University of Bath

UNIVERSITY OF
BATH

PHD

Generative Part Design for Additive Manufacturing

Goguelin, Steven

Award date:
2019

Awarding institution:
University of Bath

[Link to publication](#)

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

Generative Part Design for Additive Manufacturing

Steven Goguelin

A thesis submitted for the degree of Doctor of Philosophy

University of Bath

Department of Mechanical Engineering

November 2018

Signature of Author . . . .

Steven Goguelin

To mum and dad.

Acknowledgements

Without the monumental support of friends and family, this research would simply not have been possible. I would take this opportunity to thank you for all the help that you have provided along the way.

I would like to extend my thanks to my primary supervisor, Dr. Vimal Dhokia. From the moment that I approached you to enquire about embarking on a PhD, you have provided me with endless support and time. You gave me the freedom to explore new ideas, but always kept me on track when I travelled too far away from the research path. Without your patience and advice this thesis would never have been finished. Thank You.

Secondly, I would like to thank my second supervisor, Dr. Joseph Flynn. When I began this journey, you were busy finishing your PhD, but you still spent the time to answer my questions, discuss my research and help me integrate with the research group. Thank you for all the late nights reading, critiquing and improving my work. Finally, thanks for being a great friend outside of work and instilling in me the importance of having a work-life balance.

To Prof Stephen T Newman, thank you for creating a research group with the atmosphere that promotes inclusivity, discussion and friendship. Thank you to the researchers within AMPS, I have met some amazing friends whilst doing this research and I truly appreciate all the discussions, both productive and silly, that we have had during lunches, team meetings and Friday beers™.

To Harry Leafe, thanks for being the best distraction during my studies. Thanks for supporting me when research wasn't going well and for being there for the fun times over the last years. Finally, to my parents Graham and Judith Goguelin. You have taught me the value education over the last 21 years. Without your support, encouragement and unconditional love, this would never have been achievable. Thank You.

Abstract

Additive manufacturing has transformed from a technology primarily focused on the creation of small-scale prototypes and models to a process for the manufacture of end-use components. Additive manufacturing has a number of inherent advantages over traditional manufacturing processes. These include the ability to fabricate complex geometries at no extra production cost, creating mass customised parts without additional tooling and the ability to consolidate assemblies into single parts. However, additive manufacturing is not without limitations. There are a number of geometric constraints that limit design freedom when designing parts, for example the maximum unsupported angle that material can be printed. There are also limitations in current CAD software, which prevent designers from maximising the quality of additively manufactured parts.

Design for additive manufacturing seeks to improve the design methods or tools to help improve the functional performance, reliability, manufacturability or cost of parts produced using additive manufacturing technologies. Generative design is an emerging form of computational design in which the user provides goals and constraints to a system and generative synthesis algorithms produce a series of optimised solutions based on the input criteria. There are many limitations with current generative design systems preventing the mass adoption of the technology. These include, the lack of integration between topology optimisation synthesis algorithms and the part build orientation and, additionally, the ability to design for goals, such as part cost or build time.

To overcome these challenges, this research applies two generative design methods to design an additively-manufactured cantilever beam. The optimised beams are created by integrating a ground-structure topology optimisation with manufacturing constraints and build orientation angle information. Design performance metrics of varying degrees of abstraction are then derived from the mesh file data. These represent two common additive manufacturing business scenarios; maximum part performance and high production quantity. A data-driven generative design approach is then used to locate the top performing solutions within a solution space. This space is searched using a parametric grid-search that alters the build orientation and the overhang angle constraint. Component performance is related to the abstracted design objectives using a TOPSIS multi-criteria decision analysis.

There is an ongoing challenge associated with running many model evaluations on large mesh files. This can make generative design prohibitive in terms of computational resource. To overcome this challenge, a goal-driven generative design method is developed to solve the inverse problem of finding the optimal input parameters to the cantilever beam problem by using a Bayesian optimisation surrogate model. The data-driven and goal-driven generative design approaches are then compared for their efficiency and ability to locate optimal design solutions.

The contributions to knowledge, borne from this research are a data-driven generative design method demonstrated to be suitable for locating high-performing solutions in complex multi-dimensional solution spaces providing the number of design space dimensions is small. The novel use of Bayesian optimisation is shown to be to be 17 times more efficient than a conventional grid search for locating the top performing build orientations of two additively manufactured test parts. Finally, goal-driven generative design methods are demonstrated to locate the optimised build orientation and manufacturing constraints for the cantilever beam within 20 optimisation iterations. These outcomes demonstrate potential for future generative design CAD systems.

Contents

List of Figures

List of Tables

Glossary

3MF	3D Manufacturing Format
AM	Additive Manufacturing
AMF	Additive Manufacturing File
ASTM	American Society for Testing and Materials
AHP	Analytical Hierarchy Programming
BESO	Bidirectional Evolutionary Structural Optimisation
CSS	Cascading Style Sheets
CSV	Comma Separated Variable
CDS	Computational Design Synthesis
CAD	Computer Aided Design
CNC	Computer Numerically Controlled
DfAM	Design for Additive Manufacturing
DfMA	Design for Manufacturing and Assembly
DSV	Design Solution Volume
DED	Direct Energy Deposition
DMLS	Direct Metal Laser Sintering
EBM	Electron Beam Melting
ESO	Evolutionary Structural Optimisation
EI	Expected Improvement
XML	Extensible Markup Language
FDM	Fused Deposition Modelling
GD&T	Geometric Dimensioning and Tolerancing
GPU	Graphics Processing Unit
HPC	High Performance Computing
HTML	HyperText Markup Language
IQR	Interquartile Range
MVS	Material Void Space
MMA	Method of Moving Asymptotes
MCDA	Multi-Criteria Decision Analysis
PLA	Polylactic Acid
PBF	Powder Bed Fusion
SIMP	Solid Isotropic Material with Penalisation
SLA	Stereolithography
STL	Stereolithography File
TOPSIS	Technique for Order Preference by Similarity to Ideal Solution
UCB	Upper Confidence Bounds
WPM	Weighted Product Model
WSM	Weighted Sum Model
WAAM	Wire + Arc Additive Manufacturing

Chapter 1 - Introduction

1.1. Research Background

The American Society for Testing and Materials (ASTM) define additive manufacturing (AM) as "a process of joining materials to make objects from 3D model data, usually layer upon layer, as opposed to subtractive manufacturing methodologies" (ASTM 2012). Traditional subtractive technologies, as defined by Nassehi et al. (2012), are processes where material is removed from a single workpiece, such that the mass of the processed workpiece is less than the original. Examples include milling, turning, electro-discharge machining and water jet cutting.

The recent proliferation of AM within industry is notable. The Wohlers Associates report (Wohlers et al. 2018) highlights a 21% growth on sales revenue from AM specific parts, machines, services and materials during 2016-2017 with a total industry value of $7.3 billion. The report also details a considerable 80% increase in the sales of metal AM machines, highlighting that industry is now adopting AM as an accepted means of production for high-value end-use components. Examples of AM being used in commercial production can be found in the medical (Wang et al. 2016), footwear/fashion (Griffiths 2017), aerospace (Kellner 2015) and automotive industries (BMW Group 2018).

AM processes can be categorised by the type of material processed, the deposition technique or the way the material is fused, as shown in Table 1-1. The AM processes are placed into seven categories, namely: material extrusion, powder bed fusion, vat photopolymerization, material jetting, binder jetting, sheet lamination and directed energy deposition. Since the establishment of these categories in 2012, there has been a vast amount of development increasing the performance of AM hardware. Recent trends have focused on advancing the speed and quality of part production to ready the process for commercial implementation. Consequently, AM is now competing with casting in terms of production cost for part quantities between 50,000-100,000 parts (HP 2018; Desktop Metal 2018). To achieve this, companies have either developed new processes based on existing technology from the metal injection moulding industry (Markforged 2018), integrated parallelised processing techniques such as adding multiple lasers to machines (Renishaw plc. 2018), or incorporated the part build process, heat treatment and post-processing into the same machine (Additive Industries 2017).

Table 1-1: ASTM International classification of additive manufacturing processes (ASTM 2012).

CATEGORIES	TECHNOLOGIES	PRINTED "INK"	POWER SOURCE	STRENGTHS / DOWNSIDES
Material Extrusion	Fused Deposition Modeling (FDM)	Thermoplastics, Ceramic slurries, Metal pastes	Thermal Energy	• Inexpensive extrusion machine • Multi-material printing • Limited part resolution • Poor surface finish
	Contour Crafting			
Powder Bed Fusion	Selective Laser Sintering (SLS)	Polyamides /Polymer	High-powered Laser Beam	• High Accuracy and Details • Fully dense parts • High specific strength & stiffness • Powder handling & recycling • Support and anchor structure
	Direct Metal Laser Sintering (DMLS)	Atomized metal powder (17-4 PH stainless steel, cobalt chromium, titanium Ti6Al-4V), ceramic powder		
	Selective Laser Melting (SLM)			
	Electron Beam Melting (EBM)		Electron Beam	• Fully dense parts • High specific strength and stiffness
Vat Photopolymerization	Stereolithography (SLA)	Photopolymer, Ceramics (alumina, zirconia, PZT)	Ultraviolet Laser	• High building speed • Good part resolution • Overcuring, scanned line shape • High cost for supplies and materials
Material Jetting	Polyjet / Inkjet Printing	Photopolymer, Wax	Thermal Energy / Photocuring	• Multi-material printing • High surface finish • Low-strength material
Binder Jetting	Indirect Inkjet Printing (Binder 3DP)	Polymer Powder (Plaster, Resin), Ceramic powder, Metal powder	Thermal Energy	• Full-color objects printing • Require infiltration during post-processing • Wide material selection • High porosites on finished parts
Sheet Lamination	Laminated Object Manufacturing (LOM)	Plastic Film, Metallic Sheet, Ceramic Tape	Laser Beam	• High surface finish • Low material, machine, process cost • Decubing issues
Directed Energy Deposition	Laser Engineered Net Shaping (LENS) Electronic Beam Welding (EBW)	Molten metal powder	Laser Beam	• Repair of damaged / worn parts • Functionally graded material printing • Require post-processing machine

During the past 30 years, AM technology has transformed from a technology primarily focused on the creation of prototypes and models, to instead being used for the production of end-use functional parts (Campbell et al. 2012). Due to inherent limitations of layered manufacturing processes, not all parts are feasible or cost-effective when produced using AM. To ensure successful integration of AM into standard industrial practice, a better understanding of the design freedoms and limitations of the process is required.

The layer-wise material deposition of AM printing parts provides several inherent benefits over other traditional manufacturing methods. These include:

a) The ability to create complex part geometries that could not be produced using conventional manufacturing methods. Gibson et al. (Gibson et al. 2014) define four design complexities that can be exploited by AM. Firstly, shape complexity gives designers greater flexibility in determining the final geometry of a part. As tool-access is no longer an issue, designers have greater freedom in material placement as draft angles and undercuts no longer present difficulties. Additionally, adding extra features to AM parts has less impact on the final cost of the product compared to traditional manufacturing methods such as computer numerically controlled (CNC) machining. Secondly, hierarchical

complexity allows for multi-scale design freedom enabling the possibility of shape optimisation techniques to allocate material distribution throughout a part according to loading conditions, and lattice structures to be placed within designs. Thirdly, material complexity offers the freedom to design continuous material transitions and different materials within a single part. Finally, functional complexity allows for part consolidation or the embedding of foreign products within the final part.

b) An advantage of AM when competing with batch production casting methods is that there is no startup tooling required for production (Hopkinson et al. 2006). It is possible to create mass-customised parts across a batch as the cost of the final part is independent of tooling costs. Additionally, this lack of tooling means that AM can be used to produce on-demand, and on-site products, which reduces costs associated with inventory, supply chain and delivery.

c) Producing less waste material when compared with processes such as CNC milling, as there is no longer a requirement to machine from a large billet of material. Near net shaped fabrication processes such as directed energy deposition methods, can produce parts that are close to the final geometry. A subtractive process is then used to finish the part to the required tolerance (Ding et al. 2014). It is also possible to repair damaged parts using a combination of reverse engineering, AM and subtractive manufacturing (Teibrich et al. 2015).

d) Reducing assembly costs by part consolidation. Part count reduction has been shown to be one of the most effective ways to reduce process time and cost by reducing the number of assembly operations (Yang et al. 2015). Numerous case studies are available showing the advantages AM can have when consolidating assemblies both industrially (Autodesk 2018b) and in academia (Yang et al. 2015; Prakash et al. 2014; Sossou et al. 2018).

This list represents some of the advantages that AM provides. However, due to the nature of the AM process, there are many reasons why there may be a discrepancy between the computer-aided design (CAD) model and the final printed product. Errors can arise as a result of the discretisation of the part geometry to triangulated mesh representation, such as in a .STL file (Figure 1-1, left). Another error may be ascribed to the layered manufacturing process itself. All layered manufacturing processes require the digital model to be divided into slices before the part can be manufactured. These slices then form the basis of a material deposition plan for the part. Slices can contribute to several errors that occur when comparing the original CAD model to

the printed object. One example, termed the stair stepping effect, occurs when the discretised contours of the 2.5D layers are printed, as shown in Figure 1-1 (right). This phenomenon can significantly reduce the surface quality of the design. Furthermore, layered manufacturing can also lead to anisotropy in the mechanical properties of the part due to possible delamination of layers during loading (Stanković et al. 2017).

Figure 1-1: (left) Error associated with triangulating a sphere during the creation of STL files and (right) error associated with the stair-stepping effect attributed to the layerwise manufacturing process (images courtesy of 3D Hubs (2017)).

Alternatively, failures may occur during the build itself. The physical phase change from liquid to solid during the processing can lead to many potential failure modes. Warping and cracking can occur when there is a build-up of residual stress caused by temperature differences between prebuilt and newly heated layers. Different AM processes are more susceptible to certain failure modes. For example, directed energy deposition (DED) processes are particularly prevalent to residual stress failures, whereas powder bed fusion (PBF) processes are more susceptible to failures associated with incorrect powder recoating or impact from the recoater blade.

Print failures may dramatically increase the cost of the AM process and therefore a comprehensive understanding of the process limitations is extremely important when mitigating against print failures. Modifying the print processing parameters can dramatically change the quality of the output part. However, there is a vast amount of parameters (>100 options in Cura, a common material extrusion build preparation software) that can be modified and understanding the effect of changing a parameter, or combination of parameters, on the final part is hugely challenging.

Correct consideration of AM design heuristics can mitigate against many of the disadvantages associated with AM. To leverage the advantages of AM, it is vital to design specifically for the technology. Design for manufacturing and assembly (DfMA) has been pivotal in improving the quality of designs as the methodology has educated many designers about the importance of understanding manufacturing constraints within the design process. Boothroyd and Dewhurst (Boothroyd et al. 2011) created the seminal book on the topic, covering many traditional manufacturing methods. However, there is a need to extend DfMA to include AM processes.

Design for Additive Manufacturing (DfAM) has emerged as a critical field of research investigating design tools, guidelines and constraints on the AM process that require consideration to maximise the likelihood of achieving additively manufactured, functional, first time right parts (Thompson et al. 2016).

DfAM has been defined as "types of design methods or tools whereby functional performance and/or other key product life-cycle considerations such as, manufacturability, reliability, and cost can be optimised subject to the capability of AM technologies (Tang and Zhao 2016)". General guidelines advising on best practices for AM part design have continually been refined over the last fifteen years (Hague et al. 2003). Exploiting these guidelines is pivotal to the economic viability of the process (Atzeni and Salmi 2012) as correct implementations of DfAM can improve the functional performance of parts without increasing cost.

One of the fundamental challenges in DfAM is managing the complexity of the design opportunities and limitations. The role of the designer is to translate the part design specification into a feasible digital CAD model that can then be processed using AM. In non-trivial design problems, there are many solutions that may lead to a satisfactory result. Therefore challenges arise in optimising designs for specific attributes whilst reducing the time and cost required to develop feasible solutions. Exploiting the interaction between AM processes, machines, materials and process parameters, as well as, the many design options that arise due to the flexibility of the process is a demanding task due to the combinatorial explosion of potential options available to the designer.

CAD tools traditionally provide designers with a facility to document, simulate and visualise designs. However, synthesising geometry based on underlying design requirements remains challenging due to the multi-objective nature of most non-trivial design problems. Advances in computer hardware have meant that design simulation and computational synthesis methods have become widely accessible. The combined advantages of modern CAD systems and AM technologies has triggered a widespread rethink about the role of designers and the software tools that will support them in this increasingly complex setting.

Generative design provides a different approach to the traditional CAD process. It aims to extend CAD systems by automatically synthesising geometries based on a series of goals and constraints set by the designer. Typically, generative design systems create multiple solution variants to a given input problem. The user can then explore the generated alternatives and select the best options for further refinement. Throughout this thesis, designer and user will be used interchangeably to refer to the person using the generative design tool or method.

This research aims to overcome the challenges associated with designing parts for AM by guiding the development of future generative design CAD tools that assist designers in producing high quality AM parts based on a series of pre-defined goals and constraints. To achieve this, a manufacturability-constrained structural optimisation method will be combined with build orientation angle information to generate an AM aware constrained optimisation. This, in turn, will inform the development of two design exploration approaches, namely, data-driven and goal-driven generative design for AM. The data-driven method uses multi-criteria decision analysis (MCDA) to locate the highest performing solutions within a pre-generated solution space, derived from a parametric grid search. Goal-driven approaches optimise the input parameters to a parametric optimisation to synthesise high-performing solutions using a Bayesian optimisation surrogate model. A design solution's performance will be defined by a multi-objective weighting of high-level abstractions of low-level multi-dimensional evaluation criteria that will be outlined in the research methodology.

1.2. Thesis Outline

Chapter 1 introduces the topic of generative design and DfAM, as well as explaining the context of the research and its significance. Finally, a short introduction to the research aim is provided.

Chapter 2 reviews the literature associated with designing parts with AM. It will explore recent work governing the design of functional end-use parts and introduce the notion of the DfAM pipeline. Secondly, it examines the extent to which current CAD tools are capable of aiding designers wishing to exploit the full potential of AM. Finally, the literature review analyses state-of-the-art research in optimising geometry to fulfil user-specified functions, before a critique is presented that informs the research aims and objectives that govern this thesis.

Chapter 3 provides a generalised framework for CAD tools supporting generative design for AM. The principles for the future development of generative design systems are described, advising future researchers and generative design practitioners of their critical features.

Chapter 4 explains the research hypothesis and aim, alongside the measurable objectives and scope of the research.

Chapter 5 presents the overarching research methodology and the key methods used to achieve the research objectives.

Chapter 6 uses data-driven design methods for generative design to explore the trade-offs that occur in the AM design process. Data-driven methods search a pre-computed solution space of design solutions to find the best solution according to a series of design objectives. The solution

space is generated using a grid search approach, and then is subsequently explored a posteriori using a MCDA.

Chapter 7 evaluates the strategies for exploring the AM design space when the objective function is costly to evaluate. Through the comparison of exploration methods, namely grid search, random search and Bayesian optimisation using Gaussian processes, the optimal orientation of two AM-specific test parts will be located.

Chapter 8 examines the efficacy of goal-driven generative design methods. These methods use techniques including random search and Bayesian optimisation to efficiently explore the AM solution space by locating high-performing solutions evaluated against high-level abstracted design goals. The computational efficiency of determining these solutions is then examined based on the number of structural optimisation evaluations required. The results are then quantitatively compared to the data-driven approach and used to critique the two generative design implementations objectively.

A critical discussion of the research is given in **Chapter 9**. Finally, **Chapter 10** concludes the research, provides the contributions to knowledge and outlines future research avenues.

Chapter 2 – Literature Review

2.1. Introduction

This chapter introduces the literature that has influenced the field of DfAM. Firstly, the limitations of AM processes outlined in design guidelines are introduced. Secondly, research outlining past work in the DfAM pipeline will be reviewed. Thirdly, generative synthesis techniques that have been developed to optimise AM parts to achieve specific functionality will be reviewed. Next, an overview of the issues surrounding traditional CAD programs in designing parts for the additive manufacturing process is provided. Finally, a critique of the literature and the influence of the research gaps on the development of work in this thesis will be outlined. For clarity, a diagrammatic representation of the review is shown in Figure 2-1.

Figure 2-1: Diagrammatic representation of literature review topics.

2.2. Additive Manufacturing Guidelines

Before designing parts for AM, it is first necessary to understand any limitations that will constrain the design process. Geometric limitations in AM are affected by the process, material and the machine selected. It is essential that the designer understands these constraints in order to generate an optimal design for a particular AM combination.

In order to qualify the capabilities of AM machines, researchers have developed test parts in order to assess the geometric limitations of the machine-process-material combination. Rebaioli and Fassi (2017) provided a comprehensive review of seventy-nine geometric test parts that aim to provide a comparison between AM technologies and to quantify the limitations of AM processes. These test parts can be used to measure the geometrical accuracy, repeatability and minimum feature sizes for each of the AM machines. It is highly important to understand the

geometric capabilities of the process/material combination as this will allow the designer to decide which AM process is most suitable for a given part based on its geometric features.

By utilising these test parts to find geometric limitations, it is possible to generate design guidelines for each of the AM processes. Guidelines have been produced by many industrial AM vendors (Materialise NV 2018; Formlabs 2015) and also from academic research covering a wide range of processes including: direct metal laser sintering (DMLS) (3D Hubs 2018; Kranz et al. 2015; Thomas 2009) wire and arc additive manufacturing (WAAM) (Lockett et al. 2017; Mehnen et al. 2014), stereolithography (SLA) (Formlabs 2015), fused deposition modelling (FDM) (Teitelbaum et al. 2009) and electron beam melting (EBM) (Vayre et al. 2013).

Accurate implementation of design guidelines is fundamental to reducing print failures in AM. Failures can have a significant impact on the overall cost of the AM process, with companies having to incorporate the probability of print failure as a multiplier on the final cost of the individual part to cover overhead costs (Baumers et al. 2016). However, it is often challenging to incorporate the resulting design guidelines easily in the part development process as the discrete geometric shape guidelines are difficult to translate into the complex freeform geometries that exploit the benefits of AM. Furthermore, it is challenging to understand which of the design guidelines has the most significant impact on part performance. This is important if geometric constraints prevent the unilateral implementation of all guidelines.

2.3. The Design for Additive Manufacturing Pipeline

Having accrued an understanding of AM process constraints, a designer can embark on the process of creating a functional AM part. Most AM parts start as a 3D digital computer-aided design (CAD) file. This CAD file then follows an established AM process pipeline, including both digital and physical elements (Figure 2-2) until the manufactured part is completed. Whilst the designer no longer has a direct influence on the part after the print process begins, it is necessary to ensure post-processing in taken into account within the design stage. This ensures parts can be finished appropriately, ensuring the part design specification criteria are met. This section will discuss the role of DfAM within the AM design pipeline.
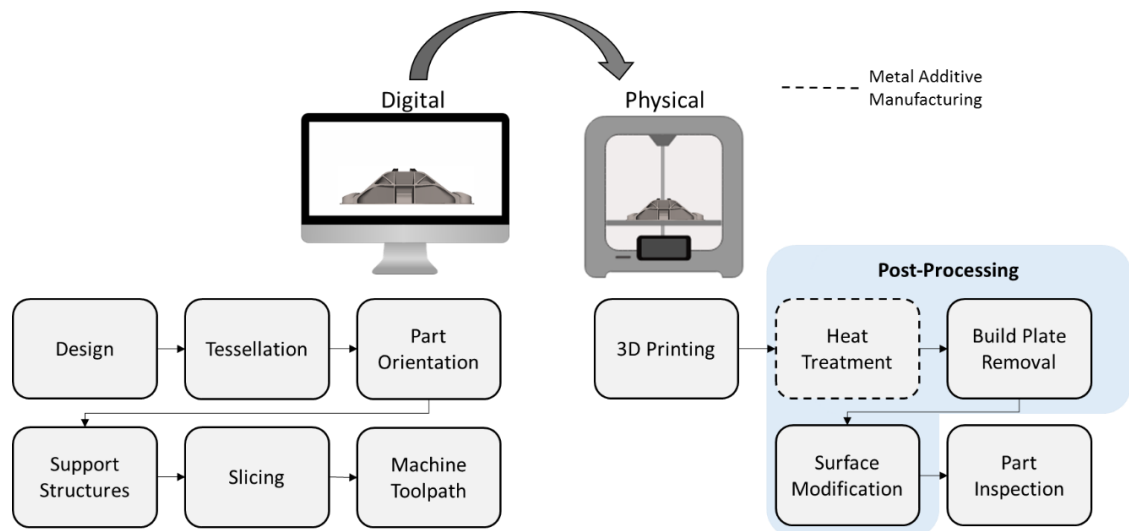
Figure 2-2: Design for additive manufacturing pipeline.

2.3.1. Design

A fundamental concept in DfAM is that of functional surfaces and volumes; these are geometrical elements that fulfil a particular functional requirement. For a structural part, a functional surface may restrict a degree of freedom, interface with another part in an assembly or act as a load bearing surface (Tang et al. 2015). The functional volume is the part volume required to connect each of the functional surfaces.

Vayre et al. (2012) proposed a design approach for AM. Initially, the part specification must be known in order to define the desired in-service behaviour. Secondly, the part is decomposed into a set of functional surfaces. A set of free parameters is then defined that can be optimised to find the best design for the part. Ponche et al. (2012) also used the notion of functional volumes to combine parts into a single assembly suitable for AM.

Salonitis (2016) applied the techniques of axiomatic design to DfAM. Axiomatic design involves mapping customer needs onto functions that the part is expected to perform. It then maps these functions onto design parameters indicating how the object will perform. Finally, these design parameters are mapped onto process variables indicating the AM method and process parameters that will be used to manufacture the part. This method has the advantage of ensuring that all functional requirements are captured by the geometry. However, there is no information about how to design the geometry to achieve the design goals.

The use of an AM feature database was explored by Bin Maidin et al. (2012). This database included 106 AM-specific design features that were identified to be uneconomical to produce using conventional manufacturing processes. The results showed that novice users of AM found

the feature database to be useful when creating novel design concepts. However, its usefulness to design professionals was less evident.

Booth et al. (2017) created a DfAM worksheet that aimed to help novice AM users improve their part designs. The sheet consists of various guidelines incorporated into eight categories. The user then assesses their design against the number of design guidelines that their part does not meet. The results show that a simple worksheet reduced the number of failed prints by 81%. However, the generalisation of the worksheet to all design processes and machines suggests that the design guidelines are quite conservative.

Kumke et al. (2016) developed a framework for DfAM derived from the VDI2221 design methodology (VDI Guideline 1993), which consists of a series of modules aimed to guide the user through the AM part development process. Whilst the framework is a useful tool for the designer, it is difficult to exploit the framework without direction integration with CAD tools.

2.3.2. Tessellation

Digital part representation is an often overlooked, highly important feature of DfAM. The .STL file format has become the de-facto standard for 3D-printing technologies. This format approximates the surfaces of the CAD model with triangles. With simple part geometries, the .STL file is typically exported in an error-free model that is suitable for AM. However, if the geometric complexity increases then occasionally the .STL file will require further processing (fixing) before the design can be printed (Oropallo and Piegl 2015).

STL files exhibit a number of potential issues including missing, degenerate and overlapping facets and non-manifold topology conditions (Leong et al. 1996). An essential requirement of DfAM tools, therefore, must be to repair the mesh before providing further insights into the overall manufacturability of the design. A further disadvantage of the STL file format is that no information is stored about the part features or print data. To overcome some of the disadvantages of the STL file, researchers have developed the AMF (additive manufacturing format) file, which extends the STL format to include colour, materials, textures and curved facets and also the 3MF (3D manufacturing format) file type which extends the XML (extensible markup language) file format for the description of 3D objects.

Another disadvantage of the .STL file format is that it is computationally inefficient to generate files with large numbers of facets, which is potentially challenging for complex part geometry. In order to combat this, researchers have developed alternative, meshless geometry representations to enable more efficient computation (Zhao et al. 2017; Pasko et al. 2011; Fryazinov et al. 2013).

However, challenges remain in slicing and generating toolpaths for direct AM of these representations.

2.3.3. Part Orientation

A well-chosen build orientation can reduce the time required to manufacture an AM part, reduce the total volume of the support structure required and alter the strength of the final part. However, as AM components become more complex, the ability to search all possible part orientations becomes computationally intractable. Therefore, there is a requirement to find more efficient methods for locating optimal part orientations. Researchers have seen the importance of optimising build orientation for AM components since the mid-1990s (Allen and Dutta 1994; Frank and Fadel 1995; Hur and Lee 1998; Lan et al. 1997; Cheng et al. 1995). After this initial surge in research activity, there was a quiet period until post-2010. At this point, there was a renewed interest in this challenge, quite possibly driven by the increased industrial interest in AM and, in particular, the proliferation of metal AM (Wohlers et al. 2018).

Morgan et al. (2016) used a gradient-based optimisation to find the optimum build orientation, where the objective function was the total support volume. Their results outperformed solutions given by commercial software; however, the authors highlighted the fact that an experienced engineer could improve upon the optimisation further. Furthermore, gradient-based approaches are subject to finding local minima and require a relatively large number of iterations. This effectively limits the practicality of this method in the presence of large mesh sizes.

Das et al. (2015) sought to exploit quadtree decomposition to improve the calculation efficiency within the optimisation scheme. Their solution optimised build orientation to minimise the volume of support structure, while also meeting specified geometric dimensioning and tolerancing (GD&T) requirements.

Al-Amari and Khan (2018) also optimised the part orientation with respect to a number of GD&T criteria, including heuristics such as maximising the number of holes that are perpendicular to the build direction. This was combined with a metric for build time comprising of the time taken to build the geometry. In this case, time was estimated by dividing the part volume by the build rate of the machine and the total support structure given by the commercial software, Magics. The orientation quality metric was only calculated for four orientations (+y,-y,+x,-x), limiting the ability to find the global optimal build orientation.

Brika et al. (2017) used a genetic algorithm to perform a multi-objective optimisation of the mechanical properties of the part, heat treatment, surface roughness, support length, build time and build cost. One of the disadvantages of using genetic algorithms is the requirement for a large

number of iterations, which can be a limiting factor if the the mesh size is large and support calculation is time-consuming.

Increasing the speed of determining optimal build orientation can be achieved by using specialist hardware such as graphics processing units (GPU's). Khardekar and McMains (2006) computed the volume of support material by using an OpenGL representation for projected pixel areas. Their results showed that using GPU's can offer speed increases of up to two orders of magnitude, which is advantageous within the setting of this optimisation problem.

A different focus was introduced in the work of Zhang et al. (2015), who used optimisation to minimise the aesthetic impact of removing support structure from visible surfaces on a component. Here, the authors trained a neural network using human preference data to determine the best orientations with respect to the volume of support in contact with visually important regions.

Zwier and Witts (2016) used the convex hull to combine the triangles from the STL file to identify suitable candidates for the lowest downward facing surface. Based on this, the amount of support structure per potential build surface was calculated. This method also had the advantage of placing the largest build surface directly onto the build plate, thereby increasing the stability of the build. Furthermore, it did not rely on testing multiple candidate build orientations. However, it must be noted that using convex hulls may not be appropriate for certain part geometries. Those geometries containing round or spherical surfaces will not be suitable, as the majority of the external surfaces will lie on the convex hull.

Surrogate optimisations are commonly used in engineering design when each objective function evaluation involves costly simulations. In these methods, the objective function is approximated with a proxy response surface. Ulu et al. (2015) used this approach to find the optimal build orientation to enhance the structural performance of parts. By using a cubic radial-basis function as a surrogate optimisation model, the authors were able to reduce the number of finite element simulations that were required to find the optimal structural geometry.

Previous research has used custom test parts to locate optimised build orientations, which may not be representative of the number of faces typically seen in engineering components produced using AM. It is therefore beneficial to use existing open-access parts that are representative of AM components to validate any algorithm. As different metrics for the total support requirement require different levels of computational complexity, the effect of simplifying these metrics on the final orientation results must be examined. Previous research has shown the hardware implementation such as GPUs can provide significant speed improvements when used effectively.

The attractiveness of hardware-based acceleration makes it essential to ensure that the methods developed within this research are compatible for extension with GPU accelerated computing.

2.3.4. Support Structures

Support structures are required in AM to ensure the successful build of parts with overhanging features. In AM, newly deposited material must be adequately supported by previously processed material. This means that some geometry features will require additional support material. These support structures are typically sacrificial and must be removed, post-build, incurring additional time and cost penalties. Furthermore, by definition, their presence consumes extra material, increases production time and decrease the surface quality of the final part. These drawbacks are particularly prevalent in metal AM, where it is estimated that between 40-70% of a component's cost may be attributed to the removal of support structure (Liu et al. 2018). An AM machine and process combination will have a critical overhang angle, θ_{crit} (Figure 2-3). Any face or facet whose normal vector exists outside the cone of critical overhang is likely to require additional support structure.
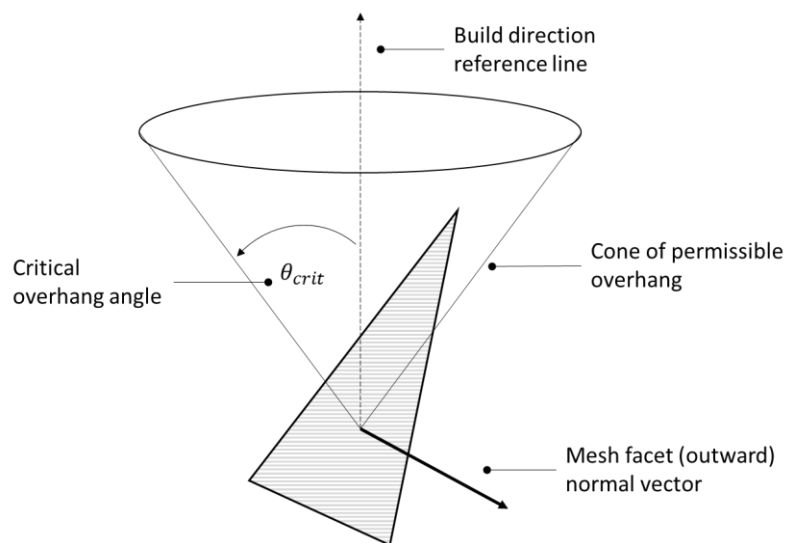
Figure 2-3: A single triangular facet, showing its outward facing normal vector. The normal vector exists beyond the limits of the critical overhang cone, declaring this facet as an overhang.

When printing metallic parts, support structures are also required to ensure that part distortion caused by residual stress build up during manufacture is limited. Support structure design is a complex process in CAD, and studies have shown that commercial software packages often overestimate the required amount of support for complex geometries (Jhabvala et al. 2012). There have been many attempts to optimise the support structure layout to minimise material usage. Vanek et al. (2014) utilised tree-like support structure, Dumas et al (2014) developed bridge-type supports, Huang et al. (2009) used a sloping wall structure and, most recently,

Mezzadri et al. (2018) and Kuo et al. (2018) used topology optimisation to minimise support volume. A full review on support structure design for AM can be found in a recent review (Jiang et al. 2018).

Research has also been targeted to improve the removal of support material. For polymer printing, it is possible to generate soluble supports by printing with a separate material (Priedeman Jr and Brosch 2004). For metallic materials, Hildreth et al. (2016) developed a method to dissolve carbon steel supports from a stainless steel DED print using electrochemical etching. Lefky et al. (2017) extended this research to powder bed fusion of stainless steels. Jhablava et al. (2012) used a pulsed laser when generating the support structures to aid with the ease of removal. Finally, Desktop Metal (Gibson et al. 2017) designed a ceramic interface layer between the support and final part that can be easily removed after sintering.

2.3.5. Additive Manufacturing Toolpath

The AM toolpath can have a significant impact on both the print times and the functional performance of the final additive part. There are advantages and disadvantages to different toolpaths, with the majority of research dedicated to optimising the tool path to reduce build time. A key issue with tool path generation in AM is the dissipation of heat. Tammas-Williams and Todd (2017) considered the future potential of AM with site-specific properties that can be generated by altering the melt strategy. Changing the laser scan strategy in powder bed processes can reduce residual stress accumulation and has been shown to make manufacturing high-temperature alloys feasible (Catchpole-Smith et al. 2017) and generate improved microstructural properties (Simonelli et al. 2014). Simulations have also been generated that can model the thermal effects of different tool paths to allow users to improve the structural properties of their parts (Parry et al. 2016) as well as to ensure that the build-up of residual stress does not lead to print failures.

The thermodynamic interactions that occur during the build can lead to residual stresses and warping within the final part. To overcome this, researchers have developed strategies to simulate and predict this deformation. The digital model can then be pre-deformed to ensure the final part meets geometric specifications (Chowdhury and Anand 2016; Schmutzler et al. 2016; Xu et al. 2017).

Ghouse et al. (2017; 2018) analysed the effect of changing laser parameters and scanning strategies on the mechanical performance of stochastic porous structures. The results demonstrated that optimisation of laser parameters can yield a 7% improvement of fatigue

strength while toolpath strategy can provide fatigue strength improvements of 8% and increase build speeds by up to 100%.

2.3.6. Post Processing

One of the primary disadvantages of AM, when compared to subtractive manufacturing, is poor quality surface finish and tolerances (Flynn et al. 2015). Due to the inherent characteristics of AM, there is often a requirement for part post-processing. This may include removal of the part from the build plate, removal of support structures, heat treatment and surface modification strategies for improved surface finish (Gordon et al. 2016). To ensure timely and cost efficient post-processing, engineers must have a full understanding of the finishing requirements for a part before it can be manufactured.

The finishing of AM parts by CNC machining is a popular method to achieve tighter dimensional and surface tolerances. Typically, additional material is added in carefully selected locations on the AM part to compensate for the material removed (Srinivasan et al. 2015). Fixturing for machining of AM parts can be challenging due to the complex part geometries that can be created using the AM process. Collision-detection algorithms can be used that determine tool access to faces that require machining (Inui et al. 2018). Frank et al. (2004) also addressed this challenge by adding sacrificial locating features in the part design. Alternatively, the Boolean difference between the AM component and a solid vice jaw can result in a conforming fixture design.

2.4. Generative Design in Additive Manufacturing

The use of computers to automatically synthesise designs is not a new idea. Herb Simon's (Simon 1996) seminal paper "The Science of Design" provided the starting point for the academic pursuit of automated engineering synthesis methods. Schon (1992) extended this work by proposing the use of artificial intelligence to assist in the design process in the early 1990's. The research questioned whether computers are better suited to being 'phenomenologically equivalent' to designers, reproducing the thoughts and methods in which designers create parts, or instead if they have the potential to be more useful when acting as design assistants.

Many researchers have since reasoned about the suitability of computational design synthesis for different areas of the design process. Computational design synthesis (CDS) is a research area focused on approaches to automate design synthesis activities (Campbell and Shea 2014). Early in the design process, CDS deals with the issue of representing design functions and predicting design performance when various parameters are undefined (Chakrabarti et al. 2011). Towards the later stages of the design process, meaningful results are achieved by interfacing fields such as design theory, artificial intelligence, computational geometry and design optimisation to analyse

the behaviour of designs. CDS differs from traditional optimisation in that the goal of synthesis is to more broadly capture, and/or utilise design decisions made by human designers. Ideally, its use is more beneficial in situations where the human designer does not have clarity of the design avenues that they may pursue (Cagan et al. 2005).

Design optimisation is typically used for later-stage design when many decisions have already been made and exploration is formulated into narrow bounds to improve specific performance aspects (Krish 2011; Snider et al. 2013). Despite a substantial research effort, commercial CAD realisations for early-stage design are limited (Kazi et al. 2017). One of the primary challenges in adopting CAD in early-stage design problems is due to iterative and chaotic way in which designs evolve. According to Dorst and Cross (2001), design is not a matter of first formulating a problem and searching for a satisfactory design concept. Rather, it is about refining both the formulation of a problem and solution ideas concurrently. It is this co-evolution of design and solution spaces that make the problem of creating a phenomenological equivalence between human designer and a machine extremely difficult.

Practical solutions have, however, been developed as computational design assistants to aid designers in finding satisfactory solutions to ill-defined problems. Optimisation has been shown to not only find high-performing solutions but also to give designers a better understanding of the relationship between geometric features and design performance (Chen et al. 2015a). In addition, in a survey of design practitioners, Bradner et al. (2014) found that professionals use design optimisation to gain understanding about the design space, and not only to generate the highest performing solution. They further stated that the computed optimum was used as the starting point for design exploration, and not the end product.

Cagan et al. (2005) provided a framework for the automatic synthesis of design components. The framework contained four major activities that must be present in all synthesis systems, namely, representation, generation, evaluation and guidance. The representation stage defines the level of detail and focus of the computational search process as well as dictating the range of candidates that can be created. Generation methods are used to synthesise the geometry. Numerous classifications of generation methods can be found in Chakrabati et al. (2011). Evaluation is used to measure the worth or potential success of a design candidate and guidance is used to provide feedback to the system that can be used to generate improved solutions. In the field of AM, a design assistant that can synthesise a large range of design possibilities, before manual exploration and assessment by a human designer, is termed '*generative design'*.

CDS and generative design have similar definitions in literature. Shea et al. (2005) defines generative design as systems that create new design processes that produce spatially novel, yet efficient and buildable designs through exploitation of current computing and manufacturing capabilities. An alternative definition, provided by Wortmann (2018), defines generative design as a process in which computers derive a design by searching a parametrically-defined design space with an optimisation algorithm according to an objective function evaluated via simulations.

An alternative definition provided by Singh and Gu (2012) states that generative design systems are generally identified as systems aiming to support human designers or automate parts of the design process through computational means. Often, generative design is linked with a performance measure that drives the design generation (Oxman 2006).

Based on the above definitions, within this thesis, generative design is defined as computational systems that extend traditional CAD by automatically synthesising geometries based on a series of goals and constraints set by the designer.

Two common terms used within generative design are the 'design space' and the 'solution space'. The design space (also known as the problem space or sample space) is a finite set that mathematically defines the design (Bradner et al. 2014). It includes the design variables, constraints, and any other bounding criteria that can either be continuous or discrete. The solution space is a subset of the design space, described as the set of all solutions computed by the design synthesis (optimisation) algorithm. This is described diagrammatically in Figure 2-4. The rectangles in the design space represent the design information required to synthesise potential design solutions, and the solution space is represented by the evaluated design instances for a given design scenario.

Generative Design Processes

Figure 2-4: A diagrammatic representation of the design and solution spaces in generative design. Input parameter sets form the design space, and the solution space is created through a synthesis method, resulting in embodied design instances that are a subset of the design space.

One of the difficulties with generative design is the computational challenges that occur with the requirement of generating large numbers of functionally optimised parts. To improve computation times, generative design systems typically run on high performance computing (HPC) hardware, either as a cloud-based service (Kazi et al. 2017; Wu et al. 2017; Aage et al. 2017) or by using accelerated computing with GPUs (Wu et al. 2016). However, there are other methods including pre-computing various points by adaptively sampling the parametric design space and interpolating the mesh between these points (Schulz et al. 2017). In computationally-expensive problems, it is common to make use of surrogate models. These are approximated models of the objective function based on previously evaluated solutions (Forrester and Keane 2009). Nourbaksh et al. (Nourbakhsh et al. 2018) use a surrogate model to approximate stress in 3D trusses, and Tang et al. (Tang et al. 2017) use a Gaussian process surrogate to predict the stress within struts of different lattice unit cells.

It is possible to algorithmically generate a design space from a user-defined expression of the problem that they are trying to solve (Chien and Flemming 2002). This computationally generated design space can be traversed, and, at each location of this high-dimensional design space, an instance in the parametric design space can be evaluated and represented by a part geometry

within the solution space. The number of solutions generated depends on the sampling frequency of the design space, which is often limited to computational resource availability. One challenge with generative design is selecting the correct sampling frequency. This broadly poses the same challenges as other multi-objective optimisation strategies due to the curse of dimensionality (Bellman 2015). Synthesis methods, especially when reliant on simulation, can be very computationally expensive and so it is necessary to optimise the search strategy where possible.

There are two perspectives when it comes to overcoming the computational challenges of generating large numbers of functionally optimised solutions that trade-off between time and financial cost. Generative design tools provide the option to compute locally or on the cloud. Local computation generates fewer solutions in a given time, but is less financially costly per solution. Alternatively, a designer may decide to utilise expensive computing resources, for example, multiple cloud servers to generate a larger numbers of solutions. Although cloud resources can be efficiently utilised, they do not come without cost. Excessive use of computation, especially when generating infeasible solutions is both costly to the end user, and the environment, with data centres creating 17% of the world's carbon footprint (Dayarathna et al. 2016). There is, therefore, a push toward the efficient use of computation resources when utilising generative design strategies. Furthermore, if many poor performing solutions are generated, it lowers the sensitivity of the solution space leading to greater challenges in finding the highest performing solution within said space. In addition, research by Fricke (1999) has shown that designers who excessively generate solutions spend more time organising solutions and losing their overview of the design problem. This can lead to a failure to generate further good solutions.

Generative design does not automate the design process. Instead it aims to remove much of the cognitive burden from the user and place it in the computational realm (Shea et al. 2005; Campbell and Shea 2014). This is achievable as computers can perform certain tasks much better than humans. Research has shown that humans struggle to comprehend more than four variables at once (Halford et al. 2005), whereas computers can store substantial amounts of data and compute billions of calculations per second. The ability to synthesise optimal geometries based on simulation and create high-dimensional representations of design solutions is only practical on computers. However, as yet, computers lack the intuition and experience of human designers. Therefore, a collaborative generative design system involving human-computer interaction is preferable. By reducing the number of tasks the designer has to perform, they can spend their cognitive capacity guiding the computer to areas of the solution space in which they are most interested (Wortmann 2017; Goguelin et al. 2017).

Generative design has gained a considerable interest from industry, with many products appearing in recent years. However, industry use of the term generative design is less specific than the definition applied within this thesis. Many commercial products conflate the terms topology optimisation and generative design. With this in mind, Figure 2-5 depicts a range of commercial applications of generative design that utilise different geometry generation algorithms (Chakrabarti et al. 2011), including nature-inspired algorithms (Von Buelow 2012), and numerous multi-scale optimisation methods. Autodesk have produced a number of generative design solutions for multiple manufacturing platforms. These include an archery bow designed to be compatible with 3-axis CNC machining (Ayres 2015), an aircraft seat that was designed for investment casting (Schwab 2017) and an aircraft partition wall that was generatively designed for the AM process (Airbus Group 2016). nTopology (2018) produced a generative design tool to develop tailored cellular structures. Nervous System has worked closely with a sports shoe manufacturer to generate a data-driven sole customised to the individual athlete's running profile (Koslow 2015). Frustum (2018) have developed a generative design tool for topology optimisation using high performance cloud computing and a proprietary implicit modelling kernel. Users are able to set up multiple instances of a test part and compare the results after the computations have completed.

Figure 2-5: Commercial Generatively Design Components.
Shown left-to-right, (top) Archery bow (Ayres 2015)*, aircraft seat* (Schwab 2017)*, aircraft partition door* (Airbus Group 2016)*, (Bottom) Cellular component* (nTopology 2018)*, aircraft bracket* (Frustum 2018)*, training shoe soles* (Koslow 2015)*.*

2.4.1. Generative Part Synthesis for Additive Manufacturing

Stouffs and Rafiq (2015) assert that optimisation is rarely intended to yield an optimal solution, and can instead, assist a designer in gaining an insight into the solution space. Exploration and

optimisation together can lead to a better understanding of the complexities of design and help designers in their decision-making process.

Goal-driven design is formulated as a parametric optimisation problem. Typically, these problems consist of maximising or minimising a function by systematically choosing the input values from within an allowed set and computing the value of the function. The function that is to be minimised or maximised is termed the objective function. It is the designer's task to satisfy a set of functional requirements by transforming them into a description of a proposed artefact comprising of its form, material composition and dimensions (Kroll and Koskela 2016). Researchers have targeted the creation of geometry synthesis algorithms that can optimise part geometry to achieve a series of specific functional objectives.

Optimisation algorithms for many objective functions have been developed, including, stiffness (Zheng et al. 2014), thermal conductivity (Pizzolato et al. 2017), thermal expansion (Sigmund and Torquato 1997), fluid permeability (Challis et al. 2012), impact resistance (Schaedler et al. 2014), Poisson's ratio, acoustics (Li et al. 2016; Umetani et al. 2016), centre of gravity (Prévost et al. 2013; Christiansen et al. 2015), rotational stability (Bächer et al. 2014), energy harvesting (Zhakeyev et al. 2017), aerodynamics (Martin et al. 2015) and buoyancy (Wang and Whiting 2016).

In the field of mechanical engineering, much of the research exists in the domain of structural optimisation. In this setting, the primary objectives is to reduce the overall mass or maximise the stiffness of the final manufactured part. This is desirable as the reduction in material usage during the manufacture of the part can lead to cost savings and performance gains during their operational life. When optimising an AM design, it is possible to modify the geometry on a number of different scale lengths (Tang and Zhao 2016), the macroscale is typically concerned with geometries greater than 10mm, the mesoscale between 1mm and 10mm and microscale geometries being produced at sub-millimetre scale. There are several review papers that give an overview of research for goal driven design on the macro- (Rozvany 2009; Sigmund and Maute 2013; Liu et al. 2018), meso- (Schaedler and Carter 2016; Osanov and Guest 2016; Dong et al. 2017) and micro- scales (Cadman et al. 2013).

2.4.2. Topology Optimisation in Additive Manufacturing

Macro-scale optimisation is typically divided into three groups: size optimisation, shape optimisation and topology optimisation. Topology optimisation is a structural optimisation method used to find the optimal material distribution within a given design domain for a given set of loads and boundary conditions (Bendsøe and Sigmund 2003).

Topology optimisation has the ability to remove the most material from the initial design volume, and as such, this method has gained much attraction from the automotive and aerospace sectors (Rozvany 2009). Topology optimisation gives answers to the fundamental engineering problem: how to place material within a prescribed design domain in order to obtain the best structural performance (Sigmund and Maute 2013). There has been a rapid development in topology optimisation research in recent years, which is partially attributed to the rise of industrial AM giving users the ability to fabricate the synthesised geometries. A number of recent industrial studies show that AM adoption is likely to increase when parts are designed specifically for AM (Eggenberger et al. 2018; Khorram Niaki and Nonino 2017).

Several optimisation methods have been developed to optimise a parts topology. These include:

- Ground Structure (Bendsøe et al. 1994; Dorn et al. 1964),

- Solid Isotropic Material with Penalisation (SIMP) (Zhou and Rozvany 1991; Rozvany et al. 1992),

- Level set (Allaire et al. 2002; Wang et al. 2003),

- Homogenisation (Bendsøe and Kikuchi 1988),

- Evolutionary Structural Optimisation (ESO) (Xie and Steven 1993) and subsequently Bidirectional Evolutionary Structural Ooptimisation (BESO) (Querin et al. 1998)

- Genetic algorithms (Rajan 1995; Ohsaki 1995)

For a detailed comparison of the methods, readers are referred to the review by Sigmund and Maute (2013). Table 2-1 describes the advantages and disadvantages of these topology optimisation methods as described by Tang and Zhao (2016).

Table 2-1: Advantages and disadvantages of different topology optimisation methods. (Adapted from (Tang and Zhao 2016)*))*

Method Name	Advantages	Disadvantages
Ground Structure	- Easy to implement; - Suitable for low volume truss structures	The initial ground structure has a significant impact on the final result
SIMP	- Easy to implement; - Requires less storage space and computational effort	- Only effective as long as the final goal is to obtain a black and white design - Result largely depends on the penalisation parameter - Easy to generate check-like structure
Level set	The optimised result has a continuous boundary	- Solutions are different for different starting points; - It is sensitive to a number of algorithmic parameters
Homogenisation	- The concept of element with intermediate density is clear; - It can be applied to optimise the structure with an intermediate density	Large storage space and computational effort
Evolution	It does not require calculating the derivative of the objective function	It is sensitive to a number of algorithmic parameters
Genetic Algorithm	It does require calculating the derivative of the objective function	- Large computation load is needed even for simple problems; - It is more likely to converge to a local optima

The listed topology optimisation approaches fit within the continuum optimisation category with the exception of the ground structure method (Deaton and Grandhi 2014) and truss-based genetic algorithms. Continuum-based approaches have been the predominant structural optimisation techniques to date, with SIMP being used in a number of commercial software packages, e.g. Optistruct and Ansys (Rozvany 2009). In addition to the disadvantages provided in Table 2-1, the output results from continuum based optimisation solutions require significant manual post-processing to provide a viable solution suitable for manufacturing.

Ground structure optimisation generates truss-like macrostructures by resizing or removing potential truss elements with the aim of maximising a particular design objective (typically compliance). Ground structure optimisation has been used in many applications including in the design of an unmanned aerial vehicle (Maheshwaraa Namasivayam and Conner Seepersad 2011), and an air bracket for a supersonic car (Smith et al. 2016). In the latter case, the optimised result provided a 69% weight saving compared to the original bracket. A further advantage of the
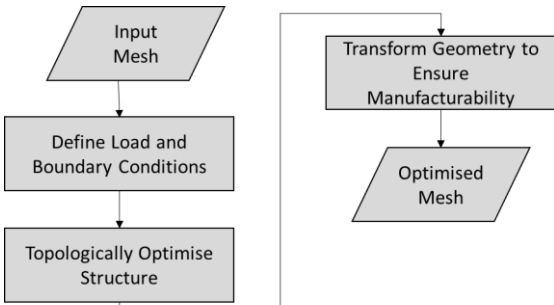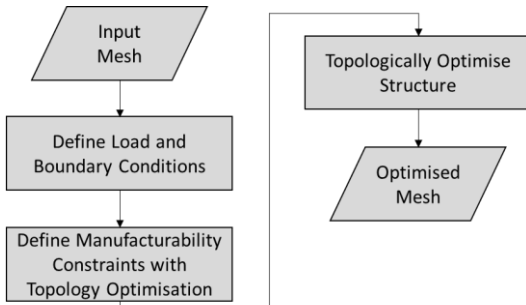
ground structure method is the ability to use beam finite elements that are faster to store and compute. Moreover, truss-based methods have been shown to be extremely efficient when the final solution occupies a small percentage of the original design volume (Smith et al. 2016). Commercial implementations of ground structure topology optimisation can also be found (Limitstate LTD 2018). Furthermore, ground-structure based structural optimisation can be constrained to have no overhanging features by removing all potential struts greater than the maximum overhanging angle from the initial ground structure (Mass and Amir 2016).

Sigmund (2011) states that gradient-based methods are far more efficient that non-gradient based optimisation approaches. As such, it is preferable to formulate the structural optimisation in such a way that gradient-based methods can be used to solve them. Evolution based methods may, however, be useful if the problem definition is unclear or unknown.

Historically, there have been two approaches incorporating AM constraints into topology optimisation processes. The first of these aims to post-process optimised solutions to ready them for additive manufacture. In the second approach, the optimisation algorithms are modified to create topology optimisation algorithms that intrinsically produce manufacturable geometries. For the most part, these constraints include prevention of excessively overhanging structures. An overview of the research is shown in Table 2-2.

Table 2-2: Classification of manufacturing constrained topology optimisation research

Approach 1: Post process topology optimisation result	**Approach 2:** Integrate manufacturing constraints into topology optimisation algorithm
References: (Brackett et al. 2011; Leary et al. 2014)	**References:** (Gaynor and Guest 2016; Carstensen and Guest 2018; Mass and Amir 2016; Mirzendehdel and Suresh 2016; Essink et al. 2017; Dhokia et al. 2017; Langelaar 2016; Langelaar 2017; Langelaar 2018)
Schematic:	**Schematic:**

Brackett et al. (2011) were early adopters of Approach 1 (Table 2-2). They used a filtering process to identify downward facing edges that extend beyond a defined 'manufacturable' length. Once identified, the authors introduce self-supporting lattice structures of differing densities to support these faces. Leary et al (2014) attempted to eliminate the support requirement by automatically identifying areas that fall outside the maximum overhang constraint. Their algorithm iteratively subdivides non-manufacturable segments until the geometry becomes manufacturable. Importantly, the authors show that manufacturing constraints inevitably increase the overall volume of the geometry for a given level of structural performance. A potential limitation of Approach 1 is that any post-processing of the as-optimised geometry may result in a contravention of the manufacturing constraints, or a reduction in the part's structural performance.

Approach 2 offers an alternative method that directly integrates manufacturing constraints into the optimisation process. Guest and Gaynor (2016) incorporated maximum overhang and length constraints into the optimisation process by using Heaviside projection. This process ensures that material can only be placed if it has supporting material beneath it, thus preventing the creation of 'unmanufacturable' geometries. Recently, this work has been extended into 3D to eliminate support structures for a given overhang angle constraint (Johnson and Gaynor 2018).

Mass and Amir (2016) used a truss-based topology optimisation, combined with a manufacturability filter, to generate an optimised, yet manufacturable, continuum structure. Their 'printability index' shows that as printability increases the performance of the part decreases. Mirzendehdel and Suresh (2016) also showed that performance decreases when a manufacturability criterion is applied to the design space.

Researchers at the University of Bath (Essink et al. 2017; Dhokia et al. 2017) proposed a bio-inspired algorithm based on termite nest building. A voxel-based approach is used to simultaneously design, optimise and appraise the structural performance and manufacturability of the part. Results show the ability to converge onto a structurally optimal and inherently manufacturable part without the requirement for an initial CAD geometry.

Langelaar (Langelaar 2016; Langelaar 2017) used an AM filter in order to generate manufacturable results by simulating the build process. The research focuses solely on constraining the design for overhang requirements. The results showed that the extra computational cost in applying a manufacturability filter is negligible compared to the optimisation time.

More recently, Langelaar (2018) integrated build orientation into the topology optimisation routine in order to minimise support structure. This work provides the first step in incorporating build orientation into the manufacturability constrained topology optimisation process, and results show that significant reductions in support structure can be made by including build orientation. However, as with other studies, this comes at the expense of part performance.

Studies that adopt Approach 2 constrain the topology optimisation using geometric manufacturing constraints. Until very recently, manufacturability constrained topology optimisation has focused on integrating overhang constraints into the optimisation, with only a single study including the effect of build orientation on the optimisation result. The initial positive results suggest it would be beneficial to examine this further by including further evaluation criteria and extending the method to 3D.

2.5. CAD Support in Design for Additive Manufacturing

Design is a complex process that often requires the ability to deal with ill-defined, 'wicked' problems (Buchanan 1992). Wicked problems cannot be definitively described because of the conflicting and ambiguous problem definition. Additionally, wicked problems do not deliver a way to state objectively when the design process is complete (*ibid.)*. There are also many cases in which CAD can be detrimental to the design process. The arguments for these will be explored in the following sections.

2.5.1. Cognitive Biases in the Design Process

By nature, human beings have a tendency toward biased decision making (Haselton et al. 2015). There is a tendency to misinterpret statistical data; make decisions according to insufficient evidence; interpret information in a way that confirms preconceptions and become fixated on information retrieved from memory (Hammond et al. 1998).

Hallihan and Shu (2013) researched the effect of confirmation bias in design. Confirmation bias is the innate human tendency to seek to validate beliefs instead of critiquing them. Confirmation bias can be problematic in the design process as it leads designers to over-rely on information that supports the notion that their designs are error-free rather than objectively assessing design information (Silverman and Mezher 1992).

Ownership bias can occur when a designer forms an association with a design that has required a considerable amount of effort and time to produce. Studies have shown a clear preference for a self-generated concepts in concept evaluation (Nikander et al. 2014). Computers have no such biases and, as such, can be used as decision support tools to augment the human designer's

ability to produce design solutions. Furthermore, simulation methods can give designers an understanding of the behaviour of the parts rather than solely the aesthetics.

2.5.2. Current CAD Usage within Additive Manufacturing

Design intent is a particularly important term when describing the design process. At the heart of design tools is the ability to communicate and represent a designer's ideas. Design representation may be thought of as a process of externalising design intent into physical forms (Atilola et al. 2015). This may include sketching, CAD models or physical prototypes.

The use of CAD has been criticised for its inability to form creative design representations in early stage design, causing designers to focus on details rather than underlying principles (Utterback et al. 2006). Based on this evidence, researchers (Lawson 2002; Veisz et al. 2012) have argued that the inexact nature of hand sketching is more advantageous than CAD modelling for early-stage design.

Current CAD tools also inhibit designers achieving optimal designs in other ways. Robertson and Radcliffe (2009) documented a number of disadvantages to using CAD systems in early design including circumscribed thinking and premature fixation. Circumscribed thinking describes the phenomenon in which the complexity of a design produced by a designer is proportional to their proficiency with the design tool. This has tremendous implications on DfAM, as the shape complexity, and part consolidation required for optimal AM parts require an expert level of CAD knowledge. Designers will have to accrue substantial experience on a CAD system before they can exploit the true potential of AM (Despeisse and Minshall 2017).

Functional fixedness is a cognitive bias defined by Duncker and Lees (1945) as being a "mental block against using an object in a new way that is required to solve a problem." This is an issue when traditional CAD tools are required to create new types of geometry that designers are not used to developing with such tools. This is compounded in AM, as designers often lack the required skills to deal with AM-specific software (Despeisse and Minshall 2017).

Building on the notion of functional fixedness, Jansson and Smith (1991) describe design fixation as designers tending to move toward a premature commitment to a particular problem solution without exploring the solution space any further. An integral part of DfAM is avoiding fixating on designs that are specific to traditional manufacturing methods. Crilly (2015) reviews design fixation and suggests factors that discourage design fixation. These include collaborative design, specific design tools including morphological analysis, and prototyping. Whilst these factors may increase the number of design concepts, novel conceptual design for AM can only be created when the designer fully understands the design requirements for their particular part and the

complete capabilities of AM. However, it also suggests that multiple ideation methods should be used when creating a diverse range of design concepts.

Abdelall et al. (2018) investigated design fixation in AM. The results showed that designers focused on non-producible features in traditional manufacturing methods rather than focusing on performance benefits. Furthermore, the designs for AM violated more design guidelines than traditional manufacturing methods, indicating that users misunderstand the notion of design freedom. In a study of 26 participants, 83% of participants stated that their modelling skills affected their ability to transfer their sketches into CAD and 36% failed to translate their designs fully into CAD (note that this was not a complex CAD modelling task). Results showed that once participants have designed for AM, it is harder to modify the design to conventional manufacturing techniques. This highlights the importance of designing specifically for a given manufacturing process. The results also document that people confuse AM manufacturing freedom capabilities with the design constraints of the actual task.

With premature fixation, the designer feels less incentive to make significant design changes as the object becomes more complex and the designer invests more time. In conceptual design there is co-evolution of the problem space and the solution space (Maher and Poon 1996; Dorst and Cross 2001). Therefore, there is a requirement to change designs based on new knowledge. Complex shapes can be difficult to manipulate in traditional CAD software. Therefore, the use of traditional CAD for optimised AM parts can be challenging.

Overcoming the aforementioned issues with CAD systems may be possible by creating generative design systems that both automate the synthesis of optimised geometry and explore trade-offs that occur within complex design spaces by evaluating designs based against multi-dimensional criteria.

2.6. Critique and Research Gaps

From the literature presented in this chapter, several DfAM research gaps have been identified and are detailed in the following paragraphs.

Firstly, to exploit the advantages of AM, it is necessary to design parts that can be fabricated without build failures. Many design guidelines have been created to inform designers about the physical limitations of the process. However, it is often challenging to translate these design guidelines into complex freeform geometry (Pradel, Zhu, et al. 2018). It would be more beneficial to integrate design constraints into generative geometry synthesis algorithms to ensure manufacturable designs are automatically produced by the CAD tools.

Automatically synthesising geometry has the further advantage of overcoming many of the challenges that users face with current CAD tools. By creating design tools that incorporate the ability to explore multiple design solutions, much greater spans of the problem space can be explored. Using generative synthesis techniques, such as topology optimisation, can reduce the time between the intention of a particular CAD geometry and the act of creating it. This gives more choice to the designer when selecting from design concepts, whilst also avoiding many issues associated with cognitive biases from the designer.

Researchers have developed geometry synthesis algorithms for many functions. However, these are typically low-level functions. For example, the mass or compliance of a part. Topology optimisation alone is not sufficient to design functional parts. This is because design is a complex and often an ill-defined process (Simon 1973). Understanding the mass or compliance of a part, while sufficient, is not enough to determine the quality of a design. Even by integrating design constraints into the optimisation it is still possible that the generative synthesis methods provide designs that will be expensive to produce. In the real world, the field of product and part development involves many compromises between high quality, low cost and high development and production speed. Originally referred to as the 'Iron triangle (Atkinson 1999),' or triple constraint, and shown in Figure 2-6, designers have been searching for ways to easily explore the trade-offs between these three criteria for many years. Future design tools should be developed that can abstract from low-level evaluation criteria to include these higher-level business goals to allow users to easily explore designs based on factors in which they are interested.
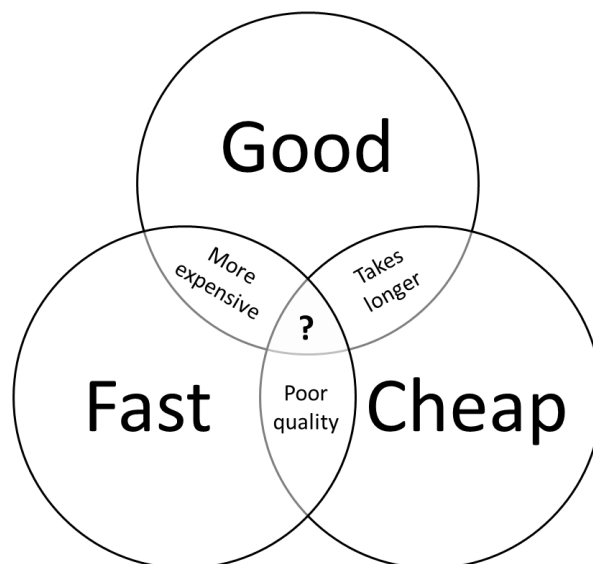
Figure 2-6: The 'Iron Triangle' of product development.

Commercial topology optimisation solutions currently lack knowledge of the design context in which they are optimised. In order to improve this, topology optimisation should be run with context of the proposed build orientation in which the part is to be manufactured. Furthermore, due to the detrimental effects of excessive support requirements, options should be included to ensure that the output solutions can be manufactured without the requirement for support structures. The time required to perform structural optimisation can limit the number of design solutions that can be generated. Therefore research should be directed towards developing efficient search methods that can find high-performing regions of the solution space without requiring large numbers of iterations.

2.7. Summary

In this chapter, a comprehensive critique of the literature outlining the development of DfAM methods and CAD tools have been outlined and the research gaps driving future research have been described. The research gaps highlight many potential avenues for future research guiding the development of future CAD tools to improve and support the way in which designers will develop and manufacture end-use parts for AM that will truly exploit the benefits of the technology. Three main research gaps were identified that will guide the development of the work research within thesis, these are:

- Incorporating of build orientation and manufacturability constraints into generative synthesis algorithms such as topology optimisation.

- Extending the range of synthesis algorithms to include high-level criteria such as business criteria and different AM production scenarios.

- Overcoming the limitations of CAD support in DfAM by creating simple and efficient methods to navigate parametric AM design spaces. These may include the use of data visualisation techniques, MCDA decision support tools, and surrogate design optimisation models.

Based on the critique of the literature, the author's view of the future perspectives for the important features of generative design systems will be formulated into a generalised framework that aims to aid researchers and generative design practitioners in the formulation of generative design enabled CAD tools. This will be realised in Chapter 3.

Chapter 3 – A Framework for Generative Design CAD Systems

3.1. Introduction

The literature review in the Chapter 2 outlined the research gaps in the field of computer-aided DfAM. In the interest of providing an overarching view on the development of next generation generative design tools, a three-stage CAD framework is established. This has been defined based on the gaps identified in the literature critique and, in the context of this thesis, acts as an anchor point from which new CAD tools can be developed. This framework has the potential to enable the creation of these new tools that are capable of overcoming current state-of-the-art CAD limitations.

3.2. A Framework for Generative Design CAD Systems

This section details a framework that consists of three interlinked stages, namely, 1) defining the design space, 2) solution development and 3) solution output. Figure 3-3 depicts the complete 3 stage framework. The figure is intended to be read from top to bottom with the central column representing actions that must occur in generative CAD systems. The left and right hand columns represent the inputs and outputs from the framework respectively.

The primary stage describes the ***defining the design space***. The proposed generative framework incorporates a series of databases of criteria that may be specific to the part. The databases include numerous design considerations extracted from part design specifications, for example, those shown in Table 3-1. Additionally, it could also contain company-specific parameters to ensure design continuity throughout an enterprise.

Table 3-1: Example of quantitative requirement inputs for generative design CAD framework.

Specification Criteria	
Quantity Required	Maximum Mass
Maximum Cost	Maximum Overall Dimensions
Thermal Conductivity	Dimensional Tolerances
Surface Roughness	Product Lifetime
Volume	Maximum Deflection
Surface Area	Minimum Stiffness

Further information regarding the processes must also be defined. For example: machines, materials and in turn the relevant process parameters, as outlined by the tree structure diagram

in Figure 3-1. In this diagram, the processes correspond to the different AM classifications as outlined in Table 1-1.The term 'machines' denotes the choice of the many available AM machinery available to the user. This might be selected from availability within a factory, laboratory, or an AM hub. Next, the available materials that can be processed on each machine must be defined, alongside their properties. This allows context of the manufacturing process to be integrated into the generative design process.
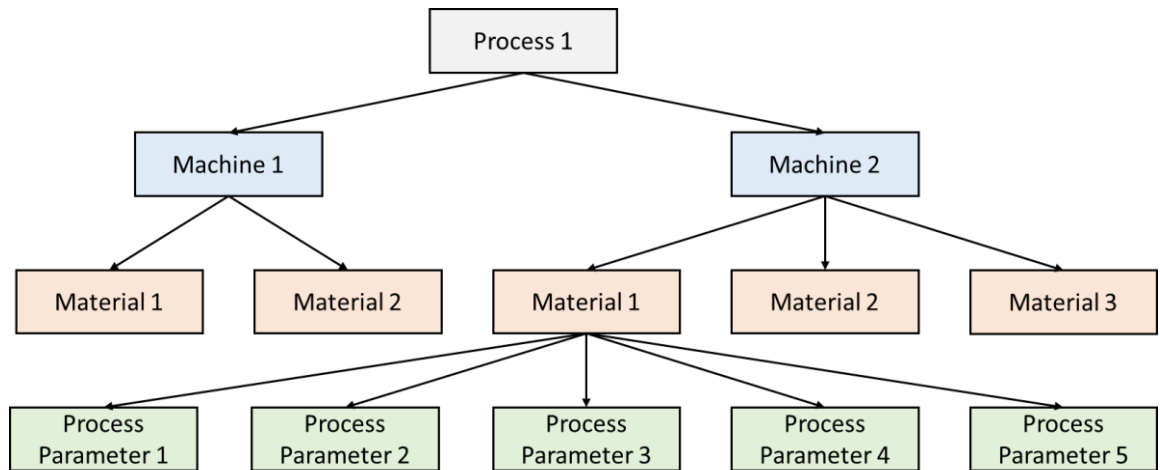
Figure 3-1: Tree structure definition of input requirements in the primary stage of the CAD framework for generative design.

Finally, the process parameters that can be altered must also be defined; examples of these can be seen in Table 3-2. It is important to ensure that all of the evaluation metrics can be defined by the information derived from the process, material and geometric information. In addition, as described within the DfAM pipeline, the importance of post-processing must not be understated. Therefore, it is prudent to include the post-processing machine availability and limitations including inspection tools in order to fully define the capability to produce end-use, production ready AM parts.

Table 3-2: Example of machine parameter inputs to generative design CAD framework.

Machine Parameters	
Layer Thickness	Recoater Time
Hatching Strategy	Deposition Rate
Laser Diameter	Scan Speed
Hatch Spacing	Build Platform Area
Argon Usage	Energy Usage
Nozzle Diameter	Bed Temperature

The designer is then required to extract the relevant design parameters from the databases. At this stage it is important that the trade-offs within the technology have been captured. Consider a

design that is optimised for weight reduction and, as such, a lattice structure is employed in the design. Whilst this is appealing and solves the mass issue, the part becomes challenging to inspect and cannot therefore be certified for use as a functional component. In this case, the designer will have to create a hierarchical structure of weighted priorities for the design.

The constraints on the design problem must also be defined. The constraints define the overall shape of the design space. Constraints can either be hard or soft. Hard constraints limit the envelope of the design spaces, therefore eliminating these solutions from ever being generated. On the other hand, soft constraints work by penalising certain solutions. This could either be in the objective function or can be used within the evaluation stage to eliminate certain design instances from the solution space.

A series of material void spaces (MVS) must also be defined that outline areas in which material cannot be placed by the algorithm. In addition, a design solution volume (DSV) is defined as the volume to which the generative synthesis method can apply the material; this could be user-defined or taken as the bounding box of the MVS features. This is shown in the example in Figure 3-2. This might include additional hardware components for attachment or for assembly tool access. From this information, a generative algorithm can be used in conjunction with optimisation techniques to deposit material within this solution volume.

Figure 3-2: Design solution volume (centre) and material void space (right) definitions, taken from initial CAD geometry (left) used in topology optimisation and generative design definitions.

The second stage of the framework is termed **solution development**. Here, solutions are optimized for the trade-offs selected from the hierarchy stated in the previous stage. At stage 2 in the concept generation, it is solely the quantitative parameters that will be considered. By taking advantage of multi-objective synthesis algorithms, the designer can program many inputs into the system. The design adapts within the solution space, converging as close as possible to objective function defined by the designer. In order to generate a series of solutions, sequences of open parametric design space variables must be selected that can be optimised during the synthesis methods. This automated synthesis of part geometries can improve the premature fixation that designers face when designing complex geometries, as well as other cognitive flaws, including ownership bias. Based on the literature gaps, the synthesis algorithms should have the facility to

incorporate the build orientation and any manufacturability constraints derived from the selected/available AM processes.

At this point that the designer has the option to impart some of the qualitative knowledge they possess. Examples include, part aesthetics and human interaction with the part. The designer will also have the opportunity to vary input parameters from the databases in the primary stage. Qualitative design information is particularly challenging to capture, as users can rarely articulate why they like a particular aspect of a design. Therefore, it is only at the comparison stage in which users can express this tacit knowledge due to the ability to compare designs within a set.

The CAD tool will then use the best solution(s) as selected by the designer from the first generative stage as a new input(s). This is achieved by generating more appropriate designs based on the modified input parameters and defined qualitative information improving the capture of the designer's initial intent. This approach is then repeated, within a user feedback loop, until a satisfactory design can be delivered by the system. The ability and speed of regenerating further design solutions are integral to the system. The overall quality of the design solution is dependent on giving the designer the ability to redefine the problem space as more knowledge is gained about the design direction.

The final stage of the framework is labelled *solution output*. For the design tool to be useful, the design would require exporting in different print files types currently used in AM, e.g. STL, AMF, or 3MF however, this should also be suitable for printing with meshless representations directly from machine code. In conjunction with the AM file formats, the tool should also give some indication of the manufacturing information including, for example, the correct build orientation and build strategy along with suitable post-processing techniques in order to finish with a functional part.
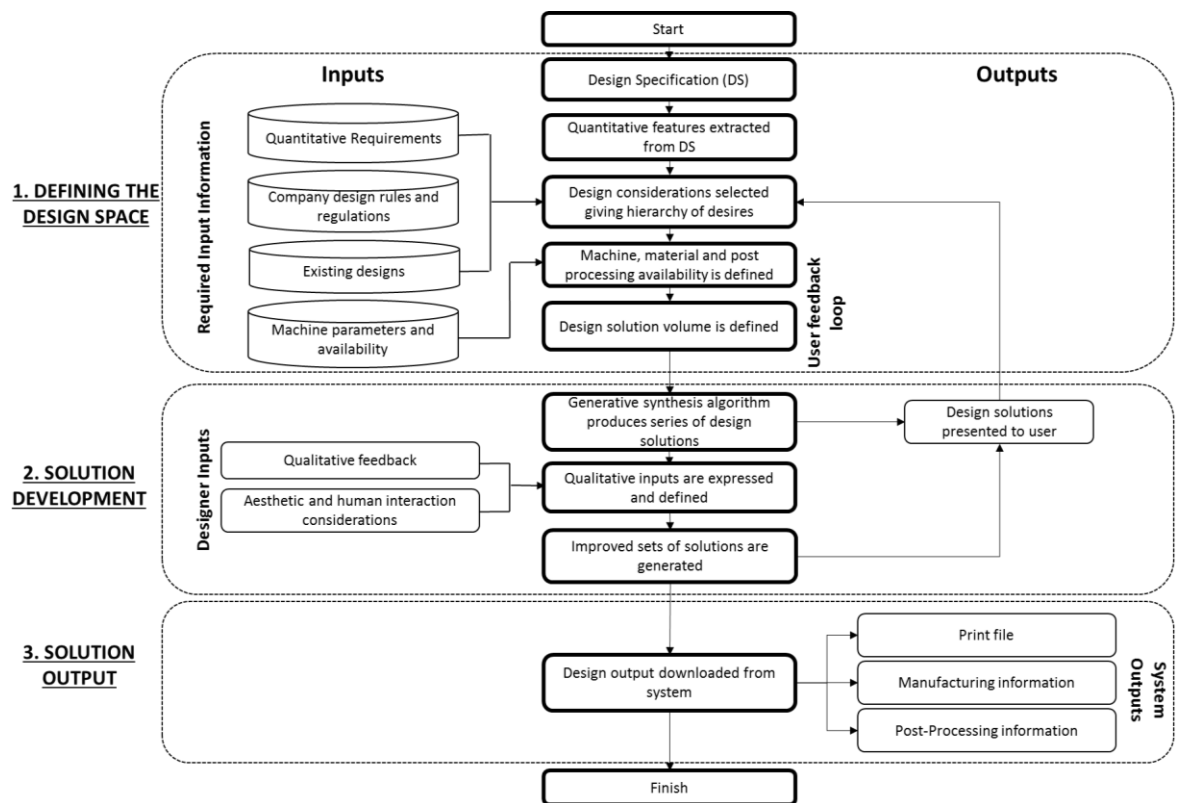
Figure 3-3: Generalised generative design CAD tool framework.

3.3. Summary

Two feasible instances of generative design methods derived from the framework, termed data-driven and goal-driven, can be used as a basis to determine the effectiveness of generative design to efficiently create AM-specific parts.

The manner in which the two approaches generate the solution space differs. Data-driven approaches generate the entire solution space before presenting the results back to the designer. These solutions can then be navigated *a posteriori*, using for example, a combination of data visualisations and decision support tools, to locate solutions according that meet the design specification.

Alternatively, the goal-driven approach aims to efficiently generate high-performing designs by learning the underlying representation of the design space. This is achieved by modifying the parametric design parameters based on the previously synthesised parts until areas of the solution space providing high-performing designs are located.

These two implementations aim to overcome a number of the research gaps, firstly, by creating geometries that are specifically designed with the manufacturing limitations of the machine. Secondly, by providing designs that can be optimised for high-level design objectives defined in stage one of the framework. The research aim, objectives and scope will now follow in Chapter 4.

36

Chapter 4 - Research Aim, Objectives and Scope

4.1. Introduction

A generalised CAD framework to support the creation of generative design tools to support DfAM is described in Chapter 3. Two implementations of the framework were created with the purpose of overcoming the gaps within the literature. The first of these is a data-driven generative design method that navigates a pre-generated solution space, a posteriori, to locate the highest performing solutions. The second implementation is a goal-driven generative design method that efficiently creates solutions by learning the underlying mathematical topography of the design space.

This, along with the research gaps in Chapter 2, lead to the following research question:

Is it preferential to use data-driven or goal-driven generative design methods when exploiting DfAM to optimise AM parts for multiple high-level design objectives?

To answer this research question, the following aim, objectives and scope have been defined.

4.2. Research Aim

This aim of this research is to identify the effectiveness of both data-driven and goal-driven approaches to generative design in developing AM-specific parts that are optimised for high-level design objectives.

The ensuing null hypothesis arising from this aim is that a goal-driven approach to generative design produces part designs with comparable performance using the same number or more model evaluations/iterations as the data-driven approach.

4.3. Research Objectives

To achieve the above aim, the following research objectives have been outlined:

1. Based on the research gaps identified in the state-of-the-art literature review, develop a generalised CAD framework that will support the development of new design tools to support generative design for AM.

2. Adapt an existing structural optimisation technique to be capable of incorporating manufacturing constraints and build orientation to give an interpretation of structural performance within an AM context.

3. Define the abstraction criteria that connect low-level evaluation criteria to high-level design objectives. Low-level evaluation criteria are closely associated with the metrics derived directly from synthesised mesh data. Alternatively, high-level design objectives are abstracted away from the part data and represent criteria such as part cost and build time. These objectives can then be used to form generative design goals that are representative of typical AM production scenarios.

4. Describe a data-driven generative design method to locate designs within a solution space based on the aforementioned, abstracted design objectives and create a series of interactive data visualisations that can present this information back to the designer.

5. Examine the efficacy of surrogate optimisation methods to efficiently explore build orientation angles with the aim of minimising the total amount of support structure required to manufacture a given mesh file.

6. Based on the results from the surrogate optimisation tests, integrate this method into a goal-driven generative design method.

7. Compare the results of data-driven and goal-driven generative part design approaches in order to identify the advantages and disadvantages of the methods in terms of their ability to explore and synthesise part geometries.

4.4. Research Scope and Boundaries

According to the aim and objectives of the research, the following areas of investigation are identified. Furthermore, due to the vastness of the DfAM research field, a series of research boundaries are defined to ensure that the research objectives can be completed. These boundaries concern DfAM and topology optimisation constraints, in addition to the evaluation metrics used to define the user-selected design objectives. A schematic of the research scope is shown in Figure 4-1. The triangle represents the research boundaries. The blue circle encloses the features that will be included in this research.
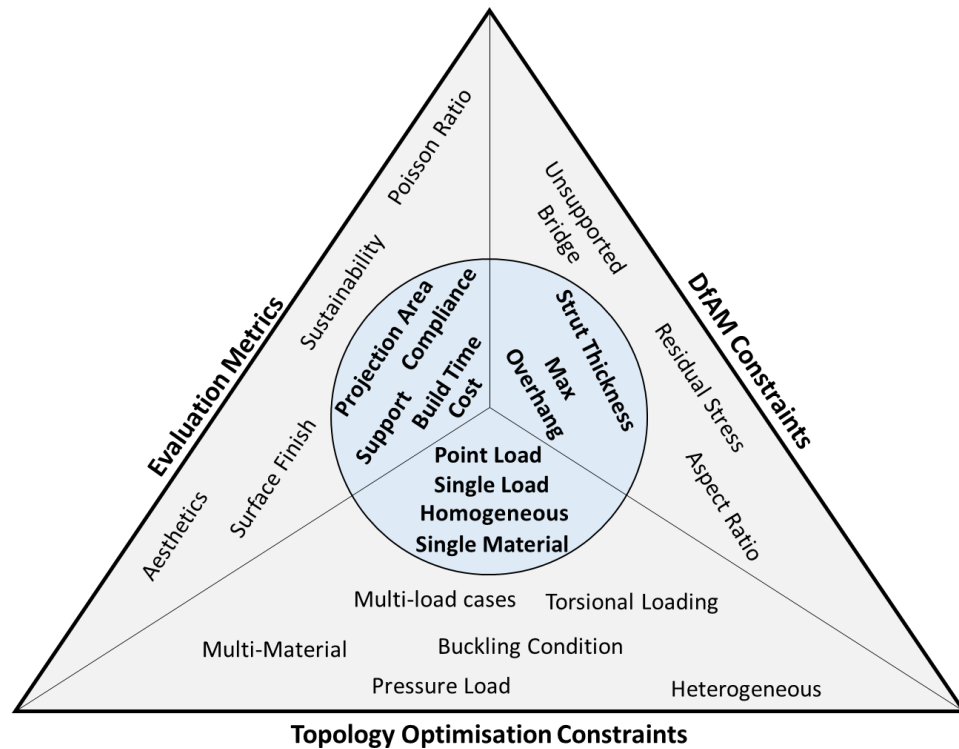
Figure 4-1: Schematic of thesis research boundaries.

4.4.1. Design for Additive Manufacturing Constraints

The literature review in Chapter 2 outlined the guidelines that have been developed to guide designers through the process of developing first-time-right designs.

These guidelines can be extensive and include information regarding the minimum thickness of supported and unsupported walls, tolerances for connected and moving parts and minimum hole and font sizes. Whilst, these are essential constraints for many designs, this research will focus primarily on two features, namely, minimum strut thickness and maximum unsupported overhang angle. This selection is attributed to the simplicity of integrating these constraints into the ground-structure topology optimisation method. Additionally, the impact that constraining the maximum overhang angle can have on reducing the volume of support structure, and in turn the final cost of the final part.

4.4.2. Adapt Existing Structural Optimisation Methods to Incorporate Additive Manufacturing Constraints and Build Orientation

Despite previous work that has begun to implement optimisation of mechanical assemblies (Zhang et al. 2017), this research will focus solely on single part design for additive manufacturing. These parts will be defined by the load and boundary conditions as well as the initial design space volume and material void spaces. Secondly, as the research aims to determine the effectiveness

of data-driven and goal-driven generative design, new topology optimisation methods will not be defined. Instead, a number of boundaries are placed on the topology optimisation. These include

- Only single load cases
- Only point loads will be applied
- Only single materials will be used.
- Only homogeneous materials will be used
- Only single parts will be optimised
- No buckling constraints will be added

In order to fully understand the manufacturability of additive parts, multi-physics simulation methods should be integrated into the system. The success of AM builds, especially in metal AM, are dependent on thermal gradient build up in the part during the build; this is of particular importance when building parts with high aspect ratio. While ample research has been undertaken into developing thermal simulations for AM, this will be outside the scope of the research within this thesis.

It is believed, that these boundaries will make for a more intuitive understanding of the results and that further integration of more advanced topology optimisation techniques can be implemented in future work as necessary.

4.4.3. Developing High-Level Objectives from Low-Level Evaluation Criteria

Depending on the stakeholders in a particular project, there may be a requirement to design parts based on different levels of abstraction. This research will only deal with design problems that can be formulated into an objective function. Therefore, for 'wicked' problems, the designer may have to perform initial work to reach this point in the design process. Within this research, the use of low-level evaluation criteria, such as facet-data from the optimised part, will be abstracted to develop evaluation criteria that are representative of high-level business goals. For example, different AM production scenarios.

One of the prominent concepts identified in the literature review with regards to generative design of AM parts is the ability to use computation to augment a designer's ability to develop optimised parts for specific design goals. The metrics used to evaluate the design should be developed to produce results that can improve the understanding the designer has about a particular design problem. To achieve this, metrics that are calculated using techniques that cannot be inferred solely by visual inspection of the design are required. For example, those resulting from simulation methods.

One of the main challenges in design is that there are many criteria, such as beauty, style or individual preference that cannot be quantified. With the advancement of machine learning algorithms, there are methods that can be employed to assert individual preference on the final design of a part. Examples include, style transfer or recommendation systems. Furthermore, human data tracking has the power to record implicit thoughts that a user may have when analysing a particular design instance, giving the ability to track the tacit elements of a design problem.

As qualitative aspects of design are paramount, reference to including qualitative quantities were integrated into the development of the generalised CAD framework in Chapter 3. However, throughout the implementation of data-driven and performance-driven generative design methods within this research, only quantitative elements of design will be assessed. This will be achieved using a combination of simulation techniques alongside measurements derived from mesh analysis techniques.

Chapter 5 - Research Methodology

5.1. Introduction

To achieve the aim and objectives set out Chapter 4, the following research methodology has been devised and implemented. The experimental approach to achieve the research aim will be explained alongside the algorithms and the software implementation used to perform this research.

5.2. Research Methodology

The purpose of this methodology is to answer the research question. To achieve this, it is necessary to outline the approaches for both data-driven and goal-driven generative methods and the approaches used to objectively compare them. The methodology follows a positivist stance. The selection of each research method has been made to follow a robust, quantitative approach to ensure the research question can be answered using inductive reasoning in an objective manner. Figure 5-1 shows the generative design methodology adopted within this research.
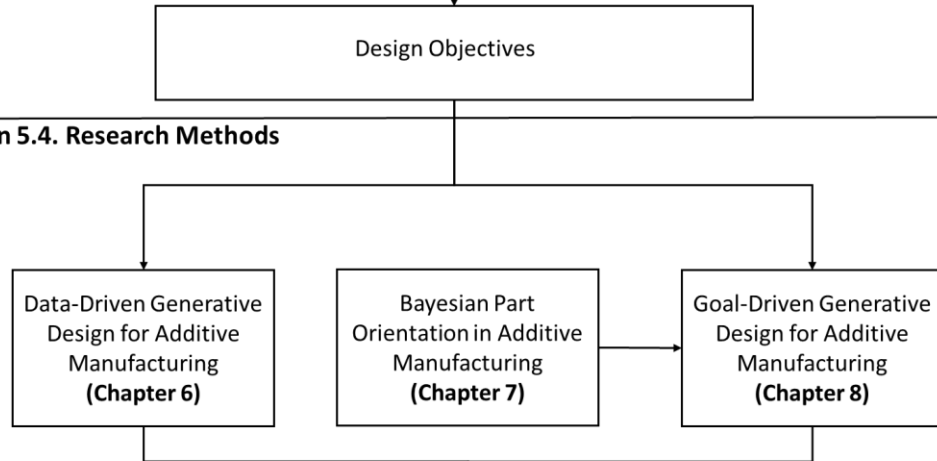
Figure 5-1: Schematic of the underlying research methodology used in this research.

This methodology will provide the knowledge required to objectively answer the research question posed in Chapter 4. This is achieved by evaluating both data-driven and goal-driven generative design methods for their capability to automatically synthesis high-performing design solutions and also their computational efficiency. Direct comparison and critique of the two methods will provide an objective view on the suitability for each method to be utilised in next generation state-of-the-art generative design CAD tools to support DfAM.

The research methodology consists of three stages. Initially, a design problem is developed that is representative of an AM design problem, consisting of a manufacturability constrained topology optimisation in conjunction with build angle. The two generative design methods will then be used to solve this design problem. The results of which will be compared in order to assess their suitability for solving generative design problems.

Design Problem: The design problem used to test the two approaches consists of a cantilever beam that is to be optimised using a ground structure-based optimisation with manufacturability

constraints, namely minimum strut thickness and maximum overhang angle. The cantilever beam will have the freedom to rotate about the y-axis between angles of -90° and 90°. Further explanation of the methods underlying the design problem can be observed in Section 5.3.

Research Methods: The use of two differing generative design methods will be explored within this thesis. Firstly, within Chapter 6, a data-driven generative design approach, where a MCDA will be used to filter the best performing solutions represented by high-level design goals.

Secondly, the inverse problem, or goal-driven approach. This method seeks to find optimal parameter settings to maximise the performance of a given set of high-level goals. A pilot study examining the ability to locate optimal build orientations for two AM specific parts, with the aim to reduce support structure requirements will firstly be tested within Chapter 7. The problem of build orientation is selected as it is representative of a two-dimensional input problem without the requirement to run costly topology optimisations.

Subject to the success of the pilot study, the goal-driven approach will be applied to the cantilever beam problem in Chapter 8. A series of objectives are combined to define user-specific AM goals which are then formed into objective functions. As the goal-driven approach described within this research is a stochastic process (attributed to the random initialisation process), it will be run for 21 repeated tests for robustness, and the median output value taken. The research methods will be explained in Section 5.4.

Comparison: The two methods will be analysed using statistical similarity metrics to determine their ability to effectively locate high-performing solutions within an AM solution space. The description of the comparison techniques can be viewed in Section 5.5.

5.3. Design Problem

As identified in the research gaps in the critique of the literature, a method of incorporating shape synthesis techniques such as topology optimisation with manufacturing constraints and build orientation is necessary to develop generative design systems capable of outputting manufacturable designs.

The design space is a mathematical representation of the underlying design problem. One of the challenges in generative design is selecting a suitable design space that can be exploited by the interactive search between the human designer and the computer. As each parameter of the design space will have an impact on the design output, the challenge for the designer is to select a set of parameters that maximises the potential advantages of computational approaches.

As each design variable is a dimension in the design space, it is possible that many design problems will be high-dimensional. If the design space is too simplistic (Figure 5-2), it is not necessary to use computational design tools and the designer is unlikely to learn anything new about the design problem by using a generative design approach. On the other hand, if the design space is chaotic as in the right hand plot of Figure 5-2, it will be practically impossible to learn the underlying design landscape making goal-driven generative design approaches difficult. Therefore, the designer's challenge is to create suitable design space landscapes that are complex enough to provide useful results whilst also being suitable for computational exploration.

Too Simple

Suitable Design Space

Too Random

Figure 5-2: Visual description of good and bad design space definitions shown in 3D space.

As detailed in the literature review, there are many different topology optimisation techniques. This research will focus on the ground structure method, which has a number of advantages over other methods. Firstly, the ability to represent the stiffness matrix using simplified bar elements. Secondly, decreasing computation times and allowing generative design to be performed without the use of HPC. Finally, it is straightforward to constrain the overhang angle of the struts by pruning the ground structure before the optimisation begins.

In truss topology optimisation, a problem can be formulated using a ground structure model. This ground structure is a highly-connected truss that connects all truss nodes to each other. The ground structure is placed within a design volume shown in Figure 5-3. The thickness of each strut is unknown before optimisation. The optimisation process then assigns thicknesses to each strut, resulting in each becoming a structural member with a set diameter or vanishing to zero thickness. The optimal truss contains a subset of the struts that exist within the ground structure. The aim of the optimisation is to minimise the compliance of the truss, i.e. $\boldsymbol{F}^T\boldsymbol{u}$. Here, $\boldsymbol{u}$ contains the displacements of the unconstrained nodes in the truss, and $\boldsymbol{F}$ contains the external forces at these nodes.

The optimization problem is defined as follows:

$$\min_{x,u} F^T u \quad \text{subject to,}$$

$$K(x)u = F \quad \text{and}$$

$$\sum_{j=1}^{n} l_j x_j \leq V_{max} \tag{1}$$

$$x \in \chi = \{x \in \mathbb{R}^n : x_j^{min} \leq x_j \leq x_j^{max}, j = 1, ..., n\}.$$

The parameter, n, is the number of struts, l_j is the length of the j^{th} strut, x_j is the cross-sectional area of j^{th} strut, and V_{max} is the maximum permissible volume of the truss. The matrix $K(x)$ is the global stiffness matrix of the structure, and x_j^{min} and x_j^{max} are the lower and upper bounds on the design variables. This configuration creates a non-linear optimisation problem that is solved using the first-order gradient based Method of Moving Asymptotes (MMA) (Svanberg 1987) within the NLOpt Python module (Johnson 2017). For a complete mathematical derivation of the truss-based optimisation, readers can refer to Haftka and Gürdal (2012) or Christensen and Klarbring (2009). A load is placed in the centre of the y-axis length creating a design problem corresponding to a typical cantilever beam.

Figure 5-3: Design volume setup for the cantilever beam design problem.

This research will take structural compliance as an indication of the structural performance of the part. Whilst compliance is not equivalent to structural strength, Aage et al. (2017) show that compliance based topology optimisation generally results in geometric structures with a favourable strength response.

5.3.1. Creating the Ground Structure
The size of the initial ground structure is selected by the designer; there are a number of options when creating a ground structure. The ground structure is classified by the valence connectivity

between the beams. A fully connected ground-structure will connect each beam to all other beams, whereas a partially connected ground structure will only connect beams that are within a pre-defined number of nodes from the initial beam. This is shown in the 2D examples in Figure 5-4.

Figure 5-4: Fully (top) connected ground-structure where all nodes are connected, and partially (bottom) connected ground-structure where only the nearest nodes are connected.

The size of the ground structure is defined by the number of potential connection nodes in the x, y, z directions such that a {3, 3, 3} grounds structure will have three nodes in each of the orthogonal axes. The size of the ground structure will have an impact on the final structural performance of the design, and should be selected based on the initial design space volume generated in the design problem to ensure a practical truss layout. Based on results from previous experiments (Smith 2016), it was decided that a {7, 3, 3} fully-connected ground structure was found to perform best for the given design problem. The length (x-axis) of the design space volume is defined to be 100mm with the y and z lengths defined to be 35mm respectively.

5.3.1.1. Integration of Additive Manufacturing Constraints

Whilst AM provides designers with increased geometric freedom, it does not provide complete design freedom, and thus it is necessary to adhere to certain manufacturing constraints (Kranz et al. 2015). Within this research, two key geometric design constraints, maximum overhang angle and minimum strut thickness are applied.

Figure 5-5 (left) shows the ground-structure for the truss without the overhang constraint applied and (right) when the ground structure has been pruned with a user-specified overhang constraint.

Figure 5-5: Ground structure shown without (left) and (right) with a manufacturability constraint applied

Maximum Overhang Constraint

A manufacturable ground structure is obtained by removing struts that have overhang angles greater than the maximum overhang angle constraint. The overhang angle, θ, is calculated for each strut by finding the angle between each strut, $\vec{u}$ and the build plate normal vector, $\vec{v}$ using (2). A filter is then applied to remove any struts that violate the conditions in (3).

$$\theta = \arccos\left(\frac{\vec{u} \cdot \vec{v}}{\|\vec{u}\| \, \|\vec{v}\|}\right) \tag{2}$$

$$\theta \in \mathbb{R}^n : \begin{cases} 0 \leq \theta_j \leq \theta_{maxoverhang} \\ 180 - \theta_{maxoverhang} \leq \theta_j \leq 180 \end{cases}, \quad j = 1, \dots, n \tag{3}$$

Minimum Strut Thickness Constraint

The minimum strut thickness is determined as the minimum strut diameter that is printable, this constraint is added into the x_j^{min} feature in the topology optimisation formulation described in (1). A value of 0.5mm is defined based on the average suggested DMLS results from the design guidelines set out in Section 2.2.

5.3.2. Build Setup

A diagram showing the two-stage rotation process can be seen in Figure 5-6. To determine the orientation of the part, a two-step procedure is implemented, as shown in Figure 5-6. Firstly, the part is rotated by an angle, θ, about the global y-axis. Then, the part height is readjusted by translating in z. The final translation in the z-direction is such that the lowest facet on the triangular mesh is aligned with the height of the build plate. In this study, θ belongs to the set, $\theta \in \mathbb{Z}: \theta \in [-90, 90]$.

Figure 5-6: Two-stage process used to define the build orientation for the data-driven and goal-driven generative design approaches.

5.3.3. Design Objectives

Design abstraction can be described as a reductionist process of removing successive levels of detail from a design representation in order to capture only the essential features of the system. Many studies support the idea that abstraction can help designers think of problem solutions more clearly (Casakin 2007; Hayes 2013). Abstraction is important in design as it allows designers to manage the complexities associated with the multi-dimensional and multi-objective nature of the design process by purposefully obfuscating unimportant design details (Hoover et al. 1991).

Research has shown that the designer's ability to abstract to alternative representations is linked to their experience, with expert designers adapting and contextualising experiential knowledge and applying it to new design problems (Kokotovich and Dorst 2016). Within this research, low-level geometric and process information will be combined to generate abstracted high-level evaluation criteria that can be used to aid designers in creating parts according to certain production objectives.

Abstraction allows designers to overcome the challenges with current optimisation tools, where structurally optimised parts described solely by minimising mass or maximising stiffness can be translated into design representations that allow designers to question the relationship between their designs and the trade-offs that occur within the iron triangle (Figure 2-6). Additionally, evaluating designs at differing abstraction levels allows different stakeholders in the design process to understand design decisions at appropriate levels of detail. Design abstraction runs the risk of oversimplifying the complexities and interactions between evaluation criteria in the design process. It is important that designers can move up and down abstraction levels to maximise the effectiveness of design abstraction in generative design, reducing the possibility of missing emergent design phenomena.

Figure 5-7 displays the abstracted evaluation criteria used in this research. Based on the synthesised geometry, a part can be evaluated against different criteria. The synthesis process converts the design space parameters to a shape that is represented by facet data. This facet data represents the lowest level data for the part, as shown in level 1. This facet data can then be used to describe further details about the part including compliance, support volume, part volume, build height and the build projection area. These can be combined with each other, to form further abstracted criteria of part structural performance, part cost, build time and build plate packing. Finally, these can be combined according to a series of user-defined weights to form the level 3 abstraction criteria of part performance and production quantity.
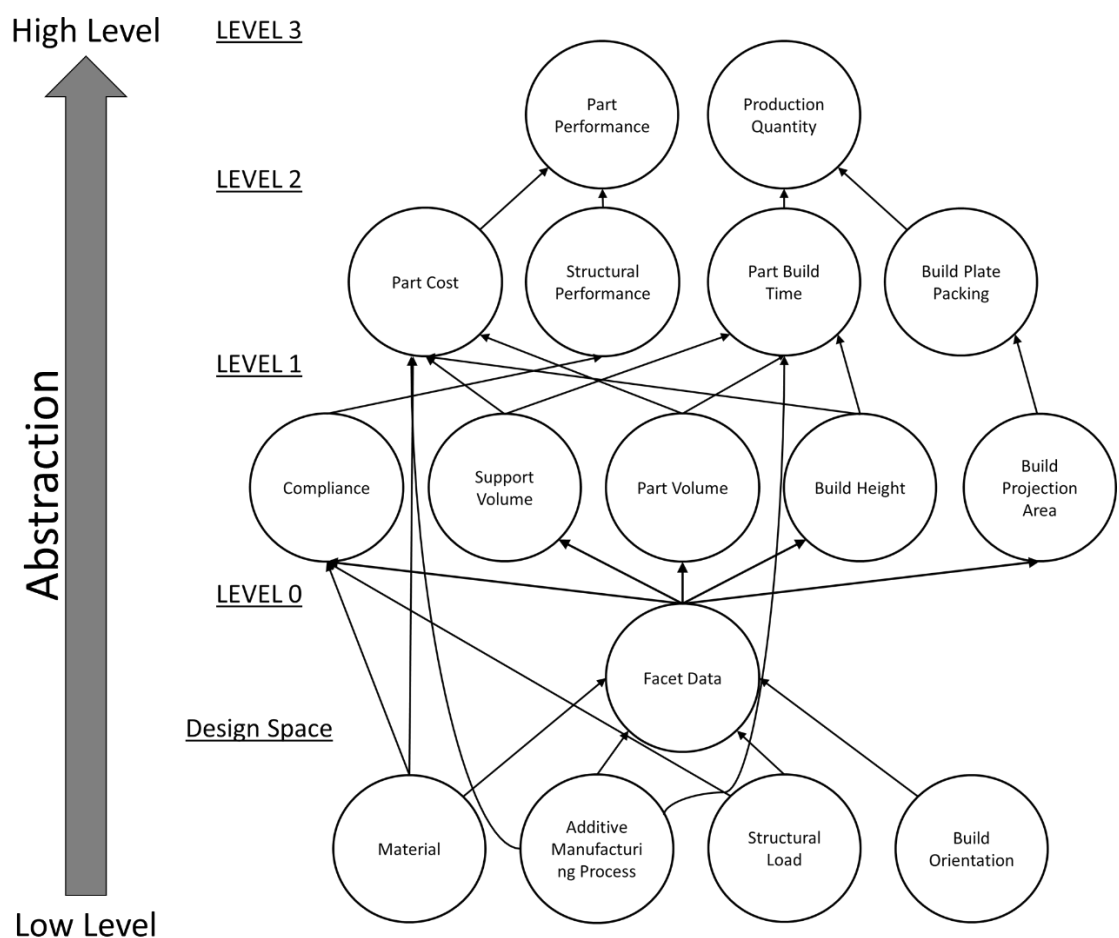
Figure 5-7: Schematic of abstraction criteria used to define high-level design objectives from part evaluation criteria.

The following evaluation criteria will be used to form the weighted model for the objective function. The evaluation criteria are dependent on the mesh data, the process used and the material selected. The evaluation criteria provide information for level 1, level 2 and level 3 abstraction criteria. Figure 5-7 shows how the low-level abstraction criteria map onto the high-level design objectives over multiple layers.

5.3.3.1. Abstraction Criteria – Design Space and Level 0

Material

As the material data is used to drive the structural optimisation process, it is integral to all further abstraction criteria. A Young's modulus is required as an input criterion for the structural optimisation.

Structural Loads

The loading conditions are also required to define the structural optimisation. The structural loads impact on the part geometry and also directly affect the part compliance in the level 1 abstraction criteria.

Build Orientation

The build orientation is required to provide an AM-aware contextualised optimisation. When combined with the manufacturing constraints it is used in order to define the initial ground structure for the structural optimisation as defined in Section 5.3.1. The build orientation is defined as an open variable within the design space and will be altered to generate the solution space.

Additive Manufacturing Process

Based on the combination of AM process, the selected material and machine, the manufacturing constraints can be expressed. These are required to define the initial ground structure and also the conditions for the inequality constraints within the structural optimisation. The overhang constraint is defined as an open variable and will be changed to generate the solution spaces for the two generative design approaches.

Facet Data

The facet data stores properties associated with the location of the coordinates of all mesh vertices and mesh faces and face normals of the synthesised mesh geometries for each solution.

5.3.3.2. Abstraction Criteria – Level 1

Mesh Volume

The mesh volume is calculated using the formula provided in Zhang and Chen (Zhang and Chen 2001).

The total volume of a triangular mesh is given as:

$$V'_{total} = \sum_i V'_i \qquad (4)$$

where;

$$V'_i = \frac{1}{6}(-x_{i3}y_{i2}z_{i1} + x_{i2}y_{i3}z_{i1} + x_{i3}y_{i1}z_{i2} - x_{i1}y_{i3}z_{i2} - x_{i2}y_{i1}z_{i3} + x_{i1}y_{i2}z_{i3}) \qquad (5)$$

Here, i stands for the index of triangles and $(x_{i1}, y_{i1}, z_{i1}), (x_{i2}, y_{i2}, z_{i2})$, and (x_{i3}, y_{i3}, z_{i3}) are coordinates of the vertices of each triangle i, which is derived from the facet data in level 0.

As the triangle mesh is an approximation of the generalised shape, the mesh volume will depend on the resolution of the triangular mesh. Selecting a resolution for the triangular mesh is a compromise between the accuracy of representation and computation. Figure 5-8 shows the effect of changing the mesh resolution on the error caused when calculating the mesh volume for a sphere with radius $r = 20$. The analytical solution, using $V = \frac{4}{3}\pi r^3 = 33510$, is represented with the horizontal red line. The default settings within Three.js, the computational framework used to generate 3D geometry, provide a sphere with 1984 faces, shown as a cross on the graph. This value represents a 0.88% error when compared to the analytical solution. Element resolutions will be selected throughout to ensure that the accuracy error for each primitive remains less than 1%.

Figure 5-8: Effect of mesh volume error caused by the triangulation of geometry.

A projection area is obtained by finding the area of the bounding rectangle (sides $BBox_x$ and $BBox_y$) that is created when the geometry is projected downwards onto the plane that represents the build plate. The build height corresponds to the z-height of the bounding box, termed $BBox_z$. These parameters are further illustrated in Figure 5-9. The bounding box is calculated using the .computeBoundingBox() method in the geometry class of the Three.js framework.

Figure 5-9: Schematic depicting derivation of build projection area and build height.

Support Structure Quantity
Three evaluation criteria are considered, namely the number of overhanging facets, the total length of support arising from each overhanging facet and the total volume of support structure required to cater for each overhanging facet. In the following paragraphs, each criterion is introduced, and the known limitations of each are discussed.

Once the geometry of a given component is represented by a triangular mesh (STL file), identifying facets that overhang beyond a critical angle is trivial. The critical overhanging angle is set at 45° as this is a common guideline for all AM manufacturing processes. It is therefore tempting to use the number of overhanging facets as a metric to measure the optimality of a particular build orientation, which can be calculated using (2).

However, the number of overhanging facets can be misleading in terms of the total volume of support material. The example given in Figure 5-10 shows two faceted surfaces of similar surface area. The triangulation of the square and octagonal surfaces clearly results in a different number of facets, determined by the shape (curvature) of the outer boundary. In reality, both surfaces would require a similar amount of support structure. This two-dimensional example extends into

three-dimensions with both the shape of the boundary, curvature of the surface and the topology.

Figure 5-10: Two polygons of approximately equal area resulting in a differing number of facets after meshing

A second argument against using the number of overhanging faces is the limitation of support structure length. Figure 5-11 displays an example of support structure length dominating the eventual support structure material volume with the leftmost orientation having almost double the support volume as a result of few but very long support columns.

Figure 5-11: An example of the support structure length becoming dominant with respect to the number of overhanging faces

To counteract the length limitation, the cumulative length of the entire support structure network can be used as an evaluation criterion. The geometric centre, v_c, of a triangular facet is identified, and a ray is cast downwards with respect to the build orientation. The length of the support is defined as the length of the ray at the point it makes its first collisions, regardless of whether it intersects another portion of the mesh or the build plate. Summing all such lengths for each overhanging facet gives an estimate of the quantity of support structure (6)

$$Total\ Support\ Length = \sum_{n=0}^{N} d_1 \qquad (6)$$

Although this is an improvement upon the number of overhanging facets, it is still susceptible to inaccuracies brought about facets area (Figure 5-12). Large facets tend to underestimate the support requirement, and smaller facets overestimate.

Figure 5-12: A triangular facet with rays cast downwards with respect to the build orientation. Rays originate at the vertices and centroid of the facet and each collides with a surface below.

It is possible to extend the complexity of the evaluation metric further to consider the volume of the support structure. Although this is a more accurate approach, it is not completely without challenges, and can dramatically increase the level of computation required to acquire a measure of orientation optimality. The calculation of support volume is almost certainly going to require the downward projection of the overhanging facet with respect to the build direction. Herein lies the first complication. Taking the example given in Figure 5-13, it can be seen that the downward projection of a single overhanging facet may collide with multiple surfaces (and multiple facets on each). This introduces an ambiguity in where to place the foundation of the support column, as it could reasonably originate from the build plate or from a lower region of the component geometry. Its origin would be determined by the degree of overlap between the projection and each surface. The interaction of the projection and multiple surfaces also significantly increases the complexity of the calculation of the support volume as the otherwise straightforward polyhedron (of projection) is interrupted.
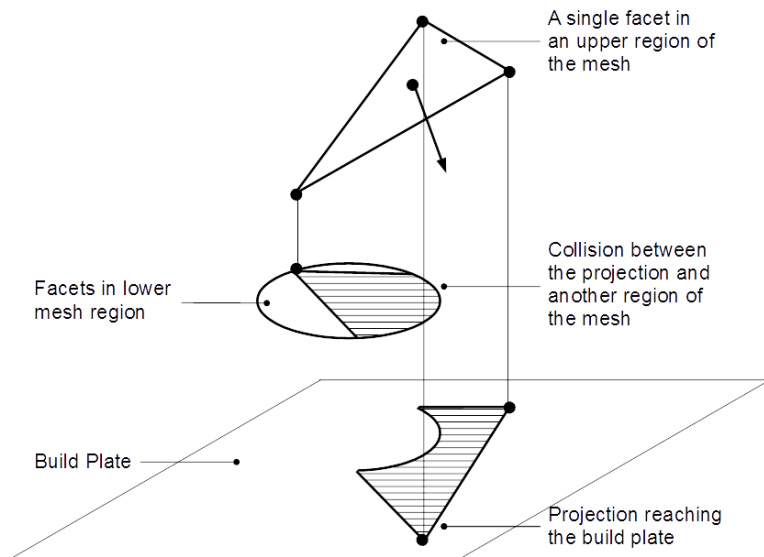
Figure 5-13: A single triangular facet, projected downwards with respect to the build direction. The projection collides partially on another region of the mesh (circular region), and the rest reaches the build plate.

A scheme has been implemented to balance the computational overhead and accuracy of approximation associated with the support volume acquisition function. Considering Figure 5-12 and Figure 5-13, a ray is cast downwards from each vertex of the triangular facet, resulting in lengths d_1, d_2 and d_3. These lengths are set according to the first collision of the ray (mesh or build plate). The area of the triangular facet, A, is also calculated from the vertices. An estimate of the support volume for the part is given by (7). In essence, this is a multiplication of the facet area by the average lengths of the rays.

$$Total\ Support\ Volume = \sum_{n=0}^{N} \frac{(d1 + d2 + d3)A}{3} \qquad (7)$$

It may seem counterintuitive to use the facet area rather than the area of the facet projection onto the xy-plane. However, the errors arising from this simplification are typically small and have the benefit of avoiding the complexities of calculating the true volume, for example, the case in Figure 5-13.

To test the assertion, above, a large number of triangles were randomly generated using the scheme depicted in Figure 5-14. The first and lowest vertex of the triangle is fixed at the origin in a Cartesian coordinate system. Two further facets are then generated on a unit square, spanning $x \in \mathbb{R} : -0.5 \leq x \leq 0.5$ and $y \in \mathbb{R} : 0 \leq y \leq 1$. These two facets are then rotated about the x-axis by the angle, α, which is drawn from a uniform random distribution between the limits of 0 and $\pi/4$ (the critical overhang angle). This process is repeated for 10^8 randomly generated triangles. The percentage error between the area of each triangle's projection onto the xy-plane

and the area of the original triangle is captured along with the aspect ratio of each triangle and its original area.

Figure 5-14: The proposed scheme to randomly generate large numbers of triangles on a variety of overhanging planes, as defined by the angle, α. One vertex of the triangle is always anchored to the origin of the coordinate space.

Figure 5-15 shows the histogram of the percentage errors between the projected and original triangles. The histogram has been normalised according to the relative probability of each bin. Errors less than 5% are dominant, thereby defending the use of the original area of the triangle in the evaluation metric as a comparison between meshes and orientations.
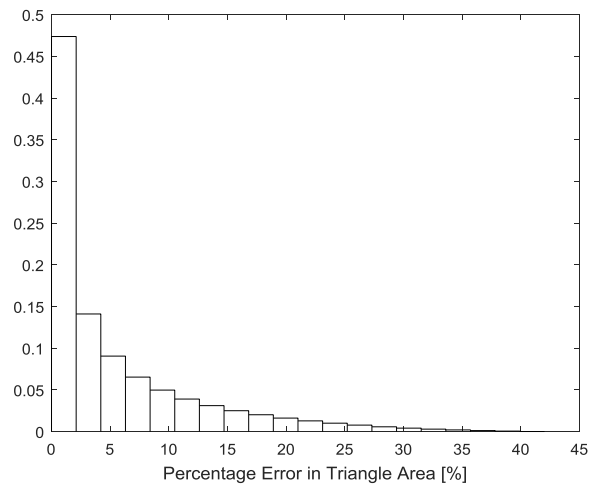
Figure 5-15: Histogram showing the distribution of percentage error between the area of a triangular facet and its projection onto the build plate. The frequency has been scaled according to the probability of each bin.

5.3.3.3. Abstraction Criteria – Level 2
Build Plate Packing and Structural Performance

The build plate packing and structural performance metrics only take in a single factor from the level below. However, the metrics of build plate packing and structural performance are easier to comprehend than build projection area and compliance. In order to refactor these metrics, the build projection area is converted from cost criteria to benefit criteria in build plate packing. This is acceptable as minimising a positive is mathematically equivalent to maximising a negative. Furthermore, this refactoring allows for greater model generalisation as future implementations may integrate further evaluation metrics into the abstraction criteria. Such examples include structural performance containing additional metrics such as Von-mises stress and maximum deflection.

Build Time

The total build time is defined by Baumers et al. (2012) to be a combination of the fixed time consumption for the operation of the machine, termed T_{job}. This includes machine atmosphere generation and machine warm up, and a print time, T_{print} which is comprised of the layer dependent time consumption, related to the part height, in conjunction with the time required to print the required volume.

Individual AM processes will have different methods for calculating T_{print} based on the deposition methods. Within this thesis the model for laser-PBF processes will be used. The model developed by Brika et al. (2017) is used and shown in (8) and (9).

$$T_{print} = \frac{Z_{part}}{L_{thick}} \times T_{recoat} + \left(\frac{V_{part}}{L_{thick}} \times D_{part}\right) + \left(\frac{V_{supp}}{L_{thick}} \times D_{supp}\right) \qquad (8)$$

Where

$$D = SS\left(\emptyset_{laser} - \frac{HS}{2}\right) \qquad (9)$$

In the above equations, Z_{part} is the build height, L_{thick} is the layer thickness, T_{recoat} is the recoater time, V_{part} is the part volume, V_{supp} is the support amount, SS is the scan spacing, $\emptyset_{laser}$ is the laser spot size, and HS is the hatch spacing.

Build Cost

According to Baumers et al. (2012) the cost for an AM build comprises of direct costs associated with the material processing during the build, the energy costs of the build and indirect costs. This indirect costs for a given include, production labour, equipment costs and maintenance (Ruffo et al. 2006). The following cost model (10) is proposed as an estimate of the cost per build:

$$C_{build} = C_{indirect} \times T_{print} + M_{build} \times Price_{material} + E_{build} \times Price_{energy} \qquad (10)$$

Where C_{build} is the build cost, $C_{indirect}$ is the indirect cost, T_{print} is the build time, M_{build} is the build mass and E_{build} is the build energy.

The mass of the part is comprised of the product of part volume, V_{part}, and the density, ρ, of the material being processed as shown in (11).

$$M_{part} = \rho V_{part} \qquad (11)$$

This can be extended to the build mass with the inclusion of the mass of the support structure using a similar method.

$$M_{supp} = \rho V_{supp} \qquad (12)$$

The material cost is given by the product of the build mass and the cost per kg of the material.

5.3.3.4. Abstraction Criteria – Level 3

Exploring high-level abstractions as goals is imperative in delivering a generative design solution, as they tend to be easily comprehensible, desirable objectives that designers seek to optimise. Three alternative theoretical production scenarios will be explored by combining four different level 2 abstraction metrics applied with differing weightings.

Table 5-1 shows three different production personas. Each persona represents a designer/production engineer from an example industry. Each of these industries have different production requirements.

Table 5-1: AM design personas representing industries with differing production requirements.

Persona	A	B	C
Design Objective	High Volume Part Production	High Performance Part Production	Maximum Overall Performance
Example Industry	Consumer goods; Consumer Automotive; Medical/Dentistry	Sporting Applications; High-Value Automotive	AM Design Consultant; Aerospace;
Weightings	Structural Performance= 0 Build Plate Packing = 80 Part Build Speed = 0 Part Build Cost = 20	Structural Performance= 80 Build Plate Packing = 0 Part Build Speed = 20 Part Build Cost = 0	Structural Performance = 25 Build Plate Packing = 25 Part Build Speed = 25 Part Build Cost = 25

Persona A: The first production scenario is a high production quantity scenario. Here the part cost is reduced, and the packing ratio is maximised. With recent advances in AM hardware improving the quality, speed and cost of the process, it is now possible to use AM within a production scenario. Figure 5-16 shows that unit costs of metal parts remain below that of casting up to a production quantity of 100,000 (Desktop Metal 2018), as shown in Figure 5-16. The successful implementation of AM as a production method, in many industries (such as automotive and consumer goods) is dependent on the efficiency of the design process. Highly optimised nesting of parts, as well as minimising the support structure requirement will reduce the overall cost of material use and the post-processing time required before shipping the final part. A typical industry might include medical/dentistry, where it is common for industries to densely pack build plates with customised parts such as dental crowns (Renishaw plc. 2017). This level of personalisation is only commercially viable using a combination of AM processing and maximisation of build plate packing. Based on this, the weightings are selected to be 80% part throughput, which is defined to be a benefit function that is to be maximised and 20% part cost that is a cost function and should be minimised. The expected output from this user weighting preference is likely to be a low build projection area and low support requirement, with the relaxed objective within the iron triangle being a lower structural performance.

Figure 5-16: Cost versus production quantity trends for Desktop Metal production additive manufacturing process compared with casting and laser powder bed fusion technologies (Adapted from Desktop Metal 2018)*.*

Persona B aims to design the overall highest performing part that is achieved when the part build speed and structural performance are prioritised. The design flexibility provided by multi-scale design methods specific to AM processes, coupled with the ability to process high-performing materials (e.g. Ti-6Al-4V and Inconel 718) has meant that AM has quickly evolved into a technology that is capable of producing parts with the same structural performance as other manufacturing processes with a much lower mass. In some cases as much as 80% lower (Carter et al. 2014). Many industries can benefit from this advantage and can justify compromising on cost in order to maximise the structural performance of the part. Examples include high-performance sporting applications, hyper/supercar parts (BMW Group 2018).

Finally, **Persona C,** is the control case and uses equal weightings. This scenario represents the compromise condition where the designer may not fully understand the problem or the impact of the design variables on the final design. It is likely that this condition, will provide a spectrum of different results, with high-performing solutions for each of the evaluation criteria for the designer to evaluate. This scenario may be representative of typical problems faced by the aerospace industries, where a reduction in part mass can lead to significant savings in the total cost of the aircraft. It is estimated that 9-17% of total typical aircraft mass may be replaced by AM components to provide a 6.4% reduction in fuel consumption (Huang et al. 2016). However, it is still necessary to produce the required volume of parts at a suitable cost if AM is going to be used in production aerospace (Uriondo et al. 2015) with large aerospace companies producing approximately 700 planes per year (Airbus Group; 2018).

Viewing this from a different perspective, Figure 5-17 demonstrates the relationship between the manufacturing personas and the iron triangle. The bottom left of the image shows regions of non-optimised designs. Designers do not want to create solutions within this region as they are outperformed by other designs. Conversely, designers wish to develop designs in the top right corner. This is infeasible due to the constraint of the iron triangle. Persona A wishes to design parts with low cost and high throughput (speed) trading off against quality. Persona B aims to maximise quality and part build speed with the relaxed constraint being part cost. Persona C desires to understand the underlying trade-offs in the design problem by equalising all values.
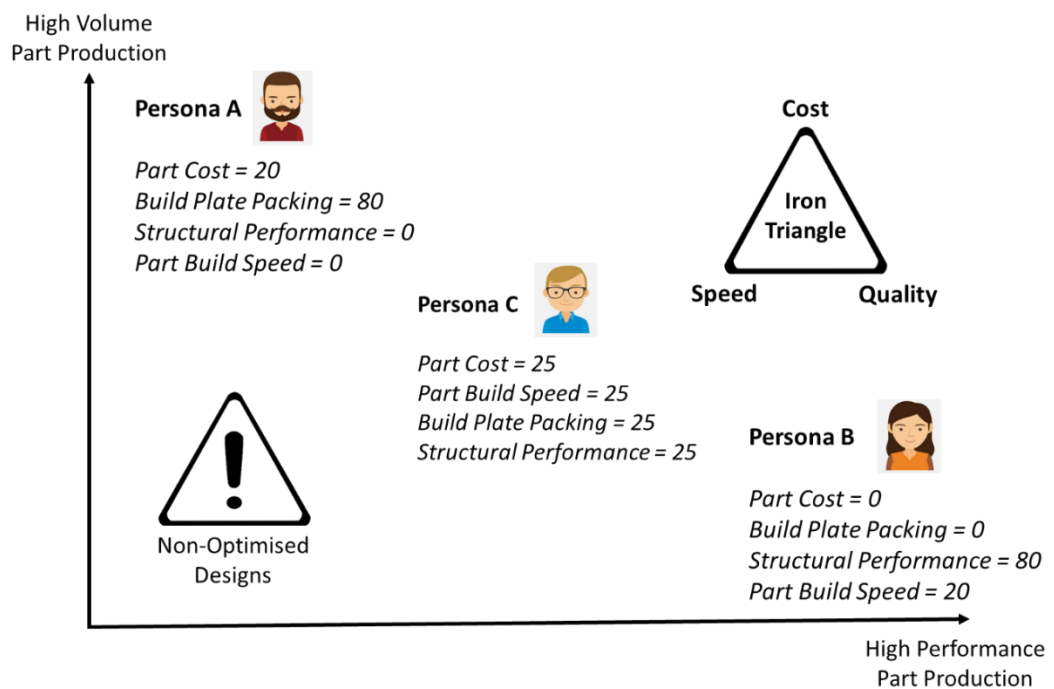
Figure 5-17: Schematic of different production scenarios generated from weighted level 2 abstraction criteria and achievable design limitations created by the iron triangle.

5.3.3.5. User Preferences

As described in the research scope in Chapter 4, this research is solely concerned with quantitative user objectives. The user inputs are acquired through an interactive user input module, shown in Figure 5-18, in which the user requirements are divided into percentage blocks and summed to 100%. The selections are made from benefit criteria (those values that are to be maximised) and cost criteria (minimised). The inputs are used to derive the objective functions in both the data-driven and goal-driven approaches to generative design.

Figure 5-18: User input selection box used to define objective function weightings and whether the criteria should be maximised (benefit) or minimised (cost).

5.4. Research Methods

The following research methods are used to implement the research methodology. The methods for data-driven and goal-driven generative are two possible realisations derived from the generative design CAD framework in Chapter 3.

5.4.1. Data-Driven Methods for Generative Design

The aim of data-driven generative design methods is to develop and implement computational solutions that locate the highest performing concepts within a previously generated AM solution space with respect to the different objectives. Often large multi-dimensional solution spaces are difficult to navigate due to difficulties in human comprehension of multi-dimensional data (Halford et al. 2005). Therefore data-driven methods aim to make sense of these datasets and feedback to the user.

Data-driven methods can be described as decision-support tools that aid the user in locating the highest-performing design solutions within a generated solution space given a set of user-defined goals. Data-driven methods can be used on existing datasets of solutions, or in the case of this research, in conjunction with a generative shape synthesis method. Decision support tools are particularly useful in making sense of complex multi-dimensional solution outputs. This allows users to rank solutions against conflicting goals and overcome the time constraints associated with examining and evaluating large numbers of solutions.

Many MCDA techniques have been developed to support decision making in engineering (Greco et al. 2016). The aim of MCDA techniques is to evaluate multiple conflicting criteria in decision making. This is imperative in design where design concept evaluation can be a complex multi-criteria decision-making problem that can combine many factors from customer needs and constraints to enterprise resources (Geng et al. 2010).

The following list includes examples of common MCDA techniques:

- Weighted sum model (WSM) (Fishburn 1967)
- Weighted product model (WPM) (Bridgman 1922)
- Analytical hierarchy programming (AHP) (Saaty 1990)
- VIKOR (Opricovic and Tzeng 2004)
- TOPSIS (Yoon and Hwang 1995)

The listed MCDA approaches share some common elements. For example, that values for alternatives are assigned for each of the criteria and then multiplied by corresponding weights and finally combined to produce a total score. However, the approaches differ in the details of how criteria values are assigned and combined. Due to the relevant strengths and weaknesses associated with the mathematical differences in the methods, typically researchers will select one of the approaches to be the most appropriate for a given problem.

The TOPSIS method (Technique for Order of Preference by Similarity to Ideal Solution) seeks to maximise the benefit criteria and minimise the cost criteria to find solutions closest to the ideal solution. Advantages of TOPSIS include being easy to implement, both positive and negative ideal solutions can be used and it can contain fuzzy numbers to deal with uncertainty (Kabir and Sumi 2012). Furthermore, TOPSIS returns a cardinal ranking of solutions making it easy for end-users to locate the top n performing solutions.

Researchers have compared different MCDA techniques (Caterino et al. 2009) and apart from the WSM and WPM methods that are more appropriate for problems involving variables with similar dimensions and criteria that require maximising or minimising, all other methods tend to be suitable for data-driven generative design problems. TOPSIS and VIKOR methods are typically the most flexible as they permit the use of variables with different units and criteria of different types.

The TOPSIS algorithm used within this research is described in the following steps.

i) Establish a decision matrix consisting of m alternatives and n criteria, the intersection of each alternative and criteria is given as x_{ij}, giving the matrix $(x_{ij})_{m \times n}$. The m alternatives correspond to each design solution in the solution space.

ii) The matrix $(x_{ij})_{m \times n}$ is then normalised to form the matrix $(r_{ij})_{m \times n}$

$$r_{ij} = \frac{x_{ij}}{\sqrt{\sum_{i=1}^{m} x_{ij}^2}}, \qquad i = 1,2,\dots,m, j = 1,2,\dots,n \tag{13}$$

iii) Calculate the weighted normalised decision matrix

$$t_{ij} = r_{ij} \cdot w_j, \qquad i = 1,2,\dots,m, j = 1,2,\dots,n \tag{14}$$

$$w_j = \frac{W_j}{\sum_{j=1}^{n} W_j}, \qquad j = 1,2,\dots,n \tag{15}$$

The criteria are weighted on a user selection input panel by allocating a percentage of resources to any given performance criteria. As described in Section 5.3.3.5, to prevent the user from attempting to maximise every performance criteria the user is allocated 100% of resources to divide between each of the performance criteria as they wish, such that

$$\sum_{j=1}^{n} W_j \leq 100 \tag{16}$$

Where W_j is the value for each of the n performance criterion.

iv) Determine the worst alternative (A_w) and the best alternative (A_b)

$$A_w = \{\langle \max(t_{ij} | i = 1,2,\dots,m \mid j \in J_-), \langle \min(t_{ij} | i = 1,2,\dots,m | j \in J_+) \}$$
$$\equiv \{t_{wj} | j = 1,2,\dots,n\},$$
$$A_b = \{\langle \min(t_{ij} | i = 1,2,\dots,m \mid j \in J_-), \langle \max(t_{ij} | i = 1,2,\dots,m | j \in J_+) \}$$
$$\equiv \{t_{bj} | j = 1,2,\dots,n\}, \tag{17}$$

Where,

$J_+ = \{j = 1, 2, \dots, n \mid j \}$, where the criteria have a positive impact (i.e benefit)

$J_- = \{j = 1, 2, \dots, n \mid j \}$, where the criteria have a negative impact (i.e cost)

v) Calculate the Euclidean distance between the alternative i and the worst condition A_w

$$d_{iw} = \left(\sum_{j=1}^{n} (t_{ij} - t_{wj})^2 \right)^{\frac{1}{2}}, i = 1, 2, \dots, m \qquad (18)$$

and the distance between the alternative i and the worst condition A_b

$$d_{ib} = \left(\sum_{j=1}^{n} (t_{ij} - t_{bj})^2 \right)^{\frac{1}{2}}, i = 1, 2, \dots, m \qquad (19)$$

Where d_{iw} and d_{ib} are the L2-norm distances for the target alternative i to the worst and best conditions respectively.

vi) Calculate the distance to the ideal condition

$$s_{iw} = \frac{d_{iw}}{(d_{iw} + d_{ib})}, 0 \leq s_{iw} \leq 1, i = 1, 2, \dots, m \qquad (20)$$

vii) Rank the alternatives according to $s_{iw} = (i = 1, 2., \dots, m)$

viii) Visualise the top n-performing solutions within the data visualisation dashboard outlined in Chapter 6.

The data-driven research method can be generalised into three distinct phases. 1) DEFINE: Defining the correct conditions for a specific structural optimisation problem, 2) GENERATE: generating a solution space of feasible design solutions and 3) LOCATE: searching this solution space to find the preferred solutions to a given set of objectives.

Within the DEFINE phase, three distinct elements of information are expressed. The first action is to define the structural constraints. This includes the positions, magnitude and direction of the loading conditions, boundary conditions for the optimisation, and any material void spaces in which material is not permitted. Secondly, performance goals are defined. These are the level 2 abstraction criteria to be maximised or minimised, including, build time, build plate packing, cost per-part or structural performance. Finally, the independent variables that are to be explored as well as the sampling frequency for each variable are defined. Examples include specific information related to the AM process, such as the build angle of the part or whether a manufacturing constraint will be applied. Next, in the GENERATE phase, the number of samples for each parameter are selected to create the solution space. Finally, in the LOCATE phase, the quantitative evaluation criteria calculated during the GENERATE phase are mapped onto the performance criteria from the DEFINE stage to identify the most appropriate designs in a solution space.

By iterating through the design parameters, a space of structurally optimised mesh solutions can be generated. The output mesh from the optimisation is subsequently scrutinised using a number of selected evaluation criteria to generate a high-dimensional representation of the final solution performance. This mapping from 3D to 'nD' is key to generative design as parts can be represented by their characteristics and behaviour rather than their CAD geometry. The criteria selected for this research are specific to DfAM and include the part compliance, the build height, width and angle, and the amount of support structure required to build the part successfully. However, there is potential for solutions to be evaluated and represented by many more evaluation criteria if required.

Finally, a mapping must be created that links the output evaluation criteria to the input performance goals. This leads to a performance-centric representation for each design solution. The solution space can then be searched using a MCDA technique that ranks the solutions for their ability to fulfil the requirements of the specified design problem, e.g. maximise structural performance or part throughput. The solution space is presented to the user using a series of data visualisations, identifying the best results in the solution space and providing further insight into the design problem.

This experimental procedure is presented in Figure 5-19.

| DEFINE | GENERATE | LOCATE |
| Input Conditions | Solution Space of Feasible Solutions | Top Solutions within Solution Space |

Define loading points and structural constraints

Define performance goals and objectives

Define independent variables and number of alternatives for design space exploration

For each alternative in variable

Generate a topology optimised design solution

Create a high-dimensional representation of solution by evaluating against many criteria.

Collate solutions into single solution space

Create mapping between evaluation criteria and performance goals

Search solution space and find best performing solutions

Highlight top performing solutions in data visualisation

Figure 5-19: Experimental procedure for data-driven generative design for additive manufacturing.

5.4.2. Bayesian Build Orientation Optimisation in Additive Manufacturing

Due to the inefficient and expensive nature of generating solution spaces of optimised parts, there is a requirement to efficiently generate and explore AM-specific design spaces. Chapter 7 compares the use of grid search, random search and surrogate models for their suitability to be used for locating the optimal build orientation for AM parts by minimising the total support structure requirements.

A surrogate model describes the relationship between inputs (input parameters to a model) and outputs (the performance measure of the simulation). Training an accurate surrogate model requires an adequate number of input and output points. The most economical way of constructing a surrogate model is to adaptively select the next set of parameters based on the responses from the previous trials (Wang et al. 2014).

Many surrogate optimisation models exist that are able to perform the function approximation; these are all commonly-used machine learning models. An overview of the most common surrogate models is outlined as follows:

- Polynomial regression, sometimes referred to as non-linear regression aims to minimise the squared error between a parametric model and training data by adjusting the numerical weights within the model.

- Random forests (Breiman 2001) are an ensemble machine learning method comprising multiple decision trees in order to perform either classification or regression. The advantage of random forests is that the algorithm is very stable when new data is added. It is also unlikely that the data will impact all of the trees. Additionally, random forests can

deal with both categorical and numerical features. One of the disadvantages of random forests is that they do not train well on small datasets. Therefore their reliability may be questionable for certain surrogate models.

- Support vector regression is an extension of support vector machines that seek to fit the error from a theoretical hyperplane, which represents the learnt function, and the data within a certain threshold.

- Artificial neural networks work by altering the weights within the layered network, it is possible to approximate the function given by the training data.

- Gaussian process regression (sometimes termed Kriging) uses a covariance function (kernel) to estimate the value for an unseen point based on training data. Gaussian Processes not only return the estimate for the unseen point but crucially also the uncertainty of the estimate, which can be useful when modelling ill-defined problems.

It has been shown the surrogate modelling using Gaussian processes consistently produces better results than other modelling techniques (Wang et al. 2014). As such Gaussian processes will be used to form the surrogate optimisation model within this research. Bayesian optimisation is a method of adaptive surrogate modelling, which typically uses a Gaussian process to form a surrogate model. This model is then updated using an acquisition function. It provides a strategy to find the extrema of objective functions that are expensive to evaluate. Bayesian optimisation techniques are particularly efficient in terms of the number of function evaluations required, which is attributed to the ability to incorporate prior belief about the problem to help direct the sampling (Brochu et al. 2010; Shahriari et al. 2016).

When attempting to locate the optimal build orientation, the target is to minimise the total amount of support structure required to support the overhanging faces on the part. Three different proxy metrics of the total support requirement are compared for their efficiency and accuracy, in order to efficiently evaluate the best possible part orientation.

To determine the orientation of the part, a three-step procedure is implemented. Firstly, the part is rotated by an angle, θ, about the global x-axis. Then, it is rotated by an angle, φ, about the global y-axis, and finally, the part height is readjusted by translating in z. The final translation in the z-direction is such that the lowest facet on the triangular mesh is aligned with the height of the build plate. In this study, θ and φ belong to the set, $\theta, \varphi \in \mathbb{Z}: \theta, \varphi \in [0, 359]$.

The following list describes the research method enacted within Chapter 7:

1. For two component designs, perform a grid search for each of the support requirement evaluation criteria, namely number of overhanging faces, support length and support volume at 15°, 20°, and 30° build angle increments. Selecting these increments will test the assumption that greater numbers of iterations will always generate better solutions and also both even and odd angle increments, whilst remaining computationally feasible.

2. Using only the first test part, run Bayesian optimisation with the Matèrn kernel and Expected Improvement (EI) acquisition function for 15, 25, 35, 50 and 100 iterations, repeating each test 30 times.

3. Using only the first test part, run random search for 15, 25, 35, 50 and 100 iterations, repeating each test 30 times.

4. Compare the grid search results to the results produced from both random search and Bayesian optimisation to determine the following.

   a. Compare the minimum support structure values obtained by each of the methods

   b. Perform the Mann-Whitney U test to determine whether differences between the results of the support structure calculation methods are statistically significantly.

5. Based on the results from 2 & 3, using the first test part, examine the effect of changing the covariance functions for n iterations depending on the balance on computational time and the performance of the Bayesian optimisation algorithm, repeating the test 30 times.

6. Compare the results of Bayesian optimisation and random search at 35 iterations

   a. Compare the best values obtained by each of the support structure calculation methods.

   b. Perform the Mann-Whitney U test to determine whether results of different methods are significantly different to one another.

7. Repeat steps 5+6 for the second test part.

A detailed explanation of the test parts, evaluation criteria and evaluation methods will be explained in the following sections.

5.4.2.1. Test Cases

As part of this experiment, two test parts have been selected to evaluate the performance of different support structure estimation metrics and optimisation approaches. Each test part was selected to be representative of engineering parts that arise from an industrial AM pipeline. In

particular, the geometries have been optimised for the manufacturing process and consist of a representative number of facets in the STL models. Furthermore, these test parts were generated using continuum based topology optimisation methods, therefore showing the ability for the test methods to generalise further than just ground structure-based optimisation techniques. Crucially, both parts are open-access, which will permit future comparisons with this study. The test parts were selected from the GrabCAD (2017) repository, a freely available database of engineering parts. Both test parts were set as community challenges with the aim of redesigning the part for AM.

The first selected part, shown in Figure 5-20, is the aeroplane bearing bracket challenge created by Alcoa. This was set as a community challenge in 2016 (GrabCAD 2016; Kurniawan 2013), and the resulting component has 14,172 faces in the STL file.

Figure 5-20: Front and side views of the winning solution to the Alcoa bracket challenge, hosted by GrabCad

The second test part, shown in Figure 5-21, is the winner of the GE jet engine bracket challenge (GrabCAD 2013; Nikol 2016). As the GrabCAD repository only contains a STEP file for this part, an STL file containing 45,432 faces was created by converting the file using Autodesk Fusion 360 with the default STL parameter settings.

Figure 5-21: Front and side views of the winning solution to the GE bracket challenge, hosted by GrabCad

5.4.2.2. Grid Search

Grid search requires that a set of values is chosen for each variable ($L^{(1)} \dots L^{(k)}$). Using grid search, the set of trials is formed by assembling every possible combination of values, which gives the total number of trials as: $S = \prod_{k=1}^{K} |L^{(k)}|$. The product over K sets makes grid search suffer from

the curse of dimensionality as the number of searches exponentially grows with the number of parameters (Bellman 2015).

5.4.2.3. Random Search

One improvement over the grid search is to sample the points within the design space randomly. Random search can be more exploratory than a grid search, as detailed in Bergstra and Bengio (2012). For each iteration, a random number is sampled from a uniform distribution to represent the rotation angle in the x- and y-axis, respectively. Figure 5-22 shows that if one of the dimensions has a greater impact on the objective function than another (horizontal axis), then grid search has the potential to miss high-performing solutions. Furthermore, by sampling at random, it is possible to explore a much greater distance than with grid search as each point searches new regions in all axes. In contrast, random search explores three times as many points in the horizontal axis than grid search, thereby increasing the likelihood of locating high-performing solutions.

Figure 5-22: Comparison of grid search and random search (Bergstra and Bengio 2012)*.*

5.4.2.4. Bayesian Optimisation

Bayes theorem, a representation of which is shown in (21), is the key to optimising an objective function. Bayes theorem states that the posterior probability of a model, A, given evidence, B, is proportional to the likelihood of A given B, multiplied by the prior probability of A. The prior represents the belief relating to the space of possible objective functions. Although this is unknown, it is possible to make assumptions about its nature, which makes some solutions more feasible than others.

$$P(A|B) \propto P(B|A)P(A) \qquad\qquad (21)$$

A Gaussian process can be used as a surrogate optimisation method. It is an extension of the multivariate Gaussian distribution to an infinite dimension stochastic process for which any finite combination of dimensions will be a Gaussian distribution. A Gaussian process is a distribution over functions that is specified by its mean function, μ, and covariance function, k, associated with two input vectors, $\boldsymbol{x}$ and $\boldsymbol{x}'$ as described in (22):

$$f(\boldsymbol{x}) \sim GP(\mu(\boldsymbol{x}), k(\boldsymbol{x}, \boldsymbol{x}')) \qquad (22)$$

The covariance function in the Gaussian process determines the smoothness of the samples that are drawn from it. The choice of the covariance function is dependent on the underlying data that is being modelled. In this research, two common covariance functions will be evaluated: the squared exponential kernel and the Matèrn kernel.

The squared exponential function is described in (23):

$$k(\boldsymbol{x}, \boldsymbol{x}') = \exp\left(-\frac{1}{2\theta^2}|\boldsymbol{x} - \boldsymbol{x}'|^2\right) \qquad (23)$$

where θ is a hyperparameter that controls the width of the kernel.

The Matèrn is another common kernel that incorporates a smoothness parameter, ς, to permit greater flexibility when modelling functions and is shown in (24).

$$k(\boldsymbol{x}, \boldsymbol{x}') = \frac{1}{2^{\varsigma-1}\Gamma(\varsigma)}\left(\frac{2\sqrt{\varsigma}}{\kappa}\|\boldsymbol{x} - \boldsymbol{x}'\|\right)^{\varsigma} \mathrm{H}_{\varsigma}\left(\frac{2\sqrt{\varsigma}}{\kappa}\|\boldsymbol{x} - \boldsymbol{x}'\|\right) \qquad (24)$$

where $\Gamma(\cdot)$ and $H\varsigma(\cdot)$ are the Gamma function and the Bessel function of order ς. Note that as $\varsigma \to \infty$, the Matèrn kernel reduces to the squared exponential kernel, and when $\varsigma = 0.5$, it reduces to the unsquared exponential kernel. As with the squared exponential function, a length factor, κ, is also included.

The aim of the GP is to estimate the value of a point based on previous training data. Given a set of inputs $\boldsymbol{x} \in \mathbb{R}^{n}$ and corresponding outputs $\boldsymbol{f} \in \mathbb{R}^{n}$, the purpose of the GP is to estimate the outputs $\boldsymbol{f}_*$ for a set of new test points $\boldsymbol{x}_*$.

The covariance matrix, $\boldsymbol{K}$, is given by all possible combinations of the samples such that

$$K = \begin{bmatrix} k(x_1, x_1) & k(x_1, x_2) & \cdots & k(x_1, x_n) \\ k(x_2, x_1) & k(x_2, x_2) & \cdots & k(x_2, x_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(x_n, x_1) & k(x_n, x_2) & \cdots & k(x_n, x_n) \end{bmatrix} \qquad (25)$$

The covariances for the unknown sample x_* are written as $k_* = [k(x_*, x_1), k(x_*, x_2), \ldots, k(x_*, x_n)^T$, and $k_{**} = k(x_*, x_*)$.

Combining the observed and unobserved data into the GP, a multivariate normal is formed with the following mean and covariance structure:

$$\begin{pmatrix} f \\ f_* \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_x \\ \mu_{x*} \end{pmatrix}, \begin{pmatrix} K & k_{x*}^T \\ k_* & k_{**} \end{pmatrix} \right) \qquad (26)$$

Forming the conditional in order to form an estimate for f_*, the expression becomes

$$p(f_* | x_*, x, f) = \mathcal{N}(\mu_{x*} + k_* K^{-1}(f - \mu_{x*}), k_{**} - k_* K^{-1} k_*^T) \qquad (27)$$

This means that for value within the specified problem domain it is possible to predict the mean and variance at that point. The full derivation for the following equations (28) and (29) can be found in Murphy (2012).

$$\mu(x_*) = \mu_{x*} + k_* K^{-1}(f - m_{x*}) \qquad (28)$$

and,

$$\sigma\big(m(x_*)\big) = k_{**} - k_* K^{-1} k_*^T \qquad (29)$$

In Bayesian optimisation, observations are accumulated such that $\mathcal{D}_{1:t} = \{x_{1:t}, y_{1:t}\}$, a prior distribution $P(f)$, is combined with the likelihood function $P(\mathcal{D}_{1:t}|f)$ to produce the posterior distribution: $P(f|\mathcal{D}_{1:t}) \propto P(\mathcal{D}_{1:t}|f)P(f)$. This posterior captures the updated beliefs about the unknown function. The algorithm for implementing Bayesian optimisation is shown below.

Algorithm 1: Bayesian Optimisation

```
FOR t = 1,2,… do
     Find x_t by optimising the acquisition function over the GP: x_t = argmax u(x|D_{1:t-1})
                                                                      x
     Sample the objective function: y_t = f(x_t)
     Augment the data D_{1:t} = {D_{1:t-1}, (x_t, y_t)} and update the GP.
ENDFOR
```

The role of the acquisition function is to guide the search for the optimum. A high value of the acquisition function typically denotes a predicted high value of the objective function. Thus, maximising the objective function provides the ability to find the next set of parameters to be used to evaluate the objective function. The maximum value of the acquisition function is found by minimising the negative of the function using the L-BFGS-B method (Byrd et al. 1995). The two acquisition functions tested in this article are the upper confidence bounds (UCB) shown in (30),

$$UCB = \mu(\boldsymbol{x}) + \kappa\sigma(\boldsymbol{x}) \qquad\qquad (30)$$

where $\mu(\boldsymbol{x})$ is the mean value at $\boldsymbol{x}$, κ is a variable used to tune the acquisition function and $\sigma(\boldsymbol{x})$ is the variance at $\boldsymbol{x}$ and the EI function shown in *(31)*.

$$EI = \frac{\mu(\boldsymbol{x}) - f(\boldsymbol{x}^+) - \xi}{\sigma(\boldsymbol{x})} \qquad\qquad (31)$$

Where $\mu(\boldsymbol{x})$ is the mean value at $\boldsymbol{x}$, $f(\boldsymbol{x}^+)$ is the best observed value, ξ is a variable used to vary the acquisition function to favour the mean or variance and $\sigma(\boldsymbol{x})$ is the variance at $\boldsymbol{x}$.

As the acquisition function is based on the expected improvement of the Gaussian process as well as the uncertainty it is possible to balance the trade-off of exploiting and exploring. By increasing the values of κ and ξ, the acquisition function favours exploring higher areas of variance, therefore exploring a greater region of the design space. Whereas, low values provide a bias toward exploiting high values of the predicted mean. The values of κ and ξ are tuned by observing a number of tests and adjusting appropriately until a suitable balance between exploration and exploitation is determined.

5.4.3. Goal-Driven Methods for Generative Design

The use of Bayesian optimisation as a surrogate model will be explored as a method to efficiently locate the highest performing solutions for a given problem definition for each of the manufacturing production scenarios outlined in Section 5.3.3.4

The goal-driven generative methods used within this research apply adaptive surrogate modelling methods to efficiently explore AM design spaces in order to locate high-performing solutions

based on user-defined goals. The experimental procedure applied when performing the goal-driven generative design experiments can be seen in Figure 5-23.
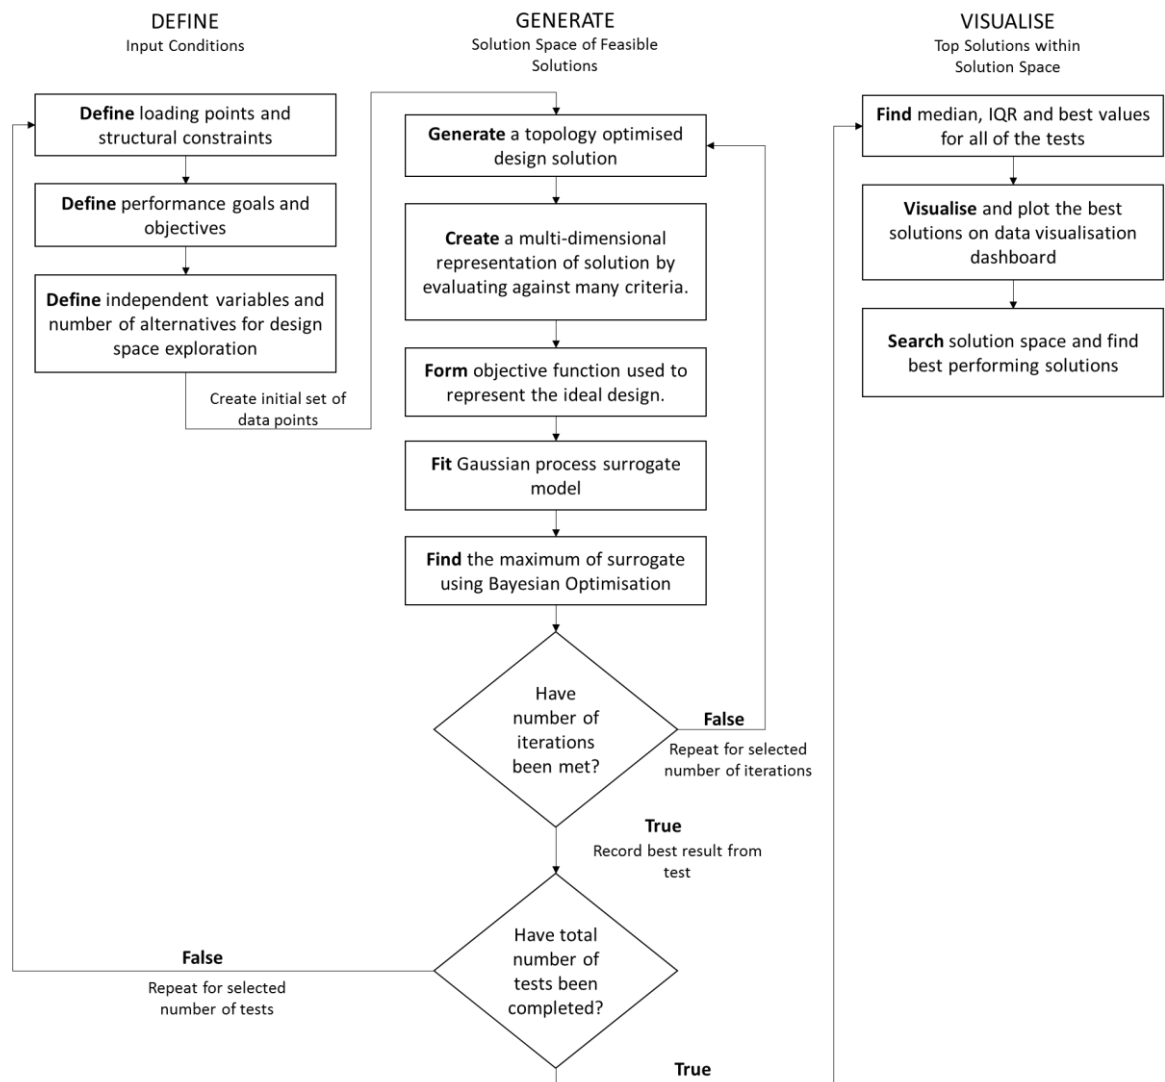
Figure 5-23: Experimental procedure employed for goal-driven approach to generative design in Chapter 8.

As with the data-driven approach, a three-stage method is created. As the goal-driven generative design approach undertaken within this research is a stochastic process, the tests are repeated 21 times, in order to gain a robust median value to assess the quality of the method. The decision to perform repetitions is based on the balance between robustness of the Mann-Whitney U test, and the computational cost of running multiple tests.
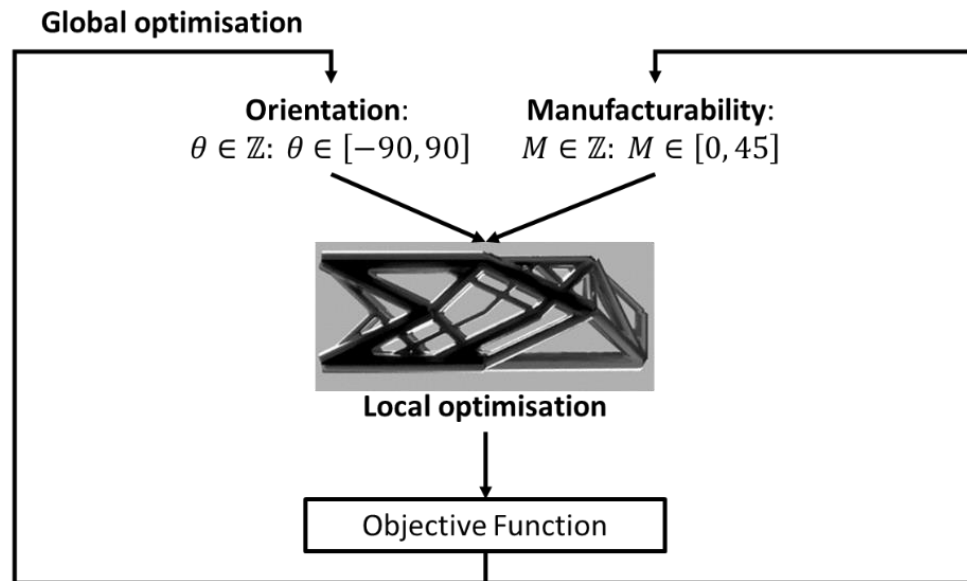
Figure 5-24: Global and local optimisation strategies used to create goal-driven generative design for additive manufacturing solutions.

Firstly, within the DEFINE stage the loading points and structural constraints, performance goals and independent variables are defined. The conditions are derived from the design problem described in Section 5.3, but can be extensible to other problems if required.

Next, within the GENERATE stage, a solution space is created by first creating an initial set of data points used to form the surrogate model. This stage consists of two different optimisation approaches, as shown in Figure 5-24. The first optimisation stage taken in the build orientation and manufacturing constraint as input parameters and performs a structural optimisation with the output being evaluated against the level 0, level 1 and level 2 abstraction criteria.

The topology optimisation is performed by rotating the part and randomly setting a manufacturability constraint between $0° - 45°$ that defines the part to be some region between non-manufacturable without support structure and fully manufacturable without support. Unlike the data-driven method, this must be continuous to form the surrogate model. The multi-dimensional representation of the design is then formed by evaluating the design against multiple evaluation criteria e.g. level 2 abstraction criteria.

The objective function will be formed of weighted level 2 abstraction criteria. The goal of the optimisation is to minimise the function, f, subject to the maximisation or minimisation of this objective function. The treatment of the evaluation criteria is dependent on whether the user desires the criteria to be maximised (benefit) or minimised (cost). As the evaluation criteria each have different units and orders of magnitude, the function must be expressed in a normalised form. This can be achieved by dividing each metric by its maximum value for objectives that are to

be minimised or by dividing the minimum value by the sum of the minimum and the evaluated metric for objectives that are to be maximised as shown in (32):

$$minimise\ f = w_1 \times \frac{f_1}{\max f_1} + w_2 \times \frac{\min f_2}{\min f_2 + f_2} + \cdots + w_m \times \frac{\min f_m}{\min f_m + f_m} + w_n \times \frac{f_n}{\max f_n} \qquad (32)$$

Once the initial points have been created, the Gaussian process surrogate model can be fitted to the data and the maximum value can be located and combined with the acquisition function to select the next points for evaluation. For n iterations, the model is adaptively updated and after ten iterations the best observed result is recorded. This optimisation is referred to as the global optimisation as this is what controls the set of parameters used to synthesise the final generative part. Once all of the tests have been completed, the median and interquartile range (IQR) values value for each of the tests is calculated and the data is plotted for comparison with the data-driven generative design method described in Section 5.4.1.

5.5. Comparison

By taking the best values from the goal-driven method, the number of iterations required to locate high-performing areas of the solution space can be compared. Furthermore, it is possible to assess the flexibility of each of the algorithms when dealing with new information or increased numbers of design parameters. An assessment will be made on the suitability of each approach for its ability to achieve optimised design results for various design problems.

The following approach will be used to compare the two generative design methods. Firstly, all of the design instances will be evaluated against the objective function for each of the three manufacturing personas and normalised against the maximum value in the dataset. Secondly, these will be ranked from smallest to largest using the default sort function in Python. Thirdly, the minimum value will be subtracted from each of the values to indicate the distance from the best solution within the dataset. The implication of this approach is that the highest performing data-driven design instance will have a value of 0. The approach can be seen in Table 5-2.

Table 5-2: Example values for calculating distances between data-driven and goal-driven generative design for additive manufacturing methods.

Design Instance	Normalised Objective Function	Design Instance	Ranked Normalised Objective Function	Ranked Normalised Objective Function - min (Ranked Normalised Objective Function) Distance
1	0.245	1	0.245	0
2	0.876	3	0.423	0.178
3	0.423	4	0.674	0.429
4	0.674	2	0.876	0.631
…	…	…	…	…
n	x	n	y	z

The first quartile, median and third quartile values can then be calculated for 21 iterations of the goal-driven method. The minimum value for the data-driven dataset can then be subtracted from these values. This value can be compared with the subtracted ranked normalised objective function list to give a value of its position within the data-driven dataset. This will provide a value of the goal-driven method to generate solutions within the top n % of solutions within the data-driven generative design method. An example of this can be viewed in Table 5-3, where the goal-driven value is labelled GD.

Table 5-3: Example comparison of goal-driven design with data-driven design instances dataset.

Design Instance	Normalised Objective Function	Design Instance	Ranked Normalised Objective Function	Ranked Normalised Objective Function - min (Ranked Normalised Objective Function) Distance
1	0.245	1	0.245	0
2	0.876	GD	0.316	0.071
3	0.423	3	0.423	0.178
4	0.674	4	0.674	0.429
…	…	2	0.876	0.631
n	x	…	…	…
GD	0.316	x	y	z

To answer the research aim, it is necessary to determine the distance for which goal-driven methods are sufficiently similar to data-driven methods to define them as having comparable performance. Within this research, goal-driven generative design solutions will be determined to have comparable performance to data-driven generative design solutions if the median

performing value lies within the top 10% or greater than the 90% percentile of the data-driven solution space.

By dividing the number of iterations required to achieve goal-driven solutions within the top 10% of those derived using the data-driven generative design method it is possible to derive a value for the efficiency of the process. This is shown in (33)

$$Efficiency = \frac{Iterations_{data-driven}}{Iterations_{goal-driven}}$$

The number of iterations was selected as the metric for comparing efficiency of the two methods rather than computational speed as it is both hardware and software agnostic. If improvements in computational hardware decrease computation time, or different implementations or algorithms for the evaluation metrics are derived, the results within this research will remain valid.

5.6. Software Implementation

It is important to note that the research reported in this thesis is based on fundamental principles and as such will work regardless of programming language or operating system. However, for the purpose of reimplementation, the following description, and schematic (Figure 5-25), of the software implementation used within this research is described.
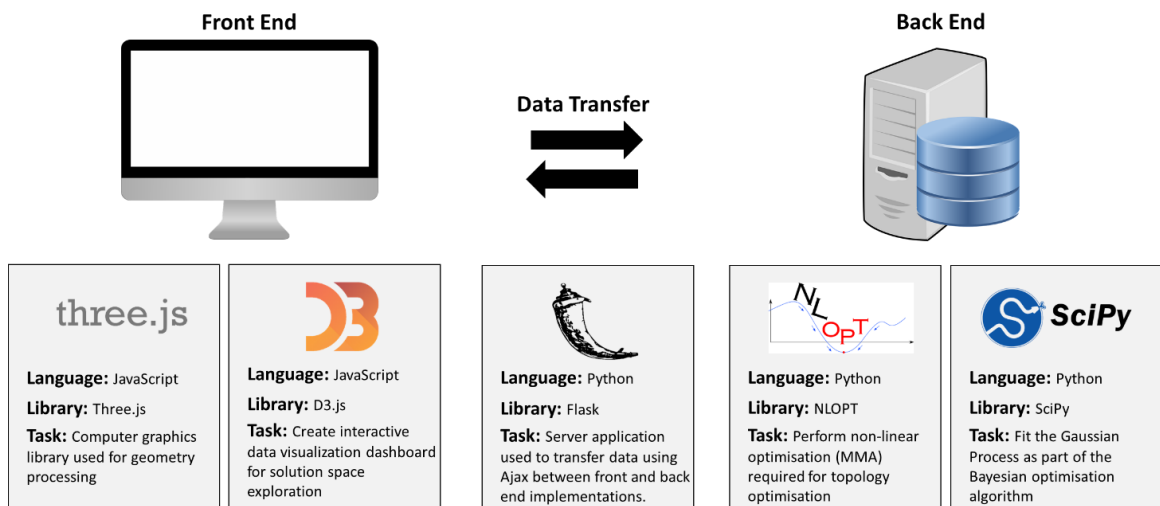
Figure 5-25: Software implementation used to implement the research methodology.

To implement the generative design methods outlined above, a software stack has to be selected that can perform the required functions. These include performing structural optimisation, visualising and analysing both 3D geometric and solution space data, and also performing the structural and Bayesian optimisations. The contrasting requirements for interactive visualisations

and 3D graphics, and scientific computing necessary for generative design systems lead to the decision to utilise multiple programming languages that would be best suited to each of the different requirements.

The generative design environment will be developed in the form of a web application with a HTML/CSS/JavaScript driven front end with a Python back-end connected using Ajax for GET and POST requests through the Flask server framework. Bootstrap and JQuery are used to facilitate the development of the frontend user interface environment. For the purposes of this research, the server is hosted locally. However, this server could be hosted on hardware dedicated for high-performance computing or alternatively, on the cloud. Google Chrome was selected as the web browser to perform the experimental work.

The front-end, referring to the layer that the user can access, is used for geometry processing and data visualisation. The geometry analysis processing is aided with the use of the THREE.js framework for 3D geometry and the D3.js framework for the development of the interactive data visualisation dashboard. The user input, solution synthesis and data visualisation environments are built using separate web-pages and data is transferred between the two using either a cloud-database (Google Firebase) or a pre-saved CSV file depending on user preference. User implementation of the TOPSIS algorithm is developed in JavaScript, for part rotation and for the evaluating the multi-dimensional evaluation criteria that are necessary for generative design.

Python is used to perform the optimisation during the synthesis stage. The truss optimisation code was extended from the 2D model developed by Aranda and Bellido (Aranda and Bellido 2016) to 3D. The NlOpt library is used to perform the non-linear MMA optimisation that is required for the ground-structure topology optimisation. Scipy is used to fit the Gaussian Process that is used for the Bayesian Optimisation.

5.7. Summary

This chapter has outlined the research methodology that will be used in order to achieve the main aim of the thesis as outlined in Chapter 4. The main outcomes from this chapter are as follows:

- The research methodology outlining the underpinning methods used to fulfil the overall objectives of this research and in turn, answer the aim of the research are detailed. The research methodology consists of three stages: the first stage outlines the design problem, a ground structure topology optimisation of a cantilever beam used in conjunction with a build orientation optimisation. The second stage outlines two different approaches to explore the design space, namely data-driven and goal-driven approaches

as well as a pilot study to test the feasibility of using Bayesian optimisation to optimise AM-specific design problems efficiently. Finally, a comparison stage that tests the efficacy of data-driven and goal-driven solutions to find high-performing solutions to various user-defined high-level design problems.

- Three levels of abstraction have been defined mapping low-level evaluation criteria to high-level design objectives. Depending on the stakeholders involved in a project and the skill of the designer, abstraction criteria from each of the levels can be used to drive the global part optimisation. The equations defining the criteria have been defined for the DMLS process. However, the methods are extensible to other AM processes.

- Weighted combinations of the high-level design objectives are used to define three different production scenarios that are common within the AM industry and are representative of the trends currently observed by the latest product releases and technical forecasts (Wohlers et al. 2018).

- A description of the three-stage experimental method used to robustly explore data-driven generative design methods for AM has been provided. A number of MCDA techniques have been analysed for the suitability for generative design. The TOPSIS method was selected and mathematically defined.

- A method for performing efficient build orientation optimisation is defined alongside three different proxy metrics for the required amount of support structure to support overhanging faces. Three methods: grid search, random search and Bayesian optimisation surrogate models are defined in order to test their effectiveness and accuracy at finding the optimal build orientation for two open-access AM parts. These part examples are typical of parts defined using continuum based topology optimisation techniques.

- Goal-driven generative design for AM is defined, and a flow-chart is provided describing the research method used to examine how Bayesian optimisation could be used to efficiently explore AM design spaces in order to locate high-performing solutions.

- A normalised objective function distance metric is defined in order to quantitatively compare the outputs from both data-driven and goal-driven approaches to generative design for AM.

- An outline of the software implementation that will be used to develop a prototype of the generative design system to achieve the above research methods is provided. This is to

aid in the replication of the results provided in the following chapters. However, it should be noted that the research methods are software agnostic, and other software stacks would also be suitable for creating new implementations of generative design tools.

The implementation of this methodology will now be presented, starting with Chapter 6, which will demonstrate the development of the data visualisation dashboard as well as results from the data-driven approach to generative design.

Chapter 6 – Data-Driven Generative Design for Additive Manufacturing

6.1. Introduction

This chapter will implement the method outlined in Chapter 5, Section 5.4.1, to develop a data-driven generative design method for additive manufacturing. A data-visualisation dashboard will be produced to navigate a solution space containing multiple design instances generated using a generative design synthesis method. The solution space will be created using a conventional grid search by altering the build orientation and overhang constraint on the parametric input to a ground structure topology optimisation. A TOPSIS MCDA support tool will be used to locate the top-performing solutions within the solution space for the different production scenarios.

6.2. Data-Visualisation Dashboard to Support Solution Space Exploration

Given that generative design methods produce a series of design solutions with multi-dimensional evaluation criteria, it is beneficial for the designer to have a method to easily interact with the solutions and quickly focus in on areas of high-performing designs. The data visualisation dashboard developed for this research is comprised of three stages, guiding the designer from the entire solution space to a single design solution by utilising different graphical representations. This dashboard has been inspired by multivariate visualisation methods from the Design Explorer (Thornton Tomasetti 2015) and Ashour and Kolarevic (2015) and extended for the specific application for AM-specific generative design problems.

6.2.1. Stage 1 – Solution Space Visualisation

A parallel coordinate plot (Figure 6-1) is used for stage one of the data visualisation. Parallel coordinate plots have been shown to be an effective way of visualising vast quantities of multi-dimensional data (Theus 2008). This gives the designer the ability to visualise the entire solution space. Studies have shown that parallel coordinate plots require a level of interactivity to be useful when large datasets are shown (Wills 2008). One technique for improving user interactivity is to apply brushing to individual axes. This allows the designer to highlight a subset of the solution space by reducing the domain of a selected axis, thus highlighting designs with certain attributes. Brushing is applied within the interactive dashboard by allowing a user to click and drag within an axis, thereby generating a rectangle describing the new axis domain. The filtered results from the parallel coordinates plot are linked to the input to stage two of the visualisation.
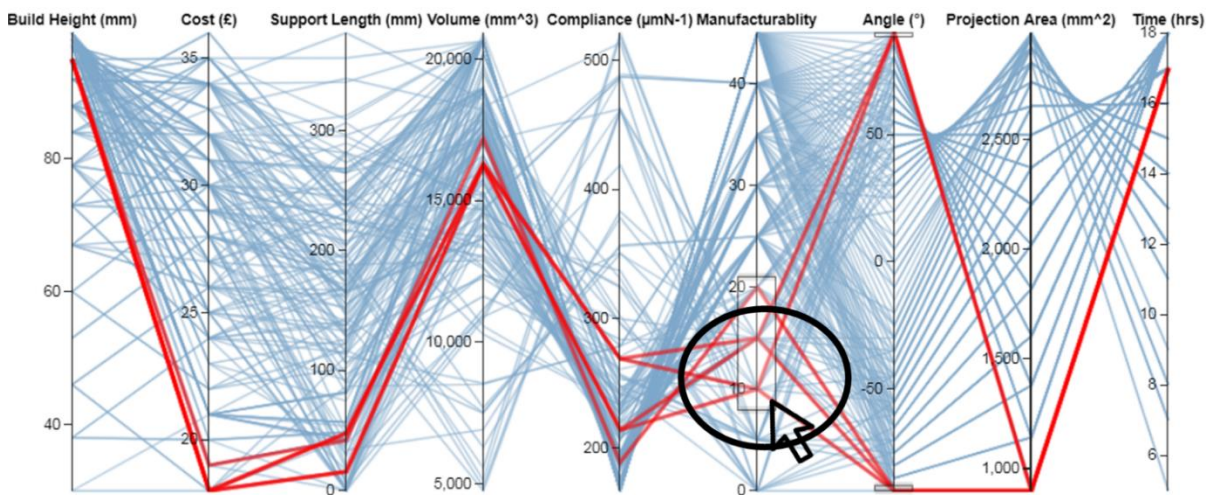
Figure 6-1: Stage one data visualisation showing parallel coordinates plot with the top 5 solutions displayed in emboldened red.

6.2.2. Stage 2 – Filtered Solution Space Visualisation

The aim for the second stage of the data visualisation is to allow the designer to understand the relationships between different evaluation criteria within the data set. A scatterplot was selected to allow the designer to easily see patterns within the data and determine any correlations between design variables. An α-transparency value of 0.7 is applied to all filtered designs to allow the user to distinguish the filtered results from the parallel coordinates plot from the rest of the solution space. Further interactivity is available to the designer within this stage, firstly, by allowing the designer to select the axis variables. By selecting different axes, the designer can view correlations between any two variables. A zoom function is also implemented to allow the designer to examine areas that show multiple design instances with similar parameter values. This is particularly important when large datasets require exploration. The designer is able to highlight particular design instances by clicking on the scatter plot circles. Upon doing this, the circle is

highlighted, and information for the particular design instance feeds into stage three of the visualisation tool. The stage 2 visualisation is shown in Figure 6-2.
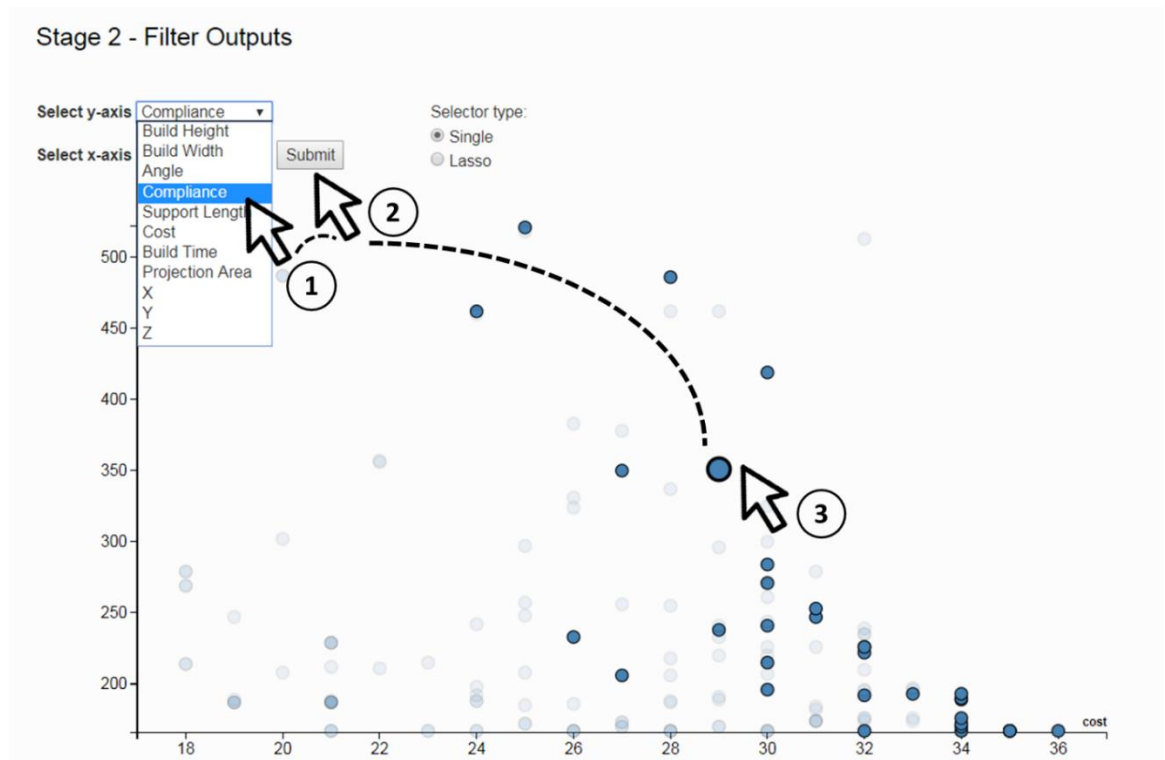
Figure 6-2: Stage 2 data visualisation. The numbered cursors show the human interaction used to navigate this visualisation stage.

6.2.3. Stage 3 – Individual Solution Visualisation

The final stage of the data visualisation involves feeding individual part data back to the user for selection of the most appropriate design solution to a given design problem. The objective of this visualisation stage is to allow the user to evaluate and select superior designs from the solution set. The part data within this stage is normalised against the entire solution set to provide a relative quality metric for each of the parts. Figure 6-3 shows the user interaction that can be used to navigate stage 3 of the data visualisation dashboard.

Two primary charts are selected for this visualisation, namely a radar chart and a bar chart. The radar chart provides a normalised comparison of the part for three user-selected criteria against the solution set. Radar charts provide a method of visualising high dimensional quantitative data. Radar charts are produced for each design comparison and overlaid using multiple colours and opacities. A bar chart is also presented that contains part specific information, including elements to indicate part manufacturability. Three manufacturability criteria are shown in Figure 6-3, notably, part volume, part surface area and part cost. However, any combination of the multi-dimensional solution space can be selected. The manufacturability data provided is scaled against

the data set to a value between 0-10. This is so that the designer can easily get numerical feedback on a parts quality relative to the entire solution set. A 3D visualisation of the design is also shown to allow the designer to view the aesthetics of the final part to determine its acceptability as a final design solution.

The designer can compare different solutions by rating individual potential solutions using interactive slider bars. The designer is given the opportunity to rate the design from 0 to 10. When the user clicks the 'Rate and Compare' button the current design image and slider bar value are stored above the radar chart; these designs can then be accessed at any point by clicking on the image. Two designs can be stored and compared at any time. The evaluation score for each part is stored within an array and can be accessed or exported from the visualisation environment at any time.
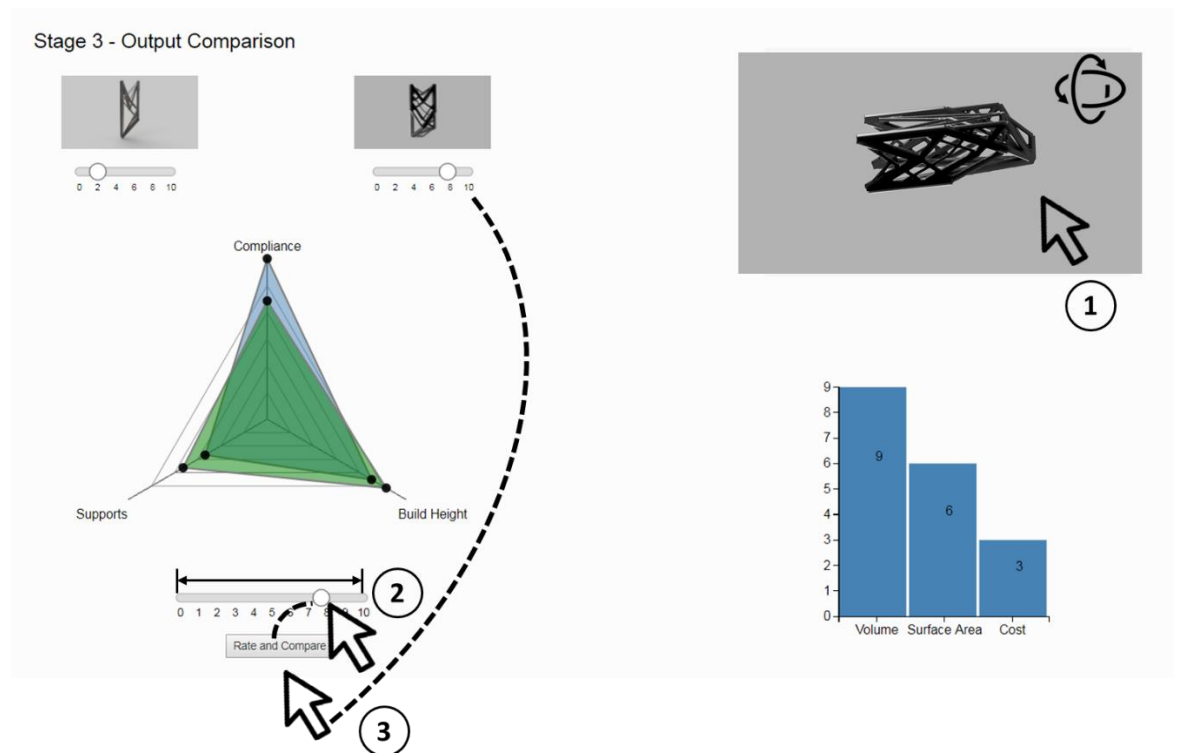
Figure 6-3: Use of stage 3 of the data visualisation dashboard to explore individual solutions within the AM solution space.

6.3. Data-Driven Generative Design for Additive Manufacturing

To generate the AM solution space, the ground structure was rotated between the angles of -90° and 90° in increments of 5°. This was repeated with varying overhang constraints applied between 0°, which leads to a design with no support structure, and 45°, where no manufacturability constraint is applied. The manufacturability constraint is again applied at increments of 5°.

The support structure metric used within this chapter is total support length. Based on the drawbacks associated with this metric discussed in Chapter 5, it is necessary to check the distribution of triangle areas in the generated truss structures. Figure 6-4 shows two histograms depicting the distribution of triangle areas in trusses generated from the ground structure topology optimisation. The results show that no significantly large triangles are present in the truss-based meshes. This is aligned with the way in which cylinders are typically triangulated in STL files. As the triangles areas are closely distributed, the potential drawbacks of the support length metric are not present for truss based geometries.

Figure 6-4: Histograms showing triangle area distribution for two typical truss based structures generated using a ground structure topology optimisation method.

Figure 6-5 shows the total length of the support structure required to print the solution for each build angle. As expected, there is no support requirement when the manufacturability constraint is applied. As the manufacturability constraint is relaxed, the length of support increases until no manufacturing constraints are applied, and the greatest amount of support structure is required. Furthermore, in line with previous studies (Mass and Amir 2016), when manufacturability constraints are applied to the optimisation, the structural performance of the part decreases. As the angle moves closer to zero, the part performance, as indicated by the compliance value, decreases dramatically. This can be attributed to a significant reduction in the structural rigidity of the ground structure. Combinations of build orientation and manufacturability constraint that have build angle in which the principle axis of the cantilever beam remains close to horizontal and a high manufacturability constraint do not produce feasible solutions, this is attributed to infeasible ground structures. These solutions were removed from the solution space leaving a total of 175 solutions for exploration. The results are shown in Figure 6-6.

Figure 6-5: Support Length vs Angle for total solution space.

Figure 6-6: Compliance vs angle for total solution space.

6.4. Exploring Different Additive Manufacturing Production Scenarios

The application of data-driven generative design for additive manufacturing methods will now be used to find the n best performing solutions within the solution space in Figure 6-5 and Figure 6-6, for each of the production scenarios defined in Section 5.3.3.4.

6.4.1. Highest Overall Performance

Consider a scenario in which the designer wants the best overall performance by equally assigning the weights to the four possible performance criteria. This is the scenario described in persona C, in the research methodology. This scenario represents a common situation in which the designer may not fully understand the problem and would like to test the general case of optimising for the overall performance of a part. The performance criteria weightings, w, are detailed below:

$$w = \begin{cases} structural\ performance = 25 \\ build\ plate\ packing = 25 \\ build\ speed = 25 \\ build\ cost = 25 \end{cases}$$

The results in Figure 6-7 demonstrate that MCDA finds preferable areas of the solution space minimising compliance, build projection and part cost. It is generally desirable to print at build angles whereby the principal axis of the cantilever beam is at 90° with respect to the build plate, as this minimises the support required for the part, maximises the number of parts that can be placed on the build plate and also maintains desirable structural performance. The output solutions and corresponding evaluation criteria for the general performance persona, which are built at 90° with respect to the build plate, are shown in Table 6-1.

Figure 6-7: Cost against build projection (left) and (right) compliance against build time plots with top solutions shown in (blue) diamonds for highest overall performance production scenario.

Optimisation Output			
Manufacturing Constraints	No Support Constraint	Partial Support Constraint	Full Support Constraint
Build Angle (°)	-90	-90	-90
Build Projection (mm²)	900	900	900
Compliance (μmN⁻¹)	167	214	247
Support Length (mm)	40,795	16,169	0

The parallel coordinates plot showing the result in conjunction with the solution space is displayed in Figure 6-8. The top five performing solutions are shown with emboldened (red) lines. The plot shows that designs with a full, partial or no support constraint may be appropriate in this case, depending on the permissible structural performance trade-offs defined by the user.
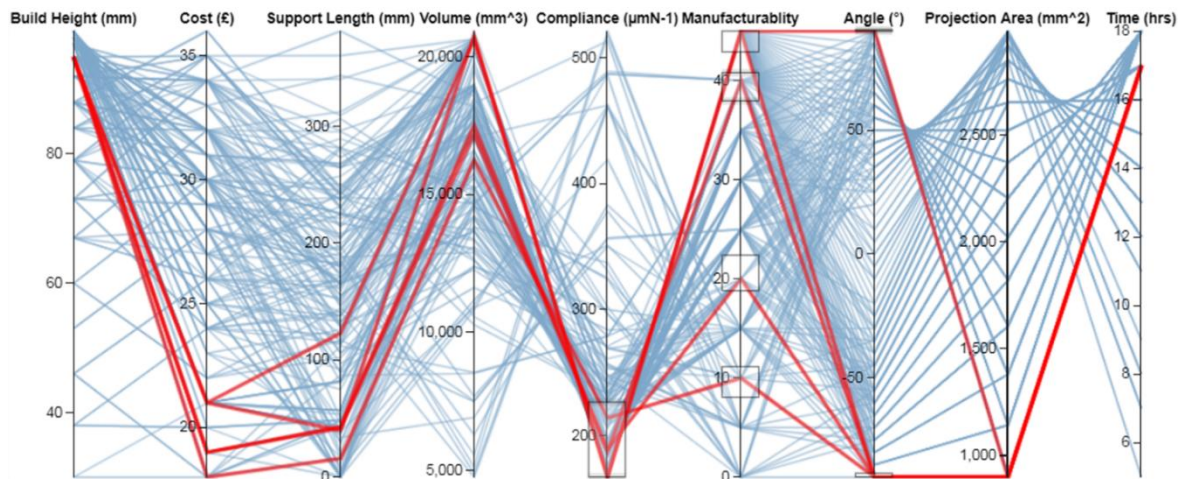
Figure 6-8: Parallel Coordinates plot showing top 5 solutions in solution space for the persona c, the highest overall performance production scenario.

6.4.1.1. Sensitivity Analysis on Input Selections

It is necessary to examine the effect of altering the weights on the output from the MCDA. To test this, each parameter is altered by 10%, and the top 5 outputs from the MCDA recorded. The following results show that for multi-objective problems with clear trade-offs, such as the overall

performance case, the MCDA is capable of finding the top results from each of the compromise conditions. This indicates that it would be beneficial to test multiple different weightings in the MCDA to broaden the number of potential outputs for exploration. The role of the designer is important in generative design. For multi-objective problems, it is the role of the designer to select the most appropriate solution given potentially conflicting outcomes (Chiandussi et al. 2012).

The following results indicate that the output from the MCDA is responsive to small changes (10%) to the user-defined input weightings. The results for the key sensitivity analysis tests are now explained. Firstly, the following weights are tested.

$$w = \begin{cases} structural\ performance = 15 \\ build\ plate\ packing = 25 \\ build\ speed = 35 \\ build\ cost = 25 \end{cases}$$

Figure 6-9: Output solution space showing top performing solutions in blue diamonds showing a preference toward high production quantity.

The results in Figure 6-9 show that this weighting scenario favours the high production quantity scenario, with a high build plate packing ratio trading off against build time and structural performance.

Secondly, the weights are set as:

$$w = \begin{cases} structural\ performance = 25 \\ build\ plate\ packing = 25 \\ build\ speed = 15 \\ build\ cost = 35 \end{cases}$$

Figure 6-10: Solution space showing top performing solutions (blue diamonds) showing preference towards high structual performance and a low build time.

The results for the following weighting, demonstrated in Figure 6-10, depict a preference towards the structural performance of the part with the trade-off occurring with poor build plate packing.

Finally, the weights are adjusted to:

$$w = \begin{cases} structural\ performance = 25 \\ build\ plate\ packing = 35 \\ build\ speed = 25 \\ build\ cost = 15 \end{cases}$$

Figure 6-11: Output solution space showing top performing solutions (blue diamonds) finding high-performing solutions for conflicting requirements.

This final result shown in Figure 6-11, displays the compromise condition with the best results favouring either the part performance or part throughput. In this trade-off situation the designer has to select which results will be the most appropriate for their given design specification.

6.4.2. High Performance Part Production

This second test case represents the high-performance environment (such as competitive motor racing) that was described by persona B in the Chapter 5. This production scenario aims to maximise structural performance with the relaxed constraints being part cost and material consumption. To locate parts within the solution space that maximise structural performance, the percentage allocation of weights, w, is given as follows:

$$w = \begin{cases} structural\ performance = 80 \\ build\ plate\ packing = 0 \\ build\ speed = 20 \\ build\ cost = 0 \end{cases}$$

Figure 6-12 shows that the MCDA analysis highlights the area of the solution space with the lowest compliance but also shows the area with the lowest amount of support structure, which although not specifically requested by the designer is clearly advantageous. The solution space is highlighted in (red) circles with the top five results from the MCDA analysis plotted in (blue) diamonds.

Figure 6-12: (Left) Part cost versus build projection area plot and (right) compliance versus part build time showing top results in diamonds for high-part performance scenario.

Figure 6-13: Parallel coordinates plot showing optimised solutions for structural performance with respect to solution space

A parallel coordinates plot is shown in Figure 6-13, which highlights the position of these solutions with respect to the entire solution space. The results show that a manufacturability constraint should not be applied and that minimal support can be achieved by building the structure at 0° or 90° to the build plate. For a single part, build time can be reduced by minimising build height, therefore the solution at 0° build angle is selected. To validate the printability, the design was exported from the design environment and printed with support structures, using an Ultimaker 2+ Extended material extrusion printer as shown in Figure 6-14 (centre). The supports were removed after the build, and the final part is shown in Figure 6-14 (right). The successful print provides two important results. Firstly, it indicates that the geometric constraints selected for the structural optimisation still lead to quality builds. Secondly, it emphasises the importance of understanding both the design requirements and the generative synthesis method as support structures can have a large contribution to the overall material usage. This is particularly challenging with ground-structure topology optimisation synthesis methods. Significant time and care must be dedicated to support structure removal to prevent compromising the target part.

Figure 6-14: Optimised design for structural performance, digital (left), with supports (centre) and after support removal (right)

6.4.3. High Volume Part Production

To maximise the economic advantage and increase the uptake of AM as a competitive manufacturing process, it is imperative that per-part costs are reduced. Pili et al (2015), show that total build cost can be reduced by up to 91.2% if the total area of the build plate is used when

compared to a single part; this could be exploited by utilising existing research undertaken in the field of build plate packing optimisation (Zhang et al. 2018). In high production quantity (represented by persona A), the overall cost of the part, including part-processing time, is extremely important. Support structure also adds significant time and cost to builds (Strano et al. 2013). Therefore, reducing support requirements is crucial for high production quantity. The weights, w, have been selected to maximise build plate packing, with some additional input from the per-part build time. As multiple parts are required, there is a trade-off between the minimum part cost and the ability to nest multiple parts on a single build plate.

$$w = \begin{cases} structural\ performance = 0 \\ build\ plate\ packing = 20 \\ build\ speed = 0 \\ build\ cost = 80 \end{cases}$$

Figure 6-15: (Left) Part cost versus Build projection area plot and (right) compliance versus part build time showing top results in diamonds for a high-production quantity scenario.

The results in Figure 6-15 (left), show that the optimal parts for high production quantity occur with a minimum build projection on the build plate, but also aim to minimise the supports that occur thus, reducing build cost and post-processing time. This can be seen to be traded off against structural performance (right). In Figure 6-16, it is clear that placing the parts on the build plate with a smaller projection area (left) yields a greater build packing ratio (2.57:1) as it is possible to fit 36 parts onto the build plate versus 14 when the build projection area is at its greatest projection area (right). Using the Cura build preparation software (Ultimaker 2017), the print time can be calculated for each production scenario. The highest part performance build plate takes approximately 3d 5h 38min to complete, whereas the build time for highest production quantity was 4d 22h 09min. Taking the packing ratio into account, the maximum production quantity scenario yields a 60.87% efficiency improvement in part throughput. This adds further credence to the importance of incorporating design objectives within structural optimisation. The parts

built in Figure 6-16 (left) representing the high-production quantity scenario perform differently to the corresponding solutions in Figure 6-16 (right), where structural performance is critical.

Figure 6-16: Comparison of packing efficiency when minimising vs maximising build projection area. (Left) High production quantity persona output compared with (right) highest structural performance output.

The highest rated design has the smallest build width and can accommodate the greatest number of parts on the build plate. Furthermore, the manufacturing constraint is applied, eliminating the requirement for support structures. The build height is relatively high, but this is traded-off against the packing volume. The position of the top 5 solutions relative to the solution space is highlighted in the parallel coordinates plot (Figure 6-17).
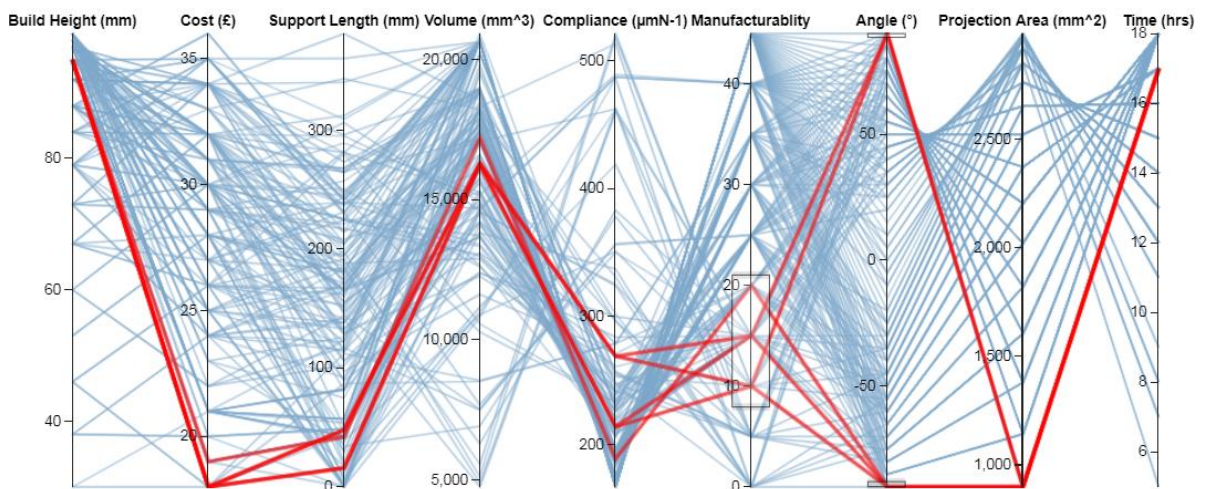
Figure 6-17: Parallel coordinates plot showing optimised solution for high production quantity with respect to solution space

The result from the constrained optimisation for the maximum production quantity goal can be seen in Figure 6-18. The as-built version is also shown demonstrating its manufacture without support structure. This highlights that a maximum overhang of 45° is suitable for manufacturing using the FDM process. Surface finish is shown to deteriorate with distance from the starting node, this indicates that the potential benefits of including an aspect ratio constraint.

Figure 6-18: Digital (left) and as printed designs (right) for optimised design for maximum production volume

6.5. Summary

By applying a novel DfAM method that combines a ground-structure topology optimisation with TOPSIS MCDA, design solutions for a cantilevered beam have been designed by optimising for three common AM-specific design personas. This chapter shows the capability of TOPSIS MCDA to accurately highlight areas of the solution space that maximise part quality with respect to abstracted design and business criteria.

The main findings from this chapter are as follows:

- The usefulness of manufacturing constraints within topology optimisation is dependent on design objectives. It is feasible that blanket utilisation of manufacturing constraints may lead to sub-optimal solutions to certain design tasks.
- Defining design objectives early in the design process and utilising generative design to explore the design space (e.g. multiple build orientations) can help to discover higher performing design solutions.
- Small changes in weightings can affect the final outcome for the MCDA. Therefore it is recommended that for multi-objective problems, the designer tests multiple different weightings to locate a broader range of feasible solutions. Thus highlighting the importance of keeping a designer within the generative part design process.
- As manufacturing constraints have a direct effect on structural performance, it is crucial to understand any machine-specific manufacturing limitations and utilise these within topology optimisation to leverage the full potential of AM.

Presently, the generated solutions space is independent of the defined goals, and the MCDA analysis is used to highlight areas of the solution space that maximise these goals. In Chapter 7,

intelligent search strategies, such as Bayesian optimisation will be introduced to reduce the computational expense of creating a solution space of structurally optimised parts.

Chapter 7 - Bayesian Optimisation in Generative Design

7.1. Introduction

Chapter 6 has shown the applicability of data-driven methods to locate the highest-performing solutions within a generated AM solution space for multiple differing objective functions. One of the disadvantages of using data-driven methods is the computational and time expense associated with generating a solution space at a suitable resolution to capture high-performing. This chapter will explore alternatives to data-driven generative design methods that can locate high-performing solutions by finding build orientations that minimise support structure requirements with much greater computational efficiency.

7.2. Alcoa Bracket

The first test part to be evaluated is the Alcoa bracket. Even with only two dimensions, due to computational limitations, it is infeasible to compute all possible angle permutations. A comparison of iterations at 15, 20, 30 degree increments are taken, and the minimum value of the objective function is found. The results for each of the evaluation metrics can be viewed in Table 7-1.

Table 7-1: Best solutions for different numbers of iterations of grid search for the Alcoa bracket.

Orientation Interval [°]	Minimum Number of Overhanging Faces	Minimum Support Length [mm]	Minimum Support Volume [mm³]
15	304	6,412	3,461
20	438	7,311	5,541
30	488	7,381	7,371

To determine the number of iterations required to approximate the Gaussian process response surface for overhanging surfaces or the amount of support, a series of tests were run with varying numbers of function evaluations. The best result from this exploration is displayed in Table 7-2.

Test Method	Minimum Number of Overhanging Faces	Minimum Support Length [mm]	Minimum Support Volume [mm³]
Random	301	6,337	2,193
Bayesian Optimisation	300	6,016	1,546

For each number of iterations explored, tests were repeated 30 times to perform a non-parametric Mann-Whitney U test. The null hypothesis of this test was that there would be no differences between the median value for the evaluation criterion found with a grid search and, independently, either random search or Bayesian optimisation. The results are presented in the boxplots of Figure 7-1 to Figure 7-3, where the horizontal red line corresponds to the best value found throughout all tests. Tables including the median value of each test and the corresponding p-value obtained from the statistical comparison and shown in Table 7-3, Table 7-4 and Table 7-5. The circles represent data points that are less than $1.5 \times IQR$ from quartile 1 (Q1) or greater than $1.5 \times IQR$ from quartile 3 (Q3), where IQR is the inter-quartile range (Q3-Q1).

Table 7-3: Iteration test values for number of overhangs for Alcoa test bracket.

Test Criteria	Random	Bayesian Optimisation	
	Median	Median	P-Value
15	604	584	1.90×10^{-1}
25	570	535	1.40×10^{-1}
35	455	399	1.90×10^{-3}
50	455	336	7.50×10^{-4}
100	410	309	4.52×10^{-6}
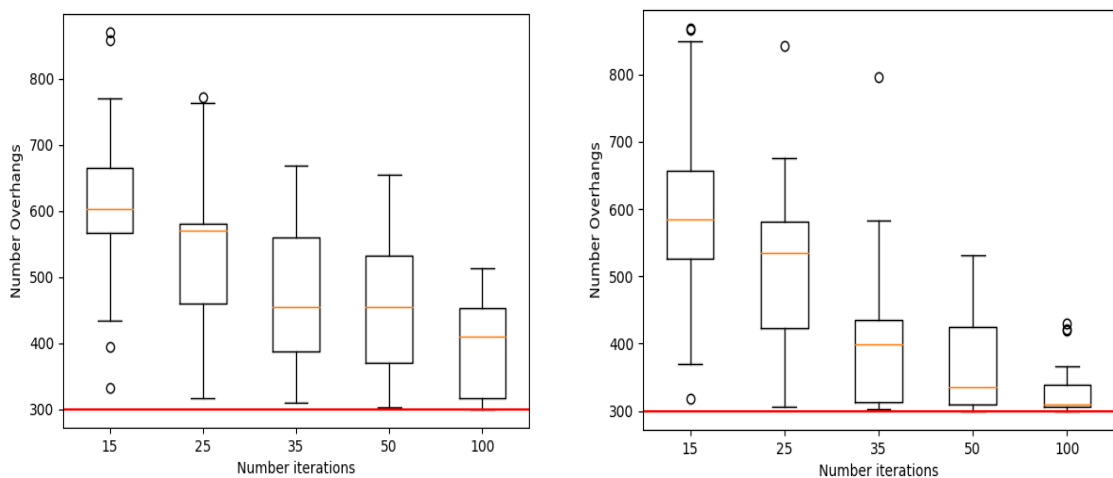
Figure 7-1: Comparison of number of iterations for random search (left) and Bayesian optimisation (right) and the number of overhanging faces (n=30 tests)

Table 7-4: Iteration test values for total support length for Alcoa test bracket.

Test Criteria	Random	Bayesian Optimisation	
	Median (mm^2)	Median (mm^2)	P-Value
15	8537	7386	1.60×10^{-2}
25	7381	7179	5.41×10^{-4}
35	7361	7155	2.05×10^{-6}
50	7315	7122	2.45×10^{-6}
100	7221	7119	9.25×10^{-4}

Figure 7-2: Comparison of number of iterations ran for random search (left) and Bayesian optimisation (right) and the output support length (n=30 tests)

Table 7-5: Iteration test values for total support volume for Alcoa test bracket.

Test Criteria	Random	Bayesian Optimisation	
	Median (mm^3)	Median (mm^3)	P-Value
15	12381	10239	3.20×10^{-1}
25	9334	8086	0.20×10^{-1}
35	7633	6779	3.50×10^{-2}
50	7393	2338	2.70×10^{-4}
100	7100	2220	1.48×10^{-5}

Figure 7-3: Comparison of number of iterations ran for random search (left) and Bayesian optimisation (right) and the output support volume (n=30 tests)

This experiment established that 35 iterations of Bayesian optimisation with the Matèrn kernel and Expected Improvement acquisition function produced improved results when compared to random search. Now, all combinations of acquisition and covariance functions are tested and compared against random search to identify their relative performances when trying to minimise support structure. The resulting box plots are shown in Figure 7-4. The median values and p-value obtained from the comparison of random search and Bayesian optimisation are found in Table 7-6, Table 7-7, and Table 7-8. The horizontal blue line indicates the best value found in all repeated tests. Comparing each combination of covariance and acquisition function to random search, shows that Bayesian optimisation provides significantly better results for 35 iterations. There is no significant difference between any of the Bayesian optimisation acquisition function and covariance functions.

Figure 7-4: Exploring the effect of changing the covariance and acquisitions functions for 35 iterations of Bayesian optimisation compared to random search for each support structure evaluation metric for the Alcoa bracket (n=30 tests).

Table 7-6: Median values for number of overhangs for Alcoa test bracket

Test Criteria	Median Value	P-Value
Random	455	-
UCB-RBF	412	2.10×10^{-2}
UCB-Matern	437	9.00×10^{-3}
EI-RBF	421	9.00×10^{-3}
EI-Matern	399	1.90×10^{-3}

Table 7-7: Median values for total support length for Alcoa test bracket

Test Criteria	Median Value (mm^2)	P-Value
Random	7361	-
UCB-RBF	7160	2.45×10^{-6}
UCB-Matern	7017	1.02×10^{-6}
EI-RBF	7190	6.44×10^{-6}
EI-Matern	6982	2.06×10^{-6}

Table 7-8: Median values for total support volume for Alcoa test bracket

Test Criteria	Median Value (mm^3)	P-Value
Random	7633	-
UCB-RBF	5039	1.10×10^{-2}
UCB-Matern	6739	3.30×10^{-2}
EI-RBF	4678	3.51×10^{-2}
EI-Matern	6779	3.60×10^{-2}

7.3. GE Bracket

The following section describes the results arising from the GE bracket experiments. A grid search was performed using 15°, 20°, and 30° intervals and the best result for each of the evaluation criteria is recorded in Table 7-9.

Table 7-9: The best solutions identified for different numbers of iterations of grid search when testing the GE bracket.

Orientation Interval [°]	Minimum Number of Overhanging Faces	Minimum Support Length [mm]	Minimum Support Volume [mm³]
15	4,036	63,718	10,181
20	3,930	66,583	22,854
30	4,101	76,029	28,142

Based on the results in Section 7.2. 35 samples are taken for both the random search and the combinations of acquisition and covariance functions in order to determine their effect on the evaluation criteria. The graphs of the results are shown in Figure 7-5, and Table 7-10 shows the

minimum values from each method. Table 7-11,Table 7-12 and Table 7-13 provide the p-values for each of the statistical tests. The horizontal blue line indicates the best results found in all repeated tests.

Table 7-10: Exploring the effect of changing the covariance and acquisitions functions for 35 iterations of Bayesian optimisation compared to random search for different performance metrics for the GE bracket (n=30 tests).

Test Method	Minimum Number of Overhanging Faces	Minimum Support Length [mm]	Minimum Support Volume [mm^3]
Random	3,896	61,058	13,401
Bayesian Optimisation	3,883	60,015	10,181

Table 7-11: Median values for total number of overhangs for GE bracket

Test Criteria	Median Value	P-Value
Random	4091	-
UCB-RBF	3961	4.61×10^{-6}
UCB-Matern	3886	6.36×10^{-9}
EI-RBF	3984	9.17×10^{-5}
EI-Matern	3944	3.02×10^{-6}

Table 7-12: Median values for total support length for GE bracket

Test Criteria	Median Value (mm^2)	P-Value
Random	83860	-
UCB-RBF	65920	2.72×10^{-6}
UCB-Matern	60665	4.37×10^{-11}
EI-RBF	65581	4.59×10^{-6}
EI-Matern	65074	2.618×10^{-6}

Table 7-13: Median values for total support volume for GE bracket.

Test Criteria	Median Value (mm^3)	P-Value
Random	37699	-
UCB-RBF	31464	9.60×10^{-2}
UCB-Matern	38570	4.30×10^{-3}
EI-RBF	32301	6.60×10^{-2}
EI-Matern	33659	4.91×10^{-2}
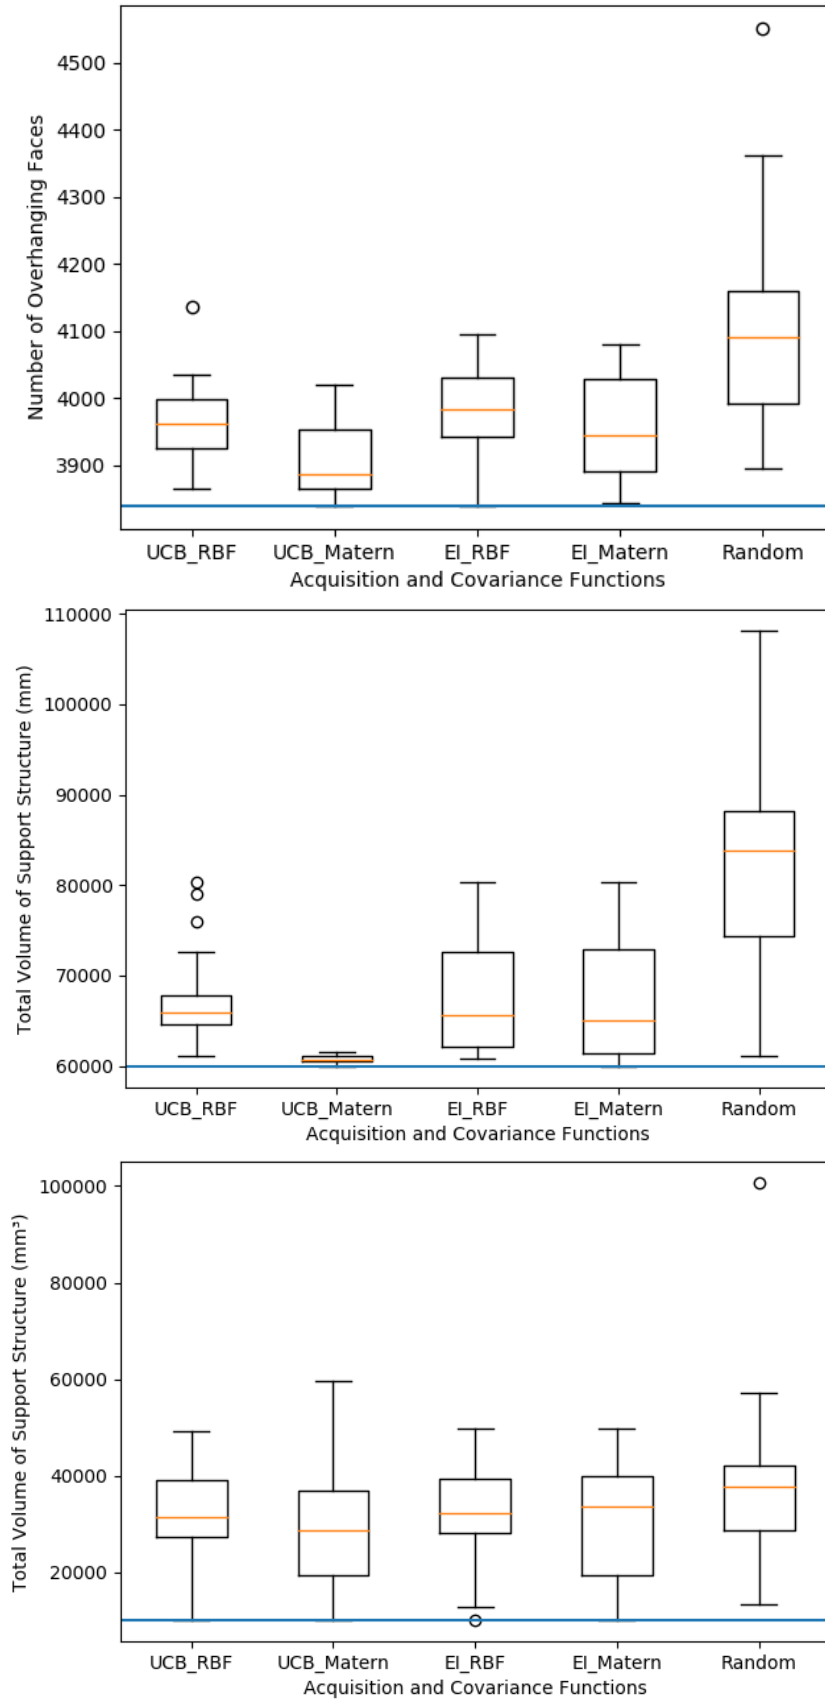
Figure 7-5: Exploring the effect of changing the covariance and acquisitions functions for 35 iterations of Bayesian optimisation compared to random search for support structure evaluation metric for the GE bracket (n=30 tests).

7.4. Key Observations

The results show that even for a small two-dimensional design problem, a grid search is an ineffective method to locate the optimal build orientation unless a fine resolution is selected. This may be computationally infeasible for large mesh sizes unless the best solution coincides with the selected resolution.

By selecting real-world examples for test parts, it has been demonstrated that Bayesian optimisation can operate alongside representative mesh sizes. The results show that good solutions can be found by forming a response surface using small numbers of iterations. Autodesk Meshmixer (Autodesk 2018a) is a commercially-available software package with the option to find optimised build orientations. The software provides multiple outputs for this optimisation, and these have been compared to the best results from the metrics for support length (Figure 7-6 (a) and Figure 7-7 (a)) and support volume (Figure 7-6 (b) and Figure 7-7 (b)). From visual inspection, it is clear that the results correspond well with the results given by the commercial software.

(a)   (b)

Figure 7-6: A comparison of results between the Meshmixer software (top) and Bayesian optimisation algorithm (bottom) for the Alcoa bracket.

Figure 7-7: A comparison of results between the Meshmixer software (top) and Bayesian optimisation algorithm (bottom) for the GE bracket.

7.4.1. Evaluation of Optimisation Methods

The Alcoa bracket results highlight that grid search is an inefficient method for finding the optimal build orientation. Firstly, using the number of overhanging facets as the evaluation criterion showed that:

- 35 iterations of random search produces a better median value than the best result from grid search with a resolution of 30°.
- 35 iterations of Bayesian optimisation produce a better result than a 20° resolution of grid search.
- Bayesian optimisation was found to produce a better result than a 15° increment in the grid search using 50 iterations; a 12.5 fold improvement in required iterations for an improved result.

Using support length as the evaluation criterion:

- 25 iterations of Bayesian optimisation produce a better median value than increments of 20° in the grid search.
- Both Bayesian optimisation and random search find better solutions than all resolutions of grid search.

Finally, using the estimated support volume as the evaluation criterion:

- 35 iterations of Bayesian optimisation and 50 iterations of random search both find better median solutions than the best solution for increments of 30° in the grid search.
- 50 iterations of Bayesian optimisation produce a better median result than increments of 20° in the grid search.

- 100 iterations of Bayesian optimisation produces a better median than all tests performed using grid search; a six-fold improvement in efficiency.

The "best value" results are more challenging to compare. The total number of iterations for the best value can be calculated as the product of the number of tests (n=30) and the number of iterations run per test. For example, the best Bayesian optimisation result for 35 iterations arose from a total of $30 \times 35 = 1{,}050$ evaluations (see Figure 7-8 and Figure 7-9). The best value for the Alcoa bracket was found with the rotation angles 303° and 31° about the x- and y-axes, respectively. As these numbers are prime, the grid search would have to have to use 1° increments to locate the same value.

Figure 7-8: Comparison between grid search, random search and Bayesian optimisation for differing numbers of iterations and support structure evaluation metrics using the Alcoa test bracket.

Figure 7-9: Comparison between different covariance and acquisition functions for the GE bracket. Showing the median and best values of support evaluation criteria metric compared to the best value found using grid search.

The grid search results from the GE bracket (using the number of overhanging facets as the evaluation criterion) show that 20° resolution performs better than 15° resolution in the grid search. This further illustrates that grid search is a challenging method for use in optimising build orientation unless the resolution is extremely small. Using the same evaluation criterion, the median value arising from 35 iterations of random search is lower than the values found using 30° increments in the grid search.

The median value for the Bayesian optimisation for all combinations of acquisition and covariance functions outperforms grid search for 15° increments after just 35 iterations. Both random search and Bayesian optimisation find lower values than all grid search tests, with Bayesian optimisation producing the best results. Bayesian optimisation performed better than grid search for 30° increments, and the combination of the UCB acquisition function and Matèrn covariance function produces a better median result than all tested increment sizes for grid search.

Using the estimated support volume criterion with the GE bracket shows that the RBF covariance function underperforms when compared against the Matèrn. Accordingly, the results yield p-values slightly above 0.05 when compared to random search. Both p-values for the Matèrn covariance function show that Bayesian optimisation performs significantly better than random search. This suggests that the addition of a smoothness factor can be important for certain geometries.

The success of Bayesian optimisation indicates that this method is effective, irrespective of which evaluation criterion is used. If fast and accurate implementations of support structure algorithms are available, they can be substituted as the evaluation metric, and Bayesian optimisation can still efficiently locate high-performing results. This shows that this method is likely to generalise to the support structure methods described in Chapter 2 and also GPU-accelerated implementations.

For both test parts, p-values from the Mann Whitney U test show that Bayesian optimisation consistently outperforms random search for thirty-five iterations or greater. This indicates that 35 iterations is a safe estimate for the minimum number of iterations required to locate high-performing results. However, it should be noted that higher numbers of iterations have been shown to produce lower median values across repeated tests. Therefore, the designer must find a balance between computational time and accuracy when parts have large mesh sizes.

7.4.2. Comparison of Evaluation Criteria

To compare each of the evaluation criteria, both test parts were processed using Ultimaker's Cura v3.4.1 build preparation software (Ultimaker 2017). The parts were orientated according to the median and best-recorded results for 35 iterations of Bayesian optimisation (in terms of support

minimisation). The Cura software was set to the default settings: 20% infill, 0.1mm layer height, PLA material, 0.4mm nozzle with a 45° max overhang angle. Table 7-14 and Table 7-15 show the Cura results. The data provided includes the support quantity in terms of filament length and mass alongside the rotation angle used to achieve these results for both the median and best observed values from the Bayesian optimisation.

Table 7-14: Comparison of different evaluation criteria using mass and filament usage length measurements given by the Cura software for the Alcoa bracket.

Evaluation Criterion	Median support quantity, length and mass [m, g]	Rotation about x- and y-axes [°]	Best support quantity, length and mass [m, g]	Rotation about x- and y-axes [°]
Number of Overhangs	4.01m/~32g	[303, 215]	3.75m/~30g	[263, 226]
Total Length of Support	3.69m/~29g	[64, 93]	3.12m/~25g	[270, 185]
Total Volume of Support	5.22m/~41g	[329, 11]	3.37m/~29g	[303, 31]

Table 7-15: Comparison of different evaluation criteria using mass and filament usage length measurements given by the Cura software for the GE bracket.

Evaluation Criterion	Median support quantity, length and mass [m, g]	Rotation about x- and y-axes [°]	Best support quantity, length and mass [m, g]	Rotation about x- and y-axes [°]
Number of Overhangs	13.33m/~105g	[189, 171]	13.04m/~103g	[8, 0]
Total Length of Support	12.89m/~102g	[16, 0]	13.12m/~104g	[194, 176]
Total Volume of Support	13.35m/~106g	[217, 177]	13.08m/~103g	[225, 0]

The test results for the Alcoa bracket (Table 7-14) show that the median value lies within 8% of the mass of the minimum value found. Using the number of overhanging facets as the evaluation criterion performs comparatively worse than the support volume, with the support length metric resulting in the most effective minimisation of support structure. The suggested Meshmixer rotations, as depicted in Figure 7-6, were imported into Cura and shown to have a mass of 28g and 24g, respectively. This indicates that for this test part, the support length and support volume criteria are accurate representations with the support length metric locating values within 3.45% of commercial software. As the Alcoa bracket has a plane of symmetry, the best values are mirrored. For example, x-axis and y-axis rotations of [270°, 185°] are equivalent to [90°, 355°]. To improve the efficiency of the search, parts with a plane of symmetry should exploit this efficiency and half the search domain.

The value of total support length for the GE bracket shows that the lowest value of support structure is also found using the support length metric. The median value for the total support volume is lower than the maximum value, indicating that the support volume metric does not accurately represent the total amount of support structure required to support the bracket within the Cura software. This can be attributed to the limitations described in Chapter 5.

The GE bracket appears to have been designed with the [90°, 0°] orientation in mind. However, if the results arising from the total support length criterion are compared against the results for the original orientation, a 17.7% saving is made in support material. The highest performing Meshmixer solutions, shown in Figure 7-7, were found to have masses of 105g and 101g, respectively. This shows that the support length criteria is within 2.86% of this value and the support volume criteria is within 1.94% of the observed value.

The results show that support length is a good approximation of the total support requirement of the part. Provided that the part is not subject to the limitations of this metric as described in Chapter 5, the support length requirement acts as a suitable and computationally efficient metric for the total support requirement. In addition, by selecting ray casting as the method to perform support length, the speed of this evaluation criterion can be improved further by parallelisation using GPU accelerated computing.

7.5. Summary

The research within this chapter presents a novel method to determine the optimal build orientation of an additively manufactured part using Bayesian optimisation. Different evaluation criteria are assessed to determine their success in finding high-quality build orientations. These are then compared against commercial software, and their performance is commensurate. The proposed method was verified using two open-access engineering parts that are typical in AM applications.

The following observations can be drawn from the research in this chapter:

- The number of overhanging facets, support structure length and support structure volume all approximate the total amount of support structure well. Based on the efficiency of calculation and the performance benchmarked against commercial software, support length is recommended as the most appropriate support proxy providing the parts do not have large flat faces. The results for each support structure metric have been shown to align within 3.5% of commercial software estimations.

- Grid search is a poor method for finding the optimal build orientation unless the angular increments are small. This is often not feasible due to computational constraints. 35 iterations of Bayesian optimisation has been shown to consistently outperform grid search at an angular increment of 20° (324 iterations) and perform significantly better than random search for the same number of iterations.

- Whilst Bayesian optimisation does not always guarantee locating the global optimum solution, the median solution is generally close to the best observed solution and has been shown to be over 17 times more efficient at finding good build orientations than grid search.

The results in this chapter have shown that Bayesian optimisation can be used to significantly improve the efficiency of locating the high-performing build orientations for additive manufacturing. The use of proxy metrics for evaluating the amount of support structure provides useful and efficient approximations of the true value (as demonstrated using build preparation software). The ability to efficiently optimise build orientation will reduce the amount of unnecessary support structure in AM builds. This, in turn, will reduce manufacturing times, waste material and the significant time and cost penalties incurred in removing support structure. This is confirmed by the results, which reduced the total mass of the GE bracket by 16.2%, by optimising build orientation.

The results from this chapter have shown that Bayesian optimisation can be used as an effective surrogate model for finding high-performing solutions to design problems which are computationally challenging to grid search. These research findings will now be used to test the goal-driven generative method, which will be described in Chapter 8, to improve the efficiency of generating optimised solutions to various design objectives. This will then be compared to the data-driven method outlined in Chapter 6.

Chapter 8 - Goal-Driven Generative Design for Additive Manufacturing

8.1. Introduction

The results from Chapter 7 show the potential to use adaptive surrogate models, in particular, Bayesian optimisation, to optimise the build orientation of structurally optimised AM parts. This chapter extends this to the design problem described in Chapter 5 by assessing goal-driven approaches to generative design for their ability to generate suitable geometry for high-level functionally optimised geometry. The results from the goal-driven generative design method will be compared against random search for its ability to converge on high-performing solutions. Secondly, the effect of changing the acquisition function will be explored, and differing numbers of iterations will be tested to determine the number of iterations required to locate high-performing solutions. Next, the design space domain will be extended to determine the ability for goal-driven methods to search larger design spaces than would be feasible with data-driven methods. Finally, the performance of solutions generated by both data-driven and goal-driven methods will be compared to answer the research question set out in Chapter 4.

8.2. Benchmarking Results against Random Search

To ensure the results from the Bayesian optimisation are converging on high median values of the objective function, the results must be benchmarked against random search. Based on the results from Chapter 7, tests are conducted for the Matèrn covariance and UCB acquisition functions. The results for each of the production scenarios described in the research methodology are run for 10 iterations of both Bayesian optimisation and random search. The results are repeated 21 times to generate a median value, as shown in Table 8-1.

Table 8-1: Output values for 21 iterations of random search and Bayesian optimisation for 10 iterations

Persona	Bayesian Optimisation (median)	Random (median)	Best result	p-value
Overall Performance	0.536	0.565	0.531	3.23×10^{-8}
Highest Structural Performance	0.341	0.427	0.316	9.82×10^{-5}
Highest Production Quantity	0.366	0.434	0.349	1.64×10^{-7}

8.3. Goal-Driven Generative Design for Additive Manufacturing

To evaluate the effectiveness of using Bayesian optimisation for locating high-performing solutions for different AM scenarios, different numbers of iterations were tested and evaluated against the multi-dimensional criteria outlined in the research methodology. The tests were repeated 21 times in order to attain a median estimate for the location of the best solution. The following graphs show the output with respect to the level 2 abstraction criteria. The red circles represent all of the solutions evaluated, the blue diamonds represent the best objective function value achieved for n-iterations, and the golden star represents the median value attained from each of the best values obtained from the 21 runs. The gold star may overlay a number of the blue diamonds if the same best value for the objective function is located on multiple runs. This value represents the most-likely value to be attained from any repeated test with n-iterations. All tests were performed using the Matèrn covariance function with a UCB acquisition function.

8.3.1. Ten Iterations of Random Search and Bayesian Optimisation

The following section contains the results for 10 iterations of both the random and Bayesian optimisation algorithms for each of the production scenarios outlined in Section 5.3.3.5.

8.3.1.1. High Overall Performance

Described by persona C in the research methodology, the following figures display the results for the overall performance metric, for both random search and Bayesian optimisation. The results show that Bayesian optimisation performs significantly better than random search for this performance metric, as shown in Table 8-1. Figure 8-1 shows the results from the random search scenario. The results show that whilst high-performing solutions are found, the high interquartile range, shown in the boxplots in Figure 8-3, means that these solutions are unlikely to be found unless many computational runs are performed. This may be computationally infeasible.

Figure 8-1: Cost-projection area and compliance-time graphs for 10 iterations of random search for the maximum overall performance metric.

Figure 8-2 shows the result from the Bayesian optimisation strategy, the results show a bias towards a low cost, low build projection solution. The best results are generated at 90°/-90° build orientation with no overhang constraint applied. The interquartile range of results is far smaller than the random case as shown in the boxplot in Figure 8-3.

Figure 8-2: Cost-projection and compliance-time plots for Bayesian optimisation using 10 iterations for the maximum overall performance production scenario.

Figure 8-3: Boxplots for maximum overall performance persona, showing 10 iterations of Bayesian optimisation and random.

8.3.1.2. High Performance Part Production

The persona B production scenario results are detailed in this section. Figure 8-4 shows the results from random search. The results show a large range in the best performing solutions, with the mean result performing poorly with respect to minimising structural performance and build time.

Figure 8-4: Cost-projection and compliance-time plots for using 10 iterations of random search with the aim of maximising part performance.

Figure 8-5 shows the results from 10 iterations of the Bayesian optimisation search strategy. The results clearly show a bias toward regions of low compliance and build time. However, whilst high-performing solutions were found in the 21 tests, the median solution was lower performing than would be satisfactory, indicating that a greater number of iterations is likely to be required to improve the median value.

Figure 8-5: Cost-projection and compliance-time plots for Bayesian optimisation using 10 iterations of UCB acquisition function with the objective of maximising part performance.

To measure the effect of increasing the number of iterations on the overall performance of the random search and Bayesian optimisation strategies, 15 and 20 iterations were tested for both methods. The following boxplots highlight the performance of the Bayesian optimisation when compared to random for 10, 15, and 20 function iterations. The results are illustrated in the boxplots in Figure 8-6. Whilst each set of tests show that the Bayesian optimisation outperforms random search, 20 iterations are required to provide a consistent median result at the highest performing value.

Figure 8-6: Trends in maximum overall part performance with the x-axis rotation set between [-90°, 90°] for different numbers of structural optimisation iterations.

Figure 8-7 shows the scatter plots comparing build cost and part projection area, part compliance and part build time for 20 iterations of Bayesian optimisation. The results show that there are a small number of solutions which deviate from the highest performing solution. The median sample is focused on the lowest compliance and lowest build time indicating that the search strategy has achieved the target from the objective function.

Figure 8-7 - Cost-projection and compliance-time plots for Bayesian optimisation using 20 iterations of UCB acquisition function with the objective of maximising part performance.

7.3.1.3. High Volume Part Production

The following section details the results from the highest volume part production scenario. As with the previous results, Figure 8-8 shows that the random search method at 10 iterations finds high-performing results with a large interquartile range (depicted in Figure 8-10), indicating its unreliability as a method to guarantee high-performing results.

Figure 8-8: Cost-projection and compliance-time plots for 10 iterations of random search for the highest part throughput production scenario.

Figure 8-9 shows the Bayesian optimisation results for 10 iterations. The results show a strong bias towards minimising build projection and build cost with the top performing solution found to be at -90° with a full manufacturability overhang constraint. The interquartile range of the Bayesian optimisation is also far smaller than that of the random search, as seen in Figure 8-10. It is notable that the best solutions for the highest overall part performance and the highest production quantity scenarios exist at the bounds of the design space domain. To show that the goal-driven generative design algorithm generalises, the design space will be modified so that the best performing solutions exist in a different location in the design space. This would ensure that the surrogate model would also find top performing solutions when they are not present at the bounds of the optimisation.

Figure 8-9: Cost-projection and compliance-time plots for Bayesian optimisation using 10 iterations of UCB acquisition function with the objective of maximising part throughput.

Figure 8-10: Boxplots for highest production quantity scenario for original design space problem.

8.3.2. Extending the Domain of the Design Space

The x-axis rotation angle is increased to, $\theta \in \mathbb{Z}: \theta \in [0, 359]$. This design space is important for two reasons. Firstly, it represents a design problem where the optimal solutions are not found at the edges of the design space domain. Secondly, this design problem is representative of the types of problems that would likely be infeasible using the grid search based approaches using in data-driven generative design for AM due to computational limitations.

8.3.2.1. High Volume Part Production

Initially, the experiment for highest production quantity was designed with a value of $\kappa = 3.0$, input to the acquisition function. The boxplot results of 21 iterations of this experiment are shown in Figure 8-11 with the output objective function value. The results demonstrate that there is not a statistically significant difference between the results from the Bayesian optimisation and the random search algorithm for either 15 iterations ($p = 0.232$) or 20 iterations ($p = 0.119$).
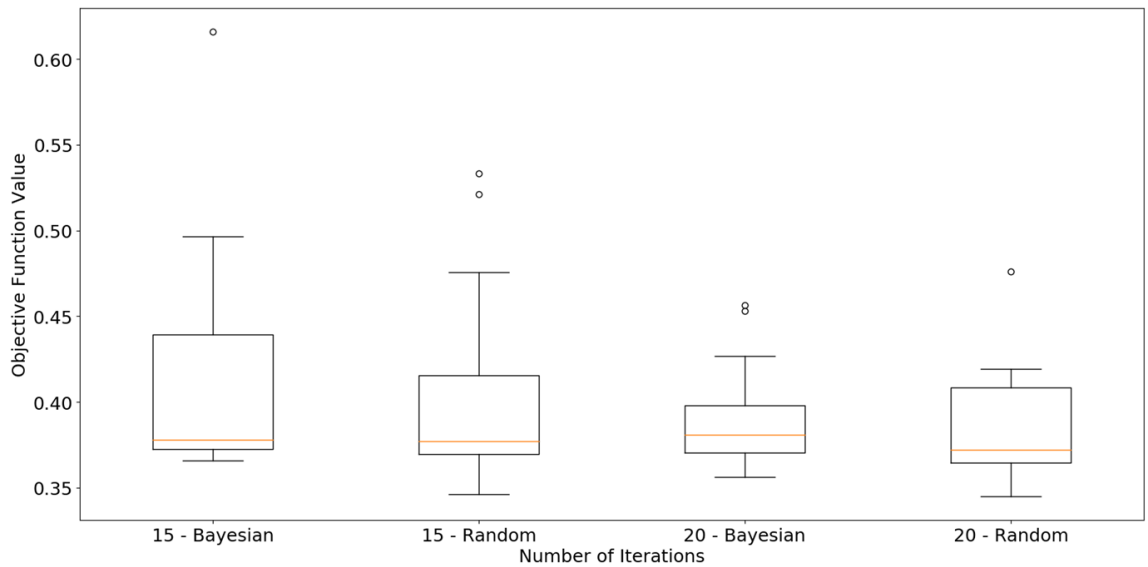
Figure 8-11: Boxplots for 15 and 20 iterations of Bayesian optimisation and random search for high production quantity persona with a UCB hyper parameter value $\kappa = 3.0$.

The following graph (Figure 8-12) shows the output from the Bayesian optimisation algorithm with 20 iterations. The results show the best results clustered around the area of lowest cost and build projection area, which is in line with the top performing results from the data-driven approaches. However, it is anticipated that this could be improved by altering the acquisition function hyper parameters with a preferential bias towards maximising the mean value.

Figure 8-12: Cost-projection and compliance-time plots for Bayesian optimisation using 20 iterations of UCB acquisition function with the objective of maximising part throughput.

Figure 8-13 depicts the contour plots for the Gaussian process output after 13 iterations. The results give an indication of the poor performance of the Bayesian optimisation. Despite the algorithm locating regions of high-performing solutions indicated with a high mean value, the

acquisition function strongly favours unexplored areas of the solution space. This process is repeated unsuccessfully as these are areas of infeasible solutions.
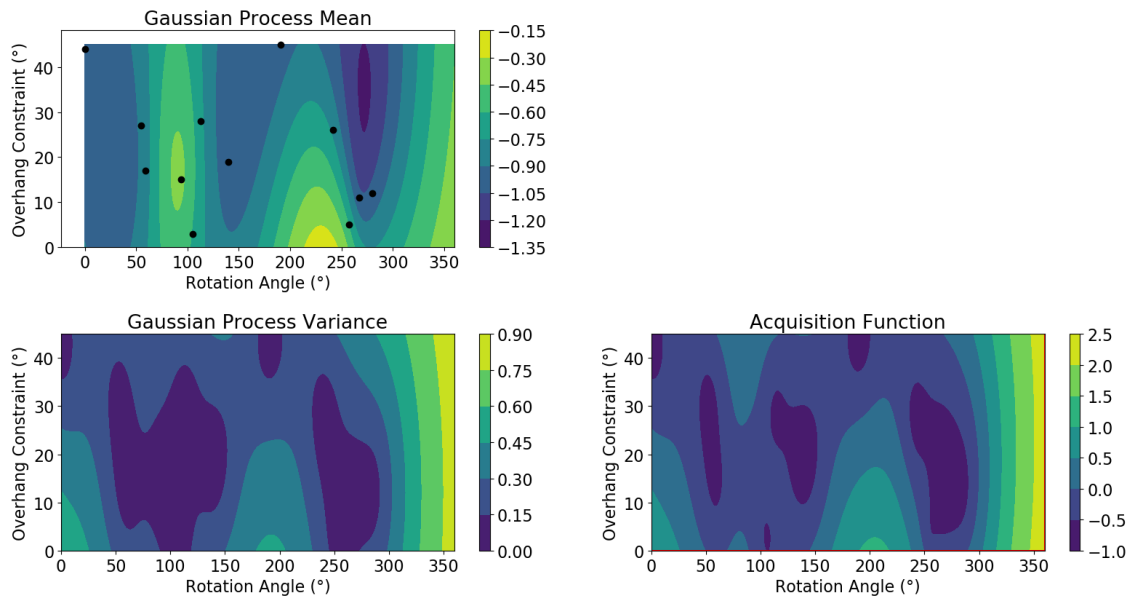
Figure 8-13: Plots for Gaussian process mean, variance and corresponding acquisition function value set with k=3.0

Figure 8-14 depicts the design space used within this problem. This plot was created by randomly sampling the design space. The areas in which no solutions are present within the plot represent areas of design space combinations that produce infeasible solutions. This is also represented in Figure 8-13 with the lighter coloured areas in the Gaussian process variance contour plot. This is attributed to high variance depicting unexplored regions of the design space. The top performing solutions are shown with blue diamonds at two distinct x-values. Discontinuities in functions can be challenging for Bayesian optimisation as they leave regions of high uncertainty which are repeatedly targeted by the Bayesian optimisation process if the hyper parameters of the acquisition function are set to favour areas of high variance. If the acquisition function favours the variance term too highly, the Bayesian optimisation will test the areas of infeasible solutions at each iteration. Upon finding an infeasible solution, the algorithm then samples a random set of design space parameters, which in turn leads to the solution space tending toward that of the random search.

Figure 8-14: Design Space for 20 iterations of random search for the highest production quantity scenario. Blue diamond's represent the best solution for each experimental run.

To improve the performance of the Bayesian optimisation, the hyper parameters of the acquisition function can be reduced to bias the output to areas of higher mean value rather than variance. To locate this value, the following plots were produced at each iteration and the values of the mean, variance and acquisition function are recorded. Figure 8-15 shows the Gaussian process contour plot after the hyper parameters have been selected that favour higher mean values rather than areas of higher variance. This agrees with the values from the design space in Figure 8-14. As the dimensionality of the design space increases, it becomes difficult to visualise the Gaussian process. However, the hyper parameters can still be selected either by manual search or alternatively, by combining the surrogate modelling with multi-dimensional visualisation techniques as shown in Wortmann (2017).
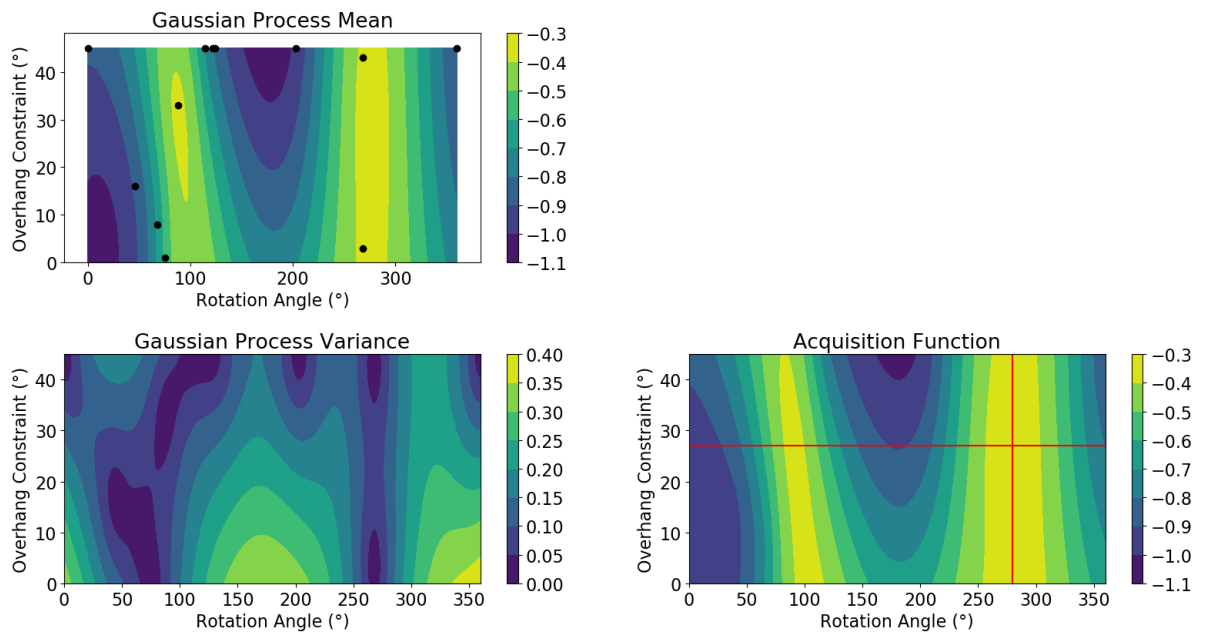
Figure 8-15: Plots for Gaussian process mean, variance and corresponding acquistion function value set with $k = 0.25$.

The results were rerun for the lower value of $\kappa = 0.25$, and the following boxplots, in Figure 8-16, depict the output results from this experiment. The results show that the reduced κ value leads to statistically significant improvement of Bayesian optimisation when compared to random search for both 15 iterations ($p = 1.34 \times 10^{-2}$) and 20 iterations ($p = 5.53 \times 10^{-3}$).

Figure 8-16: Boxplots for 15 and 20 iterations of Bayesian optimisation and random search for highest production quantity metric with a UCB hyper parameter value $\kappa = 0.25$.

The results after the parameter κ, is reduced are shown in Figure 8-17. As expected from the boxplot results, the variance in the data is much less and the performance values are closer to the data-driven generative design approach in Chapter 6.

Figure 8-17: Cost-projection and compliance-time plots for 20 iterations Bayesian optimisation for highest production quantity metric with a UCB hyper parameter value $\kappa = 0.25$

8.3.2.2. High Overall Part Performance

The results for the highest overall part performance case are detailed within this section. Figure 8-18 shows the result of 20 iterations with a $\kappa = 3.0$. As the overall performance case contains a tradeoff, it may be beneficial to have an acquisition function value which is biased towards the variance to encourage greater exploration of the design space. The results show that this works with high-performing solutions found for minimising each of the individual level 2 abstraction criteria.

Figure 8-18: Cost-projection and compliance-time plots for 20 iterations Bayesian optimisation for highest overall part performance metric with a UCB hyper parameter value $\kappa = 3.0$

An alternative scheme is to reduce the hyper parameter value in the acquisition function to $\kappa = 0.25$, in order to favour higher mean values. The results shown in Figure 8-19 highlight that this

scheme favours greater exploitation, showing a much lower variance of points surrounding the mean value. This method is arguably less useful to designers in learning the underlying structure of the design space and understanding any trade-offs that might occur within this space.

Figure 8-19: Cost-projection and compliance-time plots for 20 iterations Bayesian optimisation for highest overall part performance metric with a UCB hyper parameter value $\kappa = 0.25$

One way in which a greater diversity of solutions can be achieved is by showing solutions which are within a determined percentage from the minimum value of the objective function. Figure 8-20 shows the output when results are shown within 2% of the best observed value. These results are shown in green squares. The results show that if the experiment is run multiple times, this is an effective way of locating many possible feasible and satisfactory solutions, which can be a useful method for exploring and understanding various trade-offs in the design space.

Figure 8-20: Cost-projection and compliance-time plots for 20 iterations Bayesian optimisation for highest production quantity metric with a UCB hyper parameter value $\kappa = 0.25$ with the values increased to +/- 2% of the best observed value for each test.

By providing the designer with the fitted surrogate model, it is possible to create an interactive exploration of the design space. Figure 8-21 shows the Gaussian process output when fitted to 20

data points. The yellow regions within the Gaussian process mean plot highlight the top performing areas of the design space. This graph could be implemented as part of the interactive data visualisations to allow for a user defined design exploration.

Figure 8-21: Gaussian process mean, variance and acquisition function plots for the maximum overall performance design problem.

8.4. Comparison with Data-Driven Generative Design Methods

To compare data-driven and goal-driven generative design methods, the following plots were generated. Firstly, the normalised objective function values for each of the performance metrics were taken and sorted using the default sort function in Python. The maximum value is subtracted from each value and the absolute value is calculated, such that the best design has a value of 0. The same approach was calculated for 25th percentile, the median and 75th percentile values for goal-driven and overlaid on to the data-driven bar chart to show its relative position on the dataset. The closer the goal-driven values are to zero, the higher the solution quality.

8.4.1. Maximum Overall Performance

The median and 25th percentile solution distance for the highest overall performance metric is aligned with the 8th position in the data-driven solution space. The 75th percentile solution performs better with a value corresponding to the 7th position. This shows that the goal-driven solutions are most likely to be generated within the top 5% of the data-driven solutions within 20 iterations. However, as shown from the scatterplots in Figure 8-20, by expanding the located values to lie within 2% of the minimum value, a greater range of high performing solutions can be located. In this case, the findings are more consistent with the data-driven methods. When

comparing conflicting metrics within the objective functions, it is sensible to engineer a more exploratory search design problem which can then be optimised further by reframing the problem. In these scenarios, analysis of the Gaussian process surrogate can lead the designer to alternative regions of the design space, which may yield higher performing solutions for certain design metrics. The graph is shown in Figure 8-22. The blue bars indicate the objective values obtained by each design instance from the data-driven method. Coloured bars are overlaid that indicate the quartile 1, median and quartile 3 results obtained from the goal-driven method. Areas of interest are highlighted for clarity.
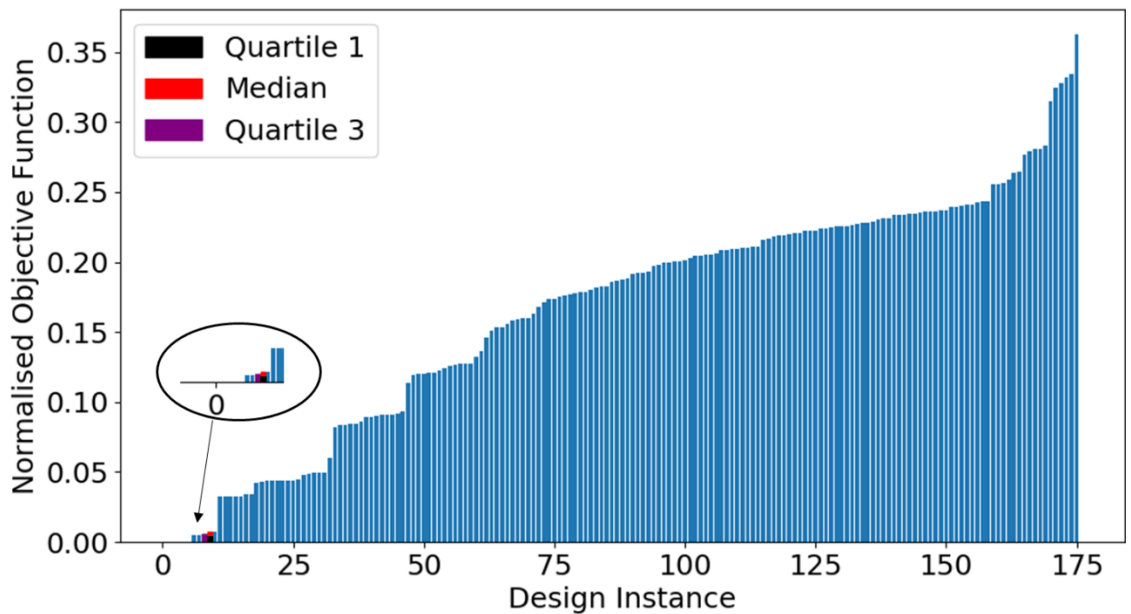
Figure 8-22: Normalised objective function distance metric comparing maximum overall performance median value for 20 iterations to each design instance created in the data-driven generative design dataset.

8.4.2. High Performance Part Production

The highest structural performance case, for the 20 iterations experimental test case, Q1, the median and Q3 all have a value of 0 indicating that the median and the interquartile range solutions are equivalent to the highest performing value from the data-driven approach. The graph showing its contextual position in the data-driven dataset is shown in Figure 8-23. The values from the goal-driven method cannot be seen on the graph due to their zero value.

Figure 8-23: Normalised objective function distance metric comparing structural performance median value for 20 iterations of each design instance created in the data-driven generative design dataset.

8.4.3. High Volume Part Production

The Q1 solution of goal-driven generative design is zero, indicating that this is equivalent to the best design in the data-driven dataset. The median value is shown in the 11[th] position and the Q3 result is at the 16[th] position. This indicates that the goal-driven method is most likely to find solutions within the top 10% of solutions in the data-driven dataset within 20 iterations. The small interquartile range also highlights the consistency of the goal-driven methods when the acquisition function is set with a bias toward the mean. The distance graph is shown in Figure 8-24.

Figure 8-24: Normalised objective function distance metric comparing highest production quantity performance median value for 20 iterations to each design instance created in the data-driven generative design dataset.

8.5. Key Observations

As with the results in Chapter 7, the results show that Bayesian optimisation is far more efficient at finding high-performing design space solutions than random search. The initial design space can have a significant impact on the performance of the Bayesian optimisation algorithm, where optimised results can be found within as few as 10 iterations in cases where the best solution exists in the boundaries of the input domain. This can be attributed to the fact that these regions are typically areas of high variance and will therefore be targeted by the Bayesian optimisation algorithm acquisition function.

As design problems become more complex and increase in dimensionality, the likelihood of the top solutions existing at the bounds of the design domain reduces. This is due to the curse of dimensionality (Bellman 2015). This states that as dimensionality increases, the probability of exploring a space for a given number of samples decreases exponentially in proportion to the

increased number of dimensions. Therefore, the designer's selection of the acquisition function hyper parameters becomes more important. For multi-dimensional problems tuning the hyper parameters can be challenging. However, based on the results in this chapter, if the design space contains discontinuities it is beneficial to select a hyper parameter value that is biased towards the mean as this prevents repeated tests within infeasible regions.

For design problems with conflicting design objectives, the designer can either select hyper parameters with a bias towards variance to create a greater diversity of solutions or alternatively include solutions that exist within a small percentage of the global maximum solutions. The results show that running the algorithm once can find high-performing solutions. Visualisation of the Gaussian process can also show other regions of the design space with high-performing solutions. The designer can select different regions of the design space in order to validate the accuracy of the surrogate model and generate further feasible solutions. This is a computationally efficient method for exploring the design space.

Comparison of the goal-driven and data-driven generative design for additive manufacturing approaches using the ranked normalised objective function distance metric shows that the 75th percentile solutions of the goal-driven method consistently find solutions in the top 10% of solutions found with the data-driven approach. The median values of goal-driven approaches sit within the top 5% of data-driven solutions for design personas B and C. The results from persona A are within at least 7%, even for an increased design space domain. This shows that goal-driven and data-driven design methods generate comparable design solutions according to the criteria defined in the research methodology.

In terms of computational efficiency, the goal-driven approaches only require 20 iterations for the same design problem. By comparing this with 175 solutions required for data-driven approaches and utilising equation 33, goal-driven methods are shown to be at least 8.75 times more efficient. The use of goal-driven approaches also allows for larger design spaces to be explored whilst still locating solutions comparable to those found using data-driven approaches. Therefore unless designers have access to a pre-existing dataset it is preferential to use goal-driven approaches over data-driven approaches if the design solutions are computationally challenging to generate.

8.6. Summary
The results of the goal-driven generative design for AM method has been implemented by using a Bayesian optimisation surrogate model to efficiently locate optimal solutions to a number of high-level evaluation criteria. The results were run using the same problem definition that was used to perform the data-driven method in Chapter 6 and comparable results were found between the

two methods. Goal-driven methods were shown to find the optimal solutions within as few as 10 iterations when compared to the dataset of 175 solutions in the case of data-driven generative design methods.

When the domain of the build orientation, and therefore, the size of the design space was increased, it was found that the quality of solutions found is dependent on the selection of hyper parameters input into the surrogate optimisation. The results show that when discontinuities are present in the design space then it is better to select hyper parameters which favour high mean values. When exploring trade-offs in the design problem, it can be beneficial to select hyper parameters that favour higher variance or to visualise the output from the surrogate and manually explore the predicted high-performing solutions.

The results between goal-driven and data-driven generative design methods were compared and the results demonstrate that goal-driven methods consistently find solutions in the top 10% of data-driven methods for the same design problem. Goal-driven design methods are shown to be at least 8.75 times more efficient at locating high-performing design solutions than data-driven methods. This highlights their suitability for use in generative design when calculating large datasets of solutions results in a high computational cost.

Chapter 9 will now provide a critical discussion of the results obtained in this research as well as any limitations of the research methodology.

Chapter 9 – Discussion

9.1. Introduction

In this chapter, the methods, observations and results from this research will be critically discussed. Discussions relate to the aim, objectives and limitations of the research and are later used to formulate the conclusions and future research.

9.2. Research Discussion

9.2.1. State-of-the-Art Literature Review and Generative CAD Framework

The review of the state-of-the-art literature identified a number of key research gaps that need to be addressed before the next generation of CAD tools for DfAM can be realised. The geometric design freedom provided by AM allows designers to design for intent, rather than DFMA limitations that constrain other traditional manufacturing technologies. However, the complexity of AM is often underestimated. Current DfAM tools often provide solutions that are infeasible to manufacture or are expensive to produce. This leads to a number of research gaps within state-of-the-art generative design that must be addressed with tools to aid designers in producing end-use, commercially-viable AM components.

The prominent research gaps include:

- Generative design synthesis methods. For example, topology optimisation has been shown to provide performance benefits when successfully applied to AM parts. Examples have shown potential weight savings of 57% (Emmelmann et al. 2011) and 60% (Tomlin and Meyer 2011) respectively for two aerospace brackets and, more recently, a 69% reduction in an vehicle airbrake hinge (Smith et al. 2016). Topology optimisation often leads to designs that require large amounts of support structure to manufacture. Liu (2018) states that between 40-70% of AM product cost can be attributed to support structure removal; therefore a large body of research has been dedicated to manufacturability constrained topology optimization. Promising studies (Langelaar 2018) have shown that this work can be extended by integrating build orientation into the manufacturability constrained optimisation to further increase the potential to create lightweight components and also to reduce part cost by reducing post-processing requirements. However, further work was required to evaluate the effectiveness of combining build orientation and manufacturing constrained generative synthesis on the development of 3D AM parts.

- Secondly, as topology optimisation has evolved to include manufacturing constraints, the output designs have moved from being impractical and cost prohibitive to fabricate, to producible with AM. However, the results do not consider the cost as part of the optimisation. It is advantageous to move away from low-level design objectives requirements traditionally found in topology optimisation results, such as minimising compliance or mass. Instead, topology optimisation should move towards high level objectives such as part throughput or part performance, incorporating considerations such as build time and cost. This would give designers a better understanding of the impact of design changes on the part development process, and ultimately increase the profitability of AM parts.

- Finally, research has shown that humans are susceptible to a number of cognitive biases that limit the use of CAD in the design process (Robertson and Radcliffe 2009). Additionally, humans can only process a maximum of four variables at once (Halford et al. 2005). By increasing the number of dimensions in which design instances are represented, design can be created that are not only judged on their appearance but also on their behaviour. To achieve the next generation of design tools it is necessary to explore methods that allow designers to easily comprehend and efficiently explore multi-dimensional design and solution spaces to generate high-performing, optimised AM part solutions.

Based on the research gaps identified in Chapter 2, a generalised CAD framework for generative design tools has been introduced containing the key features required of generative design systems. The purpose of this framework is to guide the development of generative design tools that can assist designers in producing high-quality manufacturable parts based on a series of pre-defined design goals and constraints. The research undertaken within this thesis explores two implementations of generative design systems drawn out of the framework, notably, data-driven and goal-driven generative design with a truss-based geometry synthesis method. However, many generative design methods and tools could be developed from this framework.

9.2.2. Adaptation of Structural Optimisation Incorporating Manufacturability and Build Orientation

A truss optimisation strategy was developed using the theory from Christensen and Klarbring (2009) and a ground structure pruning scheme was incorporated to eliminate potential trusses from the model that exceed a user-defined maximum overhang angle. A minimum strut thickness was also implemented, derived from design guidelines reviewed within existing literature. The

ability to rotate the ground structure is incorporated into the optimisation so that the build orientation of the part is incorporated within the structural optimisation. This ensures that the optimisation is contextualised within the layer-wise AM build process.

Within the last two years, industry has begun to integrate manufacturing constraints into commercial topology optimisation applications (Frustum 2018; Solid Thinking 2016). However, to the best of the authors' knowledge, it is only Autodesk Generative Design that aims to explore multiple manufacturing orientations into the optimisation. In addition, at the time of writing, this is still constrained to three orientations, which are 90° rotations with respect to the imported part orientation. The design problem used within this research consisting of a manufacturing constrained optimisation and additionally, incorporating the possibility to optimise for all build angles is representative of industrial state-of-the-art. The results from this research show that by modelling all build angles in the design space higher performing solutions can be obtained that would have otherwise been possible. Only a single research article (Langelaar 2018) incorporates a similar level of flexibility however, this article still only provides 2D results that do not represent, or even inspire real manufacture-ready geometries.

9.2.3. Abstracting from Low to High-Level Generative Design Objectives

A series of high-level design goals have been generated by abstracting and combining a set of evaluation metrics across four stages. The first stage combines low-level facet data with information within the design space. From this, it is possible to generate part based evaluation criteria, including the compliance, part volume, build height, build projection area and support requirements.

The next abstraction stage combines these criteria or reinterprets the criteria to make it clearer for the end user. An example of this is translating build projection area into build plate packing. It also allows for greater generalisation should further optimisations such as build plate nesting algorithms be integrated into the generative design system. Finally, the highest level of design abstraction includes a series of production AM typical production scenarios. The purpose of this abstraction is to allow for better interpretability of multi-dimensional evaluation criteria.

While scope of this research dictates that the generative design methods are only applicable to design problems that can be formed into a measurable objective function. The use of abstraction criteria allows the designer to quickly formulate multiple different design abstractions to test. In viewing design in its teleological sense, designers must *purposefully* make decisions based on current information to predict future states (Jones 1992). Therefore, design abstractions may be used, even within 'wicked' design problems, to make informed decisions about the importance of

a number of design objectives, reducing the uncertainty surrounding the initial problem. This links back to the Dorst and Cross (2001) view of design as co-evolution of the problem and solution space where the authors state that further information about a problem can only be found by actively searching for it.

As with any model, the designer must take great care to ensure that the weighted abstraction model is representative of the design problem that they are trying to solve. Designers must take great care to ensure that the computational model that is optimised maps directly onto the real-world design problem. It is likely that experimental validation of the model will be required to ensure that the weighted abstraction model correctly maps onto the actual design task therefore safeguarding against optimising for false proxy measures. Multi-criteria optimisation models have been developed for AM in previous research studies (Brika et al. 2017). However, these have not been combined with a structural optimisation and abstracted into metrics representing high-level design objectives such as different production scenarios. To, the best of the authors knowledge, this research is first attempt at incorporating the conflicting guidelines from the so-called 'iron triangle' (Atkinson 1999), exploring trade-offs between cost, quality and time which is a constant challenge and guides many industrial production strategies. As such, this research allows designers to begin understanding the trade-offs that occur in the design process when design constraints are applied, and directly map part geometry onto production scenarios, reducing the time and cost in developing AM parts.

9.2.4. Data-Driven Generative Design Method
A data-driven method was defined using MCDA TOPSIS that was used to locate the highest performing solutions within a user-generated solution space according to a series of specified abstracted design goals. Data-driven methods present a novel way of analysing CAD models and are particularly compelling in generative design due to the n-dimensional representation of design solutions. The MCDA was shown to be capable of locating high-performing solutions for each of the three design scenarios.

Research has shown that designers struggle to deal with large numbers of design concepts (Finke et al. 1992). To overcome this, an interactive data visualisation dashboard has been generated depicting the output generated solution space for user exploration. Based on previous research (Theus 2008; Ashour and Kolarevic 2015), a series of data visualisations were selected for their ability to present the design data back to the user. The three stages of visualisations are intended to allow the designers to quickly make sense of the multi-dimensional datasets that are necessary for generative design systems by assessing the location of generated points with respect to other

positions within the solution space. The creation of an interactive, manual solution space exploration tool can be considered equivalent to the state-of-the-art commercial implementations of generative design tools as described in the literature review. The solution space dashboard allows designers to find trends and correlations in the design data. Additionally, the use of dashboards decreases the time designers take to comprehend multi-dimensional design representations reducing the challenges that occur when evaluating multiple design concepts.

Previous research has shown that designers struggle to understand the relationship between geometric freedom and concepts such as build cost and time (Pradel, Bibb, et al. 2018). By integrating simulation into the n-dimensional part representations in generative design, it is feasible to partially overcome this gap. One of the key outputs from Chapter 6 is the importance of a full understanding of the design limitations of the material and process combination. This is because it is possible to create higher structurally performing solutions when overhang constraints are not applied, however, this comes at the detriment of post-process time. Furthermore, the research shows that the indiscrimindate use of manufacturing constraints during generative solution synthesis may lead to suboptimal solutions depending on the overarching design objectives. This provides two critical findings: firstly, the importance of defining design objectives clearly at the start of the design process. Secondly, it is important to provide clear feedback to the user about the results after each set of design solutions is generated. This allow the user to alter the design objectives if sensible design solutions are not being produced. In this sense data-driven generative design methods align with the view that computers act as design assistants (Schon 1992), aiding the designer in uncovering trends, while automating repetitive task such as generating geometry. However, it is the human designer that makes decisions on the next design task that must be performed.

Additionally, the results have shown that small pairwise changes (10%) in the weights can alter the top performing solutions returned from the TOPSIS MCDA. Therefore, it is recommended that designers test different input weights to understand underlying trade-offs within the design objectives and also to yield a more diverse range of feasible solutions.

If multiple search strategies are run for the same design problem, the results should be recorded to augment the data-driven algorithms. This gives the advantage of being able to adapt the algorithm for different goals. This suggests that even when using goal-driven methods, it would be beneficial to record all low-level evaluation criteria results. This data can be used to find high-

performing solutions to different design objectives within a dataset by creating different abstractions from the low-level data.

A disadvantage to using data-driven methods is the requirement to generate solution spaces that are computationally expensive when complex modelling and simulation strategies are required. Additionally, by generating large numbers of poor-performing solutions, the sensitivity of search is reduced. This makes locating solutions using the data-visualisation dashboard more challenging. However, the use of MCDA reduced the impact of this phenomenon.

Based on this information, it was deemed appropriate to pursue more efficient exploration methods that aim to minimise the number of expensive structural optimisation evaluations required to find the optimal design space parameters, i.e. build orientation and manufacturing constraint combination.

9.2.5. Surrogate Optimisation Method for Build Orientation Optimisation

Optimising the part orientation can have a significant impact on the overall cost of a part, simply by reducing the support structure requirements (Zwier and Wits 2016). However, support structures can be costly to compute (Morgan et al. 2016). The results from Chapter 7 demonstrate that a support length measure is an appropriate proxy for evaluating the total amount of support structure required for the parts tested within this research. The results for the two test parts show alignment within 3.5% of commercial software estimations while being more computationally efficient.

In addition, using surrogate optimisation models has the potential to reduce the number of iterations required to locate high-performing solutions to design problems. By testing two open-access continua-based topology optimised brackets, it was found that support structure length and Bayesian optimisation were the most effective method at locating the optimal build orientation. Thirty repeats of each test were performed to ensure robustness in the statistical analysis methods. The results were benchmarked against random search to show the efficacy of Bayesian optimisation for efficiently minimising objective functions.

Secondly, grid search is shown to be an inferior method for exploring solution spaces unless the resolution for each of the input parameters is small. Random is shown to perform significantly better than grid search. However, Bayesian optimisation is shown to outperform both other methods, in particular, grid search up to a factor of 17.

The selection of brackets developed using continuum based generative synthesis methods highlights that the surrogate optimisation methods used within the research can be generalised to

different synthesis techniques. As continuum based solutions are commonly used within commercial software, this generalisation increases the potential to integrate this research into existing commercial topology optimisation tools.

A further result of Chapter 7 is the added emphasis on the importance of combining build orientation with structural optimisation. Evidence of this is provided by showing that the total volume of material (part and support structure) for the GE bracket can be reduced by 16.2% from the as-optimised direction solely by optimising the build orientation. This demonstrated considerable potential in reducing costs associated with post-processing and support removal for AM components.

9.2.6. Goal-Driven Generative Design Method

Based on the success of the results from the surrogate based build optimisation, Bayesian optimisation was used as a method to minimise the number of expensive structural optimisation evaluations required to locate the highest performing areas of the solution space.

The Bayesian optimisation was demonstrated to locate top performing solutions when compared to data-driven methods in 10 iterations when the solutions fall on the bounds of the design spaces domain. This can be attributed to the high variance associated with these bounds, therefore, they are likely to be targeted by the Bayesian optimisation algorithm. When the top solutions were not located on the bounds of the domain, as with the maximum part performance case, 20 iterations were required to find solutions comparable to the best solution found in the data-driven solution space.

The size of the design space was increased to between 0° and 360° to ensure the optimal solutions could be located for the other two production scenarios. It is common for there to be discontinuities in the response surface output from the design space as there is going to be infeasible material and load combinations for certain ground structures. If there are discontinuities in the response surface function, then it can be challenging to locate the optimal solutions as it is impossible to explore these areas and they are always areas of high variance. As such the acquisition function may continuously try and test these areas leading to the output solutions tending toward a random search.

The use of hyper parameters in the acquisition function allows the designers to modify the balance between exploration and exploitation of the Bayesian optimisation algorithm (Brochu et al. 2010). Favouring high variance indicates greater exploration. This is often risky as it requires an investment of time that does not always equate to good solutions, however, the reward may be large. Favouring the mean reduces risk but may lead to locating local optima. Depending on the

design objective, designers may make different decisions when deciding whether to favour exploring or exploiting.

By modifying the hyper parameters in the acquisition function to favour the mean rather than the variance, it is possible to locate the optimal solutions even with discontinuities in the response surface. The hyper parameters can be determined by creating visualisations of the Gaussian process surrogate; for high-dimensional design spaces, the work within this research could be extended by using techniques developed by Wortmann (Wortmann 2017).

Gaussian process surrogates, as with many machine learning models, perform best when there are underlying trends within the design space landscape that can be modelled. This means that overly stochastic design spaces will be difficult to learn, and are unlikely to yield better results than random search. Whilst it is possible to build noise into the Gaussian process regression model to deal with some stochasticity in the system (Rasmussen 2004), if the generative synthesis method yields unpredictable results, other goal-driven generative design methods may be more suitable.

In order to answer the research aim a distance metric was created based on the normalised objective functions for each of the three performance personas detailed in Chapter 5. The results show that the goal-driven method consistently finds high-performing solutions, with the median results being either close (top 7%) or exactly the same, when compared to the data-driven methods. The best performing solution in the goal-driven case consistently locates the top solution in the data-driven dataset, even when the size of the design space is increased.

One of the benefits of the increased efficiency associated with Bayesian optimisation is the ability to optimise larger design spaces. This includes extending the bounds of the input variables or increasing the dimensionality of the space. This research has not explored large multi-dimensional design spaces and, therefore, it is not clear how many iterations would be required to locate optimal solutions within these computationally challenging generative design spaces. However, surrogate optimisation strategies have consistently been shown to be the most appropriate strategy for efficiently locating optimal solutions when function evaluations are costly (Wortmann et al. 2015). Therefore, it is suggested that high-performance computation could be combined with surrogate optimisation strategies such as goal-driven generative design to extend this work to complex multi-dimensional design spaces.

By improving the efficiency of locating optimed design space parameters, there is potential to increase the number of input dimensions in the design space. This leads to potential for future

research in optimising machine parameters alongside the build orientation and manufacturing constraints within the design problem. Further maximising the potential advantages of goal-driven generative design in industrial AM.

While it may appear that the computer is designing parts, human designers are making *all of* the creative decisions in the development of goal-driven generative design solutions. This approach to design alleviates many of the problems associated with premature fixation, circumscribed thinking (Robertson and Radcliffe 2009) and design fixation (Crilly 2015). However, the designer now faces new challenges when using generative design methods. These include: engineering the correct design space, selecting the generative synthesis method, choosing hyper parameters for surrogate models and performing the correct statistical techniques to locate the best performing solutions. New approaches to design education are required to further understand and overcome these issues. The results in Chapter 8, show the potential of generative design systems to efficiently find high-performing solutions to well-defined design spaces. This highlights the collaborative nature of generative design systems that combine the strengths of human creativity and computational speed within the design process.

Arguably, generative design methods shift the issues associated with cognitive biases into the selection of generative design methods rather than the individual part design. As designers make decisions about the synthesis methods and parameters of a design problem, they are implicitly shrinking the design space into a computationally tractable domain. However, many satisfactory designs may exist externally to the design space, in a so-called creative space that designers manually navigate, as outlined in Figure 9-1. As more synthesis methods are used, a greater span of the creative space can be covered by the design space and a greater diversity of solutions can be generated. This is attributed to the fundamental principle that algorithms follow a set of rules and the solutions derived from algorithms must exist within this rule base. Expert designers have been shown to break rules when creating innovative solutions (Cross 2004), therefore, designers must be aware of the limitations of generative design tools (and the methods used to realise them) and use them in conjunction with other creative practices to fully maximise their potential.

Figure 9-1: Schematic demonstrating how the design space acts as a subset of the creative space that designers navigate when solving wicked design problems.

The use of topology optimisation as a generative design synthesis method reduces diversity within the solution space as the fundamental principles of structural mechanics lead to solutions that display certain physical characteristics (i.e. perpendicular beams). For a more diverse exploration, designers may select more exploratory synthesis methods, for example, shape grammars (Antonsson and Cagan 2005).

This research supports previous studies that describe design optimisation as a method to develop further understanding problems rather than setting out to find the objective 'best' solution (Chen et al. 2015b; Bradner et al. 2014). By visualising the output from the design optimisation, both in the visualisation dashboard and in the Gaussian process contour plots, the designer is able to review parameter bounds and restate problems in order to continue the search for high-performing design solutions.

It cannot be said, therefore, that generative design is suitable for all design problems. Problems that cannot be mathematically defined within the design space remain outside the capability of both of the generative design methods outlined in this research. Conversely, goal-driven generative design has been found to be successful when solving problems that can be formulated with an objective function. As such, designers may benefit from using generative design methods as part of a 'toolkit' of approaches to generate ideas.

9.3. Generalising the Research Findings

Although the work presented in this thesis is focused on the methods described in Chapter 5, throughout the research, decisions were made to ensure that the research could be generalised. Furthermore, experimental findings have shown that there are areas in which the research can be

extended outside the scope of the methods outlined in Chapter 4. Firstly, the results from Chapter 7, show that goal-driven generative design methods for AM are suitable for continuum based structural synthesis methods as well as the truss based solutions shown in Chapter 8. Secondly, the results show that different combinations of level 2 abstractions can be combined and minimised. It is expected that further evaluation criteria could be added to the objective function to extend the research outcomes further. Thirdly, the laser-PBF process was used for each of the equations for the abstracted evaluation criteria. However, this research can be readily extended for other AM processes as long as the abstraction criteria (build cost, time, etc.) can be defined for these processes. Finally, the CAD framework illustrates that there are many ways of defining problems, synthesising solutions and visualising those solutions when creating generatively designed parts for AM. However, further research is required to validate these potential implementations.

9.4. Research Limitations

The work provided in synthesising geometries using data-driven and goal-driven generative design methods provides the beginnings of future design systems to support DfAM, with geometry outputs acting as inspiration to designers rather than the final solutions. There is still considerable research to be undertaken before industrial AM becomes a click-and-print solution. Therefore, the following section outlines a set of limitations associated with the research methodology.

9.4.1. Solely Quantitative Evaluation Criteria

As outlined in the literature review, the complexity of the design process means that it is often impossible for designers to explain why certain features of a design are important. For example, one person may deem a location to be cosy whereas another may describe it as being cramped. It is this so-called tacit knowledge that is very challenging to model using computers. Additionally, quantification bias can occur when users value measurable data over the immeasurable. As such, designers must be careful to ensure that correlations within data are not treated as causation. Consequently, without including qualitative design requirements into generative design tools, their use in the design of functional end-use products will be limited for many design scenarios. However, this does not mean that the use of quantitative generative design tools cannot be used to strongly influence and augment the ability for a professional design engineer to design fully functional, high-performing parts for AM. In the future it may be possible to add further levels of abstraction, to further define objectives that represent the reasons why designers are creating parts, and may include, profit, sustainability or aesthetics, as seen in Figure 9-2. Research will have to be undertaken to determine how to develop the low-level evaluation criteria and the abstraction mappings to facilitate design for qualitative objectives.
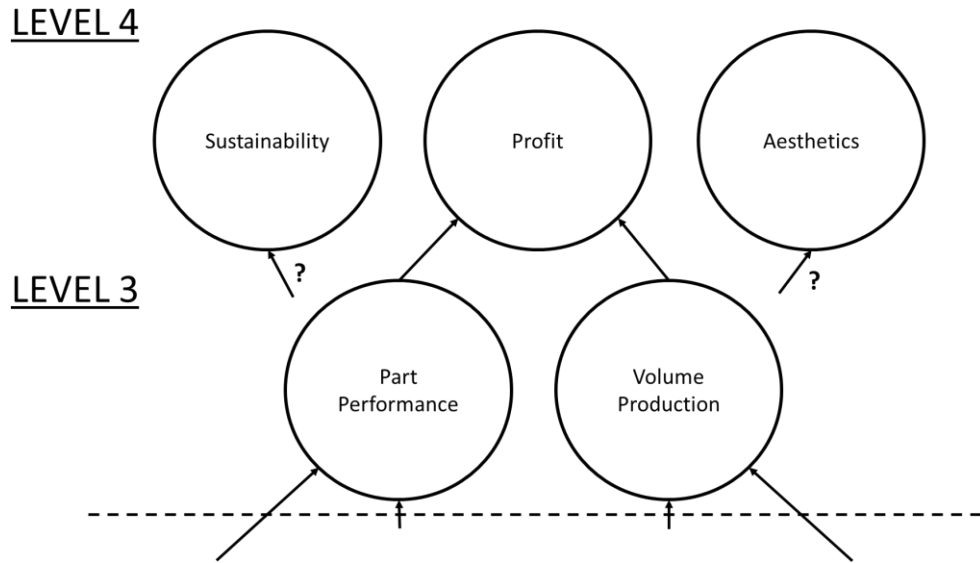
Figure 9-2: Expanding to level 4 abstraction criteria with the addition of qualitative evaluation criteria.

9.4.2. Manufacturability Criteria

Throughout this research, there has been an assertion that the ability to successfully manufacture a part was determined by the overhang angle and minimum diameter of each of the struts. Whilst these are important manufacturability constraints, typically, design guidelines consist of many more constraints, including aspect ratios, bridge lengths, minimum hole size etc. In reality, these must all be followed to ensure a successful build. Therefore, it is impossible to suggest that a design will be completely manufacturable on any process using solely thin wall and overhang constraints. Furthermore, the work undertaken in Chapter 7, could be tested on a more extensive set of geometries to generalise further than just topology optimisation based generative synthesis methods. For example, those based on meso- or micro- structures. Currently, there are limited existing research articles that extend additive manufacturing constraints further than these criteria. However, the methodology is extensible to add generative synthesis methods that can incorporate further manufacturing constraints.

9.4.3. Truss Node Valence and Stress Concentrations

The decision to utilise ground structure topology optimisation was made for its ability to perform exceptionally well when the desired output volume is a small percentage of the starting DSV and also the computational efficiency of using beam elements in the FEA. However, one of the limitations with using beam elements within the model is the failure to accurately model the as manufactured parts. Firstly, as the beam elements are rendered using 3D tube elements in the THREE.js framework, when two elements meet there is a discontinuity in the modelling, and a continuum structure is not formed, as shown in Figure 9-3. To overcome this, the elements are exported and solidified in Meshmixer before the manifold mesh file is created for 3D printing. This

could have been alternatively performed within the synthesis algorithm itself, either by using spheres at the nodes, as shown in Smith et al (2016), by solidifying the wireframe (Srinivasan et al. 2005) or by using an implicit representation of the struts (Angles et al. 2017).

Figure 9-3: Optimised truss structure showing areas with unconnected struts.

A further manufacturing issue that is not taken into account during the structural optimisation is the number of struts entering a node. When large numbers of struts enter a single node, the stress concentration at that node increases as there are no fillet radii defined at the node-strut intersections. This would suggest that manufacturing the as-manufactured trusses will have a lower stiffness when compared with the output from the truss optimisation. Therefore it is recommended that the output continuum structure as exported from Autodesk Meshmixer should be re-meshed and analysed using a 3D FEA model for structural verification.

The case study was selected for its simplicity and the ability to easily understand the output geometry. However, it should be noted that for more complex DSV's, it may be significantly more challenging to determine a suitable ground-structure and spatial distribution of nodes. More complex implementations of ground-structures should be realised to enable generalisation to all possible design space volumes.

9.4.4. Multiphysics Optimisation Tools

With industry trends showing a proliferation of metal AM (Wohlers et al. 2018), it is necessary to ensure that optimisation models can accurately model the materials and processes that are used within this field. The layered manufacturing process and the complex thermal phenomena that occur during AM builds often leads to material anisotropy in the build parts. This needs to be understood before full confidence in the structural integrity of the as-built part can be realised.

In order to comprehensively assess the manufacturability of a part, more complex simulations must be undertaken that take into account other important factors such as thermal deformations and residual stresses. Furthermore, it is impossible to assess the effect of build plate packing on

the manufacturability of a part without thermal simulation. Although it is possible that parts can be manufactured as single parts, scaling to a full build plate can have detrimental effects on the part quality due to changes in the tool/scan path and layer cooling times. This is particularly important when designing for the abstracted objectives as outlined in this thesis.

By including multiphysics simulation tools, it would also be possible to extend the design space to include the machine parameters, for example, hatching distance and laser power. This is due to the fact that the model interactions between build time, cost and thermal distortion could accurately be modelled. This would give further credence to following a goal-driven generative design approach due to the limitations of increasing the dimensionality of design space with data-driven generative design methods.

Some phenomena cannot be simulated and it may be necessary to perform empirical experimentation in order to fill knowledge gaps in understanding the design problem. It may be necessary to include hybrid methods, which incorporate experimentation and simulation in order to fully represent the AM problems that designers face. The methodology in this research is, however, fully extensible to incorporate these aspects and could be extended to include accurate AM build simulation models.

Chapter 10 - Conclusions and Future Work

10.1. Introduction

This chapter presents the conclusions obtained from this research along with the overall contribution to knowledge. Potential areas in which this research could be further investigated are covered in the final section of this chapter.

10.2. Conclusions

As a result of the research completed in this thesis, the following conclusions can be made:

- Inspired by the research gaps identified in the state-of-the-art literature review, a generalised CAD framework is defined to guide future development of generative design tools to support DfAM. The framework recognises the collaborative nature of generative design, highlighting that the best design work will be undertaken using a combination of human creativity and computational speed.

- An existing ground structure topology optimisation technique was extended from 2D to 3D and two AM specific manufacturing constraints were applied, namely maximum overhang angle and minimum strut thickness. This provided an estimation of the structural performance of a part whilst ensuring that it remains manufacturable with excessive support structure requirements.

- A series of abstraction criteria were defined that map low-level geometric evaluation criteria associated with the synthesised geometry onto high-level design criteria. For example, business and production criteria, providing simple, and easily interpretable design objectives for designers.

- A data-driven generative design approach using a TOPSIS MCDA, was defined and used to locate the top performing solutions in a predefined solution space. This solution space was defined by using grid search at 5° on the parametric design space inputs, and include a manufacturability constraint on the maximum overhang angle and the rotation angle about the x-axis.

- A novel method for optimising build orientation based on the minimisation of support structure has been defined and tested on two open-access AM parts that are typical of

part outputs from the topology optimisation process. It is performed using a Bayesian optimisation with the aim of minimising the number of evaluations required to locate the best orientation. The results show that in 35 iterations, the method can more efficiently find the same build orientations as current state-of-the-art commercial software.

- A goal-driven generative design strategy was created utilising Bayesian optimisation surrogate optimisation and is then applied to the same problem defined in the data-driven approach for comparison. Experiments show that the highest performing set of design parameter can be found with 20 structural optimisation evaluations.

- The results of the data-driven and goal-driven generative design for AM methods were compared and the results show that the more efficient method of goal-driven generative design was capable of locating solutions within the top 7% of solutions found using the data-driven methods in 8.75 times fewer iterations. This demonstrated its capability to generatively design structural components specifically for AM.

10.3. Contributions to Knowledge

The following contributions to knowledge have been realised as a result of this research:

- A goal-driven generative design for AM method comprised of a Bayesian optimisation surrogate was used to efficiently locate optimal solutions for different production scenarios. The results demonstrated that optimal solutions could be found within 20 iterations which is 8.75 times more efficient than data-driven methods. In addition, this is also extensible to larger design spaces that would not be possible for data-driven methods due to the curse of dimensionality.

- A data-driven generative design approach to AM was implemented using a TOPSIS MCDA algorithm to find the top performing solutions within a solution space defined using grid-search, in three different AM production scenarios. Data-driven approaches were shown to be an effective method to easily navigate solution spaces and understand trade-offs that occur in multi-dimensional AM evaluation criteria.

- A novel surrogate model using Bayesian optimisation and Gaussian processes is implemented as an efficient method used to find build orientations that minimise the need for support structure for two commonly used topologically optimised parts.

- A generalised CAD Framework was developed providing the basis for developing future generative design tools for AM. Many possible generative design methods can be derived

from the framework including variations on the methods outlined in Chapter 5. Examples

include modifying the surrogate optimisation models in goal-driven generative design or

using alternative MCDA techniques in data-driven generative design.

- A novel data visualisation dashboard was defined using multiple different data
visualisation methods to allow designers to explore the trade-offs and correlations that
occur between multi-dimensional AM solution space data.

10.4. Future Research

In addition to the research gaps established into the literature review, the work undertaken in this

thesis has produced knowledge that can be extended. The research limitations discussed in

Chapter 9 should be addressed to further generalise the research methodology. However, farther

reaching perspectives will be provided in this section that could be used to extend the impact of

the work provided within this thesis.

10.4.1. User Testing of Generative Design Systems

To ensure that future generative design CAD tools will be successful in aiding users in designing

parts for AM, research must be undertaken to examine the optimal user interaction within the

system. User testing should be carried out examining the efficacy of using interactive visualisation

dashboards, for example the layout developed in Chapter 6, in examining multi-dimensional

design solutions. Secondly, an evaluation of potential user interfaces should be explored to

ascertain the best method for improving design solutions within the user feedback loop defined

within the CAD framework in Chapter 4. Finally, research should be carried out examining

methods that can be utilised for incorporating qualitative metrics and extracting tacit user

knowledge to generate improved design solutions.

10.4.2. Qualitative Evaluation Metrics and Recommendation Systems

As explained in Chapter 9, a limitation of this research, is the lack of qualitative design evaluation

metrics. Therefore, future work should focus on integrating qualitative design aspects such as,

aesthetics. In order to achieve this, the feedback loop between the human designers in the

visualisation stage of the generative design framework needs to be exploited. It is believed that

conscious methods including recommendation systems, similar to those used to suggest products

to users in e-commerce marketplaces as well as unconscious techniques such as user eye tracking

could be used to ascertain the most favourable designs from a solution space based solely on

aesthetics or other qualitative design considerations. Early examples of this technique can be

explored in recent studies, for example Bylinskii et al (2017) and Mothersill and Bove Jr (2015).

This problems remains an open challenge in generative design as often designers have conflicting

opinions on similar metrics with research required in the user interface and user experience fields of generative design.

10.4.3. Generative Design Marketplaces

Practically ubiquitous internet access, increased computational power and improved mobile connectivity have allowed for an ever more connected design and manufacturing industry that has been termed industry 4.0. To exploit the generative synthesis of AM parts, it is possible to integrate generative design into a marketplace system. A generative design marketplace will consist of a generative design engine. This exists as a manifestation of a generative design approach derived out of the generative design CAD tool framework outlined in Chapter 3 and a generative design marketplace that sits on top of this engine. A schematic of this structure can be seen in Figure 10-1.

The generative design marketplace is likely to consist of a webpage in which users can purchase information that can be used to improve the performance of the algorithms within the generative design engine. This may include design guidelines for material and machine combinations, alternative generative synthesis algorithms, statistical analysis methods (i.e. TOPSIS or Bayesian optimisation) or available manufacturing machines, akin to a manufacturing hub.

In the marketplace, the machine availability will be presented containing information surrounding the location and cost of the machines. Should a machine be in use, the generative design system could modify the output geometry to be specifically created for the second-best machine-material combination in the list. This process is repeated until a satisfactory solution is met.

The information stored within the generative design marketplace can be manifested in multiple different database architectures. Two possibilities include centralised or decentralised architectures. Centralised databases are commonplace and are easy to implement. However, a single entity would have control over the marketplace. Decentralised architectures, on the other hand, often rely of peer-to-peer sharing of information and blockchain-like security and give users greater control over the data that they share in the marketplace. The advantages and disadvantages of both architectures would have to be evaluated as part of this future research.

A further advantage of the marketplace is that is democratises the process of generating information for the marketplace. This increases the number of shareholders in the design and manufacturing process. Consider, a small enterprise that can accurately define the manufacturing constraints for a given process; this information could be traded on the generative design marketplace bringing in revenue for small companies doing research and also improving the quality of the parts output from the market.
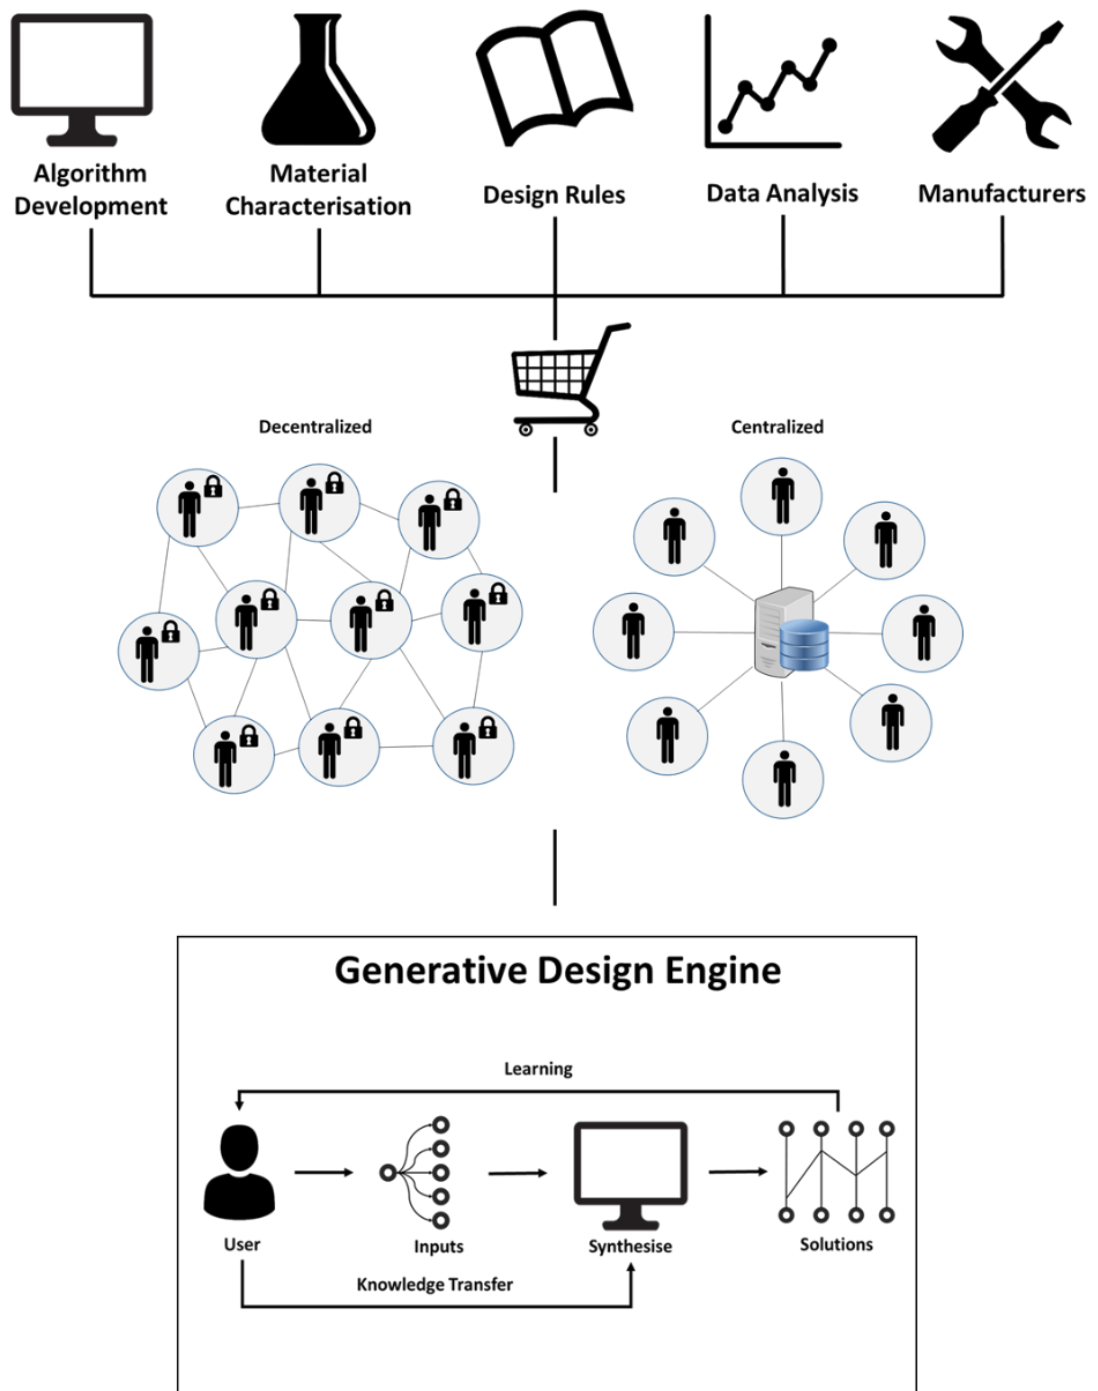
Figure 10-1: Schematic of the generative design marketplace structure.

10.4.4. Extension to Different Manufacturing Technologies

For generative design to become applicable to different industries, this research should be

extended to include a wider array of manufacturing processes. Even with advances in

manufacturing speed from new state-of-the-art AM techniques such as single pass jetting from

Desktop Metal (Desktop Metal 2018), the process remains unsuitable for production volumes greater than 100,00 parts.

Extending generative design to other manufacturing techniques involves understanding the process and constraints and encoding these constraints into the generative synthesis algorithms. This is challenging as tool access algorithms will have to be incorporated into synthesis algorithms for subtractive processes such as milling and turning. In addition, machining fixation points would have to be taken into account.

The combination of AM and investment and sand casting is becoming popular due to the relative ease of casting a vast range of metal alloys, and the ability to certify cast parts easier than those produced by AM (Schwab 2017). To achieve the production volumes required by industries such as commercial automotive, further research must be conducted to explore the production trade-offs that occur when designing for casting.

The ability for computers to synthesise shapes for multiple manufacturing processes will be extremely beneficial to designers as it has the potential to considerably reduce the development time for complex parts with conflicting design specifications.

10.4.5. Extending Generative Design with Machine Learning
This research has utilised a form of machine learning, namely Bayesian optimisation to improve the state-of-the-art in design space exploration in generative design for AM. However, there is scope for future research incorporating other machine learning techniques into the generative design process. Using generative design allows greater exploitation of machine learning due to the large amounts of data that is produced when generating many solutions with multi-dimensional design and solution spaces. This could include using machine learning for generating near real-time feedback during the design exploration and using multi-modal inputs to influence the generation of stylised or themed parts.

i)   *Real-time Generation of Design Instances*

A primary motivation of using Bayesian optimisation as a surrogate within this research was due to the computationally expensive process of performing structural optimisation simulations. This means that real-time exploration of generative design solution spaces is very challenging as exploring new areas of the solution space requires at least one further expensive computation to fit another point within the surrogate model. It would be beneficial if new methods could be developed that learn the functional mapping between the parametric design space and the output geometry from the structural optimisation. One possible method for achieving this is

through the use of generative machine learning models. These models are a subtype of statistical modelling techniques and are not related to the term 'generative design' used elsewhere in this thesis. The aim of generative models is to learn the joint probability distribution between an input and output. Popular examples include variational autoencoders (Doersch 2016) and generative adversarial networks (Goodfellow et al. 2014). There are many different types of generative models that aim to learn the true data distribution of the training set so as to generate new data points with some variations.

Recent advances in deep learning research have made it possible to learn complex functions by training neural networks on large datasets. Currently, one of the main limitations of generative design systems is their reliance on complex optimisation routines. Future research should be undertaken to attempt to map the inputs and outputs of the optimisation in order to gain real time estimates of geometry outputs from topology optimisation programs. In order to achieve this, a data set will still have to be generated. Whilst this is slow, it could be achieved using supercomputers or cloud computing clusters speeding up the generation of a sufficiently large dataset to learn the underlying generative model. Once the model is created the user could select any set of design space parameters and the model would output an approximation of the geometry as the output of the synthesis algorithm. Examples of these ideas have already been successful in the fields of material synthesis (Zsolnai-Fehér et al. 2018) and computational fluid dynamics (Umetani and Bickel 2018).

ii)   *Multi-modal Inputs to Improve User Experience when Generating Parts*

Another advancement in machine learning is the ability to transfer stylistic elements from multiple input mediums. An interesting application of this has been in style transfer (Gatys et al. 2016), here the user inputs two images. The first being the content image, and the second the style image containing stylistic elements that are to be transferred to the initial image. An example of this is shown in Figure 10-2.

Figure 10-2: Example of 2D image style transfer (Gatys et al. 2016).

At the centre of style transfer is a type of machine learning model, typically used for dealing with images, termed a convolutional neural network. Style transfer aims to minimise the error between the features in the content image and the mixed image and also to minimise the error between the style features of the style image and the mixed image. Future investigations could extend this research to 3D design. In this case the user could input a series of designs that they already favour, and use the style transfer algorithm to generate a series of new alternatives that are similar to the initial inputs.

A critical limitation of using deep-learning for 3D geometry is the requirement for large datasets of parts to train the algorithms. Databases of 3D geometry must be extended and made publicly available, through using crowdsourcing platforms, if we are to see the successful uptake of generative machine learning models extended to production ready software.

The future work has shown that this research has great potential to be further extended. It is hoped that the research provided within this thesis, alongside the suggestions for future can provide a basis for the development of next generation CAD tools that will support DfAM.

References

3D Hubs, 2017. *3D Hubs* [Online]. Available from: www.3dhubs.com [Accessed 30 August 2017].

3D Hubs, 2018. *How to design parts for Metal 3D printing* [Online]. Available from: https://www.3dhubs.com/knowledge-base/how-design-parts-metal-3d-printing#design [Accessed 13 June 2018].

Aage, N., Andreassen, E., Lazarov, B.S., and Sigmund, O., 2017. Giga-voxel computational morphogenesis for structural design. *Nature*, 550(7674), p.84.

Abdelall, E., Frank, M.C., and Stone, R., 2018. A study of design fixation related to Additive Manufacturing [Online]. *Journal of Mechanical Design*. Available from: http://dx.doi.org/10.1115/1.4039007.

Additive Industries, 2017. *Metal Fab1* [Online]. Available from: https://additiveindustries.com/systems/metalfab1 [Accessed 12 September 2018].

Airbus Group, 2016. *Pioneering bionic 3D printing - Learning from nature* [Online]. Available from: http://www.airbusgroup.com/int/en/story-overview/Pioneering-bionic-3D-printing.html [Accessed 14 June 2016].

Airbus Group;, 2018. *Airbus Commercial Aircraft delivers record performance* [Online]. Available from: https://www.airbus.com/newsroom/press-releases/en/2018/01/airbus-commercial-aircraft-delivers-record-performance.html [Accessed 28 October 2018].

Al-Ahmari, A.M., Abdulhameed, O., and Khan, A.A., 2018. An automatic and optimal selection of parts orientation in additive manufacturing. *Rapid Prototyping Journal*, (just-accepted), p.0.

Allaire, G., Jouve, F., and Toader, A.-M., 2002. A level-set method for shape optimization. *Comptes Rendus Mathematique*, 334(12), pp.1125–1130.

Allen, S. and Dutta, D., 1994. On the computation of part orientation using support structures in layered manufacturing. In: *Proceedings of solid freeform fabrication symposium, university of texas at austin, austin, TX, June*. pp.259–269.

Angles, B., Tarini, M., Wyvill, B., Barthe, L., and Tagliasacchi, A., 2017. Sketch-based implicit blending. *ACM Transactions on Graphics (TOG)*, 36(6), p.181.

Antonsson, E.K. and Cagan, J., 2005. *Formal engineering design synthesis*. Cambridge University

Press.

Aranda, E. and Bellido, J.C., 2016. Introduction to Truss Structures Optimization with Python. *Electronic Journal of Mathematics & Technology*, 10(1).

Ashour, Y. and Kolarevic, B., 2015. Optimizing creatively in multi-objective optimization. In: *Proceedings of the Symposium on Simulation for Architecture & Urban Design*. Society for Computer Simulation International, pp.128–135.

ASTM, 2012. *Standard A F2792 - Standard terminology for additive manufacturing technologies*.

Atilola, O., Tomko, M., and Linsey, J.S., 2015. The effects of representation on idea generation and design fixation: A study comparing sketches and function trees [Online]. *Design Studies*. Available from: http://www.sciencedirect.com/science/article/pii/S0142694X15000939 [Accessed 10 December 2015].

Atkinson, R., 1999. Project management: cost, time and quality, two best guesses and a phenomenon, its time to accept other success criteria [Online]. *International Journal of Project Management*, 17(6), pp.337–342. Available from: https://www.sciencedirect.com/science/article/pii/S0263786398000696 [Accessed 5 October 2018].

Atzeni, E. and Salmi, A., 2012. Economics of additive manufacturing for end-usable metal parts [Online]. *The International Journal of Advanced Manufacturing Technology*, 62(9–12), pp.1147–1155. Available from: http://dx.doi.org/10.1007/s00170-011-3878-1.

Autodesk, 2018a. Autodesk Meshmixer [Online]. Available from: www.meshmixer.com.

Autodesk, 2018b. *How GM and Autodesk are using generative design for vehicles of the future* [Online]. Available from: http://blogs.autodesk.com/inthefold/how-gm-and-autodesk-use-generative-design-for-vehicles-of-the-future/ [Accessed 6 August 2018].

Ayres, C., 2015. *Algorithmic Archery* [Online]. Available from: http://www.core77.com/posts/43879/Algorithmic-Archery [Accessed 30 August 2017].

Bächer, M., Whiting, E., Bickel, B., and Sorkine-Hornung, O., 2014. Spin-it: optimizing moment of inertia for spinnable objects. *ACM Transactions on Graphics (TOG)*, 33(4), p.96.

Baumers, M., Holweg, M., and Rowley, J., 2016. *The economics of 3D Printing: A total cost perspective* [Online]. Available from: https://www.sbs.ox.ac.uk/sites/default/files/research-projects/3DP-RDM_report.pdf.

Baumers, M., Tuck, C., Wildman, R., Ashcroft, I., Rosamond, E., and Hague, R., 2012. Combined build-time, energy consumption and cost estimation for direct metal laser sintering. In: *From Proceedings of Twenty Third Annual International Solid Freeform Fabrication Symposium—An Additive Manufacturing Conference*.

Bellman, R.E., 2015. *Adaptive control processes: a guided tour*. Princeton university press.

Bendsøe, M.P., Ben-Tal, A., and Zowe, J., 1994. Optimization methods for truss geometry and topology design [Online]. *Structural optimization*, 7(3), pp.141–159. Available from: https://doi.org/10.1007/BF01742459.

Bendsøe, M.P. and Kikuchi, N., 1988. Generating optimal topologies in structural design using a homogenization method. *Computer methods in applied mechanics and engineering*, 71(2), pp.197–224.

Bendsøe, M.P. and Sigmund, O., 2003. *Topology optimization : theory, methods and applications*. O. (Ole) Sigmund 1966-, ed. Berlin: Berlin.

Bergstra, J. and Bengio, Y., 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb), pp.281–305.

BMW Group, 2018. *Additive Manufacturing: 3D printing to perfection* [Online]. Available from: https://www.bmw.com/en/innovation/3d-print.html [Accessed 22 September 2018].

Booth, J.W., Alperovich, J., Chawla, P., Ma, J., Reid, T.N., and Ramani, K., 2017. The Design for Additive Manufacturing Worksheet [Online]. *Journal of Mechanical Design*, 139(10), pp.100904–100909. Available from: http://dx.doi.org/10.1115/1.4037251.

Boothroyd, G., Dewhurst, P., and Knight, W.A., 2011. *Product design for manufacture and assembly*. 3rd ed. P. Dewhurst & W. A. (Winston A. Knight 1941-, eds. Boca Raton, Fla.: Boca Raton, Fla.

Brackett, D., Ashcroft, I., and Hague, R., 2011. Topology optimization for additive manufacturing. In: *Proceedings of the Solid Freeform Fabrication Symposium, Austin, TX*. pp.348–362.

Bradner, E., Iorio, F., and Davis, M., 2014. Parameters tell the design story: ideation and abstraction in design optimization. In: *Proceedings of the Symposium on Simulation for Architecture & Urban Design*. Society for Computer Simulation International, p.26.

Breiman, L., 2001. Random forests. *Machine learning*, 45(1), pp.5–32.

Bridgman, P.W., 1922. *Dimensional analysis*. Yale University Press.

Brika, S.E., Mezzetta, J., Brochu, M., and Zhao, Y.F., 2017. Multi-Objective Build Orientation Optimization for Powder Bed Fusion by Laser [Online]. *Journal of Manufacturing Science and Engineering*, (50732), p.V002T01A010. Available from: http://dx.doi.org/10.1115/MSEC2017-2796.

Brochu, E., Cora, V.M., and De Freitas, N., 2010. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*.

Buchanan, R., 1992. Wicked problems in design thinking. *Design issues*, 8(2), pp.5–21.

Von Buelow, P., 2012. ParaGen: Performative Exploration of generative systems. *Journal of the International Association for Shell and Spatial Structures*, 53(4), pp.271–284.

Bylinskii, Z., Kim, N.W., O'Donovan, P., Alsheikh, S., Madan, S., Pfister, H., Durand, F., Russell, B., and Hertzmann, A., 2017. Learning visual importance for graphic designs and data visualizations. In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, pp.57–69.

Byrd, R.H., Lu, P., Nocedal, J., and Zhu, C., 1995. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5), pp.1190–1208.

Cadman, J.E., Zhou, S., Chen, Y., and Li, Q., 2013. On design of multi-functional microstructural materials [Online]. *Journal of Materials Science*, 48(1), pp.51–66. Available from: http://dx.doi.org/10.1007/s10853-012-6643-4.

Cagan, J., Campbell, M.I., Finger, S., and Tomiyama, T., 2005. A Framework for Computational Design Synthesis: Model and Applications [Online]. *Journal of Computing and Information Science in Engineering*, 5(3), pp.171–181. Available from: http://dx.doi.org/10.1115/1.2013289.

Campbell, I., Bourell, D., and Gibson, I., 2012. Additive manufacturing: rapid prototyping comes of age. *Rapid prototyping journal*, 18(4), pp.255–258.

Campbell, M.I. and Shea, K., 2014. Computational design synthesis. *AI EDAM*, 28(3), pp.207–208.

Carstensen, J. V and Guest, J.K., 2018. Projection-based two-phase minimum and maximum length scale control in topology optimization [Online]. *Structural and Multidisciplinary Optimization*. Available from: https://doi.org/10.1007/s00158-018-2066-4.

Carter, W.T., Erno, D.J., Abbott, D.H., Bruck, C.E., Wilson, G.H., Wolfe, J.B., Finkhousen, D.M., Tepper, A., and Stevens, R.G., 2014. The GE Aircraft Engine Bracket Challenge: An Experiment in

Crowdsourcing for Mechanical Design Concepts. In: *Solid Freeform Fabrication Symposium, Austin, TX*. pp.1402–1411.

Casakin, H.P., 2007. Metaphors in design problem solving: Implications for creativity. *International Journal of Design*, 1(2).

Catchpole-Smith, S., Aboulkhair, N., Parry, L., Tuck, C., Ashcroft, I.A., and Clare, A., 2017. Fractal scan strategies for selective laser melting of 'unweldable' nickel superalloys [Online]. *Additive Manufacturing*, 15, pp.113–122. Available from: https://www.sciencedirect.com/science/article/pii/S221486041630358X [Accessed 6 July 2018].

Caterino, N., Iervolino, I., Manfredi, G., and Cosenza, E., 2009. Comparative analysis of multi-criteria decision-making methods for seismic structural retrofitting. *Computer-Aided Civil and Infrastructure Engineering*, 24(6), pp.432–445.

Chakrabarti, A., Shea, K., Stone, R., Cagan, J., Campbell, M., Hernandez, N.V., and Wood, K.L., 2011. Computer-based design synthesis research: an overview. *Journal of Computing and Information Science in Engineering*, 11(2), p.21003.

Challis, V.J., Guest, J.K., Grotowski, J.F., and Roberts, A.P., 2012. Computationally generated cross-property bounds for stiffness and fluid permeability using topology optimization. *International Journal of Solids and Structures*, 49(23), pp.3397–3408.

Chen, K.W., Janssen, P., and Schlueter, A., 2015a. Analysing Populations of Design Variants Using Clustering and Archetypal Analysis. *Proceedings of the 33rd eCAADe Conference*, 1, pp.251–260.

Chen, K.W., Janssen, P., and Schlueter, A., 2015b. Analysing Populations of Design Variants Using Clustering and Archetypal Analysis. *Proceedings of the 33rd eCAADe Conference*, (September), pp.251–260.

Cheng, W., Fuh, J.Y.H., Nee, A.Y.C., Wong, Y.S., Loh, H.T., and Miyazawa, T., 1995. Multi-objective optimization of part-building orientation in stereolithography. *Rapid Prototyping Journal*, 1(4), pp.12–23.

Chiandussi, G., Codegone, M., Ferrero, S., and Varesio, F.E., 2012. Comparison of multi-objective optimization methodologies for engineering applications [Online]. *Computers & Mathematics with Applications*, 63(5), pp.912–942. Available from: http://www.sciencedirect.com/science/article/pii/S0898122111010406.

Chien, S.-F. and Flemming, U., 2002. Design space navigation in generative design systems [Online]. *Automation in Construction*, 11(1), pp.1–22. Available from: http://www.sciencedirect.com/science/article/pii/S0926580500000844 [Accessed 16 May 2016].

Chowdhury, S. and Anand, S., 2016. Artificial neural network based geometric compensation for thermal deformation in additive manufacturing processes. In: *ASME 2016 11th International Manufacturing Science and Engineering Conference*. American Society of Mechanical Engineers, p.V003T08A006-V003T08A006.

Christensen, P.W. and Klarbring, A., 2009. Sizing Stiffness Optimization of a Truss. In: P. W. Christensen & A. Klarbring, eds. *An Introduction to Structural Optimization*. Dordrecht: Springer Netherlands, pp.77–95. Available from: https://doi.org/10.1007/978-1-4020-8666-3_5.

Christiansen, A.N., Schmidt, R., and Bærentzen, J.A., 2015. Automatic balancing of 3D models [Online]. *Computer-Aided Design*, 58, pp.236–241. Available from: https://www.sciencedirect.com/science/article/pii/S0010448514001614 [Accessed 22 January 2018].

Crilly, N., 2015. Fixation and creativity in concept development: The attitudes and practices of expert designers [Online]. *Design Studies*, 38, pp.54–91. Available from: http://www.sciencedirect.com/science/article/pii/S0142694X15000137 [Accessed 6 August 2015].

Cross, N., 2004. Expertise in design: an overview. *Design studies*, 25(5), pp.427–441.

Das, P., Chandran, R., Samant, R., and Anand, S., 2015. Optimum Part Build Orientation in Additive Manufacturing for Minimizing Part Errors and Support Structures [Online]. *Procedia Manufacturing*, 1, pp.343–354. Available from: http://linkinghub.elsevier.com/retrieve/pii/S2351978915010410 [Accessed 23 June 2017].

Dayarathna, M., Wen, Y., and Fan, R., 2016. Data center energy consumption modeling: A survey. *IEEE Communications Surveys & Tutorials*, 18(1), pp.732–794.

Deaton, J.D. and Grandhi, R. V, 2014. A survey of structural and multidisciplinary continuum topology optimization: post 2000 [Online]. *Structural and Multidisciplinary Optimization*, 49(1), pp.1–38. Available from: https://doi.org/10.1007/s00158-013-0956-z.

Desktop Metal, 2018. *3D Printing at Scale* [Online]. Available from: https://www.desktopmetal.com/products/production/.

Despeisse, M. and Minshall, T., 2017. Skills and Education for Additive Manufacturing: A Review of Emerging Issues BT - Advances in Production Management Systems. The Path to Intelligent, Collaborative and Sustainable Manufacturing. In: H. Lödding, R. Riedel, K.-D. Thoben, G. von Cieminski, & D. Kiritsis, eds. Cham: Springer International Publishing, pp.289–297.

Dhokia, V., Essink, W.P., and Flynn, J.M., 2017. A generative multi-agent design methodology for additively manufactured parts inspired by termite nest building [Online]. *CIRP Annals - Manufacturing Technology*, 66(1), pp.153–156. Available from: http://linkinghub.elsevier.com/retrieve/pii/S0007850617300392 [Accessed 10 August 2017].

Ding, D., Pan, Z., Cuiuri, D., and Li, H., 2014. A tool-path generation strategy for wire and arc additive manufacturing [Online]. *The International Journal of Advanced Manufacturing Technology*, 73(1–4), pp.173–183. Available from: http://dx.doi.org/10.1007/s00170-014-5808-5.

Doersch, C., 2016. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*.

Dong, G., Tang, Y., and Zhao, Y.F., 2017. A Survey of Modeling of Lattice Structures Fabricated by Additive Manufacturing [Online]. *Journal of Mechanical Design*, 139(10), pp.100906–100913. Available from: http://dx.doi.org/10.1115/1.4037305.

Dorn, W.S., Gomory, R.E., and Greenberg, H.J., 1964. Automatic design of optimal structures. *Journal de Mecanique*, 3, pp.25–52.

Dorst, K. and Cross, N., 2001. Creativity in the design process: co-evolution of problem–solution [Online]. *Design Studies*, 22(5), pp.425–437. Available from: http://www.sciencedirect.com/science/article/pii/S0142694X01000096 [Accessed 13 January 2015].

Dumas, J., Hergel, J., and Lefebvre, S., 2014. Bridging the Gap: Automated Steady Scaffoldings for 3D Printing [Online]. *ACM Trans. Graph.*, 33(4), p.98:1--98:10. Available from: http://doi.acm.org/10.1145/2601097.2601153.

Duncker, K. and Lees, L.S., 1945. On problem-solving. *Psychological monographs*, 58(5), p.i.

Eggenberger, T., Oettmeier, K., and Hofmann, E., 2018. Additive Manufacturing in Automotive Spare Parts Supply Chains – A Conceptual Scenario Analysis of Possible Effects. In: M. Meboldt & C. Klahn, eds. *Proceedings of Additive Manufacturing in Products and Applications - AMPA2017*. Cham: Springer International Publishing, pp.223–237.

Emmelmann, C., Sander, P., Kranz, J., and Wycisk, E., 2011. Laser Additive Manufacturing and Bionics: Redefining Lightweight Design [Online]. *Physics Procedia*, 12, pp.364–368. Available

from: http://www.sciencedirect.com/science/article/pii/S1875389211001258 [Accessed 19 November 2015].

Essink, W.P., Flynn, J.M., Goguelin, S., and Dhokia, V., 2017. Hybrid Ants: A New Approach for Geometry Creation for Additive and Hybrid Manufacturing [Online]. *Procedia CIRP*, 60, pp.199–204. Available from: http://linkinghub.elsevier.com/retrieve/pii/S2212827117300239 [Accessed 10 August 2017].

Finke, R.A., Ward, T.B., and Smith, S.M., 1992. Creative cognition: Theory, research, and applications.

Fishburn, P.C., 1967. Letter to the editor—additive utilities with incomplete product sets: application to priorities and assignments. *Operations Research*, 15(3), pp.537–542.

Flynn, J.M., Shokrani, A., Newman, S.T., and Dhokia, V., 2015. Hybrid Additive and Subtractive Machine Tools-Research and Industrial Developments [Online]. *International Journal of Machine Tools and Manufacture*, 101, pp.79–101. Available from: http://www.sciencedirect.com/science/article/pii/S0890695515300894 [Accessed 28 November 2015].

Formlabs, 2015. *Formlabs Design Guide*.

Forrester, A.I.J. and Keane, A.J., 2009. Recent advances in surrogate-based optimization. *Progress in aerospace sciences*, 45(1–3), pp.50–79.

Frank, D. and Fadel, G., 1995. Expert system-based selection of the preferred direction of build for rapid prototyping processes. *Journal of Intelligent Manufacturing*, 6(5), pp.339–345.

Frank, M.C., Wysk, R.A., and Joshi, S.B., 2004. Rapid planning for CNC milling—A new approach for rapid prototyping [Online]. *Journal of Manufacturing Systems*, 23(3), pp.242–255. Available from: https://www.sciencedirect.com/science/article/pii/S0278612504800372 [Accessed 28 June 2018].

Fricke, G., 1999. Successful approaches in dealing with differently precise design problems [Online]. *Design Studies*, 20(5), pp.417–429. Available from: https://www.sciencedirect.com/science/article/pii/S0142694X99000186 [Accessed 29 October 2018].

Frustum, 2018. *Frustum Generate* [Online]. Available from: www.frustum.com [Accessed 30 August 2017].

Fryazinov, O., Vilbrandt, T., and Pasko, A., 2013. Multi-scale space-variant FRep cellular structures

[Online]. *Computer-Aided Design*, 45(1), pp.26–34. Available from:
https://www.sciencedirect.com/science/article/pii/S0010448511002405 [Accessed 2 July
2018].

Gatys, L.A., Ecker, A.S., and Bethge, M., 2016. Image style transfer using convolutional neural
networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
pp.2414–2423.

Gaynor, A.T. and Guest, J.K., 2016. Topology optimization considering overhang constraints:
Eliminating sacrificial support material in additive manufacturing through design [Online].
Structural and Multidisciplinary Optimization, 54(5), pp.1157–1172. Available from:
https://doi.org/10.1007/s00158-016-1551-x.

Geng, X., Chu, X., and Zhang, Z., 2010. A new integrated design concept evaluation approach based
on vague sets. *Expert Systems with Applications*, 37(9), pp.6629–6638.

Ghouse, S., Babu, S., Van Arkel, R.J., Nai, K., Hooper, P.A., and Jeffers, J.R.T., 2017. The influence of
laser parameters and scanning strategies on the mechanical properties of a stochastic porous
material [Online]. *Materials & Design*, 131, pp.498–508. Available from:
https://www.sciencedirect.com/science/article/pii/S026412751730624X [Accessed 25 October
2018].

Ghouse, S., Babu, S., Nai, K., Hooper, P.A., and Jeffers, J.R.T., 2018. The influence of laser
parameters, scanning strategies and material on the fatigue strength of a stochastic porous
structure [Online]. *Additive Manufacturing*, 22, pp.290–301. Available from:
https://www.sciencedirect.com/science/article/pii/S2214860418301313 [Accessed 25 October
2018].

Gibson, I., Rosen, D., and Stucker, B., 2014. *Additive Manufacturing Technologies: 3D Printing, Rapid
Prototyping, and Direct Digital Manufacturing*. Springer.

Gibson, M.A., Myerbeg, J.S., Fulop, R., Verminski, M.D., Fontana, R.R., Schuh, C.A., Chiang, Y.-M., and
Hart, A.J., 2017. Fabricating an Interface Layer for Removable Support [Online]. Available from:
https://patentimages.storage.googleapis.com/88/83/15/a0011645565356/US9833839.pdf.

Goguelin, S., Flynn, J.M., Essink, W.P., and Dhokia, V., 2017. A Data Visualization Dashboard for
Exploring the Additive Manufacturing Solution Space. *Procedia CIRP*, 60, pp.193–198.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and
Bengio, Y., 2014. Generative adversarial nets. In: *Advances in neural information processing*

systems. pp.2672–2680.

Gordon, E.R., Shokrani, A., Flynn, J.M., Goguelin, S., Barclay, J., and Dhokia, V., 2016. A Surface Modification Decision Tree to Influence Design in Additive Manufacturing. In: *Sustainable Design and Manufacturing 2016*. Springer, pp.423–434.

GrabCAD, 2013. *GE Jet Engine Bracket Challenge* [Online]. Available from: https://grabcad.com/challenges/ge-jet-engine-bracket-challenge [Accessed 4 May 2018].

GrabCAD, 2016. *Airplane Bearing Bracket Challenge* [Online]. Available from: https://grabcad.com/challenges/airplane-bearing-bracket-challenge [Accessed 4 May 2018].

GrabCAD, 2017. *GrabCAD Library* [Online]. Available from: www.grabcad.com [Accessed 29 August 2017].

Greco, S., Figueira, J., and Ehrgott, M., 2016. *Multiple criteria decision analysis*. Springer.

Griffiths, L., 2017. Carbon on mass 3D printing the Adidas Futurecraft 4d show [Online]. *TCT Magazine*. Available from: https://www.tctmagazine.com/tct-events/tct-show-uk/carbon-mass-3d-printing-adidas-futurecraft-4d-shoe/.

Haftka, R.T. and Gürdal, Z., 2012. *Elements of structural optimization*. Springer Science & Business Media.

Hague, R., Mansour, S., and Saleh, N., 2003. Design opportunities with rapid manufacturing [Online]. *Assembly Automation*, 23(4), pp.346–356. Available from: http://dx.doi.org/10.1108/01445150310698643.

Halford, G.S., Baker, R., McCredden, J.E., and Bain, J.D., 2005. How many variables can humans process? *Psychological science*, 16(1), pp.70–76.

Hallihan, G.M. and Shu, L.H., 2013. Considering confirmation bias in design and design research. *Journal of Integrated Design and Process Science*, 17(4), pp.19–35.

Hammond, J.S., Keeney, R.L., and Raiffa, H., 1998. The hidden traps in decision making. *Harvard business review*, 76(5).

Haselton, M.G., Nettle, D., and Murray, D.R., 2015. The evolution of cognitive bias. *The handbook of evolutionary psychology*, pp.1–20.

Hayes, J.R., 2013. *The complete problem solver*. Routledge.

Hildreth, O.J., Nassar, A.R., Chasse, K.R., and Simpson, T.W., 2016. Dissolvable Metal Supports for 3D

Direct Metal Printing [Online]. *3D Printing and Additive Manufacturing*, 3(2), pp.90–97. Available from: https://doi.org/10.1089/3dp.2016.0013.

Hoover, S.P., Rinderle, J.R., and Finger, S., 1991. Models and abstractions in design. *Design Studies*, 12(4), pp.237–245.

Hopkinson, N., Hague, R., and Dickens, P., 2006. *Rapid manufacturing: an industrial revolution for the digital age*. John Wiley & Sons.

HP, 2018. *HP Metal Jet* [Online]. Available from: http://www8.hp.com/us/en/printers/3d-printers/metals.html [Accessed 12 September 2018].

Huang, R., Riddle, M., Graziano, D., Warren, J., Das, S., Nimbalkar, S., Cresko, J., and Masanet, E., 2016. Energy and emissions saving potential of additive manufacturing: the case of lightweight aircraft components [Online]. *Journal of Cleaner Production*, 135, pp.1559–1570. Available from: https://www.sciencedirect.com/science/article/pii/S0959652615004849 [Accessed 28 October 2018].

Huang, X., Ye, C., Wu, S., Guo, K., and Mo, J., 2009. Sloping wall structure support generation for fused deposition modeling [Online]. *The International Journal of Advanced Manufacturing Technology*, 42(11–12), pp.1074–1081. Available from: http://dx.doi.org/10.1007/s00170-008-1675-2.

Hur, J. and Lee, K., 1998. The development of a CAD environment to determine the preferred build-up direction for layered manufacturing [Online]. *The International Journal of Advanced Manufacturing Technology*, 14(4), pp.247–254. Available from: https://doi.org/10.1007/BF01199879.

Inui, M., Nagano, S., and Umezu, N., 2018. Fast computation of accessibility cones for assisting 3+ 2 axis milling. *Computer-Aided Design and Applications*, 15(5), pp.667–676.

Jansson, D.G. and Smith, S.M., 1991. Design fixation [Online]. *Design Studies*, 12(1), pp.3–11. Available from: http://www.sciencedirect.com/science/article/pii/0142694X9190003F [Accessed 9 November 2015].

Jhabvala, J., Boillat, E., André, C., and Glardon, R., 2012. An innovative method to build support structures with a pulsed laser in the selective laser melting process [Online]. *The International Journal of Advanced Manufacturing Technology*, 59(1–4), pp.137–142. Available from: http://dx.doi.org/10.1007/s00170-011-3470-8.

Jiang, J., Xu, X., and Stringer, J., 2018. Support Structures for Additive Manufacturing: A Review.

Journal of Manufacturing and Materials Processing , 2(4).

Johnson, S.G., 2017. The Nlopt nonlinear-optimization package [Online]. Available from: https://nlopt.readthedocs.io/en/latest/NLopt_Python_Reference/.

Johnson, T.E. and Gaynor, A.T., 2018. Three-dimensional projection-based topology optimization for prescribed-angle self-supporting additively manufactured structures [Online]. *Additive Manufacturing*. Available from: https://www.sciencedirect.com/science/article/pii/S2214860417303846#bib0145 [Accessed 25 September 2018].

Jones, J.C., 1992. *Design methods*. John Wiley & Sons.

Kabir, G. and Sumi, R.S., 2012. Selection of concrete production facility location integrating fuzzy AHP with TOPSIS method. *International Journal of Productivity Management and Assessment Technologies (IJPMAT)*, 1(1), pp.40–59.

Kazi, R.H., Grossman, T., Cheong, H., Hashemi, A., and Fitzmaurice, G., 2017. DreamSketch: Early Stage 3D Design Explorations with Sketching and Generative Design. In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST '17. New York, NY, USA: ACM, pp.401–414. Available from: http://doi.acm.org/10.1145/3126594.3126662.

Kellner, T., 2015. *The FAA Cleared the First 3D Printed Part to Fly in a Commercial Jet Engine from GE* [Online]. Available from: https://www.ge.com/reports/post/116402870270/the-faa-cleared-the-first-3d-printed-part-to-fly-2/ [Accessed 22 September 2018].

Khardekar, R. and McMains, S., 2006. Fast Layered Manufacturing Support Volume Computation on GPUs [Online]. , (4255X), pp.993–1002. Available from: http://dx.doi.org/10.1115/DETC2006-99666.

Khorram Niaki, M. and Nonino, F., 2017. Impact of additive manufacturing on business competitiveness: A multiple case study. *Journal of Manufacturing Technology Management*, 28(1), pp.56–74.

Kokotovich, V. and Dorst, K., 2016. The art of 'stepping back': Studying levels of abstraction in a diverse design team [Online]. *Design Studies*, 46, pp.79–94. Available from: https://www.sciencedirect.com/science/article/pii/S0142694X16300436 [Accessed 10 September 2018].

Koslow, T., 2015. *New Balance and Nervous System Collaborate to Make Running Great* [Online]. Available from: http://3dprintingindustry.com/2015/12/07/63131/ [Accessed 21 January

2016].

Kranz, J., Herzog, D., and Emmelmann, C., 2015. Design guidelines for laser additive manufacturing of lightweight structures in TiAl6V4 [Online]. *Journal of Laser Applications*, 27(S1), p. Available from: http://scitation.aip.org/content/lia/journal/jla/27/S1/10.2351/1.4885235.

Krish, S., 2011. A practical generative design method [Online]. *Computer-Aided Design*, 43(1), pp.88–100. Available from: http://www.sciencedirect.com/science/article/pii/S0010448510001764 [Accessed 6 August 2015].

Kroll, E. and Koskela, L., 2016. Explicating concepts in reasoning from function to form by two-step innovative abductions. *AI EDAM*, 30(Special Issue 02), pp.125–137.

Kumke, M., Watschke, H., and Vietor, T., 2016. A new methodological framework for design for additive manufacturing [Online]. *Virtual and Physical Prototyping*, 11(1), pp.3–19. Available from: http://www.tandfonline.com/doi/abs/10.1080/17452759.2016.1139377.

Kuo, Y.-H., Cheng, C.-C., Lin, Y.-S., and San, C.-H., 2018. Support structure design in additive manufacturing based on topology optimization [Online]. *Structural and Multidisciplinary Optimization*, 57(1), pp.183–195. Available from: https://doi.org/10.1007/s00158-017-1743-z.

Kurniawan, M.A., 2013. *M Kurniawan GE Jet Engine Bracket Version 1.2* [Online]. Available from: https://grabcad.com/library/m-kurniawan-ge-jet-engine-bracket-version-1-2-1 [Accessed 4 May 2018].

Lan, P.-T., Chou, S.-Y., Chen, L.-L., and Gemmill, D., 1997. Determining fabrication orientations for rapid prototyping with Stereolithography apparatus [Online]. *Computer-Aided Design*, 29(1), pp.53–62. Available from: https://www.sciencedirect.com/science/article/pii/S0010448596000498 [Accessed 31 July 2018].

Langelaar, M., 2016. Topology optimization of 3D self-supporting structures for additive manufacturing. *Additive Manufacturing*, 12, pp.60–70.

Langelaar, M., 2017. An additive manufacturing filter for topology optimization of print-ready designs [Online]. *Structural and Multidisciplinary Optimization*, 55(3), pp.871–883. Available from: https://doi.org/10.1007/s00158-016-1522-2.

Langelaar, M., 2018. Combined optimization of part topology, support structure layout and build orientation for additive manufacturing [Online]. *Structural and Multidisciplinary Optimization*. Available from: https://doi.org/10.1007/s00158-017-1877-z.

Lawson, B., 2002. CAD and creativity: does the computer really help? *Leonardo*, 35(3), pp.327–331.

Leary, M., Merli, L., Torti, F., Mazur, M., and Brandt, M., 2014. Optimal topology for additive manufacture: A method for enabling additive manufacture of support-free optimal structures [Online]. *Materials & Design*, 63, pp.678–690. Available from: http://www.sciencedirect.com/science/article/pii/S0261306914004646 [Accessed 21 September 2015].

Lefky, C.S., Zucker, B., Wright, D., Nassar, A.R., Simpson, T.W., and Hildreth, O.J., 2017. Dissolvable Supports in Powder Bed Fusion-Printed Stainless Steel [Online]. *3D Printing and Additive Manufacturing*, 4(1), pp.3–11. Available from: https://doi.org/10.1089/3dp.2016.0043.

Leong, K.F., Chua, C.K., and Ng, Y.M., 1996. A study of stereolithography file errors and repair. Part 1. Generic solution [Online]. *The International Journal of Advanced Manufacturing Technology*, 12(6), pp.407–414. Available from: http://dx.doi.org/10.1007/BF01186929.

Li, D., Levin, D.I.W., Matusik, W., and Zheng, C., 2016. Acoustic Voxels: Computational Optimization of Modular Acoustic Filters [Online]. *ACM Trans. Graph.*, 35(4), p.88:1--88:12. Available from: http://doi.acm.org/10.1145/2897824.2925960.

Limitstate LTD, 2018. *LIMITSTATE:FORM* [Online]. Available from: http://limitstate3d.com/ [Accessed 14 June 2018].

Liu, J., Gaynor, A.T., Chen, S., Kang, Z., Suresh, K., Takezawa, A., Li, L., Kato, J., Tang, J., Wang, C.C.L., Cheng, L., Liang, X., and To, A.C., 2018. Current and future trends in topology optimization for additive manufacturing [Online]. *Structural and Multidisciplinary Optimization*, 57(6), pp.2457–2483. Available from: https://doi.org/10.1007/s00158-018-1994-3.

Lockett, H., Ding, J., Williams, S., and Martina, F., 2017. Design for Wire + Arc Additive Manufacture: design rules and build orientation selection [Online]. *Journal of Engineering Design*, 28(7–9), pp.568–598. Available from: https://doi.org/10.1080/09544828.2017.1365826.

Maher, M. Lou and Poon, J., 1996. Modeling design exploration as co-evolution. *Computer Aided Civil and Infrastructure Engineering*, 11(3), pp.195–209.

Maheshwaraa Namasivayam, U. and Conner Seepersad, C., 2011. Topology design and freeform fabrication of deployable structures with lattice skins. *Rapid Prototyping Journal*, 17(1), pp.5–16.

Bin Maidin, S., Campbell, I., and Pei, E., 2012. Development of a design feature database to support design for additive manufacturing. *Assembly Automation*, 32(3), pp.235–244.

Markforged, 2018. *Metal X* [Online]. Available from: https://markforged.com/metal-x/ [Accessed 12 September 2018].

Martin, T., Umetani, N., and Bickel, B., 2015. OmniAD: Data-driven Omni-directional Aerodynamics [Online]. *ACM Trans. Graph.*, 34(4), p.113:1--113:12. Available from: http://doi.acm.org/10.1145/2766919.

Mass, Y. and Amir, O., 2016. Topology optimization for additive manufacturing : accounting for overhang limitations using a virtual skeleton [Online]. *Additive Manufacturing*, 18(2014), p.2016. Available from: http://www.sciencedirect.com/science/article/pii/S2214860417301045 [Accessed 4 October 2017].

Materialise NV, 2018. *Design Guides Homepage* [Online]. Available from: https://i.materialise.com/en/3d-printing-materials/design-guides [Accessed 13 June 2018].

Mehnen, J., Ding, J., Lockett, H., and Kazanas, P., 2014. Design study for wire and arc additive manufacture. *International Journal of Product Development 20*, 19(1–3), pp.2–20.

Mezzadri, F., Bouriakov, V., and Qian, X., 2018. Topology optimization of self-supporting support structures for additive manufacturing [Online]. *Additive Manufacturing*, 21, pp.666–682. Available from: https://www.sciencedirect.com/science/article/pii/S2214860418301519 [Accessed 17 September 2018].

Mirzendehdel, A.M. and Suresh, K., 2016. Support structure constrained topology optimization for additive manufacturing [Online]. *Computer-Aided Design*, 81, pp.1–13. Available from: http://linkinghub.elsevier.com/retrieve/pii/S0010448516300951 [Accessed 12 October 2017].

Morgan, H.D., Cherry, J.A., Jonnalagadda, S., Ewing, D., and Sienz, J., 2016. Part orientation optimisation for the additive layer manufacture of metal components [Online]. *The International Journal of Advanced Manufacturing Technology*, 86(5), pp.1679–1687. Available from: https://doi.org/10.1007/s00170-015-8151-6.

Mothersill, P. and Bove Jr., V.M., 2015. The EmotiveModeler: An Emotive Form Design CAD Tool. In: *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. CHI EA '15. New York, NY, USA: ACM, pp.339–342. Available from: http://doi.acm.org/10.1145/2702613.2725433.

Murphy, K.P., 2012. *Machine Learning: A Probabilistic Perspective*. MIT Press.

Nassehi, A., Newman, S., Dhokia, V., Zhu, Z., and Asrai, R.I., 2012. Using formal methods to model hybrid manufacturing processes. In: H. A. ElMaraghy, ed. Berlin, Heidelberg: Springer Berlin

Heidelberg, pp.52–56.

Nikander, J.B., Liikkanen, L.A., and Laakso, M., 2014. The preference effect in design concept evaluation [Online]. *Design Studies*, 35(5), pp.473–499. Available from: https://www.sciencedirect.com/science/article/pii/S0142694X14000301 [Accessed 11 July 2018].

Nikol, I., 2016. *Alcoa Bracket (BRCT12a) correction* [Online]. Available from: https://grabcad.com/library/alcoa-bracket-brct-12a-correction-1 [Accessed 4 May 2018].

Nourbakhsh, M., Irizarry, J., and Haymaker, J., 2018. Generalizable surrogate model features to approximate stress in 3D trusses [Online]. *Engineering Applications of Artificial Intelligence*, 71, pp.15–27. Available from: https://www.sciencedirect.com/science/article/pii/S095219761830006X [Accessed 4 August 2018].

nTopology, 2018. *nTopology* [Online]. Available from: www.ntopology.com [Accessed 30 August 2017].

Ohsaki, M., 1995. Genetic algorithm for topology optimization of trusses. *Computers & Structures*, 57(2), pp.219–225.

Opricovic, S. and Tzeng, G.-H., 2004. Compromise solution by MCDM methods: A comparative analysis of VIKOR and TOPSIS. *European journal of operational research*, 156(2), pp.445–455.

Oropallo, W. and Piegl, L., 2015. Ten challenges in 3D printing [Online]. *Engineering with Computers*, pp.1–14. Available from: http://dx.doi.org/10.1007/s00366-015-0407-0.

Osanov, M. and Guest, J.K., 2016. Topology Optimization for Architected Materials Design [Online]. *Annual Review of Materials Research*, 46(1), pp.211–233. Available from: http://dx.doi.org/10.1146/annurev-matsci-070115-031826.

Oxman, R., 2006. Theory and design in the first digital age. *Design studies*, 27(3), pp.229–265.

Parry, L., Ashcroft, I.A., and Wildman, R.D., 2016. Understanding the effect of laser scan strategy on residual stress in selective laser melting through thermo-mechanical simulation. *Additive Manufacturing*, 12, pp.1–15.

Pasko, A., Fryazinov, O., Vilbrandt, T., Fayolle, P.-A., and Adzhiev, V., 2011. Procedural function-based modelling of volumetric microstructures [Online]. *Graphical Models*, 73(5), pp.165–181. Available from: https://www.sciencedirect.com/science/article/pii/S1524070311000087

[Accessed 2 July 2018].

Piili, H., Happonen, A., Väistö, T., Venkataramanan, V., Partanen, J., and Salminen, A., 2015. Cost Estimation of Laser Additive Manufacturing of Stainless Steel. In: *Physics Procedia*. Elsevier, pp.388–396. Available from: http://www.sciencedirect.com/science/article/pii/S1875389215015436 [Accessed 10 November 2017].

Pizzolato, A., Sharma, A., Maute, K., Sciacovelli, A., and Verda, V., 2017. Topology optimization for heat transfer enhancement in Latent Heat Thermal Energy Storage [Online]. *International Journal of Heat and Mass Transfer*, 113, pp.875–888. Available from: https://www.sciencedirect.com/science/article/pii/S0017931017303034#f0055 [Accessed 24 January 2018].

Ponche, R., Hascoët, J.-Y., Kerbrat, O., and Mognol, P., 2012. A new global approach to design for additive manufacturing. *Virtual and Physical Prototyping*, 7(2), pp.93–105.

Pradel, P., Bibb, R., Zhu, Z., and Moultrie, J., 2018. Exploring the Impact of Shape Complexity on Build Time for Material Extrusion and Material Jetting BT - Industrializing Additive Manufacturing - Proceedings of Additive Manufacturing in Products and Applications - AMPA2017. In: M. Meboldt & C. Klahn, eds. Cham: Springer International Publishing, pp.24–33.

Pradel, P., Zhu, Z., Bibb, R., and Moultrie, J., 2018. Investigation of design for additive manufacturing in professional design practice [Online]. *Journal of Engineering Design*, 29(4–5), pp.165–200. Available from: https://doi.org/10.1080/09544828.2018.1454589.

Prakash, W.N., Sridhar, V.G., and Annamalai, K., 2014. New product development by DFMA and rapid prototyping. *ARPN J Eng Appl Sci*, 9(3), pp.274–279.

Prévost, R., Whiting, E., Lefebvre, S., and Sorkine-Hornung, O., 2013. Make It Stand: Balancing Shapes for 3D Fabrication [Online]. *ACM Trans. Graph.*, 32(4), p.81:1--81:10. Available from: http://doi.acm.org/10.1145/2461912.2461957.

Priedeman Jr, W.R. and Brosch, A.L., 2004. Soluble material and process for three-dimensional modeling.

Querin, O.M., Steven, G.P., and Xie, Y.M., 1998. Evolutionary structural optimisation (ESO) using a bidirectional algorithm. *Engineering computations*, 15(8), pp.1031–1048.

Rajan, S.D., 1995. Sizing, shape, and topology design optimization of trusses using genetic algorithm. *Journal of Structural Engineering*, 121(10), pp.1480–1487.

Rasmussen, C.E., 2004. Gaussian processes in machine learning. In: *Advanced lectures on machine learning*. Springer, pp.63–71.

Rebaioli, L. and Fassi, I., 2017. A review on benchmark artifacts for evaluating the geometrical performance of additive manufacturing processes [Online]. *The International Journal of Advanced Manufacturing Technology*, 93(5), pp.2571–2598. Available from: https://doi.org/10.1007/s00170-017-0570-0.

Renishaw plc., 2017. *Renishaw previews QuantAM Dental software for additive manufacturing at IDS* [Online]. Available from: https://www.renishaw.com/en/renishaw-previews-quantam-dental-software-for-additive-manufacturing-at-ids--40927 [Accessed 8 November 2018].

Renishaw plc., 2018. *RenAM 500Q* [Online]. Available from: http://www.renishaw.com/en/renam-500q--42781 [Accessed 12 September 2018].

Robertson, B.F. and Radcliffe, D.F., 2009. Impact of CAD tools on creative problem solving in engineering design [Online]. *Computer-Aided Design*, 41(3), pp.136–146. Available from: http://www.sciencedirect.com/science/article/pii/S0010448508001334 [Accessed 2 November 2015].

Rozvany, G.I.N., 2009. A critical review of established methods of structural topology optimization [Online]. *Structural and Multidisciplinary Optimization*, 37(3), pp.217–237. Available from: https://doi.org/10.1007/s00158-007-0217-0.

Rozvany, G.I.N., Zhou, M., and Birker, T., 1992. Generalized shape optimization without homogenization [Online]. *Structural optimization*, 4(3), pp.250–252. Available from: https://doi.org/10.1007/BF01742754.

Ruffo, M., Tuck, C., and Hague, R., 2006. Cost estimation for rapid manufacturing-laser sintering production for low to medium volumes. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, 220(9), pp.1417–1427.

Saaty, T.L., 1990. The analytic hierarchy process. *European Journal of Operational Research*, 48, pp.9–26.

Salonitis, K., 2016. Design for additive manufacturing based on the axiomatic design method [Online]. *The International Journal of Advanced Manufacturing Technology*, pp.1–8. Available from: http://dx.doi.org/10.1007/s00170-016-8540-5.

Schaedler, T.A. and Carter, W.B., 2016. Architected Cellular Materials. *Annual Review of Materials Research*, (0).

Schaedler, T.A., Ro, C.J., Sorensen, A.E., Eckel, Z., Yang, S.S., Carter, W.B., and Jacobsen, A.J., 2014. Designing metallic microlattices for energy absorber applications. *Advanced Engineering Materials*, 16(3), pp.276–283.

Schmutzler, C., Zimmermann, A., and Zaeh, M.F., 2016. Compensating Warpage of 3D Printed Parts Using Free-form Deformation [Online]. *Procedia CIRP*, 41, pp.1017–1022. Available from: https://www.sciencedirect.com/science/article/pii/S2212827115011579 [Accessed 28 October 2018].

Schon, D.A., 1992. Designing as reflective conversation with the materials of a design situation [Online]. *Research in Engineering Design*, 3(3), pp.131–147. Available from: https://doi.org/10.1007/BF01580516.

Schulz, A., Xu, J., Zhu, B., Zheng, C., Grinspun, E., and Matusik, W., 2017. Interactive Design Space Exploration and Optimization for CAD Models [Online]. *ACM Trans. Graph.*, 36(4), p.157:1--157:14. Available from: http://doi.acm.org/10.1145/3072959.3073688.

Schwab, K., 2017. *Autodesk's Generatively Designed Cabit Seat Could Make Flying Cheaper* [Online]. Available from: https://www.fastcodesign.com/90124395/autodesks-generatively-designed-cabin-seat-could-make-flying-cheaper [Accessed 30 August 2017].

Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., and Freitas, N. de, 2016. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1), pp.148–175.

Shea, K., Aish, R., and Gourtovaia, M., 2005. Towards integrated performance-driven generative design tools [Online]. *Automation in Construction*, 14(2), pp.253–264. Available from: http://www.sciencedirect.com/science/article/pii/S0926580504000809 [Accessed 16 May 2016].

Sigmund, O., 2011. On the usefulness of non-gradient approaches in topology optimization [Online]. *Structural and Multidisciplinary Optimization*, 43(5), pp.589–596. Available from: https://doi.org/10.1007/s00158-011-0638-7.

Sigmund, O. and Maute, K., 2013. Topology optimization approaches [Online]. *Structural and Multidisciplinary Optimization*, 48(6), pp.1031–1055. Available from: https://doi.org/10.1007/s00158-013-0978-6.

Sigmund, O. and Torquato, S., 1997. Design of materials with extreme thermal expansion using a three-phase topology optimization method. *Journal of the Mechanics and Physics of Solids*, 45(6), pp.1037–1067.

Silverman, B.G. and Mezher, T.M., 1992. Expert critics in engineering design: Lessons learned and research needs. *AI magazine*, 13(1), p.45.

Simon, H.A., 1973. The structure of ill structured problems. *Artificial intelligence*, 4(3–4), pp.181–201.

Simon, H.A., 1996. *The sciences of the artificial*. MIT press.

Simonelli, M., Tse, Y.Y., and Tuck, C., 2014. On the Texture Formation of Selective Laser Melted Ti-6Al-4V [Online]. *Metallurgical and Materials Transactions A*, 45(6), pp.2863–2872. Available from: https://doi.org/10.1007/s11661-014-2218-0.

Singh, V. and Gu, N., 2012. Towards an integrated generative design framework [Online]. *Design Studies*, 33(2), pp.185–207. Available from: http://www.sciencedirect.com/science/article/pii/S0142694X11000391 [Accessed 8 April 2016].

Smith, C.J., 2016. *Application of layout optimisation to the design of additively manufactured metallic components*. University of Sheffield.

Smith, C.J., Gilbert, M., Todd, I., and Derguti, F., 2016. Application of layout optimization to the design of additively manufactured metallic components [Online]. *Structural and Multidisciplinary Optimization*, 54(5), pp.1297–1313. Available from: https://doi.org/10.1007/s00158-016-1426-1.

Snider, C.M., Culley, S.J., and Dekoninck, E.A., 2013. Analysing creative behaviour in the later stage design process [Online]. *Design Studies*, 34(5), pp.543–574. Available from: http://www.sciencedirect.com/science/article/pii/S0142694X13000239.

Solid Thinking, 2016. *SolidThinking Inspire* [Online]. Available from: http://www.solidthinking.com/productoverview.aspx?item=inspire overview&category=products [Accessed 8 September 2016].

Sossou, G., Demoly, F., Montavon, G., and Gomes, S., 2018. An additive manufacturing oriented design approach to mechanical assemblies [Online]. *Journal of Computational Design and Engineering*, 5(1), pp.3–18. Available from: https://www.sciencedirect.com/science/article/pii/S2288430017300659 [Accessed 6 August 2018].

Srinivasan, H., Harrysson, O.L.A., and Wysk, R.A., 2015. Automatic part localization in a CNC machine coordinate system by means of 3D scans [Online]. *The International Journal of Advanced*

Manufacturing Technology, 81(5), pp.1127–1138. Available from: https://doi.org/10.1007/s00170-015-7178-z.

Srinivasan, V., Mandal, E., and Akleman, E., 2005. Solidifying wireframes. In: *Proceedings of the 2004 bridges conference on mathematical connections in art, music, and science*.

Stanković, T., Mueller, J., and Shea, K., 2017. The effect of anisotropy on the optimization of additively manufactured lattice structures [Online]. *Additive Manufacturing*, 17, pp.67–76. Available from: https://www.sciencedirect.com/science/article/pii/S2214860417300593#! [Accessed 9 November 2018].

Stouffs, R. and Rafiq, Y., 2015. Generative and evolutionary design exploration. *AI EDAM*, 29(4), pp.329–331.

Strano, G., Hao, L., Everson, R.M., and Evans, K.E., 2013. A new approach to the design and optimisation of support structures in additive manufacturing [Online]. *The International Journal of Advanced Manufacturing Technology*, 66(9–12), pp.1247–1254. Available from: http://dx.doi.org/10.1007/s00170-012-4403-x.

Svanberg, K., 1987. The method of moving asymptotes—a new method for structural optimization [Online]. *International Journal for Numerical Methods in Engineering*, 24(2), pp.359–373. Available from: http://doi.wiley.com/10.1002/nme.1620240207 [Accessed 12 October 2017].

Tammas-Williams, S. and Todd, I., 2017. Design for additive manufacturing with site-specific properties in metals and alloys. *Scripta Materialia*, 135, pp.105–110.

Tang, T.L.E., Liu, Y., Lu, D., Arisoy, E.B., and Musuvathy, S., 2017. Lattice Structure Design Advisor for Additive Manufacturing Using Gaussian Process [Online]. , (58110), p.V001T02A020. Available from: http://dx.doi.org/10.1115/DETC2017-67282.

Tang, Y., Kurtz, A., and Zhao, Y.F., 2015. Bidirectional Evolutionary Structural Optimization (BESO) based design method for lattice structure to be fabricated by additive manufacturing [Online]. *Computer-Aided Design*, 69, pp.91–101. Available from: http://www.sciencedirect.com/science/article/pii/S0010448515000792 [Accessed 12 October 2015].

Tang, Y. and Zhao, Y.F., 2016. A survey of the design methods for additive manufacturing to improve functional performance [Online]. *Rapid Prototyping Journal*, 22(3), pp.569–590. Available from: http://www.emeraldinsight.com/doi/abs/10.1108/RPJ-01-2015-0011.

Teibrich, A., Mueller, S., Guimbretière, F., Kovacs, R., Neubert, S., and Baudisch, P., 2015. Patching

Physical Objects. In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. UIST '15. New York, NY, USA: ACM, pp.83–91. Available from: http://doi.acm.org/10.1145/2807442.2807467.

Teitelbaum, G.A., Goaer, Y., and Schmidt, L.C., 2009. Examining Potential Design Guidelines for Use in Fused Deposition Modeling to Reduce Build Time and Material Volume. *ASME 2009 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Volume 8, pp.1–10.

Theus, M., 2008. High-dimensional data visualization. In: *Handbook of data visualization*. Springer, pp.151–178.

Thomas, D., 2009. *The Development of Design Rules for Selective Laser Melting*. University of Wales.

Thompson, M.K., Moroni, G., Vaneker, T., Fadel, G., Campbell, R.I., Gibson, I., Bernard, A., Schulz, J., Graf, P., and Ahuja, B., 2016. Design for Additive Manufacturing: Trends, opportunities, considerations, and constraints. *CIRP annals*, 65(2), pp.737–760.

Thornton Tomasetti, 2015. *Design Explorer* [Online]. Available from: http://core.thorntontomasetti.com/design-explorer/ [Accessed 24 October 2018].

Tomlin, M. and Meyer, J., 2011. Topology optimization of an additive layer manufactured (ALM) aerospace part. In: *Proceeding of the 7th Altair CAE technology conference.* pp.1–9.

Ultimaker, 2017. Ultimaker Cura 3.4.1. [Online]. Available from: https://ultimaker.com/en/products/ultimaker-cura-software.

Ulu, E., Korkmaz, E., Yay, K., Burak Ozdoganlar, O., and Burak Kara, L., 2015. Enhancing the Structural Performance of Additively Manufactured Objects Through Build Orientation Optimization [Online]. *Journal of Mechanical Design*, 137(11), pp.111410–111419. Available from: http://dx.doi.org/10.1115/1.4030998.

Umetani, N. and Bickel, B., 2018. Learning Three-dimensional Flow for Interactive Aerodynamic Design [Online]. *ACM Trans. Graph.*, 37(4), p.89:1--89:10. Available from: http://doi.acm.org/10.1145/3197517.3201325.

Umetani, N., Panotopoulou, A., Schmidt, R., and Whiting, E., 2016. Printone: Interactive Resonance Simulation for Free-form Print-wind Instrument Design [Online]. *ACM Trans. Graph.*, 35(6), p.184:1--184:14. Available from: http://doi.acm.org/10.1145/2980179.2980250.

Uriondo, A., Esperon-Miguez, M., and Perinpanayagam, S., 2015. The present and future of additive

manufacturing in the aerospace sector: A review of important aspects. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 229(11), pp.2132–2147.

Utterback, J., Vedin, B.-A., Alvarez, E., Ekman, S., Walsh Sanderson, S., Tether, B., and Verganti, R., 2006. Design-inspired innovation and the design discourse. *Design-inspired innovation*, pp.154–186.

Vanek, J., Galicia, J.A.G., and Benes, B., 2014. Clever Support: Efficient Support Structure Generation for Digital Fabrication [Online]. *Comput. Graph. Forum*, 33(5), pp.117–125. Available from: http://dx.doi.org/10.1111/cgf.12437.

Vayre, B., Vignat, F., and Villeneuve, F., 2012. Designing for Additive Manufacturing [Online]. *Procedia CIRP*, 3, pp.632–637. Available from: http://www.sciencedirect.com/science/article/pii/S2212827112002806 [Accessed 29 October 2015].

Vayre, B., Vignat, F., and Villeneuve, F., 2013. Identification on Some Design Key Parameters for Additive Manufacturing: Application on Electron Beam Melting [Online]. *Procedia CIRP*, 7, pp.264–269. Available from: http://www.sciencedirect.com/science/article/pii/S2212827113002527 [Accessed 29 October 2015].

VDI Guideline, 1993. 2221: Systematic approach to the development and design of technical systems and products. *Beuth, Berlin*.

Veisz, D., Namouz, E.Z., Joshi, S., and Summers, J.D., 2012. Computer-aided design versus sketching: An exploratory case study [Online]. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 26(3), pp.317–335. Available from: https://www.cambridge.org/core/article/computeraided-design-versus-sketching-an-exploratory-case-study/E49E0BE022FD96D80F9B61691458A95B.

Wang, C., Duan, Q., Gong, W., Ye, A., Di, Z., and Miao, C., 2014. An evaluation of adaptive surrogate modeling based optimization with two benchmark problems [Online]. *Environmental Modelling & Software*, 60, pp.167–179. Available from: https://www.sciencedirect.com/science/article/pii/S1364815214001698#bib45 [Accessed 9 July 2018].

Wang, L. and Whiting, E., 2016. Buoyancy optimization for computational fabrication. In: *Computer*

Graphics Forum. Wiley Online Library, pp.49–58.

Wang, M.Y., Wang, X., and Guo, D., 2003. A level set method for structural topology optimization. *Computer methods in applied mechanics and engineering*, 192(1), pp.227–246.

Wang, X., Xu, S., Zhou, S., Xu, W., Leary, M., Choong, P., Qian, M., Brandt, M., and Xie, Y.M., 2016. Topological design and additive manufacturing of porous metals for bone scaffolds and orthopaedic implants: a review. *Biomaterials*, 83, pp.127–141.

Wills, G., 2008. Linked data views. In: *Handbook of data visualization*. Springer, pp.217–241.

Wohlers, T., Campbell, I., Diegel, O., and Kowen, J., 2018. *Wohlers Report 2018*.

Wortmann, T., 2017. Surveying design spaces with performance maps: A multivariate visualization method for parametric design and architectural design optimization. *International Journal of Architectural Computing*, 15(1), pp.38–53.

Wortmann, T., 2018. Efficient , Visual , and Interactive Architectural Design Optimization with Model-based Methods. , (July).

Wortmann, T., Costa, A., Nannicini, G., and Schroepfer, T., 2015. Advantages of surrogate models for architectural design optimization. *AI EDAM*, 29(4), pp.471–481.

Wu, D., Terpenny, J., and Schaefer, D., 2017. Digital design and manufacturing on the cloud: A review of software and services [Online]. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 31(1), pp.104–118. Available from: https://www.cambridge.org/core/article/digital-design-and-manufacturing-on-the-cloud-a-review-of-software-and-services/73F654CC4DF987217508D6D8C0C825A2.

Wu, J., Dick, C., and Westermann, R., 2016. A System for High-Resolution Topology Optimization. *IEEE Transactions on Visualization and Computer Graphics*, 22(3), pp.1195–1208.

Xie, Y.M. and Steven, G.P., 1993. A simple evolutionary procedure for structural optimization. *Computers & structures*, 49(5), pp.885–896.

Xu, K., Kwok, T.-H., Zhao, Z., and Chen, Y., 2017. A Reverse Compensation Framework for Shape Deformation Control in Additive Manufacturing [Online]. *Journal of Computing and Information Science in Engineering*, 17(2), pp.21012–21019. Available from: http://dx.doi.org/10.1115/1.4034874.

Yang, S., Tang, Y., and Zhao, Y.F., 2015. A new part consolidation method to embrace the design freedom of additive manufacturing [Online]. *Journal of Manufacturing Processes*. Available

from: http://www.sciencedirect.com/science/article/pii/S1526612515000699 [Accessed 30 November 2015].

Yoon, K.P. and Hwang, C.-L., 1995. *Multiple attribute decision making: an introduction*. Sage publications.

Zhakeyev, A., Wang, P., Zhang, L., Shu, W., Wang, H., and Xuan, J., 2017. Additive Manufacturing: Unlocking the Evolution of Energy Materials [Online]. *Advanced Science*, 4(10), p.1700187. Available from: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC5644240/.

Zhang, C. and Chen, T., 2001. Efficient feature extraction for 2D/3D objects in mesh representation. In: *Image Processing, 2001. Proceedings. 2001 International Conference on*. IEEE, pp.935–938.

Zhang, R., Auzinger, T., Ceylan, D., Li, W., and Bickel, B., 2017. Functionality-aware Retargeting of Mechanisms to 3D Shapes [Online]. *ACM Trans. Graph.*, 36(4), p.81:1--81:13. Available from: http://doi.acm.org/10.1145/3072959.3073710.

Zhang, X., Le, X., Panotopoulou, A., Whiting, E., and Wang, C.C.L., 2015. Perceptual models of preference in 3D printing direction. *ACM Transactions on Graphics (TOG)*, 34(6), p.215.

Zhang, Y., Bernard, A., Harik, R., and Fadel, G., 2018. A new method for single-layer-part nesting in additive manufacturing [Online]. *Rapid Prototyping Journal*, 24(5), pp.840–854. Available from: https://doi.org/10.1108/RPJ-01-2017-0008.

Zhao, B., Lin, Z., Fu, J., and Sun, Y., 2017. Generation of truss-structure objects with implicit representation for 3D-printing [Online]. *International Journal of Computer Integrated Manufacturing*, 30(8), pp.871–879. Available from: https://doi.org/10.1080/0951192X.2016.1224390.

Zheng, X., Lee, H., Weisgraber, T.H., Shusteff, M., DeOtte, J., Duoss, E.B., Kuntz, J.D., Biener, M.M., Ge, Q., Jackson, J.A., Kucheyev, S.O., Fang, N.X., and Spadaccini, C.M., 2014. Ultralight, ultrastiff mechanical metamaterials [Online]. *Science*, 344(6190), p.1373 LP-1377. Available from: http://science.sciencemag.org/content/344/6190/1373.abstract.

Zhou, M. and Rozvany, G.I.N., 1991. The COC algorithm, Part II: Topological, geometrical and generalized shape optimization [Online]. *Computer Methods in Applied Mechanics and Engineering*, 89(1–3), pp.309–336. Available from: https://www.sciencedirect.com/science/article/pii/0045782591900469 [Accessed 24 January 2018].

Zsolnai-Fehér, K., Wonka, P., and Wimmer, M., 2018. Gaussian Material Synthesis. *arXiv preprint*

arXiv:1804.08369.

Zwier, M.P. and Wits, W.W., 2016. Design for Additive Manufacturing: Automated Build Orientation Selection and Optimization. *Procedia CIRP*, 55, pp.128–133.