



Swansea University
Prifysgol Abertawe

Engineering Design Optimisation Using Computational Fluid Dynamics and Human-AI Collaboration

Jakub Vincalek

Submitted in fulfilment of the requirement for the degree of
Doctor of Philosophy

Faculty of Engineering and Science
Swansea University
January 2024

Declaration

I, Jakub Vincalek, hereby declare that this thesis titled 'Engineering Design Optimisation using Computational Fluid Dynamics and Human-AI Collaboration' and the work presented herein are my own, and I confirm that:

This work has not been previously accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Signed 

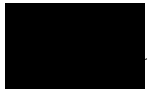
Date ..14/06/2024.....

This work is the result of my own independent study/investigations, except where otherwise stated. Other sources are clearly acknowledged by giving explicit references. I understand that failure to do this amounts to plagiarism and will be considered grounds for failure of this work and the degree examination as a whole.

Signed 

Date ..14/06/2024.....

I hereby give my consent for my work, if accepted, to be archived and available for reference use, and for the title and summary to be made available to outside organisations.

Signed 

Date ..14/06/2024.....

Abstract

The primary contributions of this thesis are a comparison of a ducted winglet to two other geometries and a user study which investigated the relationship between engineers and an optimisation algorithm. With an ever increasing emphasis on sustainable and renewable energy production, new technology can increase the efficiency of existing infrastructure such as wind turbines. Aerodynamic devices known as winglets can be fitted to the end of wind turbine blades to increase efficiency and reduce negative downstream effects. A patented winglet introduces a duct that channels freestream air from the bottom of the winglet to the top. The effectiveness of this design is simulated using computational fluid dynamics and compared to two alternative designs across a range of angles of attack. The patented winglet had lower induced drag across a range of angles of attack. Following the comparison, optimisation methods including evolutionary algorithms are employed to further increase the efficiency of the winglet. These algorithms can be used when engineers would otherwise rely on intuition, preconceptions, and theory to try and find optimal designs. The manner in which engineers utilise and engage with an evolutionary algorithm known as MAP-Elites is evaluated through a user study. 12 participants were given 20 minutes each to design a car that could travel over an inclined course with obstacles as far as possible. Their behaviour was compared to survey answers they provided before and after designing cars. Participants who engaged with the MAP-Elites

algorithm outperformed baseline designs created by the computer and participants who did not engage with the algorithm. Participants were more likely to use the MAP-Elites algorithm even if they were unaware they were using it or if they had stated that they did not trust optimisation algorithms. Shortening the optimisation cycle is explored during the optimisation of the ducted winglet. The effects of increasing uncertainty in the results are explored through two studies on the robustness of evolutionary algorithms and a study on 2-dimensional aerofoils. The studies show that the optimisation cycle can be reduced to a certain extent while maintaining the original ranking of the designs.

Abbreviations

| | |
|------------|------------------------------------|
| AI | Artificial Intelligence |
| CAD | Computer Aided Engineering |
| CFD | Computational Fluid Dynamics |
| EA | Evolutionary Algorithm |
| HCI | Human Computer Interaction |
| IEA | Interactive Evolutionary Algorithm |
| LHS | Latin Hypercube Sampling |
| MI | Mixed Initiative |
| MOO | Multi-Objective Optimisation |
| QCA | Qualitative Content Analysis |
| TSP | Travelling Salesman Problem |

Acknowledgements

Firstly, I would like to thank my family, Iva, Vaclav, Tereza, and Annie, for your everlasting support, patience, and kindness. To my beloved Cati, thank you for your love and for always being there when I needed you the most. A special thank you to the Isaacs family who have treated me like one of their own.

I would like to also thank my friends as well as the post-graduates with whom I had the pleasure of sharing my Ph.D. journey, notably Ben Smith, Daniel Fellows, Jamie Lam, Jo Parke, Richard Davies, Matthew Ritchie, Finn Monaghan, and the other members of my cohort. Especially, Connor Rees whose kindness and perseverance pushed me to accomplish even the most daunting feats.

This work was supported in part by Prof. Matt Jones and the EPIC-CDT and funded in part by the Engineering and Physical Science Research Council.



My sincerest gratitude goes to my supervisors Prof. Evans and Dr. Walton. Ben, thank you for your continued enthusiasm, your overwhelming positivity, and for guiding me along my entire journey in Swansea. Sean, this thesis along with all my work would not have been possible without your steadfast commitment to excellence and your refreshing outlook on life.



Věnováno mé nejdražší babičce a dědečkovi

All work and no play makes Jack a dull boy.

Contents

| | | |
|----------|---|----------|
| 1 | Introduction | 1 |
| 1.1 | Objectives | 4 |
| 1.2 | Layout of the Thesis | 5 |
| 2 | Novel Ducted Winglet Analysis | 7 |
| 2.1 | Introduction | 8 |
| 2.2 | Background | 9 |
| 2.2.1 | Computer Aided Design | 9 |
| 2.2.2 | Fundamental Aerodynamics of Wind Turbines | 10 |
| 2.2.3 | Induced Drag Reduction Techniques | 13 |
| 2.2.4 | Aerodynamic Testing Methods | 15 |
| 2.2.5 | Computational Fluid Dynamics | 16 |
| 2.3 | Methods and Materials | 29 |
| 2.3.1 | Novel Ducted Winglet | 29 |
| 2.3.2 | Parameterisation | 33 |
| 2.3.3 | Simulation Setup | 34 |
| 2.3.4 | Mesh Convergence Study | 40 |
| 2.4 | Results | 43 |
| 2.4.1 | Lift | 43 |
| 2.4.2 | Drag | 45 |
| 2.4.3 | Lift-to-Drag Ratio | 47 |
| 2.4.4 | Velocity and Pressure Changes Within the Duct | 47 |
| 2.4.5 | Induced Drag and the Trefftz Plane | 49 |
| 2.5 | Discussion | 53 |
| 2.5.1 | Patent Assertion 1 | 53 |
| 2.5.2 | Patent Assertion 2 | 54 |
| 2.5.3 | Patent Assertion 3 | 56 |
| 2.5.4 | Patent Assertion 4 | 56 |
| 2.5.5 | Patent Assertion 5 | 56 |

| | | |
|----------|---|------------|
| 2.5.6 | Patent Assertion 6 | 59 |
| 2.5.7 | Patent Assertion 7 | 60 |
| 2.6 | Conclusion | 60 |
| 3 | Novel Ducted Winglet Optimisation | 62 |
| 3.1 | Introduction | 63 |
| 3.2 | Background | 63 |
| 3.2.1 | Latin Hypercube Sampling | 64 |
| 3.2.2 | Evolutionary Algorithms | 64 |
| 3.2.3 | Bayesian Optimisation | 77 |
| 3.3 | Method and Materials | 78 |
| 3.3.1 | Design of Experiments | 79 |
| 3.3.2 | Evolutionary Algorithm | 80 |
| 3.3.3 | Bayesian Optimisation | 83 |
| 3.3.4 | 2D Evolutionary Algorithm Test Case | 83 |
| 3.4 | Results | 87 |
| 3.4.1 | Design of Experiments | 87 |
| 3.4.2 | Evolutionary Algorithm | 88 |
| 3.4.3 | Bayesian Optimisation | 90 |
| 3.5 | Reduced Parameter Optimisation | 92 |
| 3.5.1 | Updated Method and Materials | 93 |
| 3.5.2 | Reduced Parameter Count Results | 97 |
| 3.6 | Discussion | 100 |
| 3.6.1 | Design of Experiments | 100 |
| 3.6.2 | Evolutionary Algorithm | 101 |
| 3.6.3 | Bayesian Optimisation | 101 |
| 3.6.4 | Reduced Parameters | 102 |
| 3.7 | Conclusion | 103 |
| 4 | Effects of Uncertainty on Reducing Simulation Iterations | 104 |
| 4.1 | Introduction | 105 |
| 4.1.1 | Evolutionary Algorithm Process Affected by Uncertainty | 105 |
| 4.2 | Background | 106 |
| 4.2.1 | Travelling Salesman Problem | 107 |
| 4.2.2 | Taylor Series | 108 |
| 4.3 | Method and Materials | 109 |
| 4.3.1 | Travelling Salesman Problem | 110 |
| 4.3.2 | Effect of Equation Fidelity using Taylor Series | 114 |
| 4.3.3 | 2D CFD Simulations | 119 |
| 4.4 | Results | 120 |
| 4.4.1 | Effects of Rank Change on Optimisation Success | 120 |

| | | |
|----------|---|------------|
| 4.4.2 | Effects of Equation Fidelity on Optimisation Success | 122 |
| 4.4.3 | Effects of Reducing Simulation Iterations in 2D | 124 |
| 4.5 | Discussion | 124 |
| 4.6 | Conclusion | 127 |
| 5 | Illumination Through Collaboration | 129 |
| 5.1 | Introduction | 130 |
| 5.2 | Background | 131 |
| 5.2.1 | Human-Computer Interaction | 131 |
| 5.2.2 | Mixed Initiative Design Systems | 133 |
| 5.2.3 | Does mapping elites illuminate search spaces? and Other Questions for MAP-Elites | 135 |
| 5.2.4 | A Summary of Results from the Online Study | 137 |
| 5.3 | Method and Materials | 137 |
| 5.3.1 | Participants | 139 |
| 5.3.2 | Car Designer | 140 |
| 5.3.3 | Survey and Qualitative Content Analysis | 144 |
| 5.4 | Results | 147 |
| 5.4.1 | Fitness | 147 |
| 5.4.2 | Survey | 151 |
| 5.5 | Discussion | 173 |
| 5.5.1 | Fitness | 173 |
| 5.5.2 | Survey | 174 |
| 5.5.3 | Limitation of the Methodology | 175 |
| 5.6 | Conclusion | 175 |
| 6 | Conclusion | 177 |
| 6.1 | Contributions and Findings | 178 |
| 6.1.1 | Engineering Design Optimisation | 178 |
| 6.1.2 | Human-AI Collaboration | 179 |
| 6.2 | Future Work | 181 |
| 6.2.1 | Chapter 2 | 181 |
| 6.2.2 | Chapter 3 | 181 |
| 6.2.3 | Chapter 4 | 182 |
| 6.2.4 | Chapter 5 | 182 |
| 7 | Publications and Conference Presentations | 184 |
| A | Appendix | 203 |

Chapter 1

Introduction

Wind turbines have been in use for centuries; their original use was for grinding wheat or other similar food products. They were also commonly used for water pumping [1]. Wind turbines work on the simple principle of converting the kinetic energy available in wind traditionally to mechanical or more recently electrical energy. The amount of energy that can be converted from kinetic energy to electrical energy is dependent on the area covered by the blades of the wind turbine. So, larger blades cover more area, and can therefore convert more kinetic energy to electrical energy [2].

Based solely on this simple objective, wind turbines should have infinitely long blades to capture an infinite amount of kinetic energy. Evidently, this is impossible. As blades get longer they also get heavier; they also bend more. Additionally, the tip of the blades has a higher velocity than the root; the shape of the blade at each end of the blade must be different to account for the difference in velocity. These factors influence the design of the blade and are known as the design constraints.

The job of the engineer is to design for the objective within the constraints. A blade that is too heavy might cause issues for the supporting tower. A blade that is too short won't capture enough kinetic energy. A blade that lacks sufficient rigidity might break or hit the supporting tower. The weight of the blade, the length, and its stiffness can all be described numerically and are known as the design parameters.

Engineers have a selection of knowledge and tools to use to help them with their design objective: theoretical and mathematical works, empirical studies, and simulation software. All these tools are used to create the first iteration of the design. Something is learned and then the second design iteration is created. Then the third, and fourth, and fifth, and so on. This repetitive process can be automated to save time to find the optimum by incorporating search algorithms. The design parameters are inputs to the algorithm and the entire process is known as parametric design optimisation.

The basic design loop framework follows these steps: build, measure, learn. In the building phase, initial parameters are input and either a physical model is made or a representative model is made in engineering drawing software. In the measure phase, the design is tested and data are gathered. There are aerodynamic testing methods available to test the performance of wind turbines. The traditional method for performing aerodynamic analyses is the wind tunnel [3]. Air enters the wind tunnel, flows through a narrow chamber where the analysis is performed, and then exits out the back where the fan is located. The design of wind tunnels is rather straightforward and is an irreplaceable tool in the aerospace industry. That being said, wind tunnels have some drawbacks that limit their applicability. The first is their size they take up a significant amount of room compared to their testing area. The second is turnaround time only one design can be tested at a time and physical design have to be made for each iteration which limits how many can

be tested. Computer simulations which use computational fluid dynamics (CFD) have addressed these drawbacks and have significantly decreased design iteration turnaround time. Full scale models can now be tested without the need for manufacturing and multiple models can be tested simultaneously. As an example of the efficiency of CFD, Boeing managed to reduce their design iterations. In 1980 they tested 77 different physical models in the wind tunnel for their 767 program. By 2005, they had reduced the number down to 11 for their 787 program [4]. In the last phase, the data are analysed which guide and set the parameters for the next iteration. In the case of a wind turbine blade for example, the loop would look something like this: define an initial length, simulate the blade, then adjust the length of the blade based on the results of the simulation.

From this simplified version of the design loop, it is clear that the entire loop can be automated. The trouble with a fully automated artificial intelligence (AI) approach is that it can lead to results that are untrustworthy [5]. Moreover, the time taken to simulate each blade might be so long that the engineer would be able to find a more optimal solution in 2 iterations than an algorithm could in 10 iterations. The proposed solution in this work is to include the engineer in the automated design loop as a means of reducing overall design time. Additionally, including the engineer in the design loop provides increased engagement as a benefit [6]. Being able to explain how a design was created is equally as important as being able to explain how a design works. Uptake of automated design methods is limited in industry because engineers do not trust the results of these methods [7]. While people might not trust the results of automated methods, those who engage with and embed themselves in the design loop with optimisation algorithms create better designs than those who do not [8].

The performance of wind turbines can also be increased through the use of aerodynamic

devices known as winglets. This work investigates the working principle and effectiveness of a novel ducted winglet aimed at reducing the drag associated with increased lift by using CFD. The design incorporates a duct that channels freestream air through a traditional winglet and ejects the air at the tip of the blade. The design was created by John Smith and is patented by Fourth Dimensional Aerospace Technology. For the first time, the duct of the novel design is parameterised and tested for use on wind turbine blades. A parametric design optimisation study is also conducted to further enhance the performance of the novel design.

1.1 Objectives

There are two primary objectives of this thesis. The first is to investigate the effectiveness of a novel wingtip design concept in the context of wind turbine blade design. To do so, a manual, human-led aerospace design study of the novel ducted winglet is conducted to allow a better understanding of the underpinning aerodynamic properties of their design. The results are benchmarked against an unmodified design and a standard non-ducted winglet. The second objective of this thesis is to investigate the effectiveness of optimisation algorithms and how they affect the relationship between humans and computers in the optimisation cycle. A simple game-like car designer is used amongst a small group of participants with the algorithm compared to a control of randomly selected designs.

In addition to the primary objectives of this thesis, there are also secondary objectives relating to the performance of the novel ducted winglet and optimisation cycle turnover time. An optimisation study was conducted on the novel ducted winglet to further enhance its performance with respect to its total and induced drag. A study to measure

the effects of uncertainty in the data on the optimisation process was also conducted. By reducing the optimisation cycle time through a reduction in time spent on simulations was tested on benchmark cases.

1.2 Layout of the Thesis

The novel ducted winglet is compared to an unmodified and non-ducted winglet in Chapter 2. CFD simulations of the three geometries are evaluated based on lift and drag results from the simulations as well as a post-processing technique for induced drag. Observations about the novel ducted winglet are discussed and compared to assertions from the original patent for the design. The novel ducted winglet lowered the induced drag for most angles of attack tested but still had a lower lift-to-drag ratio than the non-ducted winglet [9].

Following the comparison among the geometries, an optimisation study is performed in Chapter 3 on the novel ducted winglet using various optimisation techniques. An initial 18-parameter design was used but long simulation times and design constraints led to a reduced 4-parameter design. The optimisation study further lowered the induced drag of the novel ducted winglet and contributed also to a reduction in total drag.

The long simulation times led to a study in Chapter 4, which investigates the effects of changing the rank order of the evaluated designs and the necessary equation fidelity needed to establish the correct rank order on optimisation success. Small changes in the rank order minimally affected the optimisation success and suggests that CFD simulation times could be drastically reduced when using evolutionary computation methods.

The relationship between engineers and algorithms meant to help in the optimisation

cycle is explored through a user study. Twelve participants were recruited to use a simple car designer and create a car which could travel as far as possible in a limited amount of time. Participants who engaged with the algorithm were more likely to create better designs than those who did not. Results from a survey taken by participants before and after using the car designer as well as their usage data is analysed in Chapter 5 [8].

Finally, the contributions and findings of this thesis are presented in Chapter 6.

Chapter 2

Novel Ducted Winglet Analysis

In this chapter, 4DAT's patented novel ducted winglet is introduced which aims to reduce the wake caused by wind turbine blades. The primary contribution of this chapter is a comparison between the novel ducted winglet and two benchmark geometries by evaluating their values of lift, drag and induced drag. The geometries were tested using in-house simulation software which has been validated on industrial applications and used commercially available software to analyse the results. A method known as the Trefftz plane analysis is used to evaluate the strength of the downstream wake and to measure to the induced drag. The novel ducted winglet reduced the induced drag and wake compared to the benchmarks. The lift-to-drag ratio of the novel ducted winglet is lower than a standard winglet though. The results from this chapter validate the patent results of the novel ducted winglet and the parameterisation of the duct provides the starting point for the following chapter which deals with optimisation.

2.1 Introduction

The goal of aerodynamic designers and engineers is to create designs that pass through air as efficiently as possible. Over the years, considerable work has been dedicated to the field of aerodynamics and with it has brought out some fascinating potential solutions. One of these designs was created by a former pilot named John Smith who embarked on the process of reducing the drag around wingtips of commercial aircraft. He focused on modifying existing wingtip attachments, commonly referred to as winglets, to reduce drag and increase their efficiency. Smith's patented solution [10] introduced a duct through a winglet as a means of further reducing the turbulence generated as air travels from regions of high pressure to low pressure. Smith refers to his invention as an infinity winglet and is referred to as the novel ducted winglet in this work. Smith incorporated Fourth Dimensional Aerospace Technology as a business venture whose aim is to investigate his novel ducted winglet. Preliminary studies indicated that the novel ducted winglet showed promise in reducing drag. To further investigate the properties of Smith's design, a comparison of the novel ducted winglet to a standard, non-ducted winglet as well as an unmodified wingtip across a range of angles of attack is needed. Additionally, Smith also believes that his novel ducted winglet could be suitable for use on wind turbine blades.

As part of the structure of the Ph.D. program, Fourth Dimensional Aerospace Technology were involved in helping decide the methodology and gave guidance on the questions they wanted answered through simulations. At the time, the CEO of Fourth Dimensional Aerospace Technology was Rebecca De Beukelaer and the COO was George Rickman. Their business objective was to determine the feasibility of retrofitting existing wind turbines with Smith's novel ducted winglet and adding the design to wind turbine blades not yet built.

This chapter, carried out in partnership with Fourth Dimensional Aerospace Technology, studies the novel ducted winglet in conditions that can be expected during normal operation of a wind turbine and compares the results to two alternative geometries. A basic introduction to the fundamental aerodynamic forces, governing equations of fluids, and software used to test the designs is given at the start. A detailed description of the novel ducted winglet is then given along with the study methodology. Results from the study are then compared to results presented in Smith's patent [10].

2.2 Background

To better understand the work presented in this and following chapters, some background information has been detailed. This section covers the fundamental aerodynamic principles that are involved in making wind turbines operate, how engineers test their ideas during the design process, and how engineers create design using computers.

2.2.1 Computer Aided Design

Engineering drawings were all done by hand prior to the invention of the computer. While many feats of engineering were designed by hand, the design iteration time was lengthy as any changes in one drawing had to be reflected in subsequent drawings. Testing these designs meant physically building them which took time and cost money. Using computers to make engineering drawings takes a different approach. By adding, removing, and altering basic shapes, or drawing and importing predefined curves, engineers are able to create and reproduce any 3D object, often referred to as computer aided design (CAD). Having a virtual representation of an object brings with it several advantages: better visualisation of the object, the ability to make changes to the object rapidly, seeing how the object fits in an assembly, and of course, simulation. There are several tests that

can be performed on a virtual object to model its behaviour under various conditions. One simple example would be a beam under a point load, or perhaps a distributed load. The shape of the beam can be changed as well to reflect a more realistic example, such as the blade of wind turbine blade. The usefulness of virtual models becomes very apparent.

2.2.2 Fundamental Aerodynamics of Wind Turbines

There are fundamental forces that act on any body in fluid including wind turbine blades in air. These forces are derived from the pressure distribution p and shear stress τ over the body surface. The surface pressure acts perpendicular to the surface of the blade while the shear stress acts parallel to the surface. The forces of interest that result from the surface pressure and shear stress on the surface are lift and drag.

To understand how a wind turbine blade's geometry affects lift and drag, the medium through which it travels through, air, must be understood. The defining properties of air are its density ρ , temperature T , viscosity μ , and pressure p [3]. When air is far away from a wind turbine, it is known as the freestream. The variables associated with the freestream are subscripted with an infinity symbol. Of course, wind turbines require the air to move so its velocity must also be considered. These five properties can be assumed to be constant in the freestream. Using both the freestream density and freestream velocity, the freestream dynamic pressure q_∞ is obtained through the following equation:

$$q_\infty \equiv \frac{1}{2} \rho_\infty V_\infty^2 \quad (2.1)$$

The dynamic pressure is used in calculating lift and drag forces and their respective coefficients. As offshore wind turbines are all located at sea-level, the standardised values at sea-level can be taken for air density, temperature, viscosity, and pressure. From these

properties two dimensionless values can be derived. The first is the Reynolds number Re which relates the inertia forces to viscous forces. The second is the Mach number M which is a ratio of the freestream velocity to the speed of sound. Both these values are important, especially when scaling down designs in a wind tunnel. The Reynolds number is also important when characterising the flow either as laminar or turbulent. Laminar flow is smooth while turbulent flow is often heavily perturbed with swirls and eddies.

Lift is the primary force responsible for rotating the blades of the wind turbine and occurs perpendicular to the freestream velocity. Lift occurs because there is a pressure difference between the top surface of the blade and the bottom surface. The top surface will have a lower pressure compared to the bottom surface, causing the blade to be “sucked” towards the direction of the lift vector. A larger pressure difference between the two surfaces will result in a higher lift. The equation for lift 2.2 is defined by a coefficient of lift C_l , the dynamic pressure and the area of the wing S .

$$L = C_l q_\infty S \quad (2.2)$$

Drag is the force that acts parallel to the direction of the freestream velocity. Unlike lift, drag is influenced by both surface pressure and shear stress. Form drag is how the shape of the blade influences its pressure distribution which affects the surrounding air. Skin friction drag is the measure of the shear stress across the surface as the air particles move over the surface of the body. Both the form drag and skin friction drag make up what is known as parasitic drag. The drag equation 2.3 is similar to the lift equation 2.2 except the coefficient of lift is replaced by the coefficient of drag C_d .

$$D = C_d q_\infty S \quad (2.3)$$

Lastly, drag also occurs increases proportionally to the amount of lift generated. The drag that is created as a result of lift is called induced drag; minimising the induced drag is key to more optimal performance. The equation for induced drag is a function of the coefficient of lift, the Oswald efficiency factor e [3], and the aspect ratio AR of the wing.

$$C_{d_i}(C_l) = \frac{C_l^2}{e\pi AR} \quad (2.4)$$

As a blade rotates relative to the oncoming air, it changes its angle of attack. Increasing the angle of attack will increase the amount of lift generated. This is true up until the flow separates in the boundary layer and causes a phenomenon called stall. Once stall occurs, the blade will experience a loss of lift and a reduction in rotational speed. Drag also slightly increases as the angle of attack increases, though a significant increase in drag occurs once stall happens. Additional details about the onset of stall and the importance of capturing the phenomenon in the boundary layer are discussed further in this chapter.

In real, 3-dimensional flows such as wind turbine blades, air not only passes over the top and bottom surfaces, but also around the tip of the blade. Air from the high pressure surface (bottom) diffuses to the low pressure surface (top) causing the air to rotate around the blade tip. The rotating air forms vortices which can be seen by the naked eye during some conditions. The vortices are a part of the turbulent downstream air referred to as the wake. A large wake can disrupt the performance of downstream wind turbines and minimising the wake is important to ensure that downstream wind turbines can still perform optimally. The generation of vortices and wake are shown in Figure 2.1.




Figure 2.1: A diagram of the formation of wingtip vortices in an idealised flow. These vortices are an indicator of the amount of induced drag caused by the wing.

2.2.3 Induced Drag Reduction Techniques

In this section, various techniques used to reduce the induced drag of an airplane are discussed. The techniques vary in their approaches and their successes. The advantages and technical challenges in realising these techniques are discussed.

Some of the first work conducted in the area of induced drag reduction with the use of winglets was done by engineers at NASA. Engineers conducted tests at Mach 0.78, which is slightly below the cruise speed for modern commercial airplanes. The experimental results showed a 20% decrease in induced drag and an improvement in the lift-to-drag ratio by 9% [11].

One concept for reducing induced drag focuses on removing wingtips entirely. The box wing concept encloses the ends of the wings in a biplane-like design. The box wing concept has been theoretically proven to be more efficient than any other wing system

under minimum induced drag conditions [12]. The trade-off between decreasing the induced drag and increasing weight as well as surface area needs to be considered when implementing a design like the box wing. Structural and safety considerations would also have to be taken into account.

Engineers work on reducing the size and duration of wingtip vortices for the safety of other aircraft. Vortices can remain in the air for several minutes. Guha & Kumar [13] were able to demonstrate that an oscillating wingtip was able to reduce the wake of wingtip vortices. The challenge faced in this application was that the oscillations of the wing became large near the natural frequency of the wing.

Biomimicry has also been employed in the area of induced drag reduction. Birds were observed to have gaps in between their feathers at the end of their wings which produced more, but smaller, vortices [14]. A similar design was used to study the effects of multiple winglets. Narayan & John [15] demonstrated that a 3-prong winglet was more effective at increasing the lift-to-drag ratio than a 2-prong or 4-prong winglet. The experiment was performed using CFD so no physical models were made, however difficulty in manufacturing a 3-prong winglet would need to be addressed.

Another case of reducing induced drag through biomimicry drew its inspiration from humpback whales. Their fins are tuberculated, meaning they have bumps along the leading edge. Bolzon et al. [16] compared the induced drag of a tubercled leading edge to a smooth leading edge at a range of angles of attack. Between 6 and 12 degrees angle of attack, the tubercled wing was able to reduce wingtip vortices. Again, manufacturing challenges would need to be investigated as well as the optimum size and distribution of the bumps along the leading edge.

Other techniques in reducing induced drag involve gaps in the wing. Aldheeb et al. [17] studied the effect of a porous wing tip. They were able to demonstrate that the size of vortices was reduced with pores running perpendicular to the freestream velocity. A reduction in the lift-to-drag ratio was noted however. A similar investigation was taken by Gunasekaran & Gerham [18] who looked at chordwise slots. Their results also showed a reduction in vortex strength while also showing a reduction in the lift-to-drag ratio. Both studies had made physical models, but the manufacturing challenges in creating these parts would need further investigation. Moreover, wings are usually filled internally with fuel, wiring and hydraulic systems which are necessary for flight.

Another approach to dealing with wingtip vortices involves disrupting the vortex with another vortex. This method was employed by Heyes & Smith [19] who conducted experiments using 6 different shapes to generate contra-rotating vortices. A reduction in wingtip vortices was achieved at the cost of decreased lift.

2.2.4 Aerodynamic Testing Methods

Aerodynamic testing was limited to wind tunnels for much of the 20th century. First used by the Wright brothers, wind tunnels have grown in size and increased wind speeds. There are some wind tunnels which are capable of testing full-scale models, though they require large amounts of energy to run. Instead, most models are scaled down to fit inside more resource-efficient wind tunnels. Changing the size of the model clearly would have an effect on the aerodynamic properties and the results gathered from the test. The solution to this dilemma is to match Re between the expected conditions and the conditions in the wind tunnel. Practically, this means increasing the velocity inside the wind tunnel proportionally to the decrease in geometry size. Wind tunnels have the advantage that once a model is built and set up in the wind tunnel, it does not take much time at all to

gather data for a range of conditions. These conditions include the air speed and angle of attack. The drawbacks of using a wind tunnel is the time it takes to manufacture models and the initial set up time. Any changes made to the geometry also take a considerable amount of time. Additionally, pressure data can only be collected from pre-determined locations on the geometry during testing.

CFD aims to solve many of the drawbacks of traditional wind tunnel testing. The geometry and simulation are all done virtually, so there is no need to build the geometry to be tested. This reduces the turnaround time while iterating through the design process. The exact conditions can also be set for the simulation, meaning there are no adjustments needed due to scaling. With enough resources, several simulations can be run simultaneously which can lead to a reduction in time for the overall design process. Of course, CFD also has some drawbacks. Generating a mesh is notoriously difficult which can lead to long delays. Simulations can take a significant amount of time to run.

2.2.5 Computational Fluid Dynamics

The following section gives a brief overview of the necessary components which are incorporated into the CFD solver used in this work, FLITE3D. Pertinent information provided here is extracted from the Flite System Theoretical Manual [20].

FLITE3D is a Reynolds-Averaged Navier-Stokes based CFD solver which has been deployed on numerous commercial applications such the Bloodhound Land Speed Record car [21]. Contained within FLITE3D [20] is a surface mesh program, a volume mesh program, a pre-processor, and the solver. The surface and volume mesh programs create high-quality, unstructured triangular and tetrahedral elements. In the boundary layer, the advancing front method is used to generate the mesh while the Delaunay method is used

in the rest of the domain. The pre-processor divides the domain into smaller portions allowing the solver to take advantage of parallel computing. At this stage, the boundary conditions are applied to the domain and the geometry. These include initial flow characteristics such as viscosity, compressibility, rotation, density, and velocity. Once the domain has been divided into subdomains and the boundary conditions have been applied, the solver calculates the aerodynamic properties within each subdomain and communicates the results from the boundary of adjacent subdomains to ensure continuity. The solver will continue to run until either the solution field has sufficiently converged, an iteration limit has been reached, or a time limit has been reached.

Meshing

A process known as meshing can begin once the computational domain has been set along with defined curves and surfaces [22]. In the case of this work, curves and surfaces were contained within the STEP file used for the 3D geometry. This process discretises the domain and the surfaces in a manner that allows aerodynamic forces to be calculated at each discretised point. In 2D, neighbouring points are connected in sets of triangles called elements; the 3D elements are called tetrahedra. The mesh used in FLITE3D is an unstructured mesh which means that the elements in the mesh are differently shaped throughout the domain. The advantage of using an unstructured mesh is that any arbitrary geometry can be discretised as long as the surfaces are congruent to the definition in the previous section, i.e. the surfaces are bounded by curves and curves lie on the intersection of two surfaces. Structured meshes are also used in other applications though the mesh generation process tends to be time-consuming when compared to unstructured meshes. The downside to unstructured meshes is an increase in computational resources needed compared to structured meshes.

The initial mesh generated by the program is known as a background mesh and is used for interpolation. This starting mesh defines the initial position and global size for the elements. From there, changes can be made to the mesh such as the sizing of elements. Smaller elements are needed in parts of the domain which are likely to encounter bigger changes in aerodynamic properties, such as the leading edge of a wind turbine blade. Regions far away from the geometry where changes to the aerodynamic properties of the air are minimal can be discretised by larger elements. The sizing of the elements can be controlled by sources: point, line, and triangular. The density δ of these sources are defined by

$$\delta(x) = \begin{cases} \delta_1 & \text{if } x \leq x_c \\ \delta_1 e^{\left| \frac{x-x_c}{D-x_c} \right| \log 2} & \text{if } x \geq x_c \end{cases} \quad (2.5)$$

where δ_1 , D , and x_c are variables set by the user. The user is able to set a radius r_1 around the source where the source is equal to the user-prescribed density δ . Another radius r_2 , which must be bigger than r_1 , sets the boundary at which the background mesh starts. Between r_1 and r_2 exists a region where the size of the elements grows from δ to the global element size.

Curves and surfaces are discretised according to the background mesh. Points are added along the curves and surfaces and then connected with straight lines. Evidently, this leads to geometric error. The size of the geometric error can be reduced by adding sources near curves in the geometry with a small radius. Recent advances reduce the geometric error to zero through the use of non-uniform rational b-splines [23] though not introduced in this work.

The advancing front method is used to generate the surface mesh. In this method, elements

are added to the mesh one at a time by a set of criteria [22]. Once the surface mesh has been generated, enhancements are made to increase the quality of the mesh. A good quality mesh has elements with approximately equilateral sides and avoids small angles. Two mesh enhancement methods are used. The first is diagonal swapping which reduces the length of a shared base between two elements. The second is mesh smoothing which helps by iteratively changing the sides of an element until they are approximately equilateral. The criteria for mesh quality are specified by the user. The volume mesh is generated after enhancements have been made to the surface mesh and it satisfies the user's criteria. The Delaunay method is employed for volume mesh generation in FLITE3D using the Bowyer algorithm [24]. Edge and boundary face recovery methods are used to recover missing edges of faces.

Governing Equations, Turbulence, and Boundary Conditions

The governing equations used in solving the flow are the Navier-Stokes equations, and were developed independently in the 19th century by physicists Claude-Louis Navier and George Gabriel Stokes [3]. There are 3 equation types that make up the set of Navier-Stokes equations: the continuity equation, the momentum equations, and the energy equation.

The continuity equation represents the mass flow rate through any arbitrary volume. Its finite space form is

$$\frac{\partial}{\partial t} \iiint \rho dV + \iint \rho \mathbf{V} \cdot d\mathbf{S} = 0 \quad (2.6)$$

where $\frac{\partial}{\partial t} \iiint \rho dV$ is the time rate increase of mass inside the volume and $\iint \rho \mathbf{V} \cdot d\mathbf{S}$ is the net mass flow out of the control volume through the surface of the volume. The continuity

equation can also be represented in its partial differential form as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{V}) = 0 \quad (2.7)$$

which relates the variables in the flow to points within the flow.

The momentum equation relates force to a change in momentum over time, most famously described in Newton's 2nd law. In the case of fluid flow, consider the body forces and the surface forces. In its finite volume form, the momentum equation is

$$\frac{\partial}{\partial t} \iiint \rho \mathbf{V} d\mathcal{V} + \iint (\rho \mathbf{V} \cdot \mathbf{dS}) \mathbf{V} = - \iint p \mathbf{dS} + \iiint \rho \mathbf{f} d\mathcal{V} + F_{viscous} \quad (2.8)$$

where $\frac{\partial}{\partial t} \iiint \rho \mathbf{V} d\mathcal{V}$ is the rate of change in flow fluctuations over time and $\iint (\rho \mathbf{V} \cdot \mathbf{dS}) \mathbf{V}$ is net flow of momentum out of the control volume; where $- \iint p \mathbf{dS} + \iiint \rho \mathbf{f} d\mathcal{V} + F_{viscous}$ are the body and surface forces, respectively. Like the continuity equation, the momentum equation can also be expressed in its partial differential form as

$$\frac{\partial(\rho u)}{\partial t} + \nabla \cdot (\rho u \mathbf{V}) = - \frac{\partial p}{\partial x} + \rho f_x + (\mathcal{F}_x)_{viscous} \quad (2.9)$$

Momentum is conserved in all three cartesian directions, yielding in the following:

$$\frac{\partial(\rho v)}{\partial t} + \nabla \cdot (\rho v \mathbf{V}) = - \frac{\partial p}{\partial y} + \rho f_y + (\mathcal{F}_y)_{viscous} \quad (2.10)$$

$$\frac{\partial(\rho w)}{\partial t} + \nabla \cdot (\rho w \mathbf{V}) = - \frac{\partial p}{\partial z} + \rho f_z + (\mathcal{F}_z)_{viscous} \quad (2.11)$$

Finally, the energy equation relates the rate of change to the heat added and work done on a fluid inside a control volume to the rate of change of energy of the fluid as it flows

through the control volume. Hence, the energy equation in its finite volume form becomes:

$$\iiint \dot{q} \rho d\mathcal{V} + \dot{Q}_{viscous} - \iint p \mathbf{V} \cdot d\mathbf{S} + \iiint \rho (\mathbf{f} \cdot \mathbf{V}) d\mathcal{V} + \dot{W}_{viscous} = \frac{\partial}{\partial t} \iiint \rho \left(e + \frac{V^2}{2} \right) \mathbf{V} \cdot d\mathbf{S} \quad (2.12)$$

where $\iiint \dot{q} \rho d\mathcal{V} + \dot{Q}_{viscous}$ is the rate of heat added, $\iint p \mathbf{V} \cdot d\mathbf{S} + \iiint \rho (\mathbf{f} \cdot \mathbf{V}) d\mathcal{V} + \dot{W}_{viscous}$ is the rate of work done on the fluid, and $\frac{\partial}{\partial t} \iiint \rho \left(e + \frac{V^2}{2} \right) \mathbf{V} \cdot d\mathbf{S}$ is the net rate of flow of total energy across the surface and the change of total energy inside the volume because of transient variations of flow-field variables.

Applying the same principles to the energy equation as the continuity and momentum equations, the partial differential form of the energy equation becomes:

$$\frac{\partial}{\partial t} \left[\rho \left(e + \frac{V^2}{2} \right) \right] + \nabla \cdot \left[\rho \left(e + \frac{V^2}{2} \right) \mathbf{V} \right] = \rho \dot{q} - \nabla \cdot (p \mathbf{V}) + \rho (\mathbf{f} \cdot \mathbf{V}) + \dot{Q}'_{viscous} + \dot{W}'_{viscous} \quad (2.13)$$

The finite space forms of the equations are useful as they provide an analytical representation of the aerodynamic properties of the flow for a given control volume. However, they provide no information for any arbitrary point within the domain of the control volume. This is a problem because aerodynamicists are concerned with what happens within the domain and how aerodynamic properties change as air travels through the domain and how those properties change over time. Therefore, the partial differential forms of the Navier-Stokes equations become useful as they do provide information at discrete points within the domain.

Turbulence in this work is approximated using the Spalart-Allmaras turbulence model [25], which relates the kinematic turbulent viscosity to an eddy viscosity variable $\tilde{\nu}$ by

$$\nu = \frac{\mu_t}{\rho} = \tilde{\nu} f_{v1} \quad (2.14)$$

where

$$f_{v1} = \frac{\chi^3}{\chi^3 + c_{v1}^3} \quad (2.15)$$

with c_{v1} being a constant and, subsequently, where

$$\chi = \frac{\tilde{\nu}}{\nu} \quad (2.16)$$

The eddy viscosity is measured using

$$\frac{\partial \tilde{\nu}}{\partial t} + \tilde{u}_j \frac{\partial \tilde{\nu}}{\partial x_j} = c_{b1} [1 - f_{t2}] \tilde{S} \tilde{\nu} + \frac{1}{\sigma} \left[\frac{\partial}{\partial x_j} ((\nu + \tilde{\nu}) \frac{\partial \tilde{\nu}}{\partial x_j}) + c_{b2} \frac{\partial \tilde{\nu}}{\partial x_j} \frac{\partial \tilde{\nu}}{\partial x_j} \right] - \left[c_{w1} f_w - \frac{c_{b1}}{\kappa^2} f_{t2} \right] \left[\frac{\tilde{\nu}}{d} \right]^2 + f_{t1} \Delta \tilde{u}^2 \quad (2.17)$$

with the respective terms in the equation representing the following:

Table 2.1: Constants used in the Spalart-Allmaras turbulence model.

| | |
|--|------------------|
| $c_{b1} = 0.1355$ | $c_{b2} = 0.622$ |
| $\omega = 2/3$ | $\kappa = 0.41$ |
| $c_{w1} = c_{b1}/\kappa^2 + (1 + c_{b2})/\omega$ | $c_{w2} = 0.3$ |
| $c_{w3} = 2$ | $c_{v1} = 7.1$ |
| $c_{t1} = 1$ | $c_{t2} = 2$ |
| $c_{t3} = 1.1$ | $c_{t4} = 2$ |

$$\tilde{S} = \tilde{\omega} + \frac{\nu}{\kappa^2 d^2} f_{v2} \quad (2.18)$$

$$f_{v2} = 1 - \frac{\chi}{1 + \chi f_{v1}} \quad (2.19)$$

$$f_w = g \left[\frac{1 + c_{w3}^6}{g^6 + c_{w3}^6} \right]^{\frac{1}{6}} \quad (2.20)$$

$$g = r + c_{w2}(r^6 - r) \quad (2.21)$$

$$r = \min\left(\frac{\tilde{\nu}}{\tilde{S}\kappa^2 d^2}, 10\right) \quad (2.22)$$

$$f_{t2} = c_{t3} \exp(-c_{t4} \chi^2) \quad (2.23)$$

$$f_{t1} = c_{t1} g_t \exp\left(-c_{t2} \frac{\tilde{\omega}_t^2}{\Delta \tilde{u}^2} [d^2 + g_t^2 d_t^2]\right) \quad (2.24)$$

$$g_t = \min\left(0.1, \frac{\Delta \tilde{u}}{\tilde{\omega}_t \Delta x}\right) \quad (2.25)$$

The values of the constant listed in the equations above is provided in Table 2.1.

Two additional terms not included in Table 2.1 is the distance d from the nearest wall to an arbitrary point in the domain and $\tilde{\omega}$ which denotes the magnitude of the vorticity and is given as

$$\tilde{\omega} = \left(\left(\frac{\partial u_3}{\partial x_2} - \frac{\partial u_2}{\partial x_3} \right)^2 + \left(\frac{\partial u_1}{\partial x_3} - \frac{\partial u_3}{\partial x_1} \right)^2 + \left(\frac{\partial u_2}{\partial x_1} - \frac{\partial u_1}{\partial x_2} \right)^2 \right)^{\frac{1}{2}} \quad (2.26)$$

Finally, boundary conditions can be set on the walls, at the inlet, the outlet, and on the surfaces of the geometry. The walls of a viscous domain set the fluid velocity as the wall velocity, and given that the walls are stationary, this means that the velocity of the air at the domain walls is zero. As there zero-only velocity components at the walls, the velocity flux is also zero. The domain walls are also isentropic, meaning that the temperature flux is set to zero. The zero velocity and zero velocity flux also hold on the surfaces. Five flow conditions are set at the inflow and four are set at the outflow walls. The Spalart-Allmaras turbulence model sets $\tilde{\nu}$ to ten percent of the laminar air viscosity. A symmetry wall is also used for some simulations in this work. The symmetry wall sets the normal velocity component with respect to the boundary to zero, and also sets the derivatives of density and total energy to zero. The normal velocity component relationship is represented as

$$u_j n_j^s = 0 \quad (2.27)$$

where n_j^s is the symmetry boundary normal and

$$\frac{\partial \rho}{\partial n^s} = \frac{\partial \epsilon}{\partial n^s} = 0 \quad (2.28)$$

represents the derivatives of density ρ and total energy ϵ .

Viscous Scheme Discretisation

The viscous scheme incorporates all terms of the Navier-Stokes equations and is therefore the most computationally expensive scheme to run. Discretising the viscous scheme allows the Navier-Stokes equations to be evaluated at each point within the domain. To do so, start with

$$\int_{\Omega_I} \frac{\partial u_i}{\partial x_j} d\mathbf{x} = \int_{\partial\Omega_I} u_i n_j d\mathbf{x} \quad (2.29)$$

which calculates the derivatives of the velocity at each point, and is

$$\frac{\partial u_i}{\partial x_j} \Big|_I \approx \partial_j^h u_i^I \equiv \frac{1}{V_I} \left[\sum_{j \in \Lambda_I} \frac{C_j^{IJ}}{2} (u_i^I + u_i^J) + \sum_{j \in \Lambda_I^B} D_j^{IJ} u_i^I \right] \quad (2.30)$$

in its discrete form. The boundary integral of the gradient

$$\int_{\partial\Omega_I} \frac{\partial u_i}{\partial x_j} n_k d\mathbf{x} \approx \sum_{j \in \Lambda_I} \frac{C_k^{IJ}}{2} (\partial_j^h u_i^I + \partial_j^h u_i^J) + \sum_{j \in \Lambda_I^B} D_k^{IJ} \partial_j^h u_i^I \quad (2.31)$$

is also evaluated. The evaluation of the derivatives results in a second order accurate scheme. A more compact way of discretising the scheme is possible by performing the discretisation at the edge and taking the derivative at the edge midpoint

$$\partial_j^h u_i^{IJ} = \frac{1}{2} (\partial_j^h u_i^I + \partial_j^h u_i^J) \quad (2.32)$$

into its tangential and normal components

$$\partial_j^h u_i^{IJ} = \partial_j^h u_i^{IJ}|_t + \partial_j^h u_i^{IJ}|_n \quad (2.33)$$

where

$$\partial_j^h u_i^{IJ}|_t = r_k^{IJ} \partial_k^h u_i^{IJ} r_j^{IJ} \quad (2.34)$$

is the component of the same direction as the edge and the normal component is

$$\partial_j^h u_i^{IJ}|_n = \partial_j^h u_i^{IJ} - r_k^{IJ} \partial_k^h u_i^{IJ} r_j^{IJ} \quad (2.35)$$

where the unit vector of the edge is

$$r_j^{IJ} = \frac{1}{l_{IJ}} (x_j^J - x_j^I) \quad (2.36)$$

The tangential component is then evaluated using the finite difference method

$$\hat{\partial}_j^h u_i^{IJ}|_t = \frac{1}{l_{IJ}} (x_k^J - x_k^I) r_j^{IJ} \quad (2.37)$$

resulting in a more compact evaluation, giving

$$\int_{\partial\Omega_I} \frac{\partial u_i}{\partial x_j} n_k d\mathbf{x} \approx \sum_{J \in \Lambda_I^A} \frac{C_k^{IJ}}{2} (\partial_j^h u_i^I|_n + \hat{\partial}_j^h u_i^I|_t) + \sum_{J \in \Lambda_I^B} D_k^{IJ} \partial_j^h u_i^I \quad (2.38)$$

as the final equation to evaluate the viscous fluxes.

The viscous terms in the Navier-Stokes equations can be approximated to

$$\int_{\partial\Omega_I} \mathbf{G}_{ij} n_j d\mathbf{x} \approx \sum_{j \in \Lambda_I^A} C_j^{IJ} \mathbf{G}_{ij}^{IJ} + \sum_{j \in \Lambda_I^B} D_j^{IJ} \mathbf{G}_{ij}^I \quad (2.39)$$

where \mathbf{G} is a matrix set as

$$\mathbf{G}_j^{IJ, I} = \begin{bmatrix} 0 \\ \tau_{1j}^{IJ, I} \\ \tau_{2j}^{IJ, I} \\ \tau_{3j}^{IJ, I} \\ u_k^{IJ, I} \tau_{kj}^{IJ, I} - q_j^{IJ, I} \end{bmatrix} \quad (2.40)$$

The deviatoric stress tensor τ_{ij}^{IJ} at the edge midpoint is

$$\tau_{ij}^{IJ} = \frac{\mu_I + \mu_J}{2} \left[-\frac{2}{3} (\partial_k^h u_k^{IJ}|_n + \hat{\partial}_k^h u_k^{IJ}|_t) \delta_{ij} + (\partial_j^h u_i^{IJ}|_n + \hat{\partial}_j^h u_i^{IJ}|_t) + (\partial_i^h u_j^{IJ}|_n + \hat{\partial}_i^h u_j^{IJ}|_t) \right] \quad (2.41)$$

with the energy dissipation term as

$$u_k^{IJ} \tau_{kj}^{IJ} = \frac{u_k^I + u_k^J}{2} \tau_{kj}^{IJ} \quad (2.42)$$

and with the heat flux as

$$q_j^{IJ} = -\frac{k_I + k_J}{2} (\partial_j^h T_{IJ}|_n + \hat{\partial}_j^h T_{IJ}|_t) \quad (2.43)$$

At the boundary, the deviatoric stress tensor is

$$\tau_{ij}^I = \mu_I \left(-\frac{2}{3} \partial_k^h u_k^I \delta_{ij} + \partial_j^h u_i^I + \partial_i^h u_j^I \right) \quad (2.44)$$

with the accompanying heat flux as

$$q_j^I = -k_I \partial_j^h T_I \quad (2.45)$$

The Spalart-Allmaras turbulence model must also be converted into its discrete form [25] using a first order upwind method for enhanced stability, set as

$$\int_{\partial\Omega_I} \tilde{u}_j \tilde{v}_j n_j d\mathbf{x} \approx \left[\sum_{J \in \Lambda_I} \frac{C_k^{IJ}}{2} (\tilde{u}_j^I \tilde{v}_J) - |C_j^{IJ} \frac{\tilde{u}_j^I \tilde{v}_J - \tilde{u}_j^I \tilde{v}_I}{\tilde{v}_I - \tilde{v}_J}| (\tilde{v}_I - \tilde{v}_J) \right] + \sum_{J \in \Lambda_I^B} D_j^{IJ} \tilde{u}_j^I \tilde{v}_I \quad (2.46)$$

which reduces the the scheme to first order.

Solving and Parallelisation

Solving involves the evaluation of the discretised domain divided into distinct time steps. The boundary conditions set the initial values for the elements and nodes in the domain. From that point forward, element and nodal values are updated as they are evaluated at each time step.

Parallelisation allows the domain to be separated into sub-domains where each sub-domain can then be evaluated on an individual CPU. This greatly increases the speed at which solutions are calculated though introduces a small cost when information from neighbouring sub-domains must be shared with each other. Each sub-domain consists of a distinct set of elements and nodes on the edge of the sub-domain are shared with neighbouring sub-domains. This introduces a layer of computational complexity as a selection of nodes are duplicated. Nodes that sit on the edge of the sub-domain are updated based on corresponding element information. The updated nodal values are then communicated to neighbouring sub-domain nodes. Once all edge nodes have been updated, the values for the internal nodes are then updated. This process allows the computation of elements and nodes as well as communication between the sub-domains to occur concurrently.

2.3 Methods and Materials

The novel winglet was patented in 2017 and describes fluid flow through a duct which runs through the winglet. The duct's inlet is on the lower surface of the blade on the high pressure side and the outlet is at the tip of the blade. While the patent describes the outlet as being on the upper surface on the low pressure side the schematics show that the outlet has its own surface. The patent describes two regions of air within the duct, each rotating in opposite directions. The patent goes on to say that these counter rotating flows inhibit the formation of vortices as air goes from the lower surface to the upper surface.

2.3.1 Novel Ducted Winglet

Winglets are a wingtip, or in this case blade tip, modification that generally extends perpendicular to the blade. They have been tested in experimental setups and proven




Figure 2.2: A cross-section of the patented ducted winglet.

in commercial applications to reduce induced drag and the size of the wake, as set out in the last section. As with any engineering challenge, solutions have some drawbacks that must be considered. There are three main points to consider in the case of winglets. The first is that the addition of a winglet increases the total surface area of a wing. This leads to an increase in skin friction drag. Another drawback of winglets has to do with the shape itself. The addition of a complex geometry to the blade increases profile drag. Lastly, the winglet adds weight to the blade and since it is located at the tip of the blade, increases the bending moment at the root of the blade. Practically, this means that the rest of the blade must be strengthened with additional internal support structures which adds more weight to the blade. As an aside, both the winglet and any necessary support structure would also add to the overall cost of the blade. For the purposes of this work, only the aerodynamic properties are evaluated. The primary reason for considering only the aerodynamic forces is that it reduces the simulation costs as the CFD does not have to be coupled with a structural solver. The blades in the simulations are assumed to be rigid and do not deflect due to aerodynamic loads.

Wind turbines and their blades provide a unique circumstance which may result in a bigger benefit when compared to airplane wings. The two primary reasons why wind turbine blades are better suited for winglets are:

1. Winglets are more effective at lower speeds because induced drag makes up a larger proportion of total drag.
2. Winglets on an airplane wing make up a larger portion of the surface area compared to winglets on a wind turbine blade. This means that the increase in skin friction and profile drag of winglets on wind turbine blades are lower than on an airplane wing. Additionally, the proportion of increased weight is smaller for a wind turbine

blade than it is on a wing.

While many studies conducted on wind turbine blades with winglets, none have yet considered a ducted design. Smith [10] claims that their novel design is able to reduce induced drag and decrease wake by channelling a portion of freestream air through a duct and ejecting the air at the tip of the winglet where it rejoins the freestream. The novel ducted winglet is originally designed for use on airplane wings, but the patent extends to all application areas including wind turbine blades. The duct in the winglet is meant to channel air and artificially extend the length of the blade and therefore increasing the aspect ratio. An increase in AR would lead to a lower coefficient of induced drag C_{D_i} , as can be seen in Equation 2.4. As previously mentioned, the addition of the winglet leads to increases in other forms of drag. To have a net decrease in drag, the novel winglet must decrease induced drag more than it increases other forms of drag. The inventor claims that the novel winglet manages to do this successfully the following sections investigate this claim along with others present in the patent.

There are 7 assertions in the patent that are going to be explored in this chapter. The claims have been paraphrased and shortened for brevity and clarity. The patent claims that the design has:

Assertion 1 Counter rotating airflow which occurs in the duct;

Assertion 2 Counter rotating airflow which occurs at the outlet of the duct;

Assertion 3 A region of high pressure in the duct which directs the air through the duct;

Assertion 4 A reduction of velocity of air as it enters the duct;

Assertion 5 An aerofoil-shaped region of air inside the duct which influence the rotation

of air;

Assertion 6 A 15% improvement in the lift-to-drag ratio;

Assertion 7 A non-symmetric velocity and pressure profile within the duct.

At the request of the inventor, the efficacy of the winglet was to be tested on wind turbine blades. The National Renewable Energy Laboratory 15 megawatt wind turbine was chosen as data for the blade was publicly available. The blade geometry was converted to 3D aerofoil coordinates using Matlab. From there, the coordinates were exported to SolidWorks [26], a commercially available CAD package, to create a full 3D model of the blade. The 3D model of the blade was exported as a STEP file.

2.3.2 Parameterisation

Parameterisation of the novel winglet followed a similar process. 5 elliptical profiles were used to create the inner part of the duct while an additional profile was located outside of the geometry to create a smooth inlet. The additional profile ensured that there would be no jagged edges near the inlet which would have increased drag. A further 11 profiles defined the outer portion of the duct. These outer profiles were symmetric NACA 0012 aerofoils; their symmetry means that the winglet will not create additional lift while minimizing its overall drag. The inner profiles were constructed of 3 parameters: the ellipse ratio, chord length, and position relative to the leading edge. The ellipse ratio defines the eccentricity of the ellipse and is given by

$$\frac{(x - h)^2}{a^2} + \frac{(y - k)^2}{b^2} = 1 \quad (2.47)$$




Figure 2.3: The parameters which define an ellipse.

with parameters shown in Figure 2.3. The chord length was estimated from the patent as was the position of each profile relative to the leading edge. A total of 18 parameters define the winglet duct. The initial design simulated in this work was approved by Smith.

2.3.3 Simulation Setup

Lift and drag values for the novel ducted winglet and the other two geometries were gathered from CFD simulations. The simulations will also provide the data necessary to evaluate the induced drag values of all three geometries during post-processing and will allow the patent assertions to be compared to simulation results. All simulations were run using Impact, a local high performance cluster. Each node on the cluster consists of 2x20 Intel(R) Xeon(R) Gold 6230/2.10 GHz CPU cores. 4 nodes were used for each simulation for a total of 160 cores. All simulations were run for 72 hours except those with the finest




Figure 2.4: The inner and outer profiles which define the shape of the recreated novel ducted winglet. The outer profiles are in black and the inner profiles are in blue.

meshes, which were run for 144 hours. The finest meshes refers to the meshes with the highest element count in Figure 2.10. Lift and drag values were calculated at each iteration with the final value represented as the mean of the last 25% of iterations. Additionally, a measure of the differences between the current and previous iterations were output called the residual. The units of the residual R correspond to the order of magnitude in the differences, such that the difference can be represented as 10^{-R} . The changes to the lift, drag, and residual values as the simulation progresses are shown in Figure 2.5. The relevant parameters needed to set up the simulation are listed in Table 2.2 and Table 2.3.

Despite the freestream velocity being in the incompressible regime, a compressible solver was used instead of an incompressible solver because the behaviour inside the duct could not be predicted prior to the start of the simulations. Evidence to support this decision a posteriori are seen in Figure 2.14 where the velocity of the air crosses the threshold

Table 2.2: Computational parameters used to set up the simulations. BC refers to boundary conditions.

| Variable | Value |
|------------------------|-----------|
| Mach number M | 0.279 |
| Reynolds number Re | 6,505,030 |
| Domain height | 75m |
| Domain width | 250m |
| Domain length | 250m |
| External walls BC | 95m/s |
| Turbulent viscosity BC | 0 |

Table 2.3: Computational parameters used to set up the simulations, specific to FLITE3D.

| Variable | Value |
|---------------------------------|-------|
| MG iterations | -3 |
| CFL number | 1.0 |
| Grids | 1 |
| Flux type | 4 |
| Viscosity scheme | 1 |
| HLLC Flux Order | 2 |
| HLLC Clipping | 1.0 |
| Gradient Method | 2 |
| Dissipation scheme | 2 |
| Coarse grid dissipation scheme | 2 |
| Second order dissipation factor | 0.5 |
| Fourth order dissipation factor | 0.4 |
| Coarse grid dissipation factor | 0.75 |
| Turbulence model | 1 |
| Number of turbulence grids | 1 |
| Turbulent CFL number | 0.5 |
| Trigger radius | 0.1 |
| Number of trigger steps | 10 |
| Turbulence trigger value | 25 |




Figure 2.5: The lift, drag, and residual values as the simulation iteration count increases. The residual values, in blue, are plotted on the left. The lift and drag values, in orange, are plotted on the right.

where compressibility effects can no longer be neglected. The simulations were run as steady simulations as opposed to unsteady simulations. Steady simulations converge faster and require less memory. This proved to be advantageous given the time it took to run a simulation as well as keeping to the memory limit given the number of simulations that needed to be run. The turbulence model was chosen for its general stability and good convergence. The model also provides dimensionality reduction which lowers the complexity and in turn lowers the overall simulation time.

The efficacy of the novel winglet needed to be compared to a benchmark. To measure the effect of the winglet and the duct, three different geometries were created: an unmodified blade, a winglet without a duct, and the novel ducted design. Each blade was tested at 8 angles of attack. Time-averaged lift and drag were used as the solver was pseudo-steady-state. The lift-to-drag ratio was calculated from these values.




Figure 2.6: CFD wind turbine domain used for all three blade configurations; reference blade tip shown.

The domain used to test the three geometries is shown in Figure 2.6. The blade sits toward the front of the rectangular domain so that downstream effects can be captured and sits far enough from the surrounding walls so that interference effects are minimised.

A method known as Trefftz plane analysis [27] was employed to measure the induced drag. The Trefftz plane analysis is a wake integration method that uses the magnitude of the downstream vorticity to give a relative indication of the induced drag. The measurement plane can be located at any arbitrary distance downstream from the blade. A distance of 1 *m* downstream was chosen for the analysis presented here. The exact value of induced drag is of minimal importance so long as the ranking with respect to the induced drag is consistent. The equation

$$D = \iint v^2 + w^2 dS \quad (2.48)$$

is the surface integral of the horizontal and vertical velocity components with respect




Figure 2.7: The cross-section of the domain where the Trefftz plane is situated.

to the direction of the freestream velocity. The entire blade will have an impact on the downstream vorticity, the difference between the different geometries should come down to the modifications made at the blade tip. The use of the Trefftz plane analysis is conducted during the post-processing stage. Results files are generated after the solver has finished and are exported to Ansys Ensign. Data from the Trefftz plane are exported then to Matlab to generate figures and calculate the induced drag.

The blade was simulated as a static entity in the same way as an airplane wing, where the freestream velocity meets the blade at the same speed. This is not the case in real life where the tip of the blade experiences a higher speed than the root due to rotation. This simplification allows the blade to be modelled while keeping computational costs down. As with the Trefftz plane analysis, because the geometry is the same for each blade except for the tips, the differences in lift and drag will be caused solely by the differences

in geometry at the tip. The entire blade was simulated rather than just a portion as it was unknown prior to the simulations how much of the blade was needed before making a significant difference to the results.

2.3.4 Mesh Convergence Study

A mesh convergence study is used to determine the appropriate mesh size for a given scenario. Complexities in the geometry and freestream conditions play a part in setting the mesh size. The objective of the mesh convergence study is to find the minimum amount of elements and nodes needed to accurately capture the aerodynamic properties within the domain. A slice of the domain is seen in Figure 2.9. Minimising the element and node count also minimises computational costs and time.

A mesh convergence study was performed on the first 5 degrees of angle of attack. Since the interest of the study is at the tip of the blade, refinements to the global mesh size were made to the blade. Further refinements were also made to the tip of the blade to properly capture the aerodynamic phenomena – especially crucial for the novel ducted winglet. Parameter values used to define the meshes at various levels of refinement are available in Table 2.4 with a sample of the meshes shown in Figure 2.8.

Due to computational resource constraints, a full convergence study could not be completed. However, the lift-to-drag ratio ranking of each angle of attack remains consistent as the element count increases. This brings confidence to the argument that the different blade tip geometries can be ranked in the correct order using CFD. Moreover, the difference in lift-to-drag ratio between one angle and the next remains approximately constant as element counts increase. Again, this supports the argument that the results gathered from the simulations will be of the correct order of magnitude for lift and drag values.



Figure 2.8: The meshes of the three different geometries that were tested: A) an unmodified blade tip, B) a standard winglet, and C) the novel ducted winglet.




Figure 2.9: A cross section of the domain which shows the refinement of the elements as they approach the wind turbine blade.

Table 2.4: Mesh parameter values used for the mesh convergence study. All values were used across all three geometries to minimise the effect of varying element counts on the simulation results. The blade line source was the same length as the blade and the winglet line source was the same height as the winglet. Both lines went through the middle of their respective geometries.

| Variable | Value |
|-----------------------------------|--------------------------|
| Background spacing | 2.9 |
| Winglet line sources element size | [0.04, 0.03, 0.02, 0.13] |
| Blade line sources element size | [0.09, 0.08, 0.07, 0.06] |
| Boundary layers | 30 |
| First layer height | 5.44×10^{-6} |
| Boundary layer growth rate | 1.2 |
| Y+ | 1.1 |
| Boundary layer element type | Hybrid |




Figure 2.10: The results from the mesh convergence study conducted on the novel ducted winglet at 5 different angles of attack.

2.4 Results

The results from the CFD simulations are presented in this section and subdivided by the aerodynamic properties that are being test: lift, drag, the lift-to-drag ratio, and the induced drag. The novel ducted winglet, the non-ducted winglet, and the reference geometries are compared across the 8 angles of attack.

2.4.1 Lift

Across all 8 angles of attack, the unmodified blade had the highest lift. The novel ducted winglet had the lowest lift generated across all angles of attack. The unmodified winglet generated more lift than the novel ducted winglet but less lift than the reference blade. One possible explanation for this is that the unmodified blade has a slightly larger lifting




Figure 2.11: Non-dimensionalised lift values for the three geometries.

Table 2.5: Non-dimensionalised lift for the novel ducted winglet, the non-ducted winglet, and the unmodified reference blade. Values in bold indicate the highest values.

| Angle | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Novel Ducted | 0.5773 | 0.6580 | 0.7387 | 0.8145 | 0.8923 | 0.9674 | 1.0394 | 1.1085 |
| Non-Ducted | 0.5787 | 0.6612 | 0.7413 | 0.8203 | 0.8963 | 0.9720 | 1.0441 | 1.1133 |
| Reference | 0.5859 | 0.6682 | 0.7492 | 0.8293 | 0.9075 | 0.9837 | 1.0576 | 1.1290 |

surface as the blade is able to stay horizontal for the full length of the blade. This is in contrast to the unmodified winglet and the novel ducted winglet which have curved blade tips and no longer act as lifting surfaces.

While both the geometries with a winglet generated less lift than the reference blade, they still maintained lift to within 1%. This is crucial to generating power and ensuring the feasibility of the design.

2.4.2 Drag

Crucially, both the novel ducted winglet and the unmodified winglet created less drag than the reference blade. Though the results are comparable, the non-ducted winglet did create the least amount of drag across all angles of attack. The decrease in drag from the reference blade to the two winglet geometries is most likely caused by improved airflow around the tip of the blade. Since the two winglet designs are identical other than the duct, the increase in drag between the two can only be attributed to the duct.

The presence of the duct decreases the freestream velocity as it enters the duct. The duct is not located parallel to the freestream velocity which means that the air must change direction quite significantly over a short distance to enter the duct. The result of this is an increase in pressure at the trailing edge of the duct as seen in Figure 2.20. An increase in pressure is directly correlated with a decrease in velocity. The change in velocity of the air as it enters the duct is also confirmed in Figure 2.21 which shows colour-coded streamlines. These are the paths that air takes as it moves through the domain. While most of the air is laminar as it travels through the duct, a turbulent region can be seen on the leading edge of the duct approximately halfway through the duct. This area of recirculating air causes an increase in drag and occurs mostly at lower angles of attack. As the angle of attack increases, the duct becomes more in line with the freestream air and the turbulent region decreases. There is also a decrease in the change in velocity as it enters the duct.

The duct suffers from an increase in surface area which in turn increases the skin friction drag. There is an additional 0.84 m^2 of surface area in the ducted version compared to the non-ducted winglet. The increase in drag due to skin friction inside the duct is very small when taking into consideration the entire length of the blade, which is 547 m^2 .

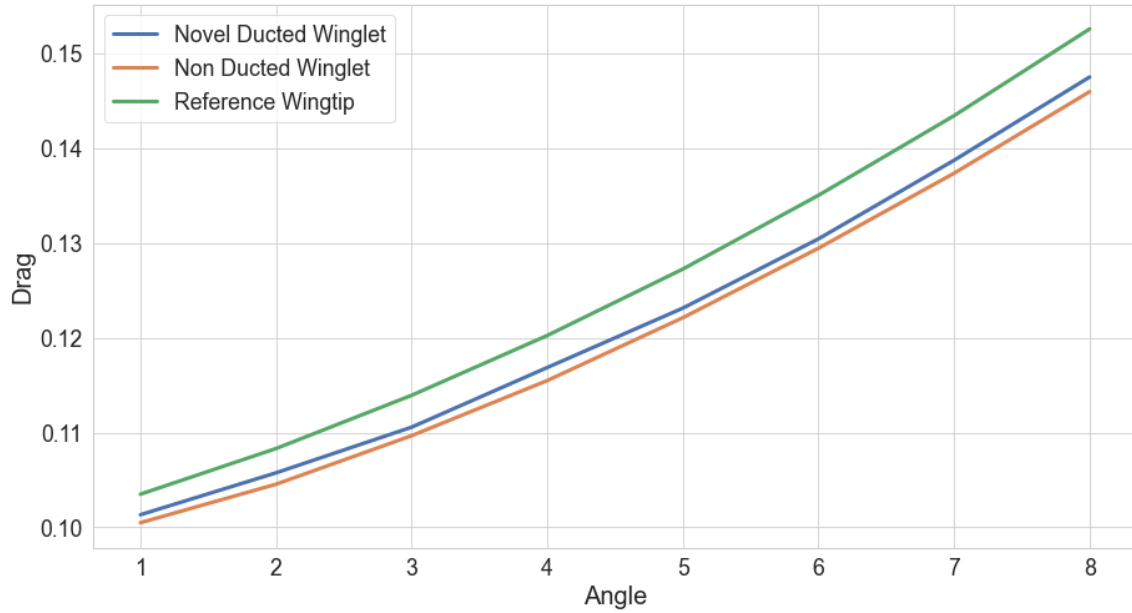


Figure 2.12: Non-dimensionalised drag values for the three geometries.

Table 2.6: Non-dimensionalised drag for the novel ducted winglet, the non-ducted winglet, and the unmodified reference blade. Values in bold indicate the lowest values.

| Angle | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Novel Ducted | 0.1013 | 0.1058 | 0.1106 | 0.1168 | 0.1231 | 0.1304 | 0.1387 | 0.1475 |
| Non-Ducted | 0.1005 | 0.1046 | 0.1097 | 0.1155 | 0.1221 | 0.1294 | 0.1373 | 0.1459 |
| Reference | 0.1035 | 0.1083 | 0.1139 | 0.1202 | 0.1272 | 0.1350 | 0.1434 | 0.1525 |

2.4.3 Lift-to-Drag Ratio

The lift-to-drag ratio is a key measure of how much lift a blade is able to generate for a given amount of drag. For wind turbine blades, this ratio is important because it determines how much power can be generated at certain wind speeds. Being able to generate more power at lower wind speeds is both financially and environmentally more efficient. Wind turbines which generate power when others are not means that electricity can be sold to the grid at a higher price compared to when all wind turbines are generating power. From an environmental standpoint, generating power at lower wind speeds means that the turbines can operate more frequently, thus reducing the dependency on other potentially polluting energy sources. The standard, non-ducted winglet has the highest lift-to-drag ratio of the three geometries at all angles of attack. The reference blade has the lowest lift-to-drag ratio while the novel ducted winglet is between the two for all angles of attack. Evidently, the increased lift of the reference blade is hindered by its drag. While both of the winglet designs suffered from a decrease in lift, the amount of drag per unit of lift means they are both more efficient than an unmodified blade. The performance of the duct of the novel design resulted in both a loss of lift and an increase in drag when compared to the non-ducted version, which led to a decreased lift-to-drag ratio.

2.4.4 Velocity and Pressure Changes Within the Duct

The path of particles of air that travel through the duct can be visualised using streamlines. The changes in velocity can be shown using streamlines by mapping the velocity at each node to the streamline. In Figure 2.14, changes in the behaviour of the air as it travels through the duct can be observed. At lower angles of attack, a region of air at the leading edge of the duct recirculates within the duct. At the trailing edge of the duct, air travels at a higher velocity and continues until it exits at the top of the winglet. The proportion of

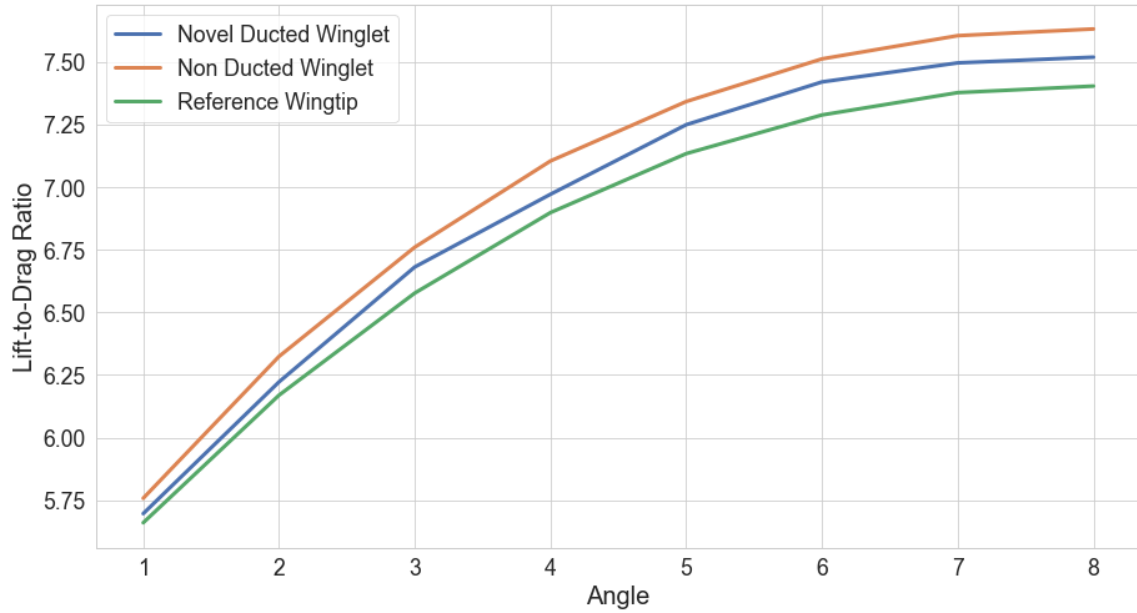


Figure 2.13: The lift-to-drag ratio across 8 angles of attack for the three different geometries.

Table 2.7: Non-dimensionalised lift-to-drag for the novel ducted winglet, the non-ducted winglet, and the unmodified reference blade. Values in bold indicate the highest values.

| Angle | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Novel Ducted | 5.6973 | 6.2224 | 6.6809 | 6.9710 | 7.2484 | 7.4191 | 7.4950 | 7.5178 |
| Non-Ducted | 5.7585 | 6.3241 | 6.7598 | 7.1037 | 7.3409 | 7.5109 | 7.6036 | 7.6303 |
| Reference | 5.6610 | 6.1695 | 6.5768 | 6.8979 | 7.1331 | 7.2873 | 7.3764 | 7.4025 |

air caught in the recirculation zone reduces as the angle of attack increases which results in a more uniform velocity profile across the duct.

The pressure on the duct wall across the range of angles of attack, shown in Figure 2.15, also help describe the flow characteristics in the duct. Lower angles of attack have a lower pressure as the mass flow rate of air decreases through the duct. A region of high pressure develops at the trailing edge of the duct and becomes larger as the angle of attack increases. Additionally, the pressure increase which corresponds to the increase in angle of attack demonstrates that the duct begins to constrict the free flow of air as it enters and travels through the duct.

2.4.5 Induced Drag and the Trefftz Plane

The novel ducted winglet was designed to specifically reduce the induced drag. Looking only at the overall drag is not completely indicative of its ability to meet its objective. The results from the Trefftz plane analysis show that the novel ducted winglet has a lower induced drag coefficient than both the reference blade and the non-ducted winglet for 7 of the 8 angles of attack. This can be seen in Figure 2.16 where z-score normalisation is used. Z-score normalisation sets the mean to 0 and standard deviation to 1; negative values are below the mean and positive values are larger than the mean. The non-ducted winglet has the highest induced drag of the three geometries. This result appears to be counter-intuitive as winglets themselves were designed to reduce the induced drag. A possible explanation for this is that because the winglet not only extends in the z-axis (upwards), it also extends in the x-axis (downstream). As mentioned before, the Trefftz plane is located in the same place for all three geometries meaning that the tip of the winglet is closer to the plane than the tip of the reference blade. The vorticity therefore has less time to dissipate. This would then lead to the higher induced drag values. An alternative would be to have

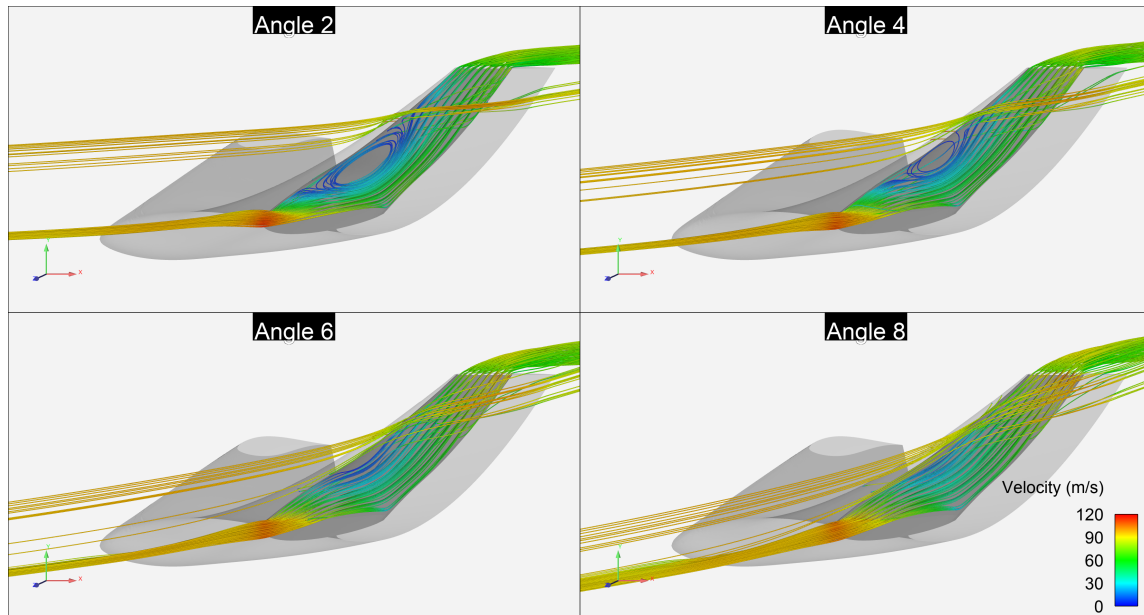


Figure 2.14: The path of air coloured by its velocity as it travels through the duct at increasing angles of attack.

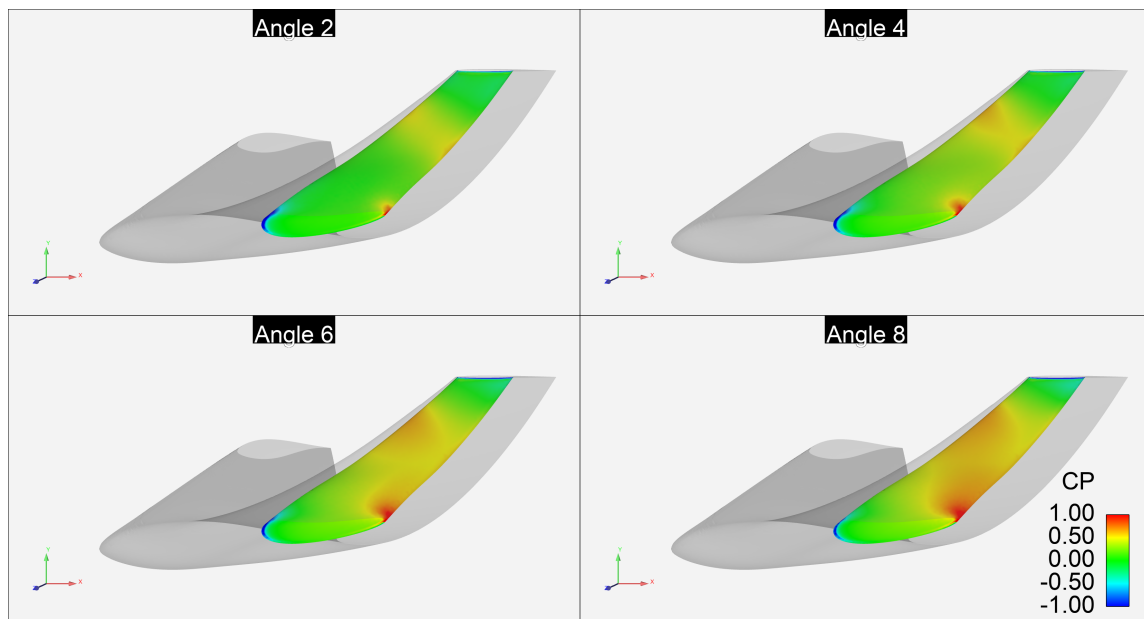


Figure 2.15: The pressure on the duct wall as air travels through the duct at increasing angles of attack.

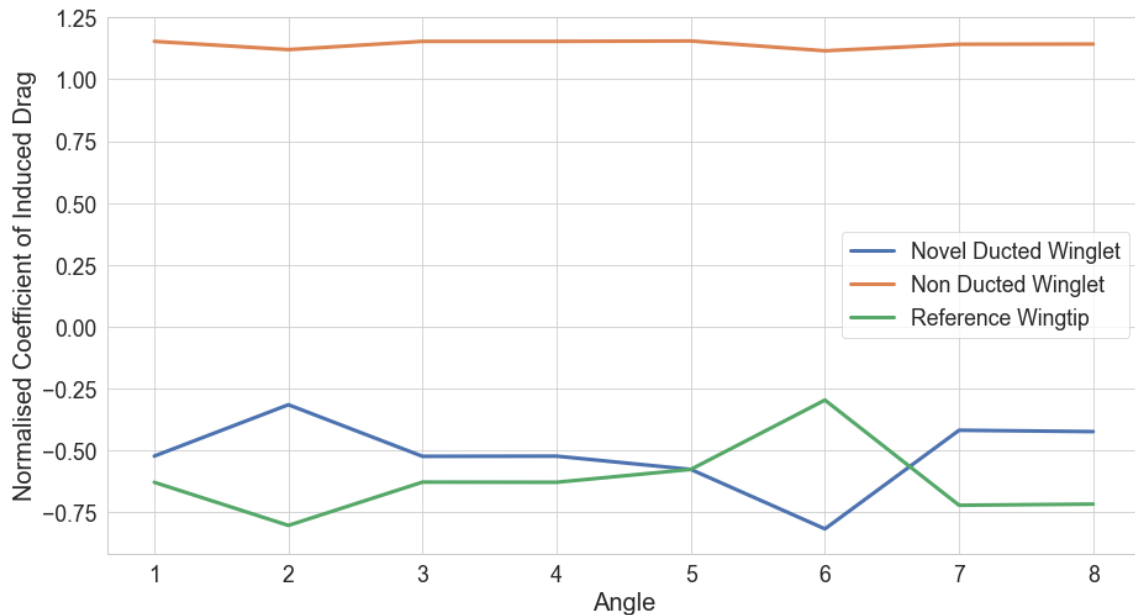


Figure 2.16: Z-score normalised coefficient of induced drag values for the three geometries.

the Trefftz plane the same distance from both the reference blade tip and the winglet tip, but then the distance between the rest of the reference blade would be larger than the distance from the wingleted blade and the plane. A larger distance between the blade and the point at which the analysis takes place would invariably lead to lower downstream vorticity for both real cases and when using CFD. In reality, vortices lose energy and dissipate while this phenomenon is captured in CFD using a dissipation factor. In either case of whether to situate the Trefftz plane in the same location or with the same distance from the blade tip, it brings confidence to the result that the novel ducted winglet can indeed reduce the induced drag.

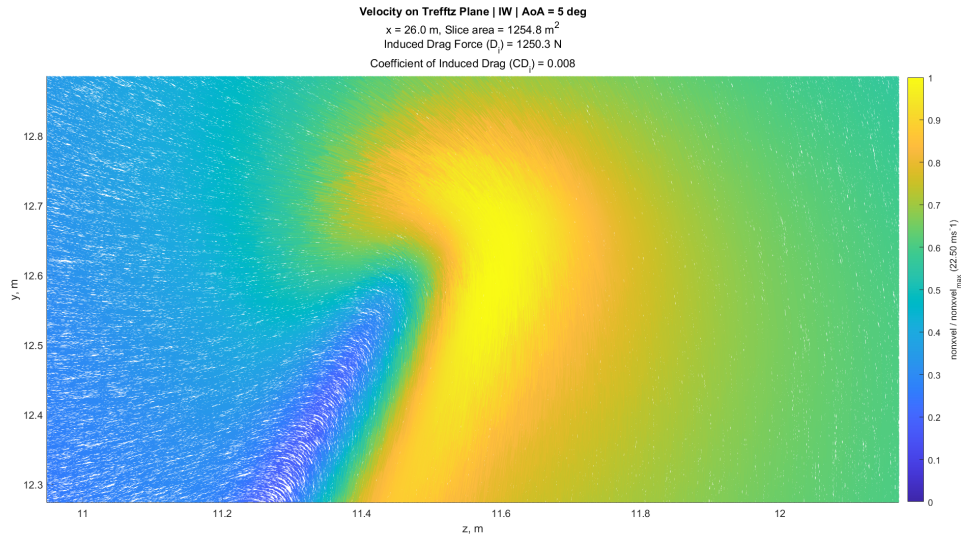


Figure 2.17: A close-up of the horizontal and vertical velocity components around the tip of the winglet.

Table 2.8: Induced drag values for the novel ducted winglet, the non-ducted winglet, and the unmodified reference blade. Values in bold indicate the lowest values.

| Angle | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Novel Ducted | 1.0129 | 0.6791 | 0.4326 | 0.2732 | 0.2016 | 0.2199 | 0.3227 | 0.5186 |
| Non-Ducted | 1.0155 | 0.6818 | 0.4373 | 0.2793 | 0.2086 | 0.2241 | 0.3311 | 0.5233 |
| Reference | 1.0130 | 0.6798 | 0.4328 | 0.2735 | 0.2016 | 0.2184 | 0.3240 | 0.5193 |

Table 2.9: Z-score normalised induced drag values for the novel ducted winglet, the non-ducted winglet, and the unmodified reference blade. Values in bold indicate the lowest values.

| Angle | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Novel Ducted | -0.524 | -0.316 | -0.525 | -0.524 | -0.577 | -0.818 | -0.419 | -0.425 |
| Non-Ducted | 1.153 | 1.112 | 1.153 | 1.153 | 1.155 | 1.115 | 1.141 | 1.142 |
| Reference | -0.629 | -0.804 | -0.629 | -0.629 | -0.577 | -0.297 | -0.722 | -0.717 |

2.5 Discussion

Contained within the patent are several claims that protect its copyright against infringement. Should anyone desire to make a similar device, it would have to differentiate itself against those claims. Within the claims lie certain assertions (to distinguish it from claim) that have been added about the performance and characteristics of the novel ducted winglet. Those assertions are investigated here; they are compared to the results from the CFD simulations using flow visualisations made during post-processing. Specifically, the following assertions are investigated: counter-rotating flows occur inside the duct (Assertion 1) and at the outlet (Assertion 2), a region of high pressure inside the duct channels the air as it passes through (Assertion 3), a reduction in velocity as air enters the duct (Assertion 4), an aerofoil-shaped region of air induces the rotation of air (Assertion 5), a 15% improvement in the lift-to-drag ratio, and that there exists non-symmetric velocity and pressure profiles within the duct (Assertion 7).

2.5.1 Patent Assertion 1

The patent asserts that within the duct are two regions of air with one rotating clockwise while the other rotates counter-clockwise. A figure in the patent accompanies this assertion. The patent goes on to state that the region of low pressure and the region of high pressure within the duct are indicative of two separate regions of flow each with their own rotational direction. However, it is not possible to determine flow rotation characteristics from the figure, or any velocity map. Without additional information, the only conclusions that can be drawn from the figure provided by the patent is that the velocity on the inboard edge of the duct is higher than the velocity on the outboard edge. Relative pressure can also be deduced from the figure, with lower pressure on the inboard side and higher pressure on the outboard edge. CFD results showing a comparable slice do not

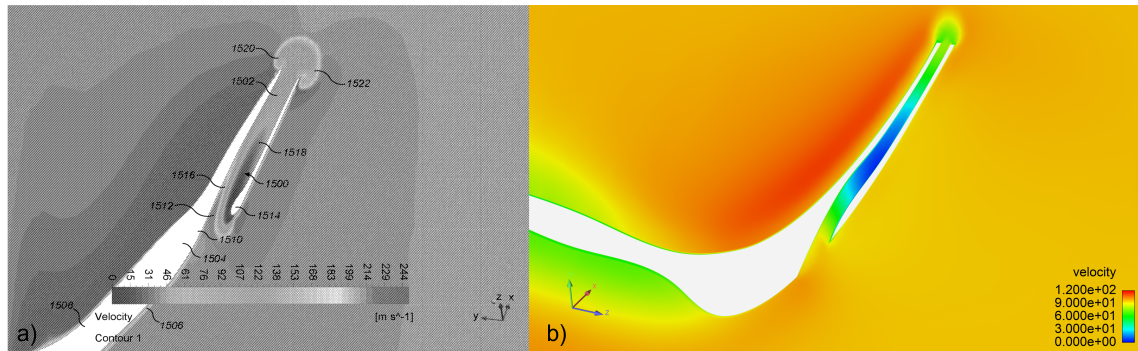


Figure 2.18: A comparison of a plane cut showing the velocity from a) the patent and b) the CFD results.

show the same differences in pressure and velocity between the two sides of the duct. As air enters the duct, it slows down, then accelerates as it rejoins the freestream. Similarities are observed on the outside of the duct and at the outlet; a region of higher velocity on the upper surface of the winglet and a ball-shaped region of lower velocity at the outlet as air from the duct rejoins the freestream.

The CFD results do not support the statement and the results from the patent could not be reproduced. Moreover, it would be difficult to have two regions of air that rotate in opposite directions in a region the size of the duct tested and especially in the direction indicated by the patent.

2.5.2 Patent Assertion 2

The patent asserts that there are two regions of air at the outlet that rotate in opposite directions continuing the on from the counter-rotating flows within the duct. The patent states that it is these two regions of air that greatly contribute to the reduction of downstream vorticity and induced drag; they are shown clearly in the schematic of Figure 2.19. The patent uses the same velocity figure as in Assertion 1 to justify the two regions of air.

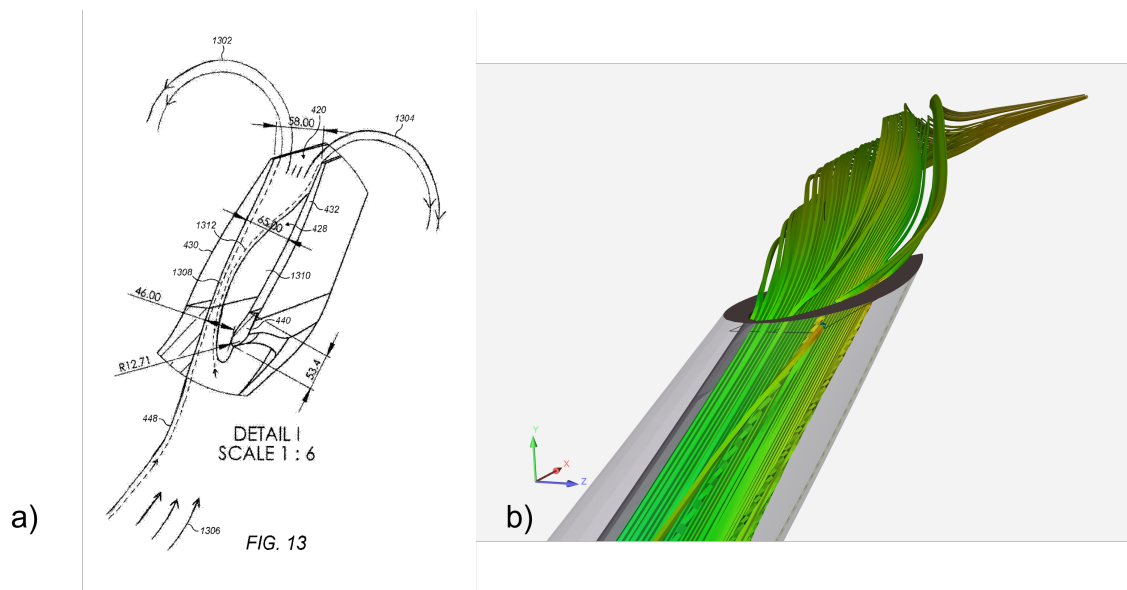


Figure 2.19: A comparison between a) how the patent describes the exhaust of air and b) the flow field results with streamlines superimposed.

While there is a different velocity at the outlet in the figure, it is not possible to determine the direction in which the air is rotating. Streamlines are needed to show how the air rotates as it travels out of the duct and rejoins the freestream. In Figure 2.19, it can be clearly seen that the air exits the duct and rotates only counterclockwise. This is also in line with results shown in Section 2.4.5, with the results from the Trefftz plane analysis. Additionally, this is expected from 3D aerofoil theory where air from the high pressure side diffuses to the low pressure side. A region of air that would rotate in the opposite direction of this would be met with lots of resistance.

As with Assertion 1, there is no evidence that there are two separate regions of air each with their own direction of rotation. This assertion can be classed as false.

2.5.3 Patent Assertion 3

Using Figure 2.18, the patent claims that a region of high pressure in the duct acts as an accelerator which channels the air through the duct. While it is possible to have two regions of air at different speeds in the same vicinity, air itself cannot act as a solid barrier to air; the two regions of air would mix. For example, air conditioning that is blown into a car travels at a different speed than the air elsewhere in the car. The air then mixes together. There is indeed a region of high pressure, visible in both the patent, Figure 2.18 A), and the CFD results, Figure 2.18 B). However, this region does not exist in isolation and it is actually a result of air entering the duct and slowing down.

This assertion is false as a region of high pressure cannot channel air. This assertion also begs the question: if there is a region of high pressure and low velocity, why not replace the region with a solid structure inside the duct and remove the region altogether? This would reduce drag and ensure a uniform flow through the duct.

2.5.4 Patent Assertion 4

The patent asserts that there is a reduction in velocity as air enters the duct. Based both on figures from the patent and from CFD results, it is easy to conclude that this assertion is true. Air does in fact slow down as it enters the duct. This assertion is proven also with Figure 2.20 and Figure 2.21.

2.5.5 Patent Assertion 5

This assertion in the patent states that an aerofoil-shaped region inside the duct induces rotation of air. Despite what is discussed in Section 2.5.1 which looked at two separate regions of rotating air, this assertion can be taken as stating that there is only one region of

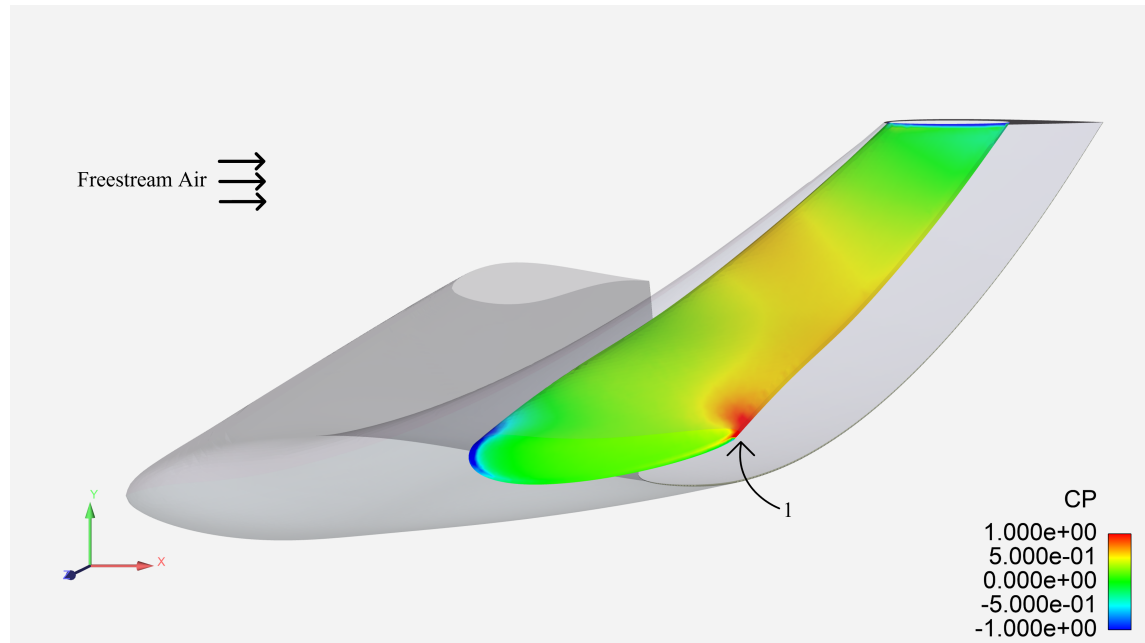


Figure 2.20: The coefficient of pressure plotted on the inside of the duct. A region of high pressure (1) develops at the rear of the duct inlet indicating a lower velocity.

air with one direction of rotation. Figure 2.19 indicates what is meant by aerofoil-shaped. It can be argued that such a region exists in the velocity field of Figure 2.18 from the patent. However, it is difficult to assess whether there is something similar in the CFD results provided in Figure 2.19. It does not appear the same distinct region as shown in the patent. As for the rotation, it cannot be concluded from the figure provided by the patent. Some form of visual or numerical analysis is needed to determine rotation direction. In this case, streamlines are used to see in which direction the air rotates, if at all. From Figure 2.21, the air travels through and does not rotate. The only rotation seems to be at the leading edge with the recirculating air that rotates about the Z axis.

There is no evidence to support the claim that there is rotating air within the duct.




Figure 2.21: The velocity (m/s) streamlines showing the path that the air takes through the duct. A small region (1) of recirculating air is visible at the leading edge of the duct near the midpoint.

2.5.6 Patent Assertion 6

While the other assertions in the patent are largely qualitative in nature, Assertion 6 provides a definitive, quantitative value which can be compared. The patent claims that there is a 15% reduction in the lift-to-drag ratio when compared to a blade without modifications. In our case, the reference blade used previously. At its best from the results gathered from CFD simulations, the ducted winglet was only able to reduce the lift-to-drag ratio by 1.8%. This is an order of magnitude out from the value in the patent. There are some reasons why there may be a difference in the two values. The first is that the design is not an exact replica of the one in the patent. Slight differences in the geometry might have led to reduced performance. Secondly, the application area is different. The patent tested the novel ducted winglet on an aeroplane and not a wind turbine blade. With respect to the application area, the freestream conditions are also different.

It is difficult to assess whether the reasons listed would be enough to make up for the 12% difference in lift-to-drag ratio reductions. While the design is not an exact replica, feedback was taken into consideration from the inventor and the design was ultimately approved for testing. By such a standard, the design must be close enough to the patented design to have all the necessary characteristics to perform. The application area is in fact different, but the method to test the winglet is identical to that of an aeroplane. The blade is stationary within the domain in the same way a wing would be. The freestream conditions is the biggest difference between the two simulations. Recall that Re to compare freestream conditions in wind tunnel testing. Re is of the same order of magnitude for both a wind turbine at sea level operating at maximum tip speed and a typical commercial airliner operating at cruise altitude and speed. This therefore brings confidence to the results provided. Compressibility effects would need to be taken into account but it is

likely that they would negatively impact the performance of the duct as regions of higher pressure would be exasperated.

2.5.7 Patent Assertion 7

The last assertion is on the velocity and pressure fields within the duct. The patent claims that there is a non-symmetrical velocity and pressure profile within the duct. Based on the results from the CFD, there is only asymmetry in the XZ plane through the duct, and not the YZ plane as indicated in the patent. It is also difficult to tell where exactly the velocity and pressure field cut was taken with respect to the duct; it could have been more towards the leading edge, the middle, or the trailing edge.

Based on the information available and the previous assertions, it is unlikely that there is an asymmetry in the duct as the patent describes it.

2.6 Conclusion

In this chapter, the fundamental aerodynamic principles that are needed to understand the phenomena experienced by objects travelling through air were discussed. The properties of air and how these set out the freestream conditions were also defined. Lift and drag were also discussed and how they can be derived from the dynamic pressure. The phenomenon of stall was discussed both as a result of an increase in angle of attack as well as boundary layer separation. The Navier-Stokes equations, the 3 sets of governing equations for physical flows, were also covered. How the Navier-Stokes equations are applied and used in FLITE3D was explained as well.

The patented ducted winglet was introduced along with the assertions made about it in the patent. Two additional designs were introduced to compare against: a unmodified

reference blade and a standard winglet. The experimental method was described to test the effectiveness of the winglet in reducing induced drag and maximising the lift-to-drag ratio. The assertions from the patent were compared to CFD results and figures from post-processing.

The novel ducted winglet had lower induced drag than both the reference blade and the standard winglet. The standard winglet had the highest lift-to-drag ratio of the three geometries, with the novel ducted winglet outperforming the reference blade.

One of the motivations for this work was the need to fill in the knowledge void left in the patent. The reasoning supporting the design in the patent is insufficient to say the least, and as has been shown in this section, at times wrong. Despite this, the novel ducted winglet still achieves its aim of reducing induced drag. The work presented in this chapter should serve as an update explanation for the behaviour of the novel ducted winglet. With evidence to suggest that the novel ducted winglet reduces induced drag, the next step is to determine if the induced drag, as well as the overall drag, can be reduced further through optimisation.

Chapter 3

Novel Ducted Winglet Optimisation

Following on from the previous chapter where the effectiveness of the ducted winglet is investigated, an optimisation study is carried out in this chapter to see if the induced drag can be decreased further. Optimisation methods are detailed with an emphasis on evolutionary algorithms and their uses. To demonstrate the optimisation effectiveness of evolutionary algorithms on aerodynamic problems, a simple 2D case is presented. The duct of the winglet is optimised to primarily reduce induced drag and secondarily reduce the overall drag to increase performance. A design of experiments approach, evolutionary algorithm, and Bayesian optimisation algorithm are employed. Following the initial study, the number of parameters are reduced to simplify the optimisation process and the number of iterations used when simulating the winglet are reduced to shorten the optimisation cycle. This chapter highlights the difficulties faced when optimising a complex engineering problem and compares the results of the different methods employed.

3.1 Introduction

Proving the effectiveness of the novel ducted winglet compared to a standard winglet and an unmodified blade was key to the industrial sponsor. An independent verification of the patent is also key to industry adoption. To further establish the case for the novel ducted winglet as a serious blade tip modification and to eke out as much performance as possible, an optimisation study is necessary. An initial design of experiments is conducted to further understand the behaviour of air as it travels through the duct. The designs created during the design of experiments are then used as the starting point for an evolutionary algorithm which is well-suited for creating several new designs at each generation. The designs from the evolutionary algorithm and the design of experiments are then used as data for Bayesian optimisation, which only creates one new design at each generation but prioritises information gain. Parameters are then reduced to shrink the parameter space and make it easier to find an optimum solution. In this chapter evolutionary algorithms and their application to aerodynamic optimisation are explained. Bayesian optimisation and Latin hypercube sampling are also discussed. A simple 2-dimensional aerofoil evolutionary algorithm is used as a test case for before being applied to the 3-dimensional novel ducted winglet. The initial optimisation process followed by the reduced parameter optimisation are detailed in this chapter as well.

3.2 Background

Several optimisation methods exist, each with their own advantages and disadvantages. At times, such as the case here of optimising a complex 3D object, it is difficult to know which optimisation method will be best suited ahead of time. The methods selected during the optimisation of the novel ducted winglet took into account the amount of

parameters, the amount of computational resources available, and the simulation time. Three optimisation methods are described here, starting from a basic method and progressively becoming more sophisticated. Details on how these methods work and how their characteristics affect optimisation performance are laid out in the following subsections.

3.2.1 Latin Hypercube Sampling

One method which samples the design space uniformly without the need for an exhaustive search is called Latin hypercube sampling (LHS) [28]. This method divides each input variable into N strata, where N is a user selected value, and allows each strata to be selected only once. In 2 dimensions, this would resemble a checkerboard where each row and each column could only have a piece each. The advantage of this method is that each input variable has a flat distribution curve leading to an even sampling distribution across the design space. LHS has been shown to be a better method than random sampling [29]. The method is suitable for problems with a single optimum.

3.2.2 Evolutionary Algorithms

Evolutionary algorithms (EA) are a series of metaheuristic algorithms which draw upon the theory of evolution proposed most famously by Charles Darwin [30]. These algorithms extract and apply the main concepts of mutation, selection, and crossover (or breeding). It is through these stages that computer scientists and engineers hope an optimal solution can be found. They are tested against a function that describes the optimum known as the fitness function. EAs are a subset of AI as they are capable of performing calculations and executing instructions faster than a human possible could.

Four nucleotide bases make the foundation of DNA: Adenine, Thymine, Cytosine, and Guanine. They provide the building blocks for all life on earth. Every species, subspecies,

and individual have a unique arrangement of these four bases. The information stored within the DNA yields the particular physical characteristics of the species or individual. According to Darwin's theory, the individuals with the best suited DNA will be able to pass it on to their offspring. As an example, only the gazelles which are faster and more agile than the fastest and most agile cheetah will survive. This optimisation will lead towards both faster gazelles and cheetahs.

EAs are inspired by the evolution of DNA. A fitness function is defined and the algorithm optimises for that function. The EA starts with an initial population size. From that starting population, the best individuals are selected. Once certain individuals have been selected, they go through a crossover process which yields new (and hopefully better) individuals. As a result, there is now a parent and a child population. To mimic nature, the child population undergoes a mutation, which makes it inherently different to the parent population. The selection, crossover, and mutation processes are repeated until a termination condition has been met. This can be either through a generation limit or through an fitness function threshold. Determining the appropriate termination condition is a challenge itself.

The process of crossover is fundamental to the working of an EA. Linking back to the building blocks of DNA, individuals in an EA population can be thought of as being made up of 0s and 1s. The 0s and 1s can appear in any order, but as mentioned with DNA, are essential in determining the characteristics of that individual. When crossover occurs, some sections in one parent are swapped with another. The result is two, new offspring individuals. There are many crossover strategies, however their objective is all the same: produce two, more optimal offspring.

Constant change is also fundamental to the behaviour of EAs. Once the parents have

reproduced and generated offspring, the offspring undergo a mutation. In most EAs, there is a very small likelihood that an offspring undergoes mutation, though it is critical to the success of the EA. By randomly changing a 0 to a 1, or vice versa, the offspring manages to introduce into the population more genetic diversity.

Lastly, the selection process identifies the fittest portion of the population for crossover and discards the rest of the individuals. The selection process can be tuned, much like the crossover strategy or the rate at which mutations occur.

Evolutionary algorithms are part of a class of gradient-free or derivative-free optimisation methods. These methods do not use information about the function at the evaluation point to determine where the next point along the function should be evaluated. Gradient-free optimisation methods differ from gradient-based methods in their approach and also in the types of problems they are well suited to solve. Some functions have multiple local optima, which are optimum points in a portion of the design space. Using gradient-based methods can cause the optimiser to converge on points in the design space which are not the global optimum. Additionally, gradient-based methods become impossible to use if the gradient cannot be calculated, thereby limiting their applicability.

Evolutionary Algorithm Strategies

There are some distinct methods which are used to change the characteristics of how EAs optimise. Some changes include tournament selection, niching, segregation through island models, and diversification with quality-diversity algorithms. Tournament selection introduces stochasticity at the selection stage by selecting multiple potential parents and choosing the top two parents through a bracket [31]. Niching preserves diversity by keeping multiple solutions at once [32]. Island models separate the domain into distinctly

separate *islands* with occasional sharing between the islands which is called migration [33]. Some algorithms, such as quality-diversity algorithms, try to maintain population diversity while presenting optimal solutions, creating a Pareto front with the number of solutions presented at once [34].

Benchmarks

To gauge the suitability of the EA, certain mathematical functions and sample engineering problems are used. These benchmarks have been used for years and across a wide variety of algorithms [35]. Four typical mathematical functions used as a benchmark are: Griewank [36], Rastrigin [37], Schaffer [38] and Ackley [39]. Each function introduced a hurdle in an attempt to slow the algorithm down and prevent it from finding the global optimum. For example, both the Rastrigin and Ackley functions contain numerous local minima.

In addition to the mathematical functions, a series of engineering problems are used to benchmark algorithms. These problems require the algorithm to define certain dimensions of the problem which are interconnected. Three engineering benchmarks are discussed here, although there are others as well. The benchmark detailed are the pressurised vessel akin to a submarine, a cantilever beam, and a spring.

First, the design of a cylindrical and domed pressure vessel is considered [40]. Here, four parameters must be found by the algorithm: the radius of the vessel, the thickness of the skin of the main body of the vessel, the thickness of the skin of the domes, and the length of the main body. Another interdependent problem is that of the welded beam [41]. In this benchmark test, the length, width, and height of the beam must be determined by the algorithm. Additionally, the length of the beam which has to be welded and

the weld height are also determined by the algorithm. There are additional engineering parameters which are introduced which constrain the problem space. Lastly, the tension and compression of a spring are used as an engineering design benchmark [42]. Here, the smallest possible spring must be designed, with the coil diameter, wire diameter, and number of coils being used as variables. Again, engineering parameters are introduced to constrain the problem space.

No Free Lunch Theorem

EA developers will evaluate their algorithms based on benchmarks and claim that their EA is better, but these developers (and algorithms) fall victim to the No Free Lunch Theory [43]. The No Free Lunch Theory states that for all problems in which a supposed *Algorithm X* outperforms *Algorithm Y*, there are an equal amount of problems in which *Algorithm Y* outperforms *Algorithm X*. In short, there does not exist an algorithm that will be the best choice in all scenarios. The development and tuning of EAs is itself an optimisation problem and EAs are developed on a case-by-case basis with their applicability to different situations being very limited.

Subjective Fitness Functions

The mathematical models, engineering benchmarks, and aerofoil optimisation can all be expressed in terms of one or more equations which the EA attempts to optimise. If this would be the case of all engineering design problems, then there would be no need for designers and the reliance on an engineer's expertise would be greatly reduce. However, engineering problems are often accompanied with constraints that cannot be modelled or would increase computation time significantly. If an aerofoil is taken as an example, additional constraints can be identified. The position of the engine or engines may affect

airflow. Components such as flaps and slats which are located at the trailing and leading edge may need a certain thickness to accompany electronics and hydraulics. Designers must consider the manufacturability of their designs as well. They may also be inclined to select more aesthetically pleasing aerofoils among a list of optimised aerofoils with similar characteristics.

Designer preference strongly influences design, even when EAs are used in areas such as non-technical art spaces. Fashion designers have used algorithms to create new arrangements of clothes [44]. In this case, there was no optimal solution, and relied entirely on the preferences of the designer. This new technique utilises a high mutation rate (20%) and low generation limit since its primary focus is that of an exploratory tool.

In Thai culture, there is a kranok pattern which appears often in art and on religious symbols. While the pattern is distinctive, there is a wide range of artistic features that can be implemented in the design. In another example of aesthetic based EA, Leelathakul & Rimcharoen [45] demonstrated a tool with which users were able to generate their own kranok patterns.

The application of EAs to art extends beyond the visual landscape, and into music. A new HCI model was built by Farzaneh & Toroghi [46] which allowed users to score EA-generated music and alter the output of the algorithm based solely on their preferences.

Lastly, user preference has also been applied to game design [47]. Users can select all outputs which suit their criteria and the algorithm will continue to generate similar levels.

Interactive Evolutionary Algorithm

An Interactive Evolutionary Algorithm (IEA) involves users in the optimisation loop. Their intersectionality lies between the user and the algorithm and strives to create something which is better than the sum of its parts. People are good at creative problem solving while machines excel at large-scale computations. Conversely, this relationship between human and computer also attempts to diminish the weaknesses of each. Humans can only remember so much at once and for so long before fatigue sets in. Computers may be good at storing a vast amount of information, but they also need that plethora of information to operate. In the context of EAs, large populations with several generations are needed before suitable convergence. This problem was highlighted by Takagi [48] and went so far as to say that reducing human fatigue is the biggest issue that needs to be resolved.

There is also a considerable amount of literature that champions the use of IEA's as a means of design optimisation. Algorithms provide an excellent starting point for designers as to what parts of their design that could be optimised. They can support the development of designs and also assist designers by fulfilling vague ideas [49]. While designing an ergonomic chair with the use of an EA, Brintrup [50] noted the efficacy of the algorithm but also stated that "a completely computer-based ergonomic design optimization framework is neither desirable nor possible."

In a side-by-side comparison of IEA and pure algorithm, the IEA was able to outperform the algorithm in a series of tasks [51]. This points to the advantage of human influence within design and with respect to overall design objectives. This can be further observed with the design of water pipes. Johns et al. [52] used an EA to determine the optimal layout of a water distribution network. The design is limited by factors such as cost, reliability and water quality. Moreover, the civil engineers who design these kinds of

networks must take into consideration of hydraulic pressure and the ability of the joints to withstand that pressure. The incorporation of the knowledge of the expert is key to the process. In the context of water pipes, an EA may produce many mathematically viable solutions, but their practical implementation would be impossible to achieve. By integrating the designer in the algorithmic loop, the EA can then generate solutions that are both mathematically sound and possible to implement. As with many of the IEA, experts must be able to determine how the algorithm is behaving and must have an easy manner in which to intervene. They must also have an easy manner in which to interact with the program.

Nature-Inspired Algorithms

One of the biggest differentiating factors of EAs is the mutation of the offspring. In addition to this, EAs will mimic the behaviour of a species. This allows the algorithm to take advantage of certain behaviours of that animal or group of animals. It uses this behaviour to modify the standard select-crossover-mutate formula.

An example of this is the Whale Optimisation Algorithm [53] which simulates the feeding patterns of humpback whales. To effectively catch their prey, whales will circle underneath their intended targets and slowly spiral inwards and upwards until they make their catch. The algorithm mimes this behaviour by converging on a solution in a spiral-like manner. A downside to this method is its susceptibility to false convergence in a local minimum.

Another, and quite popular, algorithm is the Artificial Bee Colony algorithm [54]. In this algorithm, the individuals in the population are metaphors for bees foraging for food, while the solutions are food. The better the solution, the better the so-called quality of food. This algorithm adds a layer of complexity as there are three different roles that a

bee can have: gathering, scouting, observing. A gathering bee will go to a known source of food, which in a practical sense is a solution, and returns it to the hive. A bee scout will go out of the hive and search for food. Once a source of food has been found, the scouts return to the hive and convey this information to the bee gatherers. The observing bees will take note of the many exchanges happening between gatherers and scouts and then themselves turn into gatherers by going to the source of food they believe is the best. Thus, the algorithm is able to continuously generate better and better solutions.

Another communication-based algorithm is the Firefly Algorithm [55]. The communication of fireflies is based on their brightness and the algorithm sets out three rules which take advantage of this behaviour. The first is that each firefly can attract another firefly. The second states that the attractiveness of a firefly is directly correlated to the brightness of the firefly. In effect, a brighter firefly will be more attractive than a dimmer one. As a result, the dimmer fireflies move towards the greatest source of light. The last rule determines the brightness of the fireflies. A firefly which sits in a better spot within the search space with respect to the fitness function will be brighter.

While some algorithms focus on the communal behaviours of a species, the Cuckoo Search Algorithm (CSA) [56] focuses instead on the nefarious and ultra-competitive behaviour of the cuckoo. These parasitic birds lay their eggs in other birds' nests while discarding the other eggs already in the nest. Like the Firefly Algorithm, CSA has three rules upon which the algorithm is based. Firstly, one egg can be placed in a random nest at a time. Secondly, only high-quality eggs can survive. Thirdly, the number of nests is constant. CSA also introduces a probability that the original bird will be able to find the cuckoo egg. The algorithm finds the best solution by discarding nests which are not fit and keeping the rest. This algorithm was modified by Walton et al. [57] by allowing the best solutions

to exchange information amongst each other. This change led to faster convergence.

Each evolutionary algorithm, whether nature-inspired or not, is usually developed to tackle a unique problem or category of problems. As such, to say that one algorithm is better than another would be erroneous. The efficiency and efficacy of each algorithm must be determined on a case-by-case basis. The trend of naming and developing algorithms after animals and other natural phenomena is now being actively discouraged as many of these algorithms overlap considerably and crowd the literature [58].

Multi-Objective Optimisation

Multi-objective optimisation (MOO) problems are particularly well suited to EAs. Many engineering problems can be classified as MOO; these designs involve many different trade-offs. It is often challenging for human designers to determine an optimal solution without the aid of a computer-based system. One example of an MOO problem where an EA was used successfully is presented by Tawhid & Savsani [59] who evaluated their Artificial Algae Algorithm on 20 different benchmark problems to show its suitability for engineering design optimisation problems. An industrial application of MOO engineering can be found in aerospace, specifically the design of geostationary satellites. Customer demands for higher bandwidth capabilities and longer lifespans [60] are at odds with the many constraints in place. These constraints include

- Higher power requirements;
- Thermal considerations, including hot and cold zones;
- The ability to fit in a rocket fairing; and
- The ability to withstand vibrations during launch.

Each of these variables need to be considered throughout the design lifecycle of the satellite. To illustrate the usefulness of EAs in this context, Berrezzoug et al. [61] proposed a method which applies a Gravitational Search Algorithm to the design of a satellite by considering the many design variables. The approach proved useful for determining optimal trade-offs between the variables.

Optimisation problems do not need to be strictly product design based either. The layout of a construction site is another kind of MOO problem. Offices, equipment and warehouses are among the factors which influence the layout. The objective for any layout is to maximise the time spent on value adding activities while decreasing wasteful activities like transportation or waiting. In addition to this, the safety of the construction workers and engineers cannot be compromised. One attempt at this was done by Kaveh et al. [62] who demonstrate the application of a Colliding Bodies Optimisation algorithm to a construction site layout. Their solution demonstrates that a problem with $n!$ combinations can be effectively solved with an EA.

Applications in Aerospace Engineering

Beyond benchmark problems, EAs are used to optimise real-world problems in a variety of domains, such as supply chain optimisation [63], production and manufacturing [64], and engineering [65]. Particularly, EAs have been used in aerospace engineering. The highlighted literature in this section is relevant to the work carried out in this thesis as the method of simulation and optimisation is very similar to how aiplane wings are simulated and optimised.

Joly et al. [66] proposed the successive use of two types of EAs to account for both the low-fidelity radial distributions and high-fidelity aerostructural performances during the

design of a compressor in a turbojet engine. While this is a specific example, Murman et al. [67] highlight the broad spectrum within the industry to which these algorithms could be applied. With the idea that every successive plane must be “better, faster, cheaper,” this points to the notion that incremental gains are key to achieving long term success. These incremental gains can be brought to fruition with the use of EAs. In order to achieve measurable success, the entire lifecycle of aerospace products must be considered. This includes non-destructive testing which is key to ensuring the robustness and durability of components and ultimately the safety of its users. A classification system based on EAs is proposed by D’Angelo et al. [68] to facilitate the process of finding defects in aerospace components. In addition to making components durable, they must also be made as lightweight as possible. Lighter planes require less thrust to stay aloft, and thus require less fuel to operate. This also means they are cheaper to fly. Alternatively, lighter components will allow planes to handle more cargo. A novel technique to optimise the internal structure of additively manufactured aerospace components is based on EAs [69].

Aircraft are made up of several, interconnected systems which must all work together and behave as expected in flight. One critical part of an airplane are the wings. They are the main source of lift for most aircraft, however they are also a big source of drag. Modifying their shape to suit the mission criteria can be difficult as there are several competing constraints: weight, structural integrity, desired speed and climb rate, etc. Traditionally, aerofoils were tested in a wind tunnel to determine their characteristics [70]. With advances in computational power, the next stage in testing was done using CFD software. An open-source, 2-dimensional program called XFOIL is able to accurately define the characteristics of any aerofoil and is just one of many software which can produce accurate results. However, the shape was mainly determined by the engineer [71]. More recently, the shape of the aerofoil has been influenced with EAs [72], as with the work done by Mirjalili

et al. and work done within XFOIL too [73]. An additional study furthered the usefulness of EAs with respect to aerofoil optimisation by demonstrating that an S1223 airfoil's lift and drag coefficients could be increased by 2.27% and decreased by 1.4% respectively [74]. It is worth noting that while these algorithms perform well in a 2D space, real wings are more complicated. Their shape usually changes along the span and they are also subject to wingtip aerodynamic forces which affect their performance.

While aerofoils development using an EA is usually related to airplane design, these advances can be applied to other domains as well. The blades of a wind turbine and helicopter can also benefit from aerofoil optimisation [75]. Moreover, propellers for marine use would benefit too [72].

EAs have been used extensively to optimise the shape of wings. Their use has often been in combination with other tools or methods to deliver better results. Initial 3D optimisation research conducted by Oyama et al. [76] showed the capabilities of real-coded adaptive range EAs for transonic wings. Obayashi et al. [77] demonstrated a MOO EA for optimisation of supersonic wings.

EAs have also been used for new wing concepts. Morphing wings, which change their shape based on current flight conditions, will tune themselves for the optimum lift to drag ratio. Koreanschi et al. [78] employed a custom-made EA to a morphing wingtip to study the effects on drag. They found their algorithm to be better than a standard gradient decent approach.

Utilising the strength of EAs and bringing in other tools to enhance EAs has been shown to be an effective manner of reducing optimisation time. Boutemedjet et al. [79] couple an EA with an artificial neural network to optimise the wing for an unmanned aerial vehicle.

Isakhani et al. [80] combined PARSEC, Nash strategy and an EA to improve the lift to drag ratio of a NREL S809 aerofoil by 77%.

Wing optimisation has also occurred outside the scope of airplane wings. EA optimisation has been applied to helicopter wings [81] and to wind turbine blades [82, 83, 84]. Given the vast range of EAs to choose from and taking into consideration the No Free Lunch Theorem, a simple EA with tournament selection was chosen for this work.

3.2.3 Bayesian Optimisation

The function that relates input variables to their outputs in complex optimisation problems are usually non-linear, high dimensional, and are computationally expensive to evaluate. This is the case for optimising the novel ducted winglet using CFD. One of the difficulties during the optimisation process is identifying where in the parameter space to look next. LHS will explore the parameter space but does not update the model to find out where to look next. Evolutionary algorithms take it a step further and will update the scope of the parameter space with each generation. As discussed earlier though, they can be prone to getting stuck in local minima or maxima. Additionally, they require more computational resources as many more evaluations must be run. This is where Bayesian optimisation becomes a better suited optimisation method. It identifies where to look next in the parameter space after every evaluation.

Bayesian optimisation relies on 3 functions:

1. The objective function;
2. The surrogate function;
3. The acquisition function.

The role of the objective function is to set out the variables that are going to be used as inputs and dictates the output. Importantly, the objective function is unknown prior to starting the optimisation process. A guess as to what the objective function might be is set by the surrogate function. It acts as the best representation of the objective function for the model. Lastly, the acquisition function determines the best place to sample the parameter space based on the information in the surrogate model. Bayesian optimisation is effective because the acquisition function maximises information gain during each sample. The acquisition function will either choose to select sparsely sampled areas of the parameter space or areas where a minimum could exist.

3.3 Method and Materials

In the previous chapter, the effectiveness of the novel ducted winglet was confirmed through the lift-to-drag ratio and the Trefftz plane analysis. However, no work was done to further reduce the drag. The motivations and decisions behind the optimisation process are explored in the following sections. The initial attempt at optimising the winglet involved using 6 elliptical profiles which were defined by their eccentricity, chord length, and relative position to the leading edge of the winglet. This gives 18 parameters in total to optimise. The first stage of the optimisation process was a design of experiments in which 5 designs were tested. Those designs, along with the original design, were used as input to an evolutionary algorithm where 12 more unique designs were tested. Finally, the 18 unique designs were used as input to a Bayesian optimiser, which produced 3 additional designs.

3.3.1 Design of Experiments

Each elliptical profile used to define the duct could be manipulated in three different manners. Changing the eccentricity - how far from a circle the ellipse becomes - would allow the duct to determine how much mass air flow rate would be possible to go through. The maximum mass flow rate can also be altered by changing the chord length of the ellipse. However, these two parameters were semi-dependent on each other. Though each could be set independently, an ellipse with a low eccentricity and large chord would sit outside the walls of the winglet. Contrarily, a profile with higher eccentricity and smaller chord would fit inside the winglet. It would be possible to have two profiles with different properties that could still achieve the same mass flow rate. Obviously these two profiles would be differently shaped and as a result would influence the manner in which the air flows through the duct. The last parameter determines where the profile is located along the y-axis. Profiles could be closer to the leading edge, the trailing edge, or in the middle. Again, this influenced the previous two parameters as profiles near edges had to have higher eccentricity. The effects of the semi-dependence of all three of these parameters becomes apparent particularly when using the Bayesian optimisation and later when employing LHS.

After the initial study using the design approved by the inventor, some questions were raised regarding the characteristics of the duct and how they affected the results. How did the size of the duct affect performance? Should the profiles be uniformly sized and shaped throughout the duct? Did the duct benefit from contracting or expanding halfway through the duct?

5 designs were created to answer those questions and to potentially guide the optimisation process further. The different geometries each had their own unique characteristic with

Table 3.1: 3 parameter groups which defined the duct along with their respective ranges.

| Parameter | Range |
|--------------|--------------|
| Eccentricity | [0.05, 0.50] |
| Y-axis shift | [1.00, 2.35] |
| Chord length | [0.18, 0.90] |

parameters available in Table 3.1:

- Maximising the chord length.
- Minimising the chord length.
- Minimising eccentricity.
- Contracting the chord length halfway.
- Expanding the chord length halfway.

In each of these designs, shown in Figure 3.1, changes were minimised from the original design so that changes could be attributed to a single change rather than a combination of two or more factors. Considerations had to be taken when setting up new designs as the interdependence among parameters affected the viability of the design. For example, the placement of the profile along the y-axis influenced the eccentricity of the profile, which in turned influenced the chord length.

3.3.2 Evolutionary Algorithm

The code for the EA used to optimise the novel ducted winglet is based on the 2D Xfoil example, outlined in Section 3.3.4 with the steps outlined in Algorithm 1. Using the 6 designs from the Design of Experiments as the first generation, the geometries were ranked based on their induced drag. The breeding process was done by tournament

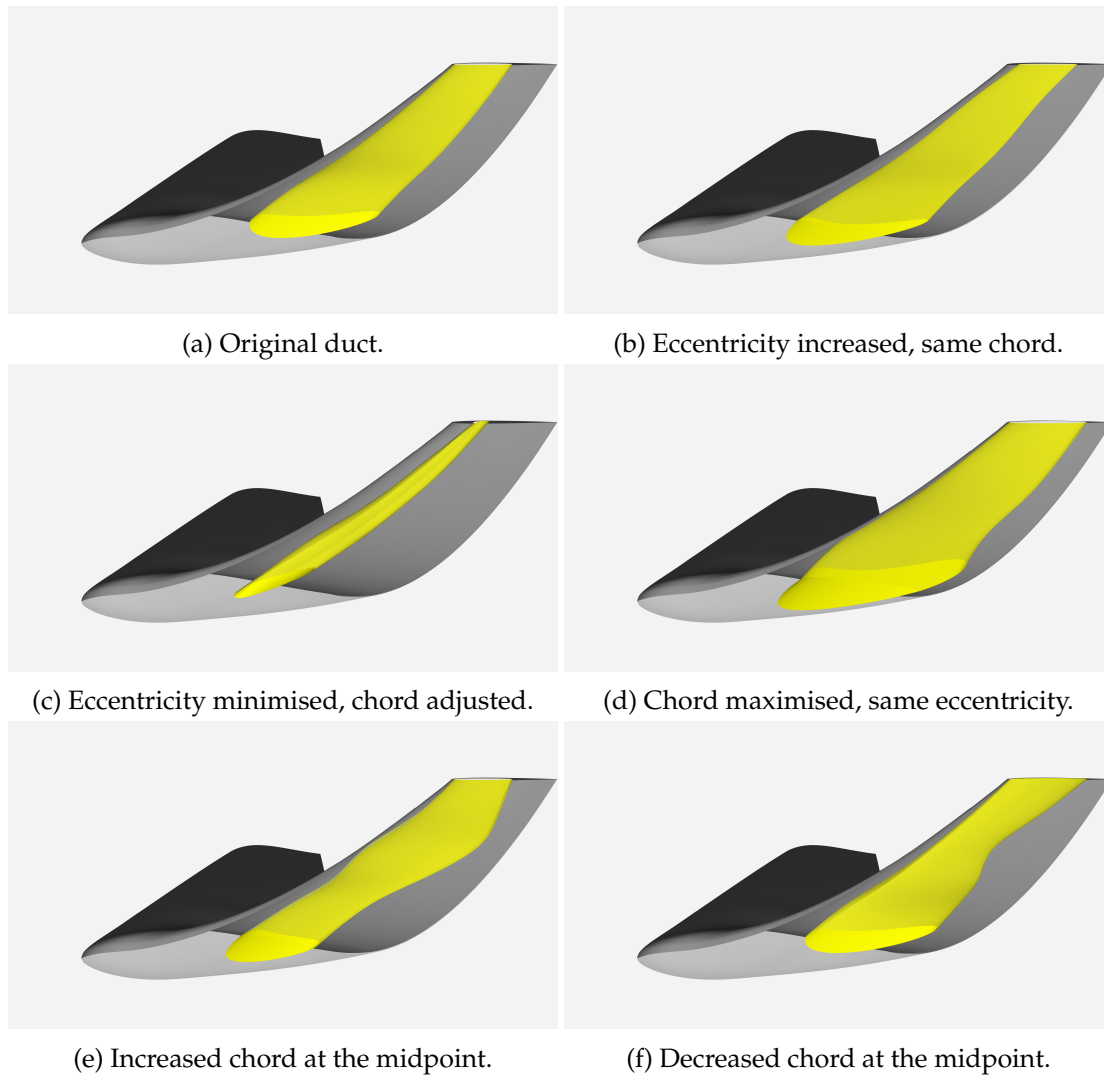


Figure 3.1: The five different designs compared to the original duct. One parameter was changed at a time to monitor the effects of each change.

Algorithm 1 3D evolutionary algorithm.

```
1: Initialise population, evaluate and rank.
2: for Number of generations do
3:   Remove bottom 50% performers from population.
4:   Select parents using tournament selection.
5:   Choose crossover point from 1 to 5.
6:   Duct profiles before the crossover point came from the first parent and after the
   crossover point came from the second parent.
7:   for Each Profile do
8:     if Profile to be Mutated then
9:       Choose one among eccentricity, chord length and y-axis shift for mutation.
10:    end if
11:    Exit if mutation has occurred.
12:  end for
13:  Evaluate and rank the new population.
14:  Repeat
15: end for
```

selection. Given the nature of the tournament selection, the lowest ranked geometry could not be selected for breeding. Only one child was made per pair of parents. A single crossover point was selected where a random number of duct profiles were used from the first parent, with the rest being from the second parent. At least one duct profile was used from each parent. A mutation could occur on one of the parameters of a single duct once the crossover had been completed. Mutations could affect only one of the duct profiles at a time, and the mutation only changed one of the three parameters defining the duct profile. The chord and eccentricity could be increased or decreased while the position of the duct could be moved forward or backward with respect to the leading edge of the winglet. Only one mutation was allowed per child. It was also possible that no mutations took place. The EA was run for 3 generations with explanations for the limited number of generations explained in Section 3.6.2.

3.3.3 Bayesian Optimisation

The Bayesian optimisation code used in this work was provided by Rahat et al. [85]. The code provides a variety of methods to use for optimisation as well as control over a variety of different parameters though the majority of the parameters were kept at their default values. Like the EA, the Bayesian optimiser was run for 3 iterations.

The variables for the optimiser were input using a CSV file which was updated after each run. In summary the process followed the steps in Algorithm 2.

Algorithm 2 Steps for new geometry using Bayesian optimisation.

- 1: 18 duct parameters plus their respective induced drag values were read into the program.
 - 2: 18 new parameters were output by the program.
 - 3: The new parameters were used as input to a Matlab program which created the individual duct profile text files.
 - 4: The text files were uploaded to Solidworks as curves.
 - 5: A boundary cut was made using the lofted section from the curves.
 - 6: If the duct went beyond the walls of the winglet the process was restarted; otherwise the geometry was exported to Flite3D.
 - 7: The geometry was simulated and the results added to the dataset.
 - 8: Repeat.
-

3.3.4 2D Evolutionary Algorithm Test Case

To determine if an EA can work on a small population, a simple EA was made using R and simple aerodynamic testing software called Xfoil [86] to evaluate 2D aerofoils.

Method and Materials

Xfoil uses a panel method to calculate the pressure across the aerofoil. The panels can be seen in Figure 3.2c and the pressure distribution can be seen in Figure 3.2a. Lift and drag values are then calculated based on the aerofoil pressure distribution. The algorithm steps

are detailed in Algorithm 3. The breeding process took place using tournament selection and crossover utilised the blending capability of Xfoil. The blending ratio was set to 0.5, meaning that coordinates of the child aerofoil were set halfway between the two parent. Afterwards, there was a possibility that a mutation could occur on the upper surface of the aerofoil. Panels on the upper aerofoil that were not near the leading edge or the trailing edge could be moved either up or down. The number of panels to be moved along with the direction were randomly chosen.

Algorithm 3 2D evolutionary algorithm.

- 1: Initialise population, evaluate and rank.
 - 2: **for** Number of generations **do**
 - 3: Remove bottom 50% performers from population.
 - 4: Select parents using tournament selection.
 - 5: Crossover using the Blend method within Xfoil.
 - 6: Mutate the aerofoil by raising or lowering a range of 2 to 160 nodes on the upper surface.
 - 7: Evaluate and rank the new population.
 - 8: Repeat
 - 9: **end for**
-

The EA could optimise for maximising lift, minimising drag, or maximising the lift-to-drag ratio. The results shown in this work are from maximising the lift-to-drag ratio and were run for 5 generations.

Results

The evolution of the top four aerofoils are shown in Figure 3.3 with the best, mean and worst fitnesses plotted in Figure 3.3f. The fitness of the top aerofoil improved slightly over the generations while the mean and the worst improved considerably. While an improving fitness is the goal of the optimisation, the diversity of the population decreases as seen in Figure 3.3e.

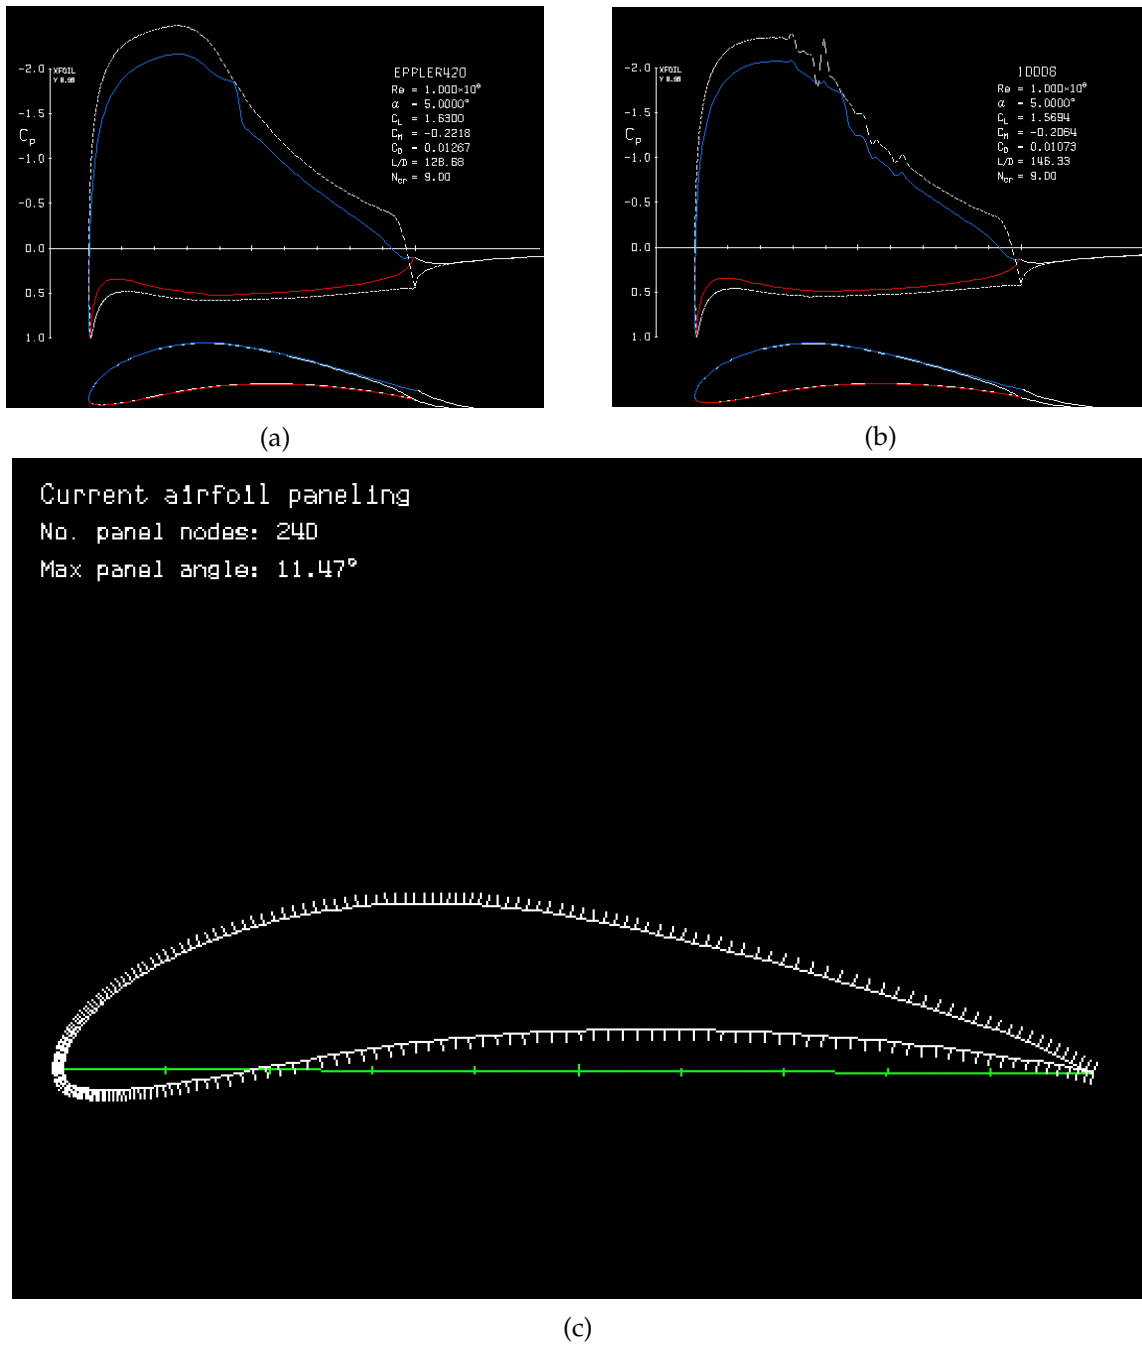


Figure 3.2: Pressure distributions for two different aerofoils and the associate panel distribution across the upper and lower surface. Panels are concentrated at the leading edge where changes in pressure are higher.

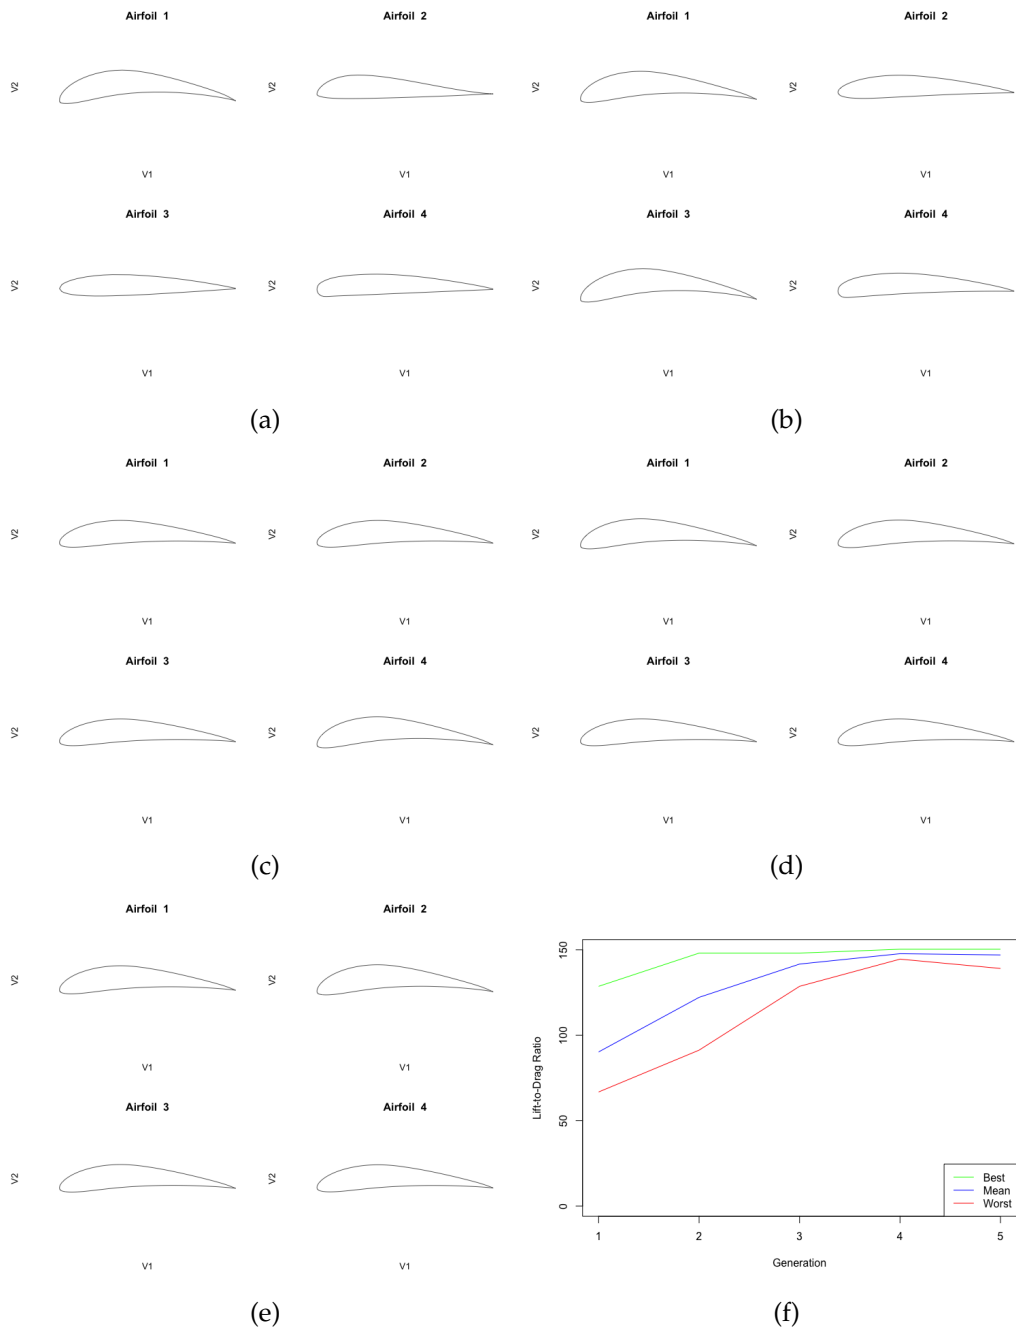


Figure 3.3: Evolution of the top four aerofoils for lift-to-drag ratio. The best, mean, and worst fitnesses are plotted (f) across each generation to show an upward trend in the lift-to-drag ratio. The diversity of the population is reduced as the generations increase though subtle changes are still visible in the last generation (e).

Conclusion

The lack of diversity is a challenge that plagues simple EAs and is addressed by quality-diversity algorithms. Another challenge when using Xfoil was the method of mutation. Even though the panels were raised or lowered very slightly, the bump that resulted from the change in height resulted in discontinuous pressure distribution. Occasionally, the discontinuity prevented the solution from converging and producing a result. Some discontinuities are shown in Figure 3.2b. This was handled in the EA by disregarding the aerofoil when breeding though it could affect the ability for the EA to optimise.

3.4 Results

The results from the three different methods are given in this section. Flow visualisations and pressure coefficient figures are provided to show a comparison among the various designs.

3.4.1 Design of Experiments

3 of 5 designs created during this portion of the optimisation study had a higher induced drag compared to the original design, while 2 designs had a lower induced drag. The two designs, F and B, managed to lower the induced drag by 0.28% and 0.51% respectively. 4 of the 5 designs had a lower overall drag when compared to the original novel ducted design, with the best performing design (C) lowering the overall drag by 0.65%. Results from the design of experiments show that minimising the chord reduced the total drag, but increasing the eccentricity reduced the induced drag. Increasing the chord overall and at the midpoint both performed poorly when considering the induced drag. The results are available in Table 3.2 and are shown in Figure 3.5 as the 0th generation.

Table 3.2: Overall drag and induced drag values for the five designs tested during the design of experiments along with the original novel ducted design. Values in bold indicate the lowest value in each row.

| Parameter | Original | Design B | Design C | Design D | Design E | Design F |
|------------------|----------|---------------|----------------|----------|----------|----------|
| Drag | 0.6022 | 0.06015 | 0.05983 | 0.06076 | 0.06006 | 0.06020 |
| Induced Drag (N) | 1243.5 | 1237.2 | 1245.0 | 1254.2 | 1254.3 | 1240.0 |




Figure 3.4: The best performing duct from the design of experiments. The duct is (a) highlighted in yellow to better show its profile along with (b) velocity streamlines and (c) pressure as air travels through the duct.

The best design from the ones tested had an increased eccentricity, making the duct less circular. Like the original design, there is still a region of recirculating air at the leading edge while the velocity of the air increases towards the trailing edge of the duct. The streamlines and pressure within the duct are shown in Figure 3.4.

3.4.2 Evolutionary Algorithm

3 generations of the EA were run with most of the designs producing an increase in induced drag, with the exception of 2 designs. Overall, the designs produced much higher values of induced drag compared to the Design of Experiments and compared to the original design. The average value of induced drag across each generation also increased with each generation. Results from the EA are shown in Figure 3.5 with the results of the designs created during the EA stage shown in Generations 1 to 3. Limitations of the approach and reasons why the EA performed poorly are discussed in Section 3.6.2.




Figure 3.5: Results from the evolutionary algorithm with the red line serving as a baseline set by the original design. The 0th generation are the designs used in the Design of Experiments stage, used as parents for the EA. Generations 1 to 3 are designs created using the EA.




Figure 3.6: The best performing duct from the evolutionary algorithm. The duct is (a) highlighted in yellow to better show its profile along with (b) velocity streamlines and (c) pressure as air travels through the duct.

The geometry with the lowest induced drag had a very unique design with a large inlet and a narrow choke point near the outlet as shown in Figure 3.6. This caused a large region within the duct to recirculate air but did also channel air through the choke point and increase the velocity of the air as it was exhausted.

3.4.3 Bayesian Optimisation

The three designs tested during the Bayesian optimisation process did not improve on the original design, with one design performing much worse than the rest. The second design produced during the Bayesian optimisation stage increased the induced drag by 1.74% when compared to the original novel ducted winglet.

The designs created through the Bayesian optimisation process are identifiable through the reduced radii on the curves. While the thinking beforehand among those involved in the project was to keep the duct as smooth as possible, this preconception did not influence the optimiser. The designs simulated during this stage of the process are seen in Figure 3.8. The first design has a narrow inlet and outlet, with a large chamber at the midpoint. The streamlines in Figure 3.8b show that at the inlet and the outlet, the air travels smoothly through the duct. The velocity greatly increases as it leaves the duct




Figure 3.7: Results from the Bayesian optimisation added to the evolutionary algorithm results with the red line serving as a baseline set by the original design.

through the narrow point. Between the inlet and the outlet, the chamber of the duct causes a large region of air to recirculate much like the the original design. The pressure in the chamber is also increased compared the the rest of the duct. The second design has a more uniformly shaped duct with a slight expansion near the outlet. Unlike the previous design, the air appears to travel through the duct without an area of recirculation, though Figure 3.8e shows that the curvature of the duct causes the velocity of the air to greatly reduce. The third and final design created during the Bayesian optimisation process has a similar shape to the second design with two key differences: the inlet area is much larger and the duct cross-sectional area increases only at the outlet. The third design has an area of recirculation at the inlet, a theme among designs. Before the midpoint of the duct, the velocity of the air increases significantly before then decreasing again at the outlet. The angle at which the air exits the duct in the third design is unique. The air travels primarily

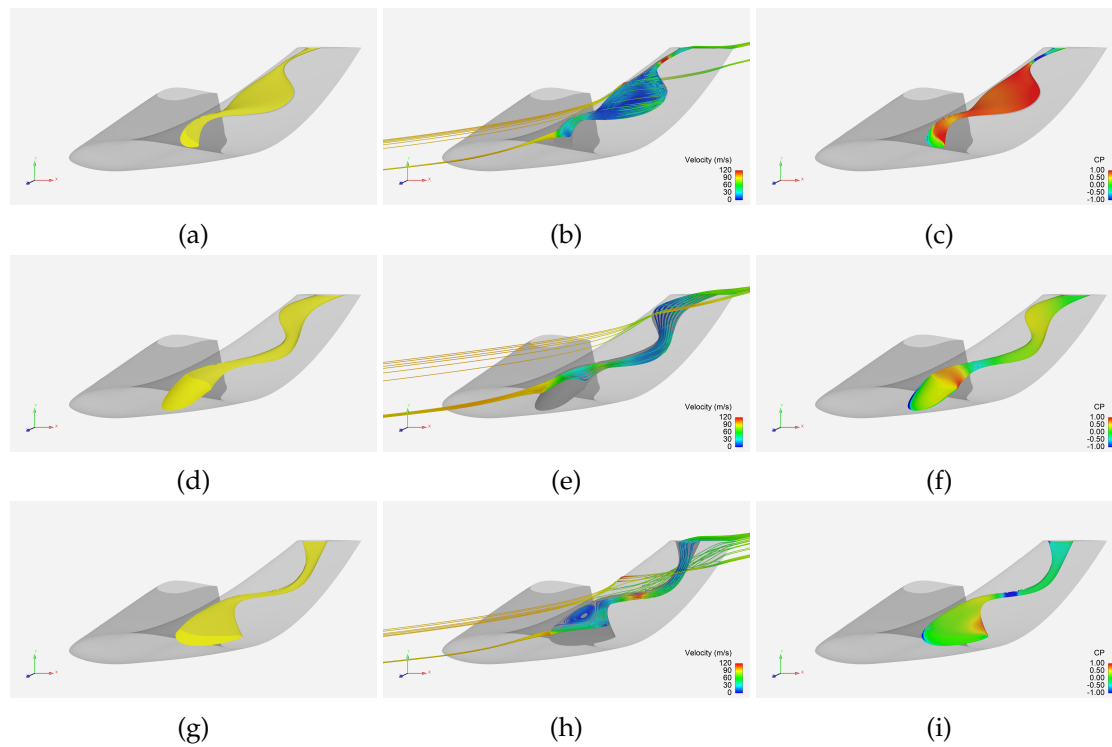


Figure 3.8: The three simulated geometries from the Bayesian optimisation process. The ducts are (a, d, g) highlighted in yellow to better show its profile along with (b, e, h) velocity streamlines and (c, f, i) pressure as air travels through the duct.

in the $Y+$ direction before rejoining the freestream air which travels at an angle of 5 degrees relative to the $X+$ direction. Abrupt changes in the direction of air are not considered by the optimiser but would be considered by an aerodynamicist. The novelty of the designs and the limited number of designs at this stage of the optimisation process are discussed in Section 3.6.3.

3.5 Reduced Parameter Optimisation

The 18 parameters that create the shape of the duct in addition to the long simulation time lead to a slow optimisation lifecycle. In order to speed up this process, reducing the

parameters and simulation time are essential. In the case of the novel ducted winglet, reducing the duct to an entry and an exit would reduce the parameter count and allow the optimisation algorithms to find a more optimal solution sooner. Eliminating the ducts in the middle of the duct would reduce the number of profiles to 2 and reduce the parameter count from 18 to 6. The y-axis shift can also be eliminated as a parameter further reducing the total parameter count to 4. The disadvantage of reducing the parameter count is the diminished control over the shape of the duct.

3.5.1 Updated Method and Materials

With a reduction in parameters, some of the methodology needed to be updated. The design of experiments is swapped with an LHS method and updates to the designs constraints is added at the Bayesian optimisation stage. At the same time, a reduction in the number of iterations performed per simulation is proposed to reduce the overall optimisation time.

Reducing Simulation Iteration Count

A reduction in simulation iteration count could be achieved in addition to reducing the parameters. In Section 2.3.3, it was stated that the simulations were run for a total of 72 hours. This was done to ensure the simulations had sufficiently converged and was measured through the residual. Once the simulations had completed their run times, the lift, drag, and lift-to-drag ratios could be compared among the different geometries. The rank order of the novel ducted winglet, non-ducted winglet, and reference blade were based on running the simulations to completions. Suppose that the simulations were run by one minute less, would the rank order be different among the designs? It is unlikely that a 0.02% reduction in simulation time would alter the ranking. Nor, would reducing

Table 3.3: Nondimensionalised drag results at various iteration intervals for the five designs tested during the design of experiments. The rank order of each iteration count is given in brackets.

| Iterations | Design B | Design C | Design D | Design E | Design F |
|------------|-------------|-------------|-------------|-------------|-------------|
| 30000 | 0.05896 (2) | 0.05876 (1) | 0.05968 (5) | 0.05924 (4) | 0.05903 (3) |
| 20000 | 0.05978 (2) | 0.05953 (1) | 0.06037 (5) | 0.06004 (4) | 0.05980 (3) |
| 10000 | 0.06184 (2) | 0.06170 (1) | 0.06263 (5) | 0.06224 (4) | 0.06191 (3) |
| 5000 | 0.05723 (2) | 0.05680 (1) | 0.05814 (5) | 0.05751 (4) | 0.05750 (3) |
| 2500 | 0.06045 (3) | 0.05960 (1) | 0.06052 (5) | 0.06047 (4) | 0.05986 (2) |

the wall clock limit by two minutes. Continuously applying this logic, however, there will eventually be a point at which the ranking will change. There must therefore exist some minimum wall clock limit or iteration count at which it can be certain that the ranking remains the same. The five designs tested in the design of experiments can be used to determine at which iteration count the ranking changes among them. In Table 3.3, the time-averaged drag values of the five designs is taken at different iteration counts and the rankings are compared. Ranking among the designs remains constant until 2500 iterations. Importantly, the best design simulated at 30000 iterations remains the best at 5000 iterations. Likewise, the worst design remains the worst through at all the iteration intervals. As mentioned in the last chapter, the drag values are averaged to account for small differences as the lift and drag values converge. In Figure 3.9, the drag values for each of the five designs tested are shown as they converge. Note that the designs with higher drag values are consistently above those with lower drag values across the range of simulation iterations. Using the data provided in Table 3.3, the iteration count can be reduced to 5000 for subsequent simulations. This method is appropriate for the initial stage of optimisation. Given that designs at this points are farther apart, ranking is appropriate. Once designs are more converged, iterations can be increased again for higher fidelity. This reduces overall optimisation time.




Figure 3.9: The drag values for five different designs across the entire range of simulation iteration counts.

Constrained Latin Hypercube Sampling

Previously mentioned in Section 3.3.1 was the interdependency of chord size and eccentricity on duct size which lead to a design and parameter space constraint. The combination of both values could lead to designs that were not feasible as the duct would exceed the walls of the winglet even if both values were within their respective ranges.

Applying the design constraint to the LHS method proved difficult because it necessitated a violation of the method. Removing an area of the parameter space will lead to an equally-sized area within the parameter space to not be sampled due to the orthogonal nature of LHS. This means that for every design that is invalid, another valid design won't be tested. To avoid this, the boundaries of the hypercube could be altered. The downside to this is that the size of each mini-hypercube is different. Again, this could be solved by

further altering the boundaries of each smaller hypercube to be equally sized. Even with these two fixes, this still causes a violation of the Latin hypercube method because the distribution of samples over the input variables' range is no longer flat.

Rather than complicating the process further, standard LHS was employed and results that exceed the winglet walls were not considered. The LHS code was run 10 000 times and the run with the most amount of viable samples was used. The seed was started at 0 and incremented by 1 after each iteration. In total 21 samples were identified, though the interdependency between the inlet and the outlet caused issues when building the geometry so it was discarded. The duct profile generated by the one discarded design went beyond the walls of the winglet. 20 geometries were then simulated and the results analysed using Trefftz plane analysis.

Bayesian Optimisation

Bayesian optimisation was revisited with the reduced number of parameters. The increase in the number of samples along with less parameters gave the model a better chance to find optima. One of the challenges with using Bayesian optimisation at this stage was the constraints placed upon the designs in the Latin hypercube stage were not being respected by the model. Of course, the model aimed to maximise the information gain and wanted to explore areas of the parameter space without samples. Despite the best efforts to guide the model to keep the duct within the winglet walls, it proved increasingly difficult as more samples were added. Instead, it was decided to let the duct extrusion exceed the winglet which created holes on the walls of the winglet. The effects of this were to be measured in the same way as before using Trefftz plane analysis.

When introducing designs that had additional holes some modifications had to be made.

During the CAD process, the holes created an infinitesimally thin edge which connected the inside of the duct to the outside; the edge was flattened to 0.01 mm to separate the two regions. In one instance, the addition of this edge added a surface when imported to FLITE3D so appropriate changes had to be made to the input files. The effects of human intervention on this stage of the optimisation process are discussed in Section 3.6.4.

3.5.2 Reduced Parameter Count Results

The twenty samples created during the constrained LHS allowed the design space to be sample better than before, when more parameters were used. Three of the designs had a lower induced drag, with the best design creating 0.64% less induced drag than the original novel ducted winglet design. Two design created a similar amount of induced drag with the rest of the designs creating more induced drag. Four of the designs had a lower overall drag compared to the original design with the other 16 creating more drag. The design which had the lowest induced drag also had the lowest overall drag, an improvement of 0.34% compared to the original novel ducted winglet design. The designs are plotted in Figure 3.10 using their induced drag and overall drag values and divided into quadrants using the original design to separate the quadrants. Only one design had lower induced and overall drag. Two designs had lower induced drag but higher overall drag while three designs had lower overall drag but higher induced drag. The remaining fourteen designs had higher induced and overall drag.

The Bayesian optimisation process did produce one design that had a lower induced drag, and produced two designs that improved the overall drag, with the best performing design improving the overall drag by 0.71%. The best performing design with respect to induced drag had an improvement of 0.20%. This stage of the optimisation process did not create any designs that improved the induced drag and overall drag. The remaining




Figure 3.10: Results from the LHS and Bayesian optimisation proces with the red line serving as a baseline set by the original design for both overall nondimensionalised drag and induced drag (N).

seven designs created higher induced drag and overall drag compared to the original novel ducted winglet design.

Of the thirty designs created during the reduced parameter stage of the optimisation process, only one design managed to lower the induced drag and the overall drag simultaneously. Five designs lowered the induced drag, of which four had a higher overall drag compared to the original design. Six designs lowered the overall drag, five of which increased induced drag. The remaining twenty designs increased the induced drag and total drag.

The best and worst performing designs from the LHS stage and the Bayesian optimisation stage are compared in Figure 3.11. The lowest induced drag and overall drag designs are compared to the highest induced drag and overall drag designs.

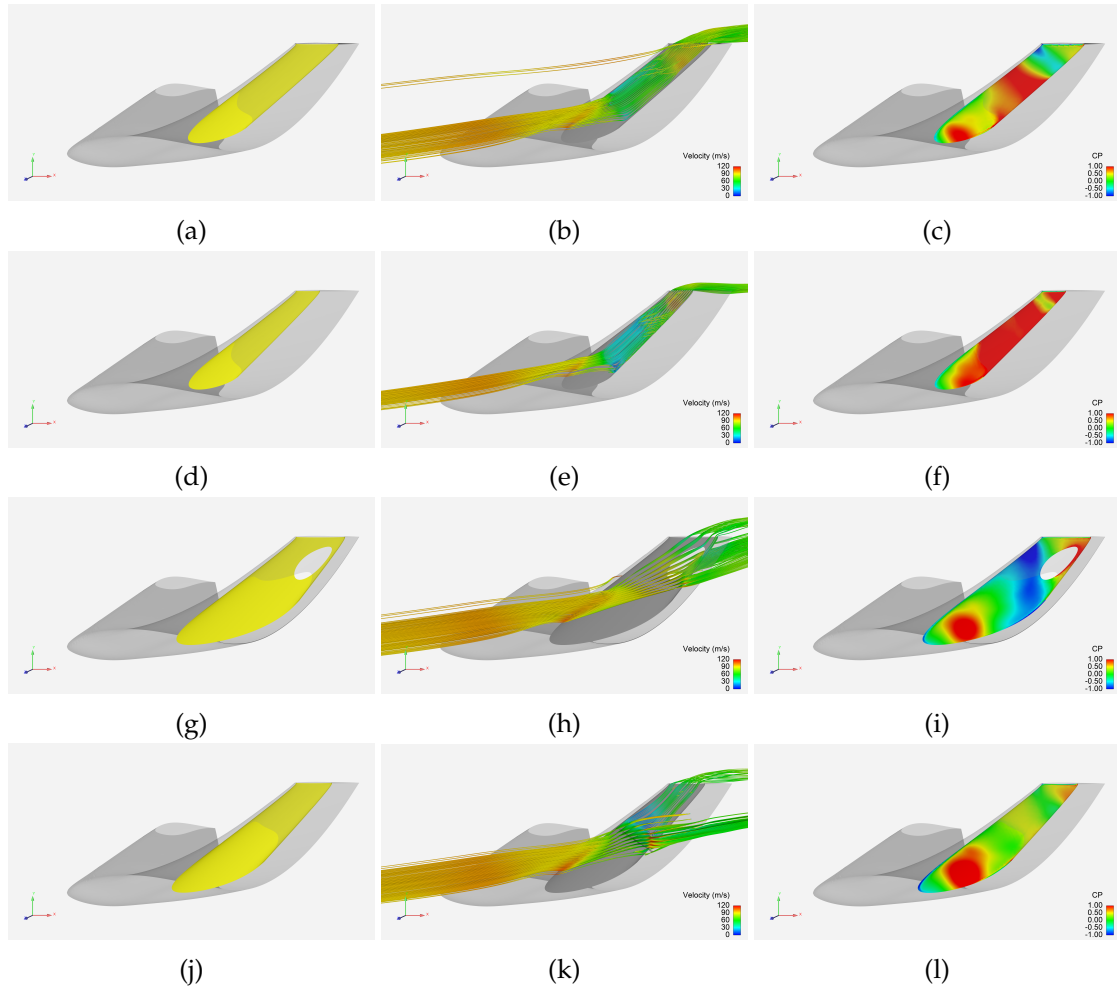


Figure 3.11: The first row (a, b, c) shows the design with the lowest induced drag. The second row (d, e, f) shows the design with the lowest overall drag. The third row (g, h, i) shows the design with the highest induced drag and the fourth row (j, k, l) shows the design with the highest overall drag. The ducts (a, d, g, j) are highlighted in yellow to better show its profile along with (b, e, h, k) velocity streamlines and (c, f, i, l) pressure as air travels through the duct.

3.6 Discussion

Different optimisation methods were applied to the novel ducted winglet with varying levels of success. While a direct comparison among the optimisation methods is perhaps not entirely possible as the later methods were able to use data from previous methods, some key strengths and weaknesses are identified. Each method produced a unique set of designs inherent to the underlying method used to create new designs. The challenges associated with each method are discussed in this section. The merits of changing the number of parameters and the effects on the optimisation process are also discussed.

3.6.1 Design of Experiments

The five designs created during this stage provided valuable insight into the effects of changing the duct on the induced drag. Preconceptions going into the first optimisation study were to keep the duct as smooth as possible without excessive curvature on along the leading and trailing edges of the duct and to angle the duct such that the freestream air could enter the duct as smoothly as possible, avoiding an area of recirculation as seen with the original novel ducted winglet design. One of the interesting findings was that ducts with a smaller cross-sectional area performed better with respect to reducing the induced drag than those with a larger cross-sectional area. This is consistent with findings which compared the non-ducted winglet to the novel ducted winglet and from the theory which discussed that increasing the surface area leads to higher skin friction drag. The most interesting design to reduce the induced drag was Design F, which decreased the chord at the midpoint along the duct. The overall drag was also reduced compared to the original design but only marginally.

LHS was considered at this stage the challenge of constraining the ducts to fit inside the

winglet would have remained. It was more time-effective to create designs that were valid than to sample areas of the design space that were not valid.

3.6.2 Evolutionary Algorithm

Twelve new designs were created using an EA during the second part of the optimisation study, using the 5 designs from the design of experiments plus the original design as the starting point. The advantage of using an EA is that algorithm is unaware of the preconceptions relating to the specific design challenge, so design are not limited by potentially obstructive design principles. However, the designs created using the EA suffered from a design flaw that hampered one of designs from the Design of Experiments. Given the nature of the crossover method, which took some duct profiles from one parent and the other duct profiles from another parent, the resulting children had either an expansion or reduction in cross-sectional area. As seen in Figure 3.5, changing the duct profiles did not reduce the induced drag. Changing the method of crossing over would have been possible, but creating additional designs took time and resulting ducts did not always stay within the confines of the outer wall of the winglet. A mutation only EA was also suggested by an expert in the field, though again this would be a time consuming manner of optimisation as manual intervention and CAD processing was still a significant portion of the design process.

3.6.3 Bayesian Optimisation

The Bayesian optimisation stage used the 12 EA designs, 5 design of experiments designs, and the original design in the dataset. The three designs created during this stage of the optimisation study were only a small sample of the total number of designs that had been attempted to be made. The parameters prescribed by the Bayesian optimiser would often

lead to the holes in the walls of the winglet as the duct would either be too large or the duct would take a path that would guide it beyond the walls of the winglet. It was difficult to assess whether geometry was valid when new values were produced simply because of the complex interdependencies between all three duct variables. Additionally, even if two duct profiles sat within the walls of the winglet, it was still possible for the region between the two profiles to emerge outside of the walls. Finding out if a geometry was valid was a time consuming process that ultimately led to only three designs being tested using this method. The increases in the induced drag also did not suggest that continuing using the Bayesian optimiser would be an effective use of resources. The exploratory nature of the optimiser through the acquisition function, along with the high number of parameters and long evaluation times also hindered progress.

3.6.4 Reduced Parameters

The reduction in simulation time helped reduce the turnaround time for a new design from ≈ 74 hours to ≈ 25 hours. This was beneficial in terms of being able to create more designs for the optimisation study. One of the challenges with reducing the simulation iteration count was that the lift and drag values output by FLITE3D were higher than those which ran for 72 hours. This made comparing the results from different stages of the study directly impossible, so percentage improvements compared to the original novel ducted winglet design were used instead.

Choosing a manner to constrain the LHS method was the subject of many internal discussions. The constraints invalidated the method and testing which method is appropriate could easily be its own chapter. Having opted for a method which was simpler to implement, the induced drag and total drag results bring some level of confidence that the despite the constraints, the method still manages to sample the design space adequately.

The decision to let the duct create additional holes in the sides of the winglet led to some interesting designs. This decision could have been applied to the designs with 18 parameters, though results from the reduced parameter study indicate that it would not have reduced the induced drag. When the duct did go beyond the walls of the winglet, additional steps were needed to allow the geometry to be meshed in FLITE3D. This highlights the need for human intervention during the design process, whether its during the parameterisation stage or the CAD stage.

3.7 Conclusion

The aim of this chapter was to optimise the novel ducted winglet. The EA was first tested on a 2D case to demonstrate that it could optimise for the lift-to-drag ratio. The EA was then implemented on a fully parameterised novel ducted winglet. A design of experiments was used to set an initial population for the EA, with one design performing better than the original design. Challenges arose using the EA on the fully parameterised model given the number of parameters, the computation time, and the crossover method which led to irregularities in the flow within the duct. A Bayesian optimiser was then used as an alternative to the EA, though it also struggled to optimise the novel ducted winglet. A decision was then made to reduce the number of parameters in the duct from eighteen to four, and to reduce the number of simulation iterations. LHS was applied to create twenty novel ducted winglets with reduced parameters. These twenty designs were then used in another round of Bayesian optimisation where ten more designs were made. Despite the reduced parameters, the Bayesian optimiser still struggled to find designs that lowered the induced drag. The effects of lowering the simulation iteration count is explored in the next chapter.

Chapter 4

Effects of Uncertainty on Reducing Simulation Iterations

The idea of reducing the simulation iteration count is proposed in the previous chapter. Two studies are conducted to gauge the effects of increased uncertainty in the results on an evolutionary algorithm by changing the order in which results are ranked and changing the fidelity of the governing equation. The travelling salesman problem is used to test the effects of rank change while Taylor series approximations of benchmark functions are used to test the effect of equation fidelity. An additional study is done on 2D aerofoils to see the effects of reducing the number of elements and the residual tolerance. The changes to the rank order and equation fidelity are run numerous times with evidence that small changes to either have minimal effect on the success of the optimisation. The ranking of the 2D aerofoils also remains constant for most of the element count and residuals. These results support the decision to reduce the simulation iteration count.

4.1 Introduction

In the previous chapter a simple evolutionary algorithm was used to optimise a 2D aerofoil and the 3D novel ducted winglet. The algorithm demonstrated it was able to optimise the lift-to-drag ratio of the 2D aerofoils but struggled to optimise the novel ducted winglet. The increased number of parameters and increased simulation time contributed to optimisation difficulties. Another difference between the 2D case and the 3D case are the governing equations. While all varying geometries were tested under the same conditions, it is not possible to determine if there was substantial error introduced during the simulation which could affect the ranking of the geometries. For example, it is possible that the geometry which performs best in real life was ranked second-best or third-best (or ostensibly worse) in the simulations. If the optimal design is not ranked first, how does this affect the optimisation success and optimisation time? Additionally, could the use of the Navier-Stokes equations introduce false optima when evaluating the novel ducted winglet geometries?

4.1.1 Evolutionary Algorithm Process Affected by Uncertainty

One of the fundamental pillars/steps of evolutionary algorithms is the selection process. The best individuals are chosen from the population to be bred while the others are discarded. In analytical fitness functions, the individuals will be ranked correctly; of course, this does not mean that the top-ranked individual from the population is the optimal solution. However, in fitness functions based on partial differential equations, the ranking is only as good as the how well the partial differential equations represent the analytical equation they represent. The individuals that get selected in the analytical equation situation might not be selected in the partial differential equation situation. This would lead to sub-optimal individuals getting to the next generations. This would

theoretically affect the algorithms ability find the optimum and how long it takes to find the optimum.

During the optimisation process in the previous chapter, designs were ranked based on their induced drag. As previously mentioned, the order of the ranking plays a part in the success and length of the optimisation process. However, it does not necessarily mean that the optimisation process will definitely not succeed. Testing this hypothesis using CFD is difficult because it takes hours or days to generate new generations. The process can be mimicked on a simplified case to better understand the effect that rank order has on optimisation. A simple evolutionary algorithm created by Robert Kübler [87] provided a good starting point as the rank of individuals could be changed easily. The code was also already set up for a benchmark case.

4.2 Background

Two principle methods are used in this chapter to test the effect of rank on optimisation success and the effect of governing equation fidelity on optimisation success: the travelling salesman problem and Taylor series. The former is a benchmark problem used to test optimisation methods and the latter is a mathematical approximation of non-linear functions. A brief mathematical description of each method is given in this section. The travelling salesman problem is used to measure the effect of changing the rank order in which solutions are ordered while the Taylor series is used to vary the approximation of benchmark functions to mimic the varying fidelity of governing equations.

4.2.1 Travelling Salesman Problem

One of the classic applied NP-hard benchmark problems is the travelling salesman problem (TSP). Imagine a salesman that must travel to various cities to sell his goods. He wants to visit all the cities while travelling the shortest possible distance though he is constrained by only being able to visit each city once. He also has to return home to the city where he started from. The problem is trivial for a low number of cities as the number of travel permutations between cities is low. However, as the number of cities increases, so does the number of travel permutations. In a strictly brute force algorithm, the time complexity is $O(n!)$ where n is the number of cities.

The formulation used by Miller et al. [88] of the TSP sets the starting city i , arrival city j , the distance d_{ij} , and the boolean x_{ij} which is set to 1 if travel occurs between the starting city and ending city. An additional continuous variable u is used to represent the position of the salesman. The function to minimise is

$$\sum_{i,j}^n d_{ij}x_{ij} \quad (4.1)$$

which is subject to

$$x_{ij} \in \{0, 1\} \quad i, j = 1, \dots, n; \quad (4.2)$$

$$\sum_{j=1}^n x_{ij} = 1 \quad i = 1, \dots, n; \quad (4.3)$$

$$\sum_{j=1}^n x_{ji} = 1 \quad i = 1, \dots, n; \quad (4.4)$$

$$u_i - u_j + (n - 1)x_{ij} \geq n - 2 \quad (2 \geq i, j \geq n; i \neq j) \quad (4.5)$$

$$1 \geq u_i \geq n - 1 \quad (2 \geq i \geq n) \quad (4.6)$$

The constraint 4.2 sets x_{ij} as a boolean and prevents the salesman from leaving and arriving from the same city. The constraints 4.3 and 4.4 are used to ensure that the salesman visits every city only once. Constraint 4.5 set the arrival city j to be one more than the starting city i while Constraint 4.6 makes sure that each city is unique; the first city is exempt as it is also the final city.

The TSP has been used extensively in EAs to test their ability to find an optimal solution. In this work, the algorithm is only tested against itself and whether or not this specific algorithm is effective at solving the problem is immaterial.

4.2.2 Taylor Series

Smooth and differentiable functions can be represented by the infinite sum of its derivatives about a point. These infinite sums are called Taylor series. The number of terms in the Taylor series determines the degree to which the original function is represented. The more terms that are added to the Taylor series, the more closely the functions is represented. Polynomial functions can be represented completely if the number of terms in the Taylor series is equal to the highest order polynomial. The general formulation for a Taylor series

denotes the point a at which the series is evaluated:

$$f(a) + \frac{f'(a)}{1!}(x-a) + \frac{f''(a)}{2!}(x-a)^2 + \frac{f'''(a)}{3!}(x-a)^3 + \dots \quad (4.7)$$

and is rewritten to

$$\sum_{n=0}^{\text{inf}} \frac{f^{(n)}(a)}{n!}(x-a)^n \quad (4.8)$$

in its general form. $f^{(n)}(a)$ determines how many terms are added to the Taylor series. The first term in Equation 4.7 becomes $f(a)$ as f^0 is simply the original function with $n!$ and $(x-a)^0$ both equal to 1.

4.3 Method and Materials

Two EAs were used to test the effects of rank order and equation fidelity. One used the TSP to compare how often and how fast the algorithm was able to find an optimum solution under increasing changes in rank, while the other used Taylor series. Both of the algorithms used to test the effects of rank change and equation fidelity using the TSP and Taylor series were run multiple times to account for randomness within the algorithms. The EAs were run on the same high performance cluster used to simulate the novel ducted winglet, with each simulation assigned to a single CPU.

In addition to the EAs, three 2D aerofoils are compared across a range of element counts and then for different levels of residual tolerance. The 2D aerofoil cases are tested under similar conditions to the cases seen in Chapter 3.

4.3.1 Travelling Salesman Problem

An EA was run on an instance of a TSP 10000 times without any changes to the code to see how often the evolutionary algorithm was able to find the optimal value. The seed was set at 0 for the first run and then incremented by 1 until all the runs had been completed. A total of 5735 returned solutions had a result within 0.5% of the known-optimal value. The 0.5% leeway was used as cutoff to include any results which may have been missed due to machine precision. The seeds of the successful runs were stored and used in subsequent runs to show the effects of changing the rank. The rank change was determined randomly by assigning a value from 1 to the maximum rank change value. The likelihood of changing the rank was set by the probability limit which ranged from 0 to 100% in increments of 10%. The range of values taken on by both the maximum rank change and the rank change probability are seen in Table 4.1. The rank was changed during the *Selection* phase of the EA set out in Algorithm 1. The code cycled through the population from highest to lowest rank when evaluated against the fitness function. For each individual of the population, a check done to see if the rank should change. If it should change, the rank of that individual was lowered in the range of 1 to the maximum change in rank. The steps are outlined in Algorithm 4.

Algorithm 4 Rank change in EA.

- 1: **for** Each Individual in Population **do**
 - 2: **if** Rank Change = True **then**
 - 3: Determine rank change value from 1 to maximum rank change.
 - 4: Lower rank by rank change value.
 - 5: **end if**
 - 6: **end for**
-

Table 4.1: The increments of the maximum changes in rank value as well as the probability that a change in rank would occur. The probabilities increase from 0 to 100% in steps of 10%.

| | Maximum Rank Change | Rank Change Probability (%) |
|-----------------|-----------------------|-----------------------------|
| Range of Values | 5, 15, 25, 35, 45, 55 | 0, 10, . . . , 90, 100 |







Figure 4.1: All runs of the TSP which found a solution within 0.5% of the best found solution.




(a) 10% rank change probability.




(b) 20% rank change probability.




(c) 30% rank change probability.



(d) 40% rank change probability.




(e) 50% rank change probability.




(f) 60% rank change probability.


Figure 4.2: All runs of the TSP with an increasing probability that a change in rank would occur, from 10% to 60%.




(a) 70% rank change probability.



(b) 80% rank change probability.



(c) 90% rank change probability.



(d) 100% rank change probability.

Figure 4.3: All runs of the TSP with an increasing probability that a change in rank would occur, from 70% to 100%.

Table 4.2: 1D benchmark functions used.

| Name | Function | Bounds | Taylor Series Terms Tested |
|------|--|--------------------------|----------------------------|
| BM1 | $f(x) = \sin(x) + \sin(\frac{10}{3}x)$ | [-2.5, 7.5] | 51-55 |
| BM2 | $f(x) = -x \sin(x)$ | [0, 10] | 15, 17, 19, 20, 21 |
| BM3 | $f(x) = 2 \cos(x) + \cos(2x)$ | $[-\frac{\pi}{2}, 2\pi]$ | 19, 20, 23, 24, 25 |
| BM4 | $f(x) = \sin^3(x) + \cos^3(x)$ | [0, 2 π] | 30-35 |

4.3.2 Effect of Equation Fidelity using Taylor Series

Four 1D multimodal benchmark functions were used to test the effects of equation fidelity on optimisation success. These multimodal benchmark functions have been used in previous literature [89]. The functions are listed in Table 4.2 along with the input bounds of the function and the number of Taylor series terms used. The functions, along with their Taylor series relatives are graphed in Figure 4.4 through 4.8 inclusive. The functions were chosen as they each contained multiple sub-optima, a pitfall of many EAs. The functions also had their optimal values located at various distances from the evaluation point, which in this case was always set to 1.

The EA was run 10000 times using the each Taylor series. The number of times that the optimum was found to $\pm 1\%$ was recorded as well as the lowest generation at which the optimum was found. The number of generations was limited to 100 to save on computational resources. An additional run for each Taylor series was run using the original function to compare the effects of using a Taylor series approximation on the number of successful runs and the generations to optimum count.

The given functions were approximated to Taylor series using a script provided by Bernardo Costa [90].

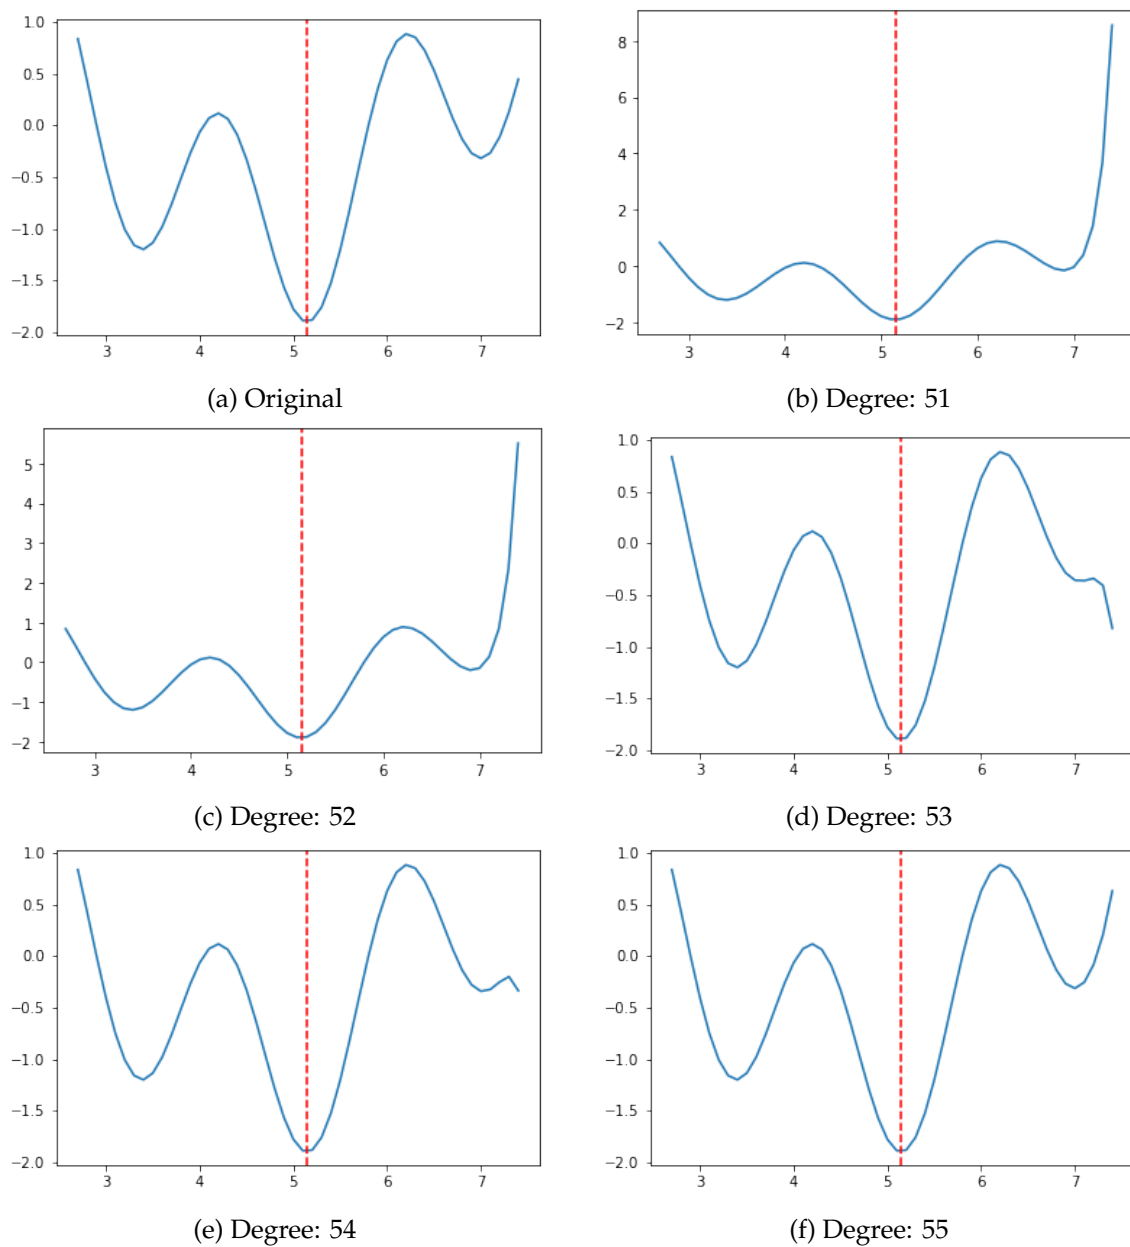


Figure 4.4: Benchmark 1, with the red dotted line indicating the optimal non-dimensional value. The degree denotes the degree of the Taylor polynomial.

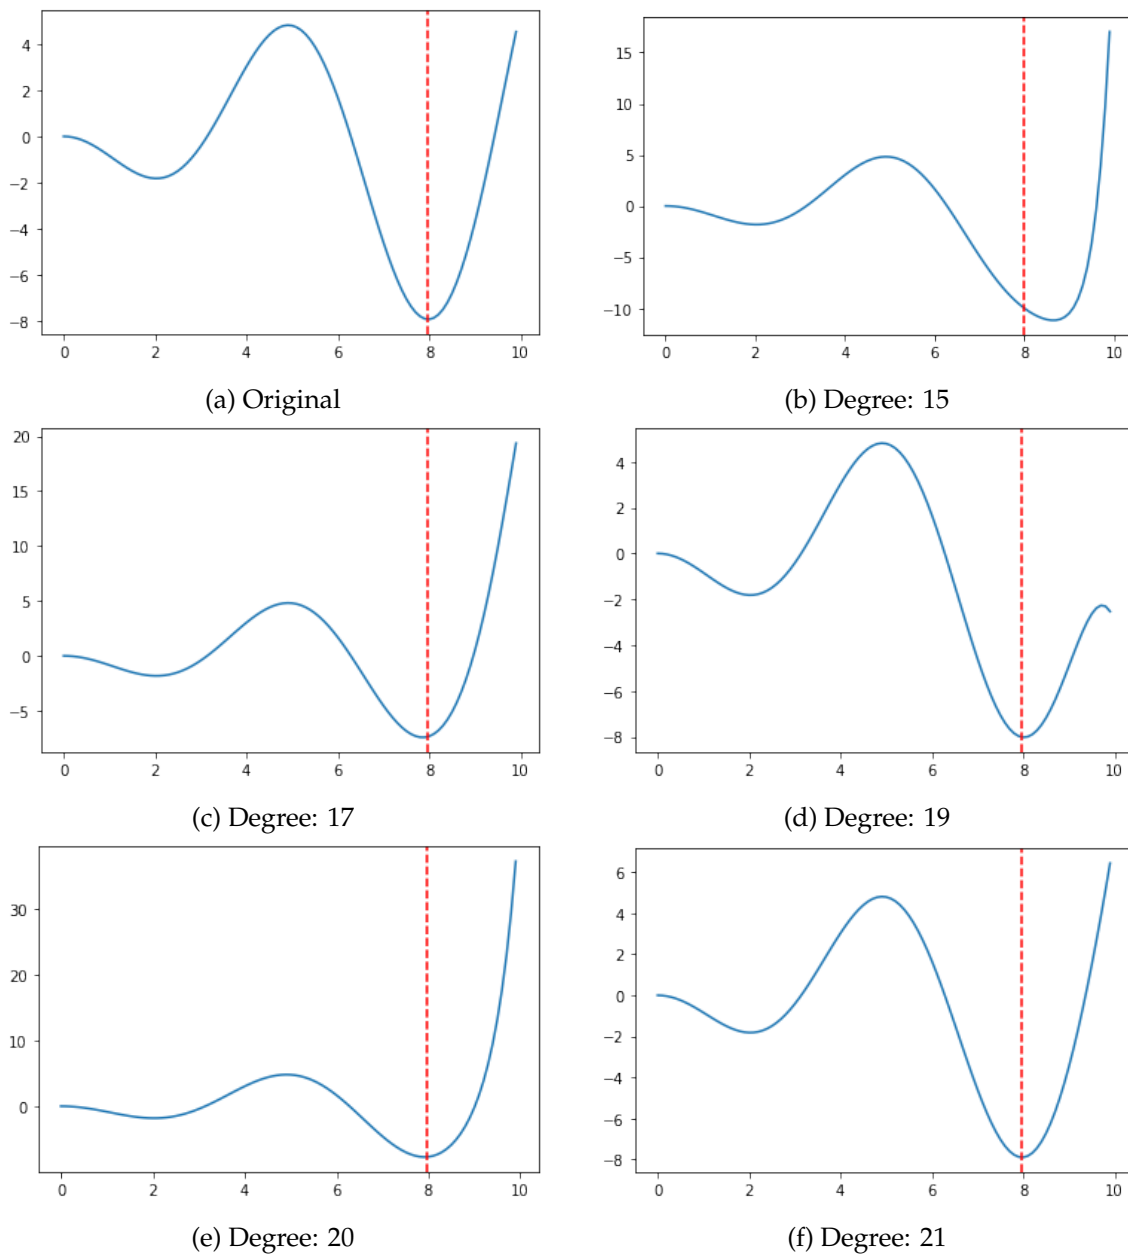


Figure 4.5: Benchmark 2, with the red dotted line indicating the optimal non-dimensional value. The degree denotes the degree of the Taylor polynomial.

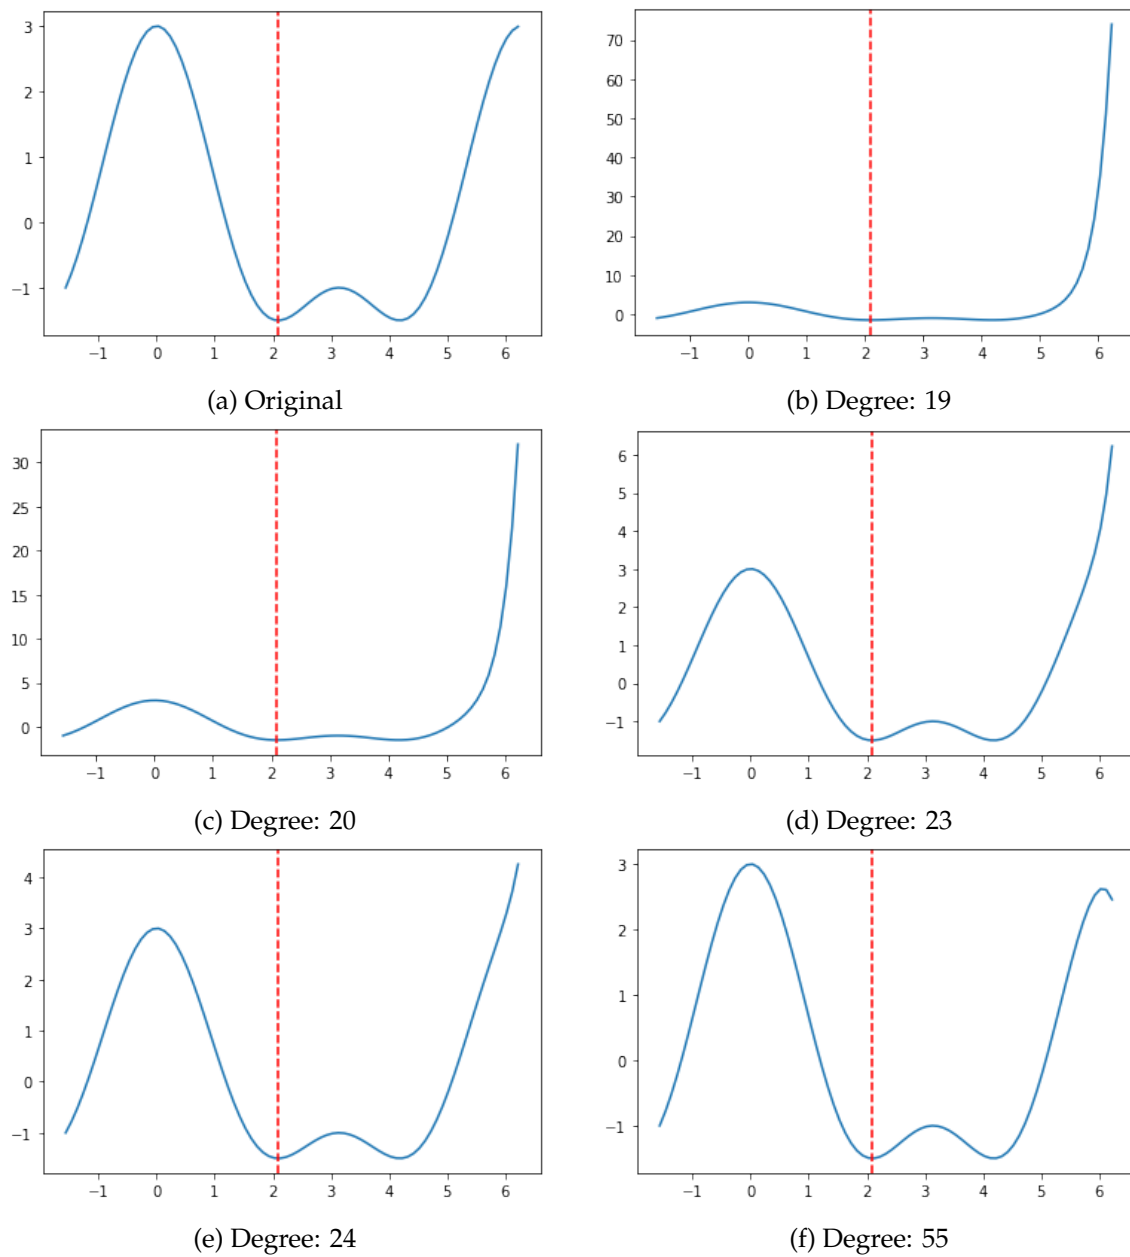


Figure 4.6: Benchmark 3, with the red dotted line indicating the optimal non-dimensional value. The degree denotes the degree of the Taylor polynomial.

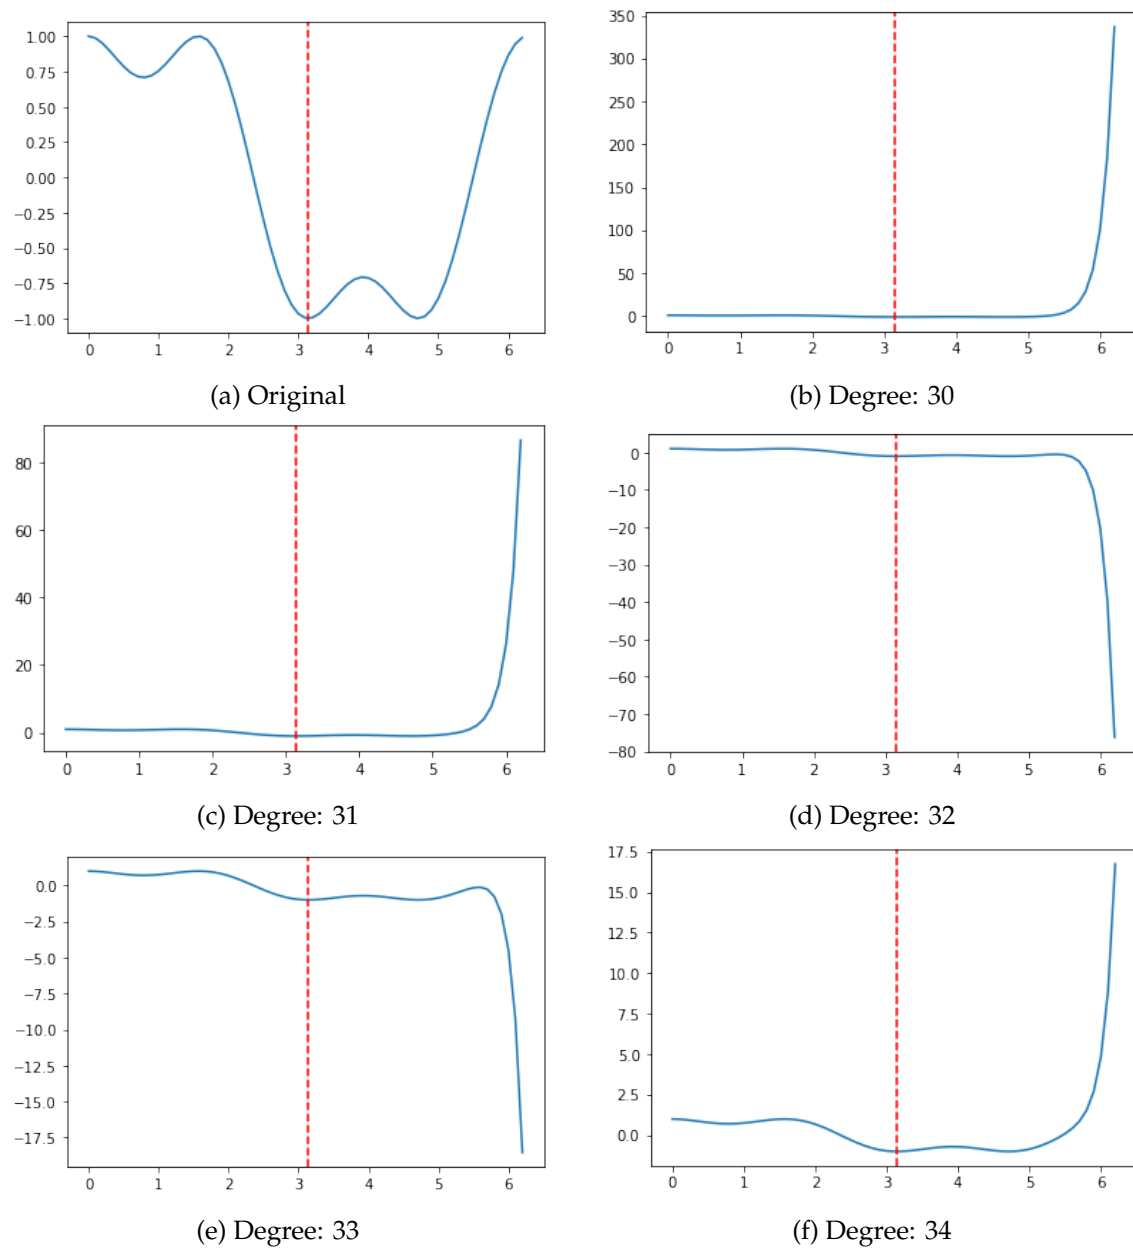


Figure 4.7: Benchmark 4, with the red dotted line indicating the optimal non-dimensional value. The degree denotes the degree of the Taylor polynomial.




Figure 4.8: Benchmark 4 continued, with the red dotted line indicating the optimal non-dimensional value. The degree denotes the degree of the Taylor polynomial.

4.3.3 2D CFD Simulations

To test the assumptions made in the previous chapter, the lift-to-drag ration of three 2D aerofoils are computed and compared. First, a mesh convergence study is done to see if the rank remains constant across the range of element counts. Second, the aerofoils' lift-to-drag ratios are compared as the residual tolerance is decreased. Both of these studies mimic the method used previously to reduce the simulation iteration count.

The simulations were run using the 2D version of FLITE as described in Chapter 2. The hyperparameters were matched where appropriate to those found in Table 2.2 and Table 2.3. Importantly, the simulations were run as steady, compressible, and turbulent cases. The aerofoils were chosen based on their prevalence in literature [91, 92, 93]. The aerofoils were standardised to a length of 1 and were located in a circular domain with a radius of 13.25. The aerofoil was defined by 174 nodes while the domain was defined by 28 points. Mesh refinement was defined by parameters listed in Table 4.3 which contained the x and y coordinates, the prescribed spacing, growth rate, and radius of effect. All mesh refinements took the form of point sources. Mesh refinements were also controlled

with a variable which dictated element refinement; the values ranged from 0.2 to 1.2 with increasing element refinement leading to higher element counts. Unlike the 3D cases, the 2D aerofoils were not split into subdomains during pre-processing for parallel computing; the computational time for the 2D cases was significantly lower than the 3D cases. The lift and drag values were calculated in the same manner as described in Section 2.3.3. All cases were run at 5 degrees angle of attack, like in the 3D cases.

4.4 Results

The effects of rank change and equation fidelity on the success of optimisations are detailed below. Changes in rank began to affect optimisation success only once the changes became very frequent and much larger. The effect of equation fidelity depended mostly on how close the approximated optimum was to the real function.

4.4.1 Effects of Rank Change on Optimisation Success

Generally, any perturbation to the algorithm caused a reduction in the number of runs which were able to find the optimal TSP route. This can be seen in Tables 4.4 and 4.5. Figure 4.9 also shows the ratio of runs which found the optimal value to the total number of runs as the maximum rank change and change probability increase. Low rank change probabilities did not affect the algorithm's ability to find the optimum beyond the initial drop in successful runs. Likewise, low maximums of change in rank also did not affect the algorithm's ability to find the optimum. Only at higher rank change probabilities for a maximum rank change of 35 and 45 did the number of successful runs begin to decrease. At a maximum rank change of 55, the decrease in number of successful runs is more evident as it begins at lower rank change probabilities and continues to the higher probabilities.

Table 4.3: The mesh refinement sources used for the 2D aerofoils. The mesh refinements were needed to capture some of the changes in flow over the aerofoil.

| Source | X | Y | Spacing | Growth Rate | Radius |
|--------|---------|---------|---------|-------------|--------|
| 1 | -0.0640 | -0.0770 | 0.0300 | 0.0100 | 0.0100 |
| 2 | 0.0375 | -0.0100 | 0.0300 | 0.0200 | 0.0150 |
| 3 | 0.900 | -0.0170 | 0.0150 | 0.0200 | 0.0150 |
| 4 | 0.0206 | 0.0160 | 0.0150 | 0.0100 | 0.0100 |
| 5 | -0.0178 | -0.0740 | 0.0150 | 0.0100 | 0.0100 |
| 6 | 0.900 | 0.0176 | 0.0150 | 0.0100 | 0.0100 |
| 7 | 0.677 | -0.0206 | 0.0150 | 0.0100 | 0.0100 |
| 8 | 1.215 | -0.118 | 0.0150 | 0.0100 | 0.0100 |
| 9 | -0.0630 | -0.0790 | 0.0300 | 0.0100 | 0.100 |
| 10 | 0.0206 | 0.0160 | 0.0300 | 0.0100 | 0.100 |
| 11 | 0.0375 | -0.0100 | 0.0300 | 0.0100 | 0.100 |
| 12 | 0.900 | 0.000 | 0.0300 | 0.0100 | 0.100 |
| 13 | 0.900 | -0.0170 | 0.0300 | 0.0100 | 0.100 |
| 14 | 1.215 | -0.118 | 0.0300 | 0.0100 | 0.100 |
| 15 | -0.0630 | -0.0790 | 0.0300 | 0.0100 | 0.500 |
| 16 | 0.0206 | 0.0160 | 0.0300 | 0.0100 | 0.500 |
| 17 | 0.0375 | -0.0100 | 0.0300 | 0.0100 | 0.500 |
| 18 | 0.900 | 0.0176 | 0.0300 | 0.0100 | 0.500 |
| 19 | 0.900 | -0.0170 | 0.0300 | 0.0100 | 0.500 |
| 20 | 1.215 | -0.118 | 0.0300 | 0.0100 | 0.500 |
| 21 | 0.0206 | 0.0160 | 0.0300 | 0.0100 | 0.100 |
| 22 | 0.0500 | 0.0360 | 0.0300 | 0.0100 | 0.100 |
| 23 | 0.900 | 0.0176 | 0.0300 | 0.0100 | 0.100 |
| 24 | 0.940 | 0.0100 | 0.0300 | 0.0100 | 0.100 |
| 25 | 1.215 | -0.118 | 0.0300 | 0.0100 | 0.100 |
| 26 | 1.335 | -0.170 | 0.0300 | 0.0100 | 0.100 |

Table 4.4: Total number of runs which successfully found the optimal value to the TSP. Max Rank Change is represented as MRC.

| MRC | Rank Change Probability (%) | | | | | | | | | | |
|-----|-----------------------------|------|------|------|------|------|------|------|------|------|------|
| | 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 | 100 |
| 5 | 5735 | 3578 | 3622 | 3560 | 3541 | 3571 | 3516 | 3583 | 3570 | 3564 | 3547 |
| 15 | 5735 | 3536 | 3510 | 3531 | 3560 | 3529 | 3554 | 3564 | 3530 | 3619 | 3549 |
| 25 | 5735 | 3534 | 3519 | 3573 | 3525 | 3525 | 3533 | 3489 | 3538 | 3594 | 3554 |
| 35 | 5735 | 3464 | 3465 | 3519 | 3507 | 3499 | 3589 | 3468 | 3406 | 3455 | 3399 |
| 45 | 5735 | 3469 | 3505 | 3471 | 3455 | 3472 | 3501 | 3424 | 3402 | 3320 | 3224 |
| 55 | 5735 | 3507 | 3541 | 3464 | 3380 | 3386 | 3353 | 3257 | 3139 | 3061 | 2975 |

Table 4.5: Ratio of runs which successfully found the optimal value to the TSP to the total number of runs. Max Rank Change is represented as MRC.

| MRC | Rank Change Probability (%) | | | | | | | | | | |
|-----|-----------------------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| | 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 | 100 |
| 5 | 1 | 0.624 | 0.632 | 0.621 | 0.617 | 0.623 | 0.613 | 0.625 | 0.622 | 0.621 | 0.618 |
| 15 | 1 | 0.617 | 0.612 | 0.616 | 0.621 | 0.615 | 0.62 | 0.621 | 0.616 | 0.631 | 0.619 |
| 25 | 1 | 0.616 | 0.614 | 0.623 | 0.615 | 0.615 | 0.616 | 0.608 | 0.617 | 0.627 | 0.62 |
| 35 | 1 | 0.604 | 0.604 | 0.614 | 0.612 | 0.61 | 0.626 | 0.605 | 0.594 | 0.602 | 0.593 |
| 45 | 1 | 0.605 | 0.611 | 0.605 | 0.602 | 0.605 | 0.61 | 0.597 | 0.593 | 0.579 | 0.562 |
| 55 | 1 | 0.612 | 0.617 | 0.604 | 0.589 | 0.59 | 0.585 | 0.568 | 0.547 | 0.534 | 0.519 |

The effect of altering the rank change probability can be compared between Figure 4.1, which shows the best fitness of 5735 runs as they approach the optimum, and Figure 4.2, which shows an increase in the probability of a change in rank for a maximum change in rank of 25.

4.4.2 Effects of Equation Fidelity on Optimisation Success

Generally, the EAs run using the Taylor series were able to find the optimum within 1% a similar amount of times as running the EA on the function itself. The Taylor series runs were also able to find the optimum using a similar number of generations. There are two cases which do not follow this trend. The first is in Benchmark 2 for Taylor series orders




Figure 4.9: The percentage of TSP runs which found a solution within 0.5% of the best found value across an increase in rank change probability, separated by the maximum value of a rank change ranging from 5 to 55.

15 and 17. A decrease by two orders of magnitude in number of successful runs is seen and can be attributed most probably to the function's representation in Figures 4.5b and 4.5c. The second is degree 32 and 33 in Benchmark 4 which again can be attributed to the shape of the function.

4.4.3 Effects of Reducing Simulation Iterations in 2D

In the mesh convergence, as seen in Figure 4.10, the aerofoils maintained their rankings aside from the lowest element count. This shows that in regions where the results begin to converge, element count begins to have less effect on the overall rank. As the residual tolerance is decreased, the difference between the aerofoils is reduced, specifically between the 4412 and 6409. The ranking is correct across the range of residual tolerances except at 1.5. Compared to the mesh convergence, Figure 4.11 shows that decreasing tolerance is more susceptible to deviations than number of elements.

4.5 Discussion

The effects of rank order change are apparent especially in runs that had a high probability of a rank change and a larger magnitude in rank change. This is perhaps unsurprising as the algorithm increasingly struggles to rely on its outputs to continue the optimisation process. What is perhaps surprising is that even with the increased volatility in the rank order, the EA still managed to find optimum routes for the TSP. This suggests that the EA used in this work is robust and is able to deal with imperfect data. This may be in part due to the re-evaluation of the population during each generation. The change in rank is only valid for one generation. One of the limitations of this work is that it was only tested on a single benchmark, rather than a series of benchmarks. Including more examples would have increased the rigour of this work by eliminating the possibility that the results are

Table 4.6: Benchmark 1 successful runs.

| Degree | Number of Successful Runs | Average Generation Where Optimum Found |
|----------|---------------------------|--|
| Baseline | 2564 | 53.50 |
| 51 | 2534 | 53.75 |
| 52 | 2524 | 54.50 |
| 53 | 2513 | 53.54 |
| 54 | 2554 | 53.42 |
| 55 | 2521 | 53.91 |

Table 4.7: Benchmark 2 successful runs.

| Degree | Number of Successful Runs | Average Generation Where Optimum Found |
|----------|---------------------------|--|
| Baseline | 5612 | 62.07 |
| 15 | 30 | 42.93 |
| 17 | 27 | 15.22 |
| 19 | 5703 | 62.15 |
| 20 | 5187 | 60.99 |
| 21 | 5527 | 61.99 |

Table 4.8: Benchmark 3 successful runs.

| Degree | Number of Successful Runs | Average Generation Where Optimum Found |
|----------|---------------------------|--|
| Baseline | 4444 | 56.99 |
| 19 | 4529 | 58.08 |
| 20 | 4410 | 57.30 |
| 23 | 4313 | 58.26 |
| 24 | 4287 | 57.03 |
| 25 | 4275 | 57.27 |

Table 4.9: Benchmark 4 successful runs.

| Degree | Number of Successful Runs | Average Generation Where Optimum Found |
|----------|---------------------------|--|
| Baseline | 4389 | 55.43 |
| 30 | 4632 | 55.92 |
| 31 | 4478 | 55.66 |
| 32 | 3550 | 55.91 |
| 33 | 3833 | 55.70 |
| 34 | 4363 | 55.47 |
| 35 | 4466 | 56.05 |




Figure 4.10: The lift-to-drag ratio of three 2D NACA aerofoils as the number of elements is increased. The rank of the three aerofoils remains constant as the number of elements are increased.




Figure 4.11: The lift-to-drag ratio of three 2D NACA aerofoils as the residual tolerance is increased. At lower tolerances, the differences between the three aerofoils decreases though the ranking remains the same for most of the range.

problem-specific.

One of the difficulties in modelling benchmark equations was the time taken to meaningfully represent the original function through a Taylor series while still introducing some variability. Equations with polynomials in the 2nd or 3rd order could be represented exactly with very few terms though any decrease in the number of terms would skew the function too far. Time was another factor when considering which problems were to be tested. Benchmark problems involving the exponential function e^x took a substantial amount of time (≥ 60 minutes) to model the equation with higher order series. Including log functions in the benchmark problem also took a substantial amount of time to convert from benchmark to Taylor series. Another consideration to take into account when interpreting the results is to note how close the optimum is to the expansion point. Changing the expansion point could influence the shape of the function and ultimately how close the Taylor series optimum lies to the function's optimum. A limitation of this approach is that it is only a one-dimensional example, whereas the Navier-Stokes equations used previously are three-dimensional. The simplicity of the benchmark functions may hinder the ability for the study to accurately reflect the complexity of equations such as the Navier-Stokes equations.

4.6 Conclusion

Both changes to the rank order and changes to the equation fidelity affected the success of the optimisation when changes became large and frequent. This suggests that EAs are robust and are resilient to small inaccuracies in the data.

In Chapter 2, three different stopping criteria were mentioned when determining sufficient convergence. The first is a solution-based criterion which uses the residual to determine




Figure 4.12: The change in rank of the 5 designs used in the design of experiments as the iteration count increases. Rank change volatility decreases as iterations increase and eventually do not change.

when to terminate the simulation. Time and iteration count are the second and third criteria, which are solution-independent. These are usually imposed, as in this work, when convergence to a specific residual may take hundreds of hours or may not be possible to achieve at all. It has also been shown here in Figure 4.12 that after a certain iteration count, the ranking of the designs does not change. Taking advantage of this is possible by implementing a convergence criteria that is semi-solution-dependent and terminates when the ranks of the designs that are to be tested in the EA do not change for a given amount of time or number of iterations.

Chapter 5

Illumination Through Collaboration

In Chapter 3, a design optimisation study was undertaken using some methods that are broadly categorised as AI. A controlled environment is necessary to better understand the interaction between the designer and the optimisation methods. The relationship between people and AI is investigated through a small scale user study involving 12 undergraduate and postgraduate students. The user study provides essential qualitative context and supports quantitative findings from a large online user study investigating the use of a support algorithm called MAP-Elites. A simple car designer application was used which presented participants with either curated designs from the MAP-Elites algorithm or randomly selected designs. A mixed-methods questionnaire was given to participants before and after using the car designer. Participants answers from multiple choice and open-ended questions are evaluated alongside their performance during the study. While participants were more likely to design better cars than computer-led designs and were more likely to engage with MAP-Elite designs than random designs, the effect of MAP-Elites on the design process is still in need of further investigation.

5.1 Introduction

The novel ducted winglet was originally designed without intervention from a computer. The novel winglet then underwent parametrisation so that it could be set up for a computer-led optimisation. Through different optimisation methods, the drag of the novel ducted winglet and the induced drag were lowered. However, there was very limited opportunity for designer intervention in the optimisation cycle. The new parameters were all set by algorithms and a constrained design of experiments. As was seen in Section 3.3.1 where the elliptical profiles was constrained to fit inside the winglet, human intervention and editing of the profiles deviated from standard algorithm practices. That is not to say that human intervention is not capable of finding an optimal design, but it certainly must be compared to grounded theory to see its effects. With simulation times exceeding 24 hours, it would be difficult to determine if designers are capable of inserting themselves in the design process and having a positive impact on the objective function. Likewise, it would be difficult to assess whether an algorithm is suitable to facilitate designers in the optimisation cycle. An interactive tool with shorter simulation turnaround times that is optimises something more common than a ducted winglet is needed to evaluate the effects of human-algorithm interaction on the design process.

These challenges are addressed in this chapter. First, similar work that is done in the field of human-computer interaction is explored. A focus is brought to a variety of works that discuss the benefits and difficulties of integrating humans in the design loop. An evolutionary algorithm that categorises the elites based on certain features is also discussed. The results from the motivating study for this chapter are discussed and how it necessitated further investigation. Participants with engineering backgrounds take part in a user study answering a survey and using a 2D car designer game. The results from the

study are presented and compared to simulations without human intervention. Categories for how participants interacted with the designer are also set and discussed.

5.2 Background

This chapter introduces the concept of Human-Computer Interaction (HCI). It is a vast subject area and very often crosses domain boundaries for interdisciplinary research such as counter-extremism [94], social cohesion [95], and as in our case, engineering [7]. Some work in this area focuses heavily on the human aspect, while others focus more on the technical side. It is the **relationship between the human and the computer** that is fundamental to the domain.

5.2.1 Human-Computer Interaction

User-centred design can help improve usability and satisfaction when dealing with software designed to assist engineers [96]. Fundamentally, the interaction that users have with the software is just as important as the software itself. People must be willing to take new technology onboard. Their lack of cooperation and willingness can be a barrier to adopting a new technology, such as EAs [97]. However, properly implemented HCI can lead to significant, positive changes in the design process. This extends beyond the siloed work of some engineers to cover cross-discipline processes, such as that between HCI and software engineering [98]. By considering the entire design process and workflow, there is an increase in value added by designers while also a shorter turnaround time between design iterations.

The symbiotic relationship between designer and software must be carefully fostered. In metacreative systems [99], such as EAs, there are key elements from an HCI standpoint

that must be taken into consideration in order to increase the probability of success if adoption is to be used as a criterion. Firstly, the computer program must be able to actively contribute to the design process. For EAs, the designer sets a fitness function and the algorithm proceeds to optimise the design with respect to that objective. While the program may ultimately yield the final result, it would not have done the task without the input of the designer. Likewise, the designer would not have been able to generate such a solution in a meaningfully short amount of time. Secondly, and as an extension of the first point, both the designer and the program must work together in real-time to find a solution. Only the designer truly knows what is needed from the program and so continuous input from the designer is needed.

Current software and hardware realities limit the applicability of some of the visualisation and intervention techniques for CFD engineers and aerodynamic designers. This is primarily due to the complex nature of 3D flow simulation and the large amount of data that needs to be processed to create visualisations. The intervention of the designer in the optimisation process can still occur though at lower frequencies compared to those optimising simpler problems.

Thorough considerations for HCI within the domain of design engineering is not without merit or precedent. As previously mentioned, the challenges and implications of HCI and software engineering were considered. In Walenstein's exploratory paper [98], boundary objects in HCI and in software engineering could be used to bridge the gap between the two domains. Boundary objects are general concepts around which a domain or group of experts centre themselves around. Often when two domains are brought together, new techniques for working emerge that are more efficient and also self-perpetuating. In a review done of advances in material engineering being adopted by researchers in HCI, it

was noted that material scientists and engineers benefited from the advances made in HCI [100]. It follows that a discovery made in material science leads to a new HCI applications, which again leads to new discoveries in material science.

5.2.2 Mixed Initiative Design Systems

The aim of Mixed Initiative (MI) systems is allow people to be more involved in the algorithmic process, and for algorithms to optimise and find solutions faster than a traditional black box system. Horvitz proposed a set of requirements in his paper [101] in which the principles of MI user interfaces are discussed. 12 factors are said to be critical for MI user interfaces, two of which are highlighted here: developing significant value-added automation and considering uncertainty about a user's goal. These two factors are chosen as they exemplify the difficulty faced by engineers in industry. The first of these factors is important from a business standpoint as value-adding activities serves as the justification for engineers' salaries. Having a system that increases the value added per task also increase the value of the engineer to that company. The second factor focuses on the discretion of the engineer and the complexity of many engineering problems. Certain parameters cannot be input or would add an unnecessary layer of complexity to the problem, such as aesthetic preferences. Horvitz also maintains that an MI system should act as a "benevolent assistant" to its users. This stance is contested by Chanel et al. [102] who argue that humans and MI systems should be seen as "teammates" instead. In this context, a benevolent assistant would simply obey the commands of the user while a teammate would offer input much the same way another engineer would.

MI systems aim to evoke natural human-human collaborations. A 2020 study of 42 engineering and science students conducted by Chen et al. [103] investigated how humans collaborate with other humans and how humans collaborate with computers. While peo-

ple should be aware that they are using a computer-based system, 35% of users incorrectly assumed they were collaborating with a human rather than a computer. This indicates that already computer systems are already able to replicate some human elements with respect to collaborations.

A human-centric approach can be taken for MI systems in a manner that builds trust in the system. For a human-centered system, Zhu et al. [104] present axioms that differentiate between an explainable model and an observable model. If the MI system is to truly be a companion, then it must be able to justify its decision in the same way engineers must also be able to justify their decisions. MI systems capable of justifying their choices will build trust with their human counterparts. Nguyen et al. [105] mention transparency, auditability, and how an algorithm came to a result as criteria for users to be able to understand and trust models. Nguyen et al. applied this to fact-checking, but the statement is applicable to MI systems, EAs, and algorithms broadly. Not only is trust important as a sentimental value, it also increases productivity as per Saeidi et al. [106]. They found a 31% increase in performance and a reduction in workload of 23.9% when dealing with a human-robot interaction. Explainability was also a key pillar in a user study of 11 energy engineers by Degen et al. [107]; they found that understandability and actionability were prerequisites for trustworthiness.

A key element to consider when designing an MI system is the relative weight given to decisions made by people and computers while considering the balance between creativity and productivity. That isn't to say there is a fixed point at which an optimum exists, but rather a dynamic range dependent on the task. Liapis et al. [108] discuss the changing influence between users and the computer with the use of 4 MI tools to assess user engagement and creativity. The paper by Liapis et al. [108] argues that MI systems

enhance creativity and lateral thinking in humans by disrupting the often logical thought process when tackling problems. The authors also discuss giving users the ability to decide the degree to which an MI system can influence user decisions by not simply splitting the decision-making process as a 50/50 split. There exist systems now that are influenced more by computer actions while others are more heavily influenced by human interaction. A computer-guided MI program that only lets users act on the final results of the program was proposed by Swearngin et al. [109]. Contrarily, Sarr et al. [110] present a MI game for young children learning how to write, where the computer acts in a reduced role in a human-led environment.

Creativity can also be a standalone metric that does not need to be evaluated alongside productivity, although the role of MI systems in this aspect is contested. Creativity was discussed by Samuel et al. [111] in their paper where they present an MI program that assists writers by providing writers with prompts they would have not otherwise had. The idea that MI systems provide an opportunity to capitalise on the advantages of both humans and computers is not shared universally. In the context of procedural content generation of video game levels, Baldwin et al. [112] say that computers are limited by human imagination.

5.2.3 Does mapping elites illuminate search spaces? and Other Questions for MAP-Elites

Multi-Dimensional Archive of Phenotypic Elites [34], shortened to MAP-Elites, categorise and store (archive) the best results (elites) from an EA by many (multi-dimensional) of their observable characteristics (phenotypic). This approach provides a more nuanced description of the successful agents in the EA. Importantly, the characteristics that are chosen to categorise the elites may not be linked to the objective function.

Given that MAP-Elites provides a range of optimal solutions, it easily allows for human intervention in the selection phase of the EA. Interactive MAP-Elites have been explored by those in the procedural content generation domain, like Alvarez et al. [113] who demonstrated the algorithm on a five-dimension dungeon level designer. Similarly, Balla et al. [114] used MAP-Elites on a rope-swing game with 5 varying dimensions and demonstrated that it could find a relationship between the parameter space and the behaviour space. *Baba is Y'all* is a game developed by Charity et al. [115] that also relies on MAP-Elites for its procedural content generation. The authors of *Baba is Y'all* ran two user studies, one which was an informal online study and another formal detailed study. Charity et al. found that users preferred to have complete control over the design process and did not want to involve the computer.

Several examples in the literature demonstrates that MAP-Elites can outperform non-diversity based algorithms but do not have a connected user study supporting their argument. Work by Charity et al. [115] on *Baba is Y'all* indicate that it is not something users want though they recognise that further work needs to be done to better understand this stance.

In the paper *Does mapping elites illuminate search spaces?*, Walton et al. [8] conducted an online user study to see if users were more likely to use the MAP-Elites algorithm over a random selection of 2D car designs. The purpose of this comparison was to determine if MAP-Elites were a preferable method to optimisation for a large number of participants. The results from the online study show that users were more likely to use the MAP-Elites algorithm, though the reasons as to why that was the case were unknown. To find out, a separate in-person study was developed. The details of this detailed study are laid out in the following section.

5.2.4 A Summary of Results from the Online Study

The following results were generated from the large scale online study [8]. This short summary is necessary to provide proper context for some of the results from the targeted user study. I contributed to the data analysis of the online study and contributed during the writing of the respective article.

Participants in the online study had 3 MAP-Elite views to choose from and one control view which had the randomly selected designs. The participants who interacted with the editor and the MAP-Elite views were statistically more likely to make improvements compared to those who did not engage with either. Additionally, participants who viewed at least one MAP-Elites view did better than those who did not. Participants who engaged with the car designer for a longer period of time were also more likely to improve, as well as those participants who made more selections. Most importantly, participants who spent more time in the MAP-Elites view compared to the control view performed better. The statistical significance was set at $p < 0.002$.

While a significant amount of data was gathered during the online study, it lacked context. The background of participants was unknown as well as their thoughts during the optimisation process. To better understand the behaviour of participants who interacted with the car designer and to provide more context, a small scale, in-person user study was proposed.

5.3 Method and Materials

The survey conducted as part of the work by Walton et al. [8] was developed entirely independently by myself. Feedback was given by the other authors to ensure clarity and

scientific rigour.

For both the work done by Walton et al. and the field in general, a user study was needed to determine if MAP-Elites were suitable to use for optimisation. This leads to the following research question:

Does the user-centred illumination algorithm have a positive influence on the optimisation process?

The MAP-Elites will be compared to randomly generated designs, referred simply as random designs. Both lists will have a non-descriptive name to prevent participants from being biased by the heading. In this case, the lists will be classed under *Insights 1/2*. The order in which they appear on screen will also be randomised so that *Insights 1* appears before *Insights 2* half of the time and vice versa for the other half. Participants were also unaware that two different methods for selecting designs was employed prior to starting the study.

The number of generation was limited to 40 as a way of ensuring that each participant spent an equal amount of time using the car designer, and to keep the user study sessions within a reasonable timeframe. Simulating 40 generations takes approximately 20 minutes, with an additional 20 minutes given to answer the survey. 5 minutes was also allocated for instruction prior to using the car designer. This totals 45 minutes maximum for each participant.

The performance of participants is compared to a baseline of 20 computer-only simulations which received no input during the 40 simulations. The number of computer-only runs was suggested by an expert in the field.

5.3.1 Participants

Participants for this study were all recruited from Swansea University. Ethical approval for this study was granted by the Faculty of Science and Engineering Ethics Board (SU-Ethics-Student-140323/6255). The two criteria for being selected were that a participant:

1. is currently studying towards or has completed an undergraduate degree in engineering; and
2. has not used the car designer before.

The reason for imposing the requirement for an engineering undergraduate degree was because the work in this thesis targets the interaction between engineers and their design tools. The second condition was imposed to prevent previous experience influence the interaction with the car designer [116]. Participants in similar user studies are often recruited because of their educational backgrounds as it implies that they have a base-level of knowledge appropriate to understand and complete the task at hand. Given the nature of the car designer and the aim of the work, students currently enrolled in an engineering undergraduate degree or those with an undergraduate engineering backgrounds were deemed suitable for this user study.

12 participants were recruited in total, 8 of which were male. At the time of the study, 4 participants were studying undergraduate degrees while the other 8 were postgraduate students. The age of participants ranged from 20 to 28 with a mean of 23.8 and a standard deviation of 2.15.

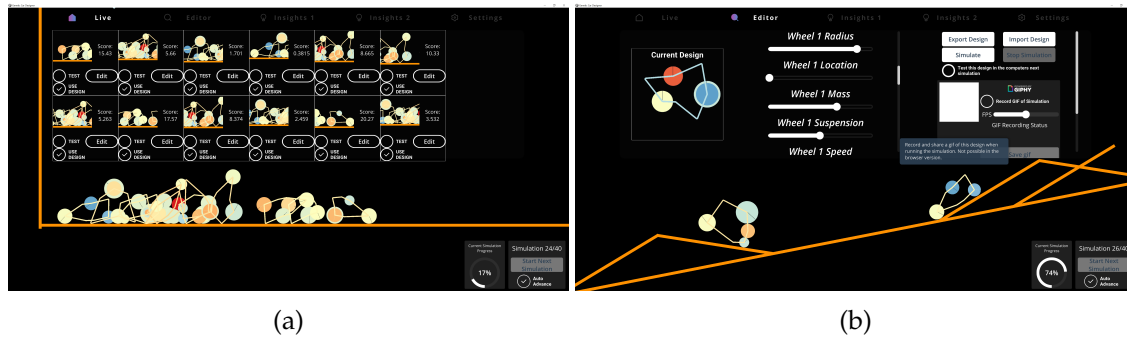


Figure 5.1: Two of the views that participants could interact with during the user study. The *Live* view (a) showed all the cars currently running and was centred on the car which had progressed the furthest. Participants could change the parameters which defined a car in the *Editor* view (b).

5.3.2 Car Designer

The car designer is a 2D game where cars travel from left to right. The objective of the game is to get a car to travel as far right as possible as it goes over triangular obstacles. The cars must also travel up a hill which gets progressively steeper the further that the cars travel. There are 4 different views that can be selected when using the car designer. The first is the *Live* view, which shows all the cars currently running. Initially, 12 cars travel the track simultaneously. Each car has its own mini-window within the *Live* view with a real-time score which corresponds to the distance travelled by the car. A button is also available to edit the designs. Pressing this button brings the participant to the *Editor* view, where users can manipulate a selected design. Sliders are used to vary the following parameters:

Wheel speed and direction were set on the same slider. Sliding from the middle to the right increased the speed clockwise, while sliding from the middle to the left increased the speed counter-clockwise. This did cause some confusion among participants but was pointed out during the study as to not artificially hamper the efforts of the participants.

Table 5.1: Design parameters for the simulated cars. Positive values for *Wheel Speed* indicate clockwise rotation while negative values indicate count-clockwise rotation.

| Parameter | Data Type | Range |
|---------------------------|------------|-----------------|
| Vertex radius | Continuous | [0.1, 3.0] |
| Body mass | Continuous | [100.0, 1000.0] |
| Wheel location | Discrete | [1, N_w] |
| Wheel mass | Continuous | [1.0, 100.0] |
| Wheel radius | Continuous | [0.1, 1.0] |
| Wheel speed and direction | Continuous | [-2500, 2500] |
| Wheel suspension | Continuous | [0.1, 100.0] |

The implications of this are discussed in the last section.

The last two views are the randomly selected designs, which are used as a control, and the MAP-Elite designs; these are labelled as *Insights 1* and *Insights 2* respectively. Each *Insights* view had an image of the previous designs along with the score. Participants can edit the designs, select the design to be simulated again in the next simulation, or use as a parent when generating designs for the next generation. The designs run again are later referred to as keeps, and those used for generating new designs are called breeds.

Prior to using the car designer, the participants are shown screenshots of the different views, shown in Figures 5.1 and fig:participantinsights. This is done to reduce the amount of questions associated with trying to understand how to use the car designer. Each participant is briefed on all the views, the simulation progress bar, the generation count, and how to interact with the editor sliders.

The car designer itself was made in Unity3D¹ and used the provided 2D rigid-body physics engine to simulate the cars as they travelled. The fitness of each car, which was available to users via the score associated with each design, was a measure of the horizontal

¹<https://unity.com/>

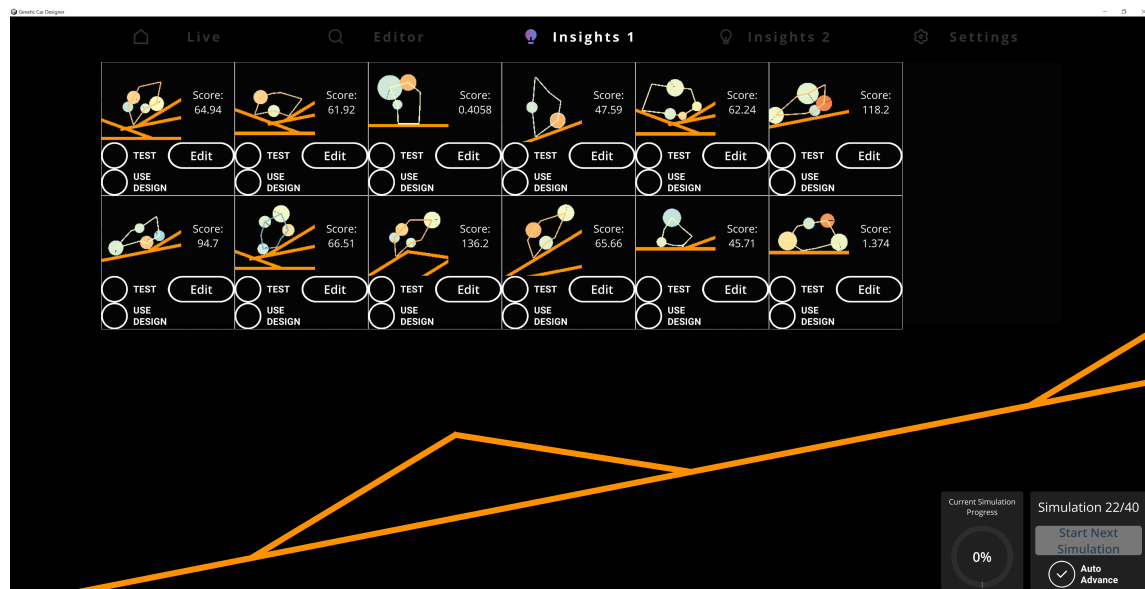


Figure 5.2: One of the *Insights* views shown to participants, which displayed either randomly generated designs or designs from the MAP-Elites algorithm. Participants could choose to use selected designs in the next generation or use a selected design as inspiration for a new car in the next generation.

distance travelled and was available for the running cars in the *Live* view and alongside the designs in the *Insights 1* and *Insights 2* views. Wheels were added using the 2D wheel joint component in Unity3D and simulated both the suspension and rotational speed. The suspension was modelled as a spring-damper system using a damped oscillation frequency and the rotational speed was modelled as motor with a given torque.

Algorithm 5 Process in which cars are optimised in the car designer.

```

Initialise starting population  $c$ 
for generation = 1,2, ... 40 do
  Simulate cars on course
  Add simulated cars to parent population  $p$ 
  for Each Car in  $p$  do
    if Breed = True then
      Add Car to parent population
    end if
  end for
  Employ EA steps as per 1
  Update Insights views
  for Each Car in  $p$  do
    if Keep = True then
      Add Car to  $c$ 
    end if
  end for
  Restart simulation
end for

```

The cars were defined as a closed polygon with a set number of vertices N_v and wheels N_w . The vertices were positioned at (r_i, ϕ_i) where $i \in 1 \dots N_v$, r is a participant-defined radius and $\phi = 360(1 - i)/N_v$. The coordinates of the closed polygon defined a mesh that would collide with the track but not the wheels. The centre of mass for the closed polygon was calculated based on the vertex coordinates and affected the behaviour of the car when coupled with the total mass of body M_b , which could be set by the participant. Given that each wheel had 5 dimensions which could be changed, plus the number of vertices and

the mass of the body, the total number of dimensions D for the design space was defined as $D = 1 + N_v + 5N_w$. In the case of this work, the number of vertices was set to 7 and the number of wheels was set to 5, leading to $D = 33$. The arbitrary number of dimensions was set so that participants could not find a trivial solution to the problem with a low number of dimensions nor could they be overwhelmed with too many dimensions. In the online study, users could select the number of vertices from 3 to 24 and the number of wheels from 1 to 12. There were also 4 levels to choose from, including “Bumps up Hill” which was described in this section.

With regards to the MAP-Elites algorithm, there are many different observable characteristics (phenotypes) that could be used to define each car. For the purposes of this work, the centre of mass was chosen as the characteristic to be used in the *Insights* view.

The seed will also be the same for all runs to make sure that participants and any runs that don't involve human input start from the same point. This should reduce the likelihood that a participant benefits from designs which are relatively high scoring compared to the average score. Participants with designs that have low initial fitness also have a greater potential for improving their fitness.

5.3.3 Survey and Qualitative Content Analysis

The survey consists of 4 parts. 2 parts were completed before interacting with the car designer, and 2 others were completed afterwards. The first part of the survey asked generic demographic questions which included age and gender. The second part of the survey asked participants to either agree or disagree with the following statements:

1. I prefer to do calculations by hand.

2. I have no problem trusting my life to technology.
3. I always double-check results from automated methods.
4. I am confident using most technology.
5. The usefulness of technology is overrated.

Participants could select “Indifferent.” These statements were designed to reveal biases of participants towards technology and are derived from former work by Vincalek et al. [5].

The third and fourth parts of the survey were specific to their interaction with the lists which showed the random designs and the MAP-Elite designs. The third part was a series of multiple choice question based on a similar study done by Knijnenburg et al.[116], asking participants if they preferred *Insights 1* or *Insights 2* when it came to areas of accuracy, diversity, satisfaction, and novelty; again, participants could choose neither. The questions posed are shown in Table 5.2.

The final part of the survey are 5 open-ended questions, once again based on the user study done by Knijnenburg et al. [116]:

1. How easy or difficult was the car designer to use?
2. How easy or difficult was it to make a decision from Insights 1 and 2?
3. Was the decision process frustrating for Insights 1 or 2?
4. How much effort did you need to invest while using Insights 1 or 2?
5. Do you trust the results from the car designer?

Participants were encouraged to answer in as much detail as possible, though short an-

Table 5.2: Questions posed to participants after they had completed the study. Participants could either choose *Insights 1*, *Insights 2*, or *Neither*. Along with their respective target areas, questions are also in the order they were presented to participants.

| Subjective Area | Question |
|-------------------------|--|
| Subjective Accuracy | Which presented better solutions? |
| | Which allowed you to select more optimal recommendations? |
| | Which showed too many poor recommendations? |
| Subjective Diversity | Which presented a variety of choices? |
| | Which showed a bigger difference in recommendations? |
| Subjective Satisfaction | Which was more valuable with respect to your time? |
| | Which had more satisfying recommendations? |
| | Which would you trust more to provide you with recommendations? |
| | Which made finding a new solution easier? |
| Subjective Novelty | Which gave you more recommendations you would not expect? |
| | Which gave you recommendations you would not have thought of yourself? |
| | Which allowed you to explore new ideas better? |

swers were still accepted if they believed that was appropriate.

The answers for the open-ended questions were interpreted using Qualitative Content Analysis (QCA)[117]. QCA was chosen over methods such as thematic analysis used in previous work [7] because the questions asked participants specific questions about key aspects of the car designer. The first step in QCA is to familiarise oneself with the data. Codes are created to group answers together with all codes brought together in a codebook. To reduce the potential for misinterpreted results, two researchers will code the answers independently. The two separate codebooks are compared and discrepancies between the codebooks are discussed and resolved.

5.4 Results

The results in this section are separated into quantitative and qualitative results and comparing both to how well each participant performed during the study. Their performance was measured by the increase in fitness from the starting design to the best design. First, the participants' performances are compared to where they spent their time and how often they chose designs from either *Insights* view. The answers participants gave to the multiple choice questions are also compared to their performance. Participants had their open-ended answers coded using quantitative content analysis which is used to give context to the behaviour they exhibited during the study. Finally, a concept of personas is used to define participant behaviour and group them according to how they chose to engage in the design optimisation process.

5.4.1 Fitness

In the following section, we'll be looking at the effect that the participants have on the fitness. The best fitness attained by each participant is compared to the fitness of runs that did not receive any input during the run. Where participants spent their time in the car designer and how often they would keep or breed designs is also compared to the maximum fitness value. Lastly, links between participants' preconceptions about technology and fitness are explored. All values relating to fitness are expressed as an improvement relative to the best design in the first generation.

Human vs Computer

In Table 5.3, the performance of the 12 participants is compared to the 20 computer runs. In each of the three metrics, the participants outperformed the computer. Notably, the average fitness of the participants is higher than the best run by the computer. Moreover,

Table 5.3: A comparison of the range of design fitnesses between participants and computer-only runs. The standard deviation is included alongside the average results. In all three categories, the participants had higher improvements in fitness.

| Fitness Improvement | Participants | Computer |
|---------------------|------------------|-----------------|
| Range | [59.8, 206.3] | [43.1, 136.4] |
| Average | 137.3 \pm 39.4 | 93.0 \pm 22.8 |
| Median | 145.5 | 94.9 |

only two participants had scores lower than the median for the computer-only simulations.

Views

Knowing where participants spend their time is important when determining whether the MAP-Elites were more useful than the random selection. On average, participants spent 11.6% of their time in the MAP-Elites view and 9.3% of their time in the random view. The minimum time spent in amount of time was also marginally higher in the MAP-Elites view, 2.4%, than the random view, 2.3%. There is, however, a large difference between the maximum amount of time spent in each view: 52.9% in MAP-Elites and 25.1% for random designs. The ratio between the percentage of time spent in each view shows that participants spent 45% more time in the MAP-Elites view than they did in the random view.

The graph presented in this section, Figure 5.5, shows the percentage of time spent in each of the 4 views with the fitness improvement overlaid. Generally, participants spent most of their time interacting with the editor. The amount of time spent using the MAP-Elites view and the random designs view generally increased during the run. Remarks about where time is spent and how it affects fitness are made as well for each participant, and are available in Appendix A.




Figure 5.3: All runs of the car designer with the best fitness of each run shown per generation, separated into participant and non-participant runs.




Figure 5.4: The average fitness of participants and computer-only runs grouped together where the shaded region indicates a confidence interval of 95%.




Figure 5.5: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for all participants where the shaded regions indicate a confidence interval of 95%.

Keeps and Breeds

In the *Insights* views, participants were able to select a design either to be used in the upcoming generation or as a parent in the next iteration. These keep and breed options were selected by participants using a checkbox next to the designs. The editor had similar functionality, but the design could only be used in the upcoming generation and not as a parent. The amount of keeps and breeds that were performed by a participant are plotted against the improvement over the course of their runs.

The figures show the cumulative number of times that a view's keeps or breeds have been selected. Many participants have steady increases in editor keeps. The toggle for this was persistent throughout their runs. The maximum number of editor keeps was 38 and the minimum was 31, with the median being 35.5 editor keeps. Generally, participants were much more likely to use the MAP-Elites for keeps and for breeds. The amount of keeps and breeds for the MAP-Elites ranged from 0 to 34 and 0 to 94 respectively. For the control keeps and breeds, the number of selections ranged from 0 to 6 and 0 to 29 respectively. The average number of keeps and breeds for all participants is compared to the average best fitness in Figure 5.6. Individual comparisons for all participants are available in Appendix A.

5.4.2 Survey

This section highlights the findings from the QCA and compares it to the fitness and the multiple choice questions. Participants are also categorised into personas based on their use of the car designer.




Figure 5.6: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for all participants where the shaded regions indicate a confidence interval of 95%.

Table 5.4: Responses to statements before using the car designer where participants could agree (✓), disagree (✗), or remain indifferent (–) to the listed statements.

| ID | Hand Calculation | Trust | Double-check | Confident | Usefulness | Fitness |
|----|------------------|-------|--------------|-----------|------------|---------|
| 1 | ✓ | ✓ | ✗ | ✗ | – | 125.35 |
| 2 | ✗ | ✓ | ✓ | ✓ | ✗ | 163.99 |
| 3 | ✗ | ✗ | ✓ | ✓ | ✗ | 137.57 |
| 4 | ✓ | – | ✓ | ✓ | ✗ | 155.13 |
| 5 | ✗ | ✓ | ✓ | – | ✗ | 155.41 |
| 6 | ✗ | – | ✗ | ✓ | ✗ | 206.23 |
| 7 | ✗ | – | ✓ | – | ✗ | 153.40 |
| 8 | ✓ | ✗ | ✗ | ✓ | ✗ | 121.69 |
| 9 | ✓ | – | – | ✓ | ✗ | 168.15 |
| 10 | ✗ | ✓ | ✓ | ✓ | ✗ | 59.82 |
| 11 | ✗ | – | – | ✓ | ✗ | 81.26 |
| 12 | ✓ | ✓ | ✓ | ✗ | – | 121.53 |

Multiple Choice Survey Questions

5 multiple choice questions were asked prior to interacting with the car designer.

5 participants said they preferred to do calculations by hand while 7 disagreed with the statement. 5 participants had no problem trusting their lives to technology, 5 were indifferent, and 2 had reservations about the statement. 7 participants double-check the results from automated methods, 3 did not, and 2 were indifferent. 8 participants agreed that they were confident using technology, with 2 participants disagreeing and 2 others being indifferent. 10 participants disagreed that the usefulness of technology was overrated, though 2 did say they were indifferent.

Generally, participants are confident using technology and recognise its usefulness. However, the majority of participants still verify the results from automated methods.

In addition to the 5 questions asked before using the car designer, the participants also answered 12 questions regarding their preference between the MAP-Elites and the ran-

domly selected designs. The questions were centred on their subjective accuracy, diversity, satisfaction and novelty. With respect to satisfaction, participants were split between the MAP-Elites and the randomly selected designs, though they showed a slight preference to the MAP-Elites. More participants (N=6) said that the MAP-Elites presented better solutions than the randomly selected designs (N=4), while 2 participants chose neither. The MAP-Elites and random designs both had 6 participants say that the view gave them more optimal solutions. 6 participants also said that the randomly selected designs showed too many poor recommendations compared to 4 who said the same thing of the MAP-Elites; 2 participants said neither. As for diversity, 6 participants said that the random designs showed a variety of choices and a bigger difference in recommendations compared to 5 participants who said the same about the MAP-Elites. In both questions, one participant chose neither. Subjective satisfaction was measured using 4 questions revolving around the ideas of time, trust, ease of use, and overall satisfaction. A total of 5 participants said that the view with the MAP-Elites was more valuable with respect to their time while 4 said it was the view with the randomly selected design, and 3 said neither. Trust amongst participants for both views was even with 5 participants choosing MAP-Elites and 5 choosing the random designs; 2 chose neither. An even amount of participants also chose the MAP-Elites (N=6) and random designs (N=6) when it came to ease of use. Overall satisfaction leaned toward the MAP-Elites (N=6) compared to the randomly selected designs (N=4) and 2 said neither. Finally, the novelty of the MAP-Elites and the randomly selected designs showed a slight preference for the MAP-Elites, especially when it came to exploring new ideas. 7 participants said the MAP-Elites allowed them to better explore new ideas while 4 participants said the randomly selected designs did this better. One participant said neither view was better. The same amount of participants (N=6) said that the MAP-Elites view and the random designs view gave them recommen-

dations they would not expect. 5 participants said the MAP-Elites were better at showing recommendations they would not have thought of themselves while 4 participants said the randomly selected design were better than the MAP-Elites in this respect. The results are tabularised with participants' answers to each question in Table 5.5.

Qualitative Content Analysis

Following the method described in Section 5.3.3, a codebook was created with Walton as the second researcher which is available in Table 5.6. A total of 24 different codes were applied which can be categorised under 7 different themes. The combined codebook has applied to the open-ended answers shown in Table 5.7. Some of codes and themes are question-dependent, such as codes in the *Difficulty* theme which are specific to the first two questions. A breakdown of how many codes are applied to each question is available in Table 5.8.

A small summary of the code frequency gives insight into how participants approached the design challenge. Participants noted 11 times that the car designer and choosing a design was easy as opposed to only 3 mentions that the car designer and choosing a design were difficult. 5 participants also mentioned that the car designer got easier to use with time. Most participants (N=9) were not frustrated by using the car designer and most (N=7) said they put in minimal effort into using the car designer. Additionally, most participants (N=9) trusted the output of the car designer. When it came to designing their cars, participants expressed interest in correlating their designs to the score a total of 3 times. Participants also mentioned 4 times that their preconceptions played a role in their design process. The most frequently coded observation, mentioned 14 times across all questions, indicated that participants were very aware of the score and utilised this when taking into account which designs to select as either a keep or breed. Participants also mentioned

Table 5.5: The responses from participants for the A/B style questions which were completed after using the car designer. Participants were given the choice of *Insights 1*, *Insights 2*, or Neither. The results were then mapped to MAP-Elites ★, randomly selected designs □ and neither –.

| Participant | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|--|---|---|---|---|---|---|---|---|---|----|----|----|
| Which presented better solutions? | – | ★ | □ | – | □ | ★ | ★ | □ | ★ | ★ | □ | ★ |
| Which allowed you to select more optimal recommendations? | □ | ★ | □ | □ | □ | □ | ★ | □ | ★ | ★ | ★ | ★ |
| Which showed too many poor recommendations? | ★ | □ | – | ★ | ★ | □ | □ | ★ | □ | □ | – | □ |
| Which presented a variety of choices? | □ | ★ | □ | – | ★ | □ | □ | ★ | ★ | □ | ★ | □ |
| Which showed a bigger difference in recommendations? | – | ★ | □ | ★ | ★ | □ | □ | ★ | □ | □ | □ | ★ |
| Which was more valuable with respect to your time? | – | ★ | □ | □ | □ | □ | ★ | – | ★ | ★ | – | ★ |
| Which had more satisfying recommendations? | □ | ★ | □ | □ | □ | – | ★ | ★ | ★ | ★ | – | ★ |
| Which would you trust more to provide you with recommendations? | □ | ★ | □ | □ | □ | – | ★ | □ | ★ | ★ | – | ★ |
| Which made finding a new solution easier? | □ | ★ | □ | ★ | □ | □ | ★ | □ | ★ | ★ | □ | ★ |
| Which gave you more recommendations you would not expect? | ★ | ★ | ★ | □ | ★ | □ | □ | □ | □ | □ | ★ | ★ |
| Which gave you recommendations you would not have thought of yourself? | □ | ★ | – | ★ | ★ | □ | □ | – | ★ | ★ | □ | – |
| Which allowed you to explore new ideas better? | □ | ★ | – | □ | ★ | □ | ★ | □ | ★ | ★ | ★ | ★ |

Table 5.6: Codebook which was used to analyse the responses for open-ended questions.

| ID | Theme | Code | Definition |
|----|----------------|------------------------|--|
| D1 | Difficulty | Easy | A participant found one or more aspects of the car designer easy to use. |
| D2 | | Neutral difficulty | A participant found it was neither easy nor difficult to use one or more aspects of the car designer. |
| D3 | | Difficult | A participant found one or more aspects of the car designer difficult to use. |
| D4 | | Easier with time | A participant found one or more aspects of the car designer easier to use as time went on. |
| F1 | Frustration | Not frustrating | A participant did not find one or more aspects of the car designer frustrating. |
| F2 | | Neutral frustration | A participant was somewhat frustrated with one or more aspects of the car designer. |
| F3 | | Frustrating | A participant did find one or more aspects of the car designer frustrating. |
| F4 | | Equal | A participant found the same level of frustration with both insights. |
| E1 | Effort | Minimal effort | A participant used little to no effort when using one or more aspects of the car designer. |
| E2 | | Some effort | A participant used some effort when using one or more aspects of the car designer. |
| E3 | | Lots of effort | A participant used lots of effort when using one or more aspects of the car designer. |
| T1 | Trust | Trust | A participant trusted the results of the car designer. |
| T2 | | Neutral trust | A participant trusted some of the results of the car designer, but had reservations. |
| T3 | | Distrust | A participant did not trust the results of the car designer. |
| S1 | Design | Seeing + Understanding | A participant expressed interest in finding the correlation between changing parameters and the car's performance. |
| S2 | | Preconceptions | A participant's preconceptions were mentioned as an influence on the design process. |
| S3 | | Score based | A participant used the scores provided to direct their design process. |
| S4 | | Vision based | A participant used visual information to direct their design process. |
| S5 | | Time | A participant mentioned time as a motivating factor in the design process. |
| U1 | User Interface | Sliders | A participant mentioned aspects of interacting with the editor. |
| U2 | | Clutter | A participant mentioned separating out cars while the car designer was running. |
| I1 | Insights | Insights unknown | A participant was confused about the role of insights in the design process or how to use it. |
| I2 | | Insights known | A participant identified one or more characteristics of the insights and/or used the insights in the design process. |

Table 5.7: Responses for the open-ended questions coded according to Table 5.7.

| Participant | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|--|----------|----------|----------|----------|----------------|----------|----------------|----------------|----|----------|----------------|----------------|
| How easy or difficult was the car designer to use? | D1 D4 | D4 | D1 | D1 U2 | D4 U1 S1 | D1 U1 | D1 U1 | D3 S2 | D4 | D1 | D1 U1 U2 | D4 U2 |
| How easy or difficult was it to make a decision from Insights 1 and 2? | D3 S3 | D1 S4 | D1 | D3 | D2 S3 S2 | D1 I2 | D1 S3 I2 | S2 S3 S4 | D1 | I1 | S3 | D1 S3 |
| Was the decision process frustrating for Insights 1 or 2? | F2 S3 | F1 | F1 | F1 | F1 | F1 I2 | F1 S3 S4 | F1 I2 | F1 | F1 | F4 | F2 S5 |
| How much effort did you need to invest while using Insights 1 or 2? | E3 I1 | E1 | E1 I2 | S4 | E1 S3 | E1 S3 | I2 S5 U2 | E1 S3 | E1 | E2 I1 | E1 S3 | E2 |
| Do you trust the results from the car designer? | T3 S1 | T1 | T1 S4 | T1 | T1 | T1 S3 | T1 S3 S4 | T1 | T1 | T1 | T2 | T1 S2 S1 |

using visual information to direct their design process 6 times. Time, or lack of time, was noted as a contributing factor twice. Participants had some trouble understanding the UI, specifically mentioning the sliders 5 times and the amount of information on screen at once 4 times. Given that most participants said they were not frustrated nor did they find the car designer difficult to use, the UI most likely did not play a significant role in a participant's ability to accomplish the task. The *Insights* views themselves were mentioned as well, 4 times by participants who, by their own admission, were unaware of their role, and 6 times by participants who seemed to understand their role in the design process to some extent.

A somewhat mixed view on the effectiveness of MAP-Elites with respect to the four subjective areas investigated was noted in a similar manner to the multiple-choice questions. One group of participants considered the MAP-Elites to be superior while another group thought the randomly selected designs were better. A third group of participants did not

show a preference to either. Participants were grouped by the frequency to which they sided with their preferred design selection method. These groups are compared to the behaviour they exhibited while interacting with the car designer. While some participants said that they preferred the randomly selected designs, their behaviour with regards to their selections and time spent in the MAP-Elites view contradicts their answers. In total, 5 distinct groups are detailed with a brief summary of their preferences and behaviours.

In their post-task answers, Participants 2, 9, 10 and 12 showed a preference for the MAP-Elites. Generally, these participants spent more time using the MAP-Elites view than the random designs view, and also made more selections from the MAP-Elites list than the randomly selected designs. These participants tended to find the car designer easy to use.

Participants 1 and 6 thought the randomly selected designs were better, though they did not make any selections from the available designs. These two participants did not select any designs from the MAP-Elites view either. More time was spent amongst the random designs than the MAP-Elites, indicating that, while they may not have made any selections, they used the random designs as inspiration when interacting with the Editor.

Unlike the first two groups, the behaviour of Participants 3, 4, and 8 did not match their perceived preference. All three of these participants made selections more often from the MAP-Elites list and Participants 3 and 4 spent more time looking at MAP-Elite designs than random designs. All three participants found the car designer easy to use and mentioned that they relied on visual information to guide their design process.

Participants 5 and 11 had mixed opinions when trying to determine their preference for either the MAP-Elites or the randomly selected designs. Despite this, both participants showed an overwhelming preference for the MAP-Elite designs by the number of selec-

tions, and interestingly by the time spent looking at random designs and **not** selecting designs to keep or breed. This could perhaps be due to the random designs and their features taking more time to interpret. Both participants were amongst the top 3 participants in terms of selections from the MAP-Elites.

Lastly, Participant 7 also did not show an overwhelming preference for either the MAP-Elites or the random designs. From their open-ended answers, they are the only participant to compare the views as well, calling the MAP-Elites view “simpler to keep track of” while noting the randomness of the other view. This participant chose not to make any selections from either view and spent most of their time in the Editor.

Open-Ended Questions vs Fitness

A summary of participants’ answers is provided in this section with context to their final fitness improvement as well as performance amongst other participants.

Participant 1; Fitness 125.35; Rank 8

This participant found the car designer neither easy nor difficult, though they did mention that it did get easier as the run progressed. Additionally, they were neutral when it came to frustration. Twice this participant mentioned using the score to optimise their design. This did contradict their statement in the last question when they said that there was no quantitative data to back up the results from the car designer. Their trust was based on being able to see how the changes they made affected the result. Participant 1 was unaware of the role of *Insights* and felt as though they were spending more effort than expected. The participant said they spent more time optimising a single recommendation.

Participant 2; Fitness 163.99; Rank 3

Table 5.8: A breakdown of the number of codes per open-ended question as well as the total across all questions.

| Code | Question 1 | Question 2 | Question 3 | Question 4 | Question 5 | Total |
|------|------------|------------|------------|------------|------------|-------|
| D1 | 5 | 6 | 0 | 0 | 0 | 11 |
| D2 | 2 | 1 | 0 | 0 | 0 | 3 |
| D3 | 1 | 2 | 0 | 0 | 0 | 3 |
| D4 | 5 | 0 | 0 | 0 | 0 | 5 |
| F1 | 0 | 0 | 9 | 0 | 0 | 9 |
| F2 | 0 | 0 | 2 | 0 | 0 | 2 |
| F3 | 0 | 0 | 0 | 0 | 0 | 0 |
| F4 | 0 | 0 | 1 | 0 | 0 | 1 |
| E1 | 0 | 0 | 0 | 7 | 0 | 7 |
| E2 | 0 | 0 | 0 | 2 | 0 | 2 |
| E3 | 0 | 0 | 0 | 1 | 0 | 1 |
| T1 | 0 | 0 | 0 | 0 | 9 | 9 |
| T2 | 0 | 0 | 0 | 0 | 2 | 2 |
| T3 | 0 | 0 | 0 | 0 | 1 | 1 |
| S1 | 1 | 0 | 0 | 0 | 2 | 3 |
| S2 | 1 | 2 | 0 | 0 | 1 | 4 |
| S3 | 0 | 6 | 2 | 4 | 2 | 14 |
| S4 | 0 | 2 | 1 | 1 | 2 | 6 |
| S5 | 0 | 0 | 1 | 1 | 0 | 2 |
| U1 | 5 | 0 | 0 | 0 | 0 | 5 |
| U2 | 3 | 0 | 0 | 1 | 0 | 4 |
| I1 | 1 | 1 | 0 | 2 | 0 | 4 |
| I2 | 0 | 2 | 2 | 2 | 0 | 6 |

Participant 2 found the car designer easier to use as the simulation progressed like Participant 1. This participant did not find the *Insights* frustrating and said that the amount of effort invested in both views was minimal. They also said that the cars progressed well even without their input and could see which cars were progressing well. This participant did trust the results from the car designer.

Participant 3; Fitness 137.57; Rank 7

This participant said the car designer was easy to use and commented on the simplicity of the editor. Making a decision from either *Insights* views was easy as well and did not find either to be frustrating. Their effort level when selecting a design from the *Insights* views was minimal, saying that the amount of designs presented made it easy to select an option. They also mentioned “viable designs” indicating that they used visual cues to optimise designs which led them to trust the results from the car designer.

Participant 4; Fitness 155.13; Rank 5

While this participant found the car designer easy to use, they did comment on the number of cars running simultaneously. Participant 4 found it difficult to choose from both the MAP-Elites and random designs, though, they were not frustrated by the selection process. They called their optimisation process “experimental” focusing mostly on visual information all while trusting the results of the car designer.

Participant 5; Fitness 155.41; Rank 4

As other participants, Participant 5 found the car designer easier to use as time went on. They did mention trying to get used to the user interface in the editor and commented on trying to correlate features with car behaviour. This participant found it neither difficult nor easy to select as design from the *Insights* views; difficulties arose from the amount of

choices available and their preconceptions on what a successful car should look like while the scores beside the designs made it easier to pick. The decision-making process was not frustrating for this participant saying they put in minimal effort and could trust the final results.

Participant 6 Fitness 206.23; Rank 1

Participant 6 found the car designer easy to use with an initial difficulty interacting with the editor. They also said it was easy to select designs from the MAP-Elites and the random lists. Selecting from either list was not frustrating and required minimal effort for them, capitalising on the accompanying score. This participant twice mentioned characteristics of both *Insights* views and trusted the results from the car designer.

Participant 7; Fitness 153.40; Rank 6

This participant found it easy to use the car designer but got confused by which slider in the editor corresponded to which part on the car in much the same way as the previous participant. They said it was easy to choose designs from both lists. Participant 7 is the only participant that specifically (and correctly) differentiated between the two *Insights* views, saying the random designs were “random and messy” while the MAP-Elites had “less choice but seemed better designs.” They also said that the random list required more effort as there were more designs available to choose from and mentioned that the time limit further increased the difficulty of finding a design. This was opposed to the MAP-Elites which were simple to track. Choosing from either list was not frustrating for this participant as the scores helped and they used visual information to see if the car seemed sensible. This same thinking was used to support their trust in the results of the car designer.

Participant 8; Fitness 121.69; Rank 9

Preconceptions made it more difficult for Participant 8 when using the car designer. Initially, they used those preconceptions to edit cars, though changed their tactics a quarter of the way through to simply look at the highest scoring cars. Occasionally, this participant would select cars that were not performing as well to keep up the diversity of the population. This participant used minimal effort and was not frustrated when choosing designs from the MAP-Elite and random lists, opting to focus on the scores for both and saying that both lists worked in tandem with each other. The ability of the best found cars to handle the course was the reason given by Participant 8 for trusting the results.

Participant 9; Fitness 168.15; Rank 2

Despite Participant 2's ranking, they had little to say in each of the open ended questions. They mentioned that the car designer was easier to use as the run progressed, the selection process from both *Insights* views was easy, required minimal effort, and was not frustrating. They also trusted the results from the car designer but did not give a reason.

Participant 10; Fitness 59.82; Rank 12

Likewise, this participant had little to say though ranked last among participants. They found the car designer neither easy nor difficult to use but were not frustrated by the selection process. Twice they mentioned that they were unsure how either view was helping the optimisation process or how they affected the design of the cars. Participant 10 was not sure if they could trust the result from the car designer but did not specify a reason as to why.

Participant 11; Fitness 81.26; Rank 11

Participant 11 found the car designer easy to use but did have some difficulty with the user interface in the editor. They also found that it was difficult to focus on a specific car given the number of cars running at the same time. The role of the *Insights* views were questioned as well. This participant used the scores next to the designs in both lists to select cars and said the frustration was equal for both. Minimal effort and a preference for designs that needed minimal editing was used in selecting designs. Participant 11 did trust the results from the car designer as the cars that they thought would work did indeed work well, but their preconceptions were challenged by some “surprising” designs.

Participant 12; Fitness 121.53; Rank 10

This participant found it difficult to keep track of their design as it was running amongst the other designs, but did find it easier as time went on. The user interface in the editor did also prove to be a little difficult. The score helped Participant 12 in choosing designs from the *Insights* views, though they were frustrated at times as the options in the lists changed. The time pressure also added slightly to the frustration. This participant said that they invested more effort into the both lists near the end, particularly as the fitness stagnated. Participant 12 trusted the car designer some of the time and other times would not as the results would change even though there were no differences to the cars. The trust was based on how the changes made to the car could change the behaviour of the car during the simulation.

Multiple Choice vs Open-Ended

Prior to starting the task of interacting with the car designer, participants answered 5 multiple-choice questions. They answered a further 5 open-ended questions after completing the task. Comparisons between the two answer sets are presented in this section.

Despite Participant 1 not checking the results from automated methods, they did question the results from the car designer. They stated they were not confident using technology which may explain why it took them many generations to understand how the changes affect the car's behaviour.

Both Participant 2 and 3 agreed they are confident using technology, and both found the car designer easy to use by the end. They also both used minimal effort when selecting designs from the *Insights* views while neither preferred to do calculations by hand.

According to their answer to the fourth question in the open-ended question, Participant 4 took an experimental approach to optimisation. This may have to do with their preference to do calculations by hand and double-checking the results from automated methods. They did state, however, that they trusted the results from the car designer.

Participant 5 found the car designer easier to use over time while stating they felt indifferent about their confidence when using technology. Participant 6 is confident using technology, does not double-check results nor preferred to do calculations by hand. This was reflected in their answer to the last open-ended question regarding trust, in which they said the results could not be "falsifiable."

Participant 7, who was able to differentiate between the MAP-Elites and the randomly selected designs, said they were indifferent about their confidence using technology and did double-check the results from automated methods. It is possible that an increased level of scrutiny led them to correctly differentiate between the two *Insights* views.

Participant 8 had preconceptions about how a well-performing should look and initially opted to focus on editing a design to fit those preconceptions. Their confidence using technology and preference to do calculations by hand may indicate an aversion to new

methods.

Participant 9's confidence using technology probably contributed to how easy they found using the car designer and how little effort they needed to invest in selecting designs from both lists despite having the highest number of breeds.

Despite Participant 10 trusting technology with their life, they were not sure they could trust the results from the car designer.

Participant 11 leaned heavily on the score to dictate the optimisation process, while not preferring to do calculations by hand and being confident using technology. It could be inferred that Participant 11 is more trusting of technology, though they were still surprised by some of the well-performing car designs.

Participant 12 also expressed some doubt when looking at the behaviour of the cars that they were simulating. They were not confident using technology but had no problem trusting technology with their life indicating that further explanation would have been helpful in understanding the results coming from the car designer.

Personas

Participants could be categorised from the way users interacted with the car designer. Participants were measured on two separate scales. The first 1D scale is based on how participants went about optimising the cars. Those that focused on one single car, akin to high use of the editor, would be on the exploit side of the scale. The participants that focused on getting as many high-quality designs running simultaneously were on the explore side of the scale. Participants that were on this side of the scale would use the keep and breed features of both *Insights*. The second scale is about whether participants used

visual information to optimise their designs or if they used quantitative data to optimise the design. Participants that used visual information to optimise are on the aesthetic side, while those that used the score to optimise the cars are on the fitness side of the scale. On both scales, it is possible to be between the two extremes. Moreover, participants' personas are not static; they can move from one quadrant to another (and back again) over the course of their runs.

To illustrate where participants sit on the scales and how personas change, 4 plots are shown in this section. Each one is a snapshot of participants' personas after 10 generations. Their personas are determined using results from the survey, the views data, and the keeps and breeds data. Answers from the survey could suggest that a participant is engaged with one or more designs at a time. Their answers could also indicate whether they focused on the score next to the designs or if they chose to design based on how the cars looked. Data from the views showed where participants were spending time, either in the *Editor* on a single design or the *Insights* views on multiple designs at once. Likewise, the amount of times participants chose to keep or breed a design could indicate if they preferred exploring many designs or exploiting a single design.

Initially, participants generally focused on a single design. This is evidenced by the general trend to use the editor shortly after starting the task. The exception to this is Participant 9 who very quickly makes use of the two *Insight* views containing the MAP-Elites and the random designs. In the first 10 generations, Participants 2, 3, 4 and 8 rely mostly on how the car looks and their preconceptions to guide their design process. Participants 5, 11, 12, 6 and 1 rely on the score provided to guide their design process in ascending order of dependence. Both Participants 7 and 10 take a balanced approach between how the car looks and the score provided to edit their designs. Likewise, Participant 9 uses both

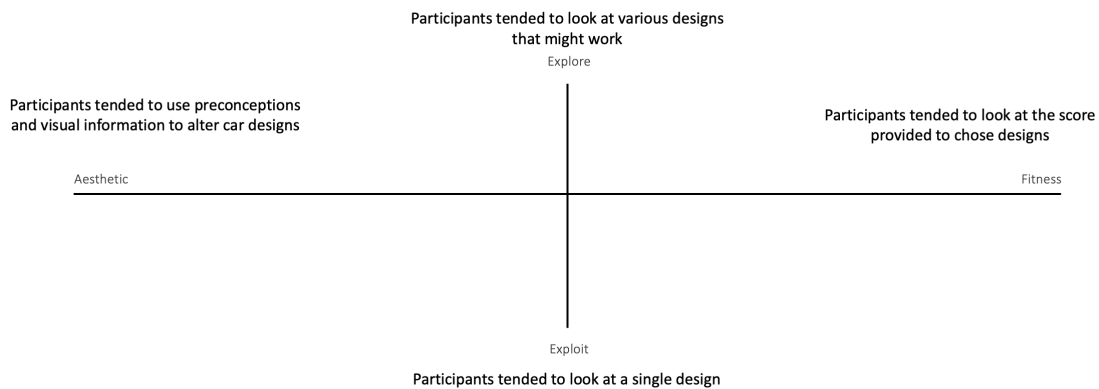


Figure 5.7: The 2D scale on which participants' personas are displayed. On the left, the participants use visual information provided by the car designer while participants on the right use the score/fitness of the designs to inform the design process. At the top, participants explore a variety of designs while participants at the bottom focus on improving a single design.

attributes to select the various designs for the next generation.

In the following 10 generations, two participants make a noticeable change in their approach to the design process. Participant 5 and 8 both stray from focusing on trying to improve a single design, and instead focus on widening their search within the design space. Their approach is guided more by the fitness of the design than how the car looks. Participant 1 also begins to shift away from focusing on one design by looking at designs from the *Insight* tabs, though does not select any of the designs to be used in future generations. The other participants maintained their design approach.

For participants that had not engaged with the *Insight* views up until the halfway point, a reminder was given to use all available views during the design process. Participants 3 and 4 adopted a more exploratory approach. Participant 3 made more selections during this time than Participant 4 but spent less time in the *Insight* views. Participant 3 also said that "it was easy to select a viable design" whereas Participant 4 looked for "overall

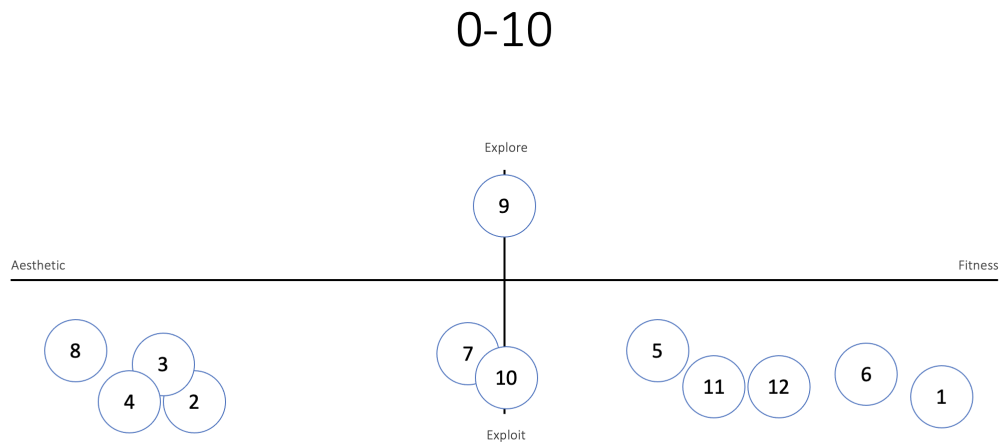


Figure 5.8: The personas of the participants during the first 10 generations. Most participants were focused on a single design with the exception of Participant 9.

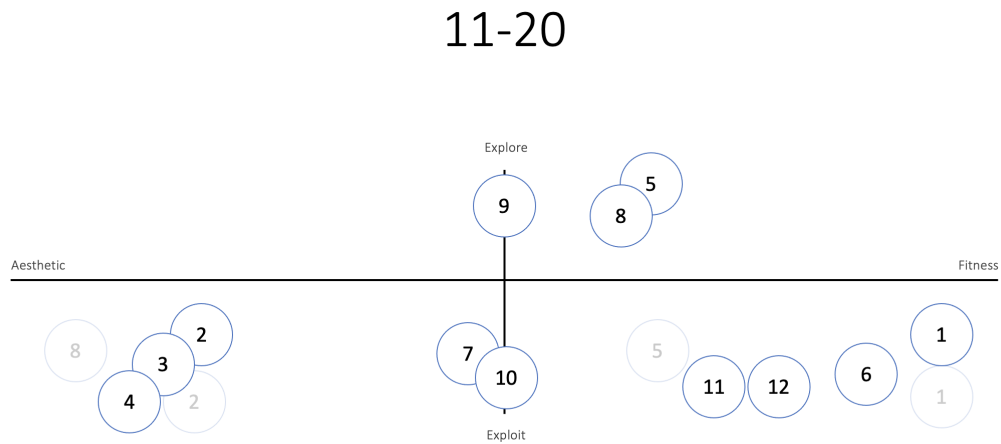


Figure 5.9: The personas of the participants between generations 11 and 20. Participants begin to use the *Insights* views more with Participants 5 and 8 shifting their design optimisation strategies considerably.

21-30

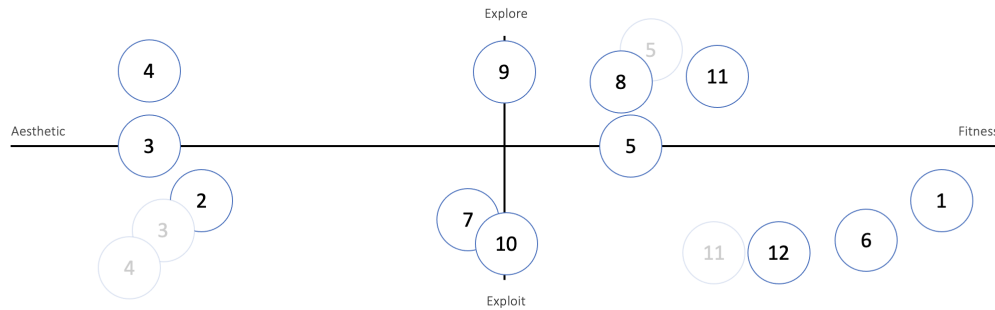


Figure 5.10: The personas of the participants between generations 21 and 30. Participants use the *Insights* views much more and begin to focus on many more designs during their design process.

ideas” to support their design process. Participant 5 shifted back towards an approach that focused more on a single design after engaging heavily with the two different design lists. A more exploratory approach was taken by Participant 11 which led to a noticeable increase in fitness improvement.

During the last 10 generations, both Participants 2 and 3 selected more designs from the MAP-Elites and the randomly selected designs. This is in contrast to Participants 1 and 4 who went back to their original design processes and focused primarily on one design. Participant 5 also had a more exploratory approach during the final generations.

Generally, participants kept their approaches to the design process when considering only whether they used visual information or the score to guide their design process. However, participants tended to shift towards an exploratory approach as they engaged with the car designer more. The participant with the biggest change in design approach

31-40




Figure 5.11: The personas of the participants in the final 10 generations. Small changes in participant behaviour are still observed, such as Participants 1-5, though they are only small changes in design philosophy.

Overall




Figure 5.12: The overall personas of participants across their entire simulation run.

was Participant 8 who started in the lower left quadrant and moved to the upper right quadrant. Other participants (2, 3, 5, 11) had gradual changes to their approaches, usually preceded by stagnation in the improvement of their best fitness. Participant 4 changed their approach temporarily as well for a short period of time though ultimately chose to focus on editing a single design. Participants 1, 6, 7, 9, 10, and 12, did not change their approaches much during the entire 40 generations. For some, such as Participants 6 and 9, sticking to their approach yielded designs with relatively high fitness values. Their approaches were diametrically opposed on the Explore-Exploit axis, but still resulted in the more optimal designs. In contrast, Participants 10 and 12 stayed with their approaches but had designs that had relatively low fitness values.

5.5 Discussion

The limitations of this approach are explored and the results are interpreted in context to the research question.

5.5.1 Fitness

Among the participant pool, there is evidence that introducing a human aspect in the design loop improves the performance of the algorithm. Participants were generally able to outperform the computer-only simulations over the 40 generations. Their performance over the first 15 generations, however, was similar. Participants were able to continue improving their designs over the entire 40 generations with a decrease in the rate of improvement while the computer-only simulations rate of improvement decreased much more. 7 participants were able to design cars that had a higher fitness than the best-performing computer-only run while the average fitness of the 20 computer-only runs was bettered by 83% of participants (N=10). Participants spent most of their time in the

Editor view and led to improvements for all participants, which emphasises the need for human input in the optimisation process. Participants were also much more likely to select the MAP-Elites algorithm to improve their design, with keeps and breeds from the MAP-Elites view contributing to an increase in fitness in all but one participants who used at least one design from the MAP-Elites view. These results are also supported quantitatively by the results from the online study which found that participants that interacted with the MAP-Elites algorithm did better than those that did not.

5.5.2 Survey

Results from the survey show that there is still some confusion about the role of algorithms like MAP-Elites in the design process. Participants had mixed views on the effectiveness of the MAP-Elites algorithm, especially when asked whether they preferred the randomly selected designs or those in the MAP-Elites view. Participants 2, 7, 9, 10 and 12 all agreed that the MAP-Elites were better than random designs across the four questions asked in subjective satisfaction, while Participants 3 and 5 said random designs were better when answering the same questions. The confusion about the role of MAP-Elites in the design process is also apparent for the 5 different participants whose behaviour when using the car designer did not match their survey answers.

The different strategies employed during the optimisation process and described by participants allowed them to be grouped together in personas. Over the 40 simulations, participants gravitated towards a more exploratory approach by using suggestions from the *Insights* in increasing frequency as the simulations progressed. Participants were more stubborn when considering which information they acted upon when optimising their designs; those who focused on the fitness value of the cars remained focused on the fitness values. Quality-diversity algorithms such as MAP-Elites should be well suited for both

of these groups of participants as it provides participants who prefer to optimise based on looks a variety of choices and ranks designs for participants who prefer to use fitness values to guide their optimisation process.

5.5.3 Limitation of the Methodology

As with all surveys, it is impossible to completely remove the possibility that questions are going to be misinterpreted. This limitation was mitigated by being present next to the participant while they answered the survey should they require clarification.

One of the assumptions made at the beginning of the study was that participants are inherently aware of the basic principles of how a car works. Of the 12 participants, none raised any concerns about understanding this part of the study. In the open-ended answers, some participants noted their preconceptions of how cars worked and how it influenced their design process.

Another assumption made was that by studying engineering or having done so in the past would bring a level of competence in design engineering and the results would therefore be more applicable for professional design engineers. The participants had a wide variety of experience in engineering and could have potentially swayed the results.

5.6 Conclusion

At the beginning of Section 5.3, the research question asked whether MAP-Elites have a positive influence on fitness. Looking at the results from the perspective of the participant runs compared to the runs which did not receive any input from a human, there is evidence to say that the MAP-Elites do indeed have a positive influence. Half of the participants were able to design a car with a higher fitness than the best run from the computer-only

runs, while 10 of 12 participants were able to design a car with a higher fitness than the average of the computer-only runs. The best performing participant, however, did not make any selections from either the MAP-Elites or the random designs which puts into question the necessity for MAP-Elites. While this may indicate that the MAP-Elites were not necessary, this participant did spend time viewing the MAP-Elites which was followed by an improvement in fitness. This trend was not unique to this participant, as many others also had increases in fitness. Likewise, increases in fitness were also seen following keeps and breeds.

Participant behaviour did not always match their perceptions about the *Insights* views as stated in their post-task survey answers, as participants tended to select from the MAP-Elites for breeds and keeps more often than the random designs, and spent more time in the MAP-Elites view than amongst the random designs.

The online study [8] provided much needed statistical evidence that MAP-Elites do indeed have a positive influence on the fitness. The targeted user study presented in this work gives context to the behaviour of participants and demonstrates that MAP-Elites are a preferred method to design optimisation than random designs.

Chapter 6

Conclusion

The objectives of this thesis were to apply optimisation techniques to an engineering design problem and to investigate the relationship between engineers and algorithms. A novel ducted winglet to be used on wind turbines was selected as a design case for the engineering design problem. The relationship between engineers and algorithms was explored using a car designer game.

The novel ducted winglet was compared to a non-ducted winglet and an unmodified reference blade in Chapter 2 to determine the effect of the duct on the total drag as well as the induced drag. The novel ducted winglet had the lowest induced drag across the range of angles of attack compared to the other two geometries though the non-ducted winglet had lower overall drag.

Once the novel ducted winglet had shown a reduction in drag compared to the reference blade and lower induced drag compared to the non-ducted winglet, an optimisation study was performed in Chapter 3 to further reduce the total drag and induced drag.

The optimisation study led to marginal improvements in both the total drag and induced drag, though limitations of the techniques used on such a complex design problem raised concerns about the robustness of these techniques. The effects of accurate rank order and fidelity of the governing equations were studied in Chapter 4. Small changes in both the rank order and the equation fidelity did not affect optimisation success significantly and brings confidence to the techniques used in Chapter 5.

A user study investigating the relationship between engineers and an algorithm called MAP-Elites was conducted as part of a larger study. Participants who used the MAP-Elites algorithm were more likely to find better designs though participants were not always aware of the role of the underlying algorithm.

6.1 Contributions and Findings

The contributions and findings of this thesis are summarised in this section. The contributions are separated into each objective.

6.1.1 Engineering Design Optimisation

The novel ducted winglet's effectiveness as an aerodynamic device on a wind turbine blade was investigated as part of the work initiated by Fourth Dimensional Aerospace Technology. Prior to this work, the novel ducted winglet had not been compared to an unmodified blade or a standard non-ducted winglet. An in-house CFD solver was used to simulate the expected conditions to be experienced by the winglet and Trefftz plane analysis was used to compare the induced drag of the three tested geometries. The novel ducted winglet had a lower induced drag for 7 out of 8 angles of attack tested. Furthermore, an updated explanation from the patent is given for some of the phenomena

which occur at the inlet, outlet, and throughout the duct.

An optimisation study of the novel ducted winglet was conducted. Prior to this work, an optimisation study had not been conducted on the novel ducted winglet, though it would have been necessary to conduct the optimisation study regardless as the application domain had changed to wind turbines. Originally, the novel ducted winglet was intended to be used on commercial aircraft where the flow characteristics differ from wind turbines.

A shortened simulation cycle was also proposed based on results from two algorithm studies. Small changes in the rank order have a limited effect on the success of the optimisation. EAs only require the rank order to be correct to find a more optimal solution and do not need solution values to be accurate to small orders of magnitude. By taking advantage of the manner in which EAs optimise, the simulation time needed is reduced to how long it takes to establish the correct rank.

6.1.2 Human-AI Collaboration

The MAP-Elites algorithm was shown to help participants of a user study create better designs compared to computer-created designs. A small scale user study with 12 participants was created as part of a larger study investigating the effectiveness of MAP-Elites in the design optimisation process.

Participants were slightly more likely to look at MAP-Elite designs than randomly selected designs, but preferred to spend much more time editing a single design; half of the participants' time was spent using the editor functionality of the car designer. Participants spent 45% more time in the MAP-Elites view compared to the control. Generally, the proportional amount of time spent looking at both the MAP-Elites and random designs increased as the number of iterations increased.

MAP-Elite breeds were much more likely to be chosen by participants than MAP-Elite keeps as well as keeps and breeds from randomly selected designs. Participants selected their own edited designs to be kept for the next generation, though the checkbox for this was persistent and was selected until explicitly deselected by participants. The other keeps and breeds were non-persistent and had to be selected anew each time, indicating a conscious decision by participants to select the MAP-Elite breeds.

Participants' preference between MAP-Elites and the control did not always align with their behaviour. Five participants preferred the randomly selected designs according to their survey answers, but the time they spent looking at the MAP-Elites and the number of selections in terms of MAP-Elite keeps and breeds was higher than the control. This suggests that while MAP-Elites are a preferred method for optimisation, more work is needed to influence the subjective accuracy, diversity, satisfaction, and novelty of the MAP-Elites algorithm.

As participants interacted with the car designer, they were more likely to shift away from exploiting one design to exploring multiple designs. Participants could be grouped into different personas based on their optimisation strategies. Some participants focused on exploiting a single design while others preferred to explore the design space by looking at multiple designs at once. Likewise, some participants focused solely on the fitness value provided for each car as a measure of success, while other participants focused on the car's look and their preconceptions to guide their optimisation process. The strategies of participants also evolved over time, with participants trending towards an exploratory approach.

6.2 Future Work

Key areas of future work have been identified in each main chapter. These ideas are discussed in this section.

6.2.1 Chapter 2

In Chapter 2, the effectiveness of a novel ducted winglet was tested and compared to the effectiveness of two other geometries using CFD. It would be beneficial to further investigate the effectiveness of these different geometries in a wind tunnel. Additional data can be gathered from physical models and the results can be compared to CFD results. Validation of the CFD results would also help support findings in Chapter 3 where different geometries were tested in an effort to optimise the shape of the novel ducted winglet. Ideally, physical models of the three different geometries would be tested on wind turbines to capture more data and determine key characteristics of each geometry. These key characteristics include the minimum air velocity needed for power generation, minimum air velocity for peak power generation, and overall efficiency.

6.2.2 Chapter 3

Different optimisation methods were explored in this chapter and additional work could be done for each method. Additional work with regards to the parameterisation can also be done to further investigate different optimisation approaches when using CFD as well as to see the degree to which humans can be implemented in the optimisation loop. The constraints implemented on Latin hypercube sampling were explored with tradeoffs explained about various approaches. A study investigating the validity of non-uniformly distributed sampling compared to non-uniformly sized hypercubes would be a novel contribution and would be especially beneficial when applied to optimisation

problems such as the ducted winglet. The effects of non-uniformity should be tested against random sampling for rigour and completeness. Further work with regards to the use of evolutionary optimisation can also be done to test the robustness of EAs and to further optimise the ducted winglet. Increasing the population and the number of generations would be a logical start. Hyperparameter tuning could also yield interesting results. Bayesian optimisation would benefit greatly from hyperparameter tuning as well and could be performed in conjunction with additional iterations and an increased starting population as well. Parameters can also be reintroduced. Originally, 18 parameters were considered which were then revised down to 4. Steadily increasing the number of parameters as more information is acquired through additional iterations could further increase the efficiency of the novel ducted winglet. A challenge may be to properly parameterise the previous designs.

6.2.3 Chapter 4

To better represent the case study seen in Chapter 2 and 3, a closer representation of the underlying physics would be beneficial. Extending the rank change study to 2-dimensional and 3-dimensional flow would allow the results to be more closely linked to the results of the previous chapters. In the same way, the hyperparameters of the CFD code which influence how the governing equations are solved could be changed to see how they affect the optimisation process in both 2 and 3-dimensional flow.

6.2.4 Chapter 5

The user study provided key insights into how people interact with optimisation algorithms. There are two routes which would be interesting to explore in the future. The first would be to apply the same user study in a more applied context. For example, the

behaviour of experienced engineers and designers working on real-world optimisation problems could be investigated in a longitudinal study. Giving participants more time with a problem and to those with more experience might prevent account for unfamiliarity bias. Secondly, developing a better algorithm which is better suited for human intervention is worth exploring. This would address some of the shortfalls of MAP-Elites seen in the online and in-person user studies.

Chapter 7

Publications and Conference

Presentations

1. Jakub Vincalek, Sean Walton, and Ben Evans. It's the journey not the destination: Building genetic algorithms practitioners can trust. In *14th International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control*, 2021
2. Jakub Vincalek, Sean Walton, and Ben Evans. It's the journey not the destination: building genetic algorithms practitioners can trust. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO '21*. ACM, July 2021
3. Jakub Vincalek, Sean Walton, and Ben Evans. A user-centered approach to evolutionary algorithms and their use in industry. *Cogent Engineering*, 8(1), November 2021

4. Jakub Vincalek, Sean Walton, and Ben Evans. On the use of CFD for the comparison of a novel wing tip design to a standard wing tip in the field of wind turbine aerodynamics. In *UK Association for Computational Mechanics School and Annual Conference, 2022*
5. Sean P. Walton, Ben Evans, Ben Smith, and Jakub Vincalek. How engineers use evolution to invent things. *Frontiers for Young Minds*, 10, August 2022
6. Jakub Vincalek, Sean Walton, and Ben Evans. Optimising a novel winglet for use on a wind turbine blade. In *AI-OPT. The University of Melbourne and Monash University, ARC Training Centre in Technologies, Integrated Methodologies and Applications (OPTIMA), 2022*
7. Jakub Vincalek, Sean Walton, and Ben Evans. Evaluating the effect of a ducted winglet on the induced drag of wind turbine blade using CFD and trefftz plane analysis. *Engineering with Computers*, April 2023
8. Jakub Vincalek, Sean Walton, and Ben Evans. Understanding the effects of a ducted winglet on wake and induced drag using CFD. In *22nd Computational Fluids Conference, 2023*
9. Jakub Vincalek, Sean Walton, and Ben Evans. Optimisation of a novel ducted winglet on a wind turbine blade for reduced wake. In *15th International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control, 2023*
10. Sean Walton, Ben Evans, Alma A. M. Rahat, James Stovold, and Jakub Vincalek. Does mapping elites illuminate search spaces? A large-scale user study of map-elites applied to human-AI collaborative design. *Submitted to: Journal of Computational Design and Engineering, 2024*

Bibliography

- [1] John K. Kaldellis and D. Zafirakis. The wind energy (r)evolution: A short review of a long history. *Renewable Energy*, 36(7):1888–1901, July 2011.
- [2] James F Manwell, Jon G McGowan, and Anthony L Rogers. *Wind energy explained: theory, design and application*. John Wiley & Sons, 2010.
- [3] John Anderson. *Fundamentals of Aerodynamics ISE*. McGraw-Hill Education, Columbus, OH, 7 edition, 2023.
- [4] Edward N Tinoco. Cfd uncertainty and validation for commercial aircraft applications. In *NATO Symposium AVT*, volume 147, pages 1–36, 2007.
- [5] Jakub Vincalek, Sean Walton, and Ben Evans. It’s the journey not the destination: building genetic algorithms practitioners can trust. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO ’21*. ACM, July 2021.
- [6] Sean P. Walton, Alma A. M. Rahat, and James Stovold. Evaluating mixed-initiative procedural level design tools using a triple-blind mixed-method user study. *IEEE Transactions on Games*, 14(3):413–422, September 2022.

- [7] Jakub Vincalek, Sean Walton, and Ben Evans. A user-centered approach to evolutionary algorithms and their use in industry. *Cogent Engineering*, 8(1), November 2021.
- [8] Sean Walton, Ben Evans, Alma A. M. Rahat, James Stovold, and Jakub Vincalek. Does mapping elites illuminate search spaces? A large-scale user study of map-elites applied to human-AI collaborative design. *Submitted to: Journal of Computational Design and Engineering*, 2024.
- [9] Jakub Vincalek, Sean Walton, and Ben Evans. Evaluating the effect of a ducted winglet on the induced drag of wind turbine blade using CFD and trefftz plane analysis. *Engineering with Computers*, April 2023.
- [10] John J. Smith. Fluid flow control duct for an aerofoil, 2017.
- [11] Richard T Whitcomb. A design approach and selected wind tunnel results at high subsonic speeds for wing-tip mounted winglets. 1976.
- [12] Luciano Demasi, Giovanni Monegato, and Rauno Cavallaro. Minimum induced drag theorems for multiwing systems. *AIAA Journal*, 55(10):3266–3287, 2017.
- [13] TK Guha and R Kumar. Characteristics of a wingtip vortex from an oscillating winglet. *Experiments in Fluids*, 58(1):8, 2017.
- [14] Marco KleinHeerenbrink, L Christoffer Johansson, and Anders Hedenström. Multi-cored vortices support function of slotted wing tips of birds in gliding and flapping flight. *Journal of The Royal Society Interface*, 14(130):20170099, 2017.
- [15] Gautham Narayan and Bibin John. Effect of winglets induced tip vortex structure on the performance of subsonic wings. *Aerospace Science and Technology*, 58:328–340,

- 2016.
- [16] Michael DP Bolzon, Richard M Kelso, and Maziar Arjomandi. Formation of vortices on a tubercled wing, and their effects on drag. *Aerospace Science and Technology*, 56:46–55, 2016.
- [17] MA Aldheeb, W Asrar, AA Omar, A Altaf, and E Sulaeman. Effect of directionally porous wing tip on tip vortex. *J Appl Fluid Mech*, 13(2):651–665, 2020.
- [18] Sidaard Gunasekaran and Tim Gerham. Effect of chordwise slots on aerodynamic efficiency and wingtip vortex. *AIAA Journal*, 56(12):4752–4767, 2018.
- [19] Andrew L Heyes and David AR Smith. Modification of a wing tip vortex by vortex generators. *Aerospace science and Technology*, 9(6):469–475, 2005.
- [20] Oubay Hassan and Ken Morgan. *FLITE SYSTEM Theoretical Manual*. Swansea University, 4 edition, 2009.
- [21] B. Evans, T. Morton, L. Sheridan, O. Hassan, K. Morgan, J. W. Jones, M. Chapman, R. Ayers, and I. Niven. Design optimisation using computational fluid dynamics applied to a land-based supersonic vehicle, the bloodhound ssc. *Structural and Multidisciplinary Optimization*, 47(2):301–316, July 2012.
- [22] J Peraire, J Peiró, and K Morgan. *Advancing Front Grid Generation*. CRC Press, December 1998.
- [23] Xi Zou, Sui Bun Lo, Ruben Sevilla, Oubay Hassan, and Kenneth Morgan. The generation of 3d surface meshes for nurbs-enhanced fem. *Computer-Aided Design*, page 103653, 2023.

- [24] N. P. Weatherill and O. Hassan. Efficient three-dimensional delaunay triangulation with automatic point creation and imposed boundary constraints. *International Journal for Numerical Methods in Engineering*, 37(12):2005–2039, jun 1994.
- [25] P. Spalart and S. Allmaras. A one-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics, January 1992.
- [26] Dassault Systemes BIOVIA. Solidworks, 2023.
- [27] Scott Monsch, Richard Figliola, Ernest Thompson, and Jose Camberos. Computation of induced drag for 3d wing with volume integral (trefftz plane) technique. In *45th AIAA Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics, January 2007.
- [28] M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239, May 1979.
- [29] J.C. Helton, F.J. Davis, and J.D. Johnson. A comparison of uncertainty and sensitivity analysis results obtained with random and latin hypercube sampling. *Reliability Engineering & System Safety*, 89(3):305?330, September 2005.
- [30] Thomas Bartz?Beielstein, Jürgen Branke, Jörn Mehnen, and Olaf Mersmann. Evolutionary algorithms. *WIREs Data Mining and Knowledge Discovery*, 4(3):178?195, April 2014.
- [31] Tobias Blickle and Lothar Thiele. A comparison of selection schemes used in evolutionary algorithms. *Evolutionary Computation*, 4(4):361?394, December 1996.

- [32] Samir Mahfoud. *Niching methods*, page 87–92. Taylor & Francis, November 2000.
- [33] Zbigniew Skolicki and Kenneth De Jong. *Improving Evolutionary Algorithms with Multi-representation Island Models*, page 420–429. Springer Berlin Heidelberg, 2004.
- [34] Danesh Tarapore, Jeff Clune, Antoine Cully, and Jean-Baptiste Mouret. How do different encodings influence the performance of the MAP-elites algorithm? In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*. ACM, jul 2016.
- [35] Ting-Song Du, Xian-Ting Ke, Jia-Gen Liao, and Yan-Jun Shen. Dslc-foa: improved fruit fly optimization algorithm for application to structural engineering design optimization problems. *Applied Mathematical Modelling*, 55:314–339, 2018.
- [36] Marco Locatelli. A note on the griewank test function. *Journal of global optimization*, 25(2):169–174, 2003.
- [37] Franciszek Seredynski, Albert Y Zomaya, and Pascal Bouvry. Function optimization with coevolutionary algorithms. In *Intelligent Information Processing and Web Mining*, pages 13–22. Springer, 2003.
- [38] Guopu Zhu and Sam Kwong. Gbest-guided artificial bee colony algorithm for numerical function optimization. *Applied mathematics and computation*, 217(7):3166–3173, 2010.
- [39] Saeed Motiian and Hamid Soltanian-Zadeh. Improved particle swarm optimization and applications to hidden markov model and ackley function. In *2011 IEEE International Conference on Computational Intelligence for Measurement Systems and Applications (CIMSA) Proceedings*, pages 1–4. IEEE, 2011.

- [40] V Alcantar, S Ledesma, SM Aceves, E Ledesma, and A Saldana. Optimization of type iii pressure vessels using genetic algorithm and simulated annealing. *International Journal of Hydrogen Energy*, 42(31):20125–20132, 2017.
- [41] Kalyanmoy Deb. Optimal design of a welded beam via genetic algorithms. *AIAA journal*, 29(11):2013–2015, 1991.
- [42] Gaganpreet Kaur and Sankalap Arora. Chaotic whale optimization algorithm. *Journal of Computational Design and Engineering*, 5(3):275–284, 2018.
- [43] David H Wolpert, William G Macready, et al. No free lunch theorems for search. Technical report, Technical Report SFI-TR-95-02-010, Santa Fe Institute, 1995.
- [44] Mahnaz Khajeh, Pedram Payvandy, and Seyyed Javad Derakhshan. Fashion set design with an emphasis on fabric composition using the interactive genetic algorithm. *Fashion and Textiles*, 3(1):8, 2016.
- [45] Nutthanon Leelathakul and Sunisa Rimcharoen. Generating kranok patterns with an interactive evolutionary algorithm. *Applied Soft Computing*, 89:106121, 2020.
- [46] Majid Farzaneh and Rahil Mahdian Toroghi. Music generation using an interactive evolutionary algorithm. In *Mediterranean Conference on Pattern Recognition and Artificial Intelligence*, pages 207–217. Springer, 2019.
- [47] Sean P. Walton, Alma A. M. Rahat, and James Stovold. Mixed-initiative procedural content generation using level design patterns and interactive evolutionary optimisation, 2020.
- [48] Hideyuki Takagi. Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE*, 89(9):1275–1296,

- 2001.
- [49] Tijana Vuletic, Alex Duffy, Laura Hay, Chris McTeague, Laura Pidgeon, and Madeleine Greal. The challenges in computer supported conceptual engineering design. *Computers in Industry*, 95:22–37, 2018.
- [50] Alexandra Melike Brintrup, Jeremy Ramsden, Hideyuki Takagi, and Ashutosh Tiwari. Ergonomic chair design by fusing qualitative and quantitative criteria using interactive genetic algorithms. *IEEE Transactions on Evolutionary Computation*, 12(3):343–354, 2008.
- [51] Guanci Yang. Game theory-inspired evolutionary algorithm for global optimization. *Algorithms*, 10(4):111, 2017.
- [52] Matthew B Johns, Herman A Mahmoud, David J Walker, Nicholas DF Ross, Edward C Keedwell, and Dragan A Savic. Augmented evolutionary intelligence: combining human and evolutionary design for water distribution network optimisation. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1214–1222, 2019.
- [53] Seyedali Mirjalili and Andrew Lewis. The whale optimization algorithm. *Advances in engineering software*, 95:51–67, 2016.
- [54] Dervis Karaboga. An idea based on honey bee swarm for numerical optimization. Technical report, Technical report-tr06, Erciyes university, engineering faculty, computer science, 2005.
- [55] Xin-She Yang. Firefly algorithm, stochastic test functions and design optimisation. *International journal of bio-inspired computation*, 2(2):78–84, 2010.

- [56] Xin-She Yang and Suash Deb. Cuckoo search via lévy flights. In *2009 World congress on nature & biologically inspired computing (NaBIC)*, pages 210–214. IEEE, 2009.
- [57] S Walton, O Hassan, K Morgan, and MR Brown. Modified cuckoo search: a new gradient free optimisation algorithm. *Chaos, Solitons & Fractals*, 44(9):710–718, 2011.
- [58] Alexander Brownlee and John Woodward. Why we fell out of love with algorithms inspired by nature. The Conversation Trust, 2015.
- [59] Mohamed A Tawhid and Vimal Savsani. A novel multi-objective optimization algorithm based on artificial algae for multi-objective engineering design problems. *Applied Intelligence*, 48(10):3762–3781, 2018.
- [60] Airbus. Telecommunications, 2020.
- [61] Salima Berrezzoug, Abdelmadjid Boudjemai, and Fethi Tarik Bendimerad. Interactive design and multidisciplinary optimization of geostationary communication satellite. *International Journal on Interactive Design and Manufacturing (IJIDeM)*, 13(4):1519–1540, 2019.
- [62] A Kaveh, M Khanzadi, M Alipour, and M Rastegar Moghaddam. Construction site layout planning problem using two new meta-heuristic algorithms. *Iranian Journal of Science and Technology, Transactions of Civil Engineering*, 40(4):263–275, 2016.
- [63] G Kannan, A Noorul Haq, and M Devika. Analysis of closed loop supply chain using genetic algorithm and particle swarm optimisation. *International Journal of Production Research*, 47(5):1175–1200, 2009.
- [64] Yongjian Li, Jian Chen, and Xiaoqiang Cai. Heuristic genetic algorithm for capacitated production planning problems with batch processing and remanufacturing.

- International Journal of Production Economics*, 105(2):301–317, 2007.
- [65] Chun Jin, Hanping Mao, Yong Chen, Qiang Shi, Qirui Wang, Guoxing Ma, and Yong Liu. Engineering-oriented dynamic optimal control of a greenhouse environment using an improved genetic algorithm with engineering constraint rules. *Computers and Electronics in Agriculture*, 177:105698, 2020.
- [66] Michael M Joly, Tom Verstraete, and Guillermo Paniagua. Integrated multifidelity, multidisciplinary evolutionary design optimization of counterrotating compressors. *Integrated Computer-Aided Engineering*, 21(3):249–261, 2014.
- [67] Earll M Murman, Myles Walton, and Eric Rebentisch. Challenges in the better, faster, cheaper era of aeronautical design, engineering and manufacturing. 2016.
- [68] Gianni D’Angelo and Francesco Palmieri. Knowledge elicitation based on genetic programming for non destructive testing of critical aerospace systems. *Future Generation Computer Systems*, 102:633–642, 2020.
- [69] Danil Nagy. Nature-based hybrid computational geometry system for optimizing the interior structure of aerospace components. In *ACM SIGGRAPH 2017 Talks*, pages 1–2. 2017.
- [70] Andrei V Popov, Lucian T Grigorie, Ruxandra Botez, Mahmood Mamou, and Youssef Mébarki. Real time morphing wing optimization validation using wind-tunnel tests. *Journal of Aircraft*, 47(4):1346–1355, 2010.
- [71] Carlos Simão Ferreira and Ben Geurts. Aerofoil optimization for vertical-axis wind turbines. *Wind Energy*, 18(8):1371–1385, 2015.
- [72] Seyedali Mirjalili, Amir H Gandomi, Seyedeh Zahra Mirjalili, Shahrzad Saremi,

- Hossam Faris, and Seyed Mohammad Mirjalili. Salp swarm algorithm: A bio-inspired optimizer for engineering design problems. *Advances in Engineering Software*, 114:163–191, 2017.
- [73] Y Rahmad, MD Robani, PS Palar, and LR Zuhail. Single-and multi-objective optimization of a low-speed airfoil using genetic algorithm. In *AIP Conference Proceedings*, volume 2226, page 020005. AIP Publishing LLC, 2020.
- [74] Xinqiang Liu and Weiliang He. Airfoil optimization design based on the pivot element weighting iterative method. *Algorithms*, 11(10):163, 2018.
- [75] Mingxu Yi, Yalin Pan, Jun Huang, Lifeng Wang, and Dawei Liu. A comprehensive optimization design method of aerodynamic, acoustic, and stealth of helicopter rotor blades based on genetic algorithm. *Mathematical Problems in Engineering*, 2019, 2019.
- [76] Akira Oyama, Shigeru Obayashi, and Takashi Nakamura. Real-coded adaptive range genetic algorithm applied to transonic wing optimization. *Applied Soft Computing*, 1(3):179–187, 2001.
- [77] Shigeru Obayashi, Daisuke Sasaki, Yukishiro Takeguchi, and Naoki Hirose. Multi-objective evolutionary computation for supersonic wing-shape optimization. *IEEE transactions on evolutionary computation*, 4(2):182–187, 2000.
- [78] Andreea Koreanschi, Oliviu Sugar Gabor, Joran Acotto, Guillaume Brianchon, Gregoire Portier, Ruxandra Mihaela Botez, Mahmoud Mamou, and Youssef Mebarki. Optimization and design of an aircraft’s morphing wing-tip demonstrator for drag reduction at low speed, part i–aerodynamic optimization using genetic, bee colony and gradient descent algorithms. *Chinese Journal of Aeronautics*, 30(1):149–163, 2017.

- [79] Abdelwahid Boutemedjet, Marija Samardžić, Lamine Rebhi, Zoran Rajić, and Takeddine Mouada. Uav aerodynamic design involving genetic algorithm and artificial neural network for wing preliminary computation. *Aerospace Science and Technology*, 84:464–483, 2019.
- [80] Hamid Isakhani, Caihua Xiong, Shigang Yue, and Wenbin Chen. A bioinspired airfoil optimization technique using nash genetic algorithm. In *2020 17th International Conference on Ubiquitous Robots (UR)*, pages 506–513. IEEE, 2020.
- [81] Alexander Gehrke, De Guyon-Crozier, Karen Mulleners, et al. Genetic algorithm based optimization of wing rotation in hover. *Fluids*, 3(3):59, 2018.
- [82] Edward M Fagan, O De La Torre, Sean B Leen, and Jamie Goggins. Validation of the multi-objective structural optimisation of a composite wind turbine blade. *Composite Structures*, 204:567–577, 2018.
- [83] Vicente Díaz-Casàs, Jose-Antonio Becerra, Fernando Lopez-Peña, and Richard J. Duro. Wind turbine design through evolutionary algorithms based on surrogate cfd methods. *Optimization and Engineering*, 14(2):305?329, February 2012.
- [84] M. Sessarego, K.R. Dixon, D.E. Rival, and D.H. Wood. A hybrid multi-objective evolutionary algorithm for wind-turbine blade optimization. *Engineering Optimization*, 47(8):1043?1062, August 2014.
- [85] Alma A. M. Rahat, Richard M. Everson, and Jonathan E. Fieldsend. Alternative infill strategies for expensive multi-objective optimisation. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '17*. ACM, July 2017.
- [86] Mark Drela. *XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils*,

- pages 1–12. Springer Berlin Heidelberg, 1989.
- [87] Robert Kubler. Simple heuristics for the traveling salesman problem in python, 2023.
- [88] Donald L. Miller and Joseph F. Pekny. Exact solution of large asymmetric traveling salesman problems. *Science*, 251(4995):754–761, February 1991.
- [89] Sunith Bandaru and Kalyanmoy Deb. A parameterless-niching-assisted bi-objective approach to multimodal optimization. In *2013 IEEE Congress on Evolutionary Computation*. IEEE, June 2013.
- [90] Bernardo Silva. Multivariate taylor series expansion, 2020.
- [91] Safayet Hossain, Muhammad Ferdous Raiyan, Mohammed Nasir Uddin Akanda, and Nahed Hassan Jony. A comparative flow analysis of naca 6409 and naca 4412 aerofoil. *International Journal of Research in Engineering and Technology*, 03(10):342?350, October 2014.
- [92] Temesgen Batu, Hirpa G. Lemu, Besufekad Negash, Eaba Beyene, Dagim Tirfe, Eyob Hailemichael, and Solomon Alemneh. Optimal airfoil selection for small horizontal axis wind turbine blades: A multi-criteria approach. *Advances in Mechanical and Materials Engineering*, 41:57?68, 2024.
- [93] Mark Jason Thomas Loutun, Djamal Hissein Didane, Mohd Faizal Mohideen Batcha, Kamil Abdullah, Mas Fawzi Mohd Ali, Akmal Nizam Mohammed, and Lukmon Owolabi Afolabi. 2d cfd simulation study on the performance of various naca airfoils. *CFD Letters*, 13(4):38?50, April 2021.
- [94] Connor Rees and Berndt Muller. All that glitters is not gold: trustworthy and ethical ai principles. *AI and Ethics*, 3(4):1241–1254, November 2022.

- [95] Anna Rose Lucy Carter, Gavin Bailey, Jennifer Pearson, Matt Jones, Simon Robinson, Dani Kalarikalayil Raju, Spencer Winter, and Jonathan Lloyd Hicks. Designing and embedding a tangible public interface in the covid era. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, CHI '22. ACM, April 2022.
- [96] Daniel A Magües, John W Castro, and Silvia T Acuna. Hci usability techniques in agile development. In *2016 IEEE International Conference on Automatica (ICA-ACCA)*, pages 1–7. IEEE, 2016.
- [97] Qing Wang, Haiwei Luo, Jian Xiong, Yanjie Song, and Zhongshan Zhang. Evolutionary algorithm for aerospace shell product digital production line scheduling problem. *Symmetry*, 11(7):849, 2019.
- [98] Andrew Walenstein. Finding boundary objects in se and hci: An approach through engineering-oriented design theories. In *ICSE Workshop on SE-HCI*, pages 92–99. Citeseer, 2003.
- [99] Oliver Bown and Andrew R Brown. Interaction design for metacreative systems. In *New Directions in Third Wave Human-Computer Interaction: Volume 1-Technologies*, pages 67–87. Springer, 2018.
- [100] Isabel PS Qamar, Rainer Groh, David Holman, and Anne Roudaut. Hci meets material science: A literature review of morphing materials for the design of shape-changing interfaces. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–23, 2018.
- [101] Eric Horvitz. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 159–166, 1999.

- [102] Caroline PC Chanel, Raphaëlle N Roy, Nicolas Drougard, and Frédéric Dehais. Mixed-initiative human-automated agents teaming: Towards a flexible cooperation framework. In *International Conference on Human-Computer Interaction*, pages 117–133. Springer, 2020.
- [103] Ting-Ju Chen and Vinayak R Krishnamurthy. Investigating a mixed-initiative workflow for digital mind-mapping. *Journal of Mechanical Design*, pages 1–40, 2020.
- [104] Jichen Zhu, Antonios Liapis, Sebastian Risi, Rafael Bidarra, and G Michael Youngblood. Explainable ai for designers: A human-centered perspective on mixed-initiative co-creation. In *2018 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–8. IEEE, 2018.
- [105] An T Nguyen, Aditya Kharosekar, Saumyaa Krishnan, Siddhesh Krishnan, Elizabeth Tate, Byron C Wallace, and Matthew Lease. Believe it or not: Designing a human-ai partnership for mixed-initiative fact-checking. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, pages 189–199, 2018.
- [106] Hamed Saeidi, Foster McLane, Bahzad Sadrfaidpour, Evan Sand, Sheng Fu, Julio Rodriguez, John R Wagner, and Yue Wang. Trust-based mixed-initiative teleoperation of mobile robots. In *2016 American Control Conference (ACC)*, pages 6177–6182. IEEE, 2016.
- [107] Helmut Degen, Christof Budnik, Gregory Conte, Andrew Lintereur, and Seth Weber. How to explain it to energy engineers? In *HCI International 2022 – Late Breaking Papers: Interacting with eXtended Reality and Artificial Intelligence*, pages 262–284. Springer Nature Switzerland, 2022.
- [108] Antonios Liapis, Georgios N Yannakakis, Constantine Alexopoulos, and Phil Lopes.

- Can computers foster human users' creativity? theory and praxis of mixed-initiative co-creativity. 2016.
- [109] Amanda Swearngin, Andrew J Ko, and James Fogarty. Scout: Mixed-initiative exploration of design variations through high-level design constraints. In *The 31st Annual ACM Symposium on User Interface Software and Technology Adjunct Proceedings*, pages 134–136, 2018.
- [110] Jean Michel A Sarr, Georgios N Yannakakis, Antonios Liapis, Alassane Bah, and Christophe Cambier. Djehuty: A mixed-initiative handwriting game for preschoolers. In *International Conference on the Foundations of Digital Games*, pages 1–4, 2020.
- [111] Ben Samuel, Michael Mateas, and Noah Wardrip-Fruin. The design of writing buddy: a mixed-initiative approach towards computational story collaboration. In *International Conference on Interactive Digital Storytelling*, pages 388–396. Springer, 2016.
- [112] Alexander Baldwin, Steve Dahlskog, Jose M Font, and Johan Holmberg. Mixed-initiative procedural generation of dungeons using game design patterns. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 25–32. IEEE, 2017.
- [113] Alberto Alvarez, Steve Dahlskog, Jose Font, and Julian Togelius. Empowering quality diversity in dungeon design with interactive constrained MAP-elites. In *2019 IEEE Conference on Games (CoG)*. IEEE, aug 2019.
- [114] Martin Balla, Adrián Barahona-Rios, Adam Katona, Nuria Pena Pérez, and Ryan Spick. Illuminating game space using map-elites for assisting video game design. In *11th AISB Symposium on AI & Games (AI&G)*, pages 1–6, 2021.

- [115] Megan Charity, Isha Dave, Ahmed Khalifa, and Julian Togelius. Baba is y'all 2.0: Design and investigation of a collaborative mixed-initiative system. *IEEE Transactions on Games*, 2022.
- [116] Bart P. Knijnenburg, Martijn C. Willemsen, Zeno Gantner, Hakan Soncu, and Chris Newell. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*, 22(4-5):441–504, March 2012.
- [117] Philipp Mayring. *Qualitative Content Analysis: Theoretical Background and Procedures*, pages 365–380. Springer Netherlands, September 2014.
- [118] Jakub Vincalek, Sean Walton, and Ben Evans. It's the journey not the destination: Building genetic algorithms practitioners can trust. In *14th International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control*, 2021.
- [119] Jakub Vincalek, Sean Walton, and Ben Evans. On the use of CFD for the comparison of a novel wing tip design to a standard wing tip in the field of wind turbine aerodynamics. In *UK Association for Computational Mechanics School and Annual Conference*, 2022.
- [120] Sean P. Walton, Ben Evans, Ben Smith, and Jakub Vincalek. How engineers use evolution to invent things. *Frontiers for Young Minds*, 10, August 2022.
- [121] Jakub Vincalek, Sean Walton, and Ben Evans. Optimising a novel winglet for use on a wind turbine blade. In *AI-OPT*. The University of Melbourne and Monash University, ARC Training Centre in Technologies, Integrated Methodologies and Applications (OPTIMA), 2022.
- [122] Jakub Vincalek, Sean Walton, and Ben Evans. Understanding the effects of a ducted

winglet on wake and induced drag using CFD. In *22nd Computational Fluids Conference, 2023*.

- [123] Jakub Vincalek, Sean Walton, and Ben Evans. Optimisation of a novel ducted winglet on a wind turbine blade for reduced wake. In *15th International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control, 2023*.

Appendix A

Appendix




Figure A.1: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 1. Through the first 11 generations, time was spent in the *Editor* view with no change in the improvement until the 12th generation. From the 12th generation onwards, time was spent in all available views. Notably, a significant amount of time was spent in *Insights 1* which showed the random designs. Steady increases to the fitness are made approximately every 5 generations with 2 consecutive improvements made in the last two generations.




Figure A.2: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 2. The majority of the time was spent in the *Editor* view from the start with several but small improvements made throughout. No improvements were made in the last 13 generations.




Figure A.3: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 3. Most of the time is spent in the *Editor* view and the *Live* view, with few but large improvements made. A limited amount of time was spent in both the *Insight* views.




Figure A.4: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 4. The *Editor* view was the most popular until the 16th generation, when an increase is observed in the other three views. Improvement is stagnant until halfway through the run, from which point small and steady increases are seen.




Figure A.5: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 5. Time is spent exclusively in the *Editor* view for the first 9 generations. This coincides with steady and large improvements in the maximum fitness. From the 10th generation to the 20th generation, time is split between the *Live* and MAP-elites view. The participant then returns to the *Editor* view for the majority of the remaining time.




Figure A.6: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 6. Exclusive use of the editor coincides with increases in the fitness in the first 15 generations. Use of the *Insights* views is seen from the 11th generation onwards.




Figure A.7: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 7. The majority of the time spent through the run is in the editor. There are few but large increases in the fitness, with 3 instances of very small improvements as well.




Figure A.8: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 8. Of all the participants, this is the only case in which the *Editor* view is the least popular of the views and is not used in the last 30 generations. Most of the time is spent in the *Live* view. This participant spends approximately a quarter of their time in the random designs view, and steadily increases time spent in the MAP-elites view from generation 15 onwards. Small and steady improvements are made throughout the run, with only one large increase at the beginning.




Figure A.9: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 9. This participant was the only one to have both *Insights* views as more popular than the *Editor* and *Live* views. An initial increase in the fitness in generation 12 was followed by exclusive use of the editor in the following generation. This was followed on by additional large increases in the fitness over the next 7 generations.




Figure A.10: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 10. Time was spent exclusively in the *Live* view for 5 generations, at which point the *Editor* view was used. Only one improvement was made throughout the entire run.




Figure A.11: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 11. Most of the time spent in the first half of the run was in the *Editor*, with no improvements to the fitness until the second half. This also coincides with increasing use of both the MAP-elite view and the random design view.




Figure A.12: A breakdown of time spent as a percentage of overall time in each view compared to the fitness improvement for participant 12. Time spent in the first half of the run alternated between the *Live* and *Editor* views. Steady improvements were made throughout. Use of the *Insights* views started in the second half of the run, though no further improvements were made.




Figure A.13: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 1. This participant only selected editor keeps over the course of their run. The increase in editor keeps coincides with improvements to the maximum fitness.




Figure A.14: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 2. This participant made use of both the keeps and breeds in the MAP-elites view, though did not use the same features in the random designs view. This participant almost exclusively selected both options for each design selected.




Figure A.15: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 3. As with the previous two participants, Participant 3 had persistent editor keeps through most of the run. From the halfway point onwards, they also started using MAP-elite breeds. Random design breeds were also selected from generation 33 to 38.




Figure A.16: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 4. Participant 4 used editor keeps until the 18th generation, at which point they switched over to MAP-elite breeds and 2 random design breeds. The switch coincides with an improvement in fitness. The participant then switches back to editor keeps for the remainder of the run with several small improvements in fitness as well.




Figure A.17: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 5. Editor keeps increased throughout the entire simulation with both MAP-elite breeds and keeps increasing from generation 11 when improvement stalled. Random design breeds were also introduced in the last 5 generations.




Figure A.18: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 6. Only editor keeps were made by this participant with improvements made as the keeps began.




Figure A.19: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 7. As with the previous participant, only editor keeps were made. Improvements to the fitness were made throughout though sparsely.




Figure A.20: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 8. A similar amount of keeps were recorded between the editor, MAP-elites and random designs with small improvements throughout the entire run. An improvement is also seen before any editor keeps indicating that the GA managed to improve the design prior to specific design changes by the participant.




Figure A.21: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 9. The most amount of MAP-elite breeds occurred during this simulation with 94 in total. Random design breeds were also used until the 34th generation.




Figure A.22: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 10. This participant relied mainly on editor keeps, 2 MAP-elite breeds at the halfway point in the run and a random design breed in the last generation.




Figure A.23: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 11. While editor keeps persisted through the entire run, fitness began improving shortly after and while MAP-elite breeds were being used. This participant had the second most amount of MAP-elite breeds by the end at 74.




Figure A.24: A breakdown of the number of keeps and breeds in the Editor and each of the *Insights* view, mapped to either the MAP-Elites or random designs, then compared to the fitness for participant 12. This participant used editor keeps almost exclusively, with the exception of a single breed of a MAP-elite in generation 27.