Machine-Learned Propulsion Strategies: From Adaptive Damage Compensation to Advanced Aeromobility

Thesis by
Meredith Leigh Hooper

In Partial Fulfillment of the Requirements for the
degree of
Ph.D. in Aeronautics

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2025
Defended May 29, 2025

Meredith Leigh Hooper
ORCID: 0009-0004-4819-9941

ACKNOWLEDGEMENTS

This endeavor would not have been possible without the guidance, mentorship, and support provided by many people along my journey. The people named and acknowledged here are only a small part of the community that has shaped me into the researcher I am today.

First, I am extremely grateful to Professor Mory Gharib for accepting me into his research group and supporting my growth as a researcher and academic throughout the last four years. Thank you for encouraging my curiosity for all sorts of different projects. I would also like to thank the members of my candidacy and defense committee, John Dabiri (chair), John Sader, Anima Anandkumar, Aaron Ames, and Yisong Yue, for their time and thoughtful feedback. Furthermore, this work would not be possible without the funding I have received from GALCIT, the Center for Autonomous Systems and Technologies (CAST), as well as the National Science Foundation (NSF GRFP).

I am deeply grateful to my collaborator Isabel Scherl, who at the time was a postdoc in the Colonius group. Isabel helped guide me through the writing and publication of my first paper, generously giving her time and feedback.

The support of the entire Gharib group has been invaluable. I have learned so much from our staff members David Jeon and especially Reza Nemovi. Reza's helicopter knowledge, help redesigning my experimental system, and connection with Kistler truly enabled Part II of this thesis. I am also indebted to the expertise of Alejandro Stefan-Zavala and Alexander Rusch, who were always willing to provide tech support from their computer science and electrical engineering backgrounds. Special thanks goes also to my office-mates throughout the years: fellow Ph.D. students Scott Bollt, Sean Devey, and Alejandro Stefan-Zavala, and visiting students Paloma Garcia Guillen and Julius Wanner. Oftentimes, our technical discussions would lend me a new lens through which to view a problem, and our non-technical discussions helped make each day of research a little less lonely. Finally, the Gharib group could not function without the efforts of our admin Sarah Pontes and CAST admin Karima Guthrie. Thank you both for handling all the minutiae, and for all the times you helped me track down Mory.

I am grateful to the GALCIT community, members of Women in GALCIT, and members of my first-year cohort. Sincere gratitude goes to Jacque Tawney and

Morgan Hooper, with whom I was able to dive deeper into educational research and even teach a class! Thank you also to members of the Dabiri lab Nina Mohebbi and Matt Fu, who generously explained the galvo-based laser scanning system I attempted to use and also allowed me to borrow optical components and even a laser for a time. This generosity is reflective of the spirit of the GALCIT community as a whole. I have been lucky to be a part of it.

Lastly, I would like to thank my family and friends. Working towards a Ph.D. is a challenging journey, filled with twists and turns, and my family and friends have provided critical support along the way. I am so grateful for my best friends Ariel Chen, Morlan Osgood, and Jessica Schreiber, whose many phone calls and visits helped keep my spirits high. Branson Davis, who I met my first year at Caltech, has become a rock for me in this journey. Thank you for your love, support, and your wonderful family. Thank you also to my parents and my sister, who have provided limitless love, support, and encouragement my entire life. I cannot overstate how lucky I am to have parents who both understand exactly what doing a Ph.D. entails. With all the love in my heart, thank you for being the best parents.

ABSTRACT

Autonomous vehicles are regularly sent into "dull, dirty, and dangerous" environments where the risk of damage is high. Avoidance or mitigation of such damage is therefore paramount to maintain effective autonomy. In this thesis, we use machine learning to investigate two different propulsive strategies that may be used by autonomous vehicles. The first, flapping propulsion, shows remarkable ability in nature to recover from damage simply by altering stroke kinematics. Using machine learning, we ask whether and how such mitigation of damage would be possible for a robotic autonomous vehicle. The second propulsive strategy we investigate is single-rotor propulsion, most commonly seen in helicopters. With this system, we seek to avoid damage before it occurs by improving mobility and control authority via thrust vectoring.

In Part I, we use an evolutionary strategy (CMA-ES) with hardware-in-the-loop to explore optimal machine-learned adaptations to propulsor damage. Experimental function evaluations are performed by a flexible propulsor actuated by a spherical parallel manipulator (SPM). The machine-learned forces and trajectory parameters are compared to *in vivo* observations in order to determine whether bio-inspired strategies to adapt to significant propulsor damage are the most efficient, or whether they may be affected by irrelevant evolutionary pressures. With amputation of approximately 50% of the propulsor, we find that a complete recovery in thrust production and fitness is made. Some characteristics of the recovered trajectory are similar to natural swimmers, while others differ. Recovery when producing side-force is even more complex. Not all trials are able to recover force production and fitness, and no clear strategy to modify amplitude or frequency is seen. We conclude Part I by using PIV measurements to detail the effect of compensatory strategies on hydrodynamics. Both amputated and intact trajectories clearly show utilization of a drag-based paddling strategy, but the hydrodynamics of the intact and amputated fins differ significantly. This suggests that the machine-learned trajectories are not simply reestablishing the same wake as the intact fin to achieve the same thrust and fitness.

Given the success in applying machine learning in-the-loop to a complex propulsive system where fluid-structure interactions are significant, we utilize the same strategy in Part II to begin to explore helicopter aeromechanics. We built an independent blade control (IBC) system that interfaces with the CMA-ES algorithm to explore

optimal blade pitch trajectories. Using this platform, we explore two preliminary optimizations designed to vector thrust: the first, for sustained thrust vectoring that might be utilized upon takeoff or landing, and the second, for short-time thrust vectoring that could be used for enhanced maneuverability. We present some preliminary results from these optimizations and lay out a foundation for future applications of this experimental system.

PUBLISHED CONTENT AND CONTRIBUTIONS

[1] M. L. Hooper, I. Scherl, and M. Gharib. “Bio-inspired compensatory strategies for damage to flapping robotic propulsors”. In: *Journal of The Royal Society Interface* 21.216 (July 2024), p. 20240141. ISSN: 1742-5662. DOI: 10.1098/rsif.2024.0141.

M.L.H.: conceptualization, data curation, formal analysis, funding acquisition, investigation, methodology, visualization, writing—original draft, writing—review and editing; I.S.: formal analysis, supervision, visualization, writing—review and editing; M.G.: conceptualization, funding acquisition, methodology, project administration, resources, supervision, writing—review and editing.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Part I

Flapping Propulsion

Chapter 1

INTRODUCTION

Flapping propulsion is an example of convergent evolution in animals as varied as fish, insects, birds, bats, and more. In contrast, the most commonly used human-invented strategy for propelling similar-sized vehicles through these fluids is propeller- or rotor-based. Rotorcraft like quadcopters perform quite well at the same scale as birds and bats, but rotary designs are more challenged by the scale and environment of fish and insects. Autonomous underwater vehicles (AUVs) using propellers are both less efficient and less maneuverable than natural swimmers [39]. Miniaturizing rotors for micro-air vehicles (MAVs) substantially reduces their efficiency [120], while implementing flapping propulsion on MAVs enhances maneuverability and increases the low-speed lift-to-drag ratio [88].

Furthermore, flapping propulsion is a remarkably robust method of propulsion. Some animals may lose up to 76% of their propulsive surface before any loss of thrust is observed, maintaining force production simply by altering their gait [43]. To maintain full autonomy, autonomous vehicles must be able to self-repair their functionality during a mission. AUVs and MAVs are often sent into "dull, dirty, and dangerous" environments [7] where the risk of damage is high. If a portion of the propulsive surface is damaged, AUVs and MAVs using flapping propulsion may be able to maintain their ability to complete their mission simply by adjusting their flapping kinematics, rather than needing recovery and/or repair. In Part I, we utilize an evolutionary strategy to determine how a flapping robot might adjust its kinematics to accommodate for loss of 50% of its propulsor. We investigate the efficiency of such adjustments and also determine the effect of these compensatory strategies on the hydrodynamics.

In Chapters 2 and 3, we give an overview of the kinematics and hydro/aerodynamics of fish swimming and insect flight. Fish swimming and insect flight represent the natural forms of flapping propulsion that are most similar to our model of flapping propulsion with a flexible flat plate. The flapping robot's kinematics are based on the key kinematic features of the natural swimmers and flyers that we describe in Chapters 2 and 3. Furthermore, the hydro/aerodynamic effects seen by natural swimmers and flyers are similar to the hydrodynamics of the flapping robot that we

discuss in Chapter 6.

Chapter 4 delineates the experimental setup of all flapping propulsion experiments. We also describe key characteristics of the covariance matrix adaptation evolutionary algorithm (CMA-ES) that forms the machine learning methodology for both Part I and Part II of this thesis. In Part I, we optimize a flexible robotic propulsor's trajectory to produce a desired force at maximum efficiency. We then amputate approximately 50% of the propulsive surface and reoptimize the trajectory, in most cases achieving full recovery in force production. In Chapter 5 we discuss the kinematics, force production, and power consumption of these amputated trajectories in comparison to amputated fish and insects. Chapter 6 provides flow visualization of the amputated optimal trajectories compared to the intact optimal trajectory and describes similar hydrodynamic features to natural swimmers and flyers. We conclude this section in Chapter 7 by discussing main takeaways and future directions.

Chapter 2

FISH SWIMMING

Fish swimming, achieved via flapping propulsion, provides a useful model for highly efficient and maneuverable underwater propulsion with minimal robotic complexity. Natural swimmers use a variety of methods to propel themselves through water. For example, organisms like jellyfish and squid have evolved a clapping mechanism for underwater propulsion, where a bell-shaped chamber ejects fluid to provide jetted thrust. We choose to focus on flapping propulsion rather than clapping propulsion because it is more efficient (controlling for actuator complexity) [98]. Here we discuss the main features of flapping propulsion in fish.

2.1 Relevant Morphology

Because fish possess a variety of different fins, each of which may contribute to thrust generation, maneuverability, or both, we must first define the nomenclature used to refer to these fins.

From tail to snout, fish may have up to five types of fins that fall into two subcategories. The tail fin, called the caudal fin, attaches to the rest of the body via the caudal peduncle. At the bottom of the fish near the tail may exist an anal fin. At the top center of the fish may exist a dorsal fin. These three fins are called median fins, because they are single fins collocated with the body midline. At the bottom center there may exist a pair of fins called the pelvic fins, and on either side of the fish rearward of the head there may exist a pair called the pectoral fins. These two types of fins fall into the second category: paired fins. Figure 2.1 shows each of these fins and the caudal peduncle on a fish.

2.2 Types of Propulsion

Some fish swim by undulating parts of their body and/or caudal fin, abbreviated as BCF. Other fish use median and/or paired fins to swim, but usually at lower speeds (MPF). BCF is by far the most prevalent; only 12% of fish utilize MPF as their main mode of propulsion [150]. Fish who use BCF may occasionally augment or replace BCF propulsion with MPF to enable greater maneuverability [130]. The categorization between BCF and MPF kinematics originates from Breder [15], who further divided these two characteristic modes of swimming into subcategories.

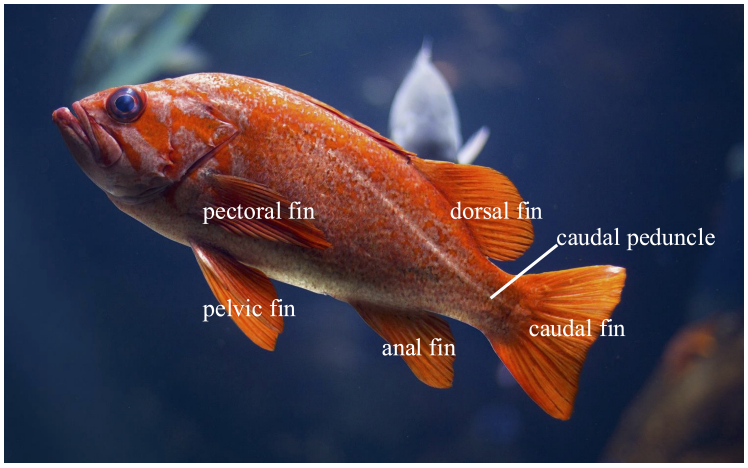

Figure 2.1: Fish with five different types of fins and the caudal peduncle labeled. The pelvic and pectoral fins are paired, while the caudal, anal, and dorsal fin are median. Original photo by Claudio Guglieri on Unsplash.

Despite not using precise kinematic analysis, Breder's categories have persisted for BCF swimming and are widely used in modern day, with some modifications [150].

2.3 Body and/or Caudal Fin (BCF)

Fish using BCF propulsion generate a backwards-moving propulsive wave to create thrust. This propulsive wave has wavelength between approximately half to one and a half times the body length during steady swimming. The body wave speed tends to be constant [154], but the amplitude may change along the body [87].

Fish who primarily swim using their body and caudal fin (BCF) are categorized into five subtypes that describe how the propulsive wave affects the fish. These subtypes span from anguilliform, where the entire body is affected by the propulsive wave, to ostraciiform, where the body is rigid and only the tail oscillates. These propulsive modes tend to move along a continuum from undulation to oscillation (i.e., increasing wavelength) as swimming speed, and therefore Reynolds number, increases [130]. The amplitude of the wave at the head also tends to decrease from anguilliform to ostraciiform [146].

2.3.1 Anguilliform

Anguilliform propulsion lies at the most undulatory end of the spectrum of modes, taking its name from the common eel *Anguilla anguilla*. In this mode, the entire body from snout to tail undulates with the propulsive wave.

Maneuverability and efficiency are enhanced in the anguilliform mode. Eels and

other anguilliform fish have been observed swimming backwards by reversing the propulsive wave [87]. Eels have also been shown to swim four to six times more efficiently than subcarangiform swimmers [148].

Anguilliform motion produces disconnected vortex rings in the wake. Two rings are shed per tail beat, one to each side of the eel. During steady motion, the ring axis is perpendicular to the path of mean motion. When accelerating, the vortex rings are reoriented slightly backward. The rings convect both rearwards and laterally in the wake, dispersing from the path of mean motion. At faster swimming speeds, the flow through the vortices is stronger [106, 147].

Fish, snakes, sharks and salamanders who commonly use the anguilliform mode for hydrodynamic propulsion include eels (*Anguilla anguilla*, [52], *Anguilla rostrata*, [147]), lampreys (ex. *Petromyzon marinus* and *Lampetra fluviatilis*, [130]) sea snakes (*Pelamis platurus*, [51]), colubrid snakes (*Pantherophis/Elaphe guttatus* and *Nerodia fasciata pictiventris*, [68]), catsharks (ex. *Figaro boardmani*, [140]), and sirenid salamanders (ex. *Pseudobranchus striatus* and *Siren lacertina*, [49]).

2.3.2 Subcarangiform

Subcarangiform is best described as the midpoint between anguilliform and carangiform propulsion, where the anterior undulations that are present in anguilliform propulsion begin to diminish in amplitude. Significant undulation only exists in the posterior half of the body [130].

Fish who are considered subcarangiform include rainbow trout (*Oncorhynchus mykiss*, [146]), swimming bluefish (*Pomotamus saltatrix*, [64]), and striped bass (*Morone saxatilis*, [64]). Some sharks are also subcarangiform swimmers, including the velvet belly lantern shark (*Etmopterus spinax*, [47]), the smallspotted catshark (*Scyliorhynus canicula*, [47]), and the blackmouth catshark (*Galeus melanostomus*, [47]).

2.3.3 Carangiform

In the carangiform mode, anterior undulation amplitude is minimal, and significant undulation only occurs in the last third of the body [130]. Due to their rigidity, carangiform swimmers can generally achieve a higher top speed than subcarangiform or anguilliform swimmers. However, this improvement in top speed is balanced by a reduction in acceleration and turning abilities [130].

Carangiform motion produces a chain of connected vortex rings in the wake parallel

to the mean direction of motion. Two rings are produced per tail beat, connecting to the chain. With increased swimming speed, the size, strength, and momentum angle of the vortex rings are all altered. The size and strength both increase, while the momentum angle is reduced [106].

Fish who fall under the carangiform category include the bluegill sunfish (*Lepomis macrochirus*, [146]), cod (ex. *Gadus morhua*, [89]), and carp (ex. *Carassius auratus*, [89]). Some species of shark also display carangiform motion, like the sandbar shark (*Carcharhinus plumbeus*, [140]), common thresher (*Alopias vulpinus*, [140]), hooktooth dogfish (*Aculeola nigra*, [140]), and spiny dogfish (*Squalus acanthias*, [140]).

2.3.4 Thunniform

Thunniform swimmers have the highest top speeds of all swimmers, mostly due to a shift in hydrodynamics. Over 90% of the thrust produced by thunniform swimmers derives from the caudal fin alone [130]. The thunniform mode is predominantly lift-based, as opposed to the added mass effects that dominate anguilliform, subcarangiform, and carangiform modes. These hydrodynamic effects will be discussed in more detail in section 2.4. Due to their reliance on the lift-based mode, thunniform swimmers often display a characteristic lunate tail. The central portion of a swept back caudal fin is empty in lunate tails (figure 2.2). The vortex sheet produced by oscillation fills the physical gap in the fin and acts as a continuous fin from the lift-producing perspective whilst minimizing surface drag [62].

Thunniform motion produces a sickle-like wake structure where the leading edge vortices are shed every half-stroke into the wake [163]. This causes formation of a reverse von Karman street, enabling a thrust-producing jet to form between the vortices. In tuna, the vortices are oriented such that the longitudinal component of the thrust-producing jet is minimized. While the shape of the wake is similar to the carangiform mode, the angular spread of the wake is narrower and the longitudinal component lessened [171].

The thunniform mode includes types of fish, sharks, and marine mammals. This includes the tuna (ex. *Thunnus thynnus*, [130]), mackerel (ex. *Scomber scombrus*, [130]), shortfin mako (*Isurus oxyrhynchus*, [47]), longfin mako (*Isurus paucus*, [8]), white shark (*Carcharodon carcharias*, [8]), porbeagle shark (*Lamna nasus*, [8]), and salmon shark (*Lamna ditropis*, [8]).

Figure 2.2: Example of fish with lunate tails. Photo by Nikola Bačanek on Unsplash.

2.3.5 Ostraciiform

Ostraciiform propulsion is at the far end of the spectrum from undulatory to oscillatory. In this mode, the fish body is inflexible, caudal peduncle abbreviated, and the caudal fin takes an extreme angle without significant curvature [15]. Although the classic example of the ostraciiform mode is the boxfish (ex. *Ostracion meleagris*), boxfishes mainly use MPF propulsion with ostraciiform motion occurring only in burst-and-coast behavior at the highest attained speeds [64]. This author has not been able to find examples of non-*Ostraciidae* ostraciiforms, nor fish that use ostraciiform propulsion as their main mode of swimming.

The ostraciiform mode has received more attention from the robotics field than biologists, due not only to its minimal biological relevance, but also to its ease of mechanical implementation. Ostraciiform propulsion requires less complex mechanical actuation than more undulatory modes, and it can be closely approximated with inflexible heaving and pitching plates.

2.4 BCF Hydrodynamics

Adult fish tend to swim at Reynolds numbers above 1,000, which is considered to be in the inertial flow regime [106]. In this regime, fish take advantage of two distinct hydrodynamic phenomena to enable BCF propulsion. These are added mass (drag-based) and lift-based (or vorticity-driven) phenomena.

2.4.1 Added Mass

Added mass effects result from the acceleration of a volume of fluid that is displaced by a moving body. It is termed "added mass" because not only must one accelerate the mass of the body itself, but also the added mass of the displaced volume of fluid, to enable motion.

To describe how added mass enables fish swimming via a propulsive wave, we consider a fish body of uniform depth broken into infinitesimal segments along its length, called propulsive elements. To best illustrate the effect of the propulsive wave, we consider a body approximately half a wavelength long with constant amplitude undulation (Figure 2.3). As a result of the propulsive wave passing through the body, each propulsive element is accelerated laterally. The lateral acceleration of the propulsive element displaces a volume of fluid that, prior to the lateral motion, was coincident with the fish's surface. This fluid contributes added mass to the system, because the acceleration of the propulsive element forcibly displaces it. The added mass effect causes high form drag on the propulsive element, thereby creating a reaction force that is perpendicular to the orientation of the propulsive element [130].

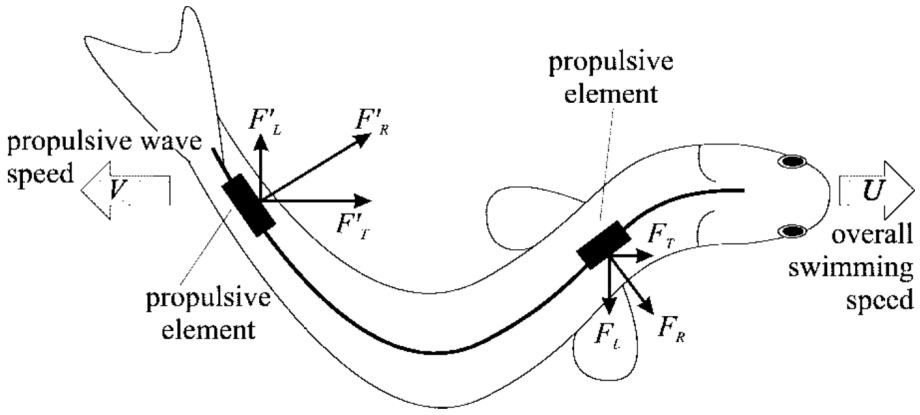

Figure 2.3: Diagram of the forces created by a propulsive wave harnessing added mass effects. ©1999 IEEE. Reprinted, with permission, from [130].

Added mass effects are primarily associated with the anguilliform, subcarangiform, and carangiform modes [130]. Subcarangiform swimmers represent the midpoint on a continuum between anguilliform and carangiform swimmers, and therefore incorporate aspects of added mass effects from both anguilliform and carangiform

modes. We thus discuss the added mass effects specific to anguilliform and carangiform modes.

Anguilliform swimmers most effectively take advantage of added mass effects. First, due to their kinematics they are less prone to sideslip and yaw. Having at least one wavelength present in the body effectively cancels lateral added mass forces [130]. Eels also show particularly high added mass coefficients, with Tytell measuring an overall added mass coefficient of 2.8 ± 0.6 (in comparison to 1.0 for a rigid cylinder accelerating perpendicular to its axis) [147].

Carangiform swimmers, on the other hand, display kinematics more prone to sideslip and yaw. There is significantly less than one waveform present, meaning that lateral forces are not effectively canceled [130]. To minimize the effects of such forces, carangiform swimmers display three distinct body shape adaptations, two of which are driven by added mass effects.

The first adaptation is a distinct caudal peduncle, or narrowing of the body depth, at the body region where the propulsive wave increases in amplitude most rapidly. Assuming a constant body depth (such as in the anguilliform mode), the largest lateral forces will be contributed when the body motion is oriented laterally. For carangiform swimmers, who move with less than one waveform, the body motion is most lateral during wave amplitude increase. To reduce the volume of fluid that must be accelerated by this lateral motion, and therefore the lateral force produced by added mass effects, the body depth during waveform amplitude increase is narrowed. This narrowing of body depth prior to the caudal fin is called the caudal peduncle.

The second is a increase in body depth prior to the caudal peduncle. Increasing anterior body depth increases the lateral forces that would be required to cause sideslip or yaw of the anterior body. This is because the increase in depth increases the added mass of the water that would need to be accelerated to yaw the anterior body. It effectively increases the resistance of the fish to the lateral force effects [87].

The third adaptation is to concentrate the carangiform swimmer's mass in the anterior body, leaving minimal mass in the oscillating caudal fin. This similarly increases the lateral forces that would be required to move the anterior body, since the mass of the anterior body is increased [130].

Added mass effects are more commonly seen in swimmers that employ lower speeds. To enable extraordinary speeds, lift-based propulsion is needed.

2.4.2 Lift-Based, or Vorticity-Driven

Lift-based phenomena are the greatest contributors to thrust in the thunniform mode, but have also been shown to contribute to the carangiform and subcarangiform modes [130]. Thunniform swimmers utilize their caudal fin similarly to a vertically oriented heaving and pitching hydrofoil. If we imagine oncoming flow from the right and viewing a cross-section of the fin, the fin attains its maximum stroke angle at positive and negative vertical extents. When beginning a fin stroke from the bottom maximum stroke angle, the motion is predominantly upward heaving. During this heave, a leading edge vortex begins to form underneath the leading edge. As the fin translates towards the midline, it pitches downwards such that zero pitch is reached at the midline. By this point, the leading edge vortex has begun to shed. After the midline, while heaving upwards towards the alternate maximum stroke angle, the fin continues to pitch down until at a significantly negative angle of attack. This fully sheds the leading edge vortex into the wake. These motions are repeated inversely to return to the bottom maximum stroke angle [4]. The leading edge vortex enhances lift on the caudal fin, of which a significant component contributes to the fish's thrust. The remaining component of lift causes a yaw or sideslip force that is effectively canceled by the mirrored caudal fin motion in the second half of the stroke.

An added benefit of the lift-based mode is that the motion of fins upstream of the caudal fin may enhance thrust production through wake interaction. Drucker and Lauder found evidence to suggest that the vortices shed from a tuna's dorsal fin positively enhance both the free-stream flow velocity and circulation around the caudal fin [33]. Further numerical modeling and experiments have confirmed that the leading-edge vortex on the tuna's caudal fin is enhanced and stabilized by dorsal fin interaction [169, 170].

Lift-based phenomena may also contribute to stability and/or maneuvering in thunniform and carangiform swimmers. Scombrid fish like tunas and mackerel (thunniform) extend their pectoral fins, acting as hydrofoils to counteract their negative buoyancy. To maintain neutral buoyancy, high speed swimming is required to produce lift [94]. The spiny dogfish, a carangiform shark [140], was observed to be incapable of swimming at a positive body pitch after pectoral fin amputation. Instead, it swam downwards continuously. Such an observation suggests that pectoral fins in carangiform sharks similarly act as lifting surfaces to maintain neutral body orientation [158]. Furthermore, when observed free-swimming in a large

tank, the spiny dogfish as well as the sandbar shark (another carangiform swimmer) maintained their pectoral fins at a mean positive angle of attack throughout motion [38].

2.5 Median and/or Paired Fins (MPF)

MPF propulsion occurs when the main contribution to thrust is derived from the motion of a median and/or paired fin. MPF propulsion is seen with four fin engagement patterns: pectoral, dorsal, anal, and combined anal and dorsal. Similar to BCF propulsion, MPF propulsion also exists on a spectrum from undulatory to oscillatory, with some engagement patterns being predominantly undulatory and others including oscillatory modes. Dorsal MPF (amiiform) and anal MPF (gymnotiform) only display undulatory modes due to those fins' wide bases precluding efficient oscillation. Pectoral fin MPF includes both the undulatory rajiform and diodontiform and the oscillatory labriform modes. Combined anal and dorsal MPF also includes both an undulatory (balistiform) and an oscillatory (tetraodontiform) mode [130]. These modes, however, are less commonly referred to in the literature. Here we follow the trend in the literature by setting aside the undulatory MPF modes, as they tend to predominantly provide low-speed propulsion [130], and instead focus on distinguishing lift-based from drag-based MPF oscillatory propulsion.

In the following subsections we also note other animals (like penguins, turtles, and otters) who implement lift- or drag-based aquatic propulsion with their appendages. It is important to state that these animals are not covered under the definition of MPF propulsion as it applies to fish; this is simply an example of evolutionary convergence and homoplasy [40]. These animals are referenced here to provide examples more familiar to the reader.

2.5.1 Lift-Based MPF

In lift-based MPF, the fins act as hydrofoils during forward movement. On the upstroke, the fin takes a negative angle of attack such that its lift is oriented down and forwards in the body frame of reference. On the downstroke, the fin takes a positive angle of attack such that its lift is oriented up and forwards. Through one complete cycle of the fin, negligible lift is generated in the body frame, as the upstroke and downstroke cancel each other. Thrust, however, is generated in both halves of the fin cycle [25]. Propulsive efficiency of the lift-based mode is greater than 80% [40].

Pectoral fin lift-based locomotion is found in many penguins (*Aptenodytes forsteri*,

Aptenodytes patagouica, *Spheniscus demersus*, *Eudyptes chrysolophus*, *Pygoscelis adeliae*, *Eudyptes crestatus*, *Eudyptula minor*, [25]) and fish such as the surf perch (*Cymatogaster aggregata*, [25]) surgeon fish (ex. *Paracanthurus hepatus* [153]), parrot fish (ex. *Chlorurus capistratoides*, [153]), and some wrasse (*Gomphosus varius*, [152, 153]).

2.5.2 Drag-Based MPF

Drag-based MPF uses the fins as paddling or rowing appendages. Paddling appendages are oriented in the vertical parasagittal plane, and rowing appendages are oriented in the horizontal plane. In contrast to lift-based MPF, thrust is only produced during half the fin cycle. This half is called the power stroke, with the non-thrust-producing half called the recovery stroke. At the middle of the power stroke, the fin normal is nearly parallel to the thrust direction. This presents the largest surface area for the posterior sweep of the fin and ensures that the resulting drag is oriented forwards in the body frame of reference. For the recovery stroke, the fin normal is nearly perpendicular to the thrust direction. The fin traverses back to its origin with a near-zero angle of attack, minimizing drag. The drag-based mode is much less efficient than the lift-based mode, scoring at 33%.

Despite the inefficiency of drag-based MPF compared to lift-based, it is still commonly seen among animals as varied as fish, turtles, and otters. It often is used at the air-water interface due to wave drag augmenting propulsion [40]. Furthermore, drag-based MPF creates higher short-term forces than lift-based MPF. This makes drag-based MPF advantageous for maneuvering behaviors [153]. Because of this, drag-based MPF is more often seen in fish who inhabit a single location (ex. coral reefs) while lift-based MPF appears more commonly in migratory species.

Drag-based appendage locomotion is found in a very wide variety of animals, from aquatic to semiaquatic to terrestrial, including the sergeant major damselfish (*Abudefduf saxatilis*, [15]), the three-spined stickleback (*Gasterosteus aculeatus*, [153]), burrfish (*Chilomycterus schoepfi*, [34]), pufferfish (ex. *Arothron hispidus*, [34]), freshwater turtles (*Mauremys caspica*, [153]), muskrats (*Ondatra zibethicus*, [40]), sea otters (*Enhydra lutis*, [40]), water opossums (*Chironectes minimus*, [40]) and even platypi (*Ornithorhynchus anatinus*, [40]).

Chapter 3

INSECT FLIGHT

Examining the kinematics and aerodynamics of insect flight provides insight to key features relevant to micro-air vehicles employing flapping propulsion. The kinematics reveal characteristic shapes of the lift and drag coefficients, which form a base that may be altered through various unsteady mechanisms. These unsteady phenomena play a crucial role in enabling insect flight, and thus are highly relevant to insect flight's robotic counterpart.

3.1 Trajectory Characteristics

Although individual kinematics differ across the wide variety of insects, with modifications to enable a range of $10 < Re < 10^5$ [126], some trajectory characteristics appear frequently. When considered in a body-centered frame, the wingtip traces a closed loop in space which we call a trajectory. This trajectory is defined by the stroke amplitude and stroke deviation. The stroke amplitude is the longest extent of the stroke, and the stroke deviation describes any deviation from a straight-line trajectory. Stroke deviations that are elliptical, U-shaped, or figure-of-eight have been described in nature [22, 36]. If the trajectory is elliptical, then the stroke amplitude is the semimajor axis while the stroke deviation is the semiminor axis.

While actuating this trajectory, the angle of attack of the wing changes. At some point between the middle of the upstroke and the middle of the downstroke, the wing pronates (the leading edge rotates downwards). Between the downstroke and the upstroke, the wing supinates (the leading edge rotates upwards). This ensures that, regardless of upstroke or downstroke, the leading edge remains the same [156].

The trajectory is not all actuated at the same velocity. Rather, the downstroke is sometimes elongated with respect to the upstroke. Furthermore, the change in angular velocity of the wing (change in angle of attack) may also not be constant. The angle of attack may change more slowly during the downstroke, likely to take advantage of its elongation [156]. Several wing trajectories from Ellington [36] are reproduced with permissions in Figure 3.1, to provide examples of the trajectory shape in different insects.

These kinematics create a characteristic shape of the lift and drag coefficients, C_L

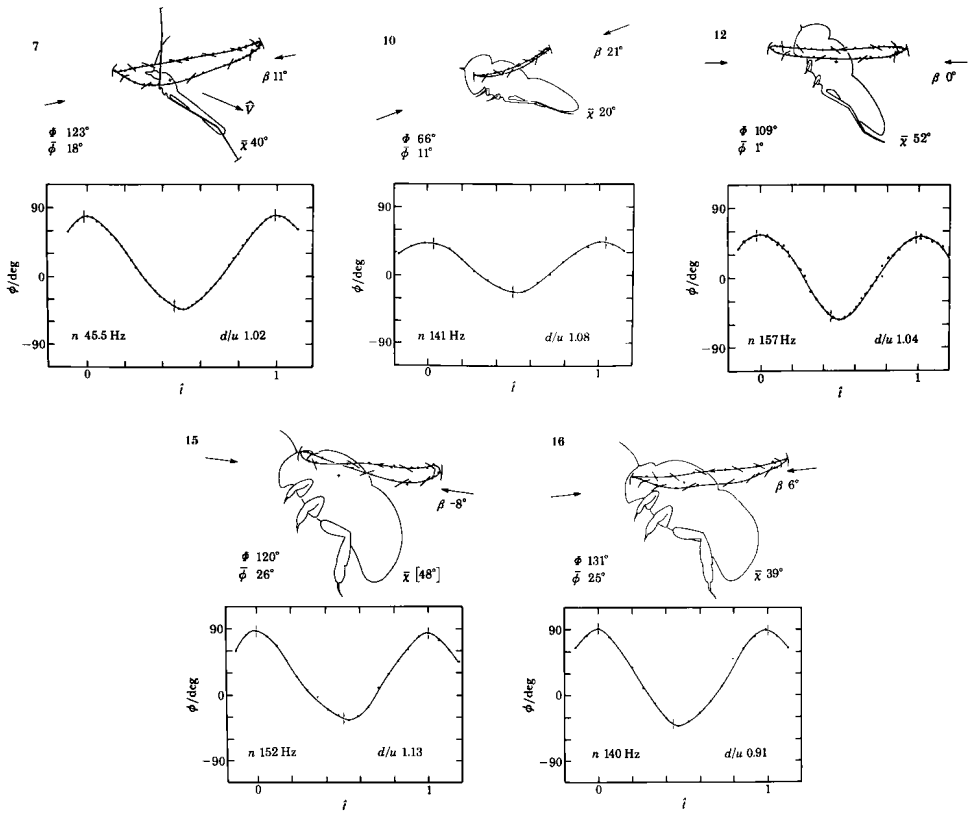

Figure 3.1: Five example wing trajectories, reproduced with permissions from [36]. Insects are, from left to right: (top row) cranefly *Tipula obsoleta*, hover-fly *Episyrphus balteatus*, drone-fly *Eristalis tenax*; (bottom row) garden bumblebee *Bombus hortorum* and white-tailed bumblebee *Bombus lucorum*.

and C_D, over a wingbeat (Figure 3.2). At the beginning of the upstroke, we see a spike in C_L to high values. From the end of that spike, when C_L has returned to lower values, until near the end of the upstroke, C_L grows at a slow rate. This near-plateau is called the "wake valley". Near the end of the upstroke, a second spike occurs to an even higher value of C_L. This peak is followed by a sharp dip into negative C_L. This pattern is qualitatively repeated for the downstroke as well as for C_D, with one exception. At the end of the downstroke, C_L does not achieve a second spike [31]. The characteristic shape of the lift and drag curves cannot be explained wholly by quasi-steady analysis, instead necessitating understanding of unsteady aerodynamics.

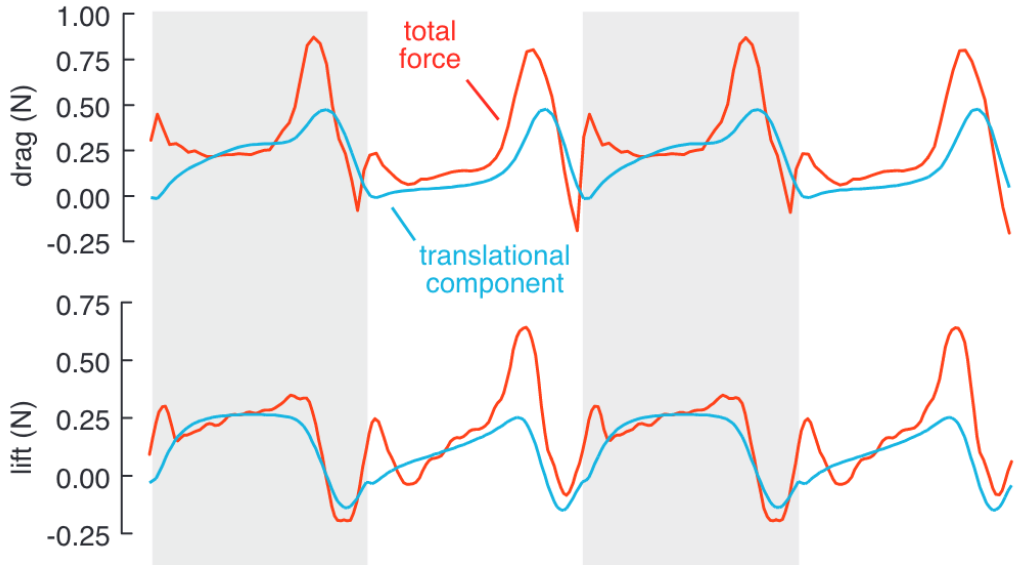

Figure 3.2: Example lift and drag for a fruit fly *Drosophila*. Gray-shaded regions are downstrokes, and white regions upstrokes. From [31]. Reprinted with permission from AAAS.

3.2 Unsteady Aerodynamics

Insect flight is dominated by unsteady aerodynamics. Although earlier research often utilized a quasi-steady approach, Ellington [35] noted that in some species the maximum lift calculated under the quasi-steady assumption was exceeded by the observed average lift. He utilized this "proof-by-contradiction" to argue that a quasi-steady analysis is insufficient to explain the observed aerodynamic forces created by insect flapping. Further studies such as Zanker and Götz [168], comparing experimentally measured lift with lift calculated by quasi-steady theory, also corroborated Ellington's argument. Insect flapping takes advantage of unsteady aerodynamics to enable both weight support and maneuverability.

Unsteady phenomena have been described in the literature for a wide variety of insects at different scales. Over the last few decades, these phenomena have been categorized into five distinct aerodynamic mechanisms: clap and fling, delayed stall, rotational circulation, added mass, and wing-wake interaction.

3.2.1 Clap and Fling, or the Weis-Fogh Mechanism

Clap and fling is a movement pattern most frequently seen in the smallest insects, whose Reynolds number is on the order of ten [103]. These insects include a variety of small beetles (*Coccinella septempunctata*, [165]), flies (*Trialeurodes*

vaporium, [156]; *Drosophila melanogaster*, [36]), thrips (*Thrips pysapus*, [36]), and damselflies (*Calopteryx splendens*, [125]). Some larger species such as butterflies and moths (Re $\approx$ 5000 [11]) have also been observed to use a clap and fling motion [69, 95] albeit intermittently and incompletely [36, 137]. The clap and fling motion is also referred to as the Weis-Fogh mechanism, since Weis-Fogh was the first to observe this phenomenon in parasitic wasps (*Encarsia formosa*) [157].

The clap and fling nomenclature refers to a distinct motion pattern that occurs in the transition from the upstroke to the downstroke. Figure 3.3 shows this motion pattern for the peacock butterfly *Inachis io* alongside fluid motion and the presumed force. As the wings come together at the end of the upstroke, the "clap" is initiated. The leading edges of the wing touch first, and the trailing edges connect last. This narrowing space between the wings forces fluid outward in a jet, resulting in an unsteady force in the direction of the jet. After clapping, the wings "fling" outward to begin the downstroke. The fling motion occurs in two distinct stages. First, the wings rotate about the trailing edge, such that the leading edge is first to open while the trailing edge is still connected. Then, the wings translate apart. The influx of fluid to the space between the wings during the first stage enhances the circulation developed around the wings in the second stage. In addition, the mirrored motion of the two wings negates the formation of trailing-edge vortices, allowing circulation to rapidly increase. By increasing the wings' bound circulation, the flinging motion also produces large unsteady lift.

Slight variations of this clap and fling motion have developed their own nomenclature. During the fling phase, the wings rotate rigidly with no additional curvature. If the wings develop additional curvature during this phase, the motion is then called clap and peel. If the wings come into contact partially, but not along the entire chord, then partial is prefixed. If the wings do not touch at all, then near is prefixed. Partial and near clap and peel are most commonly seen in butterflies, while true clap and fling is restricted to the smallest insects [36]. At the Reynolds numbers of butterflies and moths, the clap and fling motion is less energetically efficient. Instead, it is posited that butterflies and moths occasionally incorporate partial and near clap and peel movement patterns to enable erratic flight for predator evasion [69]. Indeed, in forward steady flight clap and fling are not observed in butterflies [137].

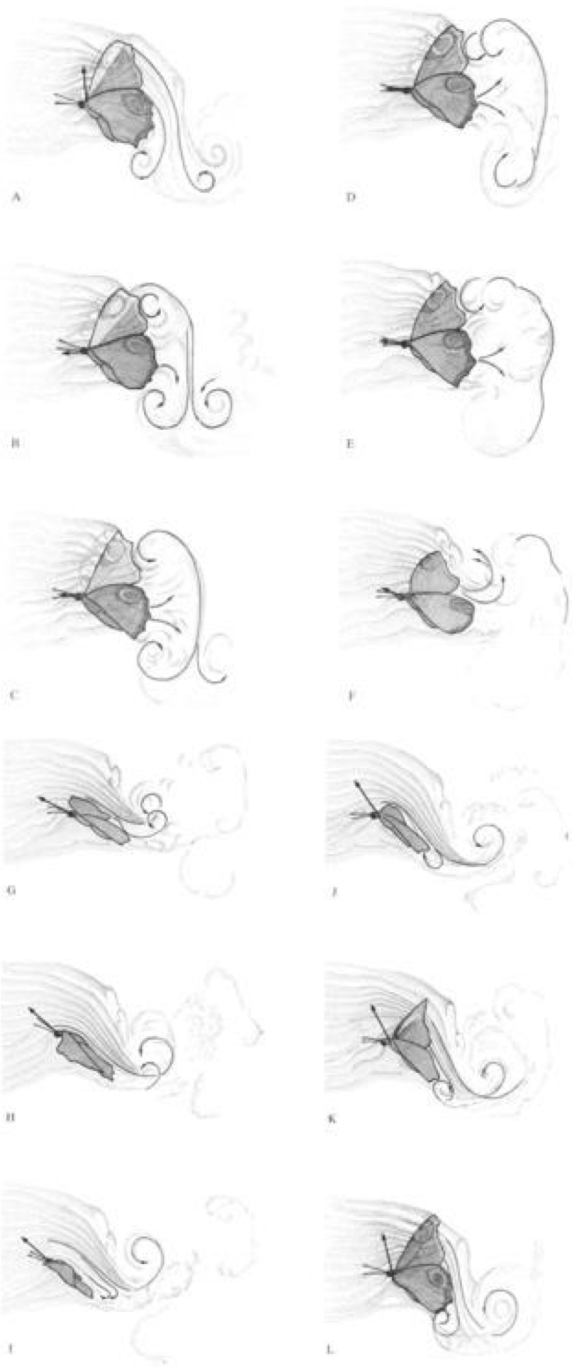

Figure 3.3: Stills of a peacock butterfly *Inachis io* drawn with fluid motion (thin arrows) and presumed force (thick arrows). A-E are pronation, E-H are the down-stroke, I,J are supination, and J-L are the upstroke. Reprinted with permissions from [17].

3.2.2 Delayed Stall

Delayed stall describes the unsteady fluid motion that occurs prior to a wing stalling, which, albeit transiently, drastically increases the lift coefficient [126]. As the angle of attack of an airfoil increases, flow separation occurs. There exists an intermediate range where flow separation at the leading edge exists, but is not paired with the trailing-edge vortex that is a hallmark of true stall. The leading edge vortex is instead contained by the separation bubble, and flow at the trailing edge is smooth. The presence of the leading-edge vortex induces a suction force on the airfoil, which enhances lift.

This leading-edge vortex (LEV), without its trailing-edge pair, also appears during the transition to stall. When an airfoil translates at its stall angle, the process of stalling is not instantaneous. For several chord lengths worth of wing translation, the LEV grows but flow has not yet separated at the trailing edge. The fluid dynamic mechanism maintaining the stability of the LEV has yet to be pinned down. At large Re, on the order 5000, spiraling axial flow convects energy to the tip vortex where it is shed to the wake— similar to the spiral burst vortices seen on delta wings. At intermediate Re, on the order 1000, spanwise flow is at its maximum at the vortex core. At low Re, on the order 100, spanwise flow is no longer spiraling and no longer confined to the vortex core. Instead, spanwise flow occurs nearly across the entire wing [11]. It is thought that these spanwise flow patterns maintain stability of the LEV by balancing the generation of vorticity at the leading edge with convection into the wake from the tip vortex [11, 164]. However, blocking such spanwise flow does not destabilize the LEV [9]. An alternate theory posits that quasi-steady centripetal and coriolis accelerations are stabilize the LEV, especially at $Re > 100-200$ [21, 86]. Note that the spanwise flow itself is driven by centripetal and coriolis accelerations, which may muddy the distinction [86]. It is only when the LEV destabilizes, sheds, and the trailing-edge vortex also sheds that an airfoil dramatically loses lift and stalls. The enhanced lift that occurs due to the stable and attached growth of the LEV is therefore called delayed stall.

Delayed stall was first identified in insects by Maxworthy [99], who posited that the "fling" phase of clap and fling employed this mechanism. Indeed, butterflies in free flight maintain an attached LEV that forms upon wing separation at the beginning of the downstroke [137]. However, it has been extended to many small insects who do not employ clap and fling such as dragonflies (*Libella luctuosa*, [136]), hawkmoths (*Manduca sexta*, [14, 160]), and even butterflies when not using clap and fling

wingbeats (*Inachis io*, [17]).

3.2.3 Rotational Circulation and the Kramer Effect

In order to maintain a positive angle of attack during both the downstroke and the upstroke, insect wings must pronate between the upstroke and the downstroke and supinate between the downstroke and the upstroke. This rapid rotation of the wing about its spanwise axis while also translating in space causes the Kramer effect [77], also called rotational lift. If the wing is rotating about a spanwise axis located away from mid-chord, a rotational drag term is also nonzero [82].

The Kramer effect occurs because of the non-zero time it takes for the Kutta condition to be enforced. When an airfoil rotates while translating, the stagnation region separates from the trailing edge. To resist this separation, and thereby attempt to reestablish the Kutta condition, a rotational circulation is imparted to the fluid in the opposite sense to the wing rotation. Until the Kutta condition is reestablished, rotational circulation will increase proportional to the angular velocity of the wing. This additional circulation causes a force on the wing [126].

The Kramer effect causes forces during supination and pronation, which may contribute to or detract from lift and drag based on timing of rotation [126]. There is evidence that insects modify rotation timing to enable control [23, 30, 31]. If the wing supinates prior to stroke reversal, lift is enhanced. If the wing supinates after stroke reversal has already occurred, the Kramer effect produces negative lift. Symmetric rotation causes a spike in positive lift prior to stroke reversal and a spike in negative lift after stroke reversal [24]. When mimicking measured fruit fly kinematics on a robotic apparatus, Dickinson et al. found that the Kramer effect accounted for 35% of the total lift produced [31].

Rotational drag can also be a key factor in producing lift for some insects. Notably, rotational drag scales with the square of the angular velocity of the wing and does not scale with translational velocity [82]. Therefore, insects with shorter wingbeat amplitudes like mosquitos tend to take advantage of rotational drag. When mimicking measured mosquito (*Culex quinquefasciatus*) kinematics using CFD, Bomphrey et al. found that rotational drag was crucial to weight support [13].

3.2.4 Added Mass

Added mass is perhaps the least clear of the unsteady mechanisms contributing to insect flight. This is due to the timing of its effects within the stroke; the majority

of added mass effects occur at the initiation of each half-stroke, because that is when acceleration is largest. Thus, its effects co-occur with initiation of the LEV, rotational circulation due to pronation/supination, and wake capture [127], all of which may themselves interact with and modify the added mass force via induced flow. Added mass studies have mainly focused on hovering or slow flight, with fast flight receiving less research [24]. Furthermore, the studies that do exist contradict each other in their findings. With experimental reproductions of flapping behavior, the added mass benefit is negligible compared to other wake capture effects for low Re, on the order 100 [127]. However, CFD modeling at the same Re, which explicitly separated wake capture effects, showed that the peak in lift was only due to added mass [83, 142]. More modern modeling includes the added mass term, but still makes assumptions on its direction [32]. CFD modeling at higher Re, on the order 1000, also showed that the lift peak produced during the initial wing acceleration phase is dominated by added-mass effects [116]. Models of flexible wings, incorporating FSI effects, showed that average lift and angle of attack at midstroke are governed by added mass effects [71]. There remains a contradiction between some experiments that show negligible added mass benefit, but many simulations that show its dominant effects.

3.2.5 Wing-Wake Interaction

When a wingbeat traverses through a region of fluid that has been altered from its free-stream characteristics by the previous wingbeat, it is called wing-wake interaction. There exists a wide spectrum of possible wing-wake interactions, which we classify into three categories. These categories are: induced jet, interaction with vortices, and interaction with downwash.

Induced jet effects were originally discovered in 2D, where the reverse von Karman vortex street formed after stall creates a strong induced jet towards the wing during stroke reversal [29]. When the wing is accelerating and the induced jet is convecting flow past the wing faster than the free-stream, this gives an added mass contribution to the forces on the wing [137]. Although the induced jet is still present in 3D [81], the induced jet magnitude has been found to be lower in 3D versus 2D. In 3D, tip vortices are shed with lateral velocity and do not form a reverse von Karman vortex street [145]. Furthermore, depending on the timing of rotation, the induced jet impinging on the surface of the wing that was on top during the immediately preceding half-stroke may incur a positive or negative spike in lift [31]. A second spike due to LEV suction may be slightly offset, depending on timing [93].

Interaction with vortices from the previous wingbeat can occur at varying points during the stroke. Traversing the wing through a previously shed vortex may allow recapture of circulation, allowing the LEV to grow more easily. Second, by recapturing this circulation, the shed vortex will reduce its velocity. This reduces the energy lost to the wake moreso than the momentum, due to the squared velocity in the kinetic energy [137]. However, vortex interactions may also have a detrimental effect on lift. If the wing encounters a vortex such that the vortex is between the wing and its counterrotating pair, vortex suction will occur, negatively impacting lift [92]. Furthermore, it is thought that vortex interaction may contribute to the "wake valley". After an initial increase in C_L, towards mid-stroke C_L becomes negative. In 2D, encountering the LEV from the previous downstroke destabilizes the upstroke LEV and results in a vortex pair being shed to the wake. Although this has been proposed as the mechanism behind the "wake valley", 3D effects cause the previous LEV to weaken and LEV destabilization does not occur [81].

The wing may interact with downwash if it traverses through the center of a vortex ring formed by the shed tip vortices, starting vortex, and LEV. This downwash reduces the wing's effective angle of attack, thereby reducing mean lift by 6-18% [145, 162]. This downwash was shown to specifically affect mid-stroke lift in experiment [10], and may therefore explain the "wake valley". However, downwash in the wake also prolongs the attachment of the LEV by reducing the effective AOA. This leads to increased lift overall when balancing the benefit from maintaining a close LEV with the downwash-induced contribution to the wake valley [81].

Different insects are affected by different combinations of the five aforementioned aerodynamic mechanisms: clap and fling, delayed stall, rotational circulation, added mass, and wing-wake interaction. Some insects like butterflies (*Vanessa atalanta*) may even choose among these mechanisms by altering wing kinematics to enable high maneuverability [137]. The unsteady phenomena may be beneficial, like delayed stall increasing the lift coefficient, or detrimental, like vortex suction reducing lift. Regardless, these mechanisms are crucial to explaining how insects fly.

Chapter 4

EXPERIMENTAL SETUP

Sections 4.1-4.3 reproduced with permission from

[1] M. L. Hooper, I. Scherl, and M. Gharib. "Bio-inspired compensatory strategies for damage to flapping robotic propulsors". In: *Journal of The Royal Society Interface* 21.216 (July 2024), p. 20240141. ISSN: 1742-5662. DOI: 10.1098/rsif.2024.0141.

4.1 Mechanical Set-Up

This work modifies the fin size, shape and flexibility of a previous experimental set-up [97]. A 200 × 50 mm flexible fin made of 0.02" thick Delrin (ρ = 1420 kgm^{-3}, E = 3.3 GPa) is attached to the end of a 3d-printed polylactic acid (PLA, ρ = 1300 kgm^{-3}, E = 3.5 GPa) rod. This fin is intended to be a simple model robotic flapping propulsor, incorporating an aspect ratio and level of flexibility inspired by fish caudal fins. The rod is actuated by a spherical parallel manipulator (SPM, designed by Steiner & Steiner [138, 139]) that moves the fin in a user-defined trajectory (Figure 4.1). Owing to its three rotational degrees of freedom, the SPM can achieve both infinite rotation about the axis of the rod and large-angle deflections along both x- and y-axes of the laboratory frame, with accuracy within 1°. The fin is submerged within a 1.0 × 2.4 × 1.2 m tank containing Chevron Superla white oil (ρ = 880 kgm^{-3}, ν = 115 cSt). Oil, as opposed to water or air, is used in this study to increase the signal-to-noise ratio, and results in a range of Reynolds number (Re) based on fin span and average velocity per cycle of Re = 444 − 954 for all optima found in the experiment. Upon removal of approximately 50% of the propulsive area, Re increases from 593 to 750 − 846 for thrust production and from 444 to 546 − 954 for side-force production. These Reynolds numbers are of the same order of magnitude, so we hypothesize that there is little change in the flow regime. For comparison in the realm of thrust production, *Spinibarbus sinensis*, *Cyprinius carpio* and *Carassius auratus* have Reynolds numbers of 535, 633 and 1264, respectively, when estimated using trajectory and fin width measurements [43]. For comparison in the realm of side force production, the Reynolds number of *Drosophila hydei* was estimated to be 221 using measurements of trajectory parameters [105] and scaling up the wing chord of *Drosophila melanogaster* [144]. Note that wing measurements

used in our estimate were scaled by the body length ratio of *D. melanogaster* to *D. hydei* [96, 113].

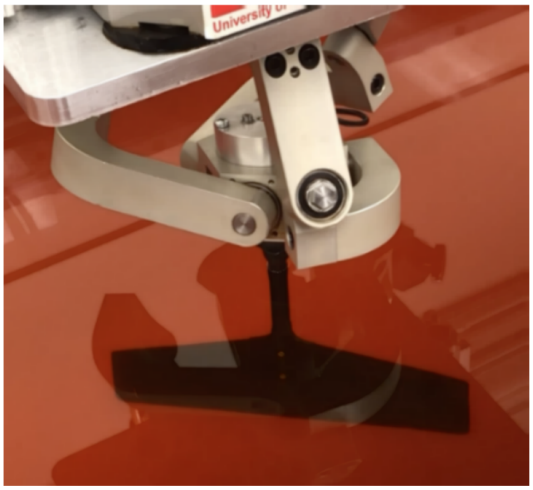

Figure 4.1: Fin attached to the spherical parallel manipulator inside the oil tank.

While moving through the working fluid, the mounted fin experiences forces and moments that are measured by a six-axis ATI Nano25 Force/Torque sensor. The rod extends 225 mm to avoid possible surface effects, and the SPM is mounted equidistant from three walls of the tank to minimize wall effects.

4.2 Parametrization of Motion

The motion of the SPM is constrained to a workspace inspired by natural fin and wingtip trajectories developed by Martin & Gharib [97]. The workspace is parametrized by the 10 variables listed in table 4.1. Since robotic propulsors can access a larger geometric space than natural swimmers and flyers, the ranges of these variables exceed what is seen in nature. We choose not to artificially constrain the robotic system to the biological range, since doing so would be unrepresentative of the capabilities of robotic propulsors in practice.

The type of trajectory (C) defines whether the base trajectory is elliptical or a figure-eight. Although some insects display figure-eight wingtip trajectories, previous work has shown that the optimization quickly disregards figure-eight trajectories [97]. For this work, we fix the type of trajectory to an ellipse, allowing nine parameters to participate in the optimization. Three parameters, the stroke angle (ϕ), thickness angle (ψ) and camber (λ), define the geometry of the trajectory. ϕ is the angle actuated by the SPM at the furthest x-excursion of the trajectory; in other words, it is a parametrization of the semi-major axis of the ellipse. ψ is the angle

	Symbol	Minimum	Maximum	Convergence Criterion	Thrust Initialization	Side-Force Initialization
Type	C	Ellipse (0)	Figure-Eight (1)	0.5	Ellipse (0)	Ellipse (0)
Stroke Angle	ϕ	15.27°	32.18°	3°	25.43°	26.12°
Thickness Angle	ψ	0°	15.27°	3°	14.29°	13.30°
Rotation Angle	χ	-70°	70°	3°	-40.46°	-62.51°
Rotation Phase	β	0	$2\pi - 1$	0.4	6.18	2.63
Speed Code	S	0	4.9	0.9	2.95	1.34
Speed-Up Value	γ	1.1	1.3	0.1	1.10	1.18
Rotation Acceleration	K_v	0	1	0.2	0.13	0.03
Camber	λ	0	1	0.2	0.20	0.18
Frequency	ω	0.7	0.9	0.01	0.71	0.72

Table 4.1: Trajectory parameters' range, convergence criteria, and initializations.

actuated by the SPM at the furthest y-excursion of the trajectory; in other words, it is a parametrization of the semi-minor axis of the ellipse. λ defines the degree of curvature of the major axis of the ellipse. Three parameters define the angle of the fin with respect to the trajectory: the rotation angle (χ), rotation phase (β) and rotation acceleration (K_v). χ is the maximum angle that the fin makes along the trajectory with a line perpendicular to the local instantaneous slope of the ellipse. β defines at which azimuthal position that maximum fin angle occurs. K_v parametrizes how quickly the fin angle changes from its maximum to minimum values. The higher the value of K_v, the closer to a square wave the fin's angular trace appears. Lastly, three parameters define the speed at which the trajectory is actuated. The trajectory is not all actuated at a constant velocity —half of the trajectory experiences a relative speed-up. This half of the trajectory is defined by the floor of the speed code (S). S equaling 1 corresponds to the section I–II–III, 2 to II–III–IV, 3 to III–IV–I and 4 to IV–I–II. The relative speed-up of this section is dictated by the speed-up value (γ), where for example a value of 1.1 would correspond to a 10% increase in speed. The frequency (ω) dictates the average speed at which a full cycle of the trajectory is executed. Selected parameters are visualized in Figure 4.2. These parameters are converted to time-resolved positions of the fin, which are executed by the SPM [97]. To best analyse fin motion, we use the angle of attack (AOA). This angle is defined azimuthally.

4.3 Artificial Evolution Strategy

The experimental apparatus and trajectory parametrization described above are paired with CMA-ES as the learning algorithm for this work [57–59, 112]. We selected the class of evolutionary algorithms for several reasons. First, they are bio-inspired: artificial evolution strategies are designed to mimic natural evolution. This allows for a closer comparison between the results of machine learning and the

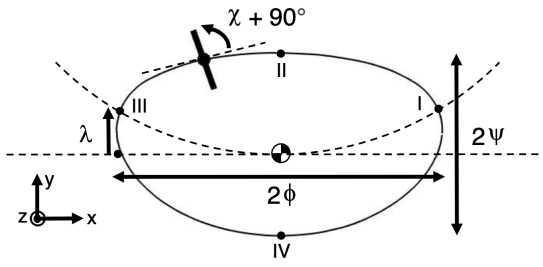

Figure 4.2: Visual representation of selected trajectory parameters.

results of natural evolution. Second, evolutionary algorithms have been successfully used across a wide variety of problems [46, 48, 72, 97, 101, 102, 115, 128, 135, 143]. Finally, CMA-ES converged more quickly than other algorithms while maintaining accuracy in our experimental system.

Here we provide a simplified process overview of the CMA-ES applied to our system. First, a set of candidate solutions (trajectories), called a 'generation', are generated from a multivariate normal distribution. The individual real-world fitnesses of this group of trajectories are assessed using the experimental set-up. Every 10 generations of an optimization takes approximately 13 hours to complete, owing to the in-the-loop evaluations of the fitness function. Each trajectory is actuated at least three times to achieve consistent results. Furthermore, the first three cycles of each trajectory are removed from the data, so as to measure only the steady-state force production. The experimentally determined fitness values of each trajectory in the current generation determine the multivariate normal distribution from which the next generation of trajectories is drawn, since these fitness values alter the mean and covariance matrix that define the probability distribution. Once a new generation has been generated from the updated probability distribution, the experimental evaluations begin again. Generations are created and evaluated until the algorithm reaches all convergence criteria. The convergence criterion for each trajectory parameter is reported in table 4.1.

CMA-ES attempts to find the global minimum of the fitness function it is given. In this work, the fitness of each trajectory is calculated by equation 4.1, where F is either the force in the z-direction (if optimizing for thrust) or the magnitude of the resultant force in the x–y plane (if optimizing for side force). For these experiments, a target force F_{target} of 1 N was defined. A trajectory with parameters near the

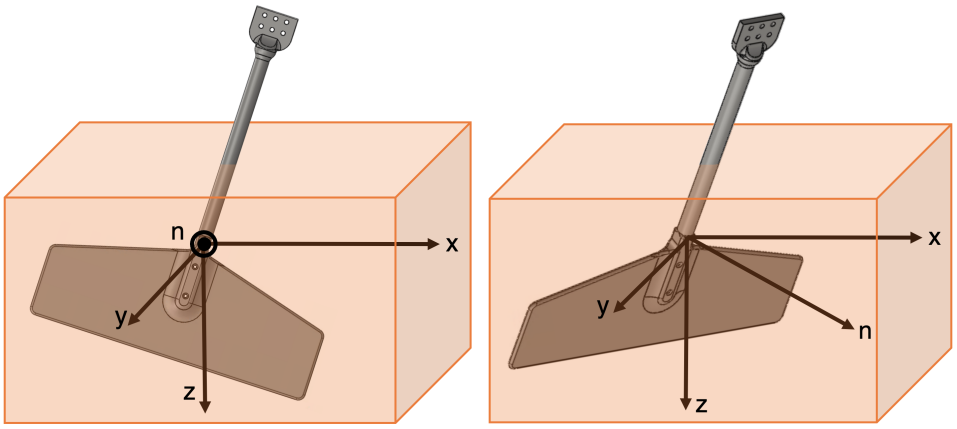

Figure 4.3: Coordinate system of the SPM.

middle of the search space can produce 1 N of thrust.

$$f = 0.8\frac{|F - F_{target}|}{F_{target}} + 0.2|1 - |\frac{F}{F_n}|| \tag{4.1}$$

The first term in equation 4.1 enforces the closeness of the produced force (F) to the target force (F_{target}) while the second term is a measure of geometric efficiency. Geometric efficiency quantifies the portion of the force produced normal to the fin ('normal force', F_n, Figure 4.3) that is in the desired direction of force production. The prefactors on each term, 0.8 and 0.2, respectively, were chosen and verified via a sensitivity analysis in prior work [97]. These prefactors remain valid because the fitness function terms have consistent definitions between intact and amputated propulsors. This fitness function will optimize for a trajectory that produces the desired amount of force in the desired direction, with high geometric efficiency. As in prior work [97], the lowest-fitness trajectory in the final generation is taken as the optimum. This is common practice among evolutionary algorithms, since the optimum over all generations is sensitive to real-world experimental noise.

In both our experiments optimizing for side force and for thrust, we introduce fin breakage during the optimization. An optimization is initiated with a full fin (Figure 1d). We initialize CMA-ES with trajectory parameter values (table 1) that we found optimal from a previous optimization run with the same fitness function for an inflexible fin (200 × 50 × 3 mm Markforged Onyx, ρ = 1200 kgm^{-3}, E = 2.4 GPa). After the actuation and evaluation of 69 generations, the optimization state is copied. The full fin then finishes optimization, which occurs 10 generations later for thrust, and 11 generations later for side force. This produces two optimization states: (i) the copied full-fin state that is close to convergence and (ii) the optimized full-fin state.

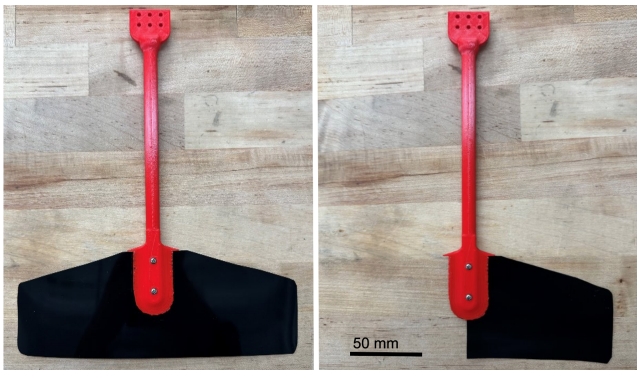

Figure 4.4: Intact and amputated flexible fins.

To compare the optimal trajectory for an unbroken fin with the learning algorithm's adaptation to a broken fin, we then resume multiple optimizations at Generation 70 from the copied full-fin state. Each of these resumed optimizations uses a fin where the left half has been removed (Figure 4.4). The half that remains is 180°from the defined AOA, as measured azimuthally. The mounting point of the rod to the fin is not altered, resulting in a total reduction in propulsive area of 44.2%.

4.4 Mathematical Treatment of CMA-ES

The following equations are taken from [57]. CMA-ES functions in a cycle of three steps until convergence. In the first step, λ candidate solutions $\boldsymbol{x}_k^{(g)}$ are sampled from a normal distribution. λ is the population size, k denotes the individual candidate solution from $k = 1, ..., \lambda$, and g is the number of the current generation. Each generation is one pass through the three-step cycle. We sample each $\boldsymbol{x}_k^{(g)}$ from

$$\boldsymbol{x}_k^{(g)} = \boldsymbol{m}^{(g-1)} + \sigma^{(g-1)} \mathcal{N}(0, \boldsymbol{C}^{(g-1)}), \tag{4.2}$$

where $\boldsymbol{m}^{(g-1)}$ is the mean, calculated from the previous generation $g-1$, $\sigma^{(g-1)}$ is the step size calculated from the previous generation, and $\boldsymbol{C}^{(g-1)}$ is the covariance matrix calculated from the previous generation. If $g = 1$, the mean, step size, and covariance are taken from initialization parameters.

The second step is to evaluate the objective function, or fitness function, on all $\boldsymbol{x}_k^{(g)}$. The fitness function f acting on $\boldsymbol{x}_k^{(g)}$ produces

$$\boldsymbol{y}_k^{(g)} = f(\boldsymbol{x}_k^{(g)}) \tag{4.3}$$

and we rank all resulting $\boldsymbol{y}_k^{(g)}$ such that

$$f(\boldsymbol{x}_{1:\lambda}^{(g)}) = \boldsymbol{y}_{1:\lambda}^{(g)} \leq ... \leq \boldsymbol{y}_{\lambda:\lambda}^{(g)}. \tag{4.4}$$

The third step uses the μ best candidate solutions, called the parent number, to recalculate and update parameters, thereby shifting the probability distribution towards lower values of $\boldsymbol{y}_k^{(g)}$. We update the mean of the probability distribution using equation 4.5:

$$\boldsymbol{m}^{(g)} = \boldsymbol{m}^{(g-1)} + c_m \sum_{i=1}^{\mu} w_i (\boldsymbol{x}_{i:\lambda}^{(g)} - \boldsymbol{m}^{(g-1)}), \tag{4.5}$$

where c_m is the mean learning rate and w_i are the recombination weights.

In order to update the step size, we need to calculate an evolution path $\boldsymbol{p}_\sigma$. The evolution path length is affected by a sum of successive steps, and is therefore referred to as cumulative step length adaptation (CSA). If all successive steps result in a path longer than that expected by random selection, the evolution path length will increase. If all successive steps result in a shorter path than that expected by random selection, then the evolution path length will decrease. If all successive steps reflect the expected path under random selection, then the evolution path length is unchanged.

$$\boldsymbol{p}_\sigma^{(g)} = (1 - c_\sigma)\boldsymbol{p}_\sigma^{(g-1)} + \boldsymbol{C}^{(g-1)^{-\frac{1}{2}}} \left(\frac{\boldsymbol{m}^{(g)} - \boldsymbol{m}^{(g-1)}}{c_m \sigma^{(g-1)}} \right) \sqrt{c_\sigma (2 - c_\sigma) \mu_{eff}} \tag{4.6}$$

We now are able to update the step size using equation 4.7:

$$\sigma^{(g)} = \sigma^{(g-1)} \exp\left(\frac{c_\sigma}{d_\sigma} \left(\frac{\|\boldsymbol{p}_\sigma^{(g)}\|}{\mathbb{E}\|\mathcal{N}(\boldsymbol{0}, \boldsymbol{I})\|} - 1 \right) \right). \tag{4.7}$$

Similarly, the covariance matrix has an evolution path $\boldsymbol{p}_c$ to utilize cumulation. We calculate $\boldsymbol{p}_c$ by

$$\boldsymbol{p}_c^{(g)} = (1 - c_c)\boldsymbol{p}_c^{(g-1)} + h_\sigma \sqrt{c_c (2 - c_c) \mu_{eff}} \sum_{i=1}^{\mu} w_i \boldsymbol{y}_{i:\lambda}, \tag{4.8}$$

where

$$h_\sigma = \begin{cases} 1 & \text{if } \frac{\|\boldsymbol{p}_\sigma\|}{\sqrt{1-(1-\sigma)^{2g}}} < \left(1.4 + \frac{2}{n+1}\right) \mathbb{E}\|\mathcal{N}(\boldsymbol{0}, \boldsymbol{I})\| \\ 0 & \text{otherwise.} \end{cases} \tag{4.9}$$

Now knowing the covariance matrix evolution path, we may update the covariance matrix using equation 4.10:

$$\boldsymbol{C}^{(g)} = \left(1 - c_1 - c_\mu \sum_{i=1}^{\lambda} w_i \right) \boldsymbol{C}^{(g-1)} + c_1 \boldsymbol{p}_c^{(g)} \boldsymbol{p}_c^{(g)\top} + c_\mu \sum_{i=1}^{\mu} w_i \boldsymbol{y}_{i:\lambda}^{(g)} \boldsymbol{y}_{i:\lambda}^{(g)\top}. \tag{4.10}$$

All parameters involved in the CMA-ES algorithm have suggested values based on the search space dimension N. These parameters are listed in table 4.2 along with their calculated values for the flapping robot. These are the optimization parameters used in this work.

Variable	Equation	Value	Description
λ	$4 + \lfloor 3 \ln N \rfloor$	10	population size
μ	$\lfloor \frac{\lambda}{2} \rfloor$	5	parent number
μ_{eff}	$\left(\sum_{j=1}^{\mu} w_j^2 \right)^{-1}$	3.1673	variance-effective parent number
c_m	1	1	mean learning rate
c_1	$\frac{2}{(N+1.3)^2+\mu_{eff}}$	0.0183	rank-one update learning rate
c_μ	$\min\left(1 - c_1, \frac{2\left(\mu_{eff}-2+\mu_{eff}^{-1}\right)}{(N+2)^2+\mu_{eff}}\right)$	0.0239	rank-mu update learning rate
c_σ	$\frac{\mu_{eff}+2}{N+\mu_{eff}+3}$	0.3407	step-size cumulation decay rate
c_c	$\frac{4N+\mu_{eff}}{N^2+4N+2\mu_{eff}}$	0.3176	rank-one cumulation decay rate
d_σ	$1 + 2\max\left(0, \sqrt{\frac{\mu_{eff}-1}{N+1}} - 1\right) + c_\sigma$	1.3407	step-size damping parameter
w_j	$\left\lvert \ln\left(\frac{\max\left(\mu, \frac{\lambda}{2}\right)+0.5}{j}\right) \right\rvert$		recombination weights

Table 4.2: CMA-ES parameters, equations, and example values (which correspond to the flapping robot).

4.5 Convergence Time and Accuracy of CMA-ES

To further justify using an algorithm that is intermediate in both complexity and time investment, we compare the CMA-ES algorithm to two other optimization algorithms: steepest descent and a Bayesian optimization.

4.5.1 Steepest Descent

Steepest descent is one of the simplest optimization algorithms. It is presented here as proof that the optimization landscape of the flapping robot is highly complex and needs a more sophisticated optimization strategy. We expect that our optimization, due to evaluating real-world fluid dynamic phenomena, will not have a simple single

minimum. Steepest descent should, therefore, struggle. Because the optimization follows the local gradient, steepest descent cannot find the global minimum of a function with multiple minima unless initial values are chosen with prior knowledge.

We implement the steepest descent algorithm with a BFGS Quasi-Newton method with a cubic line search procedure and tolerances of 1×10^{-6}. All optimizations are started from the thrust initialization values in table 4.1 and optimized for 1N thrust. Since our experiment is subject to real-world noise, we repeat the steepest descent optimization seven times. Steepest descent is highly susceptible to noise. For example, one could imagine a simple one-dimensional optimization where the fitness function takes the shape of a W. Starting slightly to the left or slightly to the right of the central peak will lead to the optimization finding two different solutions.

Each steepest descent optimization converges with between 4 and 10 evaluated trajectories. All but one optimization converge to the 1×10^{-6} step tolerance, with trial 5 converging because the predicted decrease in the fitness function (based on the gradient) is less than the function tolerance. The resulting "optimal" trajectories are reported in table 4.3.

Trial	1	2	3	4	5	6	7	Mode	CMA-ES
C	1	1	1	1	1	1	1	1	1
ϕ	25.3	25.4	25.4	25.3	25.4	25.4	25.4	25.4	24.8
ψ	14.2	14.2	14.2	14.1	14.2	14.2	14.2	14.2	12.9
χ	-40.5	-40.5	-40.5	-40.4	-40.5	-40.5	-40.5	-40.5	26.7
β	6.1	6.1	6.1	6.1	6.1	6.1	6.1	6.1	5.3
S	2	2	2	2	2	2	2	2	0
γ	1.1	1.1	1.1	1.3	1.1	1.1	1.1	1.1	1.2
K_v	0	0	0.1	0	0	0.1	0	0	0.6
λ	0	0.2	0.1	0.5	0	0.2	0.1	0.1	0.6
ω	0.7	0.79	0.71	0.7	0.7	0.7	0.77	0.7	0.7
fitness	0.2375	0.2204	0.1068	0.1204	0.1163	0.1709	0.1212		0.1157

Table 4.3: Steepest descent trial optima, mode of the optima, and comparison to the optimum found by CMA-ES.

From the seven steepest descent trials, we find the mode for each parameter. Comparing this set to the parameters found by CMA-ES, we see that the rotation angle, rotation phase, speed code, rotation acceleration, and camber are all outside of the convergence criterion. The optimum trajectory found by the steepest descent optimization is therefore distinct from the optimum trajectory found by CMA-ES. Furthermore, most of the steepest descent trials attain a fitness higher than the

fitness attained by the CMA-ES optimization. Only trial 3 attains a fitness lower than CMA-ES. However, note that the steepest descent optimizations undergo many fewer experimental actuations than CMA-ES, which means that the error in the fitness function is commensurately higher. CMA-ES takes on average approximately 800 actuations, while steepest descent only took between 4 and 10. We thus see that, while steepest descent very quickly hones in on a local minimum, the global minimum remains mysterious. The lower complexity and lower time investment do not outweigh the loss of accuracy.

4.5.2 Bayesian Optimization

Bayesian optimization is considered the state of the art for optimizing functions that are expensive to evaluate. In our case, the cost incurred by evaluating our experiment is mainly in the time domain. Note that this introduces a tension to the optimization— Bayesian optimization can struggle with computational time when large amounts of data are used due to the assumption of a Gaussian process prior. This may not be as important for experiments that are "expensive" in domains other than time.

Bayesian optimization models the fitness function f as a probability distribution. It considers any $f(\boldsymbol{x})$ that has been explicitly evaluated as an "observed variable", and those $f(\boldsymbol{x})$ still unknown as "hidden variables". Given a set of observed variables, the optimization will choose the next hidden variable to evaluate based on the conditional distribution $\boldsymbol{P}(f(\boldsymbol{x})|f(\boldsymbol{x_1}), f(\boldsymbol{x_2}), ..., f(\boldsymbol{x_n}))$. Thus, each optimization step retains all the information from each observed variable. Here we utilize an expected improvement (EI) model to choose the next hidden variable to evaluate. In an EI model, the next selected hidden variable must minimize the expected future improvement in the minimum $f(\boldsymbol{x})$ after its observation. A graphical example of the Bayesian optimization process is shown in Figure 4.5.

We choose to modify the EI model slightly to minimize the effects of evaluation time and difficult-to-escape local minima. The acquisition function is weighted inversely to the time it would take to evaluate; in other words, we bias the optimization towards time-efficient behavior. Second, we let the EI model modify its behavior when it believes it is overexploiting an area. The EI model raises its variance if it believes it is overexploiting, in order to escape the local minimum. These are both common modifications to a Bayesian optimization to maximise performance.

The Bayesian optimization results, compared to that of CMA-ES, are reported in

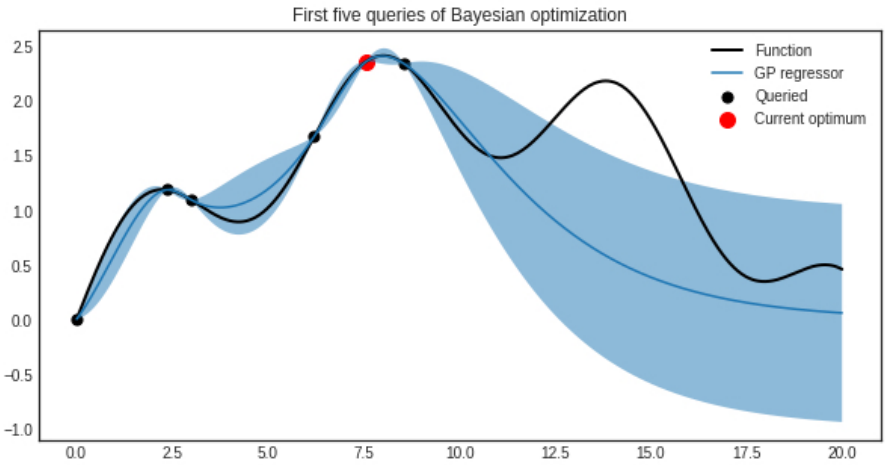

Figure 4.5: Graphical example of Bayesian optimization, reproduced from [28].

table 4.4. ϕ, χ, S, and ω are all significantly different based on their convergence criteria. The optimum Bayesian trajectory attains a lower fitness than the CMA-ES trajectory. We thus see that CMA-ES does not find the global minimum, but rather a local minimum[1]. However, we see that the CMA-ES algorithm ends up at a minimum where the desired force is matched even more closely than the optimum Bayesian trajectory. The desired force is produced to within 0.3% by CMA-ES, but to within 2% by the Bayesian optimization. This means that the small loss of fitness is due to an inefficiency in the trajectory, rather than varying force production. A small inefficiency at this local minimum is acceptable.

Optimization	C	ϕ	ψ	χ	β	S	γ	K_v	λ	ω	Fitness	Closeness to Setpoint
Bayesian	1	15.3	14.5	50.1	3.5	1	1.2	0.6	0.5	0.71	0.0615	-0.0170
CMA-ES	1	24.8	12.9	26.7	5.3	0	1.2	0.6	0.6	0.7	0.1157	0.0029

Table 4.4: Bayesian optimum versus the optimum found by CMA-ES.

The Bayesian optimization ran for six weeks until the fin broke due to fatigue. Convergence was not reached. The six-week runtime is due to the optimization taking approximately an hour to calculate the next point. At the six-week mark, 1,229 evaluations of the objective function had been performed without convergence. For a standard CMA-ES optimization, where no changes are made during a run, approximately 800 evaluations are needed for convergence. These 800 evaluations take place over the span of one week. Given this disparity in optimization time,

[1]This does not mean that the Bayesian optimization has found the global minimum either.

CMA-ES is undoubtedly the more efficient algorithm. Despite not discovering the minimum that Bayesian optimization did, we accept the small reduction in optimal fidelity for the large gains in optimization efficiency. We thus choose to implement the CMA-ES algorithm using MATLAB software written by Hansen [56], with minor modifications.

4.6 Particle Image Velocimetry (PIV)

4.6.1 Imaging Hardware

The PIV experiments use similar optics to the scanning PIV system described in [44] to enable scanning multiple planes in the y-axis. However, our setup does not utilize the galvanometer for laser beam scanning. Instead, the laser and optics are moved laterally by a Velmex Bi-Slide MN10-0200-M01-21 stage. This allows the capturing of quasi-three-dimensional, two-component, phase-averaged data. Unfortunately, several factors precluded Fu et al.'s scanning PIV system from being fully implemented. The scanning PIV system was originally chosen for its ability to return non-phase-averaged, three-dimensional, three-component vector fields. However, the wake was too large and the illumination too weak to properly utilize the technique.

In order to obtain three velocity components, a minimum sheet overlap (laser sheet thickness over step size) of 3.33 is required and 5 is preferred. Due to the power of the laser and its attenuation in the mineral oil, camera exposures below 7000 μs struggle at the required depth of field. In order to achieve an exposure of 7000 μs alongside the minimum sheet overlap, the laser sheet must scan very slowly through the volume compared to the fin's motion. Assuming the minimum three steps per laser sheet thickness, the fastest fin (0.9 Hz) will have moved approximately 7° during the scanning of one laser sheet thickness. This is too fast; quality PIV correlations for most of the field occur with fin spacings of 2.5° or less. So, we must relinquish our goal of three-component data.

Nominally, Fu et al.'s setup should be able to scan the full wake of the fin for a quasi-three-dimensional, two-component measurement (given the previously discussed restrictions on three-component data). However, the scan volume requires use of nearly the entire diameter of the condenser lens. Near the edge, the beam does not emerge normal to the flat surface of the lens. It instead is slightly tilted inwards. This angle becomes significant by the time the laser sheet reaches the desired measurement volume (away from the tank walls). Thus, Fu et al.'s setup cannot

scan the full wake of the fin either. Due to the illumination precluding three-component data, and our desired measurement volume exceeding the bounds of the optics, we chose to instead fix the optical path and move the entire laser-optics setup using a stage.

A Laserglow LRD-0655 laser, model PFR-01000-05 LabSpec, with wavelength 655nm and output power 1200mW was used for illumination. This laser beam impinges upon a condenser lens. The beam then passes through a sheet-forming optic, which is a 25cm diameter glass cylinder. The sheet-forming optic causes the beam to expand in the x-z plane. This laser sheet illuminates silver-coated glass beads (Potters Industries Conduct-o-Fil, AG-SL150-16-TRD, 0.900 g/cc average density, 100μm average diameter) that are suspended in and move with the experimental fluid. The reflection from these particles is captured by an IDT XSM-3250 with frame rate of 148 times the flapping frequency f_{fin} at 1440 by 2650 pixels. This FPS range, of approximately 104 to 133 Hz, results in fin motion of approximately 2.5° between frames and allows for an exposure of 7000 μs. The camera lens is a Nikon AF FX Nikkor with a focal length of 20mm and maximum aperture of $f/2.8$.

4.6.2 Calibration

4.6.2.1 Optics

The optics need to be calibrated such that the laser plane is parallel to the x-z plane and therefore the scanning direction is parallel to the y-axis of the spherical parallel manipulator (Figure 4.6). To calibrate the scanning optics, we take advantage of attaching a calibration plate to the spherical parallel manipulator. The calibration target consists of a symmetric array of circles.

To determine whether the laser plane is parallel to the x-z plane, we implement a two-step calibration procedure. First, the laser sheet is adjusted to impinge upon the calibration target positioned at $(\alpha, \beta, \gamma) = (90, 0, 0)°$. The stage is adjusted such that the laser sheet bisects a vertical line of circles in the center of the calibration target. If the laser sheet cannot bisect a vertical line of circles but instead is angled with reference to a vertical line of circles, the sheet-forming optic is adjusted. The angle of the sheet-forming optic with respect to the horizontal affects the inclination of the laser sheet with respect to the vertical.

Once the laser sheet bisects a vertical line of circles at $y = 0$, the calibration target is then swept through $(\alpha, \beta, \gamma) = (90, -45 - 45, 0)°$. This sweeps the target such that a vertical line of circles traverses in parallel to the x-z plane. If the laser sheet

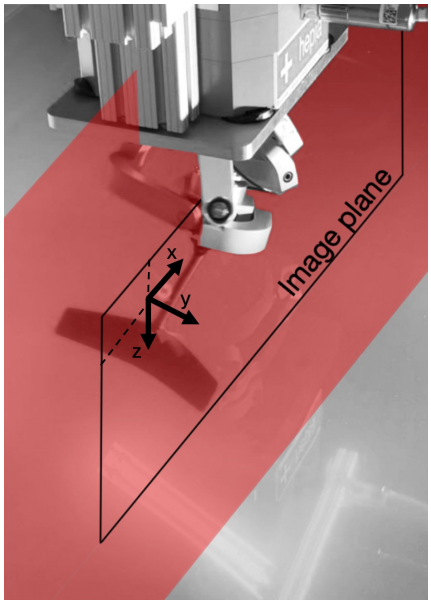

Figure 4.6: Laser plane, reproduced from [63] with permission. Scanning direction is in the y-axis.

does not bisect the vertical line of circles during the full sweep, then the angle of the full optical setup in the x-y plane must be adjusted with respect to the spherical parallel manipulator.

By calibrating the optics such that the laser sheet is parallel to the x-z plane, we inherently constrain the scanning direction to be parallel to the y-axis. To validate this, one can repeat the same calibration procedure on the leftmost and rightmost sides of the calibration target.

4.6.2.2 Image

The image is first squared to the experimental apparatus by aligning a reference line to the calibration plate. The reference line vertically bisects the camera image. The camera is spatially shifted such that this reference line bisects the central line of dots on the calibration plate.

To calibrate the resulting data for the lens distortion of the camera, we utilize the following calibration procedure. A 47 by 41 symmetric array of 2mm diameter circles spaced 5mm apart was printed on transparency paper using a Brother MFC-L2740DW. This laser printer achieves dimensional accuracy of ±0.042 mm. This is less than 1% of the spacing between circles and is equal to 0.24 pixels. Prior to printing the calibration array, the printer itself was calibrated with a 100×100 mm

square. The transparency paper was attached to a white 3D-printed plate, which itself is attached to the spherical parallel manipulator using a stem identical to the fin. This results in the ability to move the calibration array through the same 3D space as the fin.

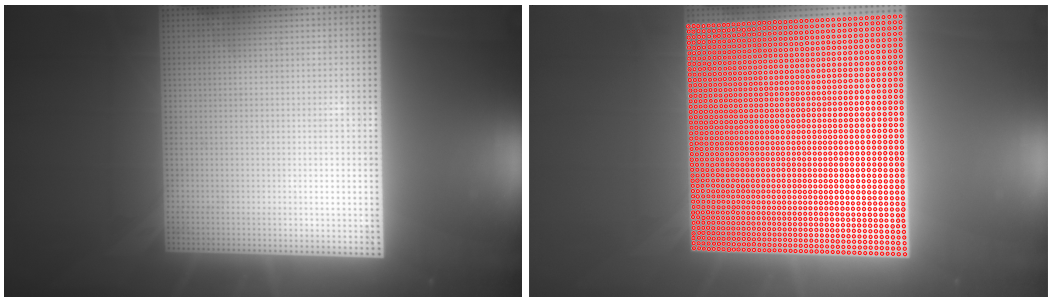

Figure 4.7: Example image from x-z calibration of camera image distortion (left), alongside computer-recognized calibration pattern (right).

We use 8 individual photos of the calibration plate at different angles about the stem to characterize the lens distortion. These angles are [-40, -30, -20, -10, 0, 10, 20, 30] degrees. An example image of the calibration plate at 20 degrees is shown in Figure 4.7. The known locations of the calibration circles in the x-z plane, when rotated at these angles about the z-axis and compared with the imaged locations of the calibration circles, quantify the lens distortion. This comparison is accomplished using MATLAB's camera calibration tools in the Image Processing and Computer Vision toolboxes. The calculated lens distortion can then be applied in the inverse to the PIV data to remove distortion. Figure 4.8 shows the original and undistorted calibration plate; minimal difference is found between the two. The distortion is less than 1.5×10^{-6} pixel at all PIV-interrogated locations, and so the data is processed as received.

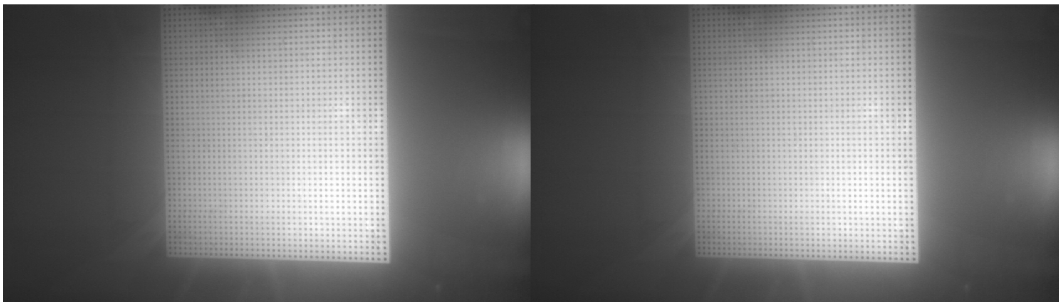

Figure 4.8: Original (left) and undistorted (right) calibration plate image.

This procedure calibrates the image projected on the x-z plane. Calibrating the

y-axis, which is the depth or scanning direction, must be integrated with the stage. To calibrate the y-axis, we position the calibration plate using the SPM at $(\alpha, \beta, \gamma) = (90°, -25°, 0°)$[2]. We then move the stage in increments of 10,000 steps and take images. Using the image locations of the calibration dots and laser sheet, combined with the known real-world measurements of the calibration dots, we can convert our desired real-world y-coordinates into stage positions. Figure 4.9 shows an example image from this calibration sequence.

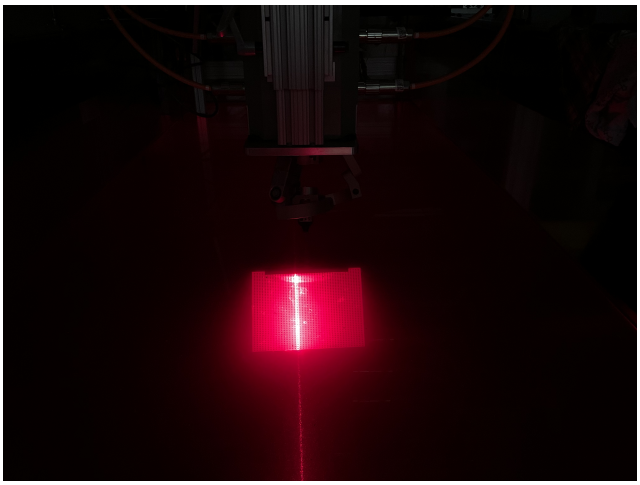

Figure 4.9: Example image from y-calibration, showing the laser intersecting the calibration plate at stage position 10,000.

4.6.3 PIV Software

To process the images obtained by the imaging hardware, PIVView from PIVTEC GmbH was used. The grid refinement was done in three passes of 96x96, 72x72, and 48x48 pixels. A standard FFT correlation with Nyquist frequency filtering and uniform window weighting was used. Multiple correlation took place four times. Sub-pixel peaks were subject to a 5-pixel nonlinear Gaussian fit. 10 validation passes took place, with a normalized median test threshold of 3.0, Z-score test threshold of 6.2, and enforced minimum correlation of 5%. Outlier replacement was with a lower-order peak or re-evaluated with a larger sample. Velocity data was taken from PIVView and processed further. Vorticity was calculated by a fourth-order centered difference scheme.

[2] $\beta = -25°$ is used to maximize the visibility of the laser sheet at the camera location.

Chapter 5

BIO-INSPIRED COMPENSATORY STRATEGIES FOR DAMAGE TO FLAPPING PROPULSORS

This chapter reproduced with permission from

[1] M. L. Hooper, I. Scherl, and M. Gharib. "Bio-inspired compensatory strategies for damage to flapping robotic propulsors". In: *Journal of The Royal Society Interface* 21.216 (July 2024), p. 20240141. ISSN: 1742-5662. DOI: 10.1098/rsif.2024.0141.

5.1 Thrust Production

5.1.1 Fitness and Force Production

Immediately upon amputation of half the fin, the CMA-ES algorithm is exploring a parameter subspace that it has found to be optimal for the whole fin while performing experimental evaluations with an amputated fin. This mismatch, where trajectories that are in the optimal basin for the whole fin are evaluated on the partial fin, causes a distinct increase in the fitness function and a distinct decrease in force produced (Figure 5.1A). A higher fitness value signals worse performance (equation 4.1). After amputation, the CMA-ES algorithm narrows in on a new optimal basin for the amputated half fin.

Remarkably, the fitness and force production of the amputated propulsors optimizing for thrust recover fully by the time of convergence, despite a propulsive area loss of almost 50% (Figure 5.1A). The fitnesses of the amputated fins are all lower than the fitness of the intact fin, and the closeness to setpoint force varies from the intact fin by at most 1.01%. Although the majority of natural swimmers in the literature are not able to maintain their thrust production after catastrophic loss of propulsor area (see appendix A, table A.1), many were subject to the removal of 100% of their propulsive area. Of those fish for whom approximately 50% was removed (*Mylopharyngodon piceus*, *Etheostoma edwini*, *Spinibarbus sinensis* and *Cyprinius carpio*), an equal split is observed between maintenance or decrease of thrust production in response to caudal fin damage. While natural swimmers, depending on the individual species, may or may not maintain efficient force production following a 50% reduction in caudal fin area, our learned approach recovers from such damage.

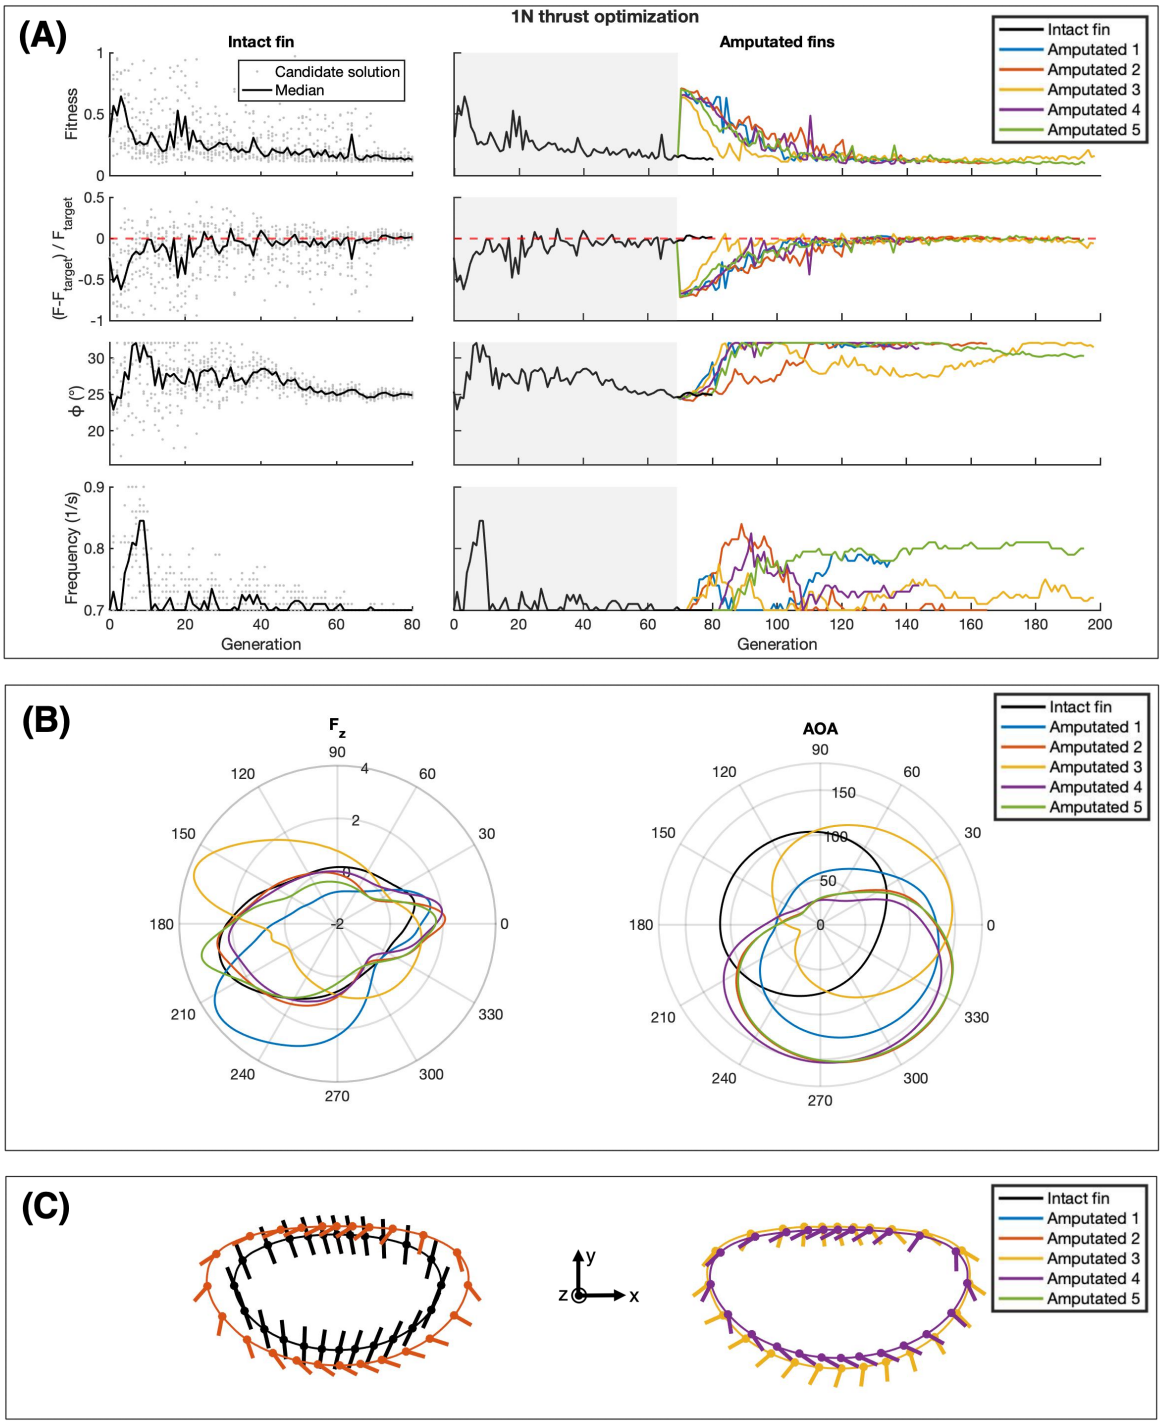

Figure 5.1: Optimization for 1N of thrust force. Caption continued on following page.

Figure 5.1: In subfigure (A), the column titled 'Intact fin' shows the intact fin optimizing for 1N of thrust (z-force). The line tracks median values for each generation, while experimentally determined fitness values for each individual candidate solution are represented by dots. The column titled 'Amputated fins' shows the progression of the learning algorithm for all fins, omitting individual candidate points for clarity. Note that Generations 0–69, shaded in grey, are identical to the intact fin. The optimization was resumed with each of the amputated fins from Generation 70 of the intact fin. Subfigure (B) shows polar plots of the force production and AOA of each trajectory. AOA is measured azimuthally. Subfigure (C) shows the optimal trajectories for thrust production, projected onto the x–y plane. The fin progresses counterclockwise, and the AOA of the fin stem is plotted every 15 azimuthal degrees. The left shows the intact fin trajectory compared with the median amputated fin solution (Amputated 2). The right shows the two most different amputated fin solutions (Amputated 3 and 5), whose spatial bounds encompass all other amputated fin solutions.

The force traces throughout the converged cycle further corroborate that force production has returned to its pre-damaged baseline. The average forces for all trials, including the intact fin, differ from the setpoint of 1N by 1.3% or less. Qualitatively, all force traces are similar (Figure 5.1B) although some phase shifts are seen in Amputated 1 and 3. These can be quantified by a Fourier decomposition (appendix A, table A.3). We see an elongated bilobed shape appear in the force traces, arising from the mechanics of the spherical parallel manipulator. Since the fin stem remains a constant radius and all force is produced normal to the fin planform, the component of normal force projected onto the z-direction is thereby largest at 0° and 180°.

5.1.2 Trajectory Parameters: Amplitude, Frequency and Angle of Attack

To maintain average force production with a smaller propulsive area, it has been posited in theory [119] that both amplitude and frequency must increase. Indeed, most *in vivo* studies of fish have found an increase in amplitude after catastrophic damage. However, *in vivo* studies of fish have not shown a consensus on frequency alteration after fin damage [15, 19, 20, 43, 53, 132, 155] (appendix A, table A.1).

In our experiments, the optimization increases both the converged amplitude and frequency after breakage to maintain efficient thrust production, matching theory [119], but not *in vivo* experiments. The converged amplitudes for the amputated fin are higher than the unbroken optimal amplitude by between 5.2° and 7.2° (Figure 5.1A). This difference is significant, as it is larger than the convergence criterion of 3°. Most of the amputated trials attained a higher converged frequency than that

of the intact fin, ranging from 0.02 to 0.08 Hz greater. This difference is larger than the convergence criterion of 0.01 Hz. However, it should be emphasized that one trial did not display an increase in the optimal frequency of motion (notably, Optimization 2; see appendix A, table A.2).

The maximum AOA is significantly changed in response to damage, both increasing and phase shifting by approximately 110° (Figure 5.1B). A paired visualization of the amplitude and AOA can be seen in Figure 5.1C. While the increase in the maximum AOA is a clear mechanism to provide higher forces, the large phase shift is intriguing. One would expect such a difference to be notable, yet as far as the authors are aware neither an increase nor phase shift of AOA have been reported *in vivo*. Thus, this may be a distinct difference between the biological and machine-learned adaptations.

5.2 Side Force

5.2.1 Fitness and Force Production

When we instead optimize for a force in the x–y plane (side force), we no longer see as straightforward a recovery. The fitness and force production of the side-force trials are worse in all cases and are not able to fully recover in Amputated 1. In contrast to the thrust trials, where all amputated fins displayed lower fitnesses than the intact fin, all amputated fins producing side force displayed higher fitness than the intact fin by 5–19 times. Their closeness to setpoint force was also highly impacted. All amputated trials except for Amputated 1 fall within 2% of the desired target force – a similar range to the thrust amputated trials. However, more of the side-force trials exist near to that 2% maximum. Finally, Amputated 1 is unable to recover its pre-damaged force production, as it converges to a solution with a 17% deficit in thrust.

Before further analysis, we must note that the desired 1N of side force was allowed to take any angle within the x–y plane during the optimization. For this reason, Figure 5.2B displays plots where the resultant forces are aligned with the x*-axis. Figure 5.2C shows the relationship between trajectories in the original x–y reference frame and the trajectories in the x*–y* reference frame. In the x*–y* reference frame, the force and AOA profiles collapse well (Figure 5.2B).

5.2.2 Trajectory Parameters: Amplitude, Frequency and Angle of Attack

The amplitude of motion increases to varying degrees when the fin is optimizing for side-force production and spontaneous amputation occurs. As seen in Figure

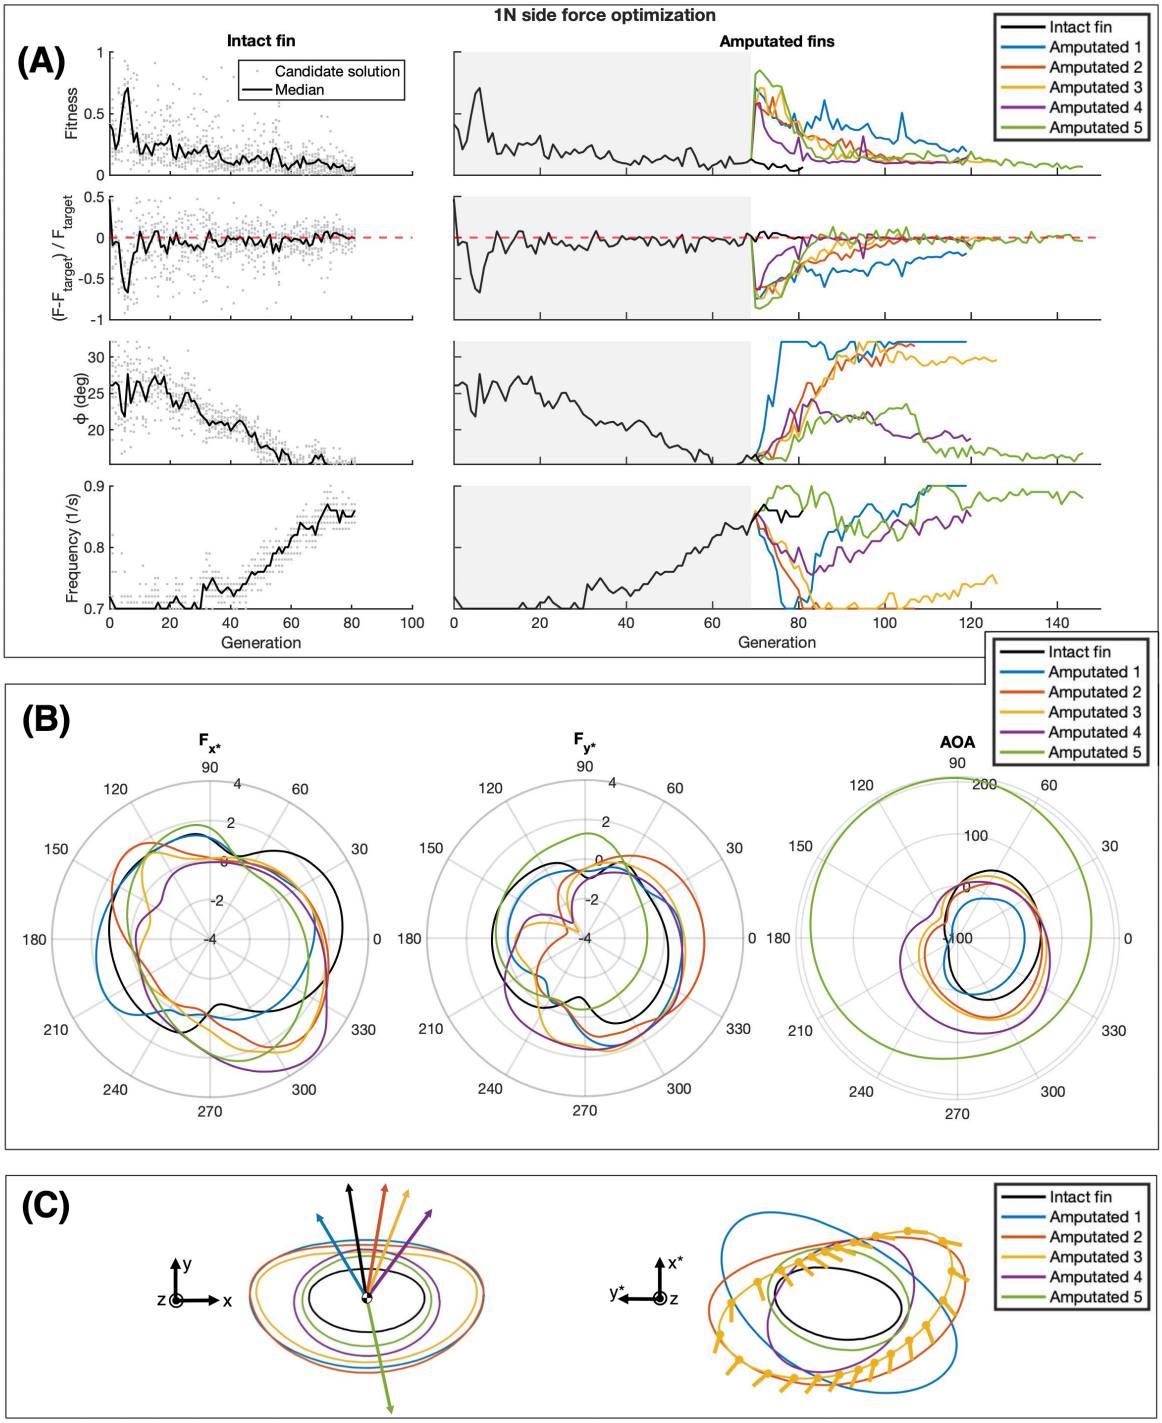

Figure 5.2: Optimization for 1N of side force. Caption continued on the following page.

Figure 5.2: In subfigure (A), the column titled 'Intact fin' shows the intact fin optimizing for 1N of side force (force in x–y plane). The line tracks median values for each generation, while experimentally determined fitness values for each individual candidate solution are represented by dots. The column titled 'Amputated fins' shows the progression of the learning algorithm for all fins, omitting individual candidate points for clarity. Note that Generations 0–69, shaded in grey, are identical to the intact fin. The optimization was resumed with each of the amputated fins from Generation 70 of the intact fin. Subfigure (B) shows polar plots of the force production and AOA of each trajectory. The forces are defined in the x*–y* directions, which are represented in (C). AOA is measured azimuthally. Subfigure (C) shows the optimal trajectories for side force production in two reference frames. The left shows the trajectories with major axes aligned and resultant forces plotted. The right shows the trajectories rotated such that their resultant forces are all in the x* direction. The AOA is plotted for the median amputated fin solution. The fin progresses counterclockwise, and the AOA of the fin stem is plotted every 15 azimuthal degrees.

5.2A, Amputated 1, 2 and 3 display significantly larger converged amplitudes than the intact fin (32.1°, 31.3° and 29.6° versus 15.2°), while Amputated 4 and 5 show a lower increase (18.7° and 16.8°). Since the optimal basin has a 3° amplitude tolerance required for convergence, we can state that Amputated 5's amplitude is within the optimal basin of the intact fin, while Amputated 4's amplitude is barely larger. We thus find that the first three amputated trials support a significant increase in amplitude to recover from a spontaneous amputation, while the last two trials do not. This learned data is consistent with the approximately equal split of data from *in vivo* experiments; for those species in which amplitude was reported, it either increased or did not change.

The effect on frequency is even less distinguishable; two out of five amputated trials achieve a higher frequency than the optimal intact fin, one achieves the same frequency and two out of five attain a lower frequency. However, the *in vivo* literature shows by a large majority that frequency tends to increase when compensating for propulsor damage.

The AOA has a similar phase between the intact fin and the amputated fins when the trajectories are rotated into a common force-defined reference frame (Figure 5.2B). This similarity stands in contrast to the production of thrust (Figure 5.1B), where the dominant Fourier mode of the intact solution displayed a 110.6° phase shift from the average of the amputated solutions' dominant Fourier mode (appendix A, table A.3). Furthermore, all trials, whether intact or amputated, show nesting of

AOA traces (Figure 5.2B). The traces are nested following the order of resultant forces in Figure 5.2C. Progressing clockwise from the resultant for Amputated 1 progresses outwards in the nesting order. The average AOA therefore correlates with the direction of force production. This implies that a simple strategy for controlling the direction of the force produced in the x–y plane would be to add or subtract AOA throughout the cycle.

5.3 Sensitivity Analysis

Examining the sensitivity reveals that the optimization landscape for side force is more challenging to navigate than for thrust. We perform principal component analysis (PCA) on the final covariance matrix calculated by the CMA-ES algorithm, scaled to unit variance. This standardized covariance matrix represents the optimal basin – a hyperellipsoid with axis directions and relative lengths given by PCA. When comparing the normalized scree plots for all thrust trials versus all side-force trials (Figure 5.3) one can see that the side-force trials experience a faster drop-off. The faster drop-off means that the side-force optimal basin has a more elongated shape than the thrust optimal basin, where several axes have roughly equal magnitude. A more elongated optimal basin is more difficult to traverse since more of its axes have a low eigenvalue and are therefore more sensitive.

By finding the radius from the origin to the hyperellipsoid's surface along the direction corresponding to a single trajectory parameter, we translate PCA-based sensitivity into a value for each real-world trajectory parameter. We normalize the calculated radius by the largest eigenvalue. The larger the normalized radius, the less sensitive the trajectory parameter is to perturbation. The normalized radii of the side-force optimization are on average lower than thrust, meaning that the sensitivity to perturbation is higher, and therefore the optimization landscape is more challenging to traverse (Figure 5.3). No other clear trends appear among the data, suggesting that all trajectory parameters are roughly equally involved in the optimization owing to their approximately equal sensitivities.

5.4 Energetics of Compensation

We estimate the energy and hydrodynamic power requirements of the machine-learned trajectory to compare with natural systems. Total work per time step is calculated by computing the dot product of the average propulsor velocity over a time step with the force measured at the end of that time step, from which we then derive the instantaneous power requirements (Figure 5.4).

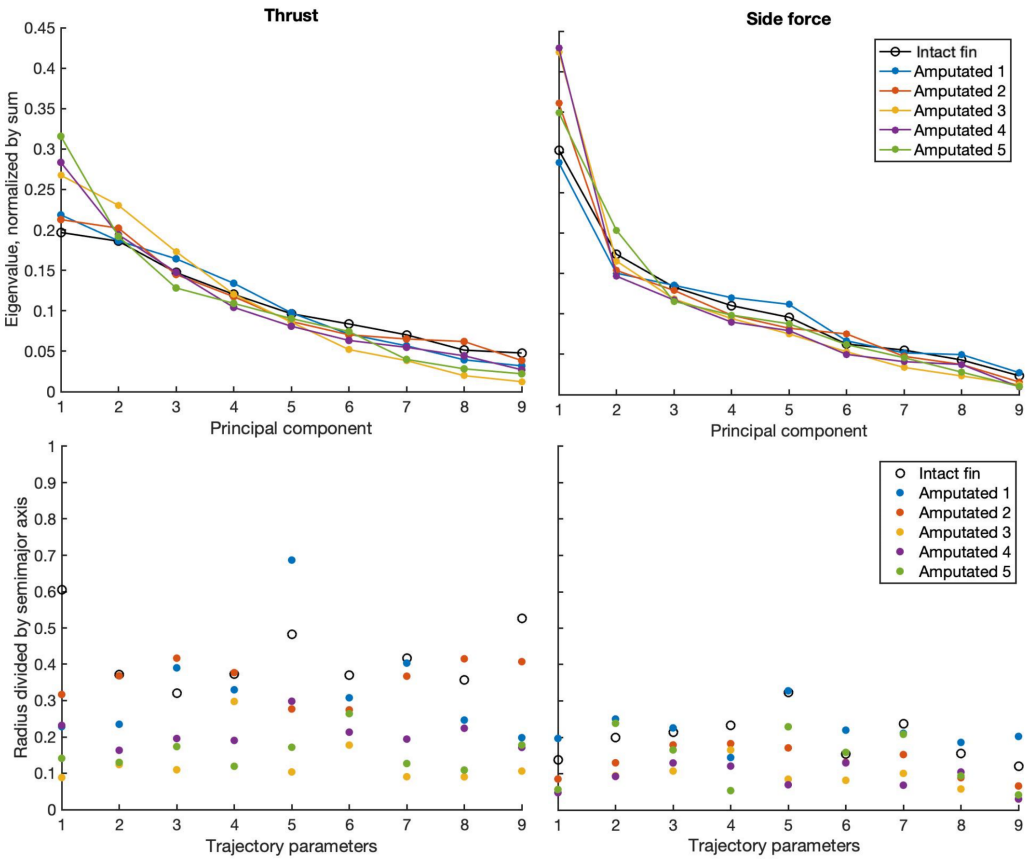

Figure 5.3: Sensitivity analysis via scree plots and hyperellipsoid radii. The left and right columns are for thrust and side force, respectively. The first row displays the PCA scree plots. The second row plots the normalized radius of each trajectory parameter, where a larger normalized radius corresponds to a lower sensitivity. In order from 1 to 9, the trajectory parameters are ϕ, ψ, χ, β, S, γ, K_v, λ and ω.

All amputated thrust trials display a higher instantaneous power requirement than the intact fin, although the average power across the trajectory is similar to within 20W. The distinction among side-force trajectories is less obvious. All trials except Amputated 5 show higher instantaneous power requirements. Furthermore, the average power across all side-force trajectories differs by almost 150W.

To compare the power requirements of robotic optimal trajectories with biological power requirements across varying working fluids, we introduce a non-dimensional force-power ratio ζ defined as

$$\zeta = \frac{F\nu}{PL} \tag{5.1}$$

using the power P, force F, kinematic viscosity ν, and a characteristic length scale L taken to be the chord. The larger the value of ζ, the more efficient the motion.

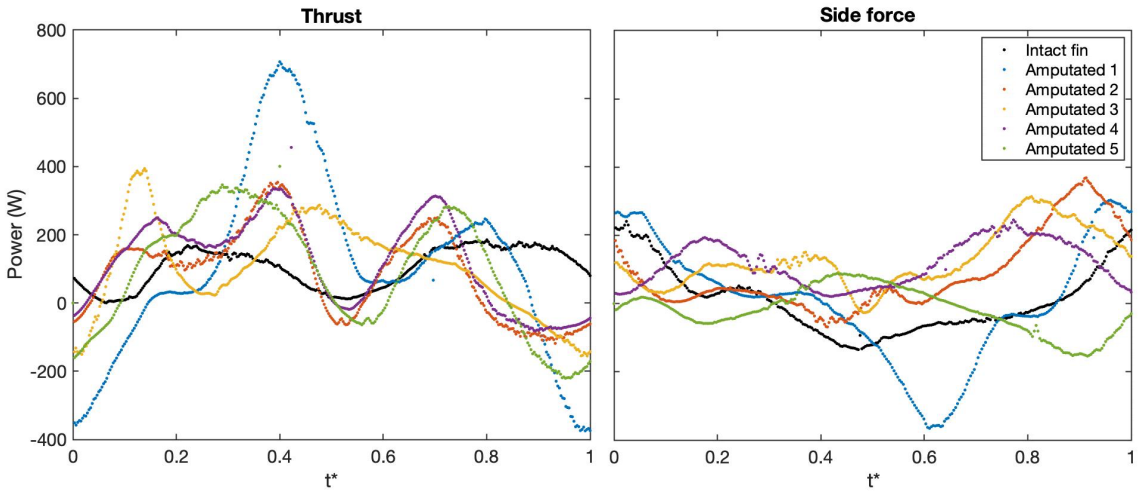

Figure 5.4: Instantaneous power requirements of optimal trajectories. The instantaneous power, estimated using force, position and time measurements, is plotted against non-dimensional time t*. A non-dimensional time of 1 is the length of one cycle.

We compare C-start power requirements with the power requirements of our thrust trajectories since C-starts use fishes' maximum possible instantaneous muscular power [151]. Using the maximum force and total power measured during C-starts in *Esox lucius* Linnaeus [42] as well as caudal fin measurements [91], we calculate a range in ζ from 6.4×10^{-6} to 2.0×10^{-5}. For our optimal machine-learned thrust trajectories, ζ was 1.2×10^{-5} for the intact fin and ranged between 3.3×10^{-6} and 6.5×10^{-6} for the amputated fins. We thus find that the optimal amputated trajectories are approximately a third as efficient as the intact machine-learned trajectory. However, both amputated and intact fins' ζ are of the same order of magnitude as the ζ of the *E. lucius* C-starts, meaning that the machine-learned thrust trajectories are probably within the energetic constraints of natural swimmers.

In addition, we compare the ζ for our optimal side-force trajectories with the ζ of *Drosophila melanogaster*. We compare power and force requirements at peak flight capacity [84] combined with wing measurements [144] to obtain a ζ of 8.4×10^{-3}. For the optimal machine-learned side-force trajectories, ζ was 2.4×10^{-6} for the intact fin and ranged between 1.5×10^{-6} and 6.7×10^{-6} for the amputated fins. While there is no decline in efficiency between amputated and optimal trajectories, we do see a decrease in efficiency of three orders of magnitude from *D. melanogaster*. Therefore, machine-learned side-force trajectories may not be achievable within the energetic constraints of natural flyers.

5.5 Discussion

When optimizing for thrust, a full recovery from breakage is seen for both fitness and force production by altering frequency, amplitude and AOA. Because they recover fully in all trials, the robotically actuated trajectories outperform the biological average. In fish, it depends on the species whether accommodation of approximately 50% damage is possible. By comparing the force traces of the intact and amputated fins, we see a qualitative return of force production to its pre-damaged baseline. Both the amplitude and frequency are increased to maintain force production. This increase matches the theory [119] but does not perfectly align with what is seen *in vivo*. Although fish do increase their tail beat amplitude upon breakage, they do not display a uniform frequency alteration after breakage. This non-uniformity may suggest that other evolutionary pressures not relevant to the robotic use case affect fishes' frequency adaptations. Furthermore, the AOA is significantly altered in the amputated fins as compared with the intact fin. The maximum AOA increases, as would be expected to increase force production, and a large phase shift is seen in the robotic optimization. Neither of these AOA modifications have been reported in the biological literature as far as the authors are aware.

Some differences between fish and machine-learned adaptations can be explained from the perspective of limited caudal fin actuation. The learned trajectory is easily able to compensate for 50% damage, while some fish may or may not be able to, owing to the larger actuation range of the SPM. The SPM is capable of actuating inside a much larger angular cone than most fish caudal fins can. This capability, combined with the result that the amputated fins converged to a stroke amplitude near the maximum mechanical constraint of the SPM, implies that the machine-learned adaptation shows improved compensation for damage because it can attain a larger range of motion.

When optimizing for side force, we encounter a more complex optimization landscape, leading to worse recoveries in both fitness and force production after breakage. In fact, one amputated trial was unable to recover to within 2% of the desired force production. These results imply that the optimization has a more difficult time accommodating the breakage when optimizing for side force as compared with thrust, probably owing to the optimization function landscape. Because the side-force optimal basins are more elongated than the thrust optimal basins, more of the principal component axes and therefore more of the trajectory parameters display high sensitivity. This high sensitivity results in more difficulty accommodating for

breakage.

A higher sensitivity, and therefore more difficult optimization, is further reflected in the effects seen on amplitude, frequency and AOA. The amputated trials show a nearly even split on whether a large increase in amplitude or an approximately equal amplitude is optimal to recover from breakage. This split is consistent with the biological literature, which similarly does not reach a consensus.

The effect on frequency is indistinguishable; amputated trials displayed higher, lower and equal frequencies compared with the intact fin. This spread contrasts with the biological literature, where most insects increase frequency in response to damage. However, one study on fruit flies noted that increasing wingbeat frequency is probably motivated by behaviours critical to predator evasion [105]. Without the observed increase in frequency, the fruit flies studied would need to use their maximum body-constrained wingbeat amplitude during nominal flight conditions, leaving no possibility for increased-speed manoeuvres. This suggests that increasing frequency is an adaptation driven by evolutionary pressures that are not relevant to the robotic use case and therefore is not optimal when only considering efficiency.

Finally, the optimal AOA trace was found to nest for all cases, when viewed in the resultant-based reference frame. This includes the intact case, which did not show a large phase shift like the thrust trials. The nesting of AOA optima is interesting, as it suggests that the directionality of force production in the x–y plane can be controlled by adding or subtracting AOA throughout the cycle.

Chapter 6

EFFECT OF COMPENSATORY STRATEGIES ON HYDRODYNAMICS

Most amputated fins showed a remarkable capacity to recover from breakage, attaining similar fitness and force production to the intact fin. It remains to be considered how this recovery is made possible. It may be that the amputated fin reestablishes similar hydrodynamics to the intact fin, thereby recovering force and fitness. However, it could also be the case that amputated fins take advantage of different hydrodynamic effects to achieve a similar result. In this chapter, we analyze flowfield measurements from PIV alongside forces and kinematics to determine what hydrodynamic mechanisms are the main contributors to force production, and whether these mechanisms shift in response to damage. We compare the amputated fin flowfields to the intact fin flowfield to determine whether the wakes are significantly different, or whether the amputated fins attempt to reestablish the same wake to attain the same fitness and force production. We note that in this chapter we primarily analyze vorticity fields to make conclusions about the flow. However, pressure also plays a role in force production in a fluid. There exists no technique to measure the pressure field with sufficient grid resolution in three dimensions within the fluid during the experiment, and so we do not consider pressure effects in this chapter.

6.1 Hydrodynamic Mechanisms Creating Thrust

6.1.1 Kinematics and Forces

The kinematics of the thrust trajectories and the bilobed shape of the force traces, where maximum thrust occurs around 10 and 190° for most trajectories[1], suggest that a drag-based mechanism is being employed (Figure 5.1). Drag-based mechanisms tend to occur when the fin's normal is parallel to the direction of fin motion. The force produced by drag will be opposite to the direction of motion. The largest force

[1]This applies to all thrust trajectories except Amputated 3. Although Amputated 1 may initially appear to be significantly different, its force and trajectory characteristics are similar. Its first maximum aligns with the majority of trajectories. Its second maximum is shifted to approximately 220° because the fin is slower to reach a high AOA. Maintenance of a higher force value after this point likely derives from flexibility effects; Amputated 1's overall lower AOA may allow the fin to bend to an optimal AOA via interaction with the fluid.

being produced at slightly greater than 0 and 180° is consistent with a drag-based explanation. When the fin traverses through 0 and 180 degrees, the fin's normal is parallel to the direction of motion. Via form drag, force would be created normal to the fin, opposing the direction of motion. Due to the mechanics of the spherical parallel manipulator, the largest component of normal force would then be directed towards positive thrust at >0 and >180 degrees, matching the observed force trends.

Amputated 3, in contrast to all other trials, shows azimuthally different maxima for its bilobed shape. In fact, Amputated 3 attains thrust maxima approximately when all other trajectories display thrust minima, and vice versa. The two minima of the force trace occur at approximately 30 and 200°. These locations are where the angle of attack of the fin are most streamlined with the direction of motion. This suggests a drag-based mechanism, because the force is minimized at a low angle of attack. However, the two maxima of the force trace occur at approximately 160 and 320°, when the constant fin stem radius is causing the fin to translate upwards. The fin's normal is approximately parallel to the direction of motion, and if the propulsion were still predominantly drag-based we would expect negative thrust to occur at these locations. It is thus less clear whether Amputated 3 is using a paddling strategy exclusively. Positive thrust occurring at 160 and 320° suggests that form drag is being overcome by some other hydrodynamic phenomenon, for example vorticity effects.

The kinematics and force trace of the movement alone cannot fully rule out lift-based or drag-based mechanisms, regardless of which trajectory we consider. We have so far assumed that the flapping motion dominates the velocity seen by the fin, causing the angle of attack to be defined with respect to the fin motion, but this is not true in all cases. For example, the motion of thunniform swimmers during heave is mostly in the direction of the fin normal, despite lift-based methods being their main mode of propulsion (section 2.4.2). During the heave motion, bulk flow velocity parallel to the fin causes lift-based forces to dominate because the effective angle of attack is small. To further rule out lift-based mechanisms, we must refer to the structure of the fin wake itself.

6.1.2 Full-Field Velocity and Vorticity

We compare the structure of the fin wake created by each trajectory using phase averages. Each phase average was calculated at 148 phase positions from 25 steady-state cycles. These phase positions are referenced to azimuthal 90° for the thrust

trajectories, as this was the phase position where the fin was most visible for referencing. Note that the phase positions do not correspond to specific azimuthal fin angles beyond this reference point. Since portions of the trajectory are subject to the speed-up value, the phase positions were instead chosen to reflect equal divisions of nondimensional time.

One distinction between lift-based and drag-based mechanisms lies in the effective angle of attack. Lift-based mechanisms use small effective angles of attack to produce forces perpendicular to the observed flow velocity. Drag-based mechanisms use large effective angles of attack to produce forces opposing the observed flow velocity. Given our measured phase-averaged velocity, we can determine the flow velocity experienced by the fin during its motion. The fin itself moves between an average 0.255 and 0.487 $\frac{m}{s}$, depending on the trajectory, but the mean wake velocity is at most 0.03 $\frac{m}{s}$. This means the fin has a high effective angle of attack, because the fin motion dominates the observed flow velocity and the fin normal is approximately parallel with the direction of motion. The high effective angle of attack, combined with the force produced opposing the direction of motion, is a hallmark of a drag-based mechanism.

The location of vortices measured in the wake can provide further clues to whether the primary mechanism is lift- or drag-based. Figure 6.3 traces the paths of the positive vortices shed from each fin. Each location at which vorticity reaches greater than 3.5 s^{-1} in a region comprising more than three pixels during the phase average is shaded. Amputated 1, 2, 4, and 5 show similar vortex structures, where the main positive vortex is shed in a roughly concave down, diagonal line that extends from the top right of the frame to the middle left of the frame. A smaller secondary region of positive vorticity exists in the top left of the frame. The main positive vortex is shed at approximately the middle of the first half-stroke, when the fin reaches its maximum z-extent. The secondary region does not grow in strength until after the fin has passed into the second half-stroke (Figure 6.1). It therefore is not attached to the fin. Amputated 3, on the other hand, does not shed its positive vortex. The vortex instead remains at the top of the frame (i.e., attached to the fin) for the full first half-stroke before being shed out of frame to the left. The intact fin's positive vortex structure, however, is the most different of all. During the first half-stroke, the intact fin sheds a positive vortex that spreads into the left of the central horseshoe shape seen in Figure 6.3. The bottom and upper structures are more persistent throughout the phase averaged vorticity.

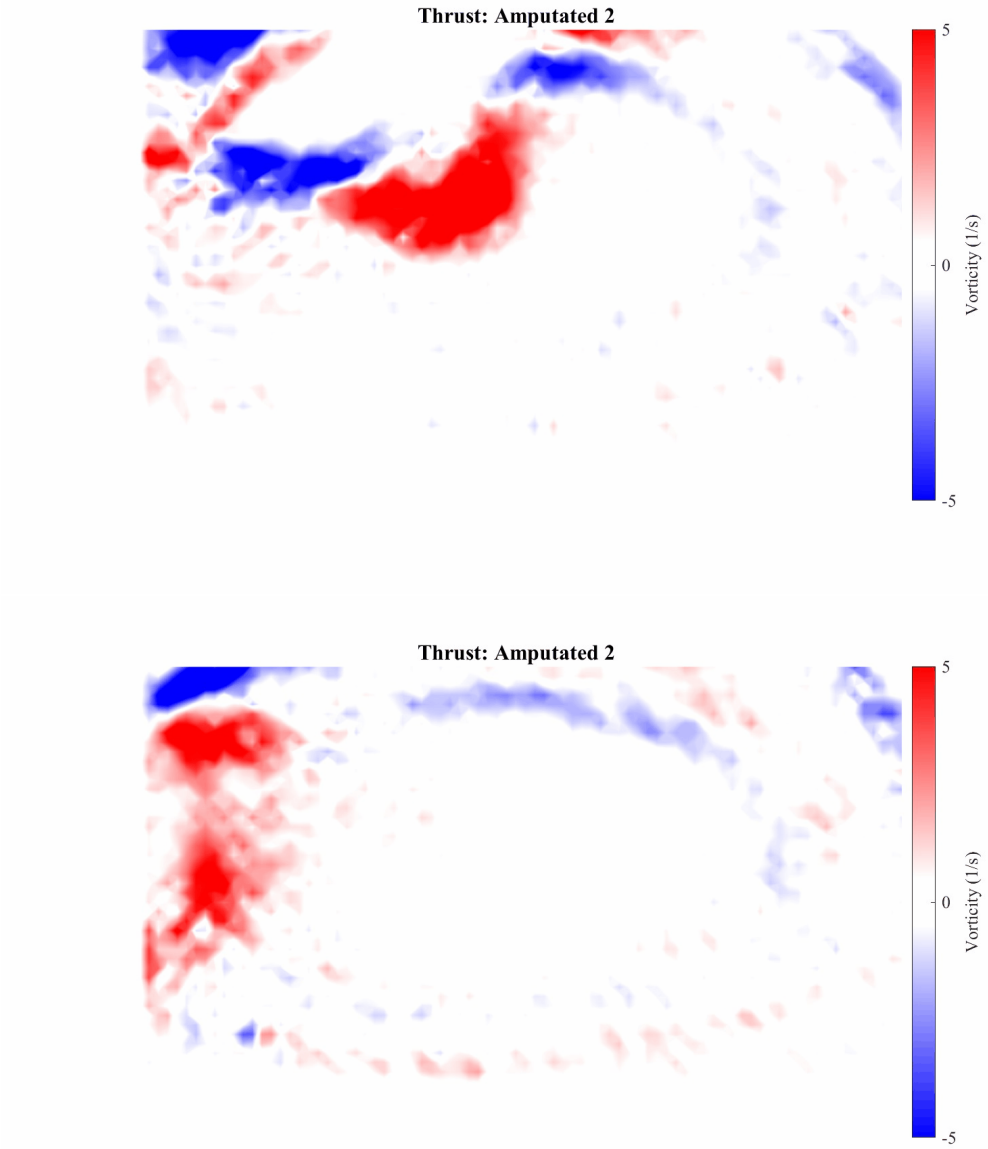

Figure 6.1: Snapshots of the $y = 0$ vorticity fields created by thrust-producing Amputated 2 at two phase positions in the cycle. The top image shows the main positive (starting) vortex, which has just begun to shed. The bottom image shows the secondary region of positive vorticity above the shed primary vortex during the second half-stroke. It also shows a region of high negative vorticity in the upper-left frame, suggesting a stopping vortex.

The negative vorticity is plotted in the same manner in Figure 6.4. All amputated cases have similar locations of high negative vorticity, while the intact case also has high negative vorticity in the middle right of the frame. This difference results from

the second half of the stroke; the amputated cases do not reinteract with the $y = 0$ plane, while the intact case does. During the second half-stroke, the intact fin sheds a negative vortex to the right, which corresponds to the right arm of the horseshoe shape. In all amputated cases, the visualized negative vorticity arc does not appear as a clear vortex whose vorticity is normal to the laser plane. Instead, there appears to be out-of-plane motion in these locations.

The structure of the wake in the intact fin and amputated fins 1, 2, 4, and 5 aligns with the drag-based mechanism that was suggested by analysis of the forces and trajectories. In drag-based paddling, a starting vortex is formed as the propulsor pitches. The starting vortex separates from the propulsor when the propulsor passes the midline and begins upward motion. When the starting vortex separates from the propulsor, the bound vorticity on the propulsor begins to coalesce. This bound vorticity, of opposite sense to the starting vortex, is shed as a stopping vortex when the propulsor reverses direction (Figure 6.2) [75]. The intact fin sheds a starting vortex at the midline of both the first and second half-strokes (positive when paddling in the negative x-direction; negative when paddling in the positive x-direction). The stopping vortex, however, is not as clearly seen. The main positive vortex is also shed at the midline for Amputated 1, 2, 4, and 5 during the first half-stroke. There is a region of strong negative vorticity that appears in the upper-left corner of frame when both the fin and the starting vortex have translated to the leftmost part of the frame (Figure 6.1). This could be the edge of the stopping vortex. The core of the stopping vortex would be out-of-frame at this point, but the strong negative vorticity suggests that it exists. The vorticity fields of the intact fin and amputated fins 1, 2, 4, and 5 are therefore consistent with utilization of a drag-based mechanism.

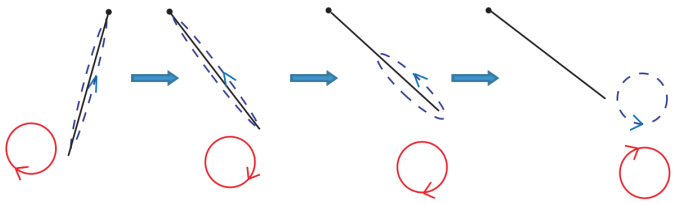

Figure 6.2: Formation and shedding of the tip vortex (continuous red lines) and the corresponding bound vortex (dashed blue lines). The shedding bound vortex forms the stopping vortex. Thin blue and red arrows indicate the rotating direction of vorticity (ω_z) on the z=0 plane. Reproduced from [75] with permission.

The structure of the wake in Amputated 3, however, contradicts the lift-based mech-

anism that was implied by analysis of the azimuthal location of maximum produced forces. At approximately 160 and 320°, despite the fin translating upwards at high angle of attack, positive thrust is produced. One explanation for this seeming contradiction is that lift-based mechanisms could be overcoming the form drag to produce positive thrust. The cross-section through $y = 0$ shows that Amputated 3 maintains an attached vortex on its bottom edge that increases in strength during the first half-stroke. It reaches its maximum strength right before being shed from the fin at approximately 160°. Because the vortex is attached to the underside of the fin, the resultant force would produce negative thrust. This is in contradiction to the measurement of positive thrust at 160°. Thus, the vorticity measured in the plane $y = 0$ suggests that a lift-based mechanism is not responsible for the thrust maximum at 160°. It does not, however, clarify how a positive thrust is produced at 160°.

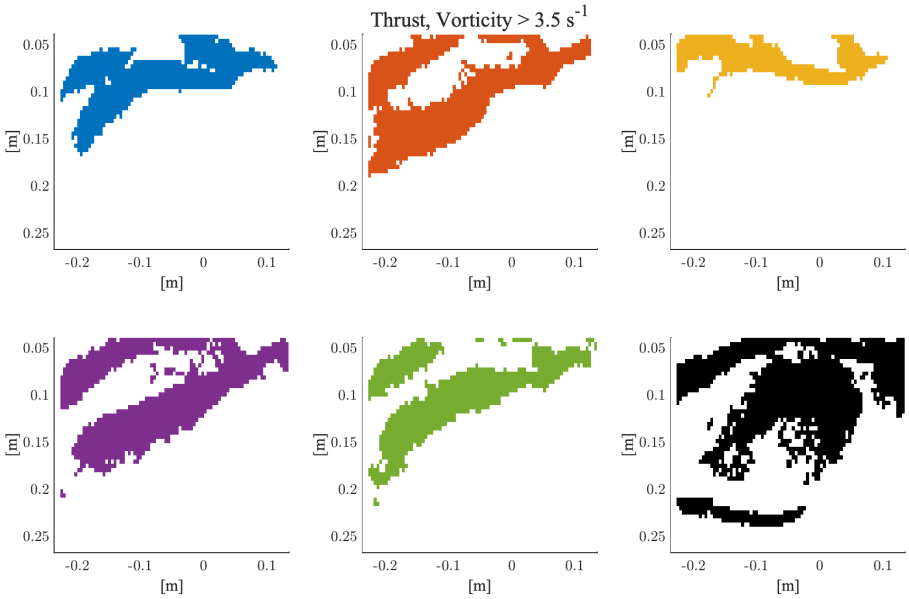

Figure 6.3: Paths of the positive vortices during the entire phase average on the $y = 0$ plane. Each spatial location which achieves vorticity higher than 3.5 s^{-1} in a region comprising more than three pixels at any point in the phase average is shaded.

6.1.3 Spatial Differences in Velocity and Vorticity

Although we have shown that the amputated and intact fin trajectories are similar in that they both utilize drag-based propulsion, it still remains to be considered how similar the overall wakes are. In this section we directly compare the velocities and vorticities attained at each spatial position to determine whether the amputated fins

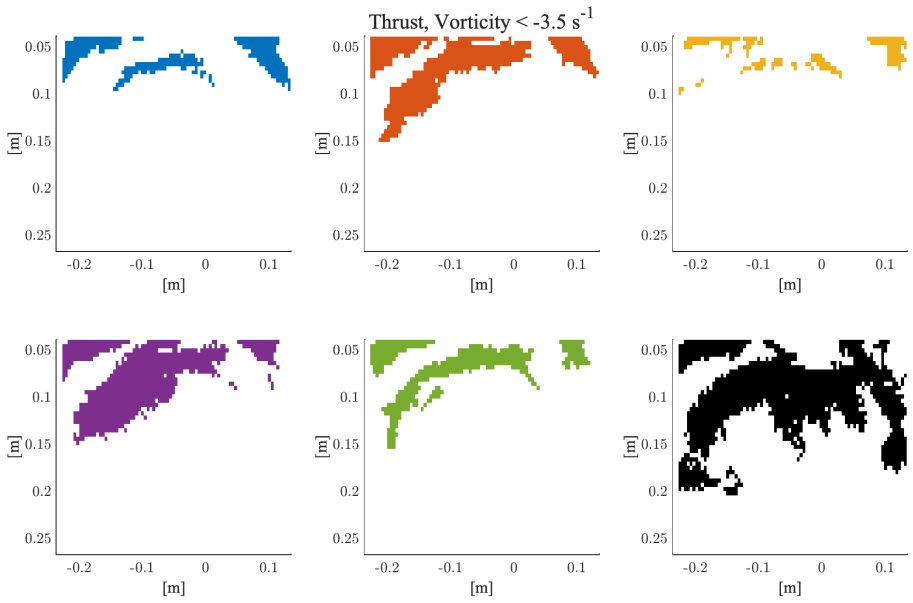

Figure 6.4: Paths of the negative vortices during the entire phase average on the $y = 0$ plane. Each spatial location which achieves vorticity more negative than -3.5 s^{-1} in a region comprising more than three pixels at any point in the phase average is shaded.

are simply reestablishing the same wake as the intact fin.

At each of the 148 phase positions, we calculate the difference in velocity magnitude and difference in vorticity at each spatial position. This difference is calculated by subtracting the second case from the first; for example, in "intact versus amputated 1" we subtract the wake of the amputated fin 1 from the wake of the intact fin. In order to represent these differences on a small set of plots, we calculate both the average velocity/vorticity difference for each spatial point as well as the maximum absolute value velocity/vorticity difference for each point.

Figure 6.5 shows the difference in hydrodynamics between the intact fin and amputated fin 1 producing thrust. The average velocity difference is highest in an ellipse in the center of the frame, with a slightly smaller ring surrounding it. We see from the maximum velocity difference plot that the greatest velocity differences form a concave down arc beginning at approximately $x = 0$. Regions of the outer ring also display a higher velocity difference, especially regions to the top and top right of the frame. In the vorticity differences, we see a similar though less pronounced ellipse and ring pattern. On the left side of the frame, the difference in vorticity is negative while the right side is positive. The maximum vorticity difference plot

shows a similar structure to the maximum velocity difference plot; the darkest region is the same concave down arc. In addition, the regions of the outer ring with higher velocity differences also show a higher vorticity difference. This general structure is repeated in all comparisons of the intact fin to amputated fins, with only slight differences (Figures A.2-A.5).

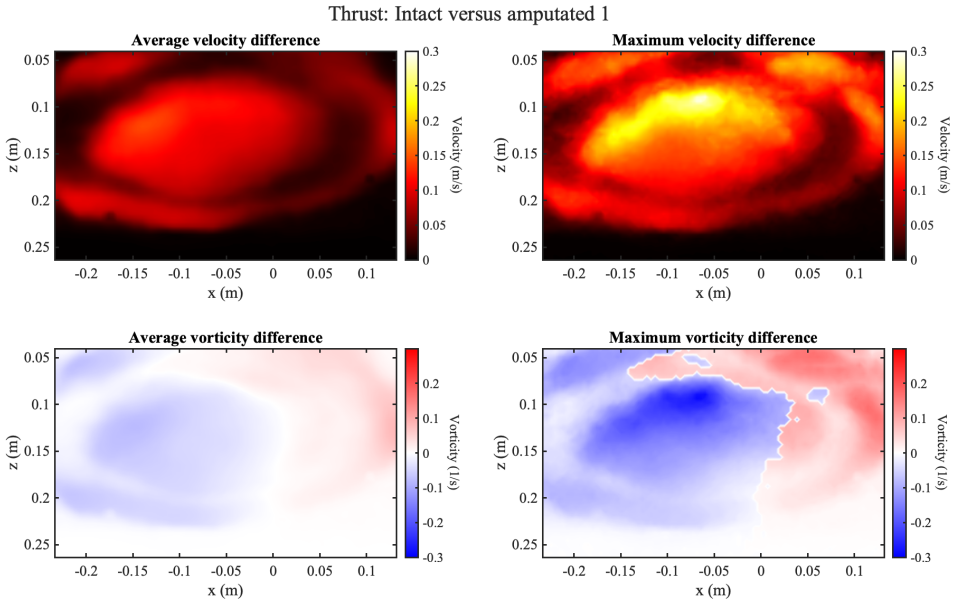

Figure 6.5: Difference in hydrodynamics between the intact and amputated fin 1 producing thrust. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

The concave down arc that appears in the maximum velocity/vorticity graphs is representative of the differences in vortex shedding. The amputated fin only intersects the $y = 0$ plane during the first half of the stroke, since it is oriented mainly in the $-y$ direction. The intact fin, on the other hand, intersects the $y = 0$ plane during the entire stroke. This means that in the $y = 0$ plane, the amputated fin's vorticity convects out of frame before the next stroke. The intact fin, however, is moving through vorticity that remains in its own wake. During the first half of the stroke, when both fins intersect the $y = 0$ plane, the amputated fin forms a positive vortex that convects in the negative x-direction with little z-motion. The intact fin, on the other hand, forms a positive vortex that sheds from the fin more rapidly, dropping in the z-axis. Negative vorticity backfills the gap between the positive vortex and the fin. This concave down arc of negative vorticity occupies the spatial location that a positive vortex convects through in most amputated cases, and thus the vorticity

and velocity difference is high in this region. Figure 6.3 shows the spatial difference in positive vortex shedding between the intact and amputated fins, and Figure 6.4 shows this difference for negative vortex shedding.

The differences in the rest of the field derive from the intact fin reinteracting with the vorticity field from the previous half-stroke. The intact fin disperses more vorticity throughout the field of view than the amputated fin. We see that while the minimum and maximum vorticities in the $y = 0$ plane are comparable to those of an amputated fin (table A.4), the overall distribution is shaped differently in the intact fin case than in the amputated fin cases (Figure 6.6). The intact fin has a higher standard deviation, leading to more vorticity of larger magnitude. In some frames, the amputated fin's distribution positively skews, representing the growth and convection of the positive starting vortex. Except for this skew, the amputated fin's distribution is always sharper than the intact fin. Figure 6.6 shows examples of the vorticity histogram at two different phase positions.

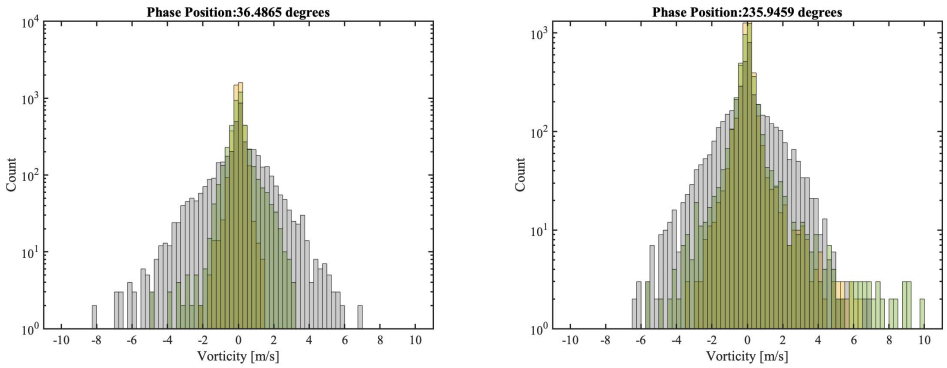

Figure 6.6: Histograms of the phase-averaged vorticity at two phase positions for the $y = 0$ plane. The gray histogram is the intact fin, the yellow histogram is Amputated 3, and the green histogram is Amputated 5. The left histogram is at 36.5° and the right histogram is at 235.9°. Again note that these phase positions are referenced to azimuthal 90°, but do not reflect the exact azimuthal position of each fin.

To consider whether the amputated fins' hydrodynamics are more similar to each other than they are to the intact fin, we compare the two most disparate amputated trajectories (Amputated 1 and Amputated 3) to each other. We repeat the spatial and histogram analyses, with the difference in hydrodynamics shown in Figure 6.7 and the histogram plotted in Figure 6.6. The average and maximum velocity differences attained are less than in the intact fin comparisons, meaning that the velocities are more spatially similar throughout a cycle. The vorticities also show a lower magnitude difference across the frame. We do see a distinct negative diagonal region

in the maximum velocity difference, resulting from the aforementioned difference in vortex shedding between Amputated 3 and all other amputated fin trajectories. However, its magnitude is still lower than the magnitudes attained in the intact fin comparisons. When considering the histogram of vorticity at different phase positions throughout the cycle, we further confirm the similarities between the amputated fins when comparing to the intact fin. Amputated 3 and 5 have a similar standard deviation, lower than the intact fin. They both display the same skewing behavior with the growth and convection of the positive vortex, despite differences in the location of such vortex. We can therefore determine that the amputated fins' hydrodynamics are more similar to each other than they are to the intact fin.

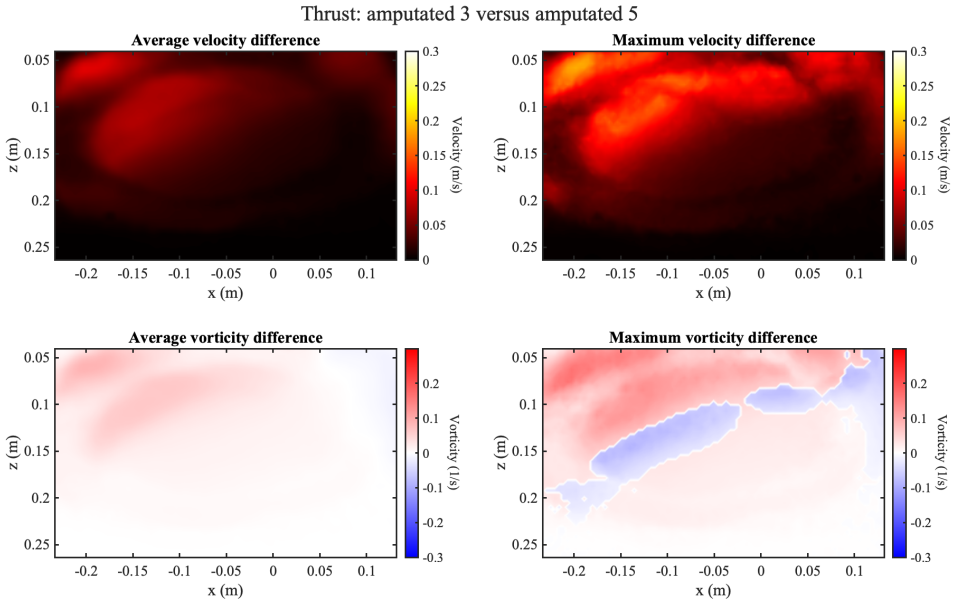

Figure 6.7: Difference in hydrodynamics between the amputated fin 3 and amputated fin 1 producing thrust. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

When comparing the intact fin with the amputated fins, we see that not only are there spatial differences in the wakes produced by these trajectories, but there are also statistical differences in these wakes. The vortices shed in different locations between the intact and amputated fins, and even among the amputated fins there are differences in wake structure. Furthermore, the intact fin creates much more vorticity in the $y = 0$ plane, as shown by the broader histogram in Figure 6.6. The

wake structures created by the amputated fins are more similar to each other than they are to the intact fin, as seen by comparing spatial and statistical measures.

6.1.4 Differences of Hydrodynamic Mechanism between Intact and Amputated Fins

Based on analyzing the forces, trajectories, velocity data, and vorticity data, it seems that the underlying hydrodynamic mechanisms are similar for most amputated fins to the intact fin. Both take advantage of drag-based propulsion, sometimes called paddling or rowing, to generate thrust force. This is the same style of propulsion employed by drag-based MPF swimmers like sea otters (section 2.5.2). However, the overall hydrodynamics are significantly different between the intact and amputated fins. Because the amputated fins do not reinteract with the central plane $y = 0$ during their second half-stroke, the vorticity and velocity fields of the central plane are significantly different. Not only does the intact fin shed a negative starting vortex in-plane during the second half-stroke, which the intact fins do not, but this negative vorticity interacts with the following (first) half-stroke to change the trajectory of the shed positive starting vortex. These significant differences show that despite reestablishment of similar fitness and force production using similar hydrodynamic mechanisms, the amputated fins are not simply reestablishing the same wake structure as the intact fin.

6.2 Hydrodynamic Mechanisms Creating Side Force

Quantitative measurements of the hydrodynamics in the side force cases do not as neatly support a drag-based versus lift-based conclusion. Instead, we identify a characteristic wake structure in most of the amputated side force trials and once again show that the amputated trials are significantly different from the intact fin. This suggests an optimal amputated wake structure and supports the conclusion that the amputated fins are not simply reestablishing the same wake as the intact fin to enable the same force and fitness characteristics.

6.2.1 Full-Field Velocity and Vorticity

Velocity and vorticity measurements from PIV data were taken in the same manner as for the thrust-producing trajectories (Section 6.1), but with the phase positions referenced to azimuthal 270° in the x-y plane for visibility reasons. In these measurements, the fin trajectories were rotated such that the x*-axis coincided with the experimental x-axis. This avoided out-of-plane motion as much as possible by

ensuring the resultant force is parallel with the PIV plane. Statistical measurements of the velocity and vorticity are reported in table A.4.

Amputated 2–4 show a distinct vortex-induced jet that is created by vortex pair shedding. When the fin passes through the central plane of the trajectory on the left side of the field of view, both a positive and negative vortex are shed to the left of the frame. The negative vortex sits atop the positive vortex and their interface is canted at a slightly positive angle (Figure 6.8). This vortex pair thus produces a strong jet flow in between the vortices from upper right to bottom left. This jet flow dominates the velocities seen in the frame. Due to conservation of momentum, we know there must be a resulting force on the fin in the opposite direction, which matches with our force measurement. We thus consider this vortex-induced jet to be the main hydrodynamic mechanism enabling recovery of side force.

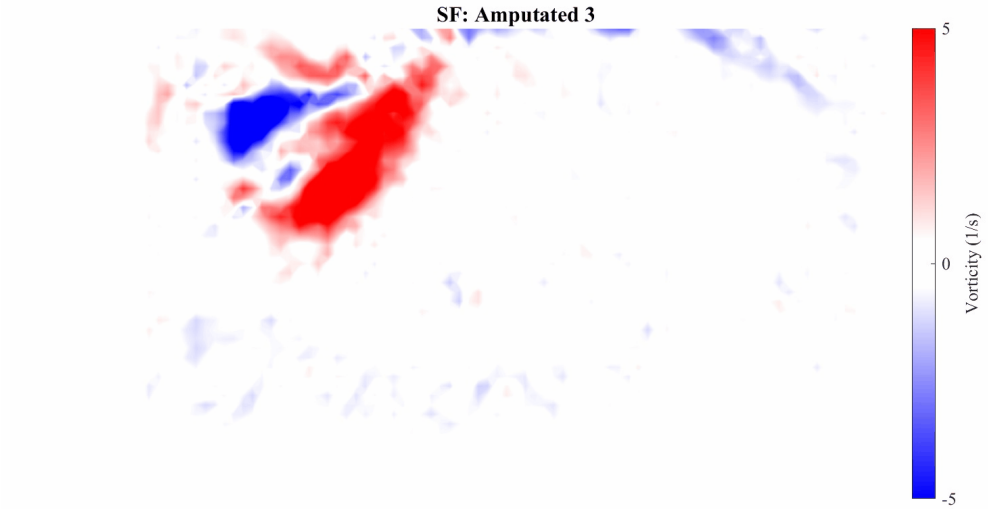

Figure 6.8: Snapshot of the $y = 0$ vorticity field created by side-force-producing Amputated 3 at one phase position in the cycle. This image shows the paired vortex that also sheds in Amputated 2 and 4.

Vorticity is much less pronounced in the visible portion of the central plane for Amputated 1 and 5. Both trials' vorticity is of smaller spatial extent and confined to the top of the frame. A snapshot of Amputated 5's vorticity is provided in Figure 6.9, showing the cycle's main visible feature: a negative vortex shed to the right and out of frame.

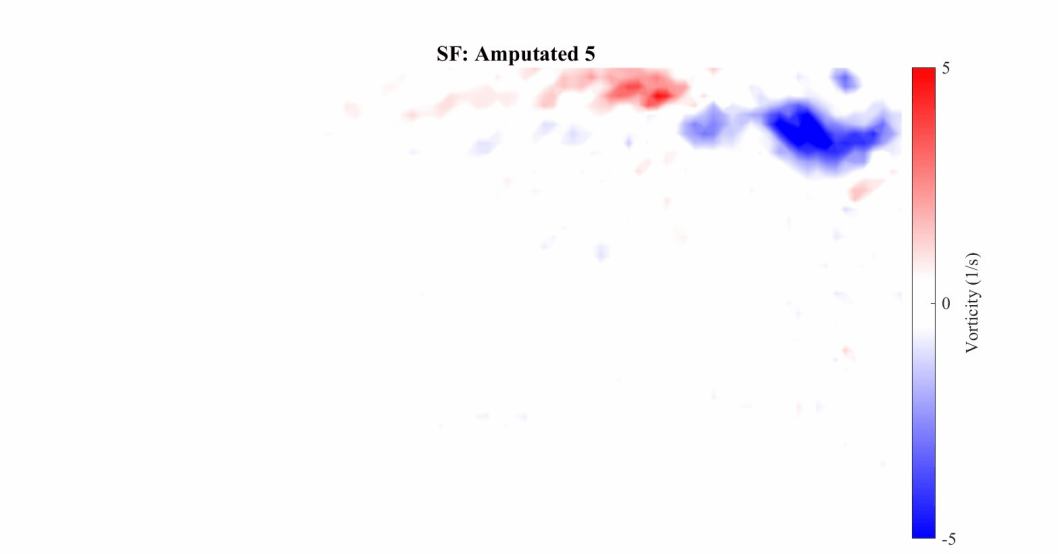

Figure 6.9: Snapshot of the $y = 0$ vorticity field created by side-force-producing Amputated 5 at one phase position in the cycle.

As was true when considering thrust, the intact fin producing side-force creates a much more complicated vorticity field. For side-force, the vortex shedding structure is significantly changed. There are two distinct structures in the intact fin vortex shedding. The first is a pair of positive vortices, which shed and translate apart in the x-direction with little z-motion before dissipation. The second is a strong negative vortex, shed to the right and out of frame. These structures are shown in Figure 6.10. The vorticity field of the intact fin does not show a clear vortex-induced jet like most of the amputated fin trials, suggesting that it may be utilizing a different hydrodynamic mechanism.

6.2.2 Spatial Differences in Velocity and Vorticity

At each of the 148 phase positions, we compare the velocity and vorticity differences as in Section 6.1. Figure 6.11 shows the difference in hydrodynamics between the intact fin and amputated fin 2 producing side force. We choose Amputated 2 as our reference because Amputated 1 converged to a solution with a 17% deficit in force production, and is therefore not representative of a full recovery after damage.

The average velocity difference is highest in an elongated region at the top right of the frame, roughly centered around $x = 0$. The maximum velocity differences show the same structure. In the vorticity differences, we see regions of positive

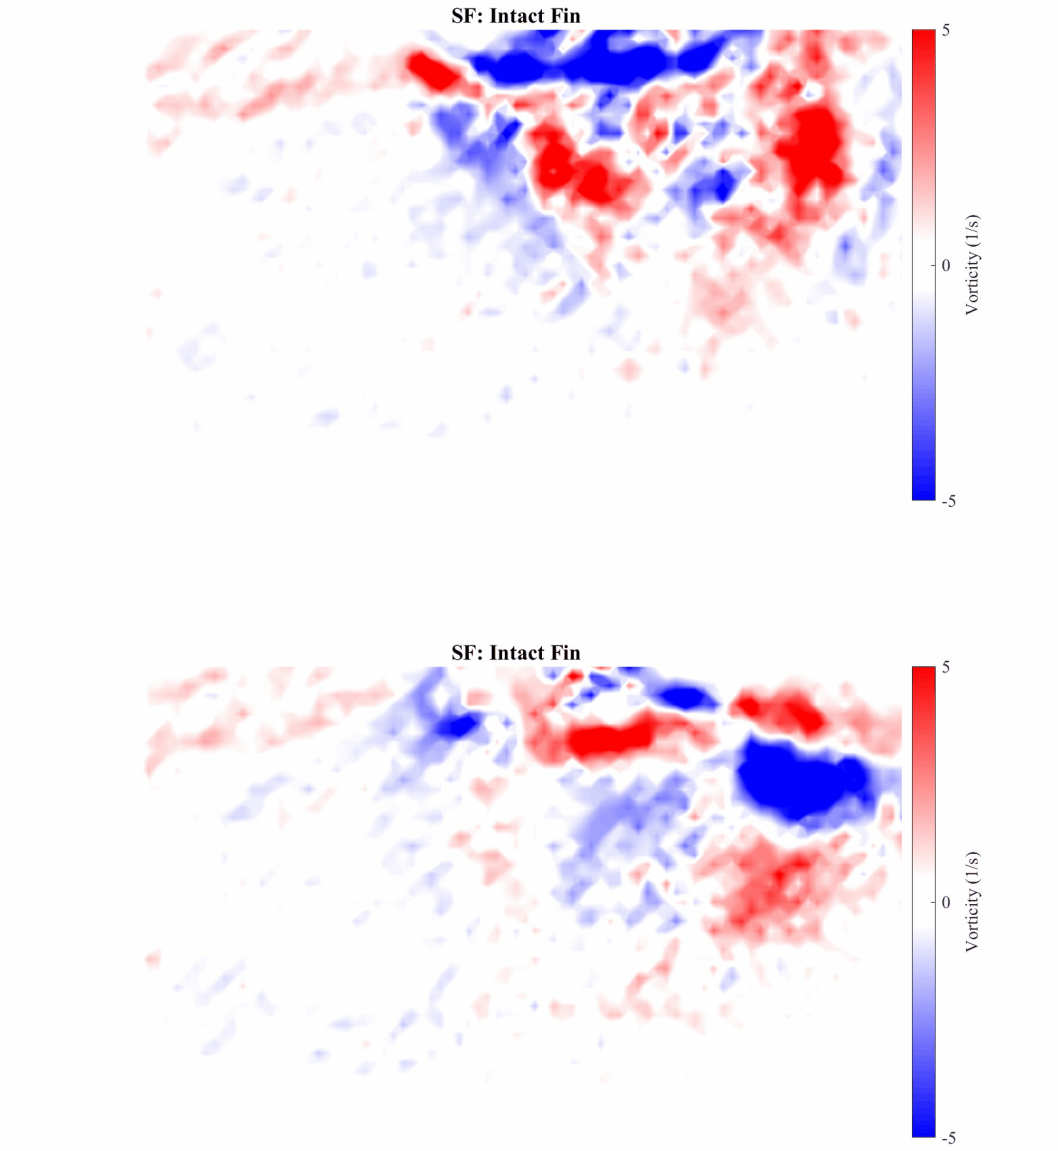

Figure 6.10: Snapshots of the $y = 0$ vorticity fields created by the side-force-producing intact fin at two phase positions in the cycle. The top image shows two positive vortices created by the stroke. The bottom image shows a large negative vortex created by the stroke.

vorticity difference at the left side of the frame as well as in the top right corner. Except for this corner, a region of negative vorticity difference occupies most of the right half of the frame. Neither the velocity nor vorticity differences display an elliptical structure as seen in the thrust measurements— rather, the magnitude of

the difference tends to uniformly decay with increasing z. This general structure is repeated in the comparisons of all amputated fins, with only slight differences. These minor differences are as follows: the region of large velocity difference may extend further in the x-direction, and the boundary between positive and negative vorticity differences may vary in shape. Plots of the other amputated fins can be found in Figures A.6–A.10.

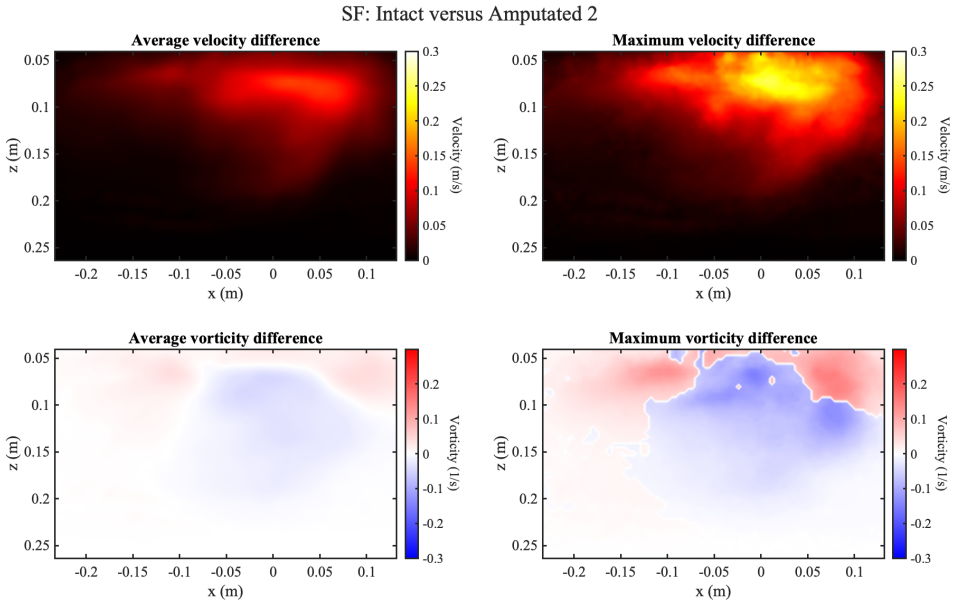

Figure 6.11: Difference in hydrodynamics between the intact and amputated fin 2 producing side force. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

To consider whether the differences between the hydrodynamics of the intact fin and the hydrodynamics of the amputated fins is greater than the differences among the amputated fins, we choose to compare Amputated 2 with Amputated 5. Both Amputated 1 and 5 do not display the characteristic paired vortex shedding described in the previous section and thus display the most different hydrodynamics to Amputated 2–4. We choose to compare Amputated 5 as it was able to maintain force production fully.

Figure 6.12 shows the difference in hydrodynamics between Amputated 2 and Amputated 5. Although the vortex shedding dynamics are different between these two trials, the overall velocity and vorticity differences support the conclusion that the amputated fins' hydrodynamics are more similar to each other than they are to the

intact fin's hydrodynamics. Both the average and maximum velocity differences attain lower values when comparing Amputated 2 to 5. The vorticity difference, while attaining similar maxima, is on average lower across the whole frame. Because the spatial velocities and vorticities are more similar between Amputated 2 and 5 than they are between the intact fin and Amputated 2, we can say that the wake structures created by the amputated fins are overall more similar to each other than they are to the intact fin.

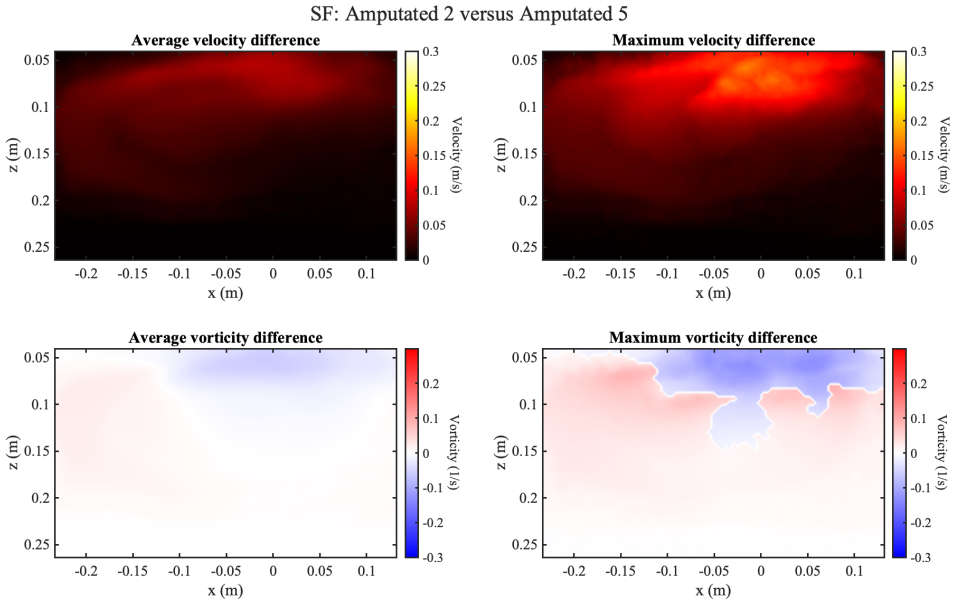

Figure 6.12: Difference in hydrodynamics between amputated fins 2 and 5 producing side force. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

Chapter 7

CONCLUSION

In Part I of this thesis, we compare trajectory adjustments as a result of catastrophic propulsor damage between machine-learned strategies and *in vivo* observations. When producing thrust, the machine-learned strategies are able to recover fully in both fitness and force production in all trials by increasing amplitude and frequency. In fish, on the other hand, not all species are able to recover force production, and not all species increase frequency as a result. One explanation as to why not all fish are able to recover force production through trajectory modification may derive from the increased actuation capabilities of the experimental setup. Although the machine-learned trajectories are within the energetic constraints of natural swimmers, the experimental setup is capable of actuating a larger amplitude than most fish. The machine-learned trajectories also show both an increase in angle of attack as well as a large phase shift, which has not been reported in the biological literature as far as we are aware.

When examining the forces and trajectories further, along with quantitative measurements of the fin wake, we find that both the intact and the amputated fins are utilizing drag-based propulsion to generate thrust force. However, the overall hydrodynamics are significantly different between the intact and amputated fins. The intact fin passes through the central plane of the trajectory on both half-strokes, while the amputated fins only pass through this plane on the first half-stroke. This results in a different spatial pattern of vortex shedding during the first half-stroke between the intact and amputated fins, and vortex shedding occurs for the intact fin in the second half-stroke that is absent in the amputated fin cases. Even the most different amputated trajectories (3 and 5) are more similar to each other than they are to the amputated fin. This supports the conclusion that the amputated fins are not reestablishing the exact same wake structure as the intact fin in order to regain fitness and force production, but they do utilize the same hydrodynamic mechanisms.

The machine-learned strategies for maintaining side force tell a more complex story. The optimization itself is more challenging to traverse than the optimization for thrust, as seen in the narrower shape of its optimal basin. In fact, one of the amputated trials was unable to recover force production to within 2%. Neither

the experimental results nor *in vivo* observations of insects show a clear strategy for modifying amplitude in response to damage. The experimental results also do not support a clear strategy for modifying frequency; but insects tend to increase frequency in response to damage, possibly for the purpose of predator evasion. The amputated cases did, however, show a clear nested trend in optimal angle of attack. This suggests that one strategy to control the direction of side force produced is to add or subtract a constant angle of attack throughout the cycle.

The hydrodynamics of the fins producing side-force do not as clearly support a drag- or lift-based propulsion strategy. They do, however, illuminate a characteristic wake structure of a vortex-induced jet. Even among the most different side-force producing trajectories (Amputated 2, which creates such a jet, and Amputated 5, the only trajectory able to recover force production without such a jet), the velocities and vorticities are more similar than they are to the intact fin's wake. This once again supports the conclusion that the amputated fins are not simply reestablishing the same hydrodynamics that were learned by their intact counterpart.

The complexity of the side force cases, as well as the inability to determine an optimal frequency and amplitude response, deserves further research. To mitigate the sensitivity of the side force optimization, one could reduce the number of variables participating in the optimization, thereby reducing the number of sensitive dimensions. Fixing the angle of attack is a good candidate, due to the observed similarity in angle of attack traces. This would reduce the complexity of the optimization from nine to six parameters and thereby minimize optimization time. Exploring other parameters with large spanned spaces, such as the magnitude and location of fin damage, may be made feasible with less experimental runtime per trial. In addition, generating a larger number of experimental trials may illuminate the optimal frequency and amplitude response, since statistical methods may be employed on a larger dataset.

Given the success of applying this machine learning technique to a complex propulsive system where fluid-structure interactions are significant, we seek to apply the same methods to understand another complex fluid-structure problem in Part II.

Part II

Individual Blade Control Rotor

Chapter 8

INTRODUCTION

Helicopter rotors pose a highly complex aeroelastic problem. Not only are the blades constantly experiencing varying velocity profiles as they translate azimuthally, but fluid-structure interactions are significant. The fluid dynamics interlink with the dynamics of the structure itself, as the flexibility of the rotor blades plays a key part in the interaction. In addition to the flexibility of the rotor blades themselves, lead-lag and flapping hinges increase the mobility of blades in response to fluid forces. While these fluid-structure interactions are often advantageous, damping unwanted vibration and mitigating unintentional dissymmetry of lift, they can equally give rise to dangerous situations like ground resonance.

The control system of helicopters, which collapses such dynamic complexity into manageable pilot input, has not significantly changed since helicopters' inception. Gaetano Crocco proposed cyclic pitch control for directional flight in 1906, with the first successful manned and semi-controlled implementations in the 1920s [85]. Most modern-day helicopters use the same system. Expanding beyond cyclic pitch control has been shown to improve a variety of desirable characteristics, yet implementation of such strategies has been slow due to the massive search space. In Part II of this thesis, we seek to use machine learning techniques to better enable exploration of individual blade control.

In Chapter 9, we provide an overview of single-rotor aeromechanics. We first describe two frequency domains and their associated effects, leading into a discussion of different methods to access these domains via modifying pitch. We explore the wide variety of benefits of these methods, including goals as various as reducing vibration, reducing blade-vortex interaction noise, delaying retreating blade stall, alleviating gusts, enhancing performance, and more. We finish this chapter with a discussion of how machine learning has been applied previously to individual blade control (IBC). To this author's knowledge, there have been no previous examples of machine learning implemented with IBC hardware-in-the-loop.

Chapter 10 begins with a brief reintroduction of the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) that has been successfully used in Part I of this thesis. We describe how CMA-ES has been implemented on the purpose-built

subscale IBC rotor and briefly touch on the software enabling integration of CMA-ES to the cyberphysical system. The hardware and electrical systems of the IBC rotor are discussed.

Chapter 10 also includes data analysis to justify the experimental setup. By stepping through nondimensionalization of the governing parameters of helicopter flight, we show that dynamic similarity cannot be achieved except with a full-scale model. Thus, we justify the departures from dynamic similarity inherent in our subscale model by calculating nondimensional parameters. We characterize the experiment data using error measures and the Allan sensor variance, and finally justify the overall experiment by running two optimizations with known outcomes.

In Chapter 11 we present two in-progress optimizations for thrust vectoring. The first optimization was designed to angle the thrust vector by approximately 5°, thereby vectoring thrust while maintaining sufficient weight support. This optimization utilized two blades with a smaller subset of the available optimization parameters. The second optimization was designed to maximize the side-force created and torque that is not normal to the rotor plane. For this optimization, we again use two blades but include more parameters to increase the search space. These optimizations provide insights into how the expanded capabilities of IBC may be leveraged for advanced aeromobility.

We conclude Part II of this thesis in Chapter 12, summarizing the preliminary results of this research and enumerating future possibilities using this setup.

Chapter 9

SINGLE-ROTOR AEROMECHANICS

Single-rotor craft, for example helicopters, rely on a single rotating element to produce variable thrust. This element is typically made up of several airfoils called blades, which each produce thrust when moving through the fluid. By altering the angles which these blades make with the flow, the helicopter is able to maneuver in three-dimensional space.

The topic of helicopter aeromechanics is a complex one, as the fluid dynamics are interlinked with the elastic behavior of the blades themselves— a classic fluid-structure interaction problem. Different blade modes (ex. torsion, bending) disturb the flow at a wide range of frequencies. For a detailed treatment of helicopter aeroelasticity, one should refer to [70]. However, it is relevant here to note that helicopter aeromechanics are separated into a low-frequency and high-frequency domain (Figure 9.1). The low-frequency domain consists of all disturbances at or below the frequency of the rotor's rotation. In the low-frequency domain, disturbances affect the helicopter's gust response, flying qualities, blade instabilities, and ground resonance. The high-frequency domain comprises all disturbances at higher frequencies than the rotor's rotation. At these frequencies, disturbances affect blade bending stress, vibration, and stall flutter [73].

9.1 Pitch Actuation

Modifying the pitch of helicopter blades allows not only for maneuverability, but in more complex systems may also mitigate some disturbances in both the low- and high-frequency domains. We discuss three common methods of controlling blade pitch in this section.

9.1.1 Swashplate

Most modern-day helicopters utilize a swashplate to change blade pitch. The swashplate is a set of two platforms that reside underneath the rotor hub. The bottom platform is non-rotating and is raised, lowered, and tilted by actuators. The second platform is stacked on top of the bottom platform and rotates with the rotor hub. When the bottom platform is raised, lowered, or tilted, these motions are transferred to the top platform.

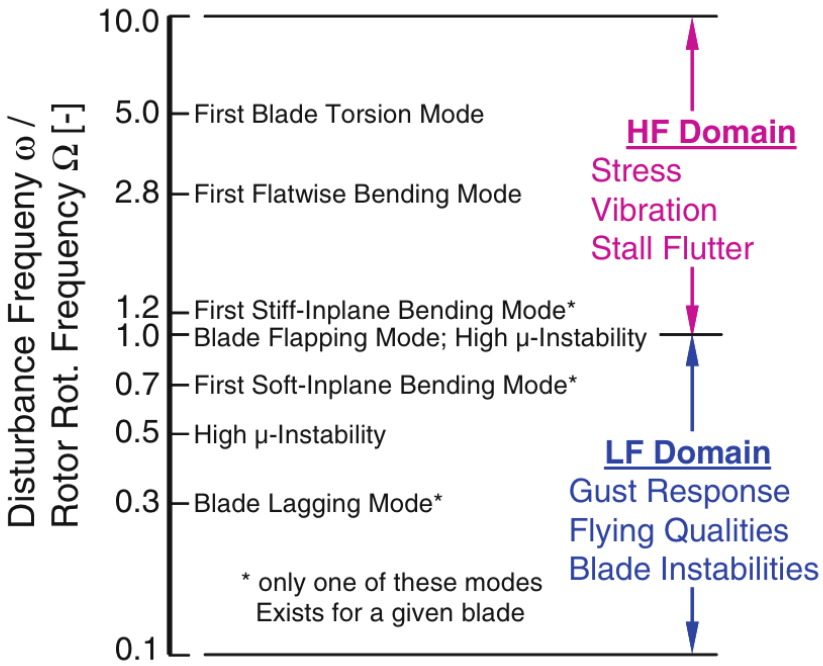

Figure 9.1: High-frequency and low-frequency domains of rotor aeromechanics, including examples of modes and instabilities within these regimes. Reproduced from [73] with permission from SNCSC.

The top platform is connected to the blades through a pitch linkage. The pitch linkage transfers the linear motion of the swashplate to an angular motion of the blade about the blade bearing. When the connection point of the pitch linkage to the swashplate is raised, the blade pitches in one direction. When the connection point is lowered, the blade pitches in the opposite sense. The direction of pitch depends on whether the pitch linkage attaches to the leading- or trailing-edge of the blade.

Thus, when the entire swashplate is raised or lowered levelly, all blades increase or decrease in pitch by the same amount. This is called collective pitch. Collective pitch enables bulk increase or decrease of thrust. This is used for vertical takeoff and landing.

When the swashplate is tilted, this results in cyclic pitch. When a blade passes through the lower side of the swashplate, its pitch value will be different from when it is located at the high side of the swashplate. Because the blades are constrained to rotate in a circle, the blade pitch must vary sinusoidally at the frequency of the rotor hub rotation. Thus, the lift of a single blade also varies sinusoidally at the rotor hub rotation frequency. This causes an imbalance in the lift generated across the rotor

disk, creating a torque about the line where average pitch is attained. However, due to gyroscopic precession, the rotor disk tilts about a line that is 90 degrees offset. When the rotor disk tilts, the overall thrust vector is tilted with it. Tilting the rotor disk allows the helicopter to move left, right, forward, or backward.

9.1.2 Higher Harmonic Control (HHC)

Simply reorienting the swashplate can only have effects at zero frequency (collective pitch) or at the rotor hub frequency (cyclic pitch). However, swashplates may be combined with Higher Harmonic Control (HHC) to enable control of higher multiples of the rotor hub frequency. In this method, the swashplate is oscillated at harmonic frequencies to excite or damp specific frequencies in the rotating frame. The oscillations may induce pitch changes of up to $\pm 3^\circ$ [66].

While HHC expands swashplate control to cover more than just two frequencies, it is still limited in range. Disturbance frequencies often occur at non-multiples of the rotor hub frequency (see Figure 9.1), which cannot be accessed by HHC. Furthermore, not all integer multiples of the rotor hub frequency can be accessed. A four-bladed rotor can only modify at 0, $4n - 1$, $4n$, and $4n + 1$ times the rotor hub frequency. This leaves gaps at critical frequencies such as two times the rotor hub frequency, control over which is required for reducing power and delaying stall [73]. The more blades a helicopter has, the more gaps in harmonic control authority HHC has [74].

HHC has mostly been studied for the application of reducing helicopter vibratory loads, comprising simulation, wind tunnel experiment, and in-flight demonstration [2, 41, 55, 104, 110, 111, 131, 134]. Flight tests showed vibration reduction between 70 and 90% [100, 117, 141, 161]. While HHC does well at this task, it often increases other undesirable characteristics. Suppressing both vibration and rotor noise has proven to be challenging except in specific cases of blade-vortex interaction [74]. Most often, the best inputs to minimize vibration increase noise and vice versa [66]. Optimizing HHC to damp vibrations can also negatively affect the flight envelope of the helicopter, increasing the likelihood of stall [131]. Vibration reduction is certainly not the only goal one can achieve by controlling higher harmonics; indeed, one might consider reversing the tradeoff and delaying retreating blade stall whilst accepting an increase in vibration [5, 114]. However, this tradeoff combined with the limited control authority of HHC has prevented it from gaining wide implementation in modern-day helicopters.

9.1.3 Individual Blade Control (IBC)

Individual blade control, or IBC, was originally proposed to mitigate the limitations of swashplated designs [78]. IBC now can be used to refer to a variety of different systems that enable the pitch of each blade to be changed separately. These range from swashplateless systems to multi-swashplate systems to active blade twist actuation and more. The method by which each blade attains an uncoupled pitch history is not crucial except for each method's limitations. Multi-swashplate systems have the greatest limitations, since each blade is still constrained to its swashplate's circularity. Swashplateless systems enable arbitrary pitch actuation, but the blade pitch is constant from root to tip. Active blade twist opens up even more degrees of freedom, since by modifying twist the pitch can change from root to tip [73]. Moving forwards, we will discuss the swashplateless IBC design, as that is the design most commonly referred to by the term IBC.

In a swashplateless IBC design, all actuators are integrated into the rotating hub. The most simple method of this is to connect each blade at its root to a servo motor that directly drives blade pitch[1]. Because the actuators reside in the rotating hub, they are able to feather the blades with any frequency. IBC can therefore access all harmonic and non-harmonic modes. Many disturbances lie outside the harmonic spectrum and therefore could be mitigated by IBC (see Figure 9.1).

IBC has been widely investigated for goals as various as noise reduction, vibration reduction, stall flutter suppression, gust alleviation, performance enhancement, and more in simulation, experiment, and in-flight testing [54, 109, 118, 121–123]. The most comprehensive study of IBC inputs was undertaken as part of the U.S./German Memorandum of Understanding on Helicopter Aeromechanics. Jacklin et al. analyzed a parameter sweep comprised of single- and multi-frequency inputs from two to six times the rotor hub frequency on a full-scale BO-105 helicopter rotor mounted in a wind tunnel. Multi-frequency inputs allowed for the creation of pulses, wavelets, and doublets. All inputs could be actuated at up to $\pm 3°$. The single-frequency input that best reduced both noise and vibration was the second-order harmonic[2]. Combining the second-order harmonic with additional frequency content reduced noise by 85% during descent, fourth-order harmonic hub vibration by 75% during de-

[1] In-flight technology demonstrators use much more complex methods to enable swashplateless IBC control with minimal increase in failure risk. Directly connecting the actuator to the blade root would create as many critical points of failure as the number of blades.

[2] Recall that the second-order harmonic is inaccessible through HHC on any rotor with four or more blades.

scent, and even improved performance by net 3% during nominal flight. Achieving low noise, low vibration, and high performance all with the same actuation was a novel finding. However, despite a very large parameter sweep (1,328 experiments in total) the curse of dimensionality means that this parameter sweep was only a small fraction of the operating space [66].

9.2 Machine Learning Applied to Helicopter Control

Determining the optimal control input for helicopters, whether utilizing HHC, IBC, or a simple swashplate, has been more recently investigated using reinforcement and machine learning (RL and ML) methods. RL and ML are well-suited to learn the complex nonlinear dynamics of aeroelasticity and mitigate IBC's curse of dimensionality.

Even the simple swashplate design poses challenges at the extremes of its flight envelope. Low-speed aerobatic maneuvers one might see in radio-controlled helicopters (like inverted flight, flips, rolls, and funnels) are highly unstable. The optimal control input to make these maneuvers autonomous on real-world radio-controlled helicopters can be determined by designing a controller using standard RL [108] or more complex RL techniques like differential dynamic programming [1]. Low-speed maneuvers also occur if the main motor fails and a safe descent under autorotation is required. In simulation, another RL technique called Q-learning [80] and a recurrent neural network [27] have both been used to optimize autonomous autorotation maneuvers.

Using IBC extends the question of what control input is optimal into intermediate flight regimes like steady cruising. While swashplated flight in these regimes is open-loop stable [108], the superposition of harmonic and/or nonharmonic pitch content may disturb such stability. Furthermore, the optimal parameters of IBC for mitigation of specific identified disturbances (such as, for example, noise or vibrations) cannot be easily determined without either costly and noncomprehensive parameter sweeps (like in [66]) or the application of ML.

Optimal IBC parameters have been successfully determined using ML techniques in order to mitigate noise from blade-vortex interaction (BVI) and to minimize third-order harmonic vertical vibrations. Gläßel et al. used neural networks to create both a direct and indirect controller to mitigate BVI noise using data from IBC test flights performed by Eurocopter Deutschland. Their models reduced the maximum of the pressure amplitudes as a proxy for noise, showing just less than a fivefold

reduction using the controllers [50]. However, this work has not been tested outside of the preexisting Eurocopter Deutschland dataset.

To mitigate third-order harmonic vertical vibration, Gao et al. applied a particle swarm optimization method to simulated IBC. They optimized the amplitude and phase of a second-order harmonic feathering, quickly converging to a result that minimized vibration by over 80% in the third-order harmonic [45]. Yang et al. investigated the same optimization, but instead using a genetic algorithm. Again, convergence was quick with 60% reduction after 30 generations and an overall 65% reduction in vibration at convergence [167]. Yang et al. repeated their work, using a steepest descent– golden section combinatorial optimization algorithm and expanding to both second- and third-order harmonic control. They found that optimal amplitude and phase varied with advance ratio and were capable of reducing vibrations by 60% [166]. Because the aeroelastic models used in each work are different, we cannot use the vibration reduction metrics to state that a specific algorithm is better for this application. Furthermore, due to the complexity of aeroelastic interactions and modeling, these results may not be transferable outside of simulation.

This author has not been able to find examples of ML implemented on IBC rotors with hardware-in-the-loop. We therefore present our experimental setup to bridge the gap between real-world IBC parameter sweeps and simulation-only ML on IBC. We apply ML directly to a real-world IBC rotor to optimize for our desired outcomes.

Chapter 10

EXPERIMENTAL SETUP

To enable efficient exploration of the IBC search space, we pair a swashplateless IBC rotor with a Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES). Given the demonstrated success of utilizing CMA-ES on a real-world, noisy optimization in Chapters 4–6, we choose to approach the optimization of IBC rotor dynamics with the same strategy.

10.1 Optimization and Software

10.1.1 CMA-ES

CMA-ES is part of a class of machine learning methods called evolutionary algorithms. Evolutionary algorithms iterate over groups of candidate trajectories called "generations". A subgroup of these candidate trajectories that have attained the lowest fitness is chosen to be the "parents" of the next generation. This process continues until convergence. CMA-ES in particular uses the fitness of the best candidate trajectories to adapt the mean and covariance of a probability distribution, using this shifted distribution to select the next generation's candidate trajectories. This probability distribution eventually converges to a minimum of the fitness function. In real-world, noisy datasets with complex optimization landscapes, CMA-ES does well at avoiding local minima [58]. For more detail, a full description and mathematical treatment of the CMA-ES algorithm can be found in Chapter 4.

The base CMA-ES algorithm has not been significantly altered between the experiments in Chapters 4–6 and the following rotor experiments. However, key parameters of the CMA-ES algorithm are defined in reference to optimization characteristics, such as the number of variables to optimize over. The number of variables is changed not only between the flapping propulsion experiments and the rotor experiments, but also among the rotor experiments. Thus, these key parameters of CMA-ES defining generation size, optimization speed, etc. will vary.

10.1.2 Optimization Variables

The motion of each blade in the assembly is defined by seven parameters: the percentage triangular Δ, the frequency f, the maximum pitch angle α_{max}, the minimum pitch angle α_{min}, the zero-crossing α_0, the speed-up value K_v, and the

phase shift θ. Δ denotes the weight of the triangular wave in a weighted average; a value of 0.25 would correspond to 75% sinusoidal and 25% triangular. α_0 defines the pitch angle of the juncture between the concave up and concave down portions of the trajectory. K_v denotes the percentage slowdown of either the concave down or concave up portion; if negative, the slowdown applies to the concave up portion and if positive, the concave down portion. Thus, at full capacity this system may optimize for 28 parameters. For all experiments in this thesis, the system utilizes two blades. The optimizations are increased in complexity by adding more parameters to participate in the optimization. It is important to note that this parameter space extends significantly beyond cyclic and collective pitch. Nonsinusoidal actuation and nonharmonic frequencies are both achievable. Figure 10.1 shows four examples of possible pitch trajectories in this parameter space.

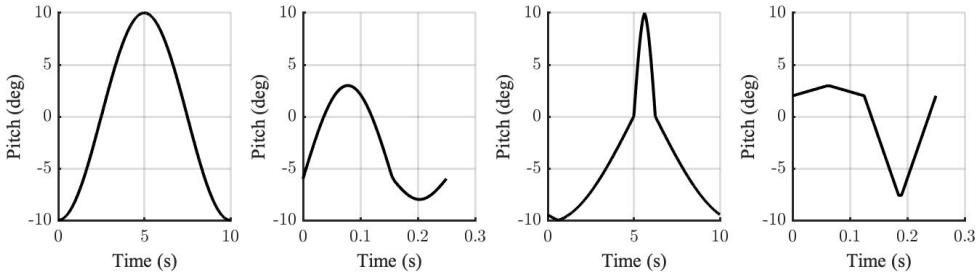

Figure 10.1: Four example pitch trajectories. The input parameters ($\Delta, f, \alpha_{max}, \alpha_{min}, \alpha_0, K_v, \theta$) from left to right are: (0, 0.1, 10, -10, 0, 0, 90), (0, 4, 3, -8, -6, 0.25, 0), (0.25, 0.1, 10, -10, 0, -0.75, 180), and (1, 4, 3, -8, 2, 0, 0).

10.1.3 Software Changes

Because this experimental setup was custom-built for the application, all experiments can be actuated with MATLAB integration. This required significant changes and addition to the software structure of the flapping robot. A structure chart for the software of the rotor experiments is in Figure 10.2. This simplified chart shows all separate code files and their roles, but does not include control flow or data flow symbols.

Filename headers in blue denote code that was written from scratch for the rotor experiments. Filename headers in green denote code that was significantly altered from its content in the flapping propulsion experiments. Filename headers in gray denote code that was minimally altered, if at all. Solid lines are internal code relationships, while dashed lines are serial connections and dotted lines are TCP

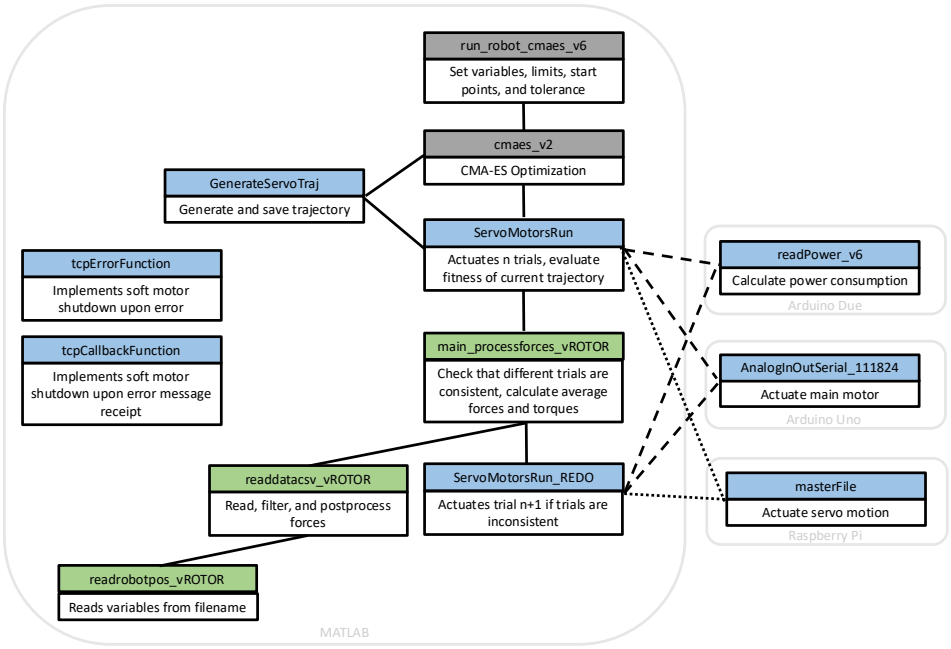

Figure 10.2: A structure chart for the rotor experiment software.

connections. Note that control of the force sensor is not integrated to the MATLAB script, and is therefore not shown in this structure chart.

10.2 Hardware

A model-scale IBC rotor was built for the purpose of this study, combining commercially-available hobby-grade components (Figure 10.3). The overall rotation of the system is actuated by a SAB DD 400 motor. This is the main motor from a Goblin RAW 420 model helicopter. It operates at 200Kv and is capable of producing approximately 10N lift at 2500rpm, using Goblin 420 blades. The blades used in our setup are Align 360mm 3G carbon fiber blades, which have a length of 360mm, chord of 33mm, and a symmetric airfoil cross-section.

The rotor hub contains four HiTec MD950TW-CAN servos compactly arranged around the central shaft. Each servo horn is attached via a ball-joint to a pitch linkage. At the other end of each pitch linkage, a second ball-joint attaches it to the pitch bearing at the root of a blade. The angular motion of the servo is transmitted through the rigid pitch linkage and results in angular motion of the blade. In this way, each servo motor individually controls the pitch of a single blade.

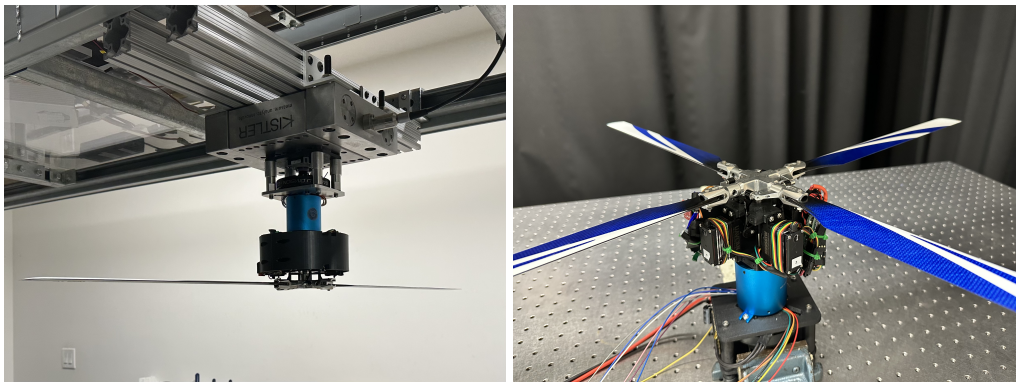

Figure 10.3: Images of the model-scale IBC rotor. The left image shows the rotor with two blades mounted to the dynamometer, as in experiment. The right image is provided for extra detail; for greater visibility in this image, the rotor is not mounted to the dynamometer, nor is the wiring shroud installed.

10.3 Electrical Systems and Microcontrollers

The experimental system is powered by two separate electrical subsystems, one to power the drivetrain and one to power the rotor hub. Due to the high-current requirements of the SAB DD 400 motor, a separate electrical subsystem is utilized to avoid any electrical interference with the servo operation.

10.3.1 Drivetrain Electrical System

The drivetrain is powered by a MeanWell UHP-1500-24 power supply, which can provide up to 1500W at 24V. The resulting 62.5A maximum current is sized to exceed the burst current rating of the electronic speed controller (ESC). We use the Castle Phoenix Edge HV40 ESC, which can provide 40A of constant current. A 60A ESC is recommended for the Goblin Raw 420 model helicopter, from which the main motor is derived. However, since our setup is encountering neither the high headspeeds nor the dynamic maneuvers of model helicopter flight, a 40A ESC is sufficient.

The ESC relays both power and speed commands to the main motor. Speed commands originate from a computer and travel over USB using serial communication to an Arduino Uno. The Arduino Uno translates this speed command to an ESC input.

The power consumption of the drivetrain electrical system is measured using a CT433-HSWF70MR TMR current sensor that is connected immediately downstream of the power supply. This electrically isolated current sensor is powered by

an Arduino Due. Its analog voltage output is measured by the Arduino Due at a rate of 1000Hz, with each measurement an average of 20 sensor readings.

10.3.2 Rotor Hub Electrical System

The rotor hub is powered by a generic 24V 8A power supply. Although this current rating is far below the stall current that can be drawn by the servos, no failure of actuation has yet been observed. This is expected, as the servos are actuating at low torque even during hub rotation.

A CT433-HSWF70MR TMR current sensor is connected inline immediately downstream of the power supply. As in the drivetrain electrical system, this current sensor is both powered and read by the Arduino Due. This allows for the calculation of the total power consumption of the experimental system, by the main motor and by the servo motors, during an experiment. This value is then transmitted to a computer via USB using serial communication.

In order to actuate the IBC servos, power and ground must be available on the rotor hub during rotation. A slip ring can be utilized to transmit signals from a non-rotating to a rotating apparatus. Because slip rings can introduce unwanted noise into any transmitted signals, only power and ground cross the slip ring. Prior to crossing the slip ring, the power and ground lines pass through a D24V150F7 voltage regulator. This 7.5V line directly powers the servos and also provides power to the onboard Raspberry Pi.

The onboard Raspberry Pi is powered by the 5V output of a D24V50F5 voltage regulator. It communicates via TCP over WiFi with a computer to receive servo commands. These commands are relayed to the servos via a generic RS485/CAN HAT.

Two power absorbing devices (PADs) are connected in-line with the hub's power and ground. PADs play a capacitative role in the system, smoothing the sudden spikes of current that occur when a servo experiences an increase in torque. These current spikes cause the Raspberry Pi to briefly lose sufficient power, resulting in unreliable behavior. The addition of the PADs ensures that the Raspberry Pi functions properly regardless of servo behavior.

10.4 Force Sensor

Force measurements are taken by a Kistler 9139AA dynamometer to which the experimental assembly is mounted (Figures 10.3 and 10.4). This piezoelectric six-

axis load cell is optimized for vibrational and dynamic measurements, which makes it ideal for observing aeroelastic phenomena.

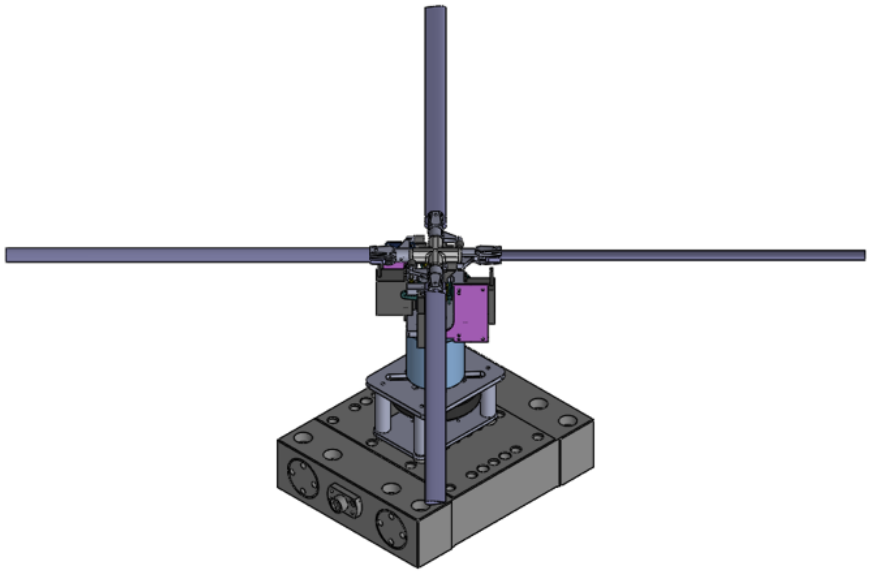

Figure 10.4: Kistler 9139AA dynamometer with experimental setup. Note that wiring, wiring shroud, and microcontrollers are not pictured.

10.5 Governing Nondimensional Parameters

In this section, we rederive the governing nondimensional parameters of helicopter flight, guided by [65]. A discussion of the relevancy of these parameters to the experimental setup follows therefrom.

10.5.1 Derivation

We begin by identifying the variables important to helicopter flight and their dimensions, where $[M]$ is mass, $[L]$ is length, and $[T]$ is time. We first list parameters that define the fluid environment through which the helicopter is flying. These include the dynamic viscosity μ, of dimension $[M][L^{-1}][T^{-1}]$; the density ρ, of dimension $[M][L^{-3}]$; the velocity V, of dimension $[L][T^{-1}]$; and the speed of sound a, of dimension $[L][T^{-1}]$. Parameters that define the rotor blades are a characteristic dimension R, of dimension $[L]$ (radius or chord, depending on convention); the material density σ, of dimension $[M][L^{-3}]$; and the elastic modulus E, of dimension $[M][L^{-1}][T^{-2}]$. In addition, we have the angular velocity of the rotor ω, of dimension $[T^{-1}]$, and gravity g, of dimension $[L][T^{-2}]$. We thus define nine

governing dimensional parameters of helicopter flight.

Some variables have been omitted from this derivation. Parameters governing structural damping are omitted from this treatment, since these are negligible in the case of a hinged rotor. The experimental setup discussed in this work incorporates lead-lag hinges and flapping blades, and therefore structural damping effects are minimal. Additional variables that more completely parametrize the structural behavior, such as the shear modulus, are also not considered. If these were to be considered, additional nondimensional groups relating such variables would appear. However, these nondimensional parameters have little effect on the main modes of aeroelastic behavior in helicopters.

We see that there are nine independent variables with three dimensions. By the Buckingham Pi theorem, there must exist six nondimensional groups that describe the system. When we replace l with the appropriate characteristic dimension (R or c), we arrive at

$$0 = f\left\{\frac{V}{\omega R}, \quad \frac{\rho V c}{\mu}, \quad \frac{V}{a}, \quad \frac{\rho}{\sigma}, \quad \frac{E}{\rho V^2}, \quad \frac{V}{\sqrt{gR}}\right\}. \tag{10.1}$$

These are, respectively, the advance ratio, the Reynolds number, the Mach number, the density ratio, the ratio of elastic force to aerodynamic inertia forces, and the Froude number. The advance ratio is calculated between the velocity of the oncoming flow V and the tip speed ωR. Matching this nondimensional parameter ensures that the relative motion between the rotor blade and the fluid is retained. The Reynolds number represents the ratio of inertial to viscous forces. It defines whether the regime of the flow is laminar, transitional, or turbulent. The Mach number dictates how important compressibility effects are in the flow. As the blade tips approach the speed of sound, compressibility effects become more important. The density ratio describes the ratio of aerodynamic to structural inertia. This ratio describes whether forces of structural origin, like centrifugal force or rotor oscillation, are important compared to forces of aerodynamic origin. Lastly, the Froude number compares the elastic force produced by the blade structure to the aerodynamic inertia forces. The Froude number determines the relative importance of elastic blade behavior.

It should be noted, provided one assumes a constant slope for the lift curve, that the Lock number is proportional to the density ratio ρ/σ. The Lock number, developed by C.N.H. Lock in 1927 [90], is more commonly utilized in the literature to describe

the ratio of aerodynamic to inertial forces. To show this proportionality, we multiply ρ/σ by a ratio of equal dimensionality:

$$\frac{\rho}{\sigma} \propto \frac{\rho}{\sigma} * \frac{[L^5]}{[L^5]}. \tag{10.2}$$

Because σ is the density of the blade, we know that

$$\sigma \propto m/[L^3] \tag{10.3}$$

and therefore

$$\frac{\rho}{\sigma} \propto \frac{\rho}{m} * \frac{[L^5]}{[L^2]}. \tag{10.4}$$

The moment of inertia of a given body of mass m about any arbitrary axis can be defined as

$$I = mk^2, \tag{10.5}$$

where k is the effective radius. We therefore replace the denominator of 10.4 with I. Furthermore, if the blade lift curve has a constant slope, which we may assume away from stall, then we can multiply the right side by $C_{L\alpha} = \frac{\partial C_l}{\partial \alpha}$ without loss of proportionality:

$$\frac{\rho}{\sigma} \propto \frac{\rho C_{L\alpha}[L^5]}{I}. \tag{10.6}$$

This results in the density ratio being proportional to the Lock number γ:

$$\frac{\rho}{\sigma} \propto \frac{\rho C_{L\alpha} c R^4}{I} = \gamma, \tag{10.7}$$

where c is the blade chord, R is the radius of the rotor, and I is the blade moment of inertia about the flapping hinge. This particular form of the Lock number is derived from a Fourier analysis of the flapping angle in [90], since the ratio of aerodynamic to inertial forces mainly affects the flapping motion of the blades. However, some controversy exists over the formulation in 10.7 as it does not accurately reflect the effect of each parameter [129]. For example, the effect of radius on Lock number is linear, whereas it appears to have a quartic relationship in the traditional definition.

10.5.2 Relevancy to Experimental Setup

It is impossible to exactly match all six nondimensional parameters in a sub- or super-scale helicopter model. This is proven in [65], and we rederive some important conclusions here.

First, the impossibility of matching all nondimensional parameters is most easily seen in the conflict between the Mach number and the Froude number. If we define

the scale factor, or the division of the model value by the full-scale value, as λ, then we have

$$\frac{V_{model}}{a_{model}} = \frac{V_{full-scale}}{a_{full-scale}}, \quad or \quad 1 = \frac{\lambda_V}{\lambda_a} \tag{10.8}$$

for the Mach number and

$$\frac{V_{model}}{\sqrt{g_{model}R_{model}}} = \frac{V_{full-scale}}{\sqrt{g_{full-scale}R_{full-scale}}}, \quad or \quad 1 = \frac{\lambda_V}{\sqrt{\lambda_g\lambda_R}} \tag{10.9}$$

for the Froude number. We assume that the model is being tested under the same gravity as the full-scale helicopter, meaning $\lambda_g = 1$. Furthermore, the speed of sound does not vary significantly under normal operating conditions. Thus $\lambda_a = 1$ as well. By substituting these values, we find that the Mach number scaling requires $\lambda_V = 1$ while the Froude number scaling requires $\lambda_V = \sqrt{\lambda_R}$. These equations cannot be simultaneously satisfied unless $\lambda_R = 1$, meaning the model and full-scale are the same size. Due to this impracticability, model-scale helicopters are designed either to match the Mach number or to match the Froude number.

An additional challenge derives from attempting to match both the Reynolds number and Mach number to ensure full aerodynamic similarity. Using scale factor analysis on the Reynolds number, we find that

$$1 = \frac{\lambda_\rho\lambda_V\lambda_R}{\lambda_\mu}. \tag{10.10}$$

This equation, combined with 10.8, results in the conclusion that

$$\lambda_R = \frac{1}{\lambda_\rho} \tag{10.11}$$

because μ is often approximated as invariant with respect to pressure (for example, the dynamic viscosity of nitrogen changes less than 1% between 1 and 10 atm [26]). This means that a $\frac{1}{10}$ scale model would require testing at ten times nominal fluid density. Some success has been found in testing subscale models in air under high pressures and/or using heavy gases like refrigerant. However, most experimental campaigns accept a relaxation in one or more nondimensional parameters, based on relative importance of their effects, to enable testing at subscale in STP air.

10.5.3 Relaxing Dynamic Similarity

In this section, we consider whether matching the value of each nondimensional parameter to that of full-scale operation is strictly necessary for our application. It is important to emphasize that the goal of this experimental setup is not to determine a

trajectory that can be utilized without modification on a real-world helicopter. That would require matching nearly all nondimensional parameters exactly. Instead, we seek a proof-of-concept solution, where the values of nondimensional parameters suggest similar acceptable aerodynamics. From the perspective of these considerations, we discuss how nondimensional parameters have been matched or made acceptable in the experimental setup.

10.5.3.1 Advance Ratio

The advance ratio is one nondimensional parameter which must be matched to full-scale operation, regardless of application. This is due to the physical significance of the advance ratio. The advance ratio determines the relative velocity direction and magnitude (apparent velocity) over a helicopter blade at any phase position during the cycle (see Figure 10.5). If the advance ratio is not matched, the apparent velocity encountered by each blade is completely different and therefore not representative of full-scale operation. In this experimental setup, we consider hovering only, i.e. $V = 0$. The advance ratio is therefore matched for hovering at any scale.

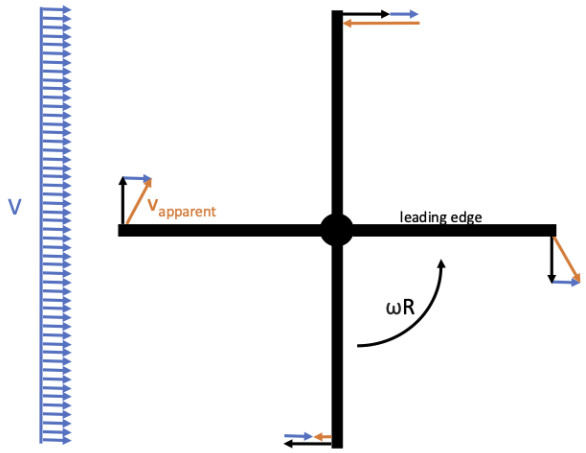

Figure 10.5: Graphical representation of how the advance ratio, $\frac{V}{\omega R}$, affects the apparent velocity over the rotor blades.

10.5.3.2 Reynolds Number

The Reynolds number, as it is described with reference to forward flight speed V, is technically matched at hover. However, a more appropriate definition of the Reynolds number in this operating condition would be in reference to the tip speed

ωR. To exactly match the Reynolds number, the quantity ωR therefore needs to be equal between the experiment and the full-scale helicopter. If we describe the ratio between the model and full-scale radius by λ_R, as in section 10.5.2, we thus determine that

$$\lambda_R = \frac{1}{\lambda_\omega}. \tag{10.12}$$

If we take the Apache helicopter as our full-scale, which has a 7.315 meter radius [12] and operates at 289 rpm [60], we would therefore need our experimental apparatus of 0.417 meter radius to operate at 5,070 rpm. This is not feasible with our experimental setup; model helicopters of this size rarely exceed 3,000 rpm, even during sprint races and 3D maneuvers. This is primarily due to structural and weight limitations. It is common practice for rotors to not accurately match the Reynolds number [133], so long as it is in the correct regime and is as close as experimentally feasible. Our experiment operates at approximately 525 rpm[1], which results in a Reynolds number of 6,527. This flow is in the turbulent regime and is therefore acceptable.

10.5.3.3 Mach Number

A similar situation exists for Mach number, where the most appropriate value to consider matching during hover is that of the tip speed Mach number, where ωR replaces V. It can often be assumed that flight dynamics and performance do not change greatly for tip speed Mach numbers at or below 0.6, and therefore in this range matching the Mach number can be relaxed [6]. Mach number effects are minimal for strictly subsonic flow, where incompressibility is a valid assumption. However, the true speed of the fluid at the blade tip depends on both the helicopter speed and the phase position of the blade. During a single cycle, the blade tip will experience a fluid speed of at most $\omega R + V$ (Figure 10.5). This means that, despite a seemingly low tip speed Mach number, transonic effects can be seen at high advance ratios. In fact, transonic effects are usually only seen in full-scale helicopters at high advance ratios [70], since full-scale helicopters operate at mid-subsonic tip speed Mach numbers (0.58-0.66) during hover [16]. Given that the experimental setup is designed to investigate a low advance ratio (hover), and that full-scale helicopters operate on the order of Mach 0.6 at hover, we may relax the

[1]The ESC maintains constant power to the motor, meaning that RPM depends on the blade pitch. At higher pitch, more torque would be needed for the same RPM. Higher torque cannot be achieved without increasing the power draw, and thus RPM decreases with increasing pitch. 525 RPM is an approximate average value seen in experiment.

Mach number restriction in our experiment. The tip speed Mach number of the experiment is below 0.6, ensuring appropriate similarity.

10.5.3.4 Density Ratio or Lock Number

The density ratio $\frac{\rho}{\sigma}$, or equivalently the Lock number, can be assumed approximately equal in both experiment and full-scale. λ_ρ is already equal to one, since we are conducting our experiment in STP air. We assume λ_σ to also be one since the materials used in the model helicopter blade are not fundamentally different than the materials in full-scale helicopter blades. Carbon-fiber composites are used in both. It should, however, be noted that the structural density does change based on composite makeup; varying proportions of carbon fiber and resin will result in varying structural density. We assume that this change is negligible given the minimal change in elastic modulus between composite model and full-scale helicopter blades [107, 159], which also depends on the proportion of fiber to resin.

10.5.3.5 Ratio of Elastic Force to Aerodynamic Force

The ratio of elastic force to aerodynamic force is redefined as $\frac{E}{\rho\omega^2 R^2}$ during hover. To match this ratio, we require

$$\lambda_E = \lambda_\omega^2 \lambda_R^2. \tag{10.13}$$

From our analysis of the Mach number, we see that we can relax strictures on λ_ω and λ_R, provided that $\lambda_\omega \lambda_R \leq 1$. Therefore, we know $\lambda_E \leq 1$. Model helicopter blades are assumed to have an elastic modulus of $69 * 10^3$ MPa [159] while full-scale helicopter blades with a balanced symmetric laminate have an assumed elastic modulus of $76 * 10^3$ MPa [107]. This results in $\lambda_E = 0.9$. However, these values may not be representative of all model helicopter blades nor all full-scale helicopter blades, due to differences in material and composite layups. It is therefore difficult to determine whether $\lambda_E = 0.9$ is accurate. Regardless, if $\lambda_E = 0.9$ then by equation 10.13, $\lambda_\omega \lambda_R \approx 0.95$. In our experimental system, ωR is 10% of the full-scale value. This does not match the λ_E scaling shown here. Therefore, the experimental system will have a different ratio of elastic force to aerodynamic inertia force.

10.5.3.6 Froude Number

The Froude number is also redefined as $\omega\sqrt{\frac{R}{g}}$ during hover. For the Froude number to be matched in an experiment under constant gravity, $\omega\sqrt{R}$ must be matched. This

results in

$$\lambda_R = \frac{1}{\lambda_\omega^2}. \tag{10.14}$$

However, this is in conflict with the Reynolds number matching requirement of equation 10.12 unless $\lambda_\omega = 1$ and thereby $\lambda_R = 1$. Because we prioritize achieving as close to an accurate Reynolds number as possible, in order to best achieve aerodynamic similarity, we must allow the Froude number to not match its full-scale value. In our experimental setup, the Froude number is 1.80 whilst at full-scale it is 4.16, once again using values from the Apache helicopter. Aerodynamic similarity is prioritized over a full treatment of elastic blade behavior, especially since the weight of the model can be shown to be negligible following a similar analysis to [65].

We have thus shown that our experimental setup produces acceptable values for the advance ratio, Reynolds number, tip speed Mach number, Lock number, ratio of elastic force to aerodynamic inertia force, and Froude number. Whilst acceptable, it should be emphasized that the disparity in degree of fidelity to each of these nondimensional parameters will undoubtedly affect transference of the experimental conclusions to full-scale helicopter operation. For example, the maximum lift coefficient can vary by 10% when the Reynolds number increases by an order of magnitude, even though both experiments are in the turbulent regime [133]. However, the goal of this work is not to prove that any experimental conclusions may be implemented immediately on full-scale helicopters. Rather, it is to provide a proof-of-concept for new modes of pitch actuation that may prove fruitful with further investigation on more accurate models.

10.6 Experimental Justification

In order to determine whether this setup can accurately optimize for the desired outcome, several preliminary experiments were run. These justification experiments and results are discussed throughout the following section.

10.6.1 Characterization

The characteristics of the data feeding into the CMA-ES optimization must be fully understood in order to quantify the quality of such data and thereby determine the necessary tolerances of the algorithm. There are two streams of data that enter into the CMA-ES algorithm. The first is data from the two current sensors, and the second is data from the force sensor. With this data, an average force and an average work are calculated. These averages are used in the fitness function of CMA-ES.

10.6.1.1 Error of Current Sensors

The error in average work can be calculated given the known resolution of the sensor readings. The average work over five seconds is calculated by a right-handed Riemann sum of power with current measurements at 1000 Hz and an assumed voltage of 24V. We thus find that, assuming no error in time measurement,

$$\sigma_W = \sqrt{\sum_{i=1}^{N} \left((V_i \ \sigma_I \ dt)^2 + (I_i \ \sigma_V \ dt)^2 \right)}, \tag{10.15}$$

where N is the total number of measurements used in the Riemann sum and dt is the time between those measurements. Each current measurement used in the Riemann sum is calculated by averaging 20 consecutive readings of the sensor (300 ns response time). Given a sensor reading resolution of $\frac{140}{4096}$A, we can determine the error of a single current measurement used in the Riemann sum:

$$\sigma_I = \frac{140}{4096\sqrt{20}} = 7.64 \times 10^{-3} A. \tag{10.16}$$

Knowing that the 24V power supply has a voltage tolerance of ±1%, we thus can substitute values into equation 10.15:

$$\sigma_W = \sqrt{2 * \sum_{i=1}^{N} \left(\left(24V * 7.64 \times 10^{-3} A * 0.001s\right)^2 + \left(I_i * 0.24V * 0.001s\right)^2 \right)}. \tag{10.17}$$

Although we do not explicitly know each I_i, the maximum value of I_i that the sensor can read is $70A$. We use this value to provide an upper bound $\sigma_W = 1.72J$. Since a single candidate trajectory is evaluated three times and averaged, $\sigma_{W,avg} = \sigma_W/\sqrt{3} = 0.993J$.

10.6.1.2 Allan Variance

The force sensor has several different sources of measurement uncertainty, including white noise, flicker noise, and drift. These uncertainty sources change in relative importance to one another depending on the timescale of measurement. For example, the difference in identical measurements that are very far apart in time will be dominated by drift, while very short time measurements will be dominated by white noise. The Allan variance uses a null data set from the force sensor to determine the stability of measurement over an observation time τ. The Allan variance σ_y^2 is

defined as

$$\sigma_y^2(\tau = nT_s) \approx \frac{1}{2(nT_s^2)(N-2n)} \sum_{i=1}^{N-2n} (x_{i+2n} - 2x_{i+1} + x_i)^2, \tag{10.18}$$

where nT_s is an integer multiple of the sampling period and N is the length of signal x. The approximate equality is due to the necessity of using numerical integration techniques for discretely sampled signals.

To calculate the Allan variance of our force sensor, 5 minutes of data was taken with no motion of the experimental setup. Because of the integrated sensor zeroing, this data should correspond to $\{F_x, F_y, F_z, T_x, T_y, T_z\} = 0$. There is no further postprocessing of the data prior to calculating the Allan variance. Figure 10.6 shows the Allan deviation σ_y for all forces and torques.

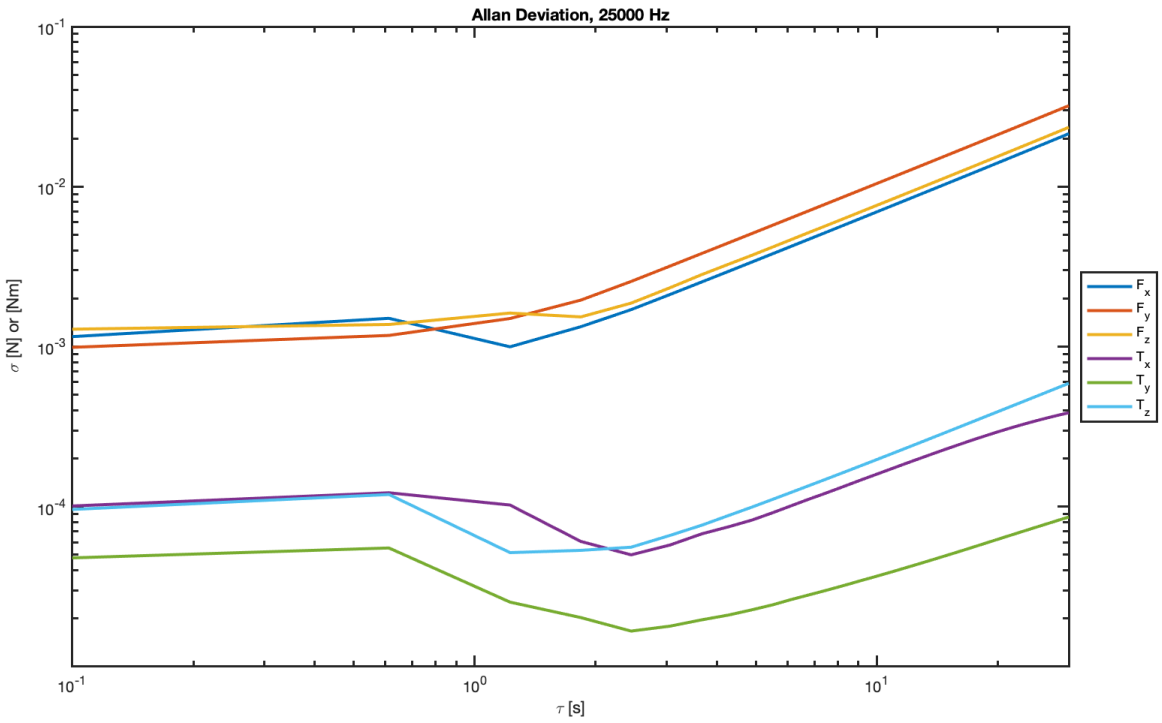

Figure 10.6: Allan deviation of all forces and torques.

The Allan deviation is useful in two regards. First, it can tell you how long you should average sensor readings to minimize error. There is a minimum value of Allan deviation where neither the white noise nor the drift are dominant. The corresponding value of τ is the optimal averaging time. In Figure 10.6, we see

that this force sensor does not experience a steep rise in noise at lower and lower averaging times due to white noise. Instead, the deviation has a flat, but slightly positive, slope below approximately $\tau = 0.61s$. This makes it more difficult to define the optimal averaging time, since it does not necessarily correspond to the global minimum. However, the slope after approximately $\tau = 2.5s$, dominated by drift, is still clearly seen. We therefore define the optimal averaging time as the minimum reached closest to this slope's inception. We can thus see slight differences between the optimal averaging times of F_x, F_y, F_z, T_x, T_y, and T_z. For example, the minimum $\sigma_{F_x} = 9.97 \times 10^{-4}N$ occurs at $\tau = 1.22s$ while the minimum $\sigma_{F_z} = 1.53 \times 10^{-3}N$ occurs at $\tau = 1.84s$.

Second, the Allan deviation can put a bound on the error at a given time τ after the sensor has reset its measurement. For example, if a measurement of T_x is taken at $\tau = 4.898s$ after sensor reset, we expect the deviation of this value to be $8.25 \times 10^{-5}Nm$ when sampling at 25,000 Hz. The Allan deviation thus reveals that the optimal thrust measurement from this sensor would be taken centered at $1.84s$ with $1.84s$ of averaging time. In practice, since our experimental setup needs to be started from rest, measurements will be 8 or more seconds from sensor reset. The calculated Allan deviation at $\tau = 8s$ is at maximum $8.3 \times 10^{-3}N$, attained by F_y, and even for a 20s averaging time, leading to $\tau = 28s$, remains below $3 \times 10^{-2}N$. This is less than 2% of the full-scale force value, and therefore the force sensor noise itself is acceptable.

10.6.1.3 Experiment Standard Deviation

Because the actuation of the IBC system will introduce noise due to vibration that will compound the sensor noise, we also measure the standard deviation of a mock single experiment. We run 52 trials with two blades at pitch angle $\alpha = 10°$ and 525 RPM. We compare the standard deviation of this repeated mock single experiment with the standard deviation of the same experimental flow, but with placement of a 200g weight instead of rotor actuation. 31 trials were measured for the 200g weight. The average values of force and torque are calculated over a range of averaging times in increments of $0.1s$ from $0.1 - 15s$ and reported on a $\sigma - \tau$ graph similarly to the Allan deviation. We investigate three different methods of selecting the averaging time; an averaging window whose center is at 17s, an averaging window whose beginning is at 8s, and an averaging window whose end is 28s. This enables us to compare the effect of a larger averaging time and the effect of a longer time until the

start of measurement, both of which we would expect to increase error due to drift.

Given the prior results of the Allan deviation, which show a significant contribution to error from drift, we also compare the raw data to two possible drift compensation protocols. In the simple drift compensation, a linear fit from the initial reading to the final reading is subtracted from the data. In the averaged drift compensation, the linear fit is taken from the average of the first 1.8s of data and the average of the last 1.8s of data. This averaging time was chosen because it has the minimum thrust error according to the Allan deviation plot. Note that we do not expect drift compensation to adequately address all forms of drift, since some types of drift are nonlinear; thus, we still expect error to overall increase with time.

Figure 10.7 compares the standard deviations of all forces and torques for each postprocessing procedure in the test matrix composed of three averaging methods, three types of drift compensation, and two different experiments.

The standard deviation values, both for the rotor experiment and 200g mass, are somewhat large compared to the measured lift at 525 RPM and the mass' known weight. Most methods have at least one measured force or torque whose minimum standard deviation is above 0.1, meaning the standard deviation is above 5%. Since these experiments were done with a pitch of 10°, the percentage of measured lift will only increase at lower angles. To improve the standard deviation further, we pass the data through a 5Hz low-pass finite impulse response (FIR) filter prior to any drift compensation. This filter was designed using the Kaiser window technique to return a minimum-order, linear-phase FIR filter satisfying 60dB stopband attenuation.

In comparing Figure 10.7 to Figure 10.8, we see that all experimental results including filtering and a method of drift compensation attain a standard deviation less than 0.1N for averaging times above approximately 0.6s. Not only that, but the simple drift compensation tends to outperform the drift compensation using a 1.8s average time. Since we average three trials to obtain the value fed into the optimization, the standard error of the mean is thus less than 0.044N. We therefore validate that an average force attained by taking a 5-second average starting at 8 seconds into the filtered force measurement will have a standard deviation less than 5% of full-scale.

10.7 Optimizations with Known Outcomes

To determine that the optimization is functioning as expected, we conduct two experiments where the outcome is known a priori. The first experiment maximizes lift with a single angle of attack for both blades as the independent variable (α_{max} =

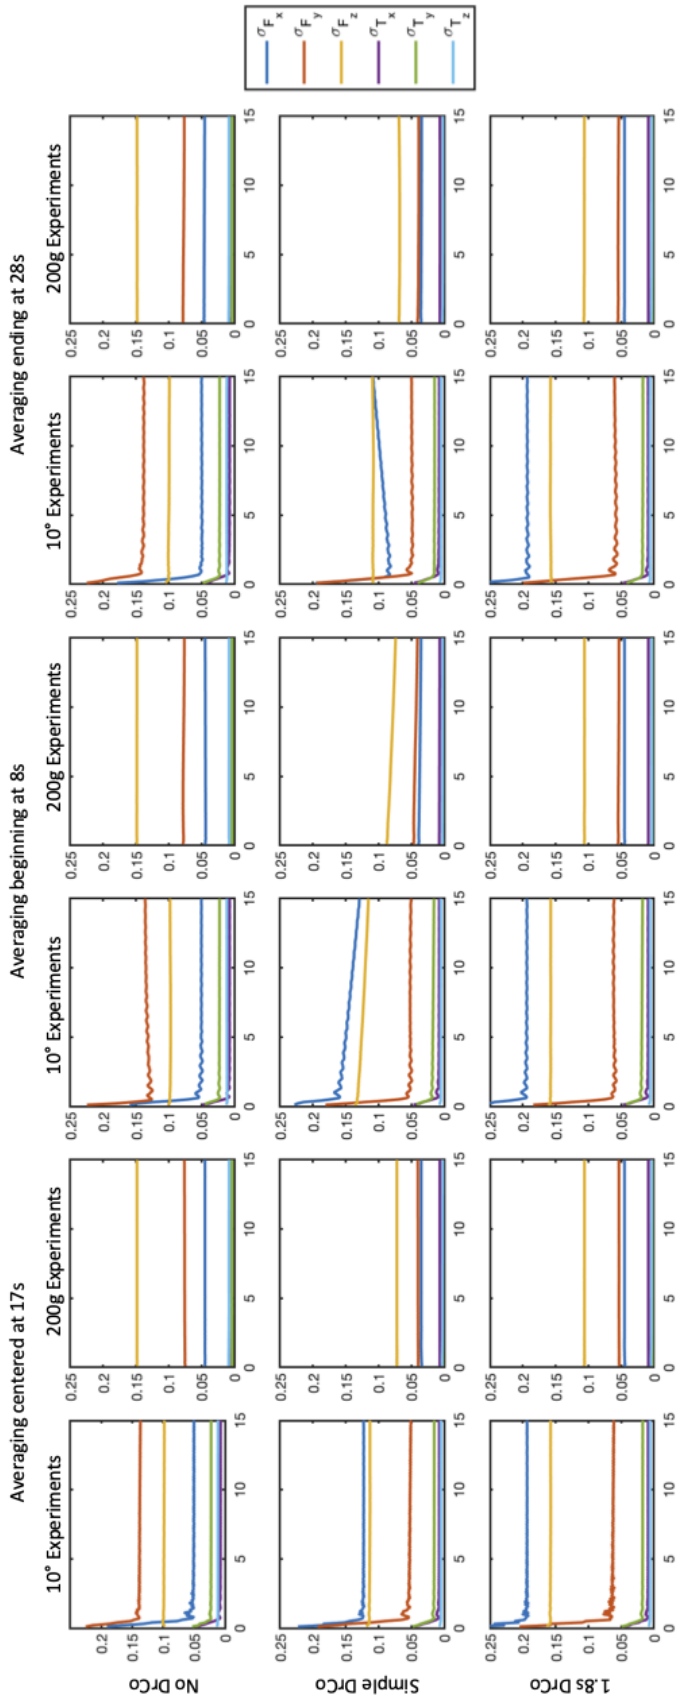

Figure 10.7: Plots of standard deviations of all forces and torques for varying postprocessing procedures, no filtering.

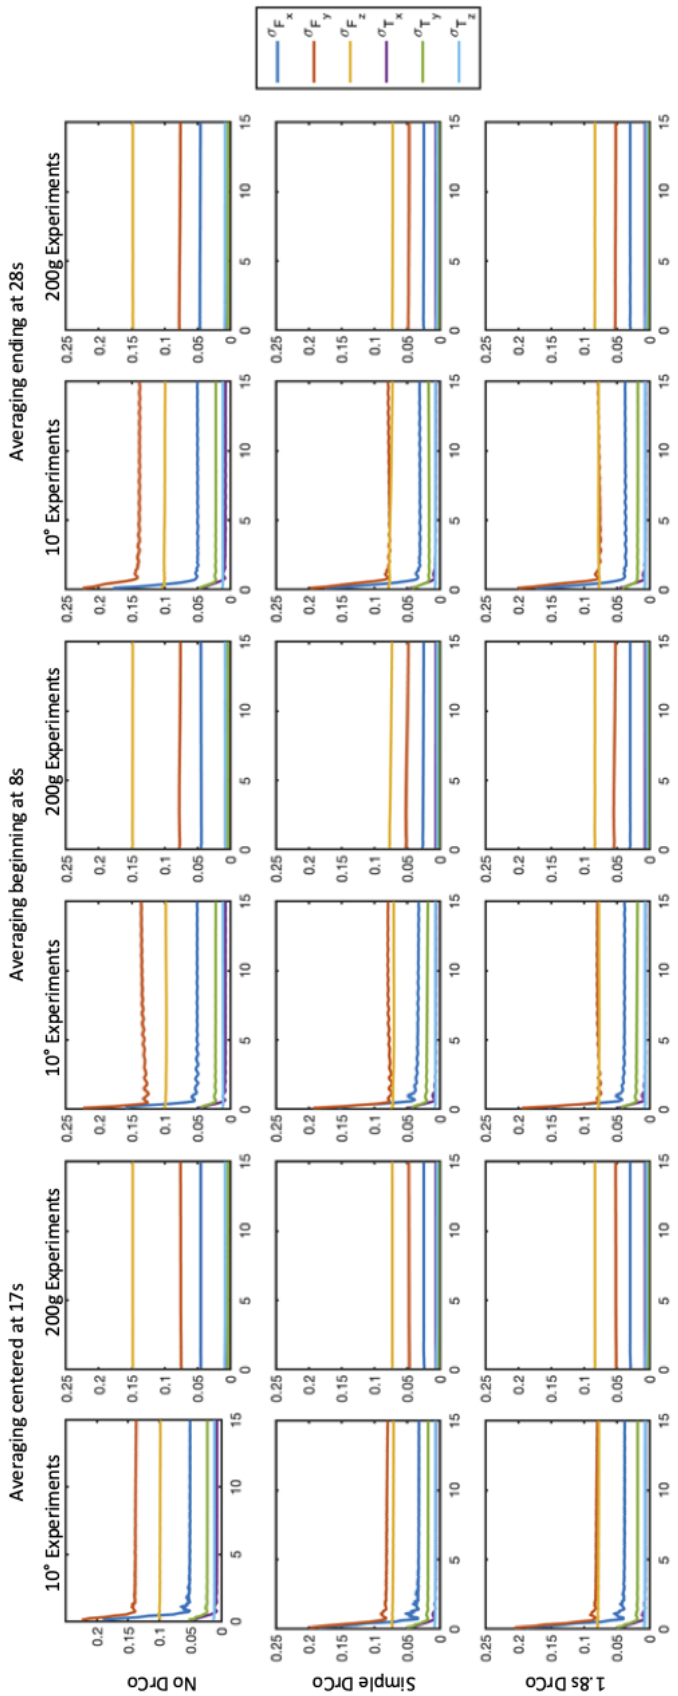

Figure 10.8: Plots of standard deviations of all forces and torques for varying postprocessing procedures, 5Hz low-pass filtered data.

$\alpha_0 = \alpha_{min}$). All other variables remain constant at 0. The second experiment maximizes lift per power with the same independent variable α.

10.7.1 Optimization for Maximum Lift

In previous testing of the experimental setup with two blades, F_z had displayed typical maximum values around 1.25N. We therefore choose the fitness function

$$f = \frac{|1.75 - F_z|}{1.75} \tag{10.19}$$

which optimizes for a lift force of 1.75N. Given that 1.75N is greater than the previously observed maximum, this fitness function will optimize for maximum lift.

In this experiment, the optimization was given a single angle of attack to optimize. This angle of attack was constant and the same for both blades. It ranged from -15 to 15° with a tolerance of 0.02°. The optimization converged to within the AOA tolerance after 43 generations with four candidate solutions per generation. The evolution of the fitness function is shown in Figure 10.9, the evolution of AOA in Figure 10.10, and the evolution of the produced forces and torques in Figure 10.11.

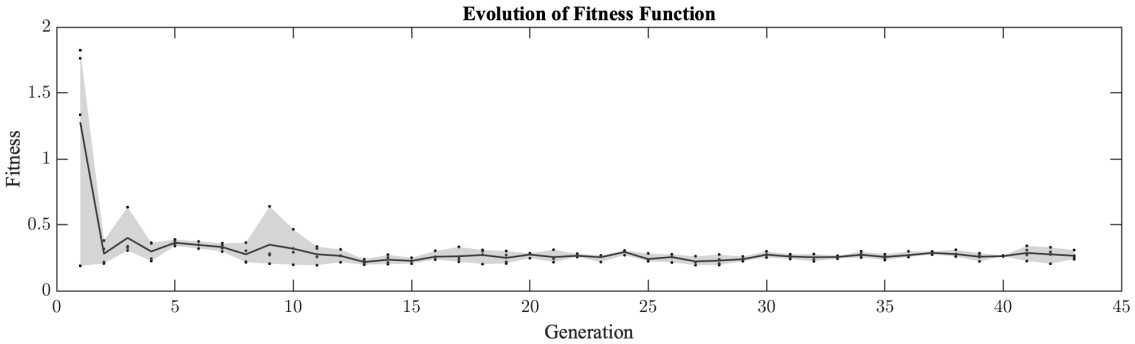

Figure 10.9: Evolution of the fitness function for the optimization for maximum lift. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

The optimization converges to an angle of attack of 9.49°. This is lower than the stall angle of the model helicopter blades, which is expected behavior. Because the ESC provides constant power to the motor, higher angles of attack cause the motor RPM to decrease and therefore the produced lift to decrease. 9.49° is consistent with the observed maximum force produced in trials outside of the optimization.

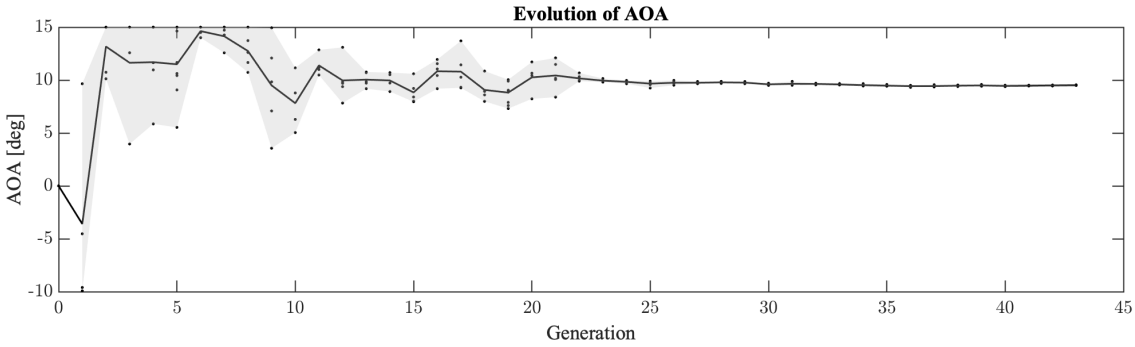

Figure 10.10: Evolution of the angle of attack for the optimization for maximum lift. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

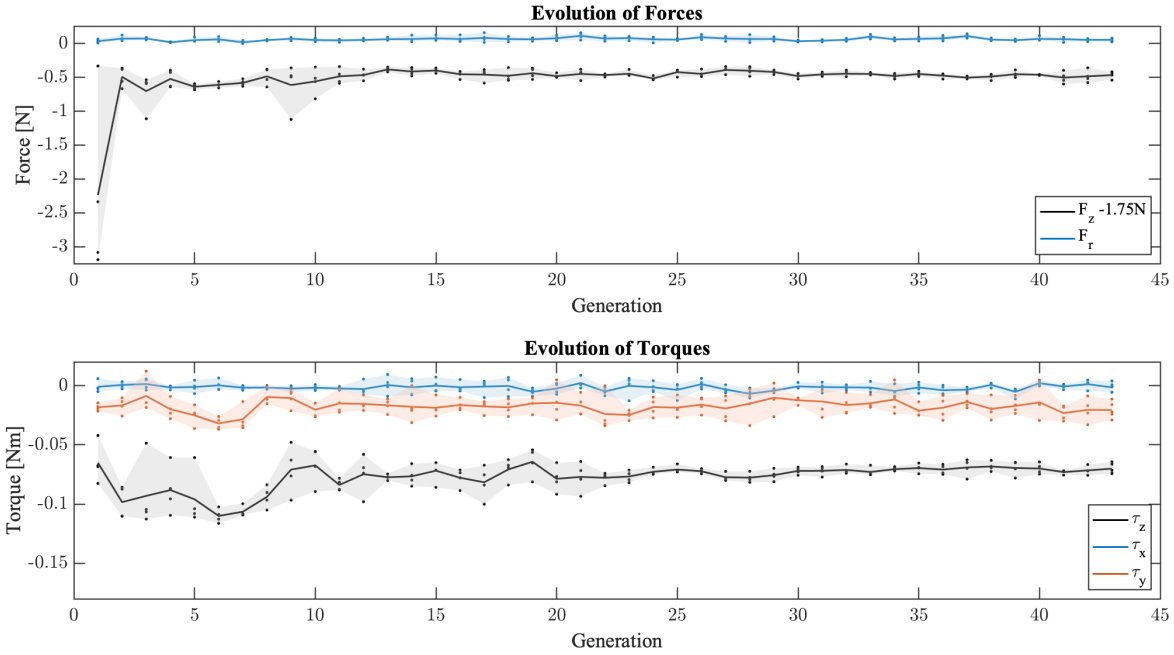

Figure 10.11: Evolution of the forces and torques for the optimization for maximum lift. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

10.7.2 Optimization for Maximum Lift-to-Power Ratio

In this experiment, we choose the fitness function

$$f = \left| \frac{W_{hub}}{F_z} \right|, \tag{10.20}$$

where W_{hub} is the total work used by the hub (mainly servo actuation) and F_z is the force produced in the z-direction (thrust). We ignore W_{motor} because the ESC in this setup provides constant power to the motor. Equation 10.20 optimizes for the lowest work-to-lift ratio (in other words, maximizing the lift per work) and is symmetric about zero pitch.

Due to the symmetry of the fitness function, we restrict the pitch to $0 < \alpha < 15$ degrees with a tolerance of 0.02°. The optimization converged to within the AOA tolerance after 52 generations with four candidate solutions per generation. The evolution of the fitness function is shown in Figure 10.12, the evolution of AOA in Figure 10.13, and the evolution of the produced forces and torques in Figure 10.14.

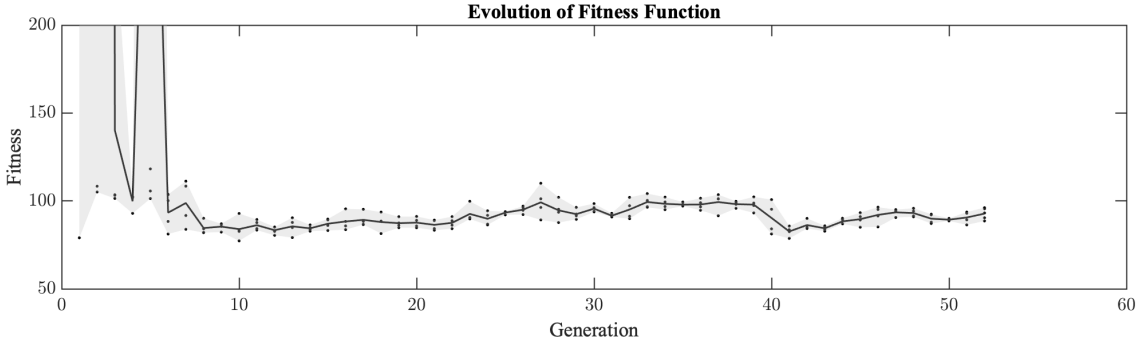

Figure 10.12: Evolution of the fitness function for the optimization for maximum lift per power. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

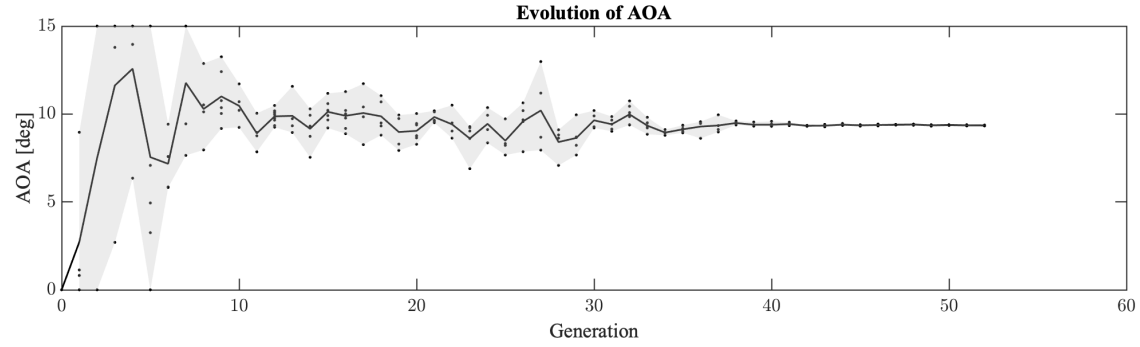

Figure 10.13: Evolution of the angle of attack for the optimization for maximum lift per power. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

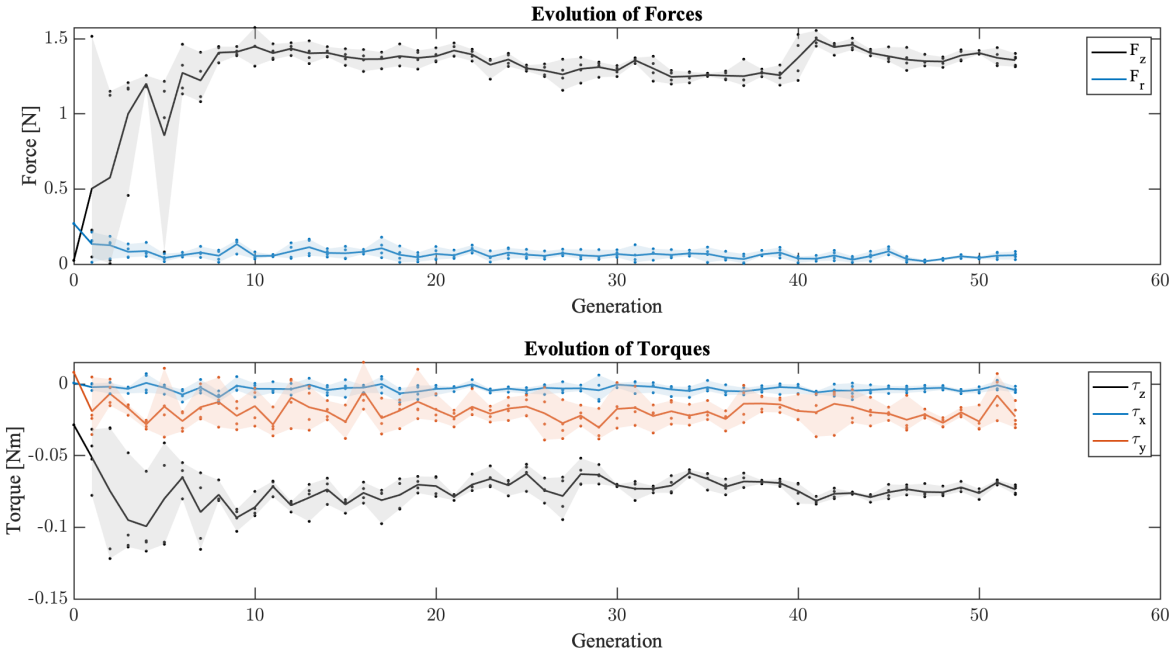

Figure 10.14: Evolution of the forces and torques for the optimization for maximum lift per power. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

The optimization converges to an angle of attack of 9.36°. This is slightly lower than the angle of attack that optimized lift alone (9.49°). Again, this result is expected because of the added influence of W_{hub} on the fitness function. The servo motors require power to maintain their position against the aerodynamic torque acting upon them. The higher the angle of attack, the higher the aerodynamic torque. It is thus expected that including the servo power would reduce the optimal angle of attack. The power does not change by a large amount with angle of attack, ranging from 115.1 to 127.7 Watts. The small decrease in angle of attack is thus reasonable.

Chapter 11

FORCE VECTORING OPTIMIZATIONS

In this chapter, we present preliminary results of two optimizations. Both optimizations are designed to vector thrust of the system, but in two different settings. The first optimization investigates small-magnitude thrust vectoring, likely to be implemented on a longer time horizon, while the second investigates large-magnitude thrust vectoring. Large-magnitude thrust vectoring is likely to be implemented on a much shorter timescale to enable high maneuverability.

11.1 Efficient Steep Takeoff/Shallow Approach Vectoring

In the first preliminary experiment, we ask the rotor to vector thrust by approximately 5°. Vectoring at this particular angle is useful for two very different scenarios in helicopter operation. The first is a steep takeoff, where the angle of ascent is near-vertical. While steep takeoffs are sometimes necessary for obstacle avoidance, they are a particularly dangerous maneuver due to lack of translational lift. The second scenario is a shallow approach angle, typically 3-5° [3]. This approach angle maximizes translational lift and is therefore considered the safest approach angle. Vectoring thrust at 5° thus is applicable to both highly risky maneuvers as well as safe, everyday maneuvers. We include a measure of efficiency in the optimization in order to select for efficient trajectories, since these maneuvers are implemented over a long timespan.

To achieve this angle of thrust vectoring, the fitness function was chosen to be

$$f = 0.8 \times \left(\frac{|F_{z,\,target} - F_z|}{F_{z,\,target}} + \frac{|F_{r,\,target} - F_r|}{F_{r,\,target}} \right) + 0.2 \times \frac{W}{1000}. \tag{11.1}$$

The first term, modified by prefactor 0.8, enforces the closeness of the produced force to a desired target force. The value of $F_{z,\,target}$ was chosen to maximize the thrust produced. The value of $F_{r,\,target}$ was chosen in reference to $F_{z,\,target}$, to achieve an approximately 5° angle between the rotor axis and the force produced. The second term, modified by prefactor 0.2, weights the optimization towards efficient behavior. To match the magnitude of the first term, such that the prefactors 0.8 and 0.2 represent the weighted average, the work was divided by 1,000.

Three variables participate in the optimization. These are the frequency of blade actuation and the phase shift of both blades (f, θ_1, θ_2). The maximum, minimum, and zero-crossing AOA are fixed at $\alpha_{min} = 4$, $\alpha_0 = 6.5$, and $\alpha_{max} = 9$. These values were chosen based on the thrust values produced by the optimizations with known outcomes, as well as the typical range of IBC actuation. The ranges, tolerances, and initial values of all variables are reported in table 11.1.

Variable	Initialization	Range	Convergence Criterion
f [Hz]	2	2–35	0.1
θ [°]	0	0–360	1
K_v	0	-0.99–0.99	0.05

Table 11.1: Initialization, range, and convergence criteria of the variables participating in preliminary force vectoring optimizations.

Although the optimization has not yet converged to within the defined experimental parameters, we comment on some of the preliminary results seen here. The evolution of the fitness function is shown in Figure 11.1, the evolution of the frequency and phase shifts in Figure 11.2, and the evolution of the produced forces and torques in Figure 11.3.

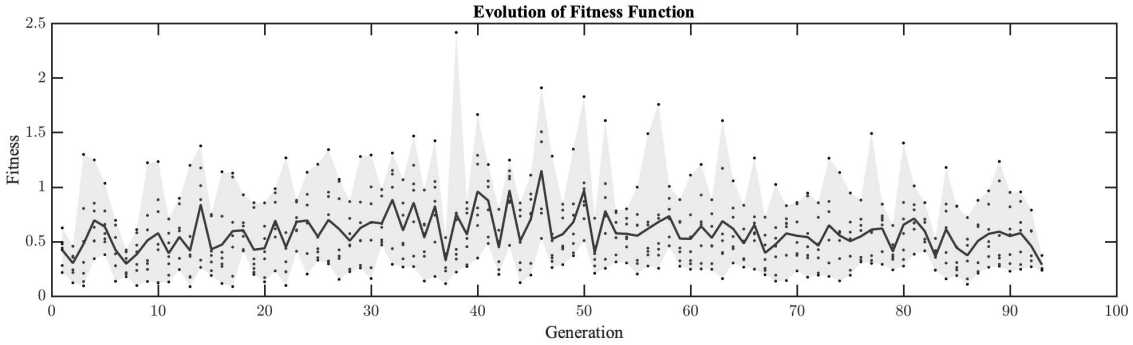

Figure 11.1: Evolution of the fitness function for the optimization for 5° of thrust vectoring. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

As seen in Figure 11.2, the optimization hones in on the optimum value for the first blade's phase shift by generation 50. This optimum value is a phase shift of 0. This is not a particularly surprising result; because the frequency of actuation is high with respect to the measurement time and the actuation is sinusoidal, the absolute phase shifts have little effect on the produced forces and torques. The relative phase shift is what is relevant to change force and torque production.

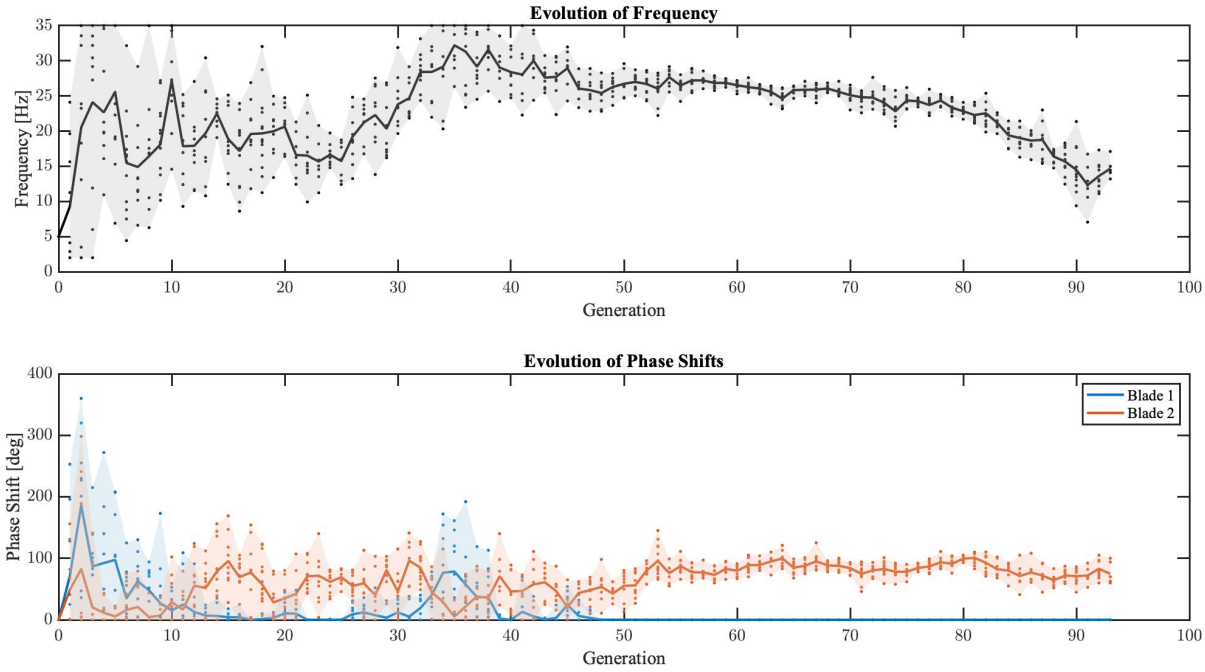

Figure 11.2: Evolution of the angle of attack for the optimization for 5° of thrust vectoring. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

From generation 54 onward, the relative phase shift has found a narrower region in which to optimize— a local minimum along the phase shift dimensions. The mean value and variance do not change by a large portion of the variable's range until at least generation 80. The relative phase shift generational average ranges from 73° to 99.75°, with an average of 85°, during this time. Although the relative phase shift does not have a distinct direction of optimizing during this period, the frequency is declining. We thus see that while a local minimum has likely been located in the phase shift plane, this local minimum is not heavily affected by the frequency of actuation. In this local minimum, the mean radial force stays relatively consistent while the thrust force changes (Figure 11.3). In comparing the average radial force over these generations to the maximum and minimum generational average thrust forces attained, we see thrust vectoring of 5-6°. The optimization is thus fulfilling our stated purpose; to achieve approximately 5° of thrust vectoring. This is achieved via phase-shifting the blades by on average 85° relative to each other, but has significant flexibility in terms of both frequency of actuation and range of relative

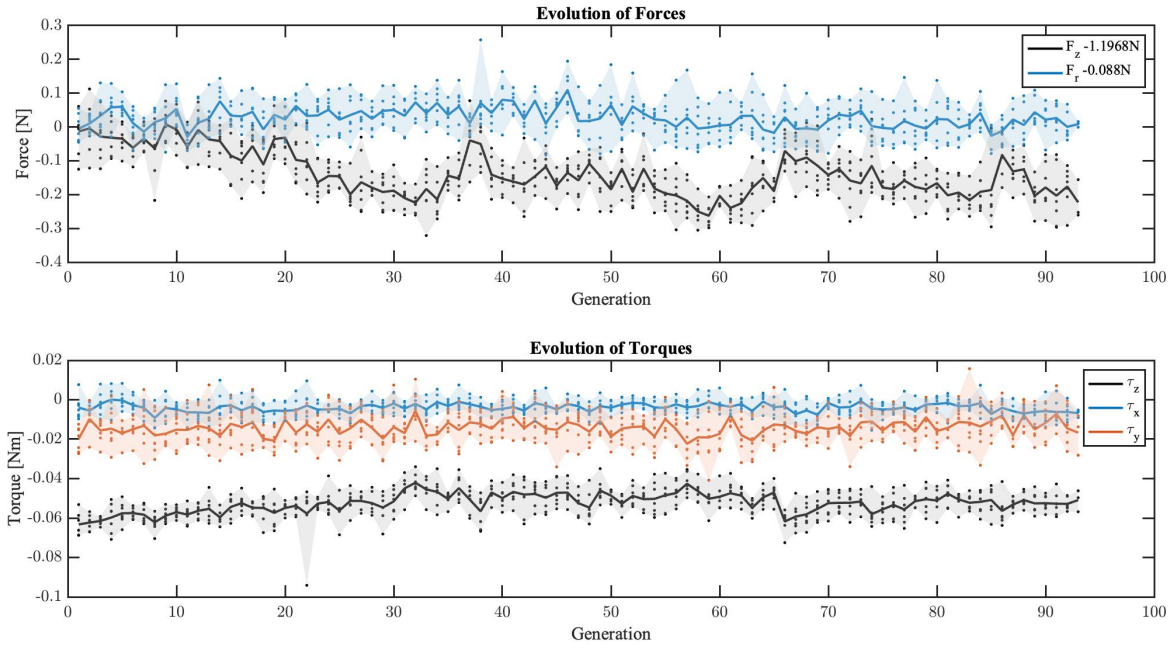

Figure 11.3: Evolution of the forces and torques for the optimization for 5° of thrust vectoring. The solid line tracks the mean of each generation, dots represent the individual values for each candidate solution, and shaded region encompasses all values per generation.

phase shift.

11.2 Maximization of Radial Force and Non-Axial Torques

A second application of thrust vectoring is in short-time, high force and torque maneuvers. These maneuvers can drastically increase maneuverability. In standard cyclic pitch, the direction of thrust is changed by torquing the rotor disk. We thus seek a pitch trajectory that can both maximally angle the direction of the thrust vector and maximize the torque reaction on the rotor disk to enable enhanced maneuverability. In this experiment we do not consider an efficiency metric, since high forces and torques will likely be implemented sparingly.

In this experiment, the fitness function was chosen to be

$$f = \frac{1}{30\sqrt{F_x^2 + F_y^2}} + \tau_x + \tau_y. \tag{11.2}$$

Because τ_x and τ_y are on average negative, this fitness function will maximize the radial force $\sqrt{F_x^2 + F_y^2}$ while also maximizing both τ_x and τ_y. The prefactor on

the force term was chosen to ensure that the magnitude of the force term would be approximately equal to the magnitude of the torque terms if thrust were vectored at 90°.

Six variables participate in the optimization. These are the frequencies of both blades, the phase shift of both blades, and the speed-up value of both blades ($f_1, f_2, \theta_1, \theta_2, K_{v,1}, K_{v,2}$). The maximum, minimum, and zero-crossing AOA are fixed at $\alpha_{min} = 4$, $\alpha_0 = 6.5$, and $\alpha_{max} = 9$. The ranges, tolerances, and initial values of all variables participating in the optimization are reported in table 11.1.

Although the optimization has not yet reached the predetermined convergence criteria, there are some interesting features of the in-progress optimization to investigate. In Figure 11.4, we see that for a period of approximately fifteen generations the optimization stagnates. Neither the mean value nor the variance of the parameters appears to change significantly during this time period. While the CMA-ES algorithm can use stagnation as a measure of convergence, this period occurs too early into the optimization to be considered true stagnation. CMA-ES stagnation convergence cannot occur before $N \times \left(5 + \frac{100}{\lambda}\right)$ generations, which in this case would be generation 97. Instead, the algorithm reacts to this period of stagnation by increasing the variance on all parameters. Increasing the variance allows the optimization to escape local minima.

This long period of stagnation implies that the optimization has found a wide minimum region that is difficult to escape. We therefore analyze some trajectories from the stagnation generation with the best fitness (generation 48, Figure 11.5) to provide insight into the characteristics of this stagnation region. Table 11.2 reports all trajectories from generation 48 with their forces, torques, and evaluated fitness and Figure 11.6 visualizes these trajectories.

The best candidate in this generation is candidate 3, which attains both the maximum F_r and greatest magnitude τ_y. This contributes to its lowest fitness. Candidate 3 produces 1.074N of thrust, leading to a thrust vectoring angle of 11.5°. The worst candidate in this generation is candidate 5, which attains the minimum F_r and the minimum thrust for a thrust vectoring angle of 7°. Its torque magnitude, however, is nearly identical to that of candidate 3. Forces and torques throughout the optimization are reported in Figure 11.7.

When comparing candidate 3 and candidate 5, we see that their trajectories for blade

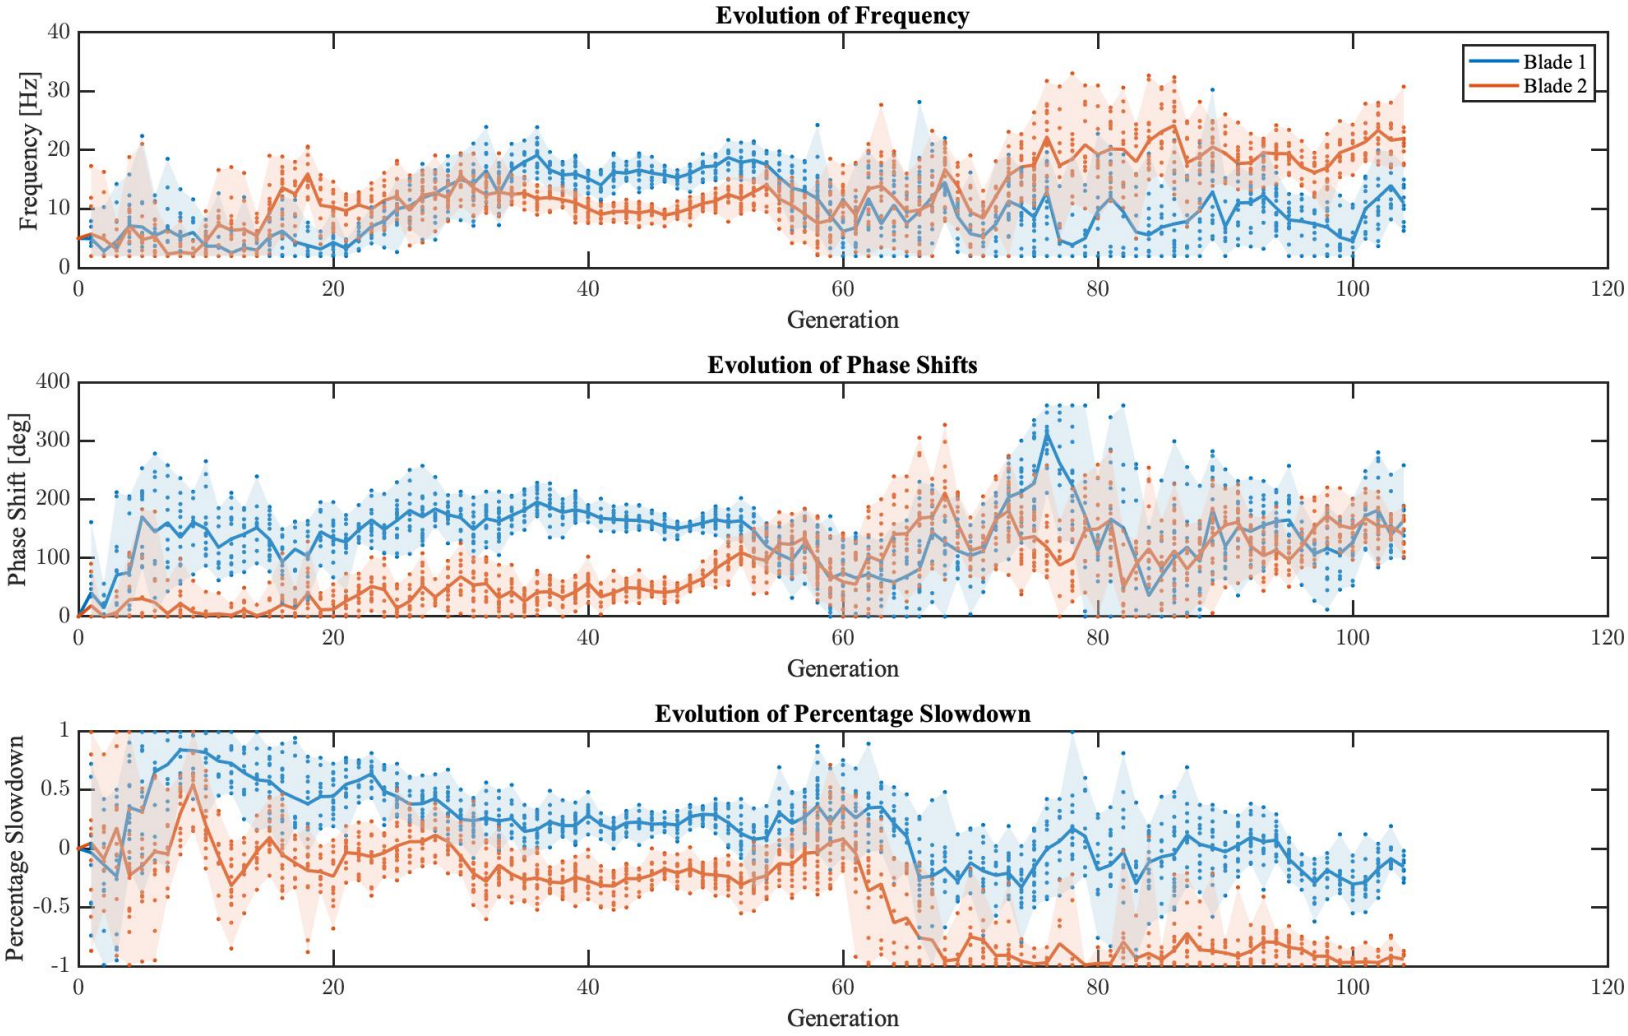

Figure 11.4: Evolution of the parameters participating in the optimization maximizing radial force and non-axial torque. Each dot represents the evaluation of one trajectory, comprising of three experiments. The shading extends from the minimum to maximum values attained per generation. The line tracks the mean per generation.

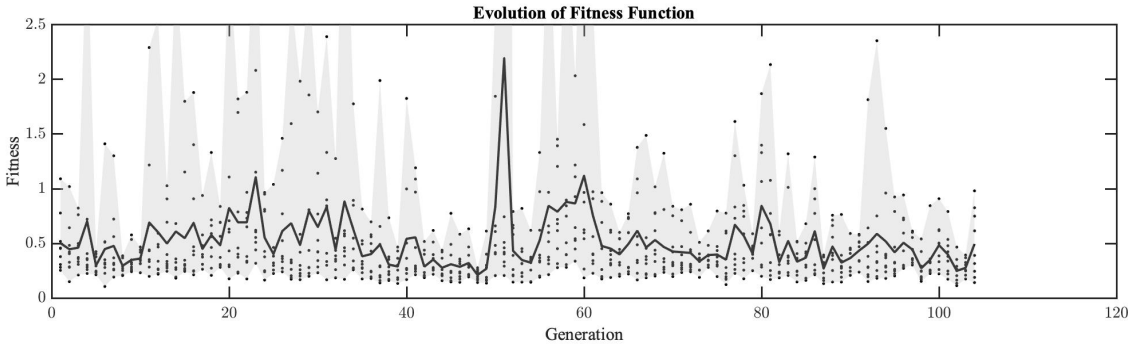

Figure 11.5: Evolution of the fitness function for the optimization maximizing radial force and non-axial torque. Each dot represents the evaluation of one trajectory, comprising of three experiments. The shading extends from the minimum to maximum values attained per generation. The line tracks the mean per generation.

Candidate	f_1 [Hz]	f_2 [Hz]	$K_{v,1}$	$K_{v,2}$	θ_1 [°]	θ_2 [°]	F_r [N]	τ_x [Nm]	τ_y [Nm]	Fitness
1	15.18	9.10	0.30	-0.19	155	63	0.166	-0.009	-0.016	0.1753
2	17.57	11.15	0.37	-0.17	161	70	0.164	0	-0.009	0.1950
3	16.39	11.19	0.29	-0.18	159	66	0.220	0.011	-0.021	0.1424
4	17.20	11.34	0.25	0.01	162	43	0.117	-0.009	-0.016	0.2602
5	16.52	9.97	0.29	-0.10	149	60	0.109	-0.011	-0.020	0.2762
6	15.93	8.53	0.30	-0.15	152	60	0.131	-0.002	-0.005	0.2480
7	16.27	12.03	0.16	-0.20	155	58	0.201	0	-0.003	0.1642
8	16.27	9.91	0.28	-0.13	146	46	0.157	0	-0.019	0.1937
9	12.94	7.21	0.24	-0.22	150	50	0.126	-0.005	-0.022	0.2387
Range	4.63	4.82	0.21	0.23	16	27	0.111	0.022	0.019	0.1338

Table 11.2: Values of the optimizable parameters along with the resultant fitness function for all candidate trajectories in generation 48.

1 are one of the most similar pairs of all candidate solutions[1] (Figure 11.6). There is less than a 0.5° difference between all points, and the frequencies are the same due to servo motor resolution. The trajectory for blade two is where candidate 3 and candidate 5 are significantly different. The frequencies of actuation are again the same due to servo motor resolution, but the phase shifts and speed-up values produce different angle of attack traces.

Candidate 5 has an overall lower average angle of attack for blade 2. This discrepancy only partially explains the resulting forces. It explains the discrepancy in lift produced between the two candidates, where candidate 5 produced 0.885 N while candidate 3 produced 1.074 N. This is because lift increases with increasing angle

[1]Candidate 1 and 8 are the most similar, but their similarity extends to blade 2. Furthermore, their fitnesses and force/torque production are also similar.

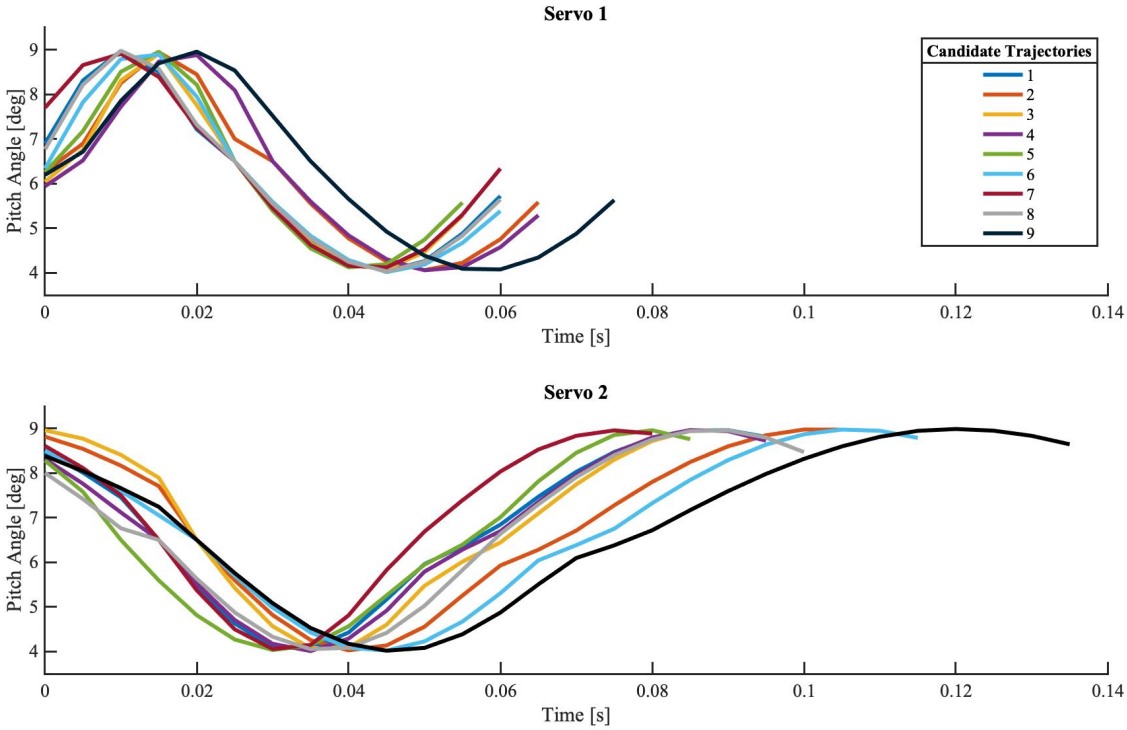

Figure 11.6: Candidate trajectories for evaluation in the 48th generation of the optimization for maximum radial force and non-axial torque.

of attack below the stall angle [79]. However, the lower average AOA cannot explain the difference in vectoring angle: 11.5 degrees for candidate 3 and 7 degrees for candidate 5. The lift-to-drag ratio increases with increasing angle of attack below the stall angle [79]. Because the angle of thrust vectoring is the arctangent of F_r (derived from drag) over F_z (lift), an increasing lift-to-drag ratio with increasing angle of attack should result in a decreasing thrust vectoring angle. Instead, candidate 3 has both a higher angle of attack and a larger thrust vectoring angle. This discrepancy suggests that a quasi-steady analysis based on average angles of attack is incomplete, and unsteady motion plays a role in the aerodynamics.

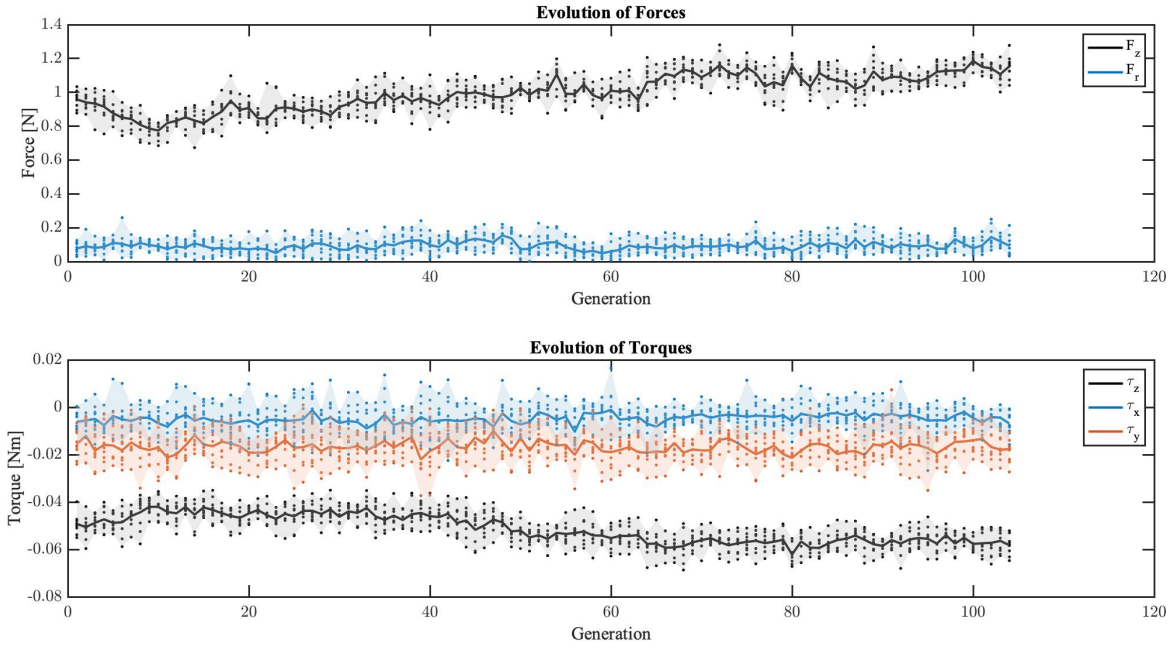

Figure 11.7: Evolution of the forces and torques for the optimization maximizing radial force and non-axial torque. Each dot represents the evaluation of one trajectory, comprising of three experiments. The shading extends from the minimum to maximum values attained per generation. The line tracks the mean per generation.

Chapter 12

CONCLUSION

In Part II of this thesis, we construct an individual blade control rotor and pair it with an evolutionary strategy (CMA-ES) to optimize for any desired measurable behavior. We discuss important experimental considerations like dynamic similarity, error measurements, and noise characterization. The experimental setup produces expected behavior when running an optimization with a known outcome.

We also present the preliminary results of two thrust-vectoring optimizations on this experimental apparatus, both to illuminate possible applications of this system, as well as to discuss some interesting features of these optimizations. These thrust vectoring optimizations are implemented on a two-blade rotor, to aid in experimental analysis. In the first optimization, where we optimize for efficient steep takeoff or shallow approach, we find a local minimum region in the phase shift plane. Within this minimum, a thrust vectoring angle of approximately 5° is achievable with a relative phase difference of approximately 85° between two blades.

The second optimization expands the optimizable variables and maximizes for both radial force and non-axial torque. This is applicable to short-time maneuvers where high maneuverability is desired. We again find a wide minimum region, in which all parameters involved in the optimization stagnate for approximately fifteen generations. When we compare the best and the worst candidate solution in the highest-fitness generation during the stagnation, we find that although one blade's trajectory is highly similar, the second blade shows significant difference in pitch angle. The candidate solution with a lower average pitch angle attains a lower lift, as expected, but displays a smaller angle of thrust vectoring. This contradicts expected behavior; since the lift-to-drag ratio increases with increasing angle of attack, a lower angle of attack should result in a larger angle of thrust vectoring. This suggests that a quasi-steady analysis is not sufficient to explain the produced forces, and that unsteady motion plays an aerodynamic role.

Although this preliminary analysis has revealed some insight into thrust vectoring on two different timescales, the final optimal results of the thrust vectoring optimizations remain to be analyzed. In addition, this experimental setup can be used in the future for a wide variety of optimization purposes. Future investigations can

range from any of the benefits of IBC discussed in Chapter 9 (like improving performance, reducing retreating blade stall, reducing blade-vortex interaction noise, etc.) to future horizons like the thrust vectoring optimizations herein.

One particularly interesting application could be using optimization to determine a pitch strategy in response to ground resonance. Ground resonance is caused by a shift in the center of mass of the system due to asymmetric movement of blades about their lead-lag hinges. This is most often a result of an unbalanced, hard touchdown. Ground resonance can only be stopped by liftoff, breaking the feedback between the helicopter and the ground. But, if a pilot does not notice the vibrations beginning to grow, liftoff quickly becomes impossible. At that point, the helicopter shakes itself to pieces via resonance. It is possible that changing the pitch of the rotor blades individually could combat the resonance by introducing an resonance-opposing time-varying force.

To enable this future experiment, one would introduce asymmetry to the rotor hub in order to shift the center of mass. This could be done in a variety of ways; the lead-lag hinges could be fixed in a ground-resonance configuration, the blades could be asymmetrically damaged, or mass could be added off-axis. A fitness function should be chosen that penalizes vibration, using a fast Fourier transform of the force data to quantify vibrational power. Because resonance phenomena are inherently dangerous to the experimental setup, a safety stop must be implemented. If, during the experiment, vibration is increasing beyond system capacity (given a safety factor), evaluation of that candidate must immediately cease. Cessation should impose an extra penalty on the fitness function to ensure that resonance-enhancing pitch alteration is avoided. With the aforementioned optimization, a varying range of asymmetries can be considered to determine the optimal pitch response to encountering ground resonance.

BIBLIOGRAPHY

[1] Pieter Abbeel et al. "An Application of Reinforcement Learning to Aerobatic Helicopter Flight". In: *Advances in Neural Information Processing Systems*. Ed. by B. Schölkopf, J. Platt, and T. Hoffman. Vol. 19. MIT Press, 2006.

[2] Mark Achache and Michel Polychroniadis. "Development of an Experimental System for Active Control of Vibrations on Helicopters". In: *Proceedings of the 12th European Rotorcraft Forum*. Vol. 64. Garmish-Partenkirchen, Germany, 1986.

[3] Federal Aviation Administration. *Helicopter Flying Handbook: Faa-H-8083-21b*. New York: Skyhorse Publishing Company, Incorporated, 2021. 1 p. ISBN: 978-1-5107-6721-8.

[4] J. M. Anderson et al. "Oscillating foils of high propulsive efficiency". In: *Journal of Fluid Mechanics* 360 (Apr. 10, 1998), pp. 41–72. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/S0022112097008392.

[5] Peter J. Arcidiacono. "Theoretical Performance of Helicopters Having Second and Higher Harmonic Feathering Control". In: *Journal of the American Helicopter Society* 6.2 (Apr. 1, 1961), pp. 8–19. ISSN: 2161-6027. DOI: 10.4050/JAHS.6.2.8.

[6] Bob Balaram et al. "Mars Helicopter Technology Demonstrator". In: *2018 AIAA Atmospheric Flight Mechanics Conference*. 2018 AIAA Atmospheric Flight Mechanics Conference. Kissimmee, Florida: American Institute of Aeronautics and Astronautics, Jan. 8, 2018. ISBN: 978-1-62410-525-8. DOI: 10.2514/6.2018-0023.

[7] Jorge M. M. Barata et al. "Propulsion for Biological Inspired Micro-Air Vehicles (MAVs)". In: *Open Journal of Applied Sciences* 06.1 (2016), pp. 7–15. ISSN: 2165-3917, 2165-3925. DOI: 10.4236/ojapps.2016.61002.

[8] Diego Bernal et al. "Review: Analysis of the evolutionary convergence for high performance swimming in lamnid sharks and tunas". In: *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology* 129.2 (June 2001), pp. 695–726. ISSN: 10956433. DOI: 10.1016/S1095-6433(01)00333-6.

[9] James M. Birch and Michael H. Dickinson. "Spanwise flow and the attachment of the leading-edge vortex on insect wings". In: *Nature* 412.6848 (Aug. 2001), pp. 729–733. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/35089071.

[10] James M. Birch and Michael H. Dickinson. "The influence of wing–wake interactions on the production of aerodynamic forces in flapping flight".

In: *Journal of Experimental Biology* 206.13 (July 1, 2003), pp. 2257–2272. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.00381.

[11] James M. Birch, William B. Dickson, and Michael H. Dickinson. "Force production and flow structure of the leading edge vortex on flapping wings at high and low Reynolds numbers". In: *Journal of Experimental Biology* 207.7 (Mar. 1, 2004), pp. 1063–1072. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.00848.

[12] Boeing. *AH-64 Apache*.

[13] Richard J. Bomphrey et al. "Smart wing rotation and trailing-edge vortices enable high frequency mosquito flight". In: *Nature* 544.7648 (Apr. 2017), pp. 92–95. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature21727.

[14] Richard J. Bomphrey et al. "The aerodynamics of *Manduca sexta* : digital particle image velocimetry analysis of the leading-edge vortex". In: *Journal of Experimental Biology* 208.6 (Mar. 15, 2005), pp. 1079–1094. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.01471.

[15] Charles M. Breder. "The locomotion of fishes". In: *Zoologica : scientific contributions of the New York Zoological Society* 4.5 (Sept. 29, 1926), pp. 159–297. ISSN: 0044-507X. DOI: 10.5962/p.203769.

[16] A. Brocklehurst and G.N. Barakos. "A review of helicopter rotor blade tip shapes". In: *Progress in Aerospace Sciences* 56 (Jan. 2013), pp. 35–74. ISSN: 03760421. DOI: 10.1016/j.paerosci.2012.06.003.

[17] A. K. Brodsky. "Vortex Formation in the Tethered Flight of the Peacock Butterfly *Inachis IO* L. (Lepidoptera, Nymphalidae) and Some Aspects of Insect Flight Evolution". In: *Journal of Experimental Biology* 161.1 (Nov. 1, 1991), pp. 77–95. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.161.1.77.

[18] R. Buchwald and R. Dudley. "Limits to vertical force and power production in bumblebees (Hymenoptera: *Bombus impatiens*)". In: *Journal of Experimental Biology* 213.3 (Feb. 1, 2010), pp. 426–432. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.033563.

[19] L Cai et al. "Effect of tail fin loss on swimming capability and tail beat frequency of juvenile black carp Mylopharyngodon piceus". In: *Aquatic Biology* 29 (Apr. 9, 2020), pp. 71–77. ISSN: 1864-7782, 1864-7790. DOI: 10.3354/ab00727.

[20] Caroline E. Champagne et al. "Effects of Fin Clipping on Survival and Position-Holding Behavior of Brown Darters, Etheostoma edwini". In: *Copeia* 2008.4 (Dec. 18, 2008), pp. 916–919. ISSN: 0045-8511, 1938-5110. DOI: 10.1643/CI-07-153.

[21] Long Chen et al. “Effects of Reynolds number on leading-edge vortex formation dynamics and stability in revolving wings”. In: *Journal of Fluid Mechanics* 931 (Jan. 25, 2022), A13. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/jfm.2021.950.

[22] Zengshuang Chen et al. “Aerodynamics of a flapping wing with stroke deviation in forward flight”. In: *Physics of Fluids* 36.5 (May 1, 2024), p. 051908. ISSN: 1070-6631, 1089-7666. DOI: 10.1063/5.0209169.

[23] Bo Cheng, Xinyan Deng, and Tyson L. Hedrick. “The mechanics and control of pitching manoeuvres in a freely flying hawkmoth (*Manduca sexta*)”. In: *Journal of Experimental Biology* 214.24 (Dec. 15, 2011), pp. 4092–4106. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.062760.

[24] Diana D. Chin and David Lentink. “Flapping wing aerodynamics: from insects to vertebrates”. In: *Journal of Experimental Biology* 219.7 (Apr. 1, 2016), pp. 920–932. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.042317.

[25] Brian D. Clark and Willy Bemis. “Kinematics of swimming of penguins at the Detroit Zoo”. In: *Journal of Zoology* 188.3 (July 1979), pp. 411–428. ISSN: 0952-8369, 1469-7998. DOI: 10.1111/j.1469-7998.1979.tb03424.x.

[26] US Department of Commerce. *NIST Chemistry WebBook, Isothermal Properties for Nitrogen*.

[27] Konstantinos Dalamagkidis, Kimon P. Valavanis, and Les A. Piegl. “Nonlinear Model Predictive Control With Neural Network Optimization for Autonomous Autorotation of Small Unmanned Helicopters”. In: *IEEE Transactions on Control Systems Technology* 19.4 (July 2011), pp. 818–831. ISSN: 1063-6536, 1558-0865. DOI: 10.1109/TCST.2010.2054092.

[28] Tivadar Danka. *modAL: A modular active learning framework for Python3*. 2018.

[29] Michael H. Dickinson. “The Effects of Wing Rotation on Unsteady Aerodynamic Performance At Low Reynolds Numbers”. In: *Journal of Experimental Biology* 192.1 (July 1, 1994), pp. 179–206. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.192.1.179.

[30] Michael H. Dickinson, Fritz-Olaf Lehmann, and Karl G. Götz. “The Active Control of Wing Rotation by *Drosophila*”. In: *Journal of Experimental Biology* 182.1 (Sept. 1, 1993), pp. 173–189. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.182.1.173.

[31] Michael H. Dickinson, Fritz-Olaf Lehmann, and Sanjay P. Sane. “Wing Rotation and the Aerodynamic Basis of Insect Flight”. In: *Science* 284.5422 (June 18, 1999), pp. 1954–1960. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.284.5422.1954.

[32] William B. Dickson, Andrew D. Straw, and Michael H. Dickinson. "Integrative Model of Drosophila Flight". In: *AIAA Journal* 46.9 (Sept. 2008), pp. 2150–2164. ISSN: 0001-1452, 1533-385X. DOI: 10.2514/1.29862.

[33] Eliot G. Drucker and George V. Lauder. "Locomotor function of the dorsal fin in rainbow trout: kinematic patterns and hydrodynamic forces". In: *Journal of Experimental Biology* 208.23 (Dec. 1, 2005), pp. 4479–4494. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.01922.

[34] Eliot G. Drucker, Jeffrey A. Walker, and Mark W. Westneat. "Mechanics of Pectoral Fin Swimming in Fishes". In: *Fish Physiology*. Vol. 23. Elsevier, 2005, pp. 369–423. ISBN: 978-0-12-350447-0. DOI: 10.1016/S1546-5098(05)23010-8.

[35] C. P. Ellington. "The aerodynamics of hovering insect flight. I. The quasi-steady analysis". In: *Philosophical Transactions of the Royal Society of London. B, Biological Sciences* 305.1122 (Feb. 24, 1984), pp. 1–15. ISSN: 0080-4622, 2054-0280. DOI: 10.1098/rstb.1984.0049.

[36] C. P. Ellington. "The aerodynamics of hovering insect flight. III. Kinematics". In: *Philosophical Transactions of the Royal Society of London. B, Biological Sciences* 305.1122 (Feb. 24, 1984), pp. 41–78. ISSN: 0080-4622, 2054-0280. DOI: 10.1098/rstb.1984.0051.

[37] María José Fernández, Dwight Springthorpe, and Tyson L. Hedrick. "Neuromuscular and biomechanical compensation for wing asymmetry in insect hovering flight". In: *Journal of Experimental Biology* (Jan. 1, 2012), jeb.073627. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.073627.

[38] F. E. Fish and L. D. Shannahan. "The role of the pectoral fins in body trim of sharks". In: *Journal of Fish Biology* 56.5 (May 2000), pp. 1062–1073. ISSN: 0022-1112, 1095-8649. DOI: 10.1111/j.1095-8649.2000.tb02123.x.

[39] Frank E. Fish. "Advantages of Natural Propulsive Systems". In: *Marine Technology Society Journal* 47.5 (Sept. 1, 2013), pp. 37–44. ISSN: 0025-3324. DOI: 10.4031/MTSJ.47.5.2.

[40] Frank E. Fish. "Transitions from Drag-based to Lift-based Propulsion in Mammalian Swimming". In: *American Zoologist* 36.6 (Dec. 1996), pp. 628–641. ISSN: 0003-1569. DOI: 10.1093/icb/36.6.628.

[41] Peretz P. Friedmann and Thomas A. Millott. "Vibration reduction in rotorcraft using active control - A comparison of various approaches". In: *Journal of Guidance, Control, and Dynamics* 18.4 (July 1995), pp. 664–673. ISSN: 0731-5090, 1533-3884. DOI: 10.2514/3.21445.

[42] H. Russ Frith and Robert W. Blake. "The Mechanical Power Output and Hydromechanical Efficiency of Northern Pike (*Esox Lucius*) Fast-Starts". In: *Journal of Experimental Biology* 198.9 (Sept. 1, 1995), pp. 1863–1873. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.198.9.1863.

[43] Cheng Fu, Zhen-Dong Cao, and Shi-Jian Fu. "The effects of caudal fin loss and regeneration on the swimming performance of three cyprinid fish with different swimming capacities". In: *Journal of Experimental Biology* (Jan. 1, 2013), jeb.084244. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.084244.

[44] Matt K. Fu, Isabel A. Houghton, and John O. Dabiri. "A single-camera, 3D scanning velocimetry system for quantifying active particle aggregations". In: *Experiments in Fluids* 62.8 (Aug. 2021), p. 168. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-021-03256-x.

[45] Yadong Gao et al. "Bionic Intelligent Algorithms Used in Helicopter Individual Blade Control Optimization". In: *Applied Sciences* 12.9 (Apr. 27, 2022), p. 4392. ISSN: 2076-3417. DOI: 10.3390/app12094392.

[46] Alexander Gehrke and Karen Mulleners. "Phenomenology and scaling of optimal flapping wing kinematics". In: *Bioinspiration & Biomimetics* 16.2 (Mar. 1, 2021), p. 026016. ISSN: 1748-3182, 1748-3190. DOI: 10.1088/1748-3190/abd012.

[47] Sven Gemballa et al. "Evolution of high-performance swimming in sharks: Transformations of the musculotendinous system from subcarangiform to thunniform swimmers". In: *Journal of Morphology* 267.4 (Apr. 2006), pp. 477–493. ISSN: 0362-2525, 1097-4687. DOI: 10.1002/jmor.10412.

[48] K.C. Giannakoglou, D.I. Papadimitriou, and I.C. Kampolis. "Aerodynamic shape design using evolutionary algorithms and new gradient-assisted metamodels". In: *Computer Methods in Applied Mechanics and Engineering* 195.44 (Sept. 2006), pp. 6312–6329. ISSN: 00457825. DOI: 10.1016/j.cma.2005.12.008.

[49] Gary B. Gillis. "Undulatory Locomotion in Elongate Aquatic Vertebrates: Anguilliform Swimming since Sir James Gray". In: *American Zoologist* 36.6 (Dec. 1996), pp. 656–665. ISSN: 0003-1569. DOI: 10.1093/icb/36.6.656.

[50] Holger Glaessel, Valentin Kloeppel, and Stephan Rudolph. "Neural control of helicopter blade-vortex interaction noise". In: SPIE's 8th Annual International Symposium on Smart Structures and Materials. Ed. by Anna-Maria R. McGowan. Newport Beach, CA, June 14, 2001, pp. 460–470. DOI: 10.1117/12.429687.

[51] Jeffrey B. Graham et al. "Surface and Subsurface Swimming of the Sea Snake *Pelamis Platurus*". In: *Journal of Experimental Biology* 127.1 (Jan. 1, 1987), pp. 27–44. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.127.1.27.

[52] J. Gray. "Studies in Animal Locomotion". In: *Journal of Experimental Biology* 10.1 (Jan. 1, 1933), pp. 88–104. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.10.1.88.

[53] James Gray. *Animal Locomotion*. London, U.K.: Weidenfeld and Nicolson, 1968.

[54] Norman. D. Ham and Robert M. McKillip Jr. "A simple system for helicopter individual-blade-control and its application to gust alleviation". In: Sixth European Rotorcraft and Powered Lift Aircraft Forum. Bristol, England, Sept. 16, 1980.

[55] C. E. Hammond. "Wind Tunnel Results Showing Rotor Vibratory Loads Reduction Using Higher Harmonic Blade Pitch". In: *Journal of the American Helicopter Society* 28.1 (Jan. 1, 1983), pp. 10–15. ISSN: 2161-6027. DOI: 10.4050/JAHS.28.10.

[56] Nikolaus Hansen. *CMA-ES Source Code*. Version 3.61.beta. Apr. 2012.

[57] Nikolaus Hansen. "The CMA Evolution Strategy: A Tutorial". In: (2016). Publisher: arXiv Version Number: 2. DOI: 10.48550/ARXIV.1604.00772.

[58] Nikolaus Hansen, Sibylle D. Müller, and Petros Koumoutsakos. "Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES)". In: *Evolutionary Computation* 11.1 (Mar. 2003), pp. 1–18. ISSN: 1063-6560, 1530-9304. DOI: 10.1162/106365603321828970.

[59] Nikolaus Hansen and Andreas Ostermeier. "Completely Derandomized Self-Adaptation in Evolution Strategies". In: *Evolutionary Computation* 9.2 (June 2001), pp. 159–195. ISSN: 1063-6560, 1530-9304. DOI: 10.1162/106365601750190398.

[60] Gregory F. Heath and Robert B. Bossler Jr. *Advanced Rotorcraft Transmission (ART) program*. Publication Title: Final Report McDonnell-Douglas Helicopter Co ADS Bibcode: 1993mcdd.rept.....H. Jan. 1, 1993.

[61] A. Hedenström, C. P. Ellington, and T. J. Wolf. "Wing wear, aerodynamics and flight energetics in bumblebees (*Bombus terrestris*): an experimental study". In: *Functional Ecology* 15.4 (Aug. 2001), pp. 417–422. ISSN: 0269-8463, 1365-2435. DOI: 10.1046/j.0269-8463.2001.00531.x.

[62] William Stewart Hoar and David John Randall. *Fish physiology*. New York San Francisco London: Academic press, 1978. ISBN: 978-0-12-350407-4.

[63] M. L. Hooper, I. Scherl, and M. Gharib. "Bio-inspired compensatory strategies for damage to flapping robotic propulsors". In: *Journal of The Royal Society Interface* 21.216 (July 2024), p. 20240141. ISSN: 1742-5662. DOI: 10.1098/rsif.2024.0141.

[64] J. R. Hove et al. "Boxfishes (Teleostei: Ostraciidae) as a model system for fishes swimming with many fins: kinematics". In: *Journal of Experimental Biology* 204.8 (Apr. 15, 2001), pp. 1459–1471. ISSN: 0022-0949. DOI: 10.1242/jeb.204.8.1459.

[65] G. K. Hunt. “Similarity requirements for aeroelastic models of helicopter rotors”. In: *Aeronautical Research Council Current Papers* 1245 (Jan. 1972).

[66] Stephen A Jacklin et al. “Investigation of a Helicopter Individual Blade Control (IBC) System in Two Full-Scale Wind Tunnel Tests: Volume I”. In: (2020).

[67] Benjamin Jantzen and Thomas Eisner. “Hindwings are unnecessary for flight but essential for execution of normal evasive flight in Lepidoptera”. In: *Proceedings of the National Academy of Sciences* 105.43 (Oct. 28, 2008), pp. 16636–16640. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.0807223105.

[68] Bruce C. Jayne. “Swimming in Constricting (Elaphe g. guttata) and Non-constricting (Nerodia fasciata pictiventris) Colubrid Snakes”. In: *Copeia* 1985.1 (Feb. 11, 1985), p. 195. ISSN: 00458511. DOI: 10.2307/1444809.

[69] L. C. Johansson and P. Henningsson. “Butterflies fly using efficient propulsive clap mechanism owing to flexible wings”. In: *Journal of The Royal Society Interface* 18.174 (Jan. 2021), p. 20200854. ISSN: 1742-5689, 1742-5662. DOI: 10.1098/rsif.2020.0854.

[70] Wayne Johnson. *Rotorcraft aeromechanics*. Cambridge aerospace series 36. Cambridge: Cambridge university press, 2013. ISBN: 978-1-107-02807-4.

[71] Chang-kwon Kang and Wei Shyy. “Analytical model for instantaneous lift and shape deformation of an insect-scale flapping wing in hover”. In: *Journal of The Royal Society Interface* 11.101 (Dec. 6, 2014), p. 20140933. ISSN: 1742-5689, 1742-5662. DOI: 10.1098/rsif.2014.0933.

[72] Stefan Kern and Petros Koumoutsakos. “Simulations of optimized anguilliform swimming”. In: *Journal of Experimental Biology* 209.24 (Dec. 15, 2006), pp. 4841–4857. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.02526.

[73] Ch. Kessler. “Active rotor control for helicopters: individual blade control and swashplateless rotor designs”. In: *CEAS Aeronautical Journal* 1.1 (Sept. 2011), pp. 23–54. ISSN: 1869-5582, 1869-5590. DOI: 10.1007/s13272-011-0001-0.

[74] Ch. Kessler. “Active rotor control for helicopters: motivation and survey on higher harmonic control”. In: *CEAS Aeronautical Journal* 1.1 (Sept. 2011), pp. 3–22. ISSN: 1869-5582, 1869-5590. DOI: 10.1007/s13272-011-0005-9.

[75] Daegyoum Kim and Morteza Gharib. “Characteristics of vortex formation and thrust performance in drag-based paddling propulsion”. In: *Journal of Experimental Biology* 214.13 (July 1, 2011), pp. 2283–2291. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.050716.

[76] Joel G. Kingsolver. "EXPERIMENTAL ANALYSES OF WING SIZE, FLIGHT, AND SURVIVAL IN THE WESTERN WHITE BUTTERFLY". In: *Evolution* 53.5 (Oct. 1999), pp. 1479–1490. ISSN: 0014-3820, 1558-5646. DOI: 10.1111/j.1558-5646.1999.tb05412.x.

[77] Max Kramer. *Increase in the maximum lift of an airplane wing due to a sudden increase in its effective angle of attack resulting from a gust*. 1932.

[78] Marcel Kretz. "Research in multicyclic and active control of rotary wings". In: 1.2 (1976), pp. 95–105.

[79] C.L. Ladson et al. *Effects of Independent Variation of Mach and Reynolds Numbers on the Low-speed Aerodynamic Characteristics of the NACA 0012 Airfoil Section*. NASA technical memorandum. National Aeronautics, Space Administration, Scientific, and Technical Information Division, 1988.

[80] D.Jin Lee, Hyochoong Bang, and Kwangyul Baek. "Autorotation of an Unmanned Helicopter by a Reinforcement Learning Algorithm". In: *AIAA Guidance, Navigation and Control Conference and Exhibit*. AIAA Guidance, Navigation and Control Conference and Exhibit. Honolulu, Hawaii: American Institute of Aeronautics and Astronautics, Aug. 18, 2008. ISBN: 978-1-60086-999-0. DOI: 10.2514/6.2008-7279.

[81] Y J Lee and K B Lua. "Wing–wake interaction: comparison of 2D and 3D flapping wings in hover flight". In: *Bioinspiration & Biomimetics* 13.6 (Sept. 14, 2018), p. 066003. ISSN: 1748-3190. DOI: 10.1088/1748-3190/aadc31.

[82] Y J Lee et al. "A quasi-steady aerodynamic model for flapping flight with improved adaptability". In: *Bioinspiration & Biomimetics* 11.3 (Apr. 28, 2016), p. 036005. ISSN: 1748-3190. DOI: 10.1088/1748-3190/11/3/036005.

[83] Fritz-Olaf Lehmann. "The mechanisms of lift enhancement in insect flight". In: *Naturwissenschaften* 91.3 (Mar. 1, 2004), pp. 101–122. ISSN: 0028-1042, 1432-1904. DOI: 10.1007/s00114-004-0502-3.

[84] Fritz-Olaf Lehmann and Michael H. Dickinson. "The Changes in Power Requirements and Muscle Efficiency During Elevated Force Production in the Fruit Fly *Drosophila Melanogaster*". In: *Journal of Experimental Biology* 200.7 (Apr. 1, 1997), pp. 1133–1143. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.200.7.1133.

[85] J. Gordon Leishman. *A History of Helicopter Flight*. 2000.

[86] David Lentink and Michael H. Dickinson. "Rotational accelerations stabilize leading edge vortices on revolving fly wings". In: *Journal of Experimental Biology* 212.16 (Aug. 15, 2009), pp. 2705–2719. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.022269.

[87] M J Lighthill. “Hydromechanics of Aquatic Animal Propulsion”. In: *Annual Review of Fluid Mechanics* 1.1 (Jan. 1969), pp. 413–446. ISSN: 0066-4189, 1545-4479. DOI: 10.1146/annurev.fl.01.010169.002213.

[88] Hao Liu et al. “Biomechanics and biomimetics in insect-inspired flight systems”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 371.1704 (Sept. 26, 2016), p. 20150390. ISSN: 0962-8436, 1471-2970. DOI: 10.1098/rstb.2015.0390.

[89] Jindong Liu and Huosheng Hu. “Biological inspiration: From carangiform fish to multi-joint robotic fish”. In: *Journal of Bionic Engineering* 7.1 (Mar. 2010), pp. 35–48. ISSN: 1672-6529, 2543-2141. DOI: 10.1016/S1672-6529(09)60184-0.

[90] C. N. H. Lock. “Further development of autogyro theory”. In: *Reports and Memoranda* 1127 (Mar. 1927).

[91] Massimo Lorenzoni et al. “The growth of pike (Esox lucius Linnaeus, 1798) in Lake Trasimeno (Umbria, Italy)”. In: *Fisheries Research* 59.1 (Dec. 2002), pp. 239–246. ISSN: 01657836. DOI: 10.1016/S0165-7836(02)00013-9.

[92] K. B. Lua et al. “Effects of pitching phase angle and amplitude on a two-dimensional flapping wing in hovering mode”. In: *Experiments in Fluids* 56.2 (Feb. 2015), p. 35. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-015-1907-9.

[93] K. B. Lua et al. “On the aerodynamic characteristics of hovering rigid and flexible hawkmoth-like wings”. In: *Experiments in Fluids* 49.6 (Dec. 2010), pp. 1263–1291. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-010-0873-5.

[94] John J. Magnuson. “COMPARATIVE STUDY OF ADAPTATIONS FOR CONTINUOUS SWIMMING AND HYDROSTATIC EQUILIBRIUM OF SCOMBROID AND XIPHOID FISHES”. In: *Fishery Bulletin* 71.2 (1973), pp. 337–356.

[95] James H. Marden. “Maximum Lift Production During Takeoff in Flying Animals”. In: *Journal of Experimental Biology* 130.1 (July 1, 1987), pp. 235–258. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.130.1.235.

[96] Therese A. Markow. “A comparative investigation of the mating system of Drosophila hydei”. In: *Animal Behaviour* 33.3 (Aug. 1985), pp. 775–781. ISSN: 00033472. DOI: 10.1016/S0003-3472(85)80010-5.

[97] Nathan Martin and Morteza Gharib. “Experimental trajectory optimization of a flapping fin propulsor using an evolutionary strategy”. In: *Bioinspiration & Biomimetics* 14.1 (Nov. 29, 2018), p. 016010. ISSN: 1748-3190. DOI: 10.1088/1748-3190/aaefa5.

[98] Nathan Martin et al. “To flap or not to flap: comparison between flapping and clapping propulsions”. In: *Journal of Fluid Mechanics* 822 (July 10, 2017), R5. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/jfm.2017.252.

[99] T. Maxworthy. “Experiments on the Weis-Fogh mechanism of lift generation by insects in hovering flight. Part 1. Dynamics of the ‘fling’”. In: *Journal of Fluid Mechanics* 93.1 (July 12, 1979), pp. 47–63. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/S0022112079001774.

[100] W Miao, Sesi B. R. Kottapalli, and H. M. Frye. “Flight demonstration of higher harmonic control (HHC) on S-76”. In: *Proceedings of the 42nd Annual Forum of the American Helicopter Society*. Washington, D.C., June 1986.

[101] Michele Milano and Morteza Gharib. “Uncovering the physics of flapping flat plates with artificial evolution”. In: *Journal of Fluid Mechanics* 534 (June 21, 2005), pp. 403–409. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/S0022112005004842.

[102] Michele Milano and Petros Koumoutsakos. “A Clustering Genetic Algorithm for Cylinder Drag Optimization”. In: *Journal of Computational Physics* 175.1 (Jan. 2002), pp. 79–107. ISSN: 00219991. DOI: 10.1006/jcph.2001.6882.

[103] Laura A. Miller and Charles S. Peskin. “Flexible clap and fling in tiny insect flight”. In: *Journal of Experimental Biology* 212.19 (Oct. 1, 2009), pp. 3076–3090. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.028662.

[104] John A. Molusis, C. E. Hammond, and John H. Cline. “A Unified Approach to the Optimal Design of Adaptive and Gain Scheduled Controllers to Achieve Minimum Helicopter Rotor Vibration”. In: *Journal of the American Helicopter Society* 28.2 (Apr. 1, 1983), pp. 9–18. ISSN: 2161-6027. DOI: 10.4050/JAHS.28.9.

[105] Florian T. Muijres et al. “Flies compensate for unilateral wing damage through modular adjustments of wing and body kinematics”. In: *Interface Focus* 7.1 (Feb. 6, 2017), p. 20160103. ISSN: 2042-8898, 2042-8901. DOI: 10.1098/rsfs.2016.0103.

[106] Ulrike K. Müller and Johan L. Van Leeuwen. “Undulatory fish swimming: from muscles to flow”. In: *Fish and Fisheries* 7.2 (June 2006), pp. 84–103. ISSN: 1467-2960, 1467-2979. DOI: 10.1111/j.1467-2979.2006.00210.x.

[107] Senthil Murugan, Ranjan Ganguli, and Dineshkumar Harursampath. “Aeroelastic Response of Composite Helicopter Rotor with Random Material Properties”. In: *Journal of Aircraft* 45.1 (Jan. 2008), pp. 306–322. ISSN: 0021-8669, 1533-3868. DOI: 10.2514/1.30180.

[108] Andrew Y. Ng et al. “Autonomous Inverted Helicopter Flight via Reinforcement Learning”. In: *Experimental Robotics IX*. Ed. by Marcelo H. Ang and Oussama Khatib. Red. by Bruno Siciliano, Oussama Khatib, and Frans Groen. Vol. 21. Series Title: Springer Tracts in Advanced Robotics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 363–372. ISBN: 978-3-540-28816-9 978-3-540-33014-1. DOI: 10.1007/11552246_35.

[109] Khanh Q. Nguyen. “Active Suppression of Stall on Helicopter Rotors”. In: 25th European Rotorcraft Forum. Rome, Italy, Sept. 14, 1999.

[110] Khanh Q. Nguyen. *Higher harmonic control analysis for vibration reduction of helicopter rotor systems*. Oct. 1994.

[111] James J. O’Leary, Sesi B. R. Kottapalli, and Mark Davis. “Adaptation of a Modern Medium Helicopter (Sikorsky S76) to Higher Harmonic Control”. In: *Proceedings of the 2nd Decennial Specialists’ Meeting on Rotorcraft Dynamics*. Moffett Field, CA, 1985.

[112] Andreas Ostermeier, Andreas Gawelczyk, and Nikolaus Hansen. “A Derandomized Approach to Self-Adaptation of Evolution Strategies”. In: *Evolutionary Computation* 2.4 (Dec. 1994), pp. 369–380. ISSN: 1063-6560, 1530-9304. DOI: 10.1162/evco.1994.2.4.369.

[113] Linda Partridge and Marion Farquhar. “Lifetime mating success of male fruitflies (Drosophila melanogaster) is related to their size”. In: *Animal Behaviour* 31.3 (Aug. 1983), pp. 871–877. ISSN: 00033472. DOI: 10.1016/S0003-3472(83)80242-5.

[114] P.R. Payne. “Higher Harmonic Rotor Control: The Possibilities of Third and Higher Harmonic Feathering for Delaying the Stall Limit in Helicopters”. In: *Aircraft Engineering and Aerospace Technology* 30.8 (Aug. 1, 1958), pp. 222–226. ISSN: 0002-2667. DOI: 10.1108/eb033000.

[115] Diing-wen Peng and Michele Milano. “Lift generation with optimal elastic pitching for a flapping plate”. In: *Journal of Fluid Mechanics* 717 (Feb. 25, 2013), R1. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/jfm.2012.614.

[116] Romain Poletti et al. “On the unsteady aerodynamics of flapping wings under dynamic hovering kinematics”. In: *Physics of Fluids* 36.8 (Aug. 1, 2024), p. 081901. ISSN: 1070-6631, 1089-7666. DOI: 10.1063/5.0215531.

[117] Michel Polychroniadis and Mark Achache. “Higher Harmonic Control: Flight Tests of an Experimental System on SA 349 Research Gazelle”. In: *Proceedings of the 42nd Annual Forum of the American Helicopter Society*. Washington, D.C., June 1986.

[118] Todd R. Quackenbush. “Testing and Evaluation of a Stall Flutter Suppression System for Helicopter Rotors Using Individual Blade Control”. In: 9th European Rotorcraft Forum. Stresa, Italy, Sept. 13, 1983.

[119] Hamed Rajabi, Jan-Henning Dirks, and Stanislav N. Gorb. "Insect wing damage: causes, consequences and compensatory mechanisms". In: *Journal of Experimental Biology* 223.9 (May 1, 2020), jeb215194. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.215194.

[120] Manikandan Ramasamy, Timothy E. Lee, and J. Gordon Leishman. "Flowfield of a Rotating-Wing Micro Air Vehicle". In: *Journal of Aircraft* 44.4 (July 2007), pp. 1236–1244. ISSN: 0021-8669, 1533-3868. DOI: 10.2514/1.26415.

[121] Peter Richter and Achim Blaas. "Full-Scale Wind Tunnel Investigation of an Individual Blade Control System for the BO-105 Hingeless Rotor". In: 19th European Rotorcraft Forum. Cernobbio, Italy, Sept. 14, 1993.

[122] Peter Richter, Hans-Dieter Eisbrecher, and Valentin Kloeppel. "Design and First Flight Test of Individual Blade Control Actuators". In: 16th European Rotorcraft Forum. Glasgow, Scotland, Sept. 1990.

[123] Peter Richter and T. Schreiber. "Theoretical Investigations and Windtunnel Tests with HHC-IBC". In: 20th European Rotorcraft Forum. Amsterdam, Netherlands, Oct. 4, 1994.

[124] Jordan C. Roberts and Ralph V. Cartar. "Shape of wing wear fails to affect load lifting in common eastern bumble bees (*Bombus impatiens*) with experimental wing wear". In: *Canadian Journal of Zoology* 93.7 (July 2015), pp. 531–537. ISSN: 0008-4301, 1480-3283. DOI: 10.1139/cjz-2014-0317.

[125] R Rudolph. "Some aspects of wing kinematics in Calopteryx splendens (Harris) (Zygoptera: Calopterigydae)". In: *Odonatologica* 5.2 (1976), pp. 119–127.

[126] Sanjay P. Sane. "The aerodynamics of insect flight". In: *Journal of Experimental Biology* 206.23 (Dec. 1, 2003), pp. 4191–4208. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.00663.

[127] Sanjay P. Sane and Michael H. Dickinson. "The control of flight force by a flapping wing: lift and drag production". In: *Journal of Experimental Biology* 204.15 (Aug. 1, 2001), pp. 2607–2626. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.204.15.2607.

[128] Michael Schmidt and Hod Lipson. "Distilling Free-Form Natural Laws from Experimental Data". In: *Science* 324.5923 (Apr. 3, 2009), pp. 81–85. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1165893.

[129] Cory Seidel and David A. Peters. "How Big Is a Lock Number?" In: *Journal of the American Helicopter Society* 66.1 (Jan. 1, 2021), pp. 1–3. ISSN: 2161-6027. DOI: 10.4050/JAHS.66.012001.

[130] M. Sfakiotakis, D.M. Lane, and J.B.C. Davies. "Review of fish swimming modes for aquatic locomotion". In: *IEEE Journal of Oceanic Engineering* 24.2 (Apr. 1999), pp. 237–252. ISSN: 03649059. DOI: 10.1109/48.757275.

[131] John Shaw et al. "Higher Harmonic Control: Wind Tunnel Demonstration of Fully Effective Vibratory Hub Force Suppression". In: *Journal of the American Helicopter Society* 34.1 (Jan. 1, 1989), pp. 14–25. ISSN: 2161-6027. DOI: 10.4050/JAHS.34.14.

[132] E. L. E. Sinclair, A. J. W. Ward, and F. Seebacher. "Aggression-induced fin damage modulates trade-offs in burst and endurance swimming performance of mosquitofish". In: *Journal of Zoology* 283.4 (Apr. 2011), pp. 243–248. ISSN: 0952-8369, 1469-7998. DOI: 10.1111/j.1469-7998.2010.00776.x.

[133] Jeffrey D. Singleton and William T. Yeager. "Important Scaling Parameters for Testing Model-Scale Helicopter Rotors". In: *Journal of Aircraft* 37.3 (May 2000), pp. 396–402. ISSN: 0021-8669, 1533-3868. DOI: 10.2514/2.2639.

[134] G. J. Sissingh and R. E. Donham. "Hingeless rotor theory and experiment on vibration reduction by periodic variation of conventional controls". In: AHS/NASA-Ames Specialists' Meeting on Rotorcraft Dynamics. Feb. 13, 1974.

[135] S.N. Skinner and H. Zare-Behtash. "State-of-the-art in aerodynamic shape optimisation methods". In: *Applied Soft Computing* 62 (Jan. 2018), pp. 933–962. ISSN: 15684946. DOI: 10.1016/j.asoc.2017.09.030.

[136] Chris Somps and Marvin Luttges. "Dragonfly Flight: Novel Uses of Unsteady Separated Flows". In: *Science* 228.4705 (June 14, 1985), pp. 1326–1329. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.228.4705.1326.

[137] R. B. Srygley and A. L. R. Thomas. "Unconventional lift-generating mechanisms in free-flying butterflies". In: *Nature* 420.6916 (Dec. 2002), pp. 660–664. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature01223.

[138] F. Steiner. "Commande d'un robot Rotule À 3 Degrés Liberté À Très haute Dynamique (hardware)". PhD thesis. Lausanne, Switzerland: MSE HES-SO, 2016.

[139] W Steiner. "Commande d'un robot Rotule À 3 Degrés de Liberté À Très haute Dynamique (Modélisation et Contrôle)". PhD thesis. Lausanne, Switzerland: MSE HES-SO, 2016.

[140] Phillip C. Sternes and Kenshu Shimada. "Body forms in sharks (Chondrichthyes: Elasmobranchii) and their functional, ecological, and evolutionary implications". In: *Zoology* 140 (June 2020), p. 125799. ISSN: 09442006. DOI: 10.1016/j.zool.2020.125799.

[141] F. K. Straub and E. V. Byrns Jr. *Application of higher harmonic blade feathering on the OH-6A helicopter for vibration reduction*. Dec. 1986.

[142] Mao Sun and Jian Tang. "Lift and power requirements of hovering flight in *Drosophila virilis*". In: *Journal of Experimental Biology* 205.16 (Aug. 15, 2002), pp. 2413–2427. ISSN: 1477-9145, 0022-0949. DOI: 10.1242/jeb.205.16.2413.

[143] A.P Thomas et al. "Synthetic jet propulsion for small underwater vehicles". In: *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. Apr. 18, 2005, pp. 181–187.

[144] Libéria Souza Torquato et al. "Cellular basis of morphological variation and temperature-related plasticity in Drosophila melanogaster strains with divergent wing shapes". In: *Genetica* 142.6 (Dec. 2014), pp. 495–505. ISSN: 0016-6707, 1573-6857. DOI: 10.1007/s10709-014-9795-0.

[145] Pat Trizila et al. "Low-Reynolds-Number Aerodynamics of a Flapping Rigid Flat Plate". In: *AIAA Journal* 49.4 (Apr. 2011), pp. 806–823. ISSN: 0001-1452, 1533-385X. DOI: 10.2514/1.J050827.

[146] Eric D. Tytell. "Do trout swim better than eels? Challenges for estimating performance based on the wake of self-propelled bodies". In: *Experiments in Fluids* 43.5 (Nov. 6, 2007), pp. 701–712. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-007-0343-x.

[147] Eric D. Tytell. "Kinematics and hydrodynamics of linear acceleration in eels, *Anguilla rostrata*". In: *Proceedings of the Royal Society of London. Series B: Biological Sciences* 271.1557 (Dec. 22, 2004), pp. 2535–2540. ISSN: 0962-8452, 1471-2954. DOI: 10.1098/rspb.2004.2901.

[148] Vincent J. T. Van Ginneken and Gregory E. Maes. "The European eel (Anguilla anguilla, Linnaeus), its Lifecycle, Evolution and Reproduction: A Literature Review". In: *Reviews in Fish Biology and Fisheries* 15.4 (Nov. 2005), pp. 367–398. ISSN: 0960-3166, 1573-5184. DOI: 10.1007/s11160-006-0005-8.

[149] Jason T. Vance and Stephen P. Roberts. "The effects of artificial wing wear on the flight capacity of the honey bee Apis mellifera". In: *Journal of Insect Physiology* 65 (June 2014), pp. 27–36. ISSN: 00221910. DOI: 10.1016/j.jinsphys.2014.04.003.

[150] John J. Videler. *Fish Swimming*. Dordrecht: Springer Netherlands, 1993. ISBN: 978-94-010-4687-9 978-94-011-1580-3. DOI: 10.1007/978-94-011-1580-3.

[151] James M. Wakeling and Ian A. Johnston. "Muscle Power Output Limits Fast-Start Performance in Fish". In: *Journal of Experimental Biology* 201.10 (May 15, 1998), pp. 1505–1526. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.201.10.1505.

[152] Jeffrey A. Walker and Mark W. Westneat. “Labriform Propulsion in Fishes: Kinematics of Flapping Aquatic Flight in the Bird Wrasse *Gomphosus Varius* (Labridae)”. In: *Journal of Experimental Biology* 200.11 (June 1, 1997), pp. 1549–1569. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.200.11.1549.

[153] Jeffrey A. Walker and Mark W. Westneat. “Mechanical performance of aquatic rowing and flying”. In: *Proceedings of the Royal Society of London. Series B: Biological Sciences* 267.1455 (Sept. 22, 2000), pp. 1875–1881. ISSN: 0962-8452, 1471-2954. DOI: 10.1098/rspb.2000.1224.

[154] C. S. Wardle, J. J. Videler, and J. D. Altringham. “Tuning In To Fish Swimming Waves: Body Form, Swimming Mode And Muscle Function”. In: *Journal of Experimental Biology* 198.8 (Aug. 1, 1995), pp. 1629–1636. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.198.8.1629.

[155] P. W. Webb. “Effects of Partial Caudal-Fin Amputation on the Kinematics and Metabolic Rate of Underyearling Sockeye Salmon (*Oncorhynchus Nerka*) At Steady Swimming Speeds”. In: *Journal of Experimental Biology* 59.3 (Dec. 1, 1973), pp. 565–582. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.59.3.565.

[156] Torkel Weis-Fogh. “Flapping Flight and Power in Birds and Insects, Conventional and Novel Mechanisms”. In: *Swimming and Flying in Nature*. Ed. by Theodore Y.-T. Wu, Charles J. Brokaw, and Christopher Brennen. Boston, MA: Springer US, 1975, pp. 729–762. ISBN: 978-1-4757-1328-2 978-1-4757-1326-8. DOI: 10.1007/978-1-4757-1326-8_18.

[157] Torkel Weis-Fogh. “Quick Estimates of Flight Fitness in Hovering Animals, Including Novel Mechanisms for Lift Production”. In: *Journal of Experimental Biology* 59.1 (Aug. 1, 1973), pp. 169–230. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.59.1.169.

[158] C. D. Wilga and G. V. Lauder. “Three-Dimensional Kinematics and Wake Structure of the Pectoral Fins During Locomotion in Leopard Sharks *Triakis Semifasciata*”. In: *Journal of Experimental Biology* 203.15 (Aug. 1, 2000), pp. 2261–2278. ISSN: 0022-0949, 1477-9145. DOI: 10.1242/jeb.203.15.2261.

[159] W. Keats Wilkie, Paul H. Mirick, and Chester W. Langston. *Rotating Shake Test and Modal Analysis of a Model Helicopter Rotor Blade*. NASA Langley Technical Report Server, 1997.

[160] Alexander P. Willmott, Charles P. Ellington, and Adrian L. R. Thomas. “Flow visualization and unsteady aerodynamics in the flight of the hawkmoth, *Manduca sexta*”. In: *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 352.1351 (Mar. 29, 1997), pp. 303–316. ISSN: 0962-8436, 1471-2970. DOI: 10.1098/rstb.1997.0022.

[161] E. Roberts Wood et al. “On Developing and Flight Testing a Higher Harmonic Control System”. In: *Journal of the American Helicopter Society* 30.1 (Jan. 1, 1985), pp. 3–20. ISSN: 2161-6027. DOI: 10.4050/JAHS.30.1.3.

[162] Jianghao Wu and Mao Sun. “The influence of the wake of a flapping wing on the production of aerodynamic forces”. In: *Acta Mechanica Sinica* 21.5 (Nov. 2005), pp. 411–418. ISSN: 0567-7718, 1614-3116. DOI: 10.1007/s10409-005-0064-4.

[163] Dan Xia et al. “The three-dimensional hydrodynamics of thunniform swimming under self-propulsion”. In: *Ocean Engineering* 110 (Dec. 2015), pp. 1–14. ISSN: 00298018. DOI: 10.1016/j.oceaneng.2015.10.008.

[164] Xi Xia and Kamran Mohseni. “A theory on leading-edge vortex stabilization by spanwise flow”. In: *Journal of Fluid Mechanics* 970 (Sept. 10, 2023), R1. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/jfm.2023.613.

[165] Lili Yang et al. “Clap-and-Fling Mechanism of Climbing-Flight Coccinella Septempunctata”. In: *Biomimetics* 9.5 (May 9, 2024), p. 282. ISSN: 2313-7673. DOI: 10.3390/biomimetics9050282.

[166] Renguo Yang et al. “Fuzzy Neural Network PID Control Used in Individual Blade Control”. In: *Aerospace* 10.7 (July 9, 2023), p. 623. ISSN: 2226-4310. DOI: 10.3390/aerospace10070623.

[167] Renguo Yang et al. “Reducing Helicopter Vibration Loads by Individual Blade Control with Genetic Algorithm”. In: *Machines* 10.6 (June 15, 2022), p. 479. ISSN: 2075-1702. DOI: 10.3390/machines10060479.

[168] J M Zanker and K G Gotz. “The wing beat of *Drosophila Melanogaster* . II. Dynamics”. In: *Philosophical Transactions of the Royal Society of London. B, Biological Sciences* 327.1238 (Feb. 28, 1990), pp. 19–44. ISSN: 0080-4622, 2054-0280. DOI: 10.1098/rstb.1990.0041.

[169] Jun-Duo Zhang, Hyung Jin Sung, and Wei-Xi Huang. “Hydrodynamic interaction of dorsal fin and caudal fin in swimming tuna”. In: *Bioinspiration & Biomimetics* 17.6 (Nov. 1, 2022), p. 066004. ISSN: 1748-3182, 1748-3190. DOI: 10.1088/1748-3190/ac84b8.

[170] Qiang Zhong, Haibo Dong, and Daniel B. Quinn. “How dorsal fin sharpness affects swimming speed and economy”. In: *Journal of Fluid Mechanics* 878 (Nov. 10, 2019), pp. 370–385. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/jfm.2019.612.

[171] Q. Zhu et al. “Three-dimensional flow structures and vorticity control in fish-like swimming”. In: *Journal of Fluid Mechanics* 468 (Oct. 10, 2002), pp. 1–28. ISSN: 0022-1120, 1469-7645. DOI: 10.1017/S002211200200143X.

Appendix A

FLAPPING PROPULSION SUPPLEMENTARY MATERIALS

Tables A.1, A.2, and A.3 reproduced with permission from

[1] Meredith Leigh Hooper, Isabel Scherl, and Morteza Gharib. *Supplementary material from "Bio-Inspired Compensatory Strategies for Damage to Flapping Robotic Propulsors"*. 2024. DOI: 10.6084/M9.FIGSHARE.C.7302782.

Organism	**Force type**	**Propulsive area removed (%)**	**Amplitude change**	**Frequency change**	**Force change**
Scardineus erythrophthalmus [15]	Thrust	100	Increase	Increase	No change
Merlangius merlangus [53]	Thrust	100	Increase	Increase	
Oncorhynchus nerka [155]	Thrust	approx. 33	No change	No change	No change
Oncorhynchus nerka [155]	Thrust	approx. 66	No change	No change	Decrease
Etheostoma edwini [20]	Thrust	approx. 50			No change
Gambusia holbrooki [132]	Thrust	1-30			Gaussian
Spinibarbus sinensis [43]	Thrust	100	Increase	No change	Decrease
Cyprinius carpio [43]	Thrust	100	Increase	Increase	Decrease
Carassius auratus [43]	Thrust	100	Increase	Increase	Decrease
Spinibarbus sinensis [43]	Thrust	59	Increase	No change	Decrease
Cyprinius carpio [43]	Thrust	53	No change	No change	Decrease
Carassius auratus [43]	Thrust	76	No change	Increase	No change
Mylopharyngodon piceus [19]	Thrust	approx. 50		No change	No change
Mylopharyngodon piceus [19]	Thrust	100		Decrease	Decrease
Pontia occidentalis [76]	Side force	approx. 15-20		Increase	No change
Bombus terrestris [61]	Side force	10	No change	Increase	
Limantria dispar [67]	Side force	44		Increase	Decrease
Pieris rapae [67]	Side force	49		No change	Decrease
Bombus impatiens [18]	Side force	approx. 25	No change	Decrease	Decrease
Manduca sexta [37]	Side force	10-36	No change	Increase	No change
Apis mellifera [149]	Side force	3-37	Increase	Increase	Decrease
Bombus impatiens [124]	Side force	11.2-23.4		Increase	Decrease
Drosophila hydei [105]	Side force	10-50	Increase	Increase	No change

Table A.1: Adaptations to damage of natural swimmers and flyers. Tabular summary of previous work on natural swimmers' and flyers' adaptations to catastrophic propulsor damage. Boxes left blank were not investigated by the cited authors. "Force type" denotes direction of force produced by the organism in the nomenclature of the robotic system. Force change is listed as an abstraction of change in critical speed, hovering ability, or other parameter revealing force trends.

	Intact Fin	**Amp 1**	**Amp 2**	**Amp 3**	**Amp 4**	**Amp 5**	**Intact Fin**	**Amp 1**	**Amp 2**	**Amp 3**	**Amp 4**	**Amp 5**
Force	Thrust	Thrust	Thrust	Thrust	Thrust	Thrust	SF	SF	SF	SF	SF	SF
Type	Ellipse	Ellipse	Ellipse	Ellipse	Ellipse	Ellipse	Ellipse	Ellipse	Ellipse	Ellipse	Ellipse	Ellipse
ϕ	24.8°	31.7°	32.0°	31.5°	31.5°	30.0°	15.2°	32.1°	31.3°	29.6°	18.7°	16.8°
ψ	12.9°	15.2°	14.8°	15.2°	13.6°	13.4°	8.1°	15.0°	15.2°	13.1°	13.7°	11.6°
χ	26.7°	51.8°	70.0°	-68.2°	-63.4°	-70.0°	-70.0°	-57.0°	-70.0°	-70.0°	-70.0°	-41.2°
β **[rad]**	5.3	2.3	2.7	4.7	6.1	5.8	2.9	2.9	4.2	4.3	4.8	4.4
S	0	1	2	2	2	1	3	4	3	2	3	4
γ	1.2	1.2	1.2	1.2	1.2	1.2	1.1	1.2	1.2	1.1	1.1	1.1
$\mathbf{K_v}$	0.6	0.8	0.5	0.5	0.9	0.5	0.2	0.4	0.4	0.4	0.3	0.3
γ	0.6	0.9	0.7	1.0	0.7	0.6	0.3	0.4	0.8	0.7	0.6	0.4
ω **[Hz]**	0.70	0.78	0.70	0.72	0.73	0.79	0.86	0.90	0.70	0.75	0.86	0.88
Closeness to setpoint force	0.0029	0.0130	0.0067	0.0032	0.0085	0.0016	-0.0062	-0.1618	-0.0091	-0.0027	-0.0189	0.0198
Fitness	0.1157	0.1088	0.0920	0.1049	0.0879	0.0810	0.0087	0.1633	0.0854	0.0933	0.1059	0.0408

Table A.2: Converged optimal trajectory parameters and fitness values. SF abbreviates side force, amp. abbreviates amputated and accel. abbreviates acceleration.

		Mode 1	**Mode 2**	**Mode 3**	**Mode 4**	**Mode 5**
Thrust, F_z (N)						
Intact Fin	**Amplitude**	0.426	0.398	0.021	0.011	0.011
	Phase	±159.56°	±48.89°	±177.34°	±120.36°	±161.67°
Amputated 1	**Amplitude**	0.635	0.621	0.218	0.056	0.047
	Phase	±103.49°	±76.55°	±20.81°	±145.21°	±142.06°
Amputated 2	**Amplitude**	0.516	0.369	0.184	0.062	0.056
	Phase	±23.48°	±125.74°	±9.98°	±93.37°	±5.65°
Amputated 3	**Amplitude**	0.529	0.247	0.212	0.151	0.125
	Phase	±62.16°	±92.58°	±153.66°	±123.06°	±28.40°
Amputated 4	**Amplitude**	0.489	0.173	0.117	0.046	0.040
	Phase	±38.45°	±127.65°	±15.77°	±21.72°	±164.29°
Amputated 5	**Amplitude**	0.623	0.413	0.086	0.059	0.046
	Phase	±15.93°	±141.96°	±58.95°	±21.01°	±163.90°
Thrust, AOA (°)						
Intact Fin	**Amplitude**	12.470	1.307	1.071	0.104	0.096
	Phase	±157.52°	±94.07°	±29.12°	±95.22°	±177.38°
Amputated 1	**Amplitude**	24.078	2.421	1.730	0.894	0.334
	Phase	±33.73°	±130.23°	±94.90°	±144.40°	±54.98°

Amputated 2	**Amplitude**	34.141	2.062	1.694	0.760	0.344
	Phase	±51.12°	±147.45°	±147.42°	±164.05°	±137.80°
Amputated 3	**Amplitude**	31.820	2.477	1.310	0.872	0.466
	Phase	±26.88°	±127.26°	±142.02°	±1.54°	±130.22°
Amputated 4	**Amplitude**	32.261	1.675	1.097	0.665	0.357
	Phase	±64.60°	±124.38°	±90.83°	±158.41°	±163.57°
Amputated 5	**Amplitude**	34.161	1.791	1.493	0.831	0.528
	Phase	±58.25°	±36.30°	±126.34°	±165.65°	±154.59°
SF, F_{x^*} (N)						
Intact Fin	**Amplitude**	0.398	0.330	0.261	0.055	0.037
	Phase	±18.06°	±9.30°	±74.55°	±7.51°	±127.28°
Amputated 1	**Amplitude**	0.327	0.205	0.120	0.085	0.062
	Phase	±26.50°	±165.29°	±15.78°	±3.84°	±143.56°
Amputated 2	**Amplitude**	0.623	0.320	0.107	0.091	0.057
	Phase	±77.79°	±23.67°	±118.16°	±89.96°	±30.22°
Amputated 3	**Amplitude**	0.632	0.423	0.124	0.094	0.080
	Phase	±43.22°	±122.70°	±31.45°	±84.18°	±131.92°
Amputated 4	**Amplitude**	1.026	0.319	0.081	0.073	0.029
	Phase	±59.48°	±140.94°	±156.39°	±8.42°	±95.61°
Amputated 5	**Amplitude**	0.507	0.236	0.108	0.049	0.020
	Phase	±129.22°	±52.89°	±32.76°	±119.62°	±85.85°
SF, F_{y^*} (N)						
Intact Fin	**Amplitude**	0.311	0.161	0.136	0.043	0.041
	Phase	±13.70°	±156.69°	±172.83°	±92.64°	±17.64°
Amputated 1	**Amplitude**	0.344	0.252	0.166	0.074	0.055
	Phase	±7.73°	±40.19°	±115.02°	±119.67°	±50.90°
Amputated 2	**Amplitude**	0.971	0.261	0.144	0.080	0.068
	Phase	±27.17°	±9.29°	±144.88°	±63.86°	±104.22°
Amputated 3	**Amplitude**	1.099	0.339	0.110	0.106	0.096
	Phase	±64.39°	±80.68°	±44.30°	±134.95°	±97.31°
Amputated 4	**Amplitude**	1.059	0.270	0.079	0.054	0.050
	Phase	±72.67°	±50.24°	±145.82°	±131.27°	±12.61°
Amputated 5	**Amplitude**	0.454	0.189	0.072	0.059	0.010
	Phase	±121.11°	±159.35°	±156.15°	±28.10°	±70.54°
SF, AOA (°)						

Intact Fin	**Amplitude**	34.840	1.137	0.443	0.246	0.174
	Phase	±11.87°	±174.42°	±6.46°	±76.56°	±147.26°
Amputated 1	**Amplitude**	27.753	2.173	1.834	0.487	0.200
	Phase	±7.86°	±97.60°	±6.82°	±89.40°	±173.98°
Amputated 2	**Amplitude**	33.648	1.865	1.722	1.028	0.433
	Phase	±42.04°	±163.96°	±43.42°	±13.47°	±66.99°
Amputated 3	**Amplitude**	33.566	1.211	1.082	0.828	0.481
	Phase	±49.85°	±54.73°	±56.28°	±0.25°	±44.98°
Amputated 4	**Amplitude**	34.395	0.990	0.724	0.479	0.221
	Phase	±60.26°	±147.83°	±19.17°	±14.79°	±45.92°
Amputated 5	**Amplitude**	19.955	1.038	0.473	0.250	0.068
	Phase	±102.73°	±157.76°	±74.68°	±134.85°	±52.74°

Table A.3: Fourier decompositions of optimal force and angle of attack traces. The first five Fourier modes are reported. SF abbreviates side force. Note that the angle of attack Fourier decompositions are reported in the x-y frame for thrust, whilst for side force they are reported in the x*-y* frame.

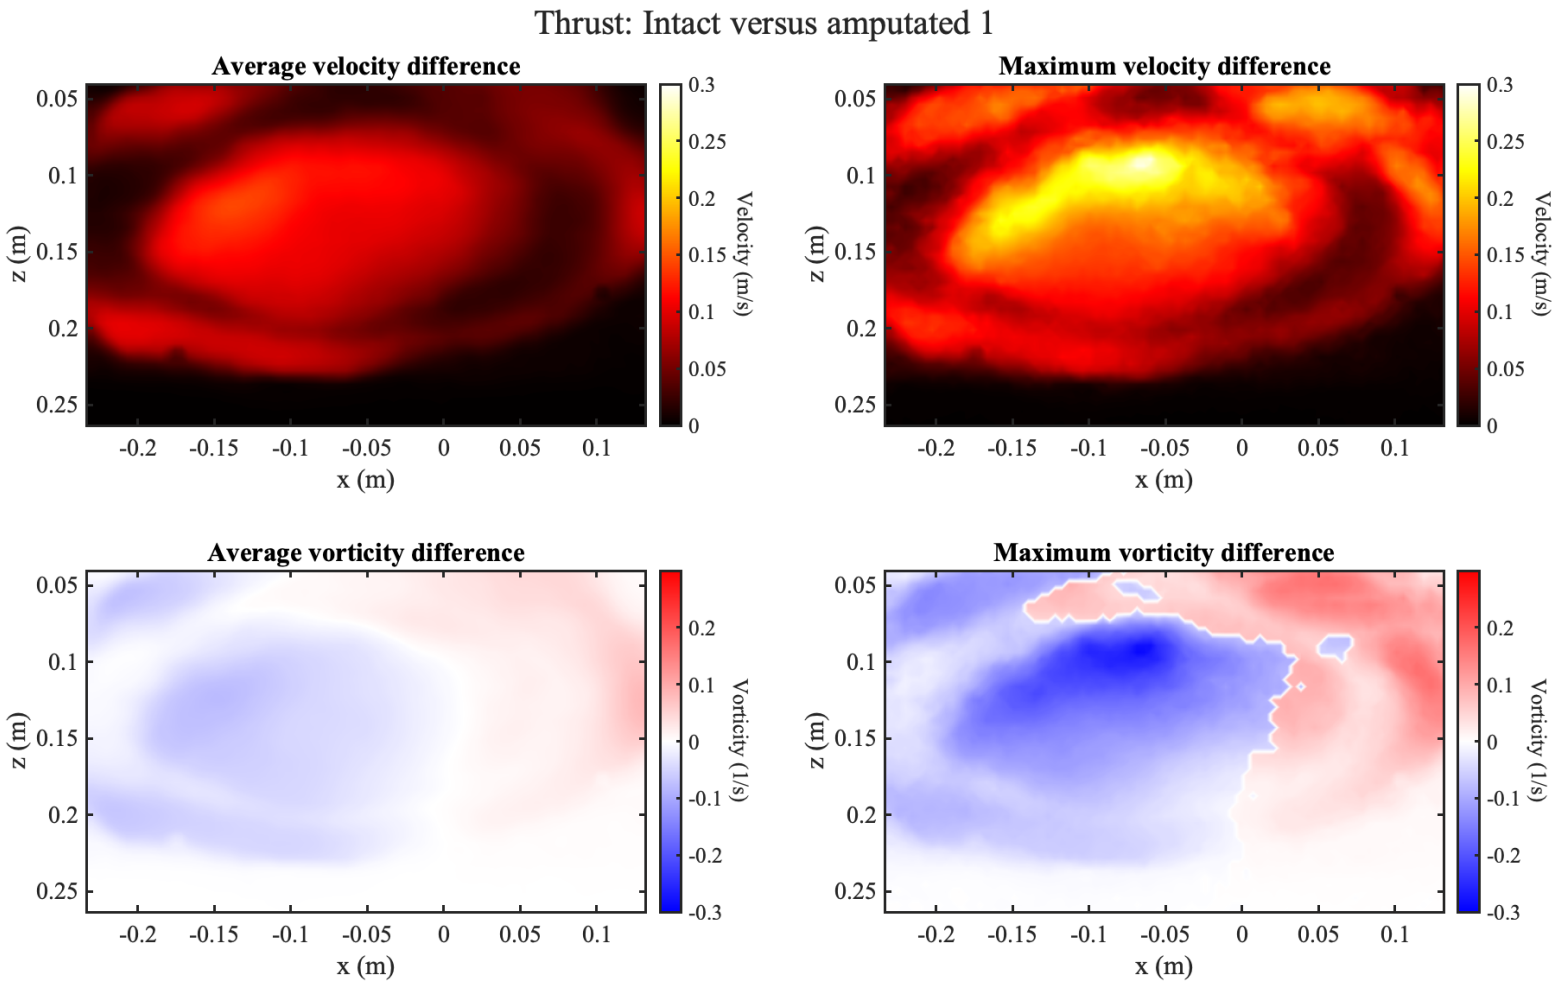

Figure A.1: Difference in hydrodynamics between the intact and amputated fin 1 producing thrust. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

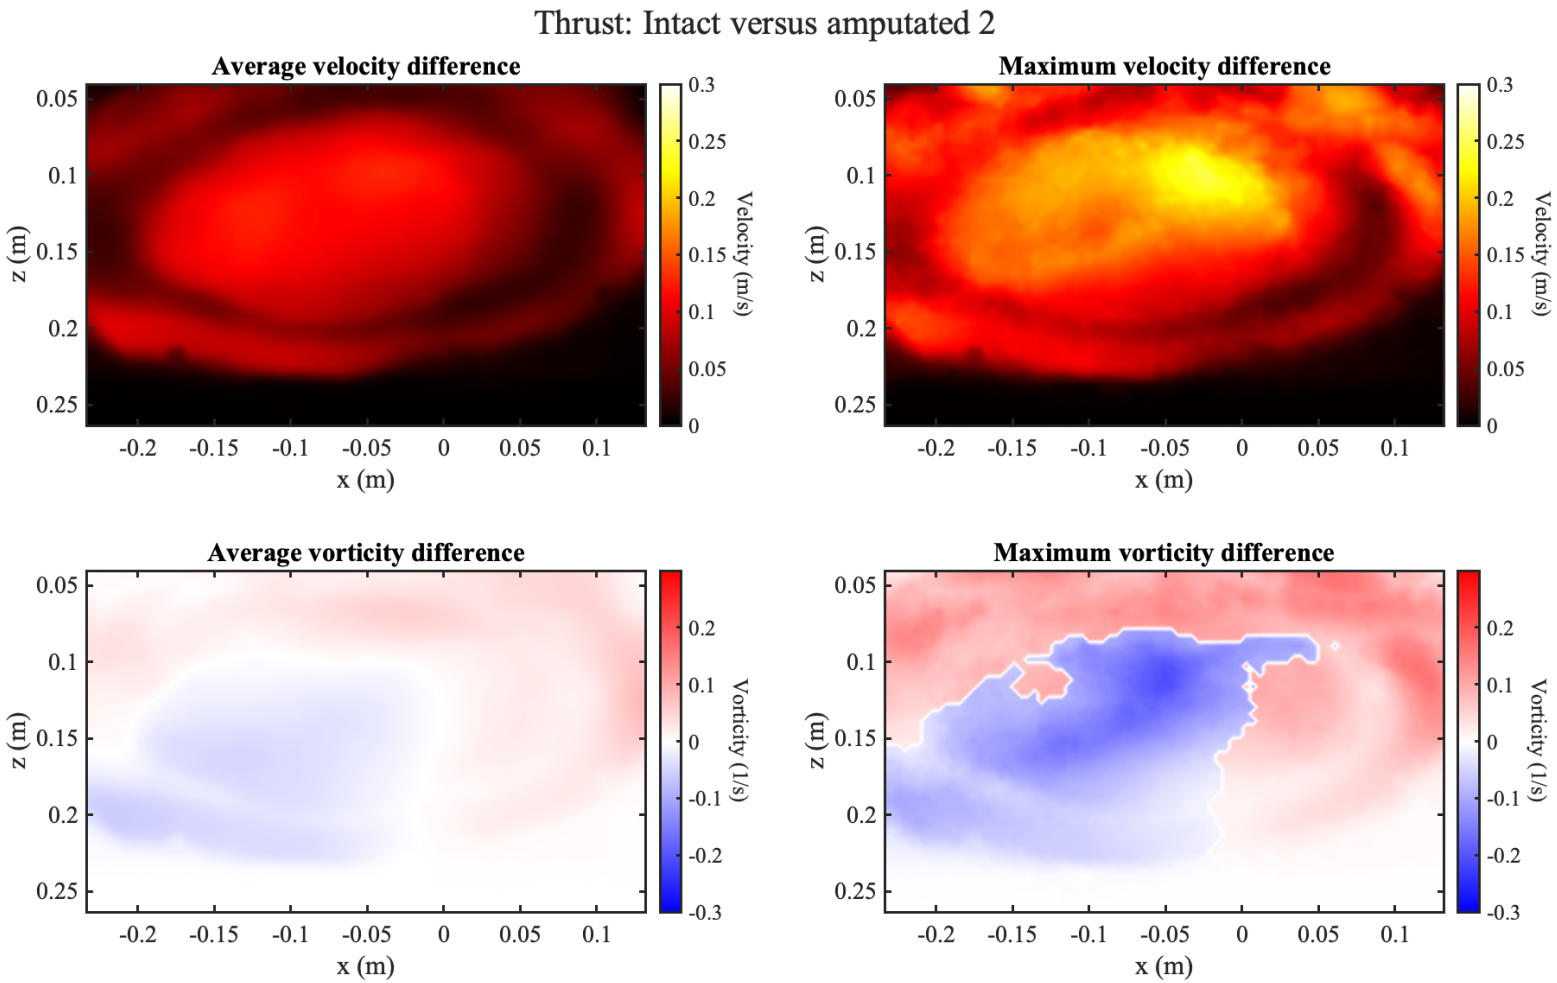

Figure A.2: Difference in hydrodynamics between the intact and amputated fin 2 producing thrust. The left column of shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

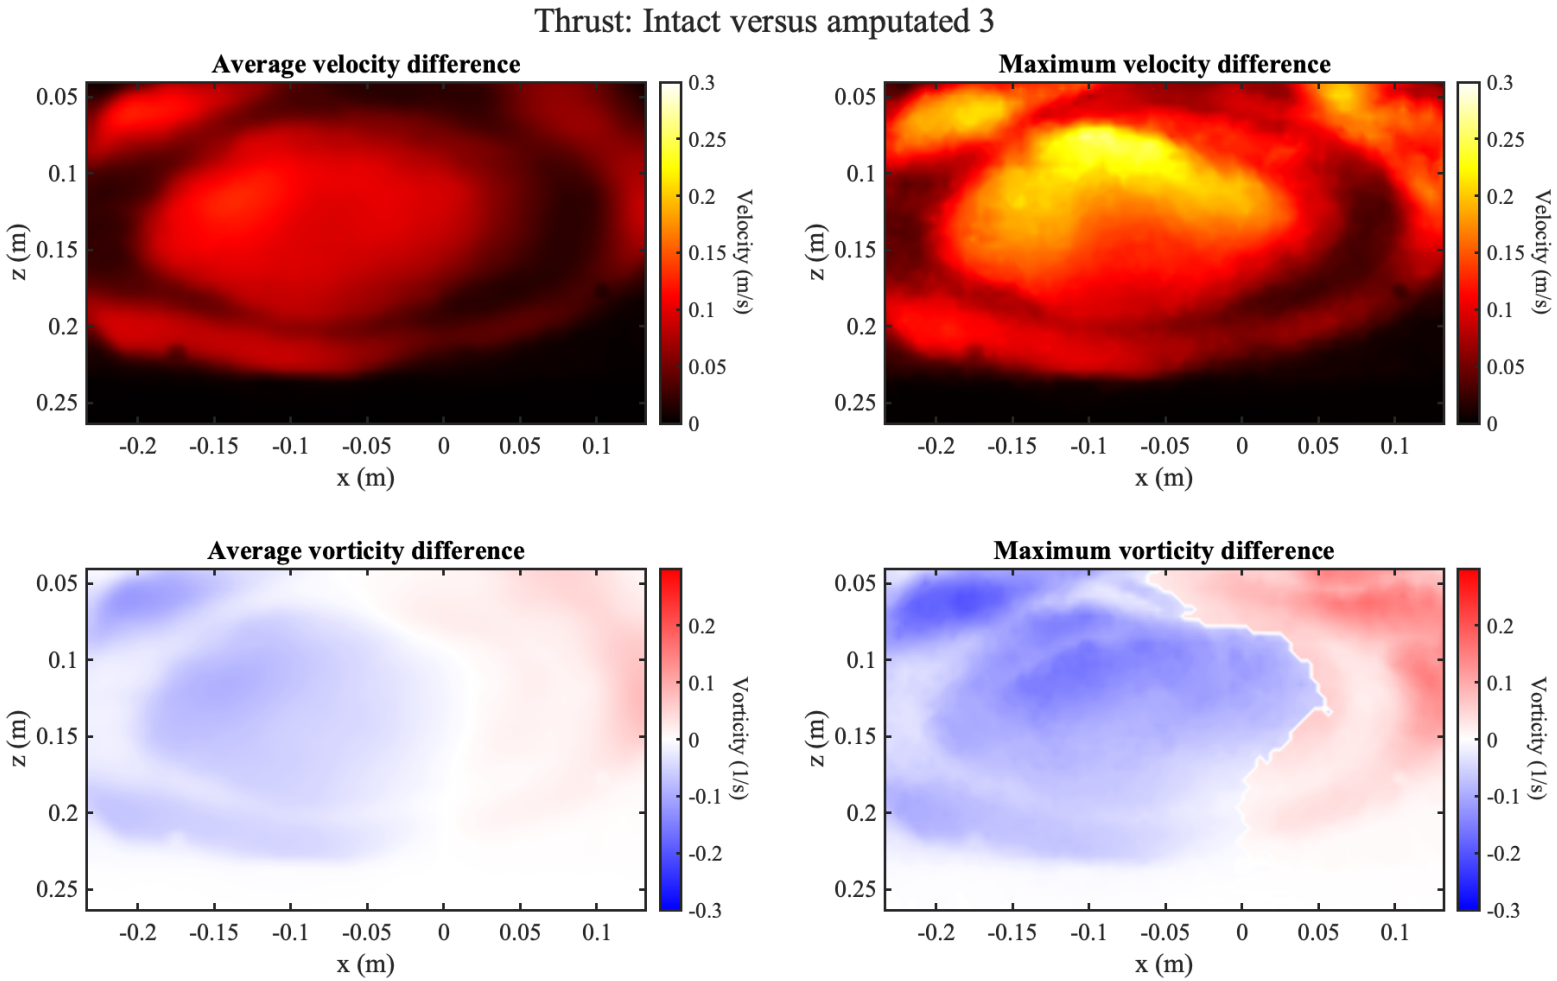

Figure A.3: Difference in hydrodynamics between the intact and amputated fin 3 producing thrust. The left column of shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

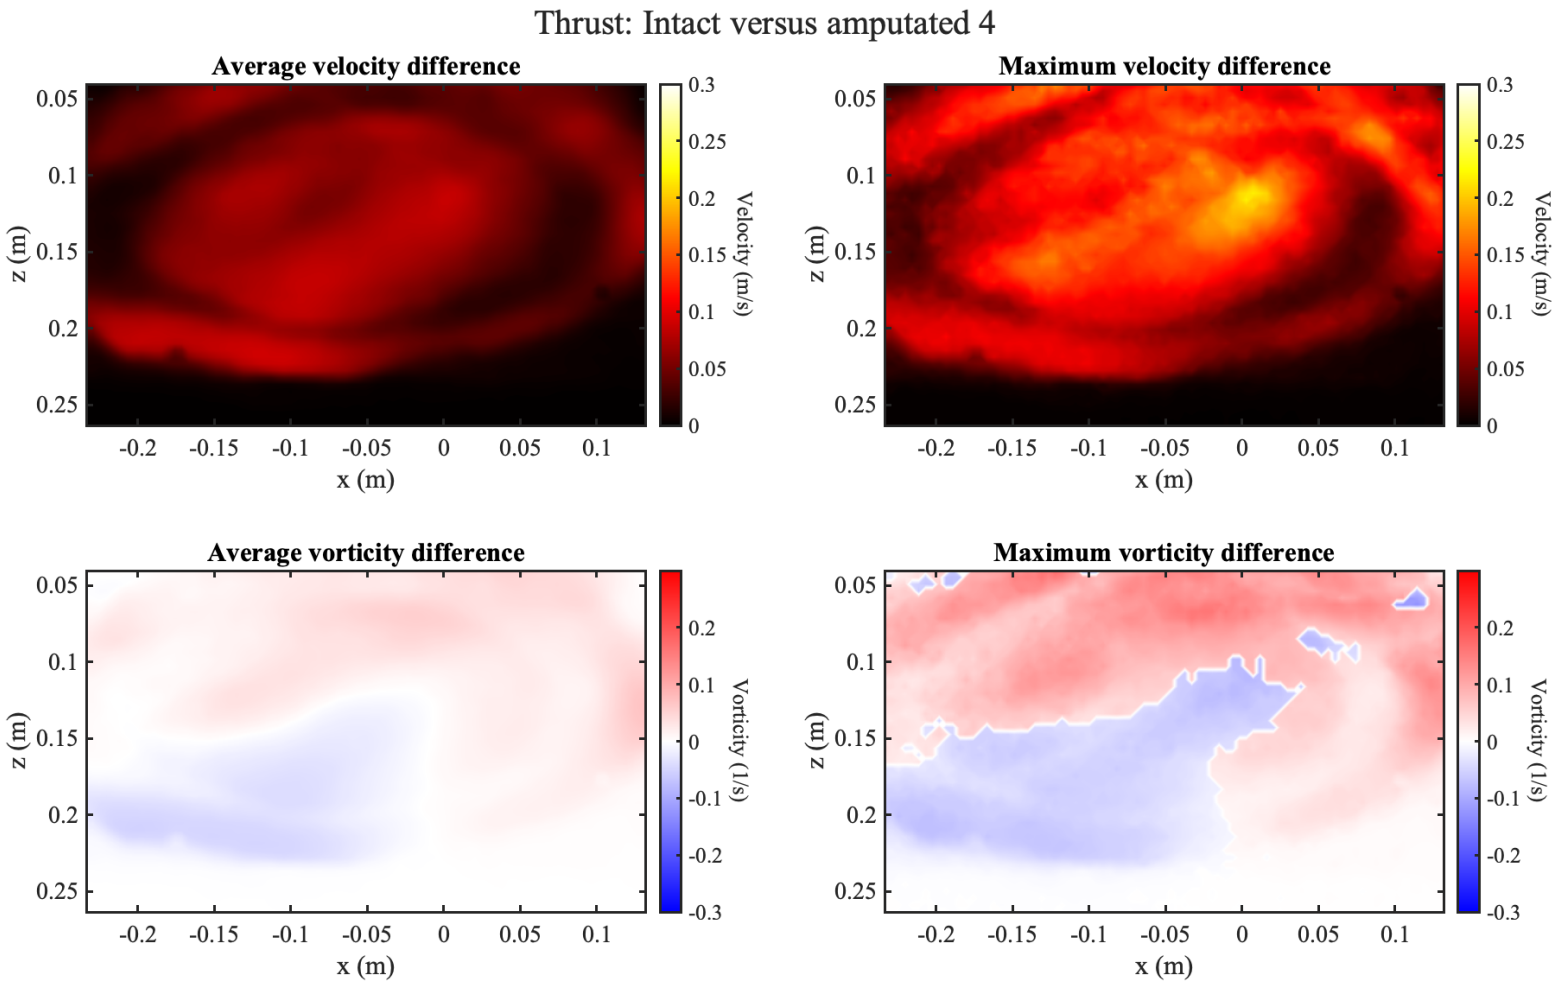

Figure A.4: Difference in hydrodynamics between the intact and amputated fin 4 producing thrust. The left column of shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

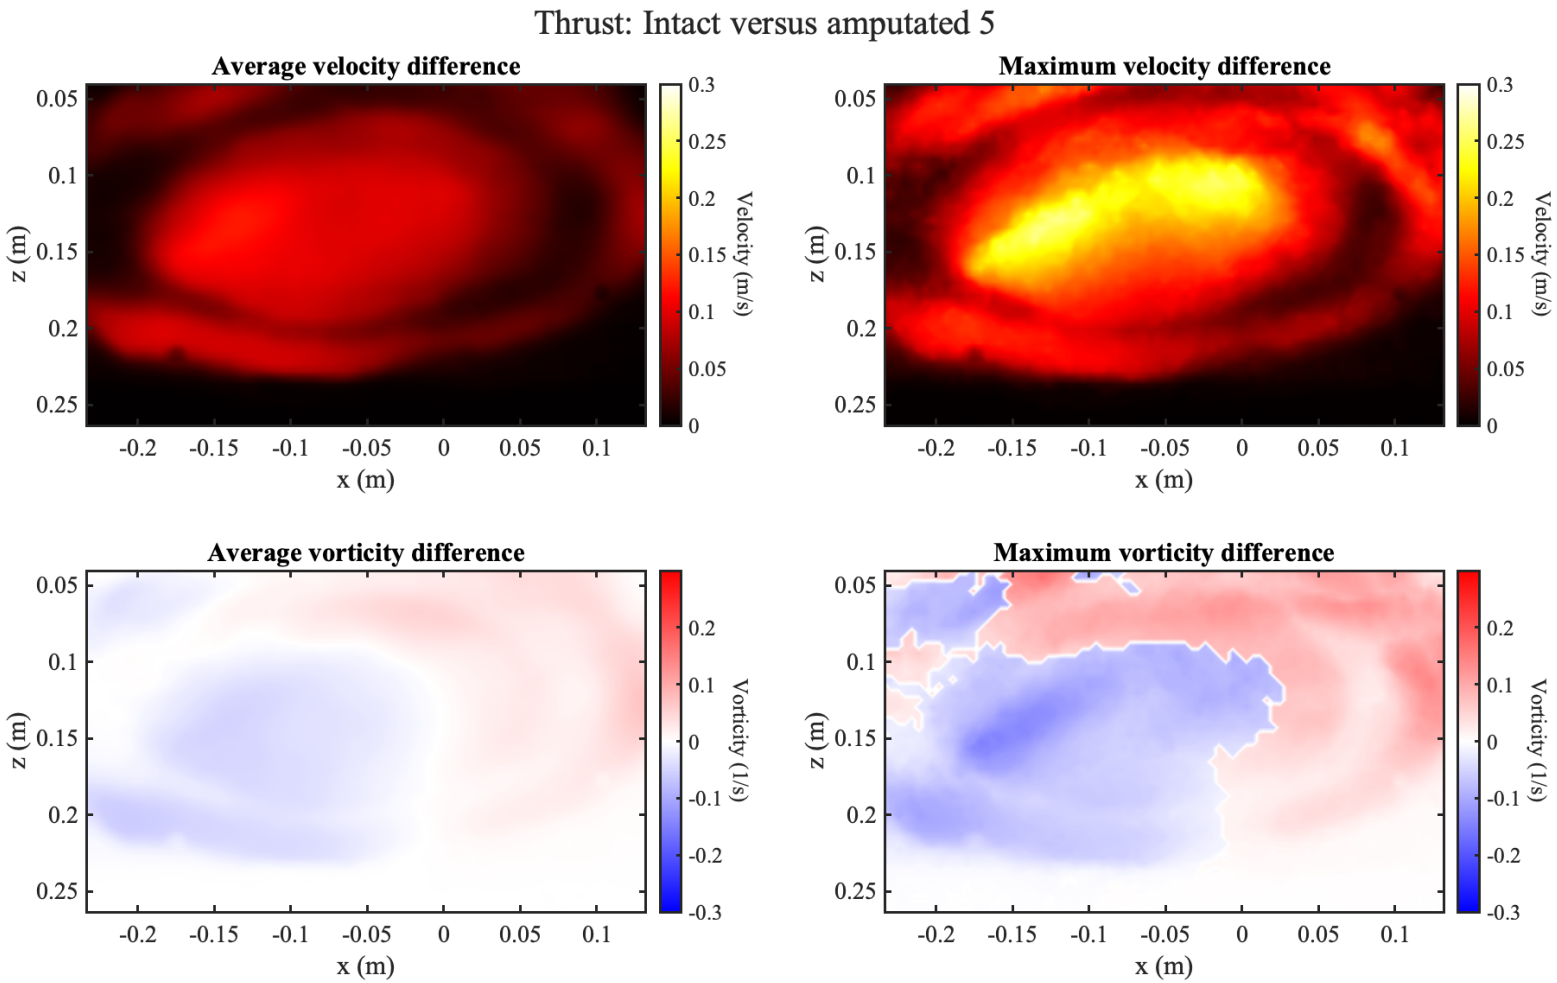

Figure A.5: Difference in hydrodynamics between the intact and amputated fin 5 producing thrust. The left column of shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

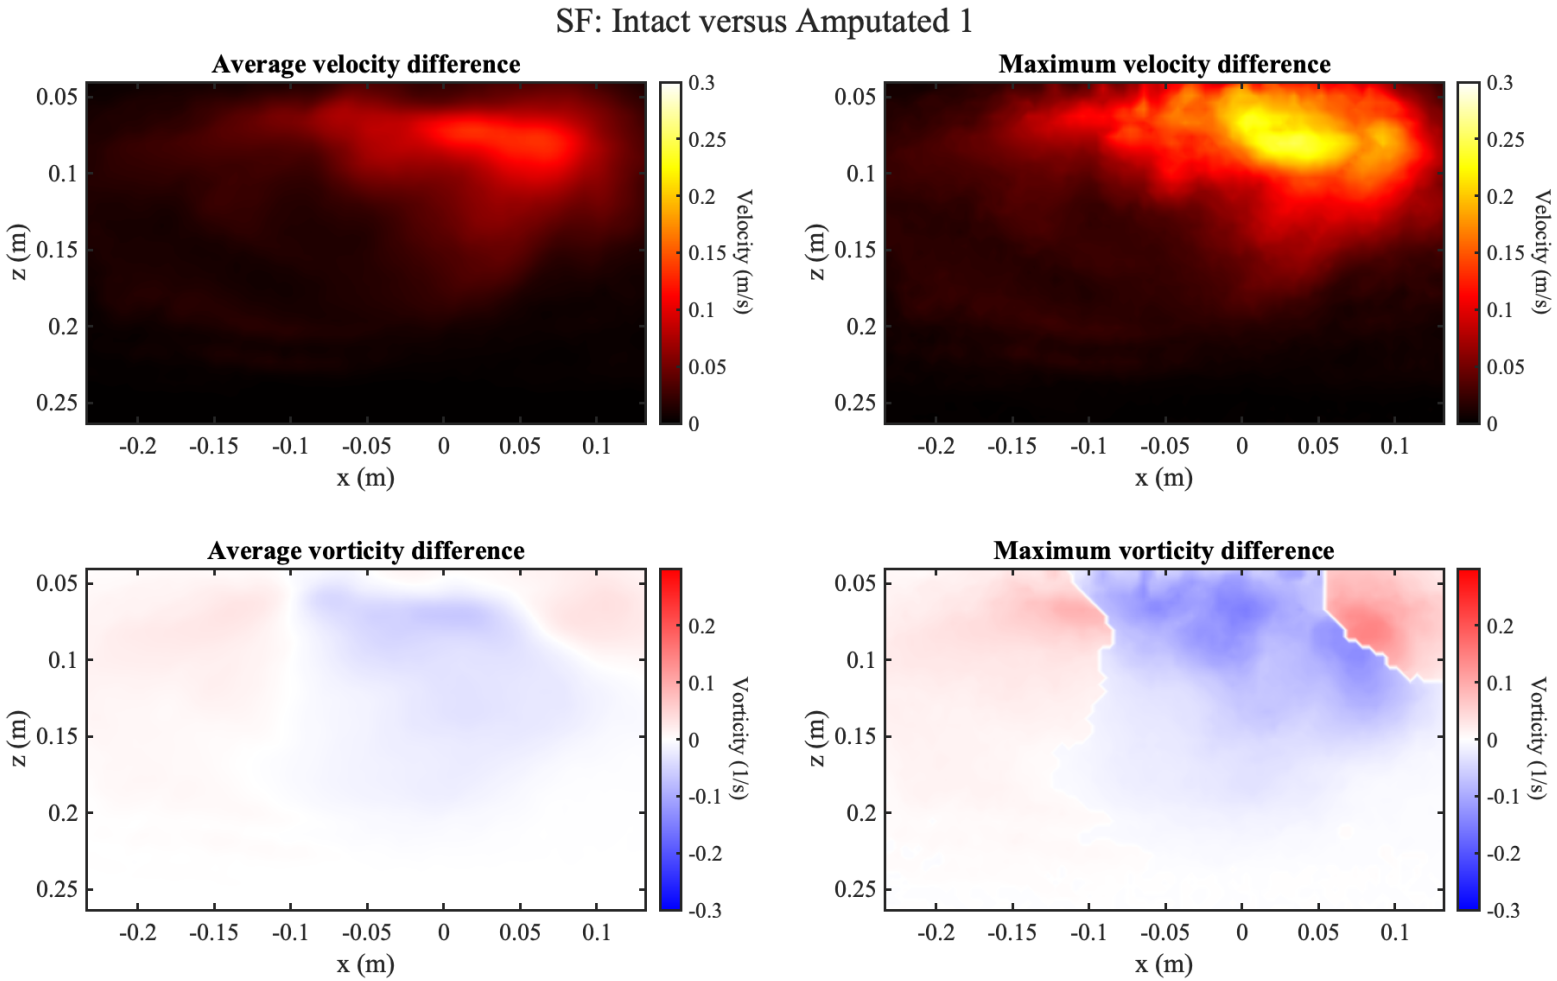

Figure A.6: Difference in hydrodynamics between the intact and amputated fin 1 producing side force. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

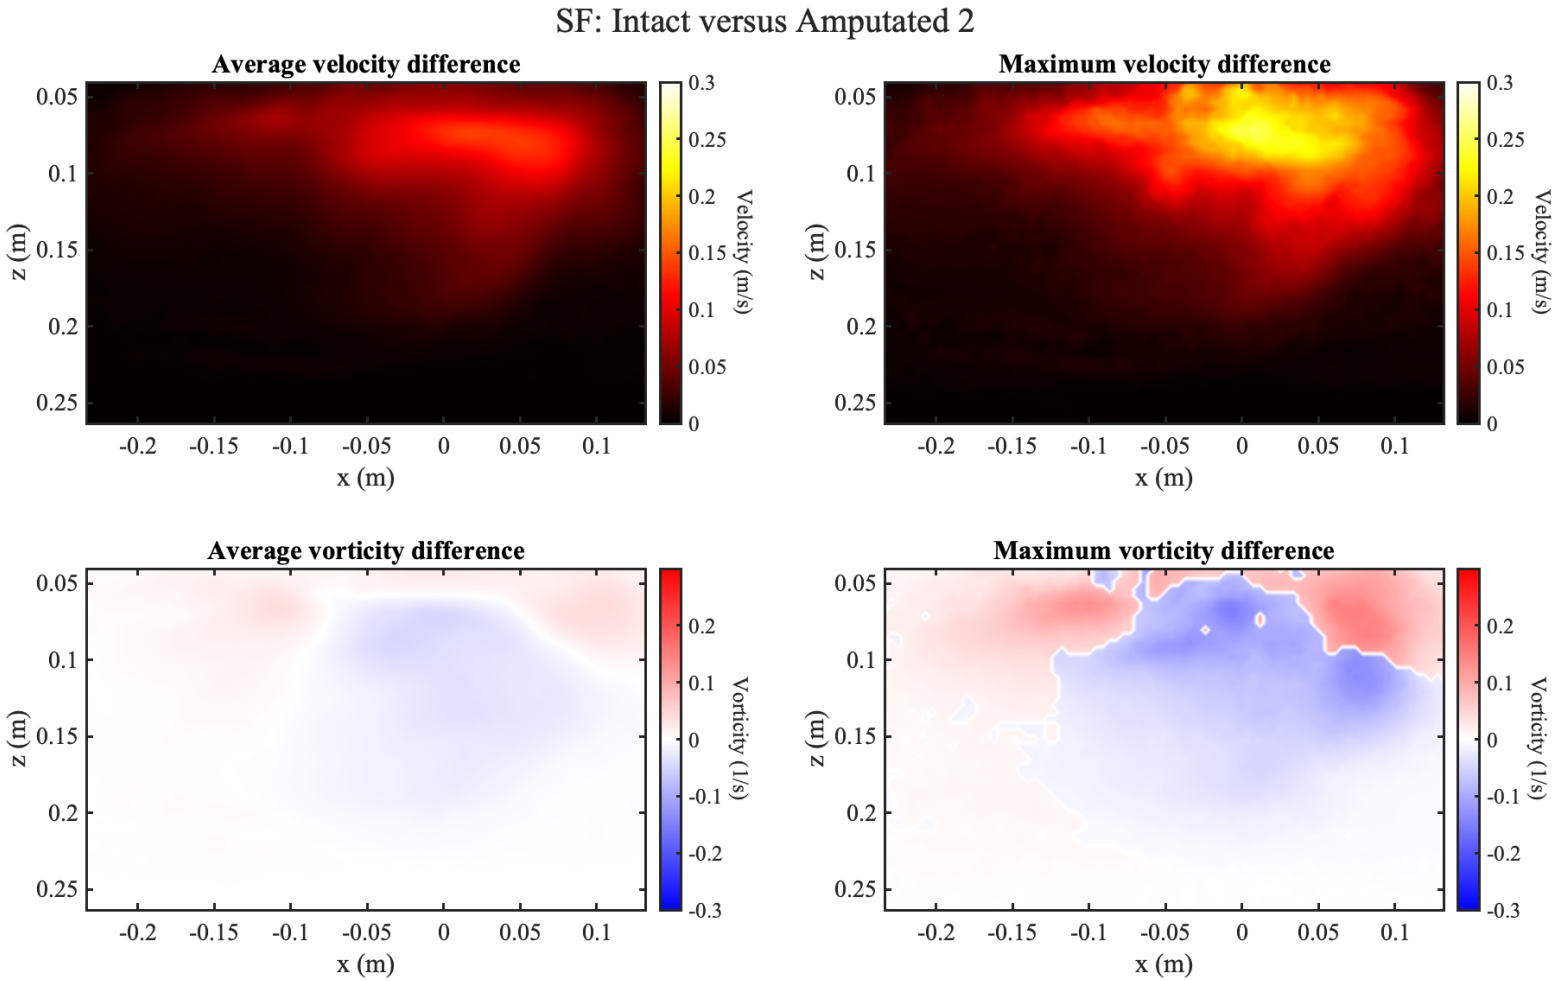

Figure A.7: Difference in hydrodynamics between the intact and amputated fin 2 producing side force. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

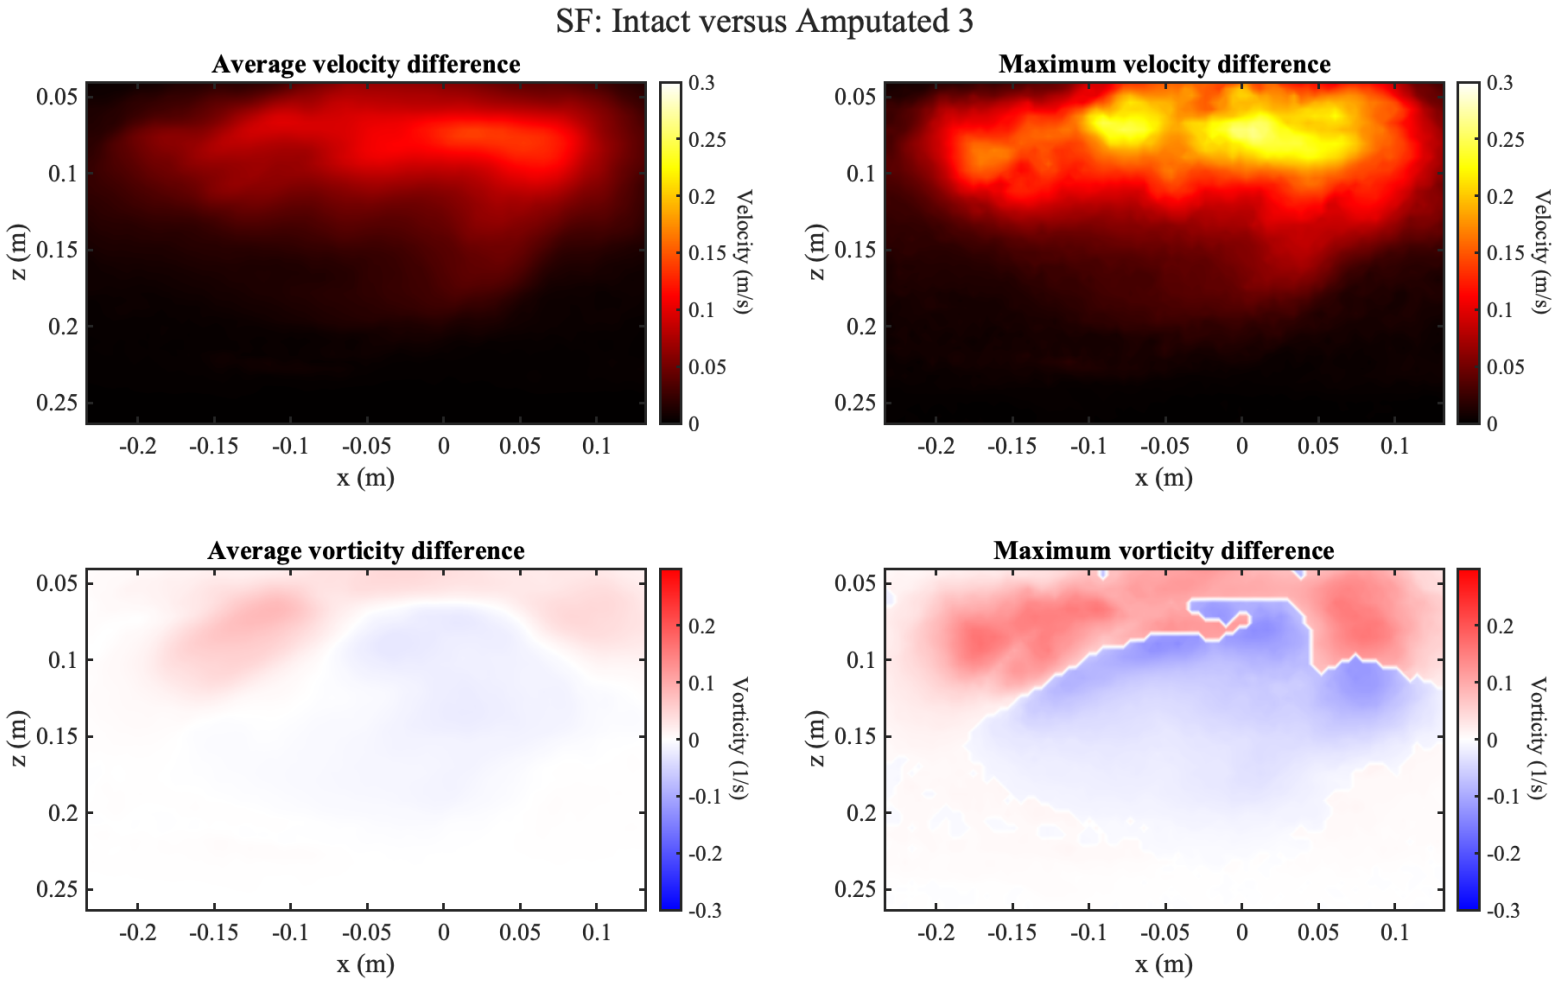

Figure A.8: Difference in hydrodynamics between the intact and amputated fin 3 producing side force. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

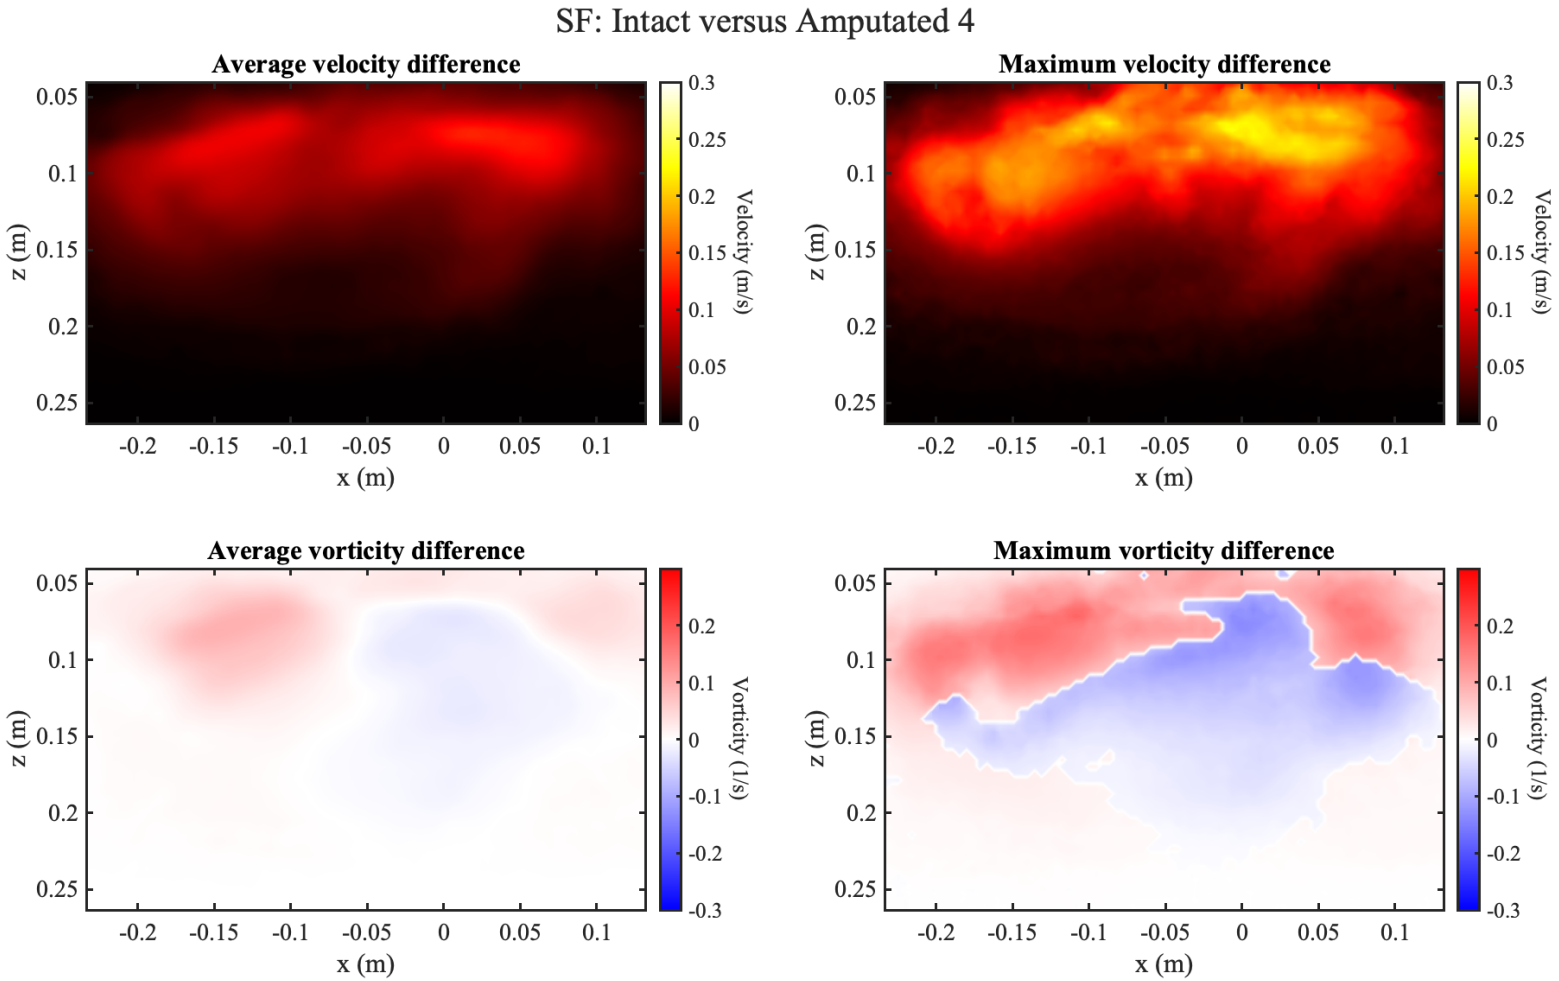

Figure A.9: Difference in hydrodynamics between the intact and amputated fin 4 producing side force. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

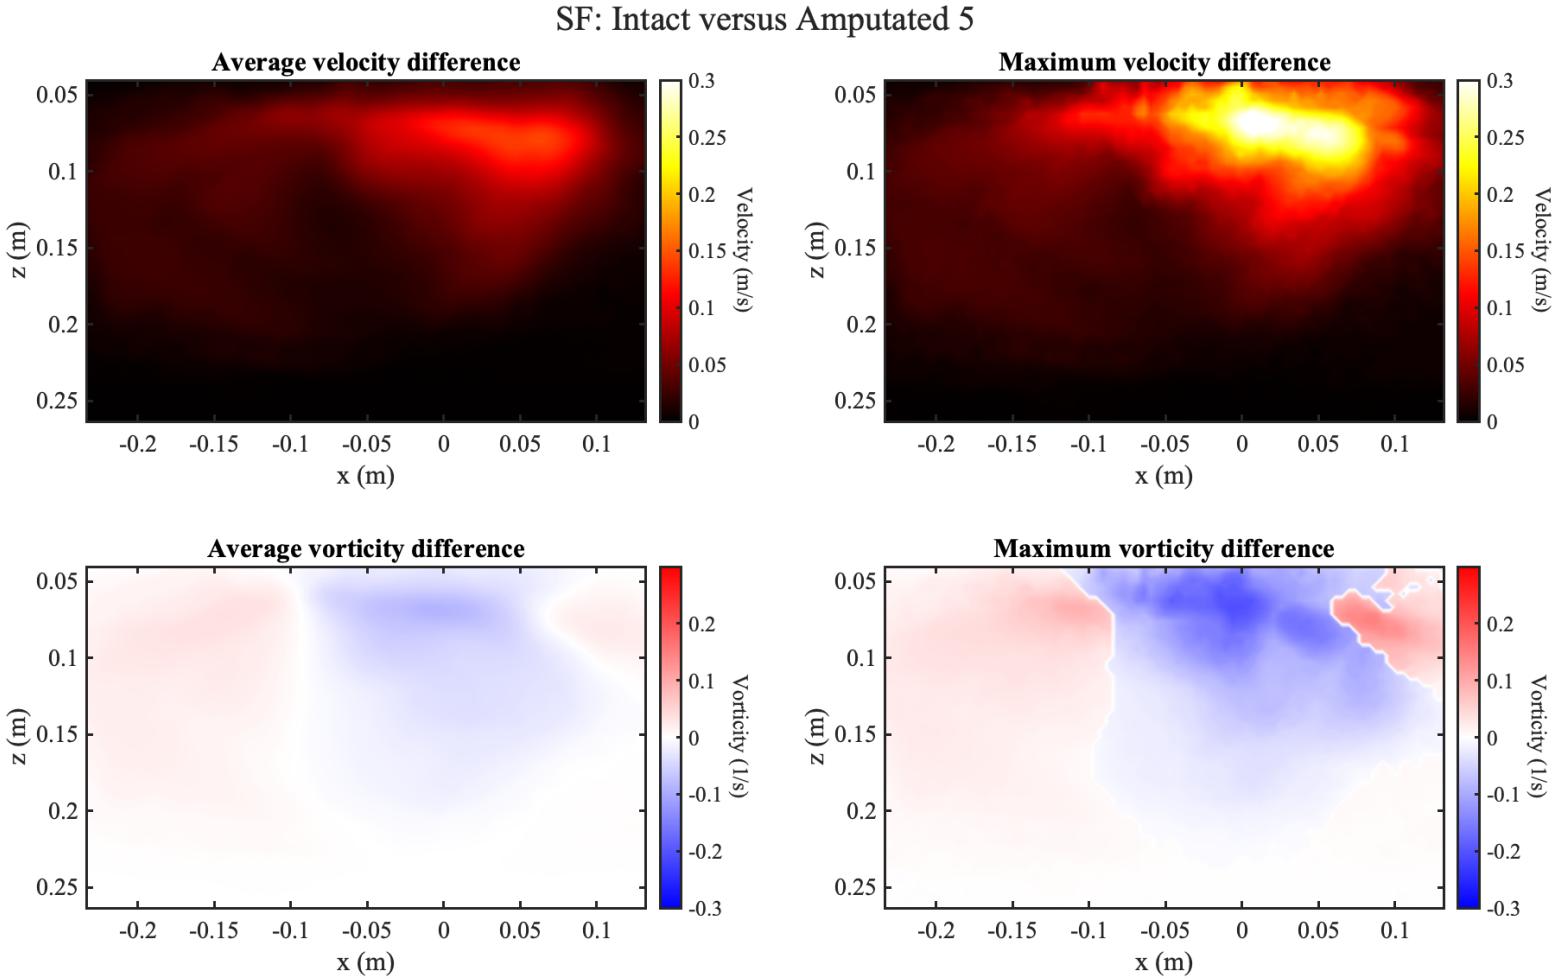

Figure A.10: Difference in hydrodynamics between the intact and amputated fin 5 producing side force. The left column shows the average difference and the right column the maximum difference for each spatial point between the 25-cycle phase averages.

Fin	Force Type	Maximum Speed [m/s]	Average Speed [m/s]	Minimum Vorticity [1/s]	Maximum Vorticity [1/s]	Average Vorticity Magnitude [1/s]	Vorticity Standard Deviation [1/s]
Intact Fin	Thrust	0.2730	0.0231	-16.02	15.70	1.1199	1.4483
Amputated 1	Thrust	0.2452	0.0190	-14.20	16.01	0.4523	0.3627
Amputated 2	Thrust	0.2104	0.0295	-15.28	17.85	0.6089	0.8476
Amputated 3	Thrust	0.2160	0.0182	-7.88	16.13	0.3194	0.2443
Amputated 4	Thrust	0.1874	0.0240	-16.17	15.14	0.8055	1.2932
Amputated 5	Thrust	0.2266	0.0234	-14.81	15.33	0.5639	0.6802
Intact Fin	Side Force	0.2421	0.0255	-18.13	15.81	0.8086	1.5370
Amputated 1	Side Force	0.1424	0.0191	-8.40	12.35	0.2947	0.6046
Amputated 2	Side Force	0.2051	0.0317	-8.91	14.15	0.3561	0.8325
Amputated 3	Side Force	0.2939	0.0358	-17.16	23.05	0.4803	1.1885
Amputated 4	Side Force	0.2539	0.0356	-17.40	16.38	0.4836	1.1573
Amputated 5	Side Force	0.2397	0.0162	-8.76	7.89	0.2466	0.5158
Intact Fin - Amputated	Thrust Difference	0.2400	0.0355	-0.2881	0.1768	0.0261	1.46–1.82
Intact Fin - Amputated	SF Difference	0.3234	0.0294	-0.2053	0.1824	0.0163	0.024–0.03

Table A.4: Distribution parameters of speed and vorticity data captured in the y=0 plane. These parameters are calculated from the 25-cycle phase-averaged data.