SPIDRS was developed as a spatial interpreter to apply additive topology construction instructions toward adaptive structural design and optimization [17].

This work was then applied to non-uniform boundaries such as airfoils [115], and the Python tools developed to apply SPIDRS as a construction tool for multi-material morphing rotor blades in Mikkelsen's work were applied in this work.

An uncoupled aerodynamic-then-structural optimization was performed to develop a morphing rotor blade mechanism based on results from the mission-driven optimization framework. The model applies aerodynamic pressures generated from the aerodynamic optimization for a rotor blade element near the center of the rotor blade. The leading edge of the rotor blade was not altered during the OML optimization, and is similarly not altered during the structural shape optimization; the trailing edge is the focus of the optimization.

Several materials were considered in the model based on Boeing fabrication capabilities and structural requirements. The outer mold line of the model would be wrapped in a carbon fiber weave around a 3-D printed internal structure. Internally, NiTiHf SMA actuators attach to the 3-D printed material via rods. Meanwhile, any other internal structure would be 3D printed, and thus rotation degrees of freedom would be restricted and tied to the external OML connections.

The procedure to tie all internal connections from SPIDRS-generated topologies can be seen in the flowchart in Fig. 6.2. The original SPIDRS instructions considered a single merged geometry, though different materials were assigned to different regions. Thus, considering SMA rotations required releasing rotation DOF connections while maintaining translation DOFs. This procedure was developed to turn SPIDRS instructions into ABAQUS degree of freedom (DOF) connections between SMA actuators and the internal plastic via a series of reference points and partitions.

Aerodynamic loading was applied to the rotor blade OML as shown in Fig. 6.3, based on the CFD pressure fields experienced by a specific rotor blade segment during the first mission stage under trim flight condition. The aerodynamic loading was calculated using the CFD pressure field applied to the rotor blade segment located 0.55 times the rotor radius down the rotor span under trim flight during the hover flight condition. The applied pressure field was selected from morphed rotor region experiencing the highest velocity field, as shown in Fig. 6.4.

A two-stage optimization was performed to match the downward trailing edge morphing target

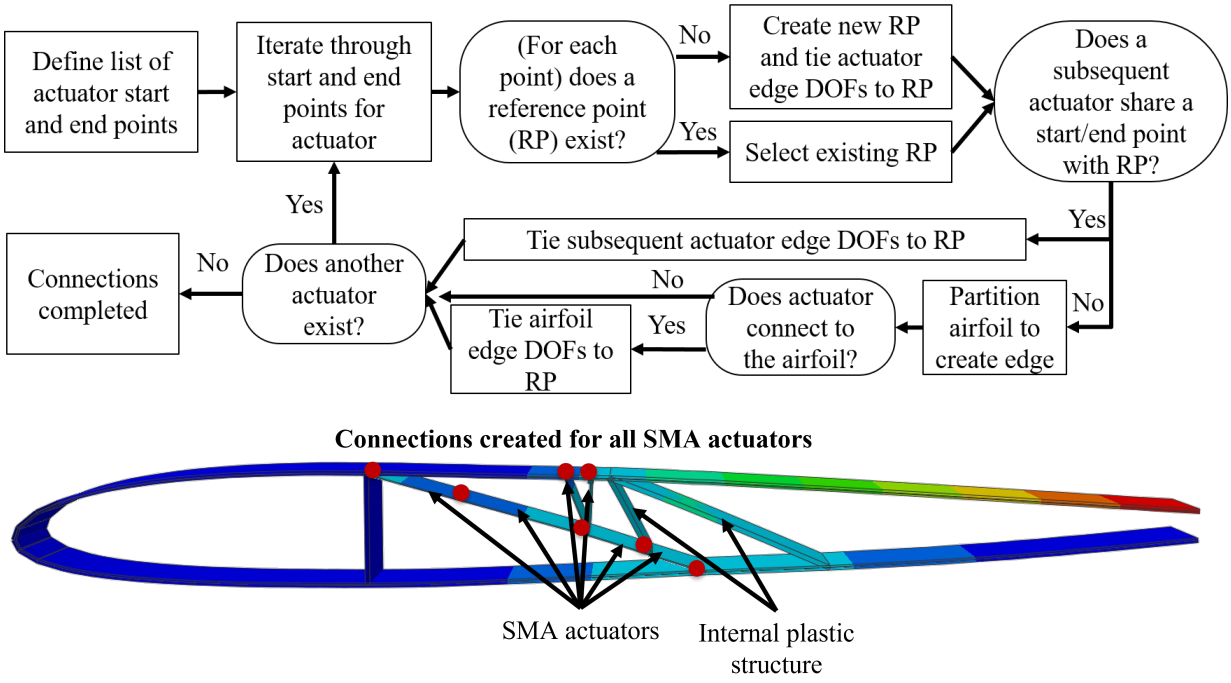

Figure 6.2: Reference points and degree of freedom constraints and connections are created for internal SMA actuators and internal plastic.

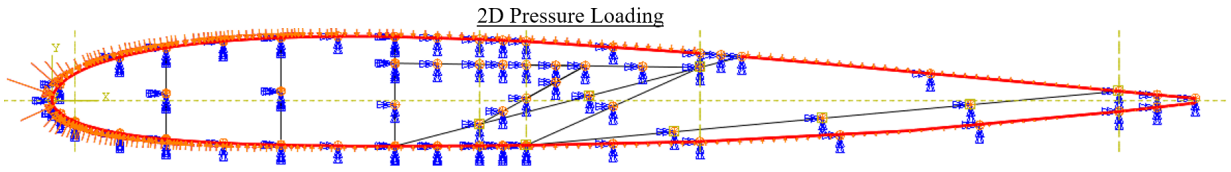

Figure 6.3: Example applied pressure loading over 2-D outer mold line based on mission pressure loading.

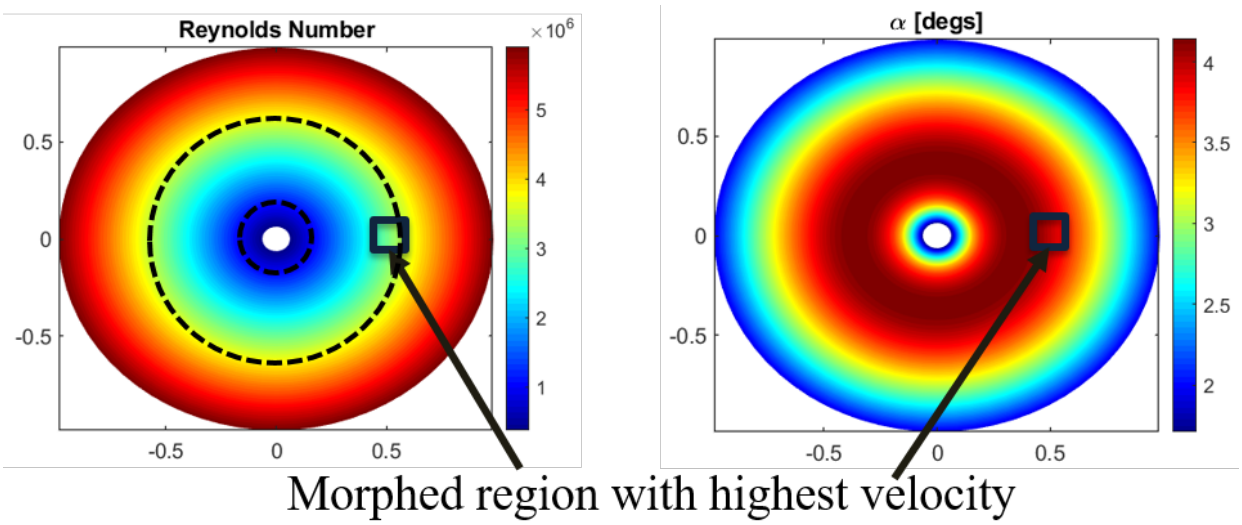

Figure 6.4: The pressure field for the highest velocity over the morphing region (between the dashed lines) of the rotor blade under the hover mission stage was applied to the structural model.

shapes while meeting structural constraints. The first optimization was performed using SPIDRS to define the internal topology and material assignment based on a series of genes that serve as creation instructions. Internal components inside the trailing edge of the rotor blade could be either 3-D printed plastic or linear actuators.

The leading edge of the rotor blade was fixed, then the SMA actuators were actuated incrementally to 1.5% actuation strain recovery. Contact constraints were applied to the outer mold line to ensure the trailing edges on the upper and lower surfaces could not pass through each other.

After the simulation was completed the OML shapes for every frame in the simulation were evaluated to determine the level of SMA actuation necessary to achieve the closest shape when compared to the best range OML shape from the aerodynamic optimization. A constraint was placed on the maximum stresses for each material to ensure no materials exceeded maximum yield stress. A minimum torsional stiffness constraint was applied to the model and measured after a twisting step, in which the blade segment was slightly twisted under a small applied moment. One out-of-plane edge of the rotor blade segment was fixed while the other was tied to an analytical plane defined by a reference point. A moment was applied to the reference point, and the rotation

Table 6.1: Internal Topology Objectives and Constraints

Objectives	Description
Error	Distance between morphed and target OML
Balanced Mass	Additional mass required to balance the rotor blade at the 1/4-chord
Constraints	
$\sigma_{max} < \sigma^y$	Maximum stress cannot exceed material yield strength
$GJ \geq GJ_{reference}$	Reference torsional stiffness must not exceed new torsional stiffness

was calculated to determine the torsional stiffness via the equation:

$$GJ = \frac{Ml}{\theta}, \tag{6.1}$$

where M is applied moment, l is out-of-plane thickness (1 element wide), and θ is the reference point rotation.

Two objectives were minimized by altering the SPIDRS inputs for constructing the internal topology: i) the 'error' between the morphed OML and the target shape, and ii) the weight necessary to balance the rotor blade at the quarter-chord. This second objective served to reduce unnecessary mass and structure in the rotor blade while also limiting changes in rotor weight balance. Constraints and objectives can be seen in Table 6.1.

A genetic optimization with a population of 100 designs was run for 250 generations to determine the best trade-offs between minimal additional mass and best shape fits, along with the best internal topology. This final topology is then used in more detailed designs for active trailing edge shape morphing. This initial optimization was performed due to the large design space of topology optimization. While aerodynamic optimizations in this work have at most 10 design variables, structural topology optimization could have thousands or hundreds of thousands of design variables.

The final Pareto frontier can be seen in Fig. 6.5. Most designs on the frontier represented relatively poor shape fits, but the best internal topologies represented shape fits very close to the

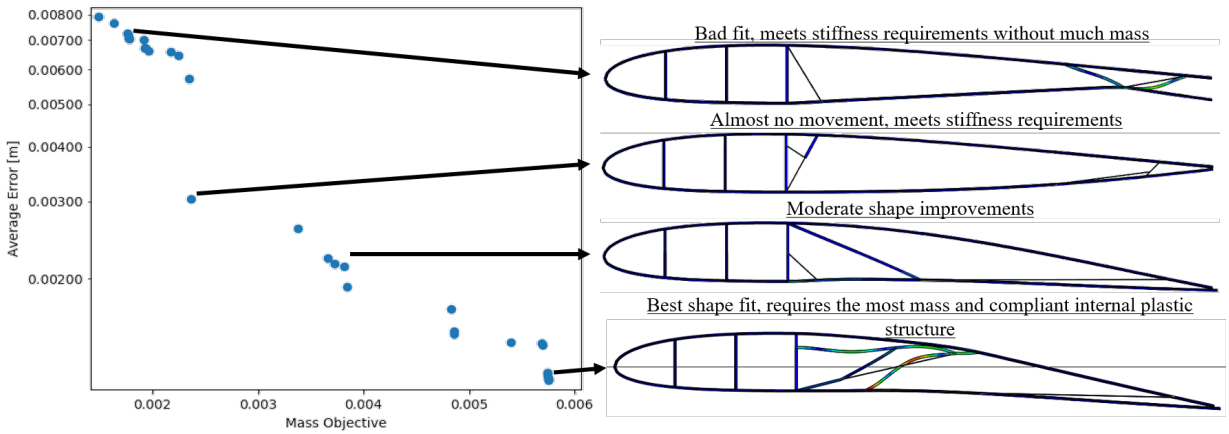

Figure 6.5: The final Pareto frontier for the rotor blade topology optimization.

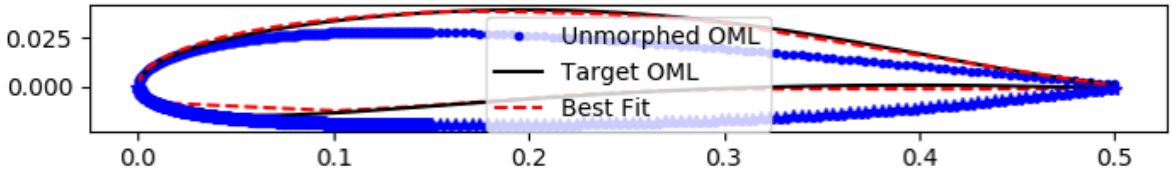

Figure 6.6: The optimized adaptive design altered the rotor blade outer mold line very closely to the target shape.

target shape, morphing the trailing edge of the rotor downward. Most designs violated constraints, either not torsionally rigid enough or generating too high of stresses. In the first generation only 11 designs satisfied constraints, which only improved to 22 designs by the 250th optimization. The best internal structure matched the target shape fairly well, as shown in Fig. 6.6

This optimization started with preferred shapes with low and higher cambered blades associated with better shapes for different mission stages based on an aerodynamic shape optimization. The best range shape had the highest cambered shape, while the high speed cruise shape had to least camber. A topology optimization was then performed to determine internal structural designs that could morph a rotor blade between the two cambered airfoil shapes while meeting structural constraints under aerodynamic loading. The final result was a set of designs with increasing shape

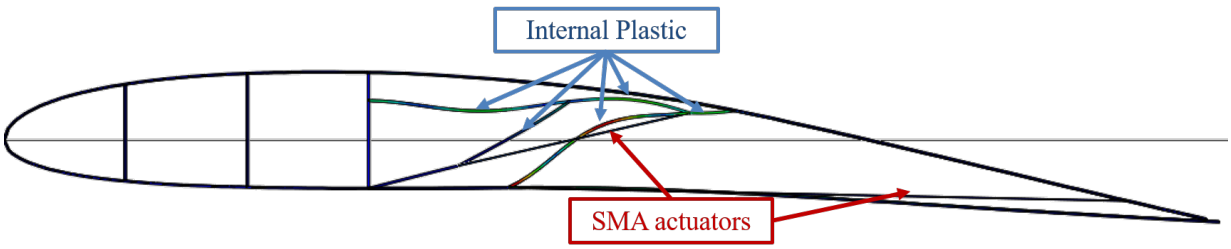

Figure 6.7: The final design selected matched the the target shape best, with two SMA actuators.

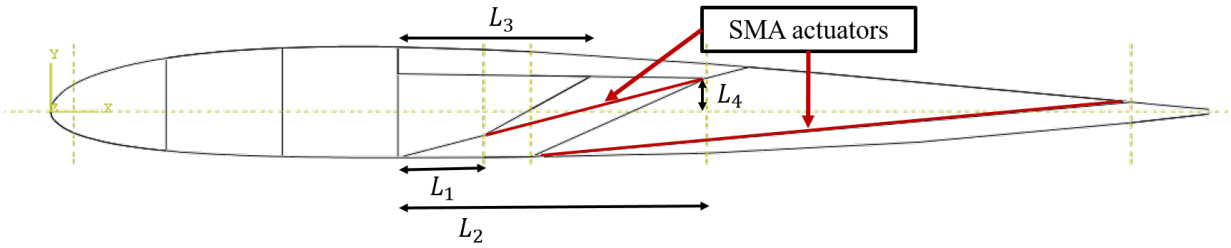

Figure 6.8: The internal topology was informed by the previous camber-morphing optimization and then parameterized.

fitting accuracy and internal mass. The selected design was the best shape fit, shown in Fig. 6.7. This design consisted of two SMA actuators attached to a linkage structure of internal plastic.

6.2 Detailed Design

A separate optimization was performed for a more detailed design for an adaptive rotor blade informed by the previous topology optimization. The internal topology could be parameterized to dramatically reduce the number of design variables while still attempting to match target trailing-edge camber morphing shapes.

A more detailed sizing optimization could be performed over a much smaller design space. The internal placement of connections between shape memory alloys and 3-D printed plastic were parameterized as shown in Fig. 6.8. These geometric parameters could alter the response of the morphing rotor blade without requiring another large topology optimization.

The OML consists of 3-D printed aluminum-filled nylon and a carbon-fiber layup. The carbon-

Table 6.2: Carbon Fiber Cured Material Properties

Material Property	Value	Units
Tensile Strength	521	MPa
Flexural Strength	777	MPa
Inter-laminar Shear Strength	64.7	MPa
Flexural Modulus	46.7	GPa
Tensile Modulus	55.1	GPa

Table 6.3: 3-D Printed Plastic Material Properties

Material Property	Value	Units
Tensile Strength	37	MPa
Tensile Modulus	3.7	GPa
Flexural Modulus	3.5	GPa

fiber layers are a weave, and were laid up in alternating 0, 45 degree orientations. The layup in was scripted in ABAQUS as shown in Fig. 6.9. The mechanical properties for the cured laminate are shown in Table 6.2 [173]. The material exhibits an extremely high flexural strength when compared with the 3-D printed nylon, whose properties can be seen in Table 6.3.

The rotor blade morphing design was optimized over 7 design variables, as shown in Table 6.4. 4 design variables define the parameterized internal topology, while the other 3 define material thicknesses: spar thickness, internal 3-D printed plastic thickness, and the number of composite layers of alternating orientation carbon-fiber weave.

Previous work considered the loading of SMA actuators as thin sheets passing all the way through a rotor blade span. The proposed prototype rotor blade actuation system developed and constructed at Boeing R&T expected to be used in the future, in contrast, would be induction-heated dog-bone thinner NiTiHf strips. Converting the 2-D adaptive model into a realizable design was done in close coordination with Boeing, particularly Michael Bass and James Mabe. The

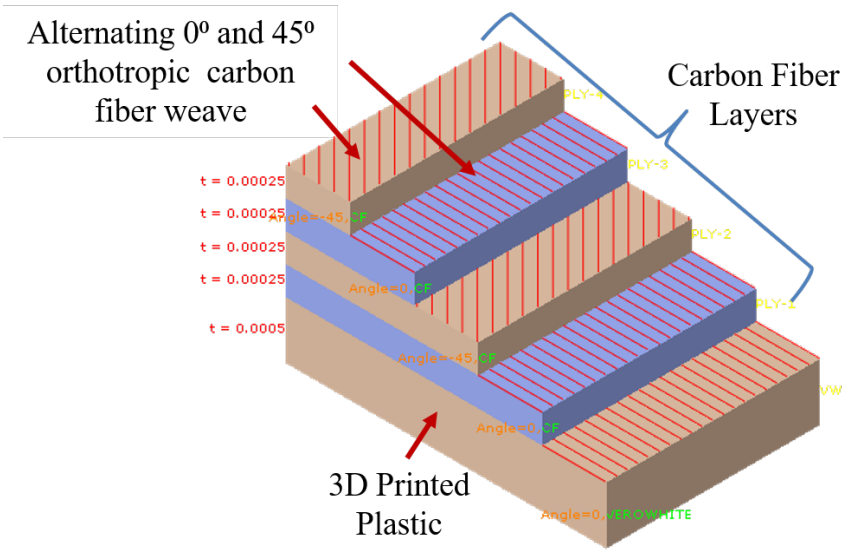

Figure 6.9: The rotor blade model OML consists of carbon-fiber weave layers or alternating 0 and 45 degree orientations overlaid on a layer of plastic.

Table 6.4: Internal Topology Objectives and Constraints

Design Variable	Description	Number of Values
L_i	Internal positioning parameters	4
t_{OML}	3-D printed plastic in OML	1
N_{layups}	Number of 0-45 degree sheet sets	1

material properties and fabrication methods were developed by Boeing, from carbon fiber to 3-D printed aluminum-reinforced plastic. As such, these actuators would be placed at periodic intervals down the rotor blade. Increasing rotor blade stiffness and required actuator loading could then be reduced by increasing the number of actuators while decreasing their spacing. Likewise, decreasing the maximum loading for actuators on a 2-D (2.5-D) model decreases the number of actuators required and increases their spacing.

The 6 design variables were optimized to minimize two objectives in a large genetic optimization. The stress constraint on SMA actuators is replaced with an objective of minimizing the

Table 6.5: Detailed Design OML Optimization Objectives and Constraints

Objectives	**Description**
Error	Distance between morphed and target OML
$\left[\frac{F}{L}\right]_{SMA}$	Required SMA force per unit length down the rotor blade span
Constraints	
$\sigma_{max} < \sigma^y$	Maximum stress cannot exceed material yield strength
$GJ \geq GJ_{reference}$	Reference rotor blade torsional stiffness may not exceed rotor blade torsional stiffness

maximum loading-per-unit-span required for each SMA actuator. Additionally, the error between target OML shapes and the closest fit OML for a given 1-way actuation was minimized as the other objective. A constraint was applied on all designs such that the stress on each component could not exceed the maximum material yield strength after applying a safety factor. The objectives and constraints can be seen in Fig. 6.5.

An optimization was performed with a population of 24 individuals over 50 generations to minimize both the shape fitting error and the maximum SMA actuator loading (per unit depth). The Pareto-optimal objective space, shown in Fig. 6.10, shows a trade-off between shape fitting errors and maximum SMA loading. A final configuration was selected from the set of Pareto-optimal designs. The design with minimal required SMA loading that had an average mean squared shape error near 0.02 mm was selected for the prototype design. The internal configuration (location of structure and SMA actuators) can be found in Appendix B.

6.3 Fabrication and Testing

The adaptive rotor blade components (actuators, 3-D printed plastic, leading edge, and carbon fiber shell) were fabricated by Boeing R&T in St. Louis using the materials previously described. The components for were shipped to College Station then fabricated and tested with the help of James Mabe and Michael Bass. The SMA linear actuators were wrapped with Kapton tape [44] then magnet wire and heated through induction up to 250 C; the temperature was monitored via

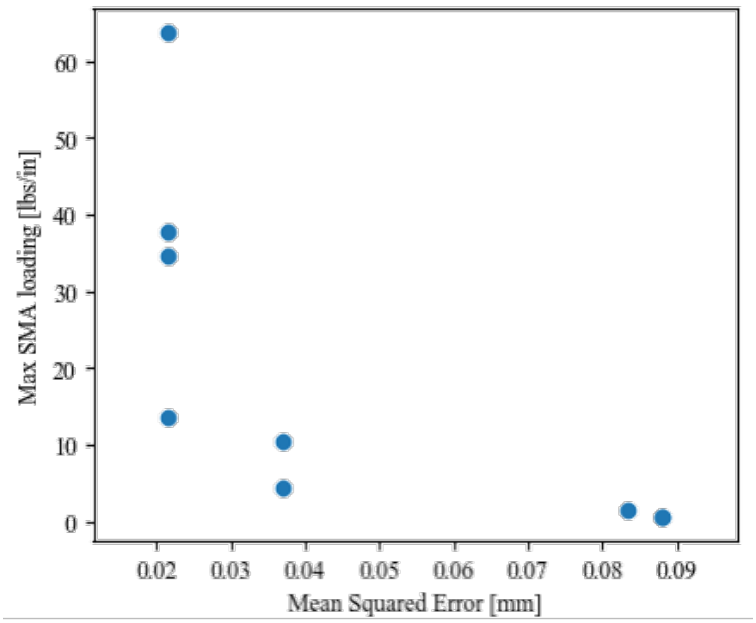

Figure 6.10: An optimization was performed to determine the best designs with respect to shape fitting accuracy and required SMA actuation force.

thermocouples placed on the actuator centers underneath the tape.

The morphing rotor blade assembly can be described as a combination of 3 sections: the leading edge assembly, the trailing edge internal structure, and two carbon fiber outer mold line pieces, one for the top and bottom. The internal trailing edge structure was assembled by sliding the internal components (3-D printed plastic, SMA actuators, connection rods, and wiring) onto rods as shown in Fig. 6.11. 3-D printed plastic and SMA actuators were slid onto each pin to distribute the load from the SMA actuators to the rest of the structure. It is worth noting that some actuation and camber was introduced into the internal structure when installing the actuators for the first time.

First, the internal system was tested by actuating the two large SMA actuators. Due to limitations in power supplies and monitoring equipment, only two SMA actuators were actuated at a time. In this respect it will be shown the two short forward-placed actuators contribute very little towards the overall shape. The actuators were heated to 200 C and the deflection near the trailing edge (along the trailing edge connection rod) with a dial gauge indicator, as shown in Fig. 6.12. A

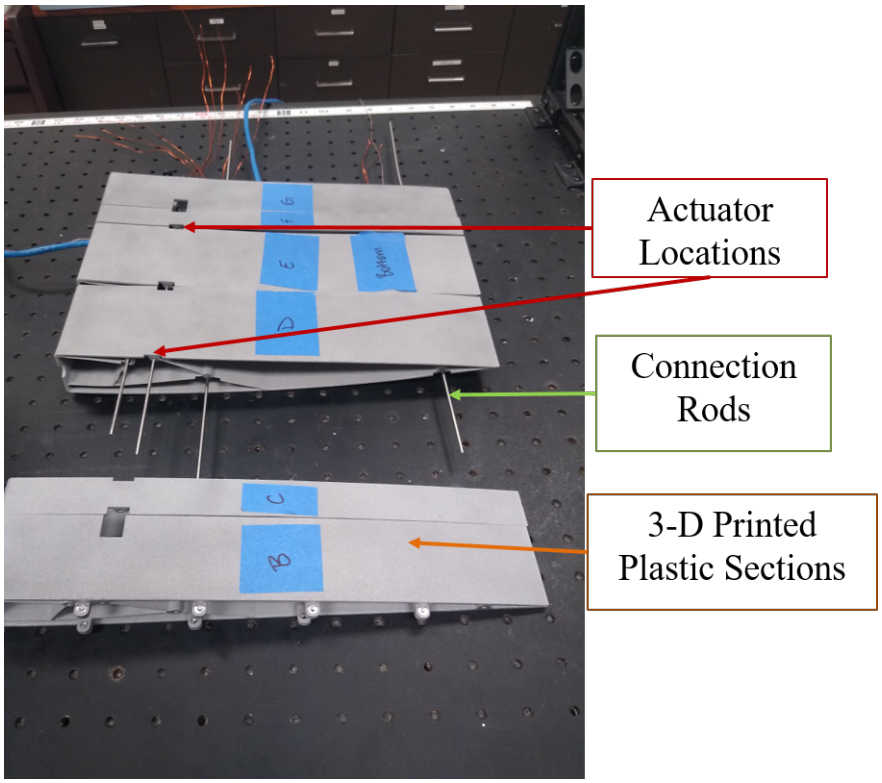

Figure 6.11: The trailing edge internal structure is assembled in the out-of-plane direction via connection rods.

maximum deflection was measured after full actuation as 0.422".

The internal structure was then mounted to the carbon fiber skin via double-sided tape for this work; an adhesive such as epoxy could be used in the future if a more permanent bond is desired. The carbon fiber OML is connected to the leading edge assembly via screws, while the leading edge is mounted to the optical table surface. The complete morphing rotor blade assembly is shown in Fig. 6.13, with a fixed and adaptive trailing edge sections. The only regions of note during the current experiments are the adaptive trailing edge region.

Next, the large actuators were cycled and the trailing edge deflections measured via a dial gauge indicator and a laser deflection near the trailing edge of the rotor blade (within an inch). Both SMA actuators were heated to over 200 C then cooled down to under 30 C. This procedure was repeated again, with the deflections recorded alongside a video of the actuation. The third and

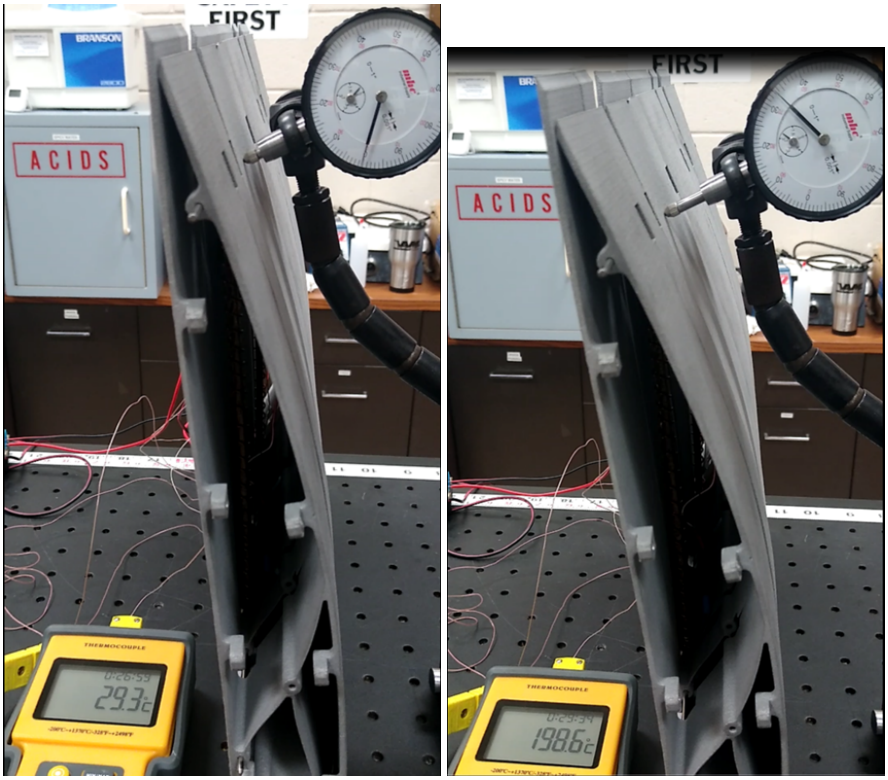

Figure 6.12: The SMA actuator was heated to 200 C, morphing the internal structure to increase the camber and deflection.

final actuation cycle increased the final temperature from 200 to 250 C, and also placed the dial gauge very close to the trailing edge.

The maximum deflections as measured via the dial gauge are shown in Table 6.6. The largest deflection was measured during the third cycle, likely due to increased maximum temperature and measuring with the dial gauge located closer to the trailing edge where maximum deflection is expected. The thermocouple readings were unable to post valid readings while actuation was underway, so at times during heating the induction circuits were briefly turned off to check actuator temperatures.

The difference in trailing edge camber is shown before and after actuation from a top view of the rotor blade in Fig. 6.14. For reference, the table mounting holes below are spaced approximately 1" (exactly 25 mm) apart. The trailing edge camber increases due to the strain recovery

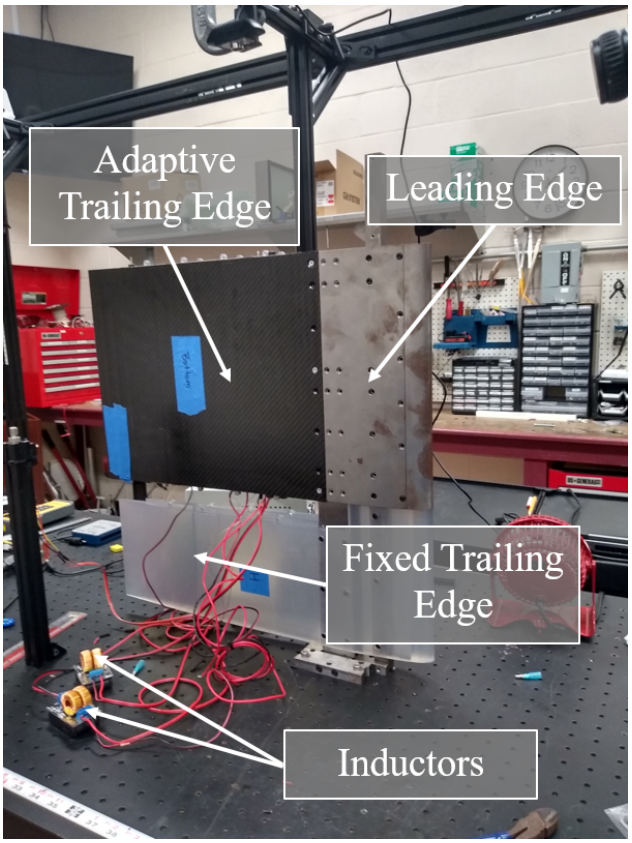

Figure 6.13: The rotor blade bench-top test stand in mounted vertically and consists of fixed and adaptive trailing edge regions.

Table 6.6: Rotor Blade Trailing Edge Deflections

Test Number	Maximum Deflection ["]
1	0.53
2	0.55
3	0.69

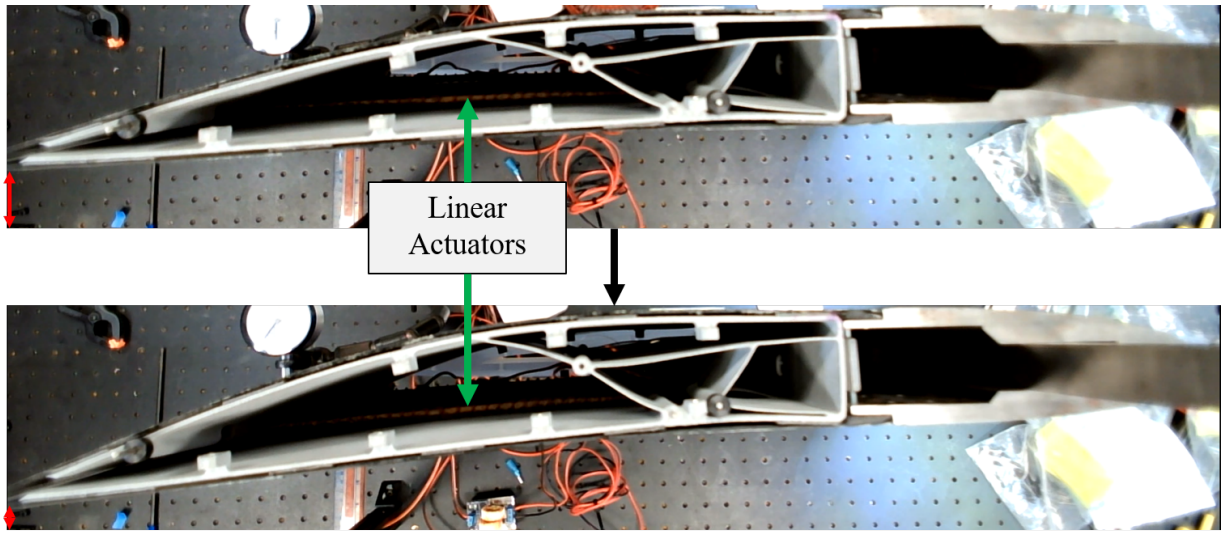

Figure 6.14: The trailing edge is morphed down due to the internal linear actuators.

from the internal linear actuators.

Next, the outer mold lines were extracted from the top view of the 3rd and final thermal cycle. The final OML was extracted via webplotdigitizer [136], with the images rotated based on the grid of mounting holes on the optical table surface. The morphed OML after full actuation and after a measured 0.5" actuation was compared with the OMLs from the reference rotor blade, aerodynamically-optimized rotor blade, and the structurally-optimized rotor blade. The OML's can be seen in Fig. 6.15.

The resulting OML comparison suggested that the SMA actuators were actuating beyond the maximum required actuation. This could be due to a higher total strain recovery than the FEA model considered, as the total strain recovery was limited in simulations to 1.5% as a conservative estimate. The trailing edge OML was also wider in the optical table testing, as least partly due to the use of tape with a notable thickness to attach the carbon fiber skin instead of a solid attachment as modeled.

Next the two shorter, forward actuators were thermally cycled, with the changes in trailing edge deflection monitored using the dial gauge. Only a change of 0.015" was measured between

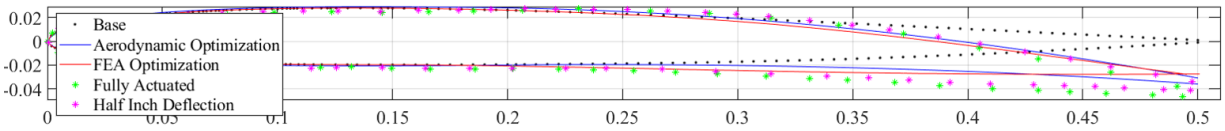

Figure 6.15: The outer mold line (OML) under full SMA actuation led to larger trailing edge deflections than required to match the target OMLs from the aerodynamic and structural optimizations under maximum deflection.

the initial and maximum trailing edge deflection, as compared with the maximum trailing edge deflection of 0.690" when the two large trailing edge actuators were heat cycled.

That the short forward actuators provide limited deflection on their own begs the question of why a very large topological optimization (a population of 100 for 1,000 generations) would select a design with unnecessary mass. The initial topology optimization discouraged additional mass based on the distance from the rotor blade quarter-chord. The forward actuators slightly improved the trailing edge camber morphing while providing structural support close to the quarter-chord, which is less penalized in the optimization than structure closer to the trailing edge. Thus the forward actuators were included in the final design and the design remained Pareto-optimal.

6.4 Chapter 6 Summary

In this chapter a number of design, optimization, and simulation tools were utilized to develop a 2-D camber morphing rotor blade. The internal topology was determined through a structural design and optimization study based on a preliminary aerodynamic optimization. The topology was then further refined using higher-fidelity modeling and a sizing optimization using the initial topology results to reduce the design space. The rotor blade was tested to compare simulated and experimental outer mold lines.

Through this process, the novel design space decomposition technique determines the placement of adaptive structures. Then an uncoupled structural optimization can determine the internal topology required to match the preferred aerodynamic shapes. Finally, a targeted sizing optimization results in a structural solution matching the ranges of trailing edge morphing aerodynamically

optimized to minimize trim power over competing individual mission stages.

The final detailed design was sent to Boeing for fabrication, then the rotor blade section was assembled and tested at Texas A&M University by the author, James Mabe, and Michael Bass. Comparable trailing edge deflections were observed between the aerodynamically-optimized, structurally optimized, and experimentally tested airfoils. This chapter details a complete design and optimization process and demonstrates how the computational tools in this work can be used in the design process; given a mission and rotorcraft, a complete detailed design for adaptive rotor blade components can be realized through a series of design and optimization procedures using the computational tools developed in this work.

7. SUMMARY, CONTRIBUTIONS, AND FUTURE WORK

Mission-adaptive aircraft offer potential performance improvements via geometric alterations between mission stages. These adaptations can be considered during the earliest stages of design, but often the computational costs are high due to a number of contributors. In early design stages the design space can be very large; each additional parameter increases the dimensionality. When considering adaptivity for missions, the performance during different mission stages must be considered, increasing the number of aerodynamic trim flight solutions. Additionally, high-fidelity aerodynamic and structural solutions often necessitate high computational costs and run times.

A large body of work has explored adaptive rotorcraft technologies, but relating these changes at the local rotor blade level to the structural design and placement necessary for improved mission performance is still a challenge. Considering the effects of adaptive structures on mission performance in early design stages can result in a very large design space. For example, consider the rotor blade twist schedule for a 5-stage mission. A static design may have 1 twist schedule, but a variable twist rotor blade may consider 5 schedules throughout the mission. This design can either be evaluated via a single mission score (in which the 5 twist schedules are evaluated as a single final performance metric) or via a multi-objective optimization (in which the twist schedule is evaluated over all 5 mission stages individually).

However, adaptive designs must also be realizable across objectives; each design must be able to morph from one configuration to another. Design and optimization methods must consider adaptive element placement to ensure feasibility. In this manner, there is a fundamental difference between design variables that can and cannot change during a mission. Consider the feasibility of the previously mentioned adaptive rotor blade twist schedule. Suppose the twist is altered via 3 torque tubes; a design and optimization problem could determine the 3 lengths and 3 twist rates resulting in the best objectives. For a 5-stage mission, the final design requires 3 lengths for the entire mission and 3 actuator twist rates for each mission stage, resulting in 3 ‘fixed’ and 15 ‘adaptive’ design variables, where ‘fixed’ means the design cannot change during the mission, as

opposed to ‘adaptive’. The single mission score approach would consider 18 design variables, while a multi-objective optimization approach could consider 6 design variables over 5 objectives, one for each mission stage.

A computational framework was developed for efficient mission-driven design and optimization of adaptive aero-structures, specifically tailored in this work toward rotorcraft. The focus considers both the computational costs to relate adaptive structures to mission-driven performance and the fundamental differences between ‘fixed’ and ‘adaptive’ design variables in the earliest stages of design. A wide range of rotorcraft aerodynamic and structural numerical methods were considered to reduce computational costs while maintaining accurate results. Rotor blade chord, twist, and shape morphing are evaluated along with structural considerations, from placement to full aero-structural coupling.

A general, modular suite of design and optimization tools is coupled with a range of adaptive rotorcraft methods and a unique design space decomposition methodology to consider which designs can morph to other similar designs to span the multi-objective space, ensuring feasibility and morphing benefits. Both multi-objective post-processing techniques and genetic algorithm selection criteria are developed and evaluated toward designing mission-driven adaptive structures via multi-objective optimizations, with good performance compared with traditional adaptive structures approaches.

7.1 Summary and Conclusions by Chapter

7.1.1 Mission-Driven Adaptive Rotorcraft Modeling

In Chapter 2 a range of adaptive rotorcraft modeling methods are outlined to cover a range of adaptive rotorcraft technologies across a range of aerodynamic conditions. At the local rotor blade level, changes in geometries are related to alterations in aerodynamic properties. These changes in geometries range from adaptive chord to variable twist and outer mold lines. Various fidelity tools are considered alongside their respective computational costs, enabling early-stage explorations without requiring cost-prohibitive numerical solutions during initial design optimizations. An ex-

isting comprehensive rotorcraft analysis tool and a BEMT code developed by the author for this computational framework are coupled with a modular set of interpolation tables to relate geometric and aerodynamic properties. First low-fidelity tools generate these aerodynamic interpolation tables, but pre-computed CFD tables are later integrated into the framework resulting in improved accuracy when compared with existing computational and experimental data.

A range of adaptive rotorcraft methods were explored for both structural consideration but primarily aerodynamic optimization. Changes in chord, twist, and outer mold line were evaluated first at the local rotor blade segment level, with discrete blade elements assigned 2-D aerodynamic properties as a function of local geometry, angle of attack, and Mach number. Altering the rotor blade twist schedule alters the local angle of attack experienced down the rotor span. Morphing the rotor blade with respect to outer mold line alters the relation between angle of attack and aerodynamic properties such as lift, drag, and moments in more complex manners, requiring aerodynamics solvers such as panel methods or computational fluid dynamics in the absence of pre-computed solutions. Once the changes in local properties are related to lift, drag, and moments, all main rotor forces are resolved around the main hub and must be balanced with the addition of other rotorcraft forces, such as vehicle weight and tail rotor contributions.

Two methods of rotorcraft modeling are implemented: a BEMT code for hover and a comprehensive vehicle analysis code. Both methods solve for the necessary pilot inputs for trim flight iteratively. The hover BEMT code determines collective pitch and considers nonlinear, axisymmetric inflow for hover and vertical climb. This code was developed for this work, while the comprehensive vehicle analysis code TRAC was taken from existing work not developed by this author but selected for its quick convergence and numerical validation data for the UH-60. Together, these two codes consider flight during both hover and forward flight.

Aerostructural modeling is a valuable but computationally expensive simulation challenge due to the coupling between aerodynamic and structural responses. In this work a fully-coupled aerostructural model is integrated into the aerodynamic simulation for hover and climb mission stages by leveraging the convergence iterations required for both the aerodynamic trim flight and

coupling iterations. The aerodynamic trim flight solutions require over 80 iterations to converge, while rotor blade twist deflections due to aerostructural loading typically only require 4 iterations for convergence. This is due to the very large inertial forces applied to the rotor blade that dominate the structural deflection. Additionally, the structural analysis takes 10-20 minutes per evaluation, orders of magnitude larger than a single aerodynamic evaluation, or even a trim flight solution for hover. These times are still much faster than other fully-coupled solutions that consider a full finite element analysis solver and computational fluid dynamics results. All of these tools are integrated into a set of aerostructural design and optimization tools that relate individual adaptive structures into mission performance metrics by solving for trim flight under most flight conditions using modeling techniques with reasonable, tested accuracy and efficient computational costs.

7.1.2 Design Space Decomposition

In Chapter 3 adaptive design and optimization is explored; the design space is decomposed into two sets of design variables: those that can and cannot change throughout a mission (or change based on different objectives). Two designs with the same fixed but different adaptive design variables can both be reached when different objectives are desired. This is especially valuable during a mission when even the same type of objective (such as power required for trim flight) may require a different configuration due to changes in other mission parameters (such as forward flight speed, numbers of passengers, etc.). A methodology is developed to optimize both fixed and adaptive design variables concurrently while considering feasibility (i.e., that the final design consists of design that can morph between each other for each mission stage/objective).

Consider a set of designs that can be reached by only altering adaptive variables in the design space. Two numerical measures are developed to evaluate the performance of any such adaptive set, keeping in mind that both adaptivity and elite performance is the objective; a large set of designs with poor objectives may be worse than a single design. The concept of a 'utopia point' is utilized to compare the approximate utopia point from a set of designs in the objective space to the approximate utopia point from a set of similar designs. In this manner only the best designs with respect to each objective (or mission stage) are considered. When the goal is continuous morphing

between stages, a distance metric is employed whereby continuous adaptivity is considered, with a priority on designs close to the approximate Pareto frontier.

Next, optimization methods are developed to determine designs that span the Pareto frontier in accordance with the adaptive multi-objective performance metrics previously described. A two-stage method optimizes all fixed and adaptive design variables to improve multi-objective mission performance. The second stage then searches through all designs to find designs that can feasibly morph between each other due to similar fixed design variables. Feasible morphing is based on a similarity parameter normalized by the design space, because in a genetic optimization each generation should generate different design variables due to mutations and genetic crossover; no designs would have identical fixed design variables.

Another method considers adaptivity during the optimization itself, encouraging designs with better adaptive performance. Elitism and rank are still considered, but designs of the same rank are selected based on the best adaptive performance metrics. This method leads to better adaptive designs for each generation in the analytical demonstrations developed, resulting in larger 'families' of similar designs but a less evenly-spread Pareto frontier. When compared with increasing the design space to consider each adaptive design variable for each objective, dimensionality is a major factor. For a two-shape Euler-Bernoulli beam optimization increasing the design space from 4 design variables to 6 (and the adaptive design space from 2 design variables to 4) led to better convergence. However, increasing the number of desired shapes from 2 to 4 resulted in the adaptive design techniques resulting in improved convergence and final designs.

The approximate local frontiers technique was also evaluated as a general selection criteria for any design exploration. The design space for an Euler-Bernoulli beam and an analytical minimization function were explored using Latin Hypercube Sampling (LHS), then the best 'adaptive' designs were determined. A final optimization could then be performed with the selected 'fixed' design variables to ensure the design is feasible and near-optimal.

Of course, the design space decomposition method is not the only valid design and optimization technique for adaptive multi-objective design and optimization, but it performs well across the

problems presented in this work and follows the established principles of rank and non-dominance often used in multi-objective design and optimization. This approach also outperforms existing adaptive design and optimization methods for larger sets of objectives by reducing the design space dimensionality during the optimization.

Shape fitting Euler-Bernoulli beams provided a comparison between existing adaptive optimization techniques and the methodologies developed in this work. The non-dominated genetic sorting algorithm for adaptive design (NSGA-AD) outperformed existing selection algorithms for determining inter-generational adaptive designs (determined via the approximate local frontiers approach in which similar designs with respect to fixed design variables are considered adaptive between each other). When compared with the traditional approach of considering every adaptive design variable for each objective as a separate dimension in the design space, the methods provided in this work performed worse for a population of 100 individuals when matching two shapes but performed much better when considering 40 and 100 individuals to match 4 shapes.

When the design space is small enough in comparison with the optimization population that considering every adaptive design variable for every objective is sufficient, the design space decomposition is less effective. However, if increasing the number of objectives from two to four shows the value of the design space decomposition for an analytical design, the applications should translate well to mission-driven optimizations, since most missions consider more than just two objectives or mission stages.

7.1.3 Computational Framework Development and Mission-Driven Adaptive Rotorcraft Design and Optimization

Continuing from Chapter 2 before applying the methods described in Chapter 3, Chapter 4 first develops a robust computational framework to relate adaptive rotorcraft elements to mission performance such that general mission-based optimizations can select preferred geometries and their placement in the rotor blade. A modular scheme starts by defining the mission, then the optimization variables describing the vehicle geometries for each mission stage. Next actuation systems are sized and aerodynamic properties assigned. Geometries are related to local aerodynamic per-

formance with various options, from experimental data to CFD interpolation tables, depending on the parameterization and desired fidelity. The rotor blade mass distribution is determined and trim flight calculated for each mission stage.

Each mission stage is defined in a trim flight state that must be determined iteratively. The required trim flight power is related to fuel burn via specific fuel consumption. Hover and cruise have specified velocities, while maximum range and endurance are determined by evaluating a range of forward velocities and selecting the best one for each objective. The velocity for maximum endurance is determined by the minimum required power for trim flight, while the velocity for maximum range is determined by the ratio of trim power to forward velocity: an increased forward flight speed with the same trim power yields a longer range.

This mission evaluation is either sequential, resulting in an overall mission score, or multi-objective, resulting in a series of trade-offs between competing objectives for each mission stage. The multi-objective approach can optimize over the entire design space, considering both fixed and adaptive design variables by post-processing the entire optimization using the design space decomposition approach developed in this work. The best set of feasible designs where adaptive structures can morph between the preferred design for each objective are determined via the design space decomposition method.

Next, a wide range of adaptive rotorcraft technologies are evaluated via parametric optimizations. The optimization framework is first demonstrated for optimizing the chord distribution over a series of four mission stages without structural feasibility constraints. The resulting preferred geometric shapes showed similarities for each objective, with a taper starting in a similar location and increasing down the rotor blade span. This initial optimization served to test the optimization framework before considering feasibility, though the final results suggest a hinge located approximately 30% down the rotor blade span that could alter the trailing edge taper.

A more detailed structural approximation for embedded SMA torque tubes was developed determining the sizing, placement and levels of twist to morph the rotor blade over a mission with four unique stages. The torque tubes alter the twist schedule to the rotor blade, changing the local

angle of attack with respect to the incoming free stream velocity and rotor blade root, which is set via pilot pitch inputs. Five torque tubes were placed inside the rotor blade for an initial optimization over all the design variables, both fixed (torque tube lengths) and adaptive (active twist rates). Next, the approximate local frontiers approach from section 3.2.2 determined feasible sets of 'similar' designs with respect to the fixed design variables, and the best set was selected using the Utopia Point performance metric from section 3.1.1. A final optimization was performed over the adaptive design variables to ensure feasibility and select a final, fully realizable adaptive design. This methodology demonstrated novel design space decomposition approach toward determining adaptive designs (and preferred fixed design variables) via similarity for multi-objective engineering design problems, in this case targeted toward mission-driven adaptive rotorcraft.

Adaptive outer mold lines (OMLs) using class shape transformations (CSTs) were then evaluated by relating alterations in airfoil shapes to changes in lift, drag, and moment profiles as a function of angle of attack and Mach number using panel method aerodynamic tools. An initial multi-objective mission-driven aerodynamic analysis revealed that structural constraints were necessary, as the 'best' shapes were much thinner airfoils than the reference rotor blade; this was likely due to structural constraints not considered in the purely aerodynamic optimization. Geometric constraints were applied to the leading edge of the rotor blade for structural consistency such that trailing edge morphing could be evaluated without altering the leading edge and spar.

Next, mission-driven sequential analysis was introduced, where each mission stage is evaluated sequentially to determine mission-wide scores, such as total fuel burn and final range. Simple missions are introduced along with the advantages of morphing between mission stages for a sequentially-defined mission. This mission score provides an alternative, but less general approach to mission-driven design and optimization when compared with the multi-objective approach. The multi-objective approach is also parallelizable, and requires less trim evaluations for mission where stage types are repeated (for example, hover under the same flight conditions for stages 1 and 4 in a mission will yield the same trim flight solution under the multi-objective approach).

The initial fuel for a given mission is also investigated. Due to the weight, and changes in

weight, of fuel there is a coupling between initial fuel, trim power, and fuel burn. Increased performance in the forms of lower power and fuel burn due to lower weight conflict performance in the forms of endurance and range. In this work the initial fuel for a given mission was determined as the fuel required to meet a common final range requirement of 30 minutes for the reference rotorcraft configuration.

Next a MEDEVAC (medical evacuation) mission was defined and optimized over the full rotor blade. The computational framework determined the best shapes for hover, range, and cruise resulting in increased final range time at the end of the mission and reduced fuel burn before the final stage. These optimizations all determined local aerodynamic properties via either existing experimental data (when the airfoil shape was not altered) or low fidelity panel methods when characterizing new airfoil shapes.

In this chapter, a computational framework is developed to design mission-driven adaptive structures over a range of mission stages and adaptive rotorcraft technologies. Two approaches to mission-based adaptive design are explored: the multi-objective and mission-score approach. The multi-objective method can run mission stages in parallel, but does not consider coupling between mission stages. This approach can also utilize the design space decomposition methodology to determine families of similar designs, denoting potential feasible morphing to span the objective space.

7.1.4 High-Fidelity Mission-Driven Adaptive Rotor Blade Technologies

Chapter 5 continues the mission-driven adaptive rotorcraft studies from the previous chapter with additional physical considerations beyond panel-method aerodynamics and analytical structural solutions. While useful for large design and optimization problems, particularly during early stages of design, it is also advantageous to consider higher fidelity optimization tools when possible.

Parameterized camber morphing airfoil tables generated via fluid dynamics were integrated into the computational framework. The tables could related single-parameter trailing morphing to lift, drag, and moments based on angle of attack and Mach number, and are also used in other

optimizations with the camber morphing parameter set to the base position. Initial comparisons show that higher cambered airfoils performed better for low lift missions on the outer rotor blade region, a similar trend to SABRE camber morphing studies [132].

A more comprehensive 14-stage Air Assault mission is presented with 5 unique mission stage types based on weight and velocity. A design and optimization problem is developed where the trailing edge camber is morphed elliptically over one region of the rotor blade, with two location variables and an adaptive maximum camber variable. The elliptical camber region placement and levels of morphing for each mission stage were determined via two optimization techniques: the multi-objective approach and the mission score approach.

The mission score approach required over twice as much time per evaluation to sequentially evaluate all 14 mission stages, while the multi-objective approach only considered the 5 unique stage types independently. Both approaches found performance increases; the mission score method improved the final range time while reducing fuel burn while the multi-objective approach reduced trim power across all mission stage types. When the mission score design was evaluated for each objective individually, the adaptive design was better, or at least Pareto-optimal, when compared with any individual designs from the multi-objective optimization. However, applying the design space decomposition post-processing to the multi-objective optimization final population resulted in a design with better multi-objective performance than the final mission score design, though it is worth noting the multi-objective performance does not consider fuel burn.

The multi-objective design space decomposition approach was then extended across multiple missions. Three missions were selected with some similar mission stage types, and both the mission score and multi-objective optimization methods were applied. The mission score approach resulted in 3 different rotor blades, each tailored for the respective mission. Meanwhile, the mission-agnostic adaptive blade (MAAB) considered ever mission stage. The MAAB required much less runtime to optimize and performed well across all 3 missions, while each mission score blade performed the best for its respective mission but worse than the MAAB for the other 2 missions.

It is worth noting that the design space for this comparison was fairly small, with one adaptive

(over 5 objectives) and two fixed design variables. In contrast, consider the five torque tube design and optimization problem statement. The multi-objective approach considered 5 fixed design variables (5 actuator lengths) and 5 adaptive design variables, the twist rates for each actuator. Without the multi-objective approach and design space decomposition methodology, two optimization strategies are typically applied to determine both the best lengths and twist rates for each torque tube over each mission stage. One solution would be to solve design and optimization problem via the mission score approach: each mission has a set of fixed design variables and alters the adaptive design variables over each mission stage or mission stage type. This would result in either 25 adaptive design variables (5 actuator twist rates for each of 5 mission stage types) or 70 adaptive design variables (5 actuator twist rates for each of 14 mission stages). Similarly, a multi-objective approach could also optimize over 5 objectives, resulting in 25 adaptive design variables (5 twist rates over 5 objectives), but with a faster evaluation time due to only solving for 5 trim flight conditions (one solutions for each mission stage type) instead of 14 (considering the entire mission sequentially).

The novel approach developed in this work and traditional methodology to determine both fixed and adaptive design variables were compared to improve the air assault mission performance via active rotor blade twist. Additionally, NSGA-AD and NSGA-II were compared as selection criteria for the design space decomposition approach. The best performance improvements were determined using the design space decomposition approach, with NSGA-AD slightly outperforming NSGA-II as the selection algorithm. **By reducing the adaptive design space using the novel methodology presented in this work, heuristic optimization algorithms more effectively determined realizable adaptive rotorcraft designs for a specific multi-stage mission. This is the most significant contribution of this work.**

The effects of ϵ as a similarity parameter were also observed. As $\epsilon \longrightarrow 0$, the assumption of feasibility between designs increases. Considering adaptive rotorcraft mission optimizations, the best adaptive designs for each mission stage when compared with the approximate utopia point, were a function of data density more than the ϵ-constraint. The best adaptive designs were similar

for a given ϵ ranging from 2%-5% for even a sparse rotorcraft adaptive twist optimization.

The computational framework was also extended to include fully coupled aero-elastic responses. While fully coupled analysis is much more time-consuming, analysis of the convergence rates for the local inflow, trim flight, and finite element structural deflection allowed for a set of nested minimizations loops which generate fully-coupled solutions for hover trim flight over the course of an hour as opposed to days, weeks, or months. A rotor blade deflection parameter was introduced to determine structural deflection down the rotor blade length and incorporate the changes into the twist schedule, iterated in loop with the trim flight calculations until convergence. The changes were relatively small (< 10 HP), though the addition of adaptive elements in the fully-coupled model were able to alter the twist and either improve or make worse the trim power depending on the mission weight. This suggests an integrated fully-coupled aerostructural optimization where the changes in twist are determined *inside* the structural analysis may yield possible morphing advantages depending on the loading conditions. Such a fully coupled trim flight solution utilizing computational fluid dynamics and in-situ finite element analysis may be feasible for design and optimization in the future either with supercomputing, parallelization, or longer dedicated runtimes.

In this chapter, high fidelity tools were developed and demonstrated for adaptive rotorcraft design. The mission score and multi-objective optimization methods were compared for a simple adaptive rotor blade example in a reasonably small objective space, and the multi-objective approach still yielded faster solutions with better multi-objective performance. The multi-objective approach can be extended beyond single missions to mission stages across missions, and appears to perform better for larger design spaces with more adaptive variables.

7.1.5 Final Adaptive Rotor Blade Structural Configuration and Prototyping

In this chapter, the mission-driven adaptive rotorcraft design and analysis process is applied to a mission with goal of developing an adaptive rotor blade section prototype at the earliest stages of the design process. Preferred trailing edge configurations for an adaptive mission from the aerodynamic analysis are set as targets for a structural topology optimization, in which the internal

structure necessary to morph between the extreme shapes is optimized. An internal linkage structure composed of both adaptive and non-morphing elements was determined that morphs the rotor blade trailing edge down when the linear actuators contract 1.5%, increasing the camber.

Higher-fidelity aerostructural design and optimization tools were employed, including the CFD interpolation tables, to determine the high-fidelity best adaptive camber region and levels of morphing. After the preferred aerodynamic shapes were selected, a more detailed optimization determined component sizing and relative placements of internal components using the previous topology designed for trailing edge camber morphing. The final structural configuration was sent to Boeing, where the parts were fabricated. The model was assembled at Texas A&M and tested to confirm similar morphing to the target shapes determined by the aerodynamic and structural optimization final configurations. This testing and fabrication combines several novel design and optimization tools that had not yet developed adaptive structures prototypes, from the design space decomposition methodology described in this work to the internal topology generation.

7.2 Contributions to the State of the Art

- An efficient mission-adaptive rotorcraft design and optimization framework was developed for early design space exploration over a wide range of morphing technologies beyond that seen in existing literature. This work considers both structural and aerodynamic effects and, while applied to demonstrate performance improvements for the UH-60, was developed for future rotorcraft design and optimization in early stages of design.

- A novel design and optimization approach toward adaptive structures using local frontiers and similarity of fixed design variables is developed in this work and outperforms traditional approaches toward determining the best designs for systems containing both fixed and adaptive design variables. This methodology can reduce the design space compared with traditional adaptive design approaches and determine better designs (for the same number of functional evaluations) or reduce the number of functional evaluations to determine the best designs.

- A novel optimization selection algorithm informed by NSGA-II was developed and tested considering both nondominance and adaptivity. This new approach toward multi-objective optimization performed well against existing optimization techniques toward adaptive structures, including a mission-driven adaptive rotorcraft design and optimization problem.

A computational design and optimization framework was developed for mission-driven design during the early stages in the design process. Existing adaptive rotorcraft modeling and design processes were integrated into a single platform. While previous modeling tools have been developed to determine advantageous adaptive structures for rotorcraft, this work focuses on improving efficiency to consider both the placement and feasibility of morphing between preferred geometries using a novel design space decomposition. The computational framework described here has flexibility between fidelity and efficiency and explores a wide range of adaptive rotorcraft technologies by leveraging pre-computed CFD, analytical structural models, and a wide range of adaptive technologies in a modular manner.

However, the largest contribution to the state of the art in this work is differentiating between fixed and adaptive design variables to generate sets of similar designs with respect to the fixed design variables. Chapter 3 details a method of determining both fixed and adaptive design variables in a multi-objective optimization without increasing the design space beyond that if all the design variables were fixed. This novel methodology and its applications enable optimizing over all fixed and adaptive design variables and maintaining feasibility without considering every adaptive design variable for every objective as different adaptive variables in the design space. This approach was shown to improve the optimization effectiveness increasingly with increasing objectives for both simple analytical solutions and more complex multi-objective mission-driven adaptive rotorcraft numerical simulations. Additionally, considering adaptivity as a genetic algorithm selection criteria (while still maintaining elitism and dominance) was shown to improve optimization results for both inter-generational designs (i.e., the best adaptive design for a given generation during an optimization) and a final adaptive design when considering every solution from the optimization. Concepts of rank,

dominance, and adaptivity were integrated into evolutionary design and optimization tools to improve multi-objective adaptive design space exploration. This methodology performed well both as a post-processing procedure and as a genetic algorithm selection criteria when compared with existing adaptive structures design approaches.

An uncoupled process was demonstrated where a prototype adaptive rotor blade section was developed based on aerodynamic and structural optimizations to determine preferred geometries, internal structure, and placement in the rotor blade for an adaptive rotor blade during a multi-stage mission. The computational framework is robust and could be expanded to other objectives with additional analysis tools. For instance, noise and blade-vortex interactions could be considered as an objective if coupled with a model that relates the trim flight conditions determined in this work to the noise generated by the rotor blade. The computational framework developed herein could serve as a powerful design and analysis tool beyond the scope of this work in the future.

7.3 Future Work

'We believe this resolves all remaining questions on this topic, no further research is needed' is a phrase rarely if ever uttered in research [122], and this work is no exception. A plethora of extensions to this work present themselves beyond simply utilizing this technology toward designing future adaptive aerostructures.

7.3.1 Adaptive Rotorcraft Modeling and Analysis

First, consider the rotorcraft modeling tools currently developed in this framework. The current tools calculate trim flight and output required power. However, other objectives such as reducing link loading and noise are not considered in this work. Pitch link loading could be calculated using the current trim solutions and some analytical structural models, but evaluating noise signatures would require additional analysis tools. Collaboration with noise signature researchers would be valuable if there is interest to integrate noise reduction into the analysis framework. Other objectives relating to life-cycle analysis could be performed based on trim flight conditions over the course of specified missions or ranges of missions.

Another region of modeling improvement concerns aero-elastic coupling. In this work a fully-coupled model requires finite element analysis (FEA) after every trim solution, which is iterated until convergence. Current work on surrogate modeling could be leveraged to dramatically reduce computational costs if the FEA simulations could be replaced with an accurate pre-computed surrogates.

Additional adaptive rotorcraft technologies could also be considered beyond chord, rotor blade outer mold lines, and twist. In forward flight the rotor blade has cyclic pilot inputs and flapping as the rotor blade rotates around the hub. In TRAC the rotor blade hinges have a constant spring loading and damping, but higher harmonic controls could govern these behaviors as described in the review of current adaptive rotorcraft technologies. Thus, the range of objectives and adaptive technologies this framework could optimize over is not limited to those evaluated in this work. Current adaptive rotorcraft technologies in this work could also be extended. For instance, models relating high-fidelity aerodynamic data to more general outer mold line morphing than the trailing edge parameterized camber could replace the current CFD interpolation tables.

Finally, the computational framework was developed not to redesign the UH-60, but to develop new, state-of-the-art adaptive rotorcraft. New objectives and general parameters for future aircraft should replace the UH-60 vehicle definitions so this set of tools can determine realizable adaptive designs for vehicles currently in early design stages. Altering the vehicle will require additional verification and validation of the computational tools as applied to new aerostructures before adaptive design optimizations can be performed.

In this work active rotorcraft morphing was enabled via shape memory alloy (SMA) actuators. These actuators required power inputs resulting in trade-offs between actuation time and maximum power. In future work these requirements could be incorporated into the design and optimization process via a few different approaches. One approach would be to fix the power during actuation and then minimize the time required to morph between configurations. Another approach would be to require actuation under a specific time, and minimize the required power for morphing. Both approaches would then consider the power requirements in the optimization loop, encouraging

designs with either faster actuation times or lower power requirements.

7.3.2 Adaptive Structures Design and Optimization

The design space decomposition methodology in this work should be further explored and scrutinized. Obviously this work is very new, and the number of studies performed using these tools limited so far this body of work. While comparing methodologies over 100-200 independent optimizations was feasible for simple analytical models, the comparisons require orders of magnitude more time and/or computational resources for similar adaptive rotorcraft mission analysis. Such analysis could be performed with relatively little additional effort beyond dedicated computational resources and time, but those resources and time would be significant. Other rotorcraft missions could be evaluated similarly, across a wide range of adaptive technologies.

Optimization parameters such as population, number of generations until convergence, and the similarity parameter ϵ could be explored in the future to determine the most efficient optimizations based on the number of adaptive design variables, fixed design variables, and objectives. The trade-offs between traditional adaptive structures design and optimization methods could be compared with this work via additional optimizations.

Other adjustments to the selection criteria NSGA-AD may yield convergence improvements. NSGA-AD does not consider diversity preservation, so additional work may be able to improve the selection algorithm via the addition of a diversity-preserving criteria. This optimization methodologies in this would could be combined with other design space exploration techniques and evaluated for further improvements, as NSGA-AD was based on the non-dominated genetic sorting approach used in NSGA-II.

REFERENCES

[1] UH-60AIRLOADS PROJECT. URL: https://rotorcraft.arc.nasa.gov/Research/Programs/uh_60_program.html.

[2] XFLR5. URL: http://www.xflr5.tech/xflr5.htm.

[3] Philip C. Abolmoali, Adam B. Donovan, Soumya S. Patnaik, Patrick McCarthy, Dominic Dierker, Nick Jones, and Robert Buettner. Integrated Propulsive and Thermal Management System Design for Optimal Hybrid Electric Aircraft Performance. In *AIAA Propulsion and Energy 2020 Forum*, VIRTUAL EVENT, August 2020. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.2020-3557, doi:10.2514/6.2020-3557.

[4] Charu C. Aggarwal, Alexander Hinneburg, and Daniel A. Keim. On the Surprising Behavior of Distance Metrics in High Dimensional Space. In Gerhard Goos, Juris Hartmanis, Jan van Leeuwen, Jan Van den Bussche, and Victor Vianu, editors, *Database Theory ICDT 2001*, volume 1973, pages 420–434. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001. URL: http://link.springer.com/10.1007/3-540-44503-X_27, doi:10.1007/3-540-44503-X_27.

[5] Pedro F. Albuquerque, Pedro V. Gamboa, and Miguel A. Silvestre. Mission-Based Multidisciplinary Aircraft Design Optimization Methodology Tailored for Adaptive Technologies. *Journal of Aircraft*, 55(2):755–770, March 2018. URL: https://arc.aiaa.org/doi/10.2514/1.C034403, doi:10.2514/1.C034403.

[6] Alexios Alexiou, Nikolaos Aretakis, Ioannis Kolias, and Konstantinos Mathioudakis. Novel Aero-Engine Multi-Disciplinary Preliminary Design Optimization Framework Accounting for Dynamic System Operation and Aircraft Mission Performance. *Aerospace*, 8(2):49, February 2021. Number: 2 Publisher: Multidisciplinary Digital Publishing In-

stitute. URL: https://www.mdpi.com/2226-4310/8/2/49, doi:10.3390/aerospace8020049.

[7] Darcy L. Allison, Michelle Boyd, and Kevin McCarthy. Transients in Early Aircraft Multidisciplinary Design. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Kissimmee, Florida, January 2018. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.2018-0656, doi:10.2514/6.2018-0656.

[8] Darin J. Arbogast, Robert T. Ruggeri, and Richard C. Bussom. Development of a 1/4-scale NiTinol actuator for reconfigurable structures. In *Industrial and Commercial Applications of Smart Structures Technologies 2008*, volume 6930, pages 148–159. SPIE, March 2008. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/6930/69300L/Development-of-a-1-4-scale-NiTinol-actuator-for-reconfigurable/10.1117/12.775929.full, doi:10.1117/12.775929.

[9] Yijie Bai, Ruisheng Sun, Wei Chen, and Bin Zhu. Design and Experiment of Aircraft Transient Performance Attitude Control System. In Meiping Wu, Yifeng Niu, Mancang Gu, and Jin Cheng, editors, *Proceedings of 2021 International Conference on Autonomous Unmanned Systems (ICAUS 2021)*, Lecture Notes in Electrical Engineering, pages 2845–2853, Singapore, 2022. Springer. doi:10.1007/978-981-16-9492-9_279.

[10] Silvestro Barbarino and Farhan Gandhi. Shape Memory Alloy Actuated Morphing Cellular Frame using Bi-stable von-Mises Trusses with Variable Length Links. In *22nd AIAA/ASME/AHS Adaptive Structures Conference*, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2014. URL: https://arc.aiaa.org/doi/10.2514/6.2014-1417, doi:10.2514/6.2014-1417.

[11] Silvestro Barbarino, Farhan Gandhi, and Steven D. Webster. Design of Extendable Chord Sections for Morphing Helicopter Rotor Blades. *Journal of Intelligent Material Systems*

and Structures, 22(9):891–905, June 2011. Publisher: SAGE Publications Ltd STM. doi:10.1177/1045389X11414077.

[12] E. Belogay, C. Cabrelli, U. Molter, and R. Shonkwiler. Calculating the Hausdorff distance between curves. *Information Processing Letters*, 64(1):17–22, October 1997. URL: https://www.sciencedirect.com/science/article/pii/S0020019097001403, doi:10.1016/S0020-0190(97)00140-3.

[13] Stephan Benz and Rohin Majeti. *Optimization of a Rotor with Morphing Blades using Particle Swarm Algorithm*. September 2019.

[14] Andreas P. F. Bernhard and Inderjit Chopra. Hover Test of Mach-Scale Active Twist Rotor Using Piezo-Bending-Torsion Actuators. *Journal of Aircraft*, 39(4):678–688, July 2002. URL: https://arc.aiaa.org/doi/10.2514/2.2981, doi:10.2514/2.2981.

[15] Robert Biedron and James Thomas. Recent Enhancements to the FUN3D Flow Solver for Moving-Mesh Applications. In *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, Aerospace Sciences Meetings. American Institute of Aeronautics and Astronautics, January 2009. URL: https://arc.aiaa.org/doi/10.2514/6.2009-1360, doi:10.2514/6.2009-1360.

[16] Brent R. Bielefeldt, Ergun Akleman, Gregory W. Reich, Philip S. Beran, and Darren J. Hartl. L-System-Generated Mechanism Topology Optimization Using Graph-Based Interpretation. *Journal of Mechanisms and Robotics*, 11(2), February 2019. doi:10.1115/1.4042512.

[17] Brent R. Bielefeldt, Gregory W. Reich, Philip S. Beran, and Darren J. Hartl. Development and validation of a genetic L-System programming framework for topology optimization of multifunctional structures. *Computers & Structures*, 218:152–169, July 2019. URL: https://www.sciencedirect.com/science/article/pii/S0045794918302207, doi:10.1016/j.compstruc.2019.02.005.

[18] Cees Bil, Ermira Abdullah, and Simon Watkins. Performance of Adaptive Airfoil Control System Using Shape Memory Alloy Actuators for UAV. In *11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, Virginia Beach, VA, September 2011. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.2011-6849, doi:10.2514/6.2011-6849.

[19] G S Bir. Structural Dynamics Verification of Rotorcraft Comprehensive Analysis System (RCAS). Technical Report NREL/TP-500-35328, 15011442, February 2005. URL: http://www.osti.gov/servlets/purl/15011442/, doi:10.2172/15011442.

[20] Kenneth Bonnema and Stephen Smith. AFTI/F-111 Mission Adaptive Wing flight research program. In *4th Flight Test Conference*, San Diego,CA,U.S.A., May 1988. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.1988-2118, doi:10.2514/6.1988-2118.

[21] Boschitsch, A, Quackenbush, T, Wachspress, and D. First-Principles Free-Vortex Wake Analysis for Helicopters and Tiltrotors. AHS International, May 2003. URL: https://vtol.org/store/product/firstprinciples-freevortex-wake-analysis-for-helicopters-and-tiltrotors.cfm.

[22] William G. Bousman. Aerodynamic Characteristics of SC1095 and SC1094 R8 Airfoils:. Technical report, Defense Technical Information Center, Fort Belvoir, VA, December 2003. URL: http://www.dtic.mil/docs/citations/ADA480203, doi:10.21236/ADA480203.

[23] Jason Bowman, Brian Sanders, Bryan Cannon, Jayanth Kudva, Shiv Joshi, and Terrence Weisshaar. Development of Next Generation Morphing Aircraft Structures. In *48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, Hawaii, April 2007. American Institute of Aeronautics and Astro-

nautics. URL: https://arc.aiaa.org/doi/10.2514/6.2007-1730, doi:10.2514/6.2007-1730.

[24] Stephen P. Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, UK ; New York, 2004.

[25] Gheorghe Bunget and Stefan Seelecke. Actuator placement for a bio-inspired bone-joint system based on SMA. In *Active and Passive Smart Structures and Integrated Systems 2009*, volume 7288, pages 259–270. SPIE, April 2009. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7288/72880L/Actuator-placement-for-a-bio-inspired-bone-joint-system-based/10.1117/12.816309.full, doi:10.1117/12.816309.

[26] Nicholas Caldwell, Ephraim Gutmark, and Robert Ruggeri. Performance Predictions of a Blade Twist Actuator System. In *45th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, January 2007. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.2007-1216, doi:10.2514/6.2007-1216.

[27] Arthur Gonçalves Carvalho and Aluizio F.R. Araujo. Improving NSGA-II with an adaptive mutation operator. In *Proceedings of the 11th annual conference companion on Genetic and evolutionary computation conference - GECCO '09*, page 2697, Montreal, Québec, Canada, 2009. ACM Press. URL: http://portal.acm.org/citation.cfm?doid=1570256.1570387, doi:10.1145/1570256.1570387.

[28] Louis Centolanza and Edward Smith. Induced-Shear Piezoelectric Actuators for Active Twist Rotor Blades. In *43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Structures, Structural Dynamics, and Materials and Co-located Conferences. American Institute of Aeronautics and Astronautics, April

2002. URL: https://arc.aiaa.org/doi/10.2514/6.2002-1446, doi:10.2514/6.2002-1446.

[29] Carlos Cesnik, Jiwon Mok, Anish Parikh, and Sangjoon Shin. Optimum Design Framework for Integrally Twisted Helicopter Blades. In *45th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics & Materials Conference*, Structures, Structural Dynamics, and Materials and Co-located Conferences. American Institute of Aeronautics and Astronautics, April 2004. URL: https://arc.aiaa.org/doi/10.2514/6.2004-1761, doi:10.2514/6.2004-1761.

[30] Nachol Chaiyaratana, Theera Piroonratana, and Nuntapon Sangkawelert. Effects of diversity control in single-objective and multi-objective genetic algorithms. *Journal of Heuristics*, 13(1):1–34, January 2007. URL: http://link.springer.com/10.1007/s10732-006-9003-1, doi:10.1007/s10732-006-9003-1.

[31] Pavel M. Chaplya, Milan Mitrovic, Gregory P. Carman, and Friedrich K. Straub. Durability properties of piezoelectric stack actuators under combined electromechanical loading. *Journal of Applied Physics*, 100(12):124111, December 2006. URL: https://aip.scitation.org/doi/full/10.1063/1.2407269, doi:10.1063/1.2407269.

[32] Abheek Chatterjee, Richard Malak, and Astrid Layton. Exploring System of Systems Resilience Versus Affordability Trade-Space Using a Bio-Inspired Metric. *Journal of Computing and Information Science in Engineering*, 21(5), May 2021. doi:10.1115/1.4050288.

[33] Robert Christie, Matthew Robinson, Fernando Tejero, and David G. MacManus. The use of hybrid intuitive class shape transformation curves in aerodynamic design. *Aerospace Science and Technology*, 95:105473, December 2019. URL: https://www.sciencedirect.com/science/article/pii/S1270963819305450, doi:10.1016/j.ast.2019.105473.

[34] Antonio Concilio and Rohin Majeti. *An Advanced Surrogate Model for Predicting the Impact of a SMA Twist System on the Helicopter Performance*. October 2019.

[35] Sarah Conley and Dorsa Shirazi. Comparing Simulation Results from CHARM and RotCFD to the Multirotor Test Bed Experimental Data. In *AIAA AVIATION 2021 FORUM*. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/abs/10.2514/6.2021-2540, doi:10.2514/6.2021-2540.

[36] Allen Davis, Bochan Lee, Moble Benedict, and Darren Hartl. Biomimetic Adaptive Airframe Technology (BAAT) for Rotorcraft Design and Optimization. In *VFS78*, Fort Worth, TX, May 2022. URL: https://vtol.org/annual-forum/forum-78-proceedings.

[37] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, April 2002. doi:10.1109/4235.996017.

[38] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable multi-objective optimization test problems. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No.02TH8600)*, volume 1, pages 825–830 vol.1, May 2002. doi:10.1109/CEC.2002.1007032.

[39] Kalyanmoy Deb and Himanshu Jain. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. *IEEE Transactions on Evolutionary Computation*, 18(4):577–601, August 2014. doi:10.1109/TEVC.2013.2281535.

[40] Matthew DiPalma, Etana Ferede, and Farhan Gandhi. Optimization of Extension-Twist Coupled Composite Blades for High-Speed Rotorcraft. page 14.

[41] Matthew DiPalma and Farhan Gandhi. Autonomous camber morphing of a helicopter rotor blade with temperature change using integrated shape memory alloys. *Journal of Intelligent*

Material Systems and Structures, 32(5):499–515, March 2021. Publisher: SAGE Publications Ltd STM. doi:10.1177/1045389X20953613.

[42] John H. Drake, Ahmed Kheiri, Ender Özcan, and Edmund K. Burke. Recent advances in selection hyper-heuristics. *European Journal of Operational Research*, 285(2):405–428, September 2020. URL: https://www.sciencedirect.com/science/article/pii/S0377221719306526, doi:10.1016/j.ejor.2019.07.073.

[43] Mark Drela. XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils. In Thomas J. Mueller, editor, *Low Reynolds Number Aerodynamics*, Lecture Notes in Engineering, pages 1–12, Berlin, Heidelberg, 1989. Springer. doi:10.1007/978-3-642-84010-4_1.

[44] Dupont. Kapton HN General-Purpose Polyimide Film. URL: https://www.dupont.com/products/kapton-hn.html.

[45] Omega East. *NPS Historical Handbook: Wright Brothers*. National Park Services, 1961. URL: https://www.nps.gov/parkhistory/online_books/hh/34/hh34c.htm.

[46] T. Fitzgibbon, M. Woodgate, and G. Barakos. Assessment of current rotor design comparison practices based on high-fidelity CFD methods. *The Aeronautical Journal*, 124(1275):731–766, May 2020. Publisher: Cambridge University Press. URL: https://www.cambridge.org/core/journals/aeronautical-journal/article/assessment-of-current-rotor-design-comparison-practices-based-on-highf
70C0CEF7522292CE1075859D6F7E9AD7#, doi:10.1017/aer.2019.162.

[47] Carlos M. Fonseca and Peter J. Fleming. Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization, 1993.

[48] Félix-Antoine Fortin, François-Michel De Rainville, Marc-André Gardner, Marc Parizeau, and Christian Gagné. DEAP: Evolutionary Algorithms Made Easy. *Journal of Machine*

Learning Research, 13:2171–2175, July 2012.

[49] Peretz P. Friedmann and Thomas A. Millott. Vibration reduction in rotorcraft using active control - A comparison of various approaches. *Journal of Guidance, Control, and Dynamics*, 18(4):664–673, July 1995. Publisher: American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/3.21445, doi:10.2514/3.21445.

[50] Edgar Galvan and Richard J. Malak. P3GA: An Algorithm for Technology Characterization. *Journal of Mechanical Design*, 137(1), January 2015. doi:10.1115/1.4028101.

[51] Farhan Gandhi, Mary Frecker, and Andrew Nissly. Design Optimization of a Controllable Camber Rotor Airfoil. *AIAA Journal*, 46(1):142–153, 2008. Publisher: American Institute of Aeronautics and Astronautics _eprint: https://doi.org/10.2514/1.24476. doi:10.2514/1.24476.

[52] Xian-Zhong Gao, Zhong-Xi Hou, Zheng Guo, Xiao-Qing Chen, and Xiao-Qian Chen. Joint optimization of battery mass and flight trajectory for high-altitude solar-powered aircraft. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 228(13):2439–2451, November 2014. doi:10.1177/0954410013518510.

[53] G. García. Approximating the Hausdorff distance by alpha-dense curves. *Boletín de la Sociedad Matemática Mexicana*, 2021. doi:10.1007/s40590-021-00382-8.

[54] W. Geisler, Guido Dietz, Holger Mai, B. Junker, and Th Lorkowski. Dynamic Stall Control Investigations on a Full Size Chord Blade Section. Marseille, France, January 2004. URL: https://www.researchgate.net/publication/224785439_Dynamic_Stall_Control_Investigations_on_a_Full_Size_Chord_Blade_Section.

[55] Alfred Gessow. Effect of Rotor-Blade Twist and Plan-Form Taper on Helicopter Hovering Performance. *National Advisory Committee for Aeronautics*, 1542:27,

1948. URL: https://ntrs.nasa.gov/api/citations/19930082219/downloads/19930082219.pdf.

[56] Alfred Gessow. FLIGHT INVESTIGATION OF EFFECTS OF ROTOR-BLADE TWIST ON HELICOPTER PERFORMANCE IN THE HIGH -SPEED AND VERTICAL-AUTOROTATIVE -DESCENT CONDITIONS. *National Advisory Committee for Aeronautics*, 1666:30, 1948.

[57] Rahman Gharari, Navid Poursalehi, Mohammadreza Abbasi, and Mahdi Aghaie. Implementation of Strength Pareto Evolutionary Algorithm II in the Multiobjective Burnable Poison Placement Optimization of KWU Pressurized Water Reactor. *Nuclear Engineering and Technology*, 48(5):1126–1139, October 2016. URL: https://www.sciencedirect.com/science/article/pii/S1738573316300493, doi:10.1016/j.net.2016.04.004.

[58] William W. Gilbert. Mission Adaptive Wing System for Tactical Aircraft. *Journal of Aircraft*, 18(7):597–602, July 1981. URL: https://arc.aiaa.org/doi/10.2514/3.57533, doi:10.2514/3.57533.

[59] Ioannis Goulos, Vassilios Pachidis, Roberto dIppolito, Jos Stevens, and Chrissy Smith. An Integrated Approach for the Multidisciplinary Design of Optimum Rotorcraft Operations. *Journal of Engineering for Gas Turbines and Power*, 134(9), July 2012. doi:10.1115/1.4006982.

[60] Boris Grohmann, Christoph Maucher, Peter Jänker, and Boris Grohmann. ACTUATION CONCEPTS FOR MORPHING HELICOPTER ROTOR BLADES. page 10, 2006.

[61] Weilin Guan, Nguyen K. Pham, and Edwin A. Peraza Hernandez. Design exploration of a tensegrity twisting wing enabled by shape memory alloy wire actuation. In *Active and Passive Smart Structures and Integrated Systems XV*, volume 11588, pages 16–32. SPIE, March 2021. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11588/1158809/

Design-exploration-of-a-tensegrity-twisting-wing-enabled-by-shape/10.1117/12.2582829.full, doi:10.1117/12.2582829.

[62] Thomas Hamilton and Brian J. German. Optimal Airspeeds for Scheduled Electric Aircraft Operations. *Journal of Aircraft*, 56(2):545–555, March 2019. URL: https://arc.aiaa.org/doi/10.2514/1.C035051, doi:10.2514/1.C035051.

[63] Dong Han, Vasileios Pastrikakis, and George N. Barakos. Helicopter flight performance improvement by dynamic blade twist. *Aerospace Science and Technology*, 58:445–452, November 2016. URL: https://www.sciencedirect.com/science/article/pii/S1270963816306708, doi:10.1016/j.ast.2016.09.013.

[64] Robert V. Hess and Clifford S. Gardner. Study by the Prandtl-Glauert method of compressibility effects and critical Mach number for ellipsoids of various aspect ratios and thickness ratios, March 1947. NTRS Author Affiliations: NTRS Report/Patent Number: NACA-RM-L7B03a NTRS Document ID: 19930093731 NTRS Research Center: Legacy CDMS (CDMS). URL: https://ntrs.nasa.gov/citations/19930093731.

[65] Desmond J Higham and Nicholas J Higham. *MATLAB guide*, volume 150. Siam, 2016.

[66] Mike Hirschberg. JMR Technology Demonstration Update: The Road to Future Vertical Lift. 62(1):6, 2016.

[67] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Complex Adaptive Systems. A Bradford Book, Cambridge, MA, USA, April 1992.

[68] Phoenix Integration. Phoenix Integration's ModelCenter is THE framework for Model Based Engineering, which reduces development costs, improves efficiency and stimulates innovation. URL: https://www.phoenix-int.com/.

[69] Stephen A Jacklin, Axel Haber, Gary deSimone, Thomas R Norman, Cahit Kitaplioglu, and Patrick Shinoda. Full-Scale Wind Tunnel Test of an Individual Blade Control System for a UH-60 Helicopter. page 12, Montreal, Canada, 2002. American Helicopter Soceity.

[70] A. Dean Jacot. Shape memory alloy consortium (SMAC). In *Smart Structures and Materials 1999: Industrial and Commercial Applications of Smart Structures Technologies*, volume 3674, pages 158–165. SPIE, July 1999. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/3674/0000/Shape-memory-alloy-consortium-SMAC/10.1117/12.351552.full, doi:10.1117/12.351552.

[71] S. Jafari and T. Nikolaidis. Meta-heuristic global optimization algorithms for aircraft engines modelling and controller design; A review, research challenges, and exploring the future. *Progress in Aerospace Sciences*, 104:40–53, January 2019. URL: https://www.sciencedirect.com/science/article/pii/S0376042118301416, doi:10.1016/j.paerosci.2018.11.003.

[72] Ruholla Jafari-Marandi and Brian K. Smith. Fluid Genetic Algorithm (FGA). *Journal of Computational Design and Engineering*, 4(2):158–167, April 2017. URL: https://www.sciencedirect.com/science/article/pii/S2288430016300458, doi:10.1016/j.jcde.2017.03.001.

[73] Rohit Jain and Hyeonsoo Yeo. Effects of torsion frequencies on rotor performance and structural loads with trailing edge flap. *Smart Materials and Structures*, 21(8):085026, August 2012. URL: https://iopscience.iop.org/article/10.1088/0964-1726/21/8/085026, doi:10.1088/0964-1726/21/8/085026.

[74] Wayne Johnson. *Helicopter Theory*. Courier Corporation, January 1994. Google-Books-ID: SgZheyNeXJIC.

[75] Wayne Johnson. Technology Drivers in the Development of CAMRAD II - Vertical Flight Library & Store. San Francisco, CA., 1994. URL: https://vtol.org/store/product/technology-drivers-in-the-development-of-camrad-ii-7645.cfm.

[76] Wayne Johnson. NDARC NASA Design and Analysis of Rotorcraft Validation and Demonstration. page 34, San Francisco, CA., 2010.

[77] Shiv Joshi, Zeb Tidwell, William Crossley, and Sekaripuram Ramakrishnan. Comparison of Morphing Wing Stategies Based upon Aircraft Performance Impacts. In *45th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics & Materials Conference*. American Institute of Aeronautics and Astronautics, 2008. _eprint: https://arc.aiaa.org/doi/pdf/10.2514/6.2004-1722. URL: https://arc.aiaa.org/doi/abs/10.2514/6.2004-1722, doi:10.2514/6.2004-1722.

[78] Massimiliano Kaucic, Mojtaba Moradi, and Mohmmad Mirzazadeh. Portfolio optimization by improved NSGA-II and SPEA 2 based on different risk measures. *Financial Innovation*, 5(1):26, June 2019. doi:10.1186/s40854-019-0140-6.

[79] Adeel S. Khalid. Development and Implementation of Rotorcraft Preliminary Design Methodology using Multidisciplinary Design Optimization. November 2006. URL: asme.

[80] Maryam Khoshlahjeh, Silvestro Barbarino, and Salvatore Ameduri. Chapter 17 - Shape memory alloy applications for helicopters. In Antonio Concilio, Vincenza Antonucci, Ferdinando Auricchio, Leonardo Lecce, and Elio Sacco, editors, *Shape Memory Alloy Engineering (Second Edition)*, pages 591–607. Butterworth-Heinemann, Boston, January 2021. URL: https://www.sciencedirect.com/science/article/pii/B9780128192641000170, doi:10.1016/B978-0-12-819264-1.00017-0.

[81] Maryam Khoshlahjeh and Farhan Gandhi. Extendable Chord Rotors for Helicopter Envelope Expansion and Performance Improvement. *Journal of the American Helicopter Society*, 59, January 2014. doi:10.4050/JAHS.59.012007.

[82] Jinhan Kim and Shin Yoo. Software review: DEAP (Distributed Evolutionary Algorithm in Python) library. *Genetic Programming and Evolvable Machines*, 20(1):139–142, March 2019. doi:10.1007/s10710-018-9341-4.

[83] Noboru Kobiki, Yasutada Tanabe, Takashi Aoyama, Do-Hyung Kim, Hee Jung Kang, Seong-Yong Wie, and Seung-Ho Kim. Design, Analysis and Prototyping of Active Tab Rotor. *Transactions of the Japan Society for Aeronautical and Space Sciences*, 62(2):64–74, 2019. doi:10.2322/tjsass.62.64.

[84] Frank Kody, Ethan Corle, Mark D. Maughmer, and Sven Schmitz. Higher-Harmonic Deployment of Trailing-Edge Flaps for Rotor-Performance Enhancement and Vibration Reduction. *Journal of Aircraft*, 53(2):333–342, 2016. doi:10.2514/1.C033335.

[85] Dominik Komp, Manfred Hajek, and Juergen Rauleder. Rotor Performance Enhancements with Spanwise Varying Active Camber Morphing. In *AIAA Scitech 2021 Forum*, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2021. URL: https://arc.aiaa.org/doi/10.2514/6.2021-0200, doi:10.2514/6.2021-0200.

[86] Dominik Komp, Sumeet Kumar, Manfred Hajek, and Jürgen Rauleder. Effect of active camber morphing on rotor performance and control loads. *Aerospace Science and Technology*, 108:106311, January 2021. URL: https://www.sciencedirect.com/science/article/pii/S1270963820309937, doi:10.1016/j.ast.2020.106311.

[87] J. Krishnamurthi and F. Gandhi. Chord Extension Morphing for Active Rotor Track and Balance. 2018. doi:10.2514/6.2018-0798.

[88] Jayanth Krishnamurthi and Farhan Gandhi. Flight Simulation and Control of a Helicopter Undergoing Rotor Chord Extension Morphing. page 23.

[89] Brenda Kulfan and John Bussoletti. *"Fundamental" Parameteric Geometry Representations for Aircraft Component Shapes*. September 2006. doi:10.2514/6.2006-6948.

[90] Devesh Kumar. *Design and Analysis of Composite Rotor Blades for Active/Passive Vibration Reduction.* Thesis, 2013. Accepted: 2013-06-12T14:17:11Z. URL: http://deepblue.lib.umich.edu/handle/2027.42/98006.

[91] P.K. Kumar and D.C. Lagoudas. Introduction to Shape Memory Alloys. In *Shape Memory Alloys: Modeling and Engineering Applications*, pages 1–51. Springer US, Boston, MA, 2008. doi:10.1007/978-0-387-47685-8_1.

[92] Sumeet Kumar, Dominik Komp, Manfred Hajek, and Jürgen Rauleder. Integrated Rotor Performance Improvement and Vibration Reduction Using Active Camber Morphing. American Society of Mechanical Engineers Digital Collection, December 2019. URL: https://asmedigitalcollection.asme.org/SMASIS/proceedings/SMASIS2019/59131/V001T04A010/1071365, doi:10.1115/SMASIS2019-5588.

[93] Dimitris Lagoudas. *Shape Memory Alloys: Modeling and Engineering Applications*. Springer, 1st edition, 2008.

[94] R. Larson. Flight control system development and flight test experience with the F-111 mission adaptive wing aircraft. In *Astrodynamics Conference*, Fluid Dynamics and Co-located Conferences. American Institute of Aeronautics and Astronautics, August 1986. URL: https://arc.aiaa.org/doi/10.2514/6.1986-2237, doi:10.2514/6.1986-2237.

[95] Pedro B. C. Leal and Darren J. Hartl. Structurally Consistent Class/Shape Transformation Equations for Morphing Airfoils. *Journal of Aircraft*, 56(2):505–516, March 2019. URL: https://arc.aiaa.org/doi/10.2514/1.C035025, doi:10.2514/1.C035025.

[96] Pedro BC Leal, Marcelo A Savi, and Darren J Hartl. Aero-structural optimization of shape memory alloy-based wing morphing via a class/shape transformation approach. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 232(15):2745–2759, December 2018. Publisher: IMECHE. doi:10.1177/0954410017716193.

[97] leal26. AeroPy, December 2021. original-date: 2015-08-02T23:25:37Z. URL: https://github.com/leal26/AeroPy.

[98] Bochan Lee. *Helicopter Autonomous Ship Landing System*. Masters of Aerospace Engineering, Texas A&M University, College Station, TX, 2018.

[99] Bochan Lee. *ON THE COMPLETE AUTOMATION OF VERTICAL FLIGHT AIRCRAFT SHIP LANDING*. Doctorate of Philosophy in Aerospace Engineering, Texas A&M University, College Station, TX, 2021.

[100] Bochan Lee and Moble Benedict. Development and Validation of a Comprehensive Helicopter Flight Dynamics Code. In *AIAA Scitech 2020 Forum*, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2020. URL: https://arc.aiaa.org/doi/10.2514/6.2020-1644, doi:10.2514/6.2020-1644.

[101] Leishman. *Principles of Helicopter Aerodynamics | Aerospace engineering*. Cambridge Univesity Press, 2000. URL: https://www.cambridge.org/us/academic/subjects/engineering/aerospace-engineering/principles-helicopter-aerodynamics-2nd-edition-1,https://www.cambridge.org/us/academic/subjects/engineering/aerospace-engineering.

[102] Miqing Li, Jinhua Zheng, and Jun Wu. Improving NSGA-II Algorithm Based on Minimum Spanning Tree. In Xiaodong Li, Michael Kirley, Mengjie Zhang, David Green, Vic Ciesielski, Hussein Abbass, Zbigniew Michalewicz, Tim Hendtlass, Kalyanmoy Deb, Kay Chen Tan, Jürgen Branke, and Yuhui Shi, editors, *Simulated Evolution and Learning*, Lecture Notes in Computer Science, pages 170–179, Berlin, Heidelberg, 2008. Springer. doi:10.1007/978-3-540-89694-4_18.

[103] Chen Liang, Frank M. Davidson, Lawrence McDonald Schetky, and Friedrich K. Straub. Applications of torsional shape memory alloy actuators for active rotor blade control: opportunities and limitations. In *Smart Structures and Ma-*

terials 1996: Smart Structures and Integrated Systems, volume 2717, pages 91–100. SPIE, May 1996. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/2717/0000/Applications-of-torsional-shape-memory-alloy-actuators-for-active-roto 10.1117/12.239049.full, doi:10.1117/12.239049.

[104] Yuanchang Liang, Minoru Taya, and Yasuo Kuga. Design of membrane actuators based on ferromagnetic shape memory alloy composite for the synthetic jet actuator. In *Smart Structures and Materials 2004: Smart Structures and Integrated Systems*, volume 5390, pages 268–275. SPIE, July 2004. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/5390/0000/Design-of-membrane-actuators-based-on-ferromagnetic-shape-memory-alloy 10.1117/12.540100.full, doi:10.1117/12.540100.

[105] Aristid Lindenmayer. Mathematical models for cellular interactions in development I. Filaments with one-sided inputs. *Journal of Theoretical Biology*, 18(3):280–299, March 1968. URL: https://www.sciencedirect.com/science/article/pii/0022519368900799, doi:10.1016/0022-5193(68)90079-9.

[106] Li Liu, Peretz Friedmann, Insung Kim, and Dennis Bernstein. Rotor Performance Enhancement and Vibration Reduction in Presence of Dynamic Stall Using Actively Controlled Flaps. *Journal of The American Helicopter Society - J AMER HELICOPTER SOC*, 53, October 2008. doi:10.4050/JAHS.53.338.

[107] Yuan Lyu and Rhea P. Liem. Flight performance analysis with data-driven mission parameterization: mapping flight operational data to aircraft performance analysis. *Transportation Engineering*, 2:100035, December 2020. URL: https://www.sciencedirect.com/science/article/pii/S2666691X20300361, doi:10.1016/j.treng.2020.100035.

[108] James Mabe, Lynn Gravatt, Glenn Bushnell, Ephraim Gutmark, Russell DiMicco, and Chris

Harris. Shape Memory Alloy Actuators for Deployable Rotor Blade Aerodynamic Devices. In *46th AIAA Aerospace Sciences Meeting and Exhibit*, Aerospace Sciences Meetings. American Institute of Aeronautics and Astronautics, January 2008. URL: https://arc.aiaa.org/doi/10.2514/6.2008-1451, doi:10.2514/6.2008-1451.

[109] Rohin Majeti and Stephan Benz. Mission-based optimal morphing parameters for rotors with combined chord and twist morphing. *Open Research Europe*, 1:121, October 2021. doi:10.12688/openreseurope.14060.1.

[110] Rohin K. Majeti, Berend G. van der Wall, and Christoph G. Balzarek. Linearly variable chord-extension morphing for helicopter rotor blades. *CEAS Aeronautical Journal*, 12(1):55–67, January 2021. doi:10.1007/s13272-020-00477-4.

[111] Pascual Marqués. Active Blade Twist in Rotary UAVs using Smart Actuation. In *Advanced UAV Aerodynamics, Flight Stability and Control*, pages 399–419. John Wiley & Sons, Ltd, 2017. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118928691.ch11, doi:10.1002/9781118928691.ch11.

[112] Mathworks. MATLAB - MathWorks, August 2019. URL: https://www.mathworks.com/products/matlab.html.

[113] J.-B. Maurice, F.a. King, and W. Fichter. Derivation and Validation of a Helicopter Rotor Model with Trailing-Edge Flaps. *Journal of Guidance, Control, and Dynamics*, 36(5):1375–1387, September 2013. URL: https://arc.aiaa.org/doi/10.2514/1.59102, doi:10.2514/1.59102.

[114] David McCullough. *The Wright Brothers*. Simon and Schuster, 1 edition, May 2016. URL: https://www.simonandschuster.com/books/The-Wright-Brothers/David-McCullough/9781476728759.

[115] Madalyn Mikkelsen, Michayal Mathew, Patrick Walgren, Brent Bielefeldt, Pedro B. C. Leal, Darren Hartl, and Andres Arrieta. Morphing Airfoil Design via L-System Generated Topology Optimization. In *ASME 2019 Conference on*

Smart Materials, Adaptive Structures and Intelligent Systems, Louisville, Kentucky, USA, September 2019. American Society of Mechanical Engineers. URL: https://asmedigitalcollection.asme.org/SMASIS/proceedings/SMASIS2019/59131/Louisville,%20Kentucky,%20USA/1071392, doi:10.1115/SMASIS2019-5695.

[116] Judah Milgram, Inderjit Chopra, and Friedrich Straub. Rotors with Trailing Edge Flaps: Analysis and Comparison with Experimental Data. *Journal of The American Helicopter Society - J AMER HELICOPTER SOC*, 43, October 1998. doi:10.4050/JAHS.43.319.

[117] Matthew Misiorowski, Michael Pontecorvo, and Farhan Gandhi. A Bi-Stable System for Rotor Span Extension in Rotary-Wing Micro Aerial Vehicles. In *23rd AIAA/AHS Adaptive Structures Conference*, Kissimmee, Florida, January 2015. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.2015-0788, doi:10.2514/6.2015-0788.

[118] Mihir Mistry. *Quasi-static rotor morphing concepts for rotorcraft performance improvements*. Ph.D., The Pennsylvania State University, United States – Pennsylvania. URL: https://www.proquest.com/docview/1355741601/abstract/E72730160B20415DPQ/1.

[119] Milan Mitrovic, Greg P. Carman, and Friedrich K. Straub. Response of piezoelectric stack actuators under combined electro-mechanical loading. *International Journal of Solids and Structures*, 38(24):4357–4374, June 2001. URL: https://www.sciencedirect.com/science/article/pii/S0020768300002730, doi:10.1016/S0020-7683(00)00273-0.

[120] J. Morgado, M. Â. R. Silvestre, and J. C. Páscoa. Validation of New Formulations for Propeller Analysis. *Journal of Propulsion and Power*, 31(1):467–477, January 2015.

URL: https://arc.aiaa.org/doi/10.2514/1.B35240, doi:10.2514/1.B35240.

[121] Patrick Moser, Silvestro Barbarino, and Farhan Gandhi. Helicopter Rotor-Blade Chord Extension Morphing Using a Centrifugally Actuated Von Mises Truss. *Journal of Aircraft*, 51(5):1422–1431, September 2014. URL: https://arc.aiaa.org/doi/10.2514/1.C032299, doi:10.2514/1.C032299.

[122] Randall Munroe. Further Research is Needed, February 2020. URL: https://xkcd.com/2268/.

[123] Khanh Nguyen, Mark Betzina, and Cahit Kitaplioglu. Full-Scale Demonstration of Higher Harmonic Control for Noise and Vibration Reduction on the XV-15 Rotor. *Journal of the American Helicopter Society*, 46:182–191, July 2001. doi:10.4050/JAHS.46.182.

[124] Khanh Nguyen and Inderjit Chopra. Application of higher harmonic control (HHC) to hingeless rotor systems. In *30th Structures, Structural Dynamics and Materials Conference*, Mobile,AL,U.S.A., April 1989. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.1989-1215, doi:10.2514/6.1989-1215.

[125] D. Norman, R. Hynes, and D. Gangsaas. An integrated maneuver enhancement and gust alleviation mode for theAFTI/F-111 MAW aircraft. In *Guidance and Control Conference*, Guidance, Navigation, and Control and Co-located Conferences. American Institute of Aeronautics and Astronautics, August 1983. URL: https://arc.aiaa.org/doi/10.2514/6.1983-2217, doi:10.2514/6.1983-2217.

[126] Jae-Sang Park and Ji-Hwan Kim. Design and aeroelastic analysis of active twist rotor blades incorporating single crystal macro fiber composite actuators. *Composites Part B: Engineering*, 39(6):1011–1025, September 2008. URL: https://www.sciencedirect.com/science/article/pii/S1359836807001539, doi:10.1016/j.compositesb.2007.11.006.

[127] M. Pettes-Duler, X. Roboam, B. Sareni, Y. Lefevre, J.-F. Llibre, and M. Fenot. Multidisciplinary Design Optimization of the Actuation System of a Hybrid Electric Aircraft Powertrain. *Electronics*, 10(11):1297 (24 pp.), 2021. Place: Switzerland Publisher: MDPI. doi:10.3390/electronics10111297.

[128] Francis R. Phillips, Trent D. White, Allen Davis, and Darren J. Hartl. Analysis of Rotor Blade Aeroelastic Deformation Utilizing the Uncoupled Static Aeroelastic Analysis Method. In *AIAA SCITECH 2022 Forum*, San Diego, CA & Virtual, January 2022. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.2022-2555, doi:10.2514/6.2022-2555.

[129] Michael E Pontecorvo, Silvestro Barbarino, Gabriel J Murray, and Farhan S Gandhi. Bistable arches for morphing applications. *Journal of Intelligent Material Systems and Structures*, 24(3):274–286, February 2013. Publisher: SAGE Publications Ltd STM. doi:10.1177/1045389X12457252.

[130] Przemyslaw Prusinkiewicz and James Hanan. *Lindenmayer Systems, Fractals, and Plants*. URL: https://link.springer.com/book/10.1007/978-1-4757-1428-9.

[131] Todd R. Quackenbush. *A comprehensive hierarchical aeromechanics rotorcraft model (CHARM) for general rotor/surface interaction : SBIR Phase II final report*. C.D.I. report 99,3. Continuum Dynamics, Inc., Princeton, NJ, 1999. Section: 166 Seiten.

[132] Jürgen Rauleder, Amine Abdelmoula, Dominik Komp, and Sumeet Kumar. Aerodynamic Performance of Morphing Blades and Rotor Systems. page 21.

[133] Kumar Ravichandran, Shreyas Ananthan, Graduate Research, Asst Scientist, Alfred Gessow, Rotorcraft Center, and Benjamin Hein. *Active Rotor Controls for Vibration Reduction and Performance Enhancement*. December 2019.

[134] Kumar Ravichandran, Inderjit Chopra, Brian Wake, and Benjamin Hein. Trailing-Edge Flaps for Rotor Performance Enhancement and Vibration Reduction. *Journal of the Ameri-*

can Helicopter Society, 58:1–13, April 2013. doi:10.4050/JAHS.58.022006.

[135] Beatrice Roget and Inderjit Chopra. Wind-Tunnel Testing of Rotor with Individually Controlled Trailing-Edge Flaps for Vibration Reduction. *Journal of Aircraft*, 45(3):868–879, May 2008. URL: https://arc.aiaa.org/doi/10.2514/1.28455, doi:10.2514/1.28455.

[136] Ankit Rohatgi. Webplotdigitizer: Version 4.5, 2021. URL: https://automeris.io/WebPlotDigitizer.

[137] Robert T. Ruggeri, A. Dean Jacot, and Dan J. Clingman. Shape memory actuator systems and the use of thermoelectric modules. In *Smart Structures and Materials 2002: Industrial and Commercial Applications of Smart Structures Technologies*, volume 4698, pages 386–394. SPIE, July 2002. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/4698/0000/Shape-memory-actuator-systems-and-the-use-of-thermoelectric-modules/10.1117/12.475086.full, doi:10.1117/12.475086.

[138] Umberto Saetti. *Rotorcraft Simulations with Coupled Flight Dynamics, Free Wake, and Acoustics*. PhD thesis, August 2016. doi:10.13140/RG.2.1.2259.3529.

[139] Rohan S. Sharma and Serhat Hosder. Investigation of Mission-Driven Inverse Aircraft Design Space Exploration with Machine Learning. *Journal of Aerospace Information Systems*, 18(11):774–789, 2021. doi:10.2514/1.I010966.

[140] John Shaw. *Higher harmonic blade pitch control : a system for helicopter vibration reduction*. Thesis, Massachusetts Institute of Technology, 1980. URL: https://dspace.mit.edu/handle/1721.1/15817.

[141] John Shaw and Nicholas Albion. Active Control of the Helicopter Rotor for Vibration Reduction. *Journal of the American Helicopter Society*, 26(3):32–39, July 1981. doi:10.4050/JAHS.26.32.

[142] Patrick M. Shinoda, Hyeonsoo Yeo, and Thomas R. Norman. Rotor Performance of a UH60 Rotor System in the NASA Ames 80 by 120Foot Wind Tunnel. *Journal of the American Helicopter Society*, 49(4):401–413, October 2004. doi:10.4050/JAHS.49.401.

[143] Jayant Sirohi and Inderjit Chopra. Compact Piezo-Hydraulic Hybrid Actuator. In *Twelfth International Conference on Adaptive Structures and Technologies*. CRC Press, 2002. Num Pages: 12.

[144] Michael Smith. *ABAQUS/Standard User's Manual, Version 6.9*. Dassault Systèmes Simulia Corp, United States, 2009.

[145] Stephen B. Smith and David W. Nelson. Determination of the aerodynamic characteristics of the mission adaptive wing. *Journal of Aircraft*, 27(11):950–958, November 1990. Publisher: American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/3.45965, doi:10.2514/3.45965.

[146] Jaroslaw Sobieszczanski-Sobieski. Optimizing Multiobjective Function Problems. In *Multidisciplinary Design Optimization Supported by Knowledge Based Engineering*, pages 98–115. John Wiley & Sons, Ltd, 2015. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118897072.ch6, doi:10.1002/9781118897072.ch6.

[147] B. K. Sreejith and Vighnesha Nayak. Capability of RANS simulation to predict laminar separation bubble on the E216 airfoil. *Materials Today: Proceedings*, 46:2884–2887, January 2021. URL: https://www.sciencedirect.com/science/article/pii/S2214785321025074, doi:10.1016/j.matpr.2021.03.453.

[148] M. Srinivas and L.M. Patnaik. Genetic algorithms: a survey. *Computer*, 27(6):17–26, June 1994. doi:10.1109/2.294849.

[149] W. Z. Stepniewski and C. N. Keys. *Rotary-Wing Aerodynamics*. Dover Publications, New York, reprint edition edition, July 1984.

[150] F. K. Straub, H. T. Ngo, V. Anand, and D. B. Domzalski. Development of a piezoelectric actuator for trailing edge flap control of full scale rotor blades. *Smart Materials and Structures*, 10(1):25–34, February 2001. doi:10.1088/0964-1726/10/1/303.

[151] Friedrich K. Straub. A feasibility study of using smart materials for rotor control. *Smart Materials and Structures*, 5(1):1–10, February 1996. doi:10.1088/0964-1726/5/1/002.

[152] Friedrich K. Straub, Dennis K. Kennedy, Alan D. Stemple, V. R. Anand, and Terry S. Birchette. Development and whirl tower test of the SMART active flap rotor. In *Smart Structures and Materials 2004: Industrial and Commercial Applications of Smart Structures Technologies*, volume 5388, pages 202–212. SPIE, July 2004. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/5388/0000/Development-and-whirl-tower-test-of-the-SMART-active-flap/10.1117/12.562645.full, doi:10.1117/12.562645.

[153] Justin K. Strelec, Dimitris C. Lagoudas, Mohammad A. Khan, and John Yen. Design and Implementation of a Shape Memory Alloy Actuated Reconfigurable Airfoil. *Journal of Intelligent Material Systems and Structures*, 14(4-5):257–273, April 2003. Publisher: SAGE Publications Ltd STM. doi:10.1177/1045389X03034687.

[154] H. Stroud and D. Hartl. Shape memory alloy torsional actuators: a review of applications, experimental investigations, modeling, and design. *Smart Materials and Structures*, 29(11):113001, October 2020. Publisher: IOP Publishing. URL: https://doi.org/10.1088/1361-665x/abbb12, doi:10.1088/1361-665X/abbb12.

[155] Daniel B. Sullivan, Florent Righi, Darren J. Hartl, and Jonathan Rogers. Shape Memory Alloy Rotor Blade Deicing. In *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Boston, Massachusetts, April 2013. American In-

stitute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.2013-1915, doi:10.2514/6.2013-1915.

[156] Joseph Totah. A critical assessment of UH-60 main rotor blade airfoil data. In *11th Applied Aerodynamics Conference*, Monterey,CA,U.S.A., August 1993. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.1993-3413, doi:10.2514/6.1993-3413.

[157] Hsue-Shen Tsien and Lester Lees. The Glauert-Prandtl Approximation for Subsonic Flows of a Compressible Fluid. *Journal of the Aeronautical Sciences*, 12(2):173–187, April 1945. URL: https://arc.aiaa.org/doi/10.2514/8.11219, doi:10.2514/8.11219.

[158] Vimal L. Vachhani, Vipul K. Dabhi, and Harshadkumar B. Prajapati. Improving NSGA-II for solving multi objective function optimization problems. In *2016 International Conference on Computer Communication and Informatics (ICCCI)*, pages 1–6, January 2016. doi:10.1109/ICCCI.2016.7479921.

[159] John Valasek. *Morphing Aerospace Vehicles and Structures*. Aerospace Series. John Wiley & Sons, Ltd, 1 edition, 2012.

[160] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.

[161] Gerhard Venter and Jaroslaw Sobieszczanski-Sobieski. Particle Swarm Optimization. *AIAA Journal*, 41(8):1583–1589, August 2003. URL: https://arc.aiaa.org/doi/10.2514/2.2111, doi:10.2514/2.2111.

[162] J. B. Vos, A. Rizzi, D. Darracq, and E. H. Hirschel. NavierStokes solvers in European aircraft design. *Progress in Aerospace Sciences*, 38(8):601–697, November 2002. URL: https://www.sciencedirect.com/science/article/pii/S0376042102000507, doi:10.1016/S0376-0421(02)00050-7.

[163] Patrick Philip Walgren. *A NONLINEAR SUBSTRUCTURE METHOD FOR EFFICIENT REDUCED-ORDER STRUCTURAL MODELING BASED ON A CLASSICAL PLASTICITY FRAMEWORK*. PhD Dissertation, Texas A&M University, College Station, TX, August 2022.

[164] Xiangyu Wang, Yanping Yang, Dong Wang, and Zijian Zhang. Mission-oriented cooperative 3D path planning for modular solar-powered aircraft with energy optimization. *Chinese Journal of Aeronautics*, 35(1):98–109, January 2022. URL: https://www.sciencedirect.com/science/article/pii/S1000936121001692, doi:10.1016/j.cja.2021.04.015.

[165] Annual Forum Washington, Jr Chester, and W. Langston. Presented at the American Helicopter Society 57, 2001.

[166] Jonathan M. Weaver-Rosen, Pedro B. C. Leal, Darren J. Hartl, and Richard J. Malak. Parametric optimization for morphing structures design: application to morphing wings adapting to changing flight conditions. *Structural and Multidisciplinary Optimization*, 62(6):2995–3007, December 2020. URL: https://link.springer.com/10.1007/s00158-020-02643-y, doi:10.1007/s00158-020-02643-y.

[167] Eric W. Weisstein. L^2-Norm. URL: https://mathworld.wolfram.com/.

[168] P. Wierach, S. Opitz, and S. Kalow. Experimental investigation of an active twist model rotor blade with a low voltage actuation system. *The Aeronautical Journal*, 119(1222):1499–1512, December 2015. URL: https://www.cambridge.org/core/product/identifier/S0001924000011374/type/journal_article, doi:10.1017/S0001924000011374.

[169] Matthew Wilbur, Paul Mirick, Yeager, Teddy William, Chester Langston, Carlos Cesnik, and Sang-Joon Shin. Vibratory loads reduction testing of the NASA/Army/MIT active twist rotor. *Journal of the American Helicopter Society*, 47, February 2001. doi:10.4050/JAHS.47.123.

[170] Matthew L. Wilbur and W. Keats Wilkie. ACTIVE-TWIST ROTOR CONTROL APPLICATIONS FOR UAVs. In *Transformational Science and Technology for the Current and Future Force*, volume Volume 42 of *Selected Topics in Electronics and Systems*, pages 185–192. WORLD SCIENTIFIC, November 2006. URL: https://www.worldscientific.com/doi/abs/10.1142/9789812772572_0024, doi:10.1142/9789812772572_0024.

[171] Ji Won Woo, Jun-Young An, Moon Gyeang Cho, and Chang-Joo Kim. Integration of path planning, trajectory generation and trajectory tracking control for aircraft mission autonomy. *Aerospace Science and Technology*, 118:107014, November 2021. URL: https://www.sciencedirect.com/science/article/pii/S1270963821005241, doi:10.1016/j.ast.2021.107014.

[172] Benjamin K. S. Woods, Curt S. Kothera, and Norman M. Wereley. Wind Tunnel Testing of a Helicopter Rotor Trailing Edge Flap Actuated via Pneumatic Artificial Muscles. *Journal of Intelligent Material Systems and Structures*, 22(13):1513–1528, September 2011. doi:10.1177/1045389X11424216.

[173] XPREG. XPREG - Prepregs for Advanced Composites. URL: http://www.xpreg.com/.

[174] Fan Yang, Zhufeng Yue, Lei Li, and Weizhu Yang. Aerodynamic optimization method based on Bezier curve and radial basis function. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 232(3):459–471, March 2018. doi:10.1177/0954410016679433.

[175] Yung H. Yu, Bernd Gmelin, Wolf Splettstoesser, Jean J. Philippe, Jean Prieur, and Thomas F. Brooks. Reduction of helicopter blade-vortex interaction noise by active rotor control technology. *Progress in Aerospace Sciences*, 33(9):647–687, January 1997. URL: https://www.sciencedirect.com/science/article/pii/S0376042197000067, doi:10.1016/S0376-0421(97)00006-7.

[176] Yasir Zahoor, Roeland De Breuker, and Mark Voskuijl. Preliminary Design of a TE Morphing Surface for Rotorcraft. In *AIAA Scitech 2020 Forum*, Orlando, FL, January 2020. American Institute of Aeronautics and Astronautics. URL: https://arc.aiaa.org/doi/10.2514/6.2020-1301, doi:10.2514/6.2020-1301.

[177] Wei Zhang and Jiao-long Zhang. Improved NSGA-II algorithm based on differential evolution mechanism. In *2017 36th Chinese Control Conference (CCC)*, pages 4334–4338, July 2017. ISSN: 1934-1768. doi:10.23919/ChiCC.2017.8028039.

[178] Eckart Zitzler. Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications. page 134.

[179] Eckart Zitzler, Marco Laumanns, and Lothar Thiele. SPEA2: Improving the Strength Pareto Evolutionary Algorithm. Technical report, 2001.

[180] Eckart Zitzler and Lothar Thiele. An Evolutionary Algorithm for Multiobjective Optimization: The Strength Pareto Approach. page 43.

APPENDIX A

UH-60 BLACKHAWK PARAMETERS AND CONFIGURATION

Table A.1: Parameters for the UH-60 Blackhawk

Main Rotor		
Parameter	Value	Units
Number of blades	4	N/A
Radius	8.17	m
Chord	0.53	m
Solidity	0.08	N/A
Rotational speed	27	$\frac{\text{rads}}{\text{s}}$
Tip speed	220.6	$\frac{\text{m}}{\text{s}}$
Mast tilt	-3.0	Degrees
Airfoil	SC-1095	N/A
Maximum manufactured twist	7.8	degrees
Mass per unit length	14	$\frac{\text{kg}}{\text{m}}$
Tail Rotor		
Parameter	Value	Units
Number of blades	4	N/A
Radius	1.68	m
Chord	0.25	m
Rotational speed	125	$\frac{\text{rads}}{\text{s}}$
Tip speed	210	$\frac{\text{m}}{\text{s}}$
Solidity	0.19	N/A

APPENDIX B

FINAL ADAPTIVE ROTOR BLADE CONFIGURATION

Table B.1: Final Adaptive Rotor Blade Model Parameters for Construction

Model Parameter	Design Value(s)	Units
Upper CST Parameters	0.171 0.181 0.287 0.096	N/A
Lower CST Parameters	0.135 0.0188 -0.0124 -0.0268	N/A
Plastic Thickness	1.01	mm
Airfoil Segment Width	5	mm
SMA Actuator Loading	15	$\frac{\text{lbs}}{\text{in}}$
Carbon Fiber Thickness	1.5	mm (6 layers)
Spar Location	0.15	m

Table B.2: Final Internal Topology Locations

	Starting Location $[x, y]$ m	**Ending Location** $[x, y]$ m
Leading Edge	$[0, 0]$	
Internal Plastic Structure	$[.15, -0.0197]$	$[.2, -0.0102]$
	$[.2, -0.0102]$	$[.1834, 0.0151]$
	$[.15, 0.0162]$	$[.1834, 0.0151]$
	$[.1843, -0.0151]$	$[.2567, 0043]$
	$[.2074, -0.0194]$	$[.2567, 0.0043]$
	$[.2567, 0.0043]$	$[.3017, 0.0191]$
SMA Actuators	$[.2, -0.0102]$	$[.2567, 0.0043]$
	$[.2074, -0.0194]$	$[.4667, 0.0041]$
Spar	$[.15, -0.0197]$	$[.15, 0.0277]$
	$[.0498577, -0.0167]$	$[.0498577, 0.0229]$
	$[.100119, -0.0192]$	$[.100119, 0.0273]$

APPENDIX C

MODELCENTER FRAMEWORK

C.1 ModelCenter Framework

The ModelCenter framework consists of pre-processing, trim flight analysis numerical solvers, and post-processing codes. The pre-processing codes define the rotorcraft geometry and flight conditions, along with any morphable parameters during the flight. The trim flight analysis codes, a hover BEMT code and TRAC, solve for trim flight based on the pilots inputs, and return the required power, velocity, collective pitch, etc. In post-processing, objectives such as trim power, fuel burn, range time, etc. are determined based on the trim flight solution(s). In the ModelCenter optimization framework, pre-processing inputs are altered under a set of flight conditions defined by specific missions to maximize the performance metrics evaluated in post-processing. In pre-processing, the rotor blade geometry, morphing parameters, and flight profile are defined. The local aerodynamic outer mold line and twist are common adaptive rotor blade parameters

C.1.1 DefineMission

The function DefineMission pulls the mission parameters from a function called packageMissionDefinitions, which defines the mission stages and other mission defined parameters. The individual phase definitions can be seen in the comments of the individual functions. Some parameters for the mission definition can be seen in Fig. C.1.

C.1.2 Switches

The function Switches defines a number of options for the ModelCenter evaluations. The defaults work for twist and camber morphing using the existing CFD tables. Other CFD tables could be generated for other airfoils and camber morphing profiles.

Matlab Commands

```
%%% DEFINE MISSION PHASES %%%
% missionPhase key:
% 1: hover (HOGE), 2: loiter (at optimal speed), 3: maximum range, 4: cruise as a specified speed,
% 10+: run at specified speed (knots), <1, climb or rate (in km/s to ensure value below 1)
missionPhase = [1 180 1 180 3];
phaseNumber = length(missionPhase);

% Length of each mission phase (in hours or nautical miles)
phaseLength = [5.0/60 50/180 5/60 50/180 30/60 ];

% Starting fuel weight converted from lbs to newtons
fuelWeight = fuelWeightInput; % 2401*4.45 (2  tanks totaling 360 gallons);

% Empty fuselage weight (no passengers, added equipment, or fuel) in Newtons
emptyWeight = ones(1,phaseNumber)*60805;

% Passenger weight (365 lb passengers) in Newtons. Input total passengers in brackets
passengerWeight = [2 2 8 8 8]*365*4.45;

% Equipment weight (converted from lbs to Newtons)
equipmentWeight = ones(1,phaseNumber)*0;

% Total Mission weight (including total fuel weight, which will be subtracted out in TRAC), in Newtons
missionWeight = emptyWeight+fuelWeight+passengerWeight+ equipmentWeight;

% Air density for each mission phase in kg/m^3
airDens = [ones(1,phaseNumber)]*0.949;

% Specify whether the mission phase is defined by time or distance (1 = time, 2 = distance)
timeFlag = [1 1 1 1 1];
```

Mission Phases

Phase Lengths

Initial Fuel

Empty Weight

Passenger Weight

Equipment Weight

Phase Air Density

Figure C.1: The parameters for each mission stage are defined and packaged at the DefineMission stage.

C.1.3 Inputs

The function Inputs defines the geometric and morphing properties. Optimizations and trade studies alter these parameters, after the mission and switches have been defined.

C.1.4 AirfoilProfiles

The function AirfoilProfiles is used if the morphing CFD tables are not. When using parameterized trailing edge camber CFD tables interpolation tables are generated for every camber morphing parameter. The outer mold lines can also be generated using Class/Shape Transformation (CST) tables, which are then discretized and the aerodynamic properties determined using AeroPy and XFOIL. The key outputs are lift, drag, and moment interpolation tables where the inputs for a 2D blade element section are local angle of attack and Mach number.

C.1.5 CallSizing

The function CallSizing sizes SMA torque tubes for twist morphing based on the geometric and material properties of the actuators and rotor blade. The added mass of the actuators is added to the rotor blade mass distribution. There are checks to ensure the torque tube achieves the require twist without violating stress and strain considerations, and the variable actFit returns whether an actuator large enough to meet the stress/strain constraints will fit inside the rotor blade based on internal clearance. The main output passed forward is actuatorMass, which will be added to the total rotor mass in the function TotalMass.

C.1.6 ProcessingTRACInputs

The folder ProcessingTRACInputs contains the definitions for change in rotor twist, local chord, and local lift curve slope. The lift curve slope is used to determine some TRAC parameters and is defined either as a constant or from the linear regions of the airfoil profiles.

C.1.7 TotalMass

The function TotalMass defines both the rotor blade mass per unit length for each rotor blade element and the total mass of the rotorcraft at takeoff.

C.1.8 DefineCamber

The function DefineCamber defines the CFD camber morphing parameter for multiple stages down the length of the rotor blade. The camber requires start and end locations and a morphing parameter from -0.1 (trailing edge down) to 0.1 (trailing edge up/reflex). The camber profile can either be an immediate, discrete jump or an elliptical camber distribution defined by the switch ellipticalCamber.

C.1.9 TrimRotorCraftvXX

The BEMT code and TRAC are called in this function to evaluate mission performance and determine the required power for trim flight over various mission stages. Missions can be run sequentially or independently

C.1.10 CalculatePerformance

The function CalculatePerformance determines performance metrics such as fuel burn, power for each mission stage, maximum range, etc. These performance metrics can then be passed to ModelCenter trade studies and optimizations. All final mission performance metrics should run through this function.

APPENDIX D

MISSION DEFINITIONS

D.1 MEDEVAC Mission

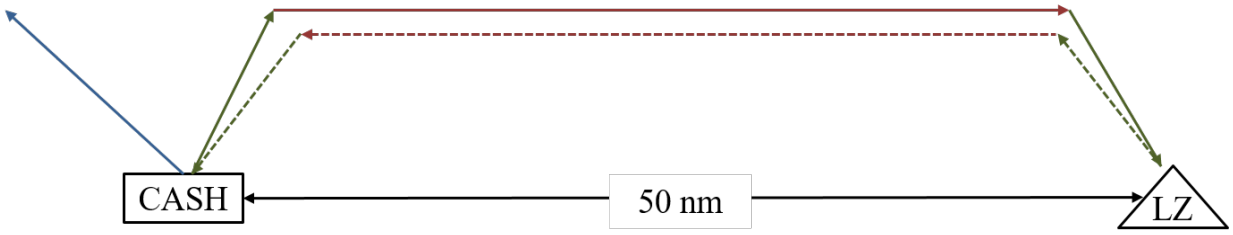

Figure D.1: The MEDEVAC mission is 5-stage mission consisting of three different mission stage types.

Figure D.2: Five Stage MEDEVAC Mission Definition

Mission Phase Description	Flight Condition	Time/Distance
1) Hover over TAA at 6,000 ft. altitude on 95 deg. day	Hover	0.5 min
2) Cruise at high continuous speed (180 knots)	Cruise	50 nm
3) Descend to add passengers (6 passengers 365 lbs)	Hover	1 min
4) Cruise at high continuous speed (180 knots) to CASH	Cruise	50 nm
5) Determine max range based on fuel reserve	Max Range	Until out of fuel

Figure D.3: MEDEVAC Mission Constraints

Description	Value	Units
Final Range Time	30	Minutes
Maximum Power (over all 5 mission stages)	3150	Horsepower

D.2 Air Assault Mission

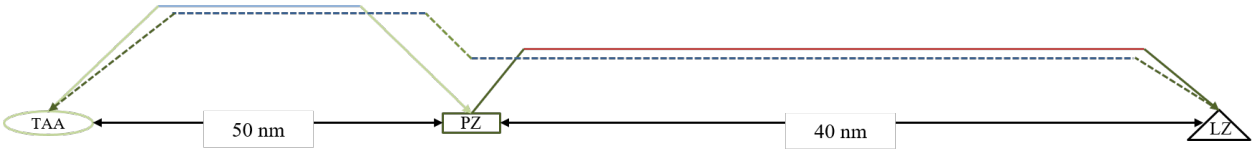

Figure D.4: A proposed 14-stage mission with 5 mission stage types. The mission starts and ends at the Terminal Arrival Area (TAA).

Figure D.5: 14-Stage Air Assault Mission

Mission Phase Description	Flight Condition	Time/Distance
1) Hover over TAA at 6,000 ft. altitude on 95 deg. day	Light Hover	0.5 min
2) Climb 200 ft.	Light Hover	2 min
3) Cruise at best range speed	Light Range	50 nm
4) Descend to land at PZ to pick up 12 soldiers	Light Hover	1 min
5) Hover before takeoff	Heavy Hover	0.5 min
6) Climb 80 ft.	Heavy Hover	1 min
7) Cruise at 160 knots	Heavy Cruise	40 nm
8) Descend to land at LZ (using hover blade shape)	Heavy Hover	1 min
9) Climb 80 ft. using hover	Heavy Hover	2 min
10) Cruise at best range speed	Heavy Range	40 nm
11) Climb 200 ft.	Heavy Hover	2 min
12) Cruise at best range speed to TAA	Heavy Range	50 nm
13) Descend to land at TAA unload soldiers	Heavy Hover	1 min
14) Determine max range based on fuel reserve	Light Range	Until out of fuel

Figure D.6: Air Assault Mission Constraints

Description	Value	Units
Final Range Time	30	Minutes
Maximum Power (across all 14 stages)	3150	Horsepower

D.3 High Lift Mission

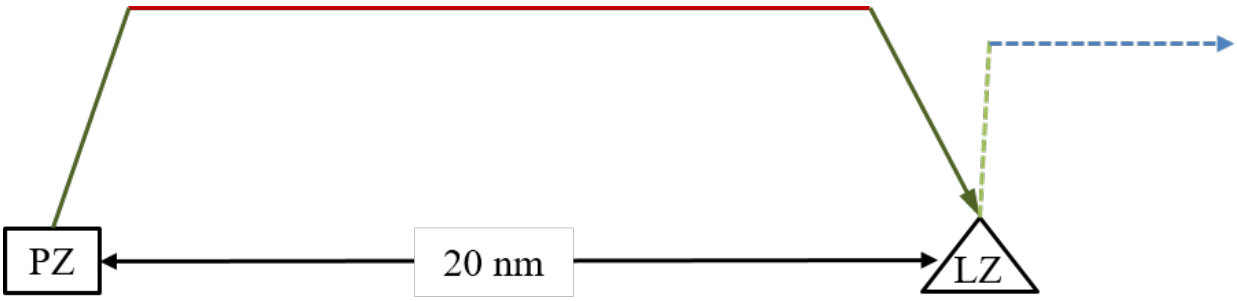

Figure D.7: 7-Stage Heavy Load Mission

Figure D.8: 7-Stage Heavy Load Mission

Mission Phase Description	Flight Condition	Time/Distance
1) Hover over Pickup Zone (PZ) at 6,000 ft. altitude, 95 deg. then add 4,000 lbs load	Light Hover	0.5 min
2) Climb 300 ft.	Heavy Hover	3 min
3) Cruise at 120 knots	Heavy Cruise	20 nm
4) Descend 300' to landing zone (LZ)	Heavy Hover	3 min
5) Hover, then release added cargo	Heavy Hover	0.5 min
6) Climb 300 ft.	Light Hover	3 min
7) Determine max range based on fuel reserve	Light Range	NA

Figure D.9: Heavy Load Mission Constraints

Description	Value	Units
Final Range Time	30	Minutes
Maximum Power (across all 7 stages)	3150	Horsepower

APPENDIX E

DISTANCE (ERROR) QUANTIFICATION

E.1 Distance (Error) Quantification

In the current work, distance (error) metrics are necessary to measure between sets individual designs, usually as represented in the design and objective space. The design methods described in section 3.1.1 require calculating distances to determine the quality of each adaptive Pareto frontier, either by comparing a set of many points (when comparing local adaptive frontiers to the approximate idealized frontier) or a set consisting of a single point (when comparing adaptive utopia points to the idealized utopia point). Meanwhile, a measure of distance in the design space is necessary to determine ϵ-similar sets of designs with respect to fixed design variables $\mathbf{x}^f$.

In the n^{th}-dimensional objective space, a distance quantification serves as an error metric to be minimized using the distance metrics described in section 3.1.1. This distance, or error, metric Dist between non-empty sets $\mathbf{X}$ and $\mathbf{Y}$ must satisfy a few conventional requirements:

- The distance must not be negative for both the total value and for any dimension, such that:

$$\mathrm{Dist}\{\mathbf{X}, \mathbf{Y}\} \geq \mathbf{0}. \tag{E.1}$$

- The distance between two sets is symmetric (i.e. the distance between any two sets $\mathbf{X}$ and $\mathbf{Y}$ is equal to the distance between sets $\mathbf{Y}$ and $\mathbf{X}$):

$$\mathrm{Dist}\{\mathbf{X}, \mathbf{Y}\} = \mathrm{Dist}\{\mathbf{Y}, \mathbf{X}\}. \tag{E.2}$$

Additionally, the error metrics used in this work should apply broadly across different types of design variables and objectives with minimal hyper-parameters. Due to the general nature of design and optimization, the methodology must remain applicable via [163]:

- Normalization: a distance metric must apply across different design variables and objectives with varying scales and units. Each dimension may be non-dimensionalized either in pre-processing or post-processing. In this work, distances are non-dimensionalized immediately before the distance calculations based on either the bounds prescribed by the user or those of the current n^{th}-dimensional space.

- Outlier-handling: In this work, designs with very poor performance in one objective may still provide valuable adaptivity for other objectives. In some cases designs with outliers in one dimension can still be considered in another dimension in the objective space. Particularly when considering continuous adaptivity, it is important to find solutions that neither exclude nor are dominated by outliers in any dimension of the objective space. In contrast with the objective space, outliers in the design space are excluded from families of similar designs.

- Dataset-agnostic: distance metrics must be able to consider different set sizes and distributions. The size of optimization datasets can vary wildly depending on the design space, objective space, computational runtime, etc. The design and objective spaces have no guarantee for normally distributed data; a distance metric should not make assumptions about the distributions. Dataset sets may range from a single point (as in the utopia point case) to an entire population of designs.

Multiple distance metrics are necessary to meet these requirements across the entire design and optimization methodology presented here, specifically because the number of dimensions in design and objective spaces must be the same for comparison, while the number of designs in a given set has no such restrictions. The first metric is a general formulation for the distance between sets of the same size with n dimensions, the *Minkowski distance* [4]:

$$Dist\{\boldsymbol{X}, \boldsymbol{Y}\} = \left(\sum_{i=1}^{n} |X_i - Y_i|^p \right)^{\frac{1}{p}}, \tag{E.3}$$

where p is the user-specified power, typically an integer. When $p \longrightarrow \infty$, the total distance

approaches the largest distance in any dimension, while at $p = 1$, the *Minkowski distance* equals the sum of distances across all dimensions. Another common value for p is 2, which is equal to the Euclidean distance or an L_2-norm [167]. In this work, distances in the design variable space are typically calculated using a *Minkowski distance* where $p = \infty$. This metric was selected because coupling between design variables, along with scaling, is not assumed. As such, the ϵ-constraint is applied to each fixed design variable individually, and designs which violate that constraint are not considered ϵ-similar.

The distance between approximate and adaptive utopia points, however, is measured in this work as a Euclidean distance, minimizing the distance where $p = 2$. This yields a final adaptive design as close in Euclidean space as possible to the approximated ideal solution. Of course, other coefficients could also be evaluated, and is a parameter that can be passed into the optimization tools as developed in this work. A distance metric previously described in this chapter used to evaluate continuous adaptivity is the *Hausdorff* distance. Thus, to determine the best continuously adaptive designs across multiple objectives as in equation 3.6, equation 3.7 is reduced to (for adaptive frontier $\hat{\mathbf{J}}^*(\mathbf{x}_i^f)$ and approximate Pareto frontier $\mathbf{J}^*$):

$$Dist_H\{\hat{\mathbf{J}}^*(\mathbf{x}_i^f), \mathbf{J}^*\} \approx \max\left\{\max_{j \in \mathbf{J}^*} d(j, \hat{\mathbf{J}}^*(\mathbb{D}_i))\right\}. \tag{E.4}$$

Substituting equation E.4 into equation 3.6 and replacing $\mathbf{x}_i^f$ with the approximate representation $\mathbb{D}_i$ (c.f. section 3.2.2) yields:

$$\hat{\hat{\mathbf{J}}}^* = \min_{\mathbb{D}_i}\left(\max\left\{\max_{j \in \mathbf{J}^*} d(j, \hat{\mathbf{J}}^*(\mathbb{D}_i))\right\}\right), \tag{E.5}$$

resulting in a final set of fixed design variables minimizing the distance between any design in $\mathbf{J}^*$ and the approximate local frontier $\hat{\mathbf{J}}^*(\mathbb{D}_i)$.