

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Geometric Manufacturability Analysis for Additive Manufacturing

Permalink

<https://escholarship.org/uc/item/9s3277jh>

Author

Budinoff, Hannah Dawes

Publication Date

2019

Peer reviewed|Thesis/dissertation

Geometric Manufacturability Analysis for Additive Manufacturing

By

Hannah Dawes Budinoff

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Mechanical Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Sara McMains, Chair

Professor Dennis Lieu

Professor Marti Hearst

Summer 2019

Geometric Manufacturability Analysis for Additive Manufacturing

Copyright 2019

by

Hannah Dawes Budinoff

Abstract

Geometric Manufacturability Analysis for Additive Manufacturing

by

Hannah Dawes Budinoff

Doctor of Philosophy in Engineering – Mechanical Engineering

University of California, Berkeley

Professor Sara McMains, Chair

During the development of a new product, it is difficult for designers to predict how their design decisions will impact manufacturability and manufacturing cost of the individual parts in their product. Additive manufacturing is increasingly becoming a viable option to produce high fidelity prototypes and even small-scale production part runs. However, as an emerging technology, there are few resources available to help designers make design decisions regarding quality and manufacturability for additive manufacturing. Most information developed to help designers ensure manufacturability is in the form of general guidelines that designers must interpret and then use their best judgment to scrutinize their design. Designers can only guess, based on previous experience, if the process can produce part features that meet their specified geometric tolerances. However, by using algorithms to analyze part geometry, it is possible to predict additive manufacturing outcomes. This thesis describes the development of two software tools to analyze part geometry in near real-time: one that predicts manufacturability, and another that predicts achievable quality.

These tools are used to explore how automated part geometry analysis influences the effectiveness of design for additive manufacturing feedback. The research hypothesis of this thesis is that part geometry analysis improves the practicality, accuracy, and usefulness of design for additive manufacturing feedback. To test this hypothesis, three research thrusts were conducted: evaluating the performance of the newly developed tools relative to existing tools, experimental verification of the predictions of the tools, and a user study evaluating usage of the manufacturability tool during a design task. Comparison with existing tools indicated that both tools described in this thesis have similar computation time as existing solutions, while providing greater potential to allow designers to analyze manufacturing trade-offs, with a more comprehensive approach to modeling sources of errors in the manufacturing process. A range of parts were printed using fused deposition modeling and then inspected. The experimental results showed that the predictions of both tools were relatively accurate, and highlighted several additional process parameters that can be included in the modeling approach to improve accuracy. Lastly, a user study demonstrated that use of the software tool reduced the number of manufacturability problems in participants' designs while requiring a similar amount of time to use, compared with using a list of design heuristics. The findings of the thesis support the practicality, accuracy, and usefulness of geometry analysis software tools to support design for additive manufacturing.

Acknowledgments

I would like to gratefully acknowledge the support and guidance of my doctoral advisor, Prof. Sara McMains. You helped me so much in developing as a teacher, communicator, and researcher. I am also grateful to Prof. David Dornfeld for starting me on my academic journey at Berkeley and to the members of my dissertation and qualifying committees for their feedback on my research. Thanks also to current and former members of the Computer-Aided Design and Manufacturing Laboratory and the Laboratory for Manufacturing and Sustainability for their collaboration and helpful discussions, especially Alberto Rinaldi, Bodi Yuan, Xiang Li, Jeronimo Mora, Sara Shonkwiler, Max Micali, Raunak Bhinge, and Amrita Srinivasan. Thank you to the staff of the ME machine shop, especially Jesse Lopez, and to Yilin Sun, who was enormously helpful in gathering experimental data. Special thanks to Sarah Frank, Julia Kramer, and Audrey Ford for the invaluable moral support and writing feedback. Finally, I would like to express my deepest appreciation to my family and my husband. I am eternally grateful for your patience and belief in me.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1752814. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

Parts of this dissertation have previously been published by ASME (Proceedings of the 2018 IDETC/CIE Conference) [1] and Elsevier (Procedia CIRP) [2].

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Research focus	4
1.2.1	Assumptions	6
1.3	Contribution	6
1.4	Thesis structure	7
2	Background and previous work	8
2.1	Design for manufacturing	8
2.1.1	Software-based geometry analysis for DFM	9
2.2	Design for additive manufacturing	13
2.2.1	Sources of geometric inaccuracy in AM	14
2.2.2	Experimental measurements of geometric accuracy in AM	18
2.2.3	Design guidelines for additive manufacturing	23
2.2.4	Existing software-based DFAM tools	25
2.3	Summary	29
3	Geometry analysis for design for additive manufacturing feedback	30
3.1	Introduction	30
3.2	Coordinate system	31
3.3	Manufacturability assessment - Will It Print	32
3.3.1	Tipping	34
3.3.2	Small features	36
3.3.3	Surface roughness	39
3.3.4	Warping	41
3.3.5	Support material	42
3.4	Implementation of Will It Print	44
3.5	Analysis of achievable tolerances - TAAM	46
3.5.1	Mathematical analysis of the impact of orientation, layer thickness, and layer deviations on geometric accuracy	48
3.5.2	Other factors impacting achievable tolerances	54
3.5.3	Refined error equations for FDM	54
3.5.4	Visualization of errors corresponding to tolerances	55
3.5.5	Tolerance allocation tool details	58
3.6	Implementation of tolerance allocation tool - TAAM	63
3.7	Conclusion	64

4	Example application of geometry analysis tools	66
4.1	Introduction	66
4.2	Materials and methods	67
4.2.1	Benchmarking performance of Will It Print	69
4.2.2	Benchmarking performance of TAAM	71
4.3	Results and Discussion - Will It Print	74
4.3.1	Results - Will It Print output	74
4.3.2	Comparison of Will It Print to other tools	78
4.3.3	Discussion	83
4.4	Results and Discussion - TAAM	85
4.4.1	TAAM Output	85
4.4.2	Comparison of TAAM to other orientation optimization tools	89
4.4.3	Discussion	91
4.5	Conclusion	95
5	Experimental verification of geometry analysis	97
5.1	Introduction	97
5.2	Materials and methods	98
5.2.1	AM equipment	99
5.2.2	Parts printed	99
5.2.3	TAAM measurement procedure	102
5.2.4	TAAM implementation and comparison	103
5.3	Results - Study 1: Will It Print	105
5.3.1	Warping	105
5.3.2	Tipping	108
5.3.3	Small features	111
5.4	Discussion - Study 1: Will It Print	116
5.5	Results - Study 2: TAAM	117
5.5.1	Flatness	118
5.5.2	Cylindricity	121
5.5.3	Orientation (angularity, parallelism, perpendicularity)	122
5.6	Discussion - Study 2: TAAM	123
5.6.1	Effect of independent variables	123
5.6.2	Effect of extraneous variables	125
5.6.3	Overall evaluation of TAAM performance	127
5.7	Conclusion	128
6	Verification of utility of geometry analysis tools	130
6.1	Introduction	130
6.2	Study method	131
6.2.1	Study participants	132
6.2.2	Description of software and worksheet manufacturability tools	132

6.2.3	Design problem	133
6.2.4	Study measurements	134
6.2.5	Statistical tests	135
6.3	Results	135
6.3.1	Evaluating study hypotheses	136
6.3.2	Differences in designer effectiveness	138
6.3.3	Effect of prior experience	140
6.3.4	Participant comments	142
6.3.5	Evaluation of the software tool	143
6.4	Discussion	145
6.5	Conclusions	147
7	Conclusion	148
7.1	Review of research objectives	148
7.2	Contribution	151
7.3	Future research	151
7.4	Conclusion	153
A	Experimental parts	154
B	Usability tool assignment	158
C	DFAM worksheet	162

Chapter 1

Introduction

In this chapter, the context of the problem that this thesis seeks to address is presented. Then, the specific research focus of the thesis is summarized. Contributions of this work are discussed. Finally, the structure of the thesis is described.

1.1 Motivation

Advances in engineering, mechanical design, and manufacturing lead to new technologies that improve our quality of life. The impact of engineered goods on human activities is dramatic and ubiquitous. Most of the objects we interact with daily—from airplanes to cellphones to medical devices—were designed, tested, and manufactured in a new-product-development process led by engineers.

Industries involved in the production of engineered products are also a fundamental part of the US and worldwide economy. The manufacturing industry contributed 12% to US GDP, compared with a worldwide average 16% of GDP in 2016 [3]. Despite a general downturn in the industry in the past few decades, between 2017 and 2018, the contribution of durable goods manufacturing to US GDP increased 5.4% [4]. It is clear from these numbers that the engineering and production of products is an essential activity with an outsized impact on our daily lives. Supporting this industry are thousands of engineers and technicians. In 2013, roughly 12 million workers (8.8% of all US employees) were employed by the US manufacturing sector [5]. In that same year, the number of engineers employed in the US was approximately 1.5 million [6]. Engineers who work on developing new products have a particularly important and challenging task.

During the design process for a new product, designers make a myriad of decisions that determine the final configuration of the product (e.g., deciding the geometry and arrangement of features, determining how individual parts will fit into an assembly, and determining what material and manufacturing process will be used to create the parts). These complex decisions are often interconnected, affecting both the functionality of the product and influencing the cost and ease of manufacturing the product. Design decisions can have a significant impact on the total product cost, and therefore, the ultimate success of a prod-

uct. Many studies have anecdotally stated that decisions made during the product design stage determine for 70% of ultimate manufacturing cost. While that exact value may be apocryphal, there is a consensus in industry that product design decisions determine a significant portion of product costs [7]. While the exact impact of design decisions on product cost depends on the industry and product type, these decisions are crucial in determining the overall cost, and designers should consider them carefully.

Even though engineering design is central to the creation of new products, the process of design itself has only somewhat recently become the focus of study [8]. There is a lack of awareness regarding best practices for engineering design, especially in the specific case of considering the interconnections between design decisions and manufacturing. When a design tool is presented in the literature, there is not typically an associated user study validating the tool's utility or its impact on the design process. There is a need to understand how to consider manufacturing constraints effectively during the design process, in order to develop and manufacture products in a cost-efficient manner.

Historically in the US, an “over-the-wall” design process was typical [9]: engineers designed a part without considering manufacturing complexity or cost and then transferred the part to manufacturing specialists to manufacture. The part was re-designed if or when those specialists determined that the manufacturing cost or complexity was too high. However, design changes are very costly to implement late in the product development process. A widely cited rule of thumb known as the “rule of 10” states that the cost to implement a design change increases by a factor of 10 between each phase of the design process [10]. So, for example, making a change to the design of a part after production tooling is procured may be ten times more costly than making the change before procuring the tooling. In addition to the cost of scrapped tooling, the increased cost is due to the large number of person-hours required to process and propagate the change [10].

In order to avoid costly changes late in the development process, there is a need to move away from the over-the-wall model. To this end, there has been an increased interest from academia and industry over the last 50 years to develop methodologies to promote the consideration of manufacturing constraints during the design process. These methodologies are often referred to as design for manufacturing (DFM). DFM is a cornerstone of concurrent engineering. Concurrent, or simultaneous, engineering is a method for developing a new product where the different stages of the design cycle (e.g., concept generation, prototyping, testing, manufacturing) occur simultaneously. In the traditional engineering design process, these stages occur sequentially. A concurrent engineering approach can lead to shortened product development lead times [11–13], improved product quality, [14], and reduced product cost [13]. Researchers have published many examples of DFM tools and strategies for traditional manufacturing processes [15, 16] that can be used as part of a concurrent engineering approach.

With advances in technology, new manufacturing processes are developed for designers to consider. Additive manufacturing (AM) is a collection of technologies that have developed in the past 35 years where a part is created by adding material in a layer-by-layer fashion. AM is fundamentally different from traditional subtractive manufacturing approaches, where a part

is made by removing extra material. AM technologies include stereolithography, selective laser melting and sintering, and fused deposition modeling. While AM was initially used primarily for prototyping, it is increasingly used for one-off and small batch production. It has started to become a viable process for designers to consider selecting for producing their products. The growth in interest and use of AM technologies has been extraordinary. AM has experienced double-digit growth most of the past 35 years, growing into a market that was worth over \$4 billion in 2014 [17].

To introduce the key terminology of AM processes, a brief overview of one AM process, namely fused deposition modeling (FDM), is presented here. Fused deposition modeling is also known as fused filament fabrication. The part geometry is represented in an STL (stereolithography) file, which contains geometry information in the form of vertex location and connectivity information between vertices. Before printing a part, a designer can rotate the STL file in 3D space and choose a build orientation. A slicer program divides the part geometry into a series of layers, based on the chosen layer thickness and build orientation. The slicer program then generates building instructions that are sent to the FDM machine. The FDM machine pulls plastic filament from a spool into a heated nozzle, known as an extruder (Fig. 1.1). The extruder (or the build platform, depending on the machine) is moved to the appropriate positions while the extruder deposits heated filament. If there are overhanging faces (i.e., faces that are downward-facing with no part geometry beneath them to support them), the extruder also deposits removable support material beneath the overhanging face. Depending on the machine, the support material may be made using a different filament than the filament used to build the part, or the same filament. Support material is typically created in a less dense scaffold pattern, meant to be easy to remove after printing. Once a single layer is complete, the build platform moves vertically, and the construction of the next layer begins. The build platform and build chamber surrounding the print area may or may not be heated above ambient temperature.

Additive manufacturing technologies give designers the ability to create complex geometry that was not possible with more traditional manufacturing processes [18]. AM presents designers with opportunities for part consolidation, weight reduction, functional customization, personalization, and new aesthetics functional customization [19]. AM technologies also bring new challenges for designers and manufacturers [20]. In AM, process parameters used during the manufacturing process can have a significant influence on the quality of the produced part [21]. As discussed in [22], designers who use AM make decisions about build orientation, layer thickness, support material, but may not understand how these decisions impact the achievable quality of their part. Designers may send a design to be manufactured only to have several unsuccessful prints due to poor selection of process parameters, build direction, or layer thickness. The impact of some of these factors may lessen as AM technology matures. However, many of these issues are inherent to the layered nature of AM and need to be considered earlier in the design process when changes are less expensive to implement. As a relatively new technology, there is still uncertainty as to the best practices of applying DFM techniques to additive manufacturing [17, 23]. The set of design methodologies to help designers consider the freedoms and constraints associated with addi-




Figure 1.1: An overview of the key parts of an extrusion-based system, such as FDM.

tive manufacturing, known as design for additive manufacturing (DFAM), requires further development.

In summary, design and manufacturing play critical roles in our economy and society. Design and manufacturing processes are also inherently linked. DFM can help designers understand the connections between these two processes so they can more objectively consider trade-offs between design and manufacturing goals, such as cost and quality, to ultimately improve design effectiveness and reduce cost. However, there are few user studies about designers using DFM tools during the design process. We have little understanding of how designers handle manufacturing trade-offs, how designer characteristics impact usage of DFM tools, or even if DFM tools are effective at reducing manufacturing issues. Further, the area of DFAM is still maturing and consequently, there are few DFAM tools.

1.2 Research focus

The research focus of this thesis is on geometry analysis for DFAM. As will be surveyed in the next chapter, researchers have developed DFAM tools addressing various needs of designers at different stages of the design process. One promising but underdeveloped area of research is applying geometry analysis to DFAM. Geometry analysis for DFAM uses algorithms to analyze part geometry (i.e., size and arrangement of different part features) to assess compatibility with AM geometry constraints. Geometry analysis can provide the designer with manufacturing feedback (e.g., this feature is too small to be successfully manufactured, or this feature will likely have a large amount of geometric error). In this thesis, we develop two geometry analysis tools. These two tools are developed to better address the needs of

designers who would like to use AM, but also to help explore what makes DFAM analysis effective. Both tools have a similar focus in that they are DFAM tools based on geometry analysis. However, the tools were developed to address designer needs at different stages in the design process.

The first tool developed in this thesis assesses the AM manufacturability of a part. Manufacturability refers to the relative ease with which a part can be manufactured. There is a pressing need to develop design guidelines for AM to help ensure parts are designed to be manufacturable [23]. Manufacturability analysis systems, which use algorithms to detect features on a part that are incompatible with a particular manufacturing process, have been developed for traditional manufacturing processes [15]. Few analogous algorithms exist for AM. A manufacturability analysis system for AM will enable designers to evaluate if their design contains any features or attributes that prevent it from being manufactured easily. If such features or attributes exist, then the designer can consider geometry or process parameter changes. As a result, designers can improve the manufacturability of their part early in the design process.

The second tool developed in this thesis assesses the achievable quality of a part. Quality is quantified using geometric tolerances, which limit the amount of geometric variation on part features. Tolerancing connects design intent and manufacturing outcomes by quantifying allowable inaccuracy on a manufactured part compared to the nominal, perfect geometry specified by the designer. Predicting achievable tolerances would aid designers by helping them determine if a particular AM process was capable of creating their part with the desired level of accuracy. However, most DFM tools do not adequately address tolerances [24]. Additionally, because AM is a new technology, there is uncertainty about the tolerances that are achievable by different AM processes [22]. Furthermore, the processes of tolerance analysis and tolerance transfer have challenges unique to AM [25]. If designers can understand what level of tolerance is achievable, they can consider changes to the geometry, the tolerance scheme, or the AM process parameters used to manufacture their part.

This thesis hypothesizes that geometric manufacturability analysis for AM is practical, accurate, and useful. The practicality of such systems is vital to consider. If geometric manufacturability analysis tools are not easy to use and feasible to implement, they will not be useful in practice. It is also essential that such tools meet the needs of designers. The second part of the hypothesis, which focuses on accuracy, is also crucial. The accuracy of the manufacturability estimates needs to be verified to ensure that the predictions are sufficiently robust to help designers understand trade-offs. Here, accuracy is measured experimentally. Finally, after verifying the practicality and accuracy of the tools, we will assess the utility of the tools. User testing and feedback from designers are used to assess the utility of the tools.

It is expected that by making more explicit and explorable connections between design and manufacturing, especially in the specifications regarding part geometry, designers can more efficiently analyze manufacturing trade-offs during the design stage. This thesis is a first step in confirming this supposition. As stated earlier, there is a dearth of research focused on engineering design. Findings of this thesis will help lay the groundwork for future

study to understand the utility of not just DFAM but also DFM feedback more generally, shedding light on what factors impact the efficacy of DFM feedback.

1.2.1 Assumptions

The tools described in this thesis evaluate individual components of parts that would be manufactured separately. The analysis of assemblies of parts is outside of the scope of this thesis (we assume assembly tolerances have already been allocated to individual parts). We assume that DFM tools would be used at a stage in the product development process when there is at least a rough computer model representing the design geometry.

As part of this thesis, some user testing was conducted with undergraduate students. We assume that students serve as representations of actual designers. Although students are likely most representative of novice designers, the student body included in the study had a wide range of experience levels with design and manufacturing, providing insights about more experienced designers.

The software tools developed in this thesis are meant to serve as a testbed to explore DFM. The ease-of-use may not be as high as for commercial software packages, but it is adequate to provide useful insights about how designers consider manufacturing constraints during the design process and how designer characteristics impact the usage of DFM tools.

The focus of this thesis is mainly on fused deposition modeling (FDM), primarily because of the accessibility of that technology. FDM printers are relatively cheap and easy to find on most college campuses and makerspaces (collaborative workspaces with a focus on enabling the general public to make new things). However, we posit that this thesis and the tools described here can be adapted to other AM processes and insights can be applied to non-additive processes as well.

1.3 Contribution

This overarching goal of this thesis is to enable the development of more effective DFAM software tools in the future. The research described in this thesis advances this goal in several key ways. The first contribution is specific to the state of the art in software-based DFAM feedback. This research moves the state of the art of DFAM feedback related to geometric accuracy and manufacturability criteria forward. Several models for predicting geometric accuracy and manufacturability of AM parts have been developed in industry (e.g., Meshmixer) and academia [26–30] but they have not been extensively evaluated. This thesis evaluates existing tools, comparing and contrasting the predictions and contributions of these tools. Also, we evaluate the accuracy of the previously developed tools and highlight opportunities for improvements based on experimental measurements.

Another contribution is the development and sharing of our software tools. The tools described here are open source and can be used by other researchers who want to conduct further DFAM research. Also, the tools can be used by designers, whether they be student

designers who are working on projects while in school or engineers working in industry. Hopefully, further use of these tools will enable better manufacturing outcomes for designers as well as an improved understanding of the DFAM process for researchers.

This thesis contributes some general insights into designer behavior for DFM for traditional manufacturing processes, as well. Although the trade-offs and phenomena captured in our tools are specific to AM, trade-off analysis is a desirable function for any DFM system. The results presented here improve our understanding of designer needs in regards to software-based DFM, through literature review, critical evaluation of existing systems, and user testing with one of our tools. In addition, the user testing described here highlights important interactions between designers and software-based DFM tools.

Taken together, the contributions of this thesis advance the current understanding of what attributes are necessary for effective DFM software tools. This improved understanding can be leveraged to make more effective DFM software tools in the future, that will help designers better understand manufacturing constraints and trade-offs early in the design process. Ultimately, improved DFM will enable concurrent engineering practices that result in higher quality products, lower scrap rates, and shorter product development lead times.

1.4 Thesis structure

This thesis consists of seven chapters. The second chapter reviews relevant literature to motivate and provide context to the research questions discussed above. Existing literature on topics such as DFM, DFAM, and design-support tools are summarized. Chapter 3 focuses on the material and methods of the thesis. Since the focus of this thesis is largely computational, this chapter details the algorithms and framework of the geometry analysis tools we developed. In Chapter 4, we evaluate the hypothesis that geometric analysis in support of DFAM is practical. The tools developed in Chapter 3 are applied to an example part, and their analyses are compared with existing DFAM software tools. Chapter 5 focuses on evaluating the hypothesis that geometric analysis in support of DFAM is accurate. This chapter summarizes experimentation that was conducted to evaluate manufacturability criteria and estimates of achievable tolerances. In Chapter 6, we employ user testing with the tool to test the hypothesis that geometric analysis in support of DFAM is useful. Chapter 7 concludes the thesis by discussing the research objectives and revisiting the main hypothesis. Particular areas in need of future research are also highlighted.

Chapter 2

Background and previous work

The goal of this thesis is to demonstrate that geometry analysis to assess the manufacturability of parts to be made using additive manufacturing is feasible, accurate, and helpful. In order to motivate and provide context for this goal, it is necessary to review the current state of research in related areas. This chapter provides a broad overview of design for manufacturing (DFM) and design for additive manufacturing (DFAM), before narrowing to more specific topics relevant to the thesis.

2.1 Design for manufacturing

DFM can refer to a diverse set of strategies that seek to help designers consider manufacturing costs and constraints during the design of new products. Some of these strategies are organizational or social, such as encouraging interaction between the often siloed teams focused on design and manufacturing. Some of these strategies are technical, such as the use of DFM software tools to assess product manufacturability [13, 14, 31].

As summarized in [14], benefits of applying DFM include: improving the quality of designs; allowing designers to understand trade-offs between product performance and manufacturing yield; reducing product development cycle time; reducing manufacturing cost and manufacturing cycle time; higher and more predictable manufacturing yields; and reduction in maintainability/serviceability efforts and warranty costs. With these benefits, it is easy to understand why many major companies now implement different DFM strategies [12, 13]. One recent study found that a DFM-focused redesign of a medical device reduced production time by 75% and production cost by 8% [32]. However, the adoption of DFM is not universal [12]. DFM usage is limited in some industries, like aerospace [9], and at smaller companies that lack the resources to put DFM strategies into practice [33]. There is more work to be done to understand the need for DFM in industry and to better address those needs through new DFM strategies and tools.

The terms “producability” and “manufacturability” were introduced in the 1960s to identify the relative ease of manufacturing a part [34]. In the late 1960s, Boothroyd, in conjunction with Redford and later Dewhurst, created guidelines to minimize the cost of

manufacturing and assembling engineered goods. In the subsequent years, researchers and companies have developed many different tools and methodologies to help designers apply DFM, as is summarized elsewhere [15, 16].

Different DFM tools and strategies are appropriate at different stages in the design process. For example, if applied during conceptual design, it is possible that DFM guidelines might be used by designers to develop the design concepts. If applied during concept evaluation, methods for developing a rough estimate of manufacturing cost maybe applied. If applied after detailed design, DFM can be used to re-design existing parts to alleviate issues that have arisen during manufacturing or to reduce manufacturing cost. Generally, more Design for X tools have been developed for later design stages (e.g., detail design) than earlier, more conceptual phases [16]. Design for X is the term used to identify the general category of tools that focus on helping designers achieve some goal, represented by the X, such as design for sustainability or design for assembly. On the spectrum of Design for X methodologies, DFM tends to be more specific and concrete and less conceptual. This characteristic is due, in part, to the fact that manufacturing cost and overall manufacturability is driven by material selection, the complexity and orientation of part features, and the manufacturing process chosen. In conceptual design, these variables are not yet set, and so it is difficult to apply DFM tools fully.

Historically, most DFM information was typically shared in the form of guidelines or heuristics [24]. However, starting in the 1990s, a number of systems were described that use geometry analysis to provide DFM feedback [35–37]. Because humans are better able to address the creative aspect of design, it makes sense to offload some rote tasks to a software tool that can automatically assess the manufacturability of a part. The efficiency of human decision making is negatively impacted by the complexity of the problem (in particular, the number of design parameters [38]). Complex problems can be difficult for designers, who may struggle to remember relationships between different parameters while trying to optimize competing design objectives, such as improving a part’s manufacturability while also decreasing part cost. Using computer-based tools to perform manufacturability analysis should help overcome the limitations of manual assessment by designers [31].

2.1.1 Software-based geometry analysis for DFM

Software tools and algorithms that perform automated geometry analysis can provide helpful design feedback, allowing designers to improve manufacturing outcomes. These types of systems are often referred to as manufacturability analysis systems. Manufacturability analysis systems automatically analyze part geometry to help designers consider manufacturing constraints, helping them to identify and rectify any potential manufacturing problems before finalizing their design [15, 39].

Many manufacturability analysis systems and algorithms exist for traditional manufacturing processes, as summarized in [15, 39, 40]. Analysis of the systems that exist for traditional manufacturing processes can provide insight into the potential impact of AM manufacturability analysis systems. There are some indications that geometry analysis for DFM can

improve designs, as opposed to a designer self-assessing their part using DFM rules or guidelines. Design-support tools can be helpful for designers because they can serve as an extension of designers short term memory [41], allowing them to focus on other more demanding tasks. Barnawal, Dorneich, Frank, and Peters [42] found that a 3D software tool helped designers to identify geometry leading to castability problems more effectively than providing the designers only with textual feedback. Riggs, Poli, and Woolf [43] also described preliminary findings that showed promise for teaching design for injection molding principles through a software package. Lynn [44] found that students who interacted with a simulation-based software tool were better able to recognize geometrical limitations of subtractive manufacturing processes compared with students who did not interact with the tool. It is expected that a DFAM manufacturability analysis system can provide a similar benefit to designers.

Before presenting a summary of DFAM, important considerations of any software-based design tool, such as a DFM software tool, will be discussed. These considerations will be used later to frame the development and evaluation of the DFAM tools presented in this thesis.

Considerations of computer-based design tools

Design, even when aided by computer-based design tools, is an intrinsically human activity. For several decades, researchers have explored “interactions between designers, between designers and computers, and between designers and users, especially in the early stages of design” [45]. The interaction between designers and computers is of particular relevance to this thesis and has been the focus of some prior research. Interactions between humans and other elements of a system, known as human factors, have been found play a crucial role in how designers use (or fail to use) design-support tools like CAD [46] and concurrent engineering tools [47]. Poor human factors can limit or prevent implementation of concurrent engineering design tools in industry [47].

Because of the relative lack of research on the efficacy of DFM, it is currently unclear how human factors may impact performance using DFM tools. However, insights can be drawn from the general field of human-computer interaction. Important human-computer interaction concepts relevant to the design of support systems are reviewed here. These concepts will inform the evaluation of existing DFM tools, described later in this introduction, and guide the development of DFM tools described in subsequent chapters of this thesis. The relevant concepts focused on here are prior experience, interactivity, and visualization.

Experience level of designers- Research suggests that a designer’s prior experience and domain knowledge can impact a designer’s performance on design tasks. Barnawal et al. [42] argued that novice designers would benefit especially from DFM systems because, as summarized by Ahmed, Wallace, and Blessing [48], they rely largely on trial and error in their design methodology. Therefore, a design-support tool that rapidly evaluates their design could help them more efficiently iterate and improve their design. Expertise in the design subject improves a designer’s ability to understand the coupling between design variables [49]. Sim-

ilarly, a designer’s domain knowledge can influence how a designer interacts and feels about a design- or decision-support tool [46,50]. In one study, experts preferred directly specifying the values of design parameters in a design-support tool, whereas novices preferred specifying the relative importance of their needs, letting an algorithm choose the best design for them [51]. These prior studies suggest that designer experience level will play an important role in terms of designer effectiveness and experience for software-based DFM tools, but this has yet to be confirmed. It is desirable to determine if software-based DFM tools are equally effective at helping novices and experienced designers alike to improve the manufacturability of their designs.

System interactivity- With increases in computing power and because of the human limitations discussed early, there is a tendency in software-based design tools to treat design problem as a mathematical problem that needs to be optimized. However, designers typically have a great deal of intrinsic knowledge that is difficult to build into an algorithm and a nuanced view of design trade-offs that is challenging to define in an optimization problem. As summarized in Chandrasegaran et al. [45], about 20 years ago, a chief engineer of Boeing stated, “Computers don’t design airplanes. We have not put the knowledge that is in the airplane designer’s head into AI [artificial intelligence] that balances all these objectives. But, someday we will continue to probably move to that. Right now, the knowledge to design airplanes is in the designer’s head.” While advances have been made using AI for mechanical engineering design, there currently is still a pressing need to have humans involved in the design process. In fact, the need for designer insight in design and design optimization processes was recently expressed by a panel of industry representatives at a design optimization workshop [52].

Design problems, such as DFM problems, often take the form of multi-criteria decision making, with several different design objectives that must be optimized simultaneously. Often, no single global optimum exists because there are conflicting objectives (e.g., maximize quality while minimizing cost), and so designers must find a compromise solution that best fits their needs. There are two general methods for the designer to convey the importance of the different objectives: a priori and a posteriori articulation of preferences.

Typically, a priori methods involve turning designer preferences into part of the algorithm, such as giving each design objective a scalar weight to reduce the multi-objective problem into a single-objective problem, but this process is flawed since it requires designers to state their preferences before they fully understand the relationships between different objectives or the range of values achievable for each objective. In a system with a priori articulation of preferences, a designer may be forced to make explicit decisions regarding the relative importance of different design variables, and this can cause a decision-maker to feel stress and decision conflict, which may negatively influence the adoption and utilization of the tool [53].

In a posteriori articulation of preferences, the designer is presented with a range of solutions, and through their interaction with those potential solutions, they formulate their preferences. As noted by Balling [54], “in many cases, people don’t know what they really

want until they see some designs.” Once designers have explored the range of options, they can find a single most preferred solution by either manually sifting through the data or by entering their now-fully-formed preferences into an optimization algorithm. This a posteriori articulation of preferences is in line with the findings of Brill et al. [55] who found that designers can best formulate their preferences when they are presented a few diverse solutions. Interaction with possible solutions is key to helping the designer refine their understanding of the design problem.

An important consideration in interactive systems is the delay between user input and system output. Real-time feedback about design decisions could be helpful for novices to understand and retain information about the impacts of differing decisions [56] (similar to the benefits of direct manipulation in software systems [57]). Specifically, in design-support tools, the delay in response time has been found to impact effectiveness [58]. For an interactive design-support tool like a DFM tool to be effective, it needs to be able to give the designer feedback quickly.

System visualization- In order for a software-based design-support tool to efficiently communicate a range of possible designs, and to allow a designer to interactively explore trade-offs among those designs, effective visualization of the design data is crucial. As summarized by Abi Akle, Minel, and Yannou [59], several research groups have built systems to allow designers to explore large, multivariate datasets [60–63]. Tweedie et al. [64] have also proposed other tools. While these systems have not been extensively validated for engineering design problems, one study found that a visualization system helped designers find a good solution from a large data set [65].

When selecting a single design out of a larger set of potential designs, whose key attributes can be summarized in a series of parameters, Abi Akle et al. [59] determined that a parallel coordinate plot, a visualization for multivariate data that uses parallel lines, was the most effective visualization to aid designers. However, the best visualization for one design problem may not be the best for another. For extremely geometric, space- and feature-dependent design problems, like DFM manufacturability feedback, it is difficult to summarize performance as a series of parameters. Although the visualization of DFM feedback has not been the subject of many studies, Barnawal, Dorneich, Frank, and Peters [42] showed that 3D graphical manufacturability feedback (features incompatible with casting) was more effective when presented in a 3D format, rather than textual or 2D feedback.

These attributes (experience level of designer, system interactivity, system visualization) are likely to influence designer experience and efficiency when using software-based, geometry analysis tools. In subsequent sections and chapters of this thesis, these attributes will be used to evaluate existing DFM tools and to guide the development of new tools.

Now, the focus of this thesis shifts to a particular process for which this thesis seeks to apply geometry analysis systems to assess manufacturability: additive manufacturing (AM). Basic concepts regarding AM were introduced in Chapter 1. Next, a summary of DFM methodologies focused on AM will be presented.

2.2 Design for additive manufacturing

There are many different sets of tools and guidelines to guide designer decisions regarding subtractive manufacturing processes, as discussed in the previous section. However, because AM is still maturing, there has been less time to develop guidelines and tools compared to other, more established manufacturing processes, like machining or injection molding. There is a need for more DFM methodologies focused specifically on AM [23]. All AM processes are similar in that they work by adding material rather than subtracting it and that they typically add material on a layer-by-layer basis. As discussed in Chapter 1, the focus of this thesis is on FDM, which uses a specific set of materials and physical processes, but we will discuss other AM processes, as well.

Different manufacturing processes use different materials, machine configurations, and physical processes for creating parts, which results in different achievable part features and tolerances. When designers are considering using AM to manufacture their part, they must consider a completely different set of constraints than if they were using any other traditional manufacturing process (e.g., injection molding). This fact is one main motivation for developing guidelines and methodologies that are specific to AM. Another reason for developing DFM specific to AM is that it is so unique and distinct from other manufacturing processes that it requires new, novel approaches to design processes and practices [17].

Design for additive manufacturing (DFAM) concepts can be roughly divided into three categories: “opportunistic,” “restrictive,” and “dual” [66]. Opportunistic DFAM seeks to inform designers about the geometric freedoms enabled by AM so they can take advantage of it in their designs. One reason that AM has achieved so much popularity is the level of geometric freedom it enables. Structures that would be infeasible to create using traditional manufacturing processes can be created with AM, giving rise to growing interest in topology optimization to create complex, lightweight structures. There is a growing need for DFAM research dedicated to helping designers explore and exploit this geometric freedom during the design of new products [19, 67]. A limited number of user studies have focused on opportunistic DFAM [66, 68–70].

Although opportunistic DFAM is an exciting area of research, it is equally important for designers to consider restrictive DFAM, which focuses on the limitations of different AM processes, so they do not create designs that are infeasible or prohibitively expensive to manufacture. Dual DFAM methods, which consider both the restrictions and opportunities of AM, provide a more useful, general framework for DFAM than opportunistic alone [66, 71]. However, dual DFAM methods are in their infancy because there is still a need to understand better the limitations of AM.

The focus of this thesis is on restrictive DFAM. As will be discussed in subsequent sections, there are still gaps in our understanding of what part geometry is suitable for manufacture using AM processes. In addition to clear guidelines about the limits of AM manufacturability, there is also a need to understand what level of quality AM can produce. Currently, designers must rely mostly on trial and error to determine if an AM process

can produce parts with the desired level of accuracy on key geometric features. Building our understanding of restrictive DFM through a better understanding of achievable part quality and geometric accuracy of parts created through AM will enable the development of dual DFAM tools and guidelines in the future. The following sections focus specifically on sources of geometric inaccuracy and summarize the most common ways that the geometric inaccuracy of parts can be quantified.

2.2.1 Sources of geometric inaccuracy in AM

Many diverse factors can influence the geometric accuracy of a part manufactured using an AM process. The material and machine used, the specific build method employed, the environmental conditions, and choices made by the designer can all impact the quality of the part. The sources of error for FDM are summarized in Fig. 2.1, an Ishikawa cause-and-effect diagram. Important sources of error can be divided into several categories: material, method, environment, machine, measurement, and human.

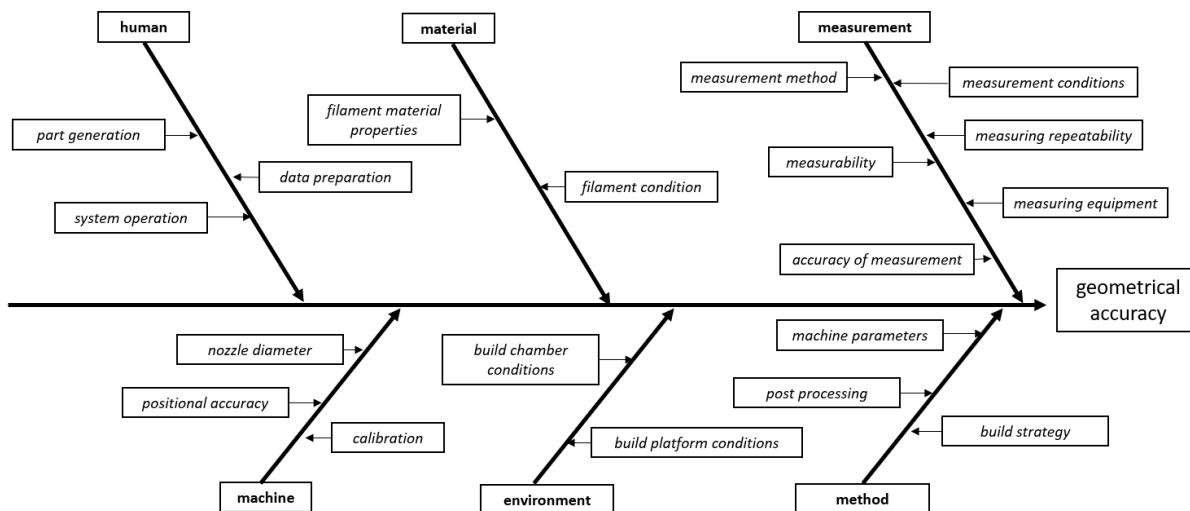


Figure 2.1: An Ishikawa diagram showing that a diverse set of factors impact geometric accuracy. (This figure was inspired by [72], with information also taken from [73].)

Material: The thermal properties of the material used, thermal gradients resulting from the build process, and geometry of the part interact, resulting in shrinkage and warpage of parts, which is especially problematic in larger parts [74]. Deformation due to thermal warping has been known to cause significant dimensional inaccuracy in AM parts [75]. For metal AM, residual stresses and warping are very problematic and have been the focus of some previous work [76,77]. In FDM, deposited material cools from extrusion temperature to the temperature of its surrounding environment. The material shrinks as it cools. However, the material adheres to the build platform and cannot shrink freely. Additional material

is deposited at the extrusion temperature on top of the cooling first layer, further creating thermal gradients in the part. In large parts, residual stresses caused by thermal gradients can lead to warping during printing or when the part is removed (Fig. 2.2). For FDM, PLA and ABS are common material choices. Due in part to its lower glass transition temperature, which results in less shrinking and warpage [73], PLA is more popular than ABS. Filament conditions, such as irregularities in diameter or exposure to high humidity, and filament color can also impact the quality of the print [78, 79].




Figure 2.2: Thermal stresses develop during deposition due to uneven cooling and constrained faces touching the build platform, causing warping, either before or after a part is removed from the build platform. Reprinted from *Robotics and Computer-Integrated Manufacturing*, 50, A. Armillotta, M. Bellotti, and M. Cavallaro, “Warping of FDM parts: Experimental tests and analytic model,” p. 141, Copyright 2018, with permission from Elsevier [80].

Method: The build method used significantly affects the accuracy and quality of the final part. As summarized by Yang and Anam [21], factors such as the speed at which the nozzle moves, the temperature of the nozzle and the build platform, raster pattern, intertrack gap distance, part orientation, and layer thickness [81–83, 83, 84], all affect the final quality of the part. The designer can adjust many of these factors.

Environment: Environmental conditions can also influence the geometric accuracy of the manufactured part. The temperature of the build chamber affects the accuracy of the part [85–87], possibly by affecting the shape of each layer [87]. Build plate and build chamber temperature are also important variables in determining the thermal history of a part [84, 88] and therefore also impact the amount of warping in the part.

Machine: The configuration and accuracy of the AM machine itself is also a key factor in determining the accuracy of the manufactured parts. Higher-end machines with more accurate positioning systems (such as controllers and motor-and-drive systems) can result in much lower geometric error [89]. The position of each layer, relative to other layers, is affected by the positioning precision of the machine [90]. The shape of the deposited track, which is controlled by the nozzle size and various machine settings, also effects geometric accuracy [21].

Human: Because the focus of this thesis is on helping designers make better DFAM design decisions, the branch of the Ishikawa diagram focusing on the effect of humans on geometric accuracy is most relevant. There are many choices made by designers that can affect their parts' geometric accuracy (Fig. 2.3).

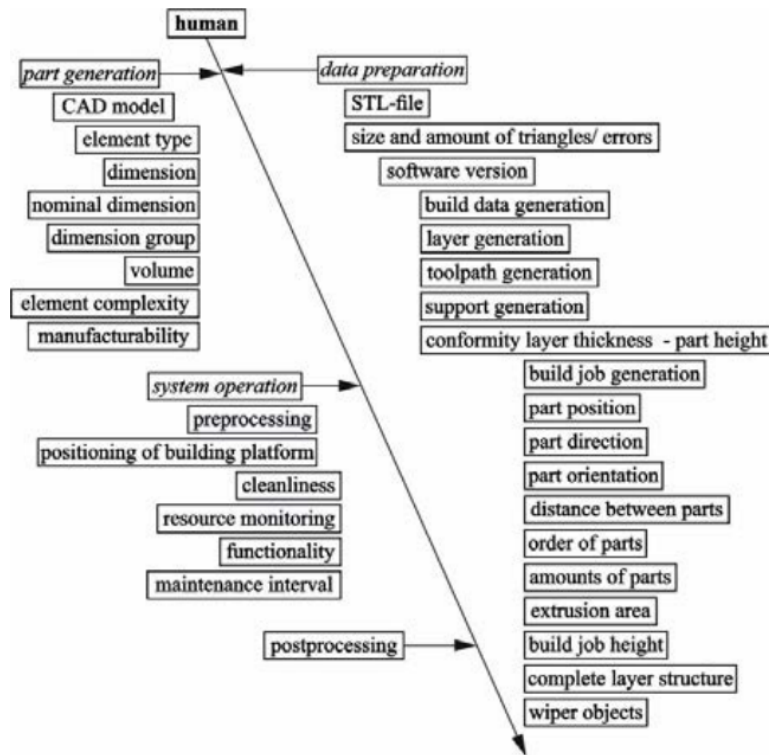


Figure 2.3: Human interaction is responsible for many sources of geometric error in FDM. Reprinted from Procedia CIRP, 43, T. Lieneke et al., “Dimensional Tolerances for Additive Manufacturing: Experimental Investigation for Fused Deposition Modeling,” p. 288 [91]. Reprinted courtesy of the Copy Right Holder under a Creative Commons License CC BY 4.0.

Under the part generation stem shown under the “human” branch in Fig. 2.3, there are many geometry-related sources of error. The size of the printed part (as measured by its volume or nominal dimensions) is important [85,92]. If it contains features below the minimum resolution of the printer, these features will not be printed or will print with

significant inaccuracy. Another important consideration is the manufacturability of the part geometry. If the part contains certain types of features or geometry, such as small bridges connecting larger geometry [29] or long faces that tend to warp [80], it will fail to print or will print poorly. The designer can limit the sources of error related to part geometry by adjusting the geometry of the part to achieve better outcomes.

Under the data preparation stem in Fig. 2.3, there is a long list of designer decisions that can impact part quality. One source of error is the creation of STL files: for parts with curved geometry, depending on the number of facets used, the part may be printed with tessellation errors [93]. Almost all of the other sources of error on this branch are related to part orientation. Build orientation can be set by either a designer or the manufacturer responsible for overseeing the manufacturing process. The build orientation cannot be considered in isolation; its effects interact with the part geometry and layer thickness used. One key way that build orientation, layer thickness, and part geometry interact to cause error is the stair-step effect. An example of the stair-step effect is shown in Fig. 2.4. This effect is clearly visible on printed parts (one example of this is Fig. 2.5).




Figure 2.4: Depending on the orientation of the part (determined by the angle between a face on the part and the build direction), the layer thickness, and the geometry of the shape, the stair-step effect can be very prominent.

Another effect of the interaction between layer thickness and part geometry is that if the part height is not a multiple of the layer thickness, the deviation of the height from the nominal dimension will likely be much higher than it would be otherwise [72]. The selection by the designer or manufacturer regarding printing process parameters can also have significant impacts. The position of the part on the build platform affects the accuracy of the part [72], possibly due to thermal effects [75]. Other variables, such as number and types of parts printed simultaneously, may also impact the accuracy of the part.

In summary, the designer has significant control over the quality of their print. To minimize geometric inaccuracy on a part, it is crucial that the designer understand the

manufacturing impacts of their design decisions. Now that the different sources of geometric error have been discussed, we will discuss some different ways in which geometric accuracy can be quantified.

2.2.2 Experimental measurements of geometric accuracy in AM

Surface roughness

The stair-step effect is a prominent feature on AM printed parts, and as a result, many previous studies have focused on measuring the stair-step effect using surface roughness measurements [94–100]. These studies are helpful for understanding the stair-step effect and what factors influence it.

Experimental studies have shown that the edge profile of each layer is not perfectly rectangular [95–98]. Approximating each layer as a rectangular prism, as shown in Fig. 2.4, is convenient, but not accurate. The shape of the edge of each layer depends on the specific AM process (e.g., FDM, SLA) used.

For FDM, the extruded filament tends to have a cross-section similar to that of an ellipse. This shape changes slightly, depending on orientation (Fig. 2.5). For stereolithography (SLA), the profile of each layer tends to be slightly triangular, as opposed to the rounded profile of FDM.




Figure 2.5: Sections of parts printed using FDM, showing the rounded edge shape of filaments. The shape changes slightly at different surface orientations, becoming more elongated at 90° . Reprinted from *Rapid Prototyping Journal*, 19(4), A. Boschetto, V. Giordano, and F. Veniali, “3D roughness profile model in fused deposition modelling,” p. 249, Copyright 2013, with permission from Emerald Publishing Limited [101].

The shape of each layer and the layer thickness size influence the surface roughness at different orientations. On surfaces printed using FDM, surface roughness tends to be lower for faces oriented around 90° from the build direction and larger for faces oriented close to 0° (upward-facing) or 180° (downward-facing) from the build direction (Fig. 2.6). When support material is present (typically for angles larger than 135° , depending on the printer), the surface roughness tends to be slightly higher. This phenomenon is due to gouging of the surface during manual removal of the support material, or due to residual support material that was not completely removed [102]. The type of support structure chosen can result in slightly different levels of surface roughness [103].

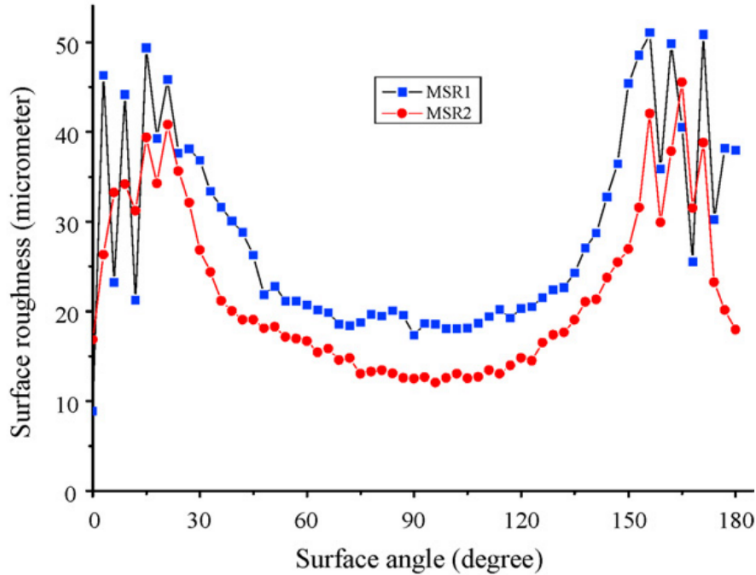



Figure 2.6: Surface roughness measurements for surfaces printed using FDM, with the surface angle indicating the angle between the normal of the printed face and the build direction. The first data set, labeled MSR1, had a layer thickness of 0.254 mm and the second data set, MSR2, had a layer thickness of 0.178 mm, indicating that roughness decreases with layer thickness. Reprinted from the Journal of Materials Processing Technology, *209*, D. Ahn, et al., “Representation of surface roughness in fused deposition modeling,” p. 5597, Copyright 2009, with permission from Elsevier [99].

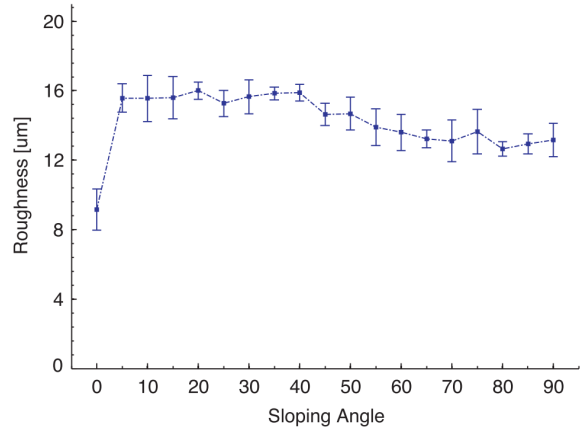
The relationship between surface angle and surface roughness is different for different AM processes. For SLA, the smoothest surface tends to be orientated around 120° from the build direction (Fig. 2.7i). For selective laser melting (SLM), the presence of partially bonded particles on the edges of layers results in relatively high surface roughness, even on surfaces oriented 90° from the build plate, when there is no error due to the stair-step effect. This trend can be seen in Fig. 2.7ii, with a slight decrease in error due to the decreasing stair-step effect at orientations close to 90° , but a relatively large roughness at 90° .

Process parameters can affect the surface roughness and the shape of the edge profile. Bacchewar, Singhal, and Pandey found that varying the laser power caused different edge profile shapes [104], which results in different relationships between surface roughness and face orientation. Similarly, different build styles and different materials were found to result in slightly different trends for how surface roughness varied with orientation (Fig. 2.6).

The surface roughness arising from other processes, including material jetting, has also been measured [94]. For most AM processes, the stair-step effect caused by layer thickness and orientation strongly affects the quality of the printed surface.



(i) SLA (the different series represent different build styles). Reprinted from *Rapid Prototyping Journal*, 31(1), P.E. Reeves and R.C. Cobb, “Reducing the surface deviation of stereolithography using inprocess techniques,” p. 25, Copyright 1997, with permission from Emerald Publishing Limited [98].



(ii) SLM. Reprinted from *Journal of Materials Processing Technology*, 213, G. Strano et al., “Surface roughness analysis, modelling and prediction in selective laser melting,” p. 591, Copyright 2013, with permission from Elsevier [100].

Figure 2.7: Experimental surface roughness measurements for surfaces printed using SLA and SLM, with the surface or sloping angle indicating the angle between the normal of the printed face and the build direction

Dimensional accuracy

Although surface roughness has been the subject of most studies that seek to evaluate the quality of parts made using AM, tolerances are the preferred method for controlling geometric errors in industry. Because no manufacturing process can produce parts that have perfect dimensional accuracy, tolerances are used by designers to specify the allowable deviation from the perfect nominal geometry they initially specified in their design. Tight tolerances for key dimensions can lead to better functionality, but also drive up the cost, as more precise and expensive processes must be used to achieve those tolerances. Tolerances play an important role in the product life cycle, enabling designers to specify a needed level of geometric accuracy and giving a framework for inspecting manufactured parts to ensure they will meet the designers’ requirements (Fig. 2.8).

There are different basic types of tolerances: conventional (which assigns a certain acceptable plus-or-minus range of deviation for each dimension on a drawing); and geometric (geometric dimensioning and tolerancing, commonly referred to as GD&T) [105]. GD&T developed out of a need by designers to have better control over the geometric variation of features, such as flatness, perpendicularity, circularity, and cylindricity [106].

Typically, designers in industry use GD&T, defined by standards established by international committees: ASME Y14.5 [107] in the US and ISO 1101 [108] internationally. These standards were established before the development of AM, and so are typically used to limit

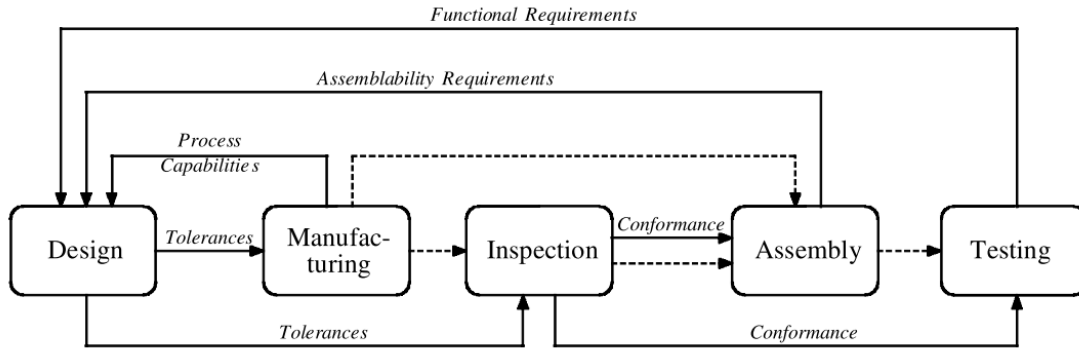


Figure 2.8: Tolerances serve an important role in communicating information between the design and manufacturing stages of the product life cycle. Reprinted from International Journal of Production Research, 40(11), Y.S. Hong and T.C. Chang “A comprehensive review of tolerancing research,” p. 2426, Copyright 2002, with permission from Taylor & Francis [105].

variation on parts produced using traditional subtractive manufacturing methods. Although these standards are mostly agnostic to the type of manufacturing process used and can be applied to parts produced using AM, there are some gaps related to new AM capabilities regarding novel geometry [25]. A draft standard, ASME Y14.46 [109], has been released to ensure that these gaps are filled. The same GD&T callouts defined in ASME Y14.5 and ISO 1101 can be used whenever possible for AM, with new additional AM controls added as needed. Because of its widespread use in industry, GD&T will remain the standard approach for specifying allowable variation on part features produced using AM in the future. Therefore, it is important to understand AM manufacturability in the context of tolerances and GD&T specifically.

Experimental work has been conducted to quantify the range of achievable tolerances for different AM processes. Mahesh, Wong, Fuh, and Loh [75] examined SLA, SLS, LOM, and FDM, printing a single benchmark part and measuring it to determine deviations from nominal dimensions. SLA was found to be very accurate and smooth on almost all features, while the other processes had lower accuracy and higher surface roughness. The range of achievable tolerances for FDM corresponds to IT-classes 11 through 16 [72, 89], similar to drop forging or casting, but less accurate than most basic machining operations. LS and LM have been found to have similar levels of accuracy to FDM [72].

The process of quantifying the achievable quality of AM by manufacturing and inspecting different features on a test artifact has been the focus of several studies. Some of these studies examined achievable geometric tolerances [21, 85, 86, 110, 111], while others examined achievable conventional plus or minus tolerances [72, 75, 86, 89, 91, 112]. Often, these test artifacts have several key attributes, in line with guidelines set forth by Richter and Jacobs, as summarized in [110]. The attributes are as follows: the parts include a variety of features (e.g., holes or thin walls) with different sizes, they are large enough to show spatial variation in printer performance, and they do not require too much time or material to manufacture.

An example of a test artifact for measuring the accuracy of AM processes is shown in Fig. 2.9.




Figure 2.9: A test artifact proposed by NIST researchers with different sized rectangular and cylindrical features, described in [110].

As summarized in the introduction to this section, process parameters can greatly affect geometric accuracy [21, 85, 92]. For FDM parts, Mahmood et al. found that the five most influential variables impacting achievable geometric tolerances were component size, extruder temperature, platform temperature, print orientation, and layer thickness [85]. Similarly, Minetola et al. found that part orientation [85] and machine used [89] significantly impacted the level of error associated with geometric tolerances like flatness (Fig. 2.10). These experimental results emphasize the point that was shown qualitatively in the Ishikawa diagram: many different variables, including several variables that designers have direct control over (e.g., component geometry, print orientation, layer thickness) can affect the quality of the parts manufactured using AM.

Experimental results indicate that part orientation influences the achievable geometric tolerances specified on that part, but most experimental research and development of test artifacts (e.g., the NIST artifact) measure geometric error at only a few orientations, typically with features oriented 0° and 90° from the build direction. One experimental study that examined error associated with geometric tolerances at different orientations indicated that flatness error can vary by as much as 100%, depending on the orientation used [21]. More study is needed to ascertain how geometric tolerances are influenced by a more complete set of build orientations.

Another limitation of prior work is that they focus on relatively simple geometry, like planar surfaces and cylindrical bosses. If a designer wanted to predict quality on a complex part with a variety of features and intricate geometry, there is currently no clear mechanism (besides repeatedly printing sample geometry of a desired part) to predict how the complex interactions between part geometry and process parameters would impact the achievable




Figure 2.10: There are different achievable tolerances at different orientations and for different FDM machines. Reprinted from *Procedia CIRP*, 41, P. Minetola, L. Iuliano, and G. Marchiandi, “Benchmarking of FDM Machines through Part Quality Using IT Grades,” p. 1031 [89]. Reprinted courtesy of the Copy Right Holder under a Creative Commons License CC BY 4.0.

geometric tolerances on the specific part geometry. There is a need to understand geometric accuracy on complex geometries, with overhanging features, curvature, and features orientated in different directions [86].

2.2.3 Design guidelines for additive manufacturing

In addition to the need for guidance on achievable geometric tolerances in AM, there is a need for general guidance and best practices for creating parts using AM processes [23, 113]. Many sets of design rules, principles, and guidelines have been developed in recent years to help designers understand best practices for DFAM, with varying levels of specificity [114]. Some of these have been developed in industry (as summarized by Thompson et al. [17]) by Stratasys, Materialise, Shapeways, and 3D Hubs [115]. Some sets of guidelines have also been developed in academia [116–120].

These sets of design guidelines typically list best practices with small pictorial examples of good and bad designs. One of these lists, the Design for Additive Manufacturing Worksheet, is shown in Fig. 2.11 as an example. The Design for Additive Manufacturing worksheet [116] was created by analyzing common guidelines cited in the literature and by consulting experts, and so it can be used as a summary of important guidelines. Briefly, the guidelines outlined

Design for Additive Manufacturing

A quick method for reducing the number of printing and prototyping failures, by Joran Booth
Instructions: Mark one for each category for the part you plan to print. Check daggers and stars first, then scores

Mark One	Complexity Simple parts are inefficient for AM	Mark One	Functionality AM parts are light and medium duty	Mark One	Material Removal Support structures ruin surface finish	Mark One	Unsupported Features Unsupported features will droop	Sum Across Rows	Totals
<input checked="" type="radio"/>	The part is the same shape as common stock materials, or is completely 2D 	<input checked="" type="radio"/>	Mating surfaces are bearing surfaces, or are expected to endure for 1000+ of cycles 	<input type="radio"/>	The part is smaller than or the same size as the required support structure 	<input type="radio"/>	There are long, unsupported features 	x5 =	
<input checked="" type="radio"/>	The part is mostly 2D and can be made in a mill or lathe without repositioning it in the clamp 	<input checked="" type="radio"/>	Mating surfaces move significantly, experience large forces, or must endure 100-1000 cycles. 	<input type="radio"/>	There are small gaps that will require support structures 	<input type="radio"/>	There are short, unsupported features 	x4 =	
<input type="radio"/>	The part can be made in a mill or lathe, but only after repositioning it in the clamp at least once 	<input type="radio"/>	Mating surfaces move somewhat, experience moderate forces, or are expected to last 10-100 cycles 	<input type="radio"/>	Internal cavities, channels, or holes do not have openings for removing materials 	<input type="radio"/>	Overhang features have a sloped support 	x3 =	
<input type="radio"/>	The part curvature is complex (splines or arcs) for a machining operation such as a mill or lathe 	<input type="radio"/>	Mating surfaces will move minimally, experience low forces, or are intended to endure 2-10 cycles 	<input type="radio"/>	Material can be easily removed from internal cavities, channels, or holes 	<input type="radio"/>	Overhanging features have a minimum of 45deg support 	x2 =	
<input type="radio"/>	There are interior features or surface curvature is too complex to be machined 	<input type="radio"/>	Surfaces are purely non-functional or experience virtually no cycles 	<input type="radio"/>	There are no internal cavities, channels, or holes 	<input type="radio"/>	Part is oriented so there are no overhanging features 	x1 =	
Mark One	Thin Features Thin features will almost always break	Mark One	Stress Concentration Interior corners must transition gradually	Mark One	Tolerances Mating parts should not be the same size	Mark One	Geometric Exactness Large, flat areas tend to warp	+	
<input type="radio"/>	Some walls are less than 1/16" (1.5mm) thick 	<input type="radio"/>	Interior corners have no chamfer, fillet, or rib 	<input type="radio"/>	Hole or length dimensions are nominal 	<input type="radio"/>	The part has large, flat surfaces or has a form that is important to be exact 	x5 =	
<input type="radio"/>	Walls are between 1/16" (1.5mm) and 1/8" (3mm) thick 	<input type="radio"/>	Interior corners have chamfers, fillets, and/or ribs 	<input type="radio"/>	Hole or length tolerances are adjusted for shrinkage or fit 	<input type="radio"/>	The part has medium-sized, flat surfaces, or forms that are should be close to exact 	x3 =	
<input type="radio"/>	Walls are more than 1/8" (3mm) thick 	<input type="radio"/>	Interior corners have generous chamfers, fillets, and/or ribs 	<input type="radio"/>	Hole and length tolerances are considered or are not important 	<input type="radio"/>	The part has small or no flat surfaces, or forms that need to be exact 	x1 =	
				Starred Ratings * Consider a different manufacturing process † Strongly consider a different manufacturing process		Total Score 33-40 Needs redesign 24-32 Consider redesign 16-23 Moderate likelihood of success 8-15 Higher likelihood of success		Overall Total <input type="text"/>	

Citation: The Design for Additive Manufacturing Worksheet, by Joran W. Booth, 2015. This work is licensed under the Creative Commons Attribution-NonDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nd/4.0/>.



Figure 2.11: The Design for Additive Manufacturing Worksheet. Reprinted from Journal of Mechanical Design, 139(10), J. Booth et al., “The Design for Additive Manufacturing Worksheet,” p. 3, [116]. Reprinted courtesy of the Copy Right Holder under a Creative Commons License CC BY 4.0.

in the worksheet are as follows:

- **Complexity:** The capabilities of AM are best used for complex geometry. Simple shapes may be better manufactured using a different process.
- **Functionality:** AM parts are not very strong and may fail if loads are applied.
- **Material removal:** Overhanging features need to be supported. Support material is difficult to remove from small or internal features and may mar the surface it was used to support.

- **Unsupported features:** Overhanging features, especially large features, need to be supported or they will droop.
- **Thin features:** Features that are smaller than 1-2 mm are likely to break. Further, features that are smaller than the printer resolution may fail to print altogether.
- **Stress concentration:** Sharp corners may serve as stress concentration. Fillets or chamfers also tend to be printed with better geometric accuracy than sharp corners.
- **Tolerances:** Parts will not print with exact, nominal dimensions. For parts that interface with other parts, this inaccuracy must be considered.
- **Geometric exactness (warping):** Large flat faces printed on the build platform tend to warp.

Outside of design guidelines, designers can also refer to academic research that discusses geometric limitations [121–123] and the expected dimensional accuracy of different AM processes, such as the research summarized in the previous section.

2.2.4 Existing software-based DFAM tools

Although many guidelines have been developed, there are few widely available automated tools for AM. The next three subsections will briefly summarize the tools that do exist, focusing on tools that assess AM manufacturability of parts and tools to help choose a build orientation. Other computer-based DFAM tools that do not fit into this category will also be discussed.

Manufacturability assessment

Several software-based versions of DFAM guidelines exist, using a GUI to guide a designer through a list of guidelines [117, 124, 125]. However, it is unclear if these computerized versions offer any significant advantage over lists of guidelines that are conveyed online or in worksheets.

Some software tools or algorithms exist that can automatically assess a specific part geometry to determine if it meets particular DFAM guidelines mentioned previously. Some algorithms have been developed that detect features on a part that are smaller than the size recommended in AM design guidelines [26–30, 126]. Some commercial alternatives with this capability include Shapeways and Meshmixer. Several tools visualize support material to help the designer assess if the support material placement is acceptable at a given orientation. [28, 127–129]. Similarly, Cura and other commercial tools highlight overhanging faces. There are only a few systems that can evaluate compliance with several design guidelines. Meshmixer can identify overhanging features, flag small features, and can examine part stability in different orientations. Ranjan et al. proposed a system to calculate a producibility

rating based on the existence of problems like small features, rough surfaces, and support material [30].

In summary, there are existing solutions that have the capability to flag one or two DFAM guideline violations, like highlighting features that are below a minimum feature size or flagging overhanging faces. However, these tools generally cannot evaluate a wide range of guideline violations, which limits their utility.

Orientation selection

A large body of research has focused on optimizing process parameters of additive manufacturing processes to ensure the best outcome. The outcome chosen varies widely (cost, time, part quality, etc.). Commonly, researchers have developed tools that evaluate different build orientations and recommend a particular orientation.

Some efforts have focused on minimizing cost or support volume [129, 130] and do not consider geometric error, despite the fact that geometric error is a key consideration of designers choosing between different process plans. When geometric error is considered, it is often evaluated as an average error metric. Some researchers [131–133] developed algorithms to optimize build direction to minimize surface roughness, while others [134, 135] have minimized other average error metrics, such as volumetric error. However, these average error metrics are not directly connected to the system designers already use to specify allowable geometric error: GD&T. Some research has specifically analyzed the connection between AM errors and GD&T [90, 93, 136–138], allowing for integration with the tolerance allocation process.

The interactivity of orientation optimization systems ranges widely. Many systems only present the designer with a single “best” orientation. A limitation of giving a single orientation is that these systems assume that a designer has well-established and fixed preferences regarding the dimensional accuracy needed on particular features. As discussed in the sections on considerations for computer-based design tools, this does not accurately reflect designers’ reality, which is filled with ambiguity and continually evolving perceptions about trade-offs between design objectives.

Systems with more interactivity will visualize not just the objectives at the “best” orientation, but also other possible orientations. This approach is taken by Savage and Cheng, in a system that enables the designer to interactively browse through different orientations and the corresponding print time and support material associated with those orientations [139]. Another common approach is to use colors plotted on a Gaussian sphere to evaluate the objective values (such as error or cost) associated with different orientations. This visualization has been employed by researchers to show how different metrics vary at different orientations by Mao et al. (Fig. 2.12) and others [90, 102, 138]. Rinaldi [140] and Savage and Cheng [139] employed a similar visualization, but with the 3D Gaussian sphere surface projected into 2D.

Another approach is to visualize the objective value for a given build orientation as a color plotted on the surface of the part CAD model. The objective values and representative

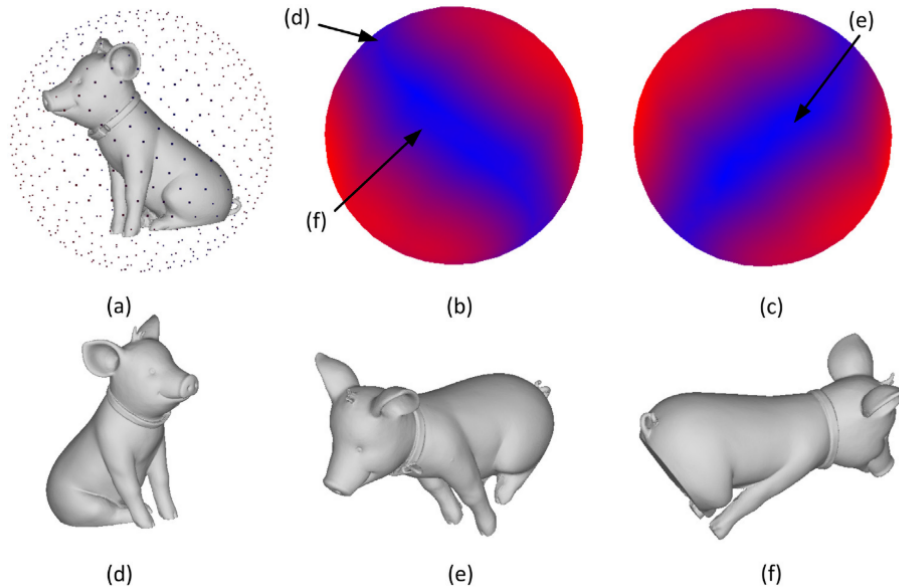


Figure 2.12: A random sampling of build orientations for an example part are displayed on the surface of a Gaussian sphere in (a), with quality metrics displayed using color on the surface of the sphere (b,c). The part is displayed in several build orientations (d-f). Reprinted from *Computer-Aided Design*, 107, H. Mao et al., “Adaptive slicing based on efficient profile analysis,” p. 97, Copyright 2019, with permission from Elsevier [141].

colors are updated when the model is rotated to a new build orientation. Ahn, Kim, and Lee created a system to plot surface roughness [131] and Boschetto, Bottini, and Veniali [142] built a system for both dimensional accuracy and surface roughness (shown in Fig. 2.13).

While various orientation optimization schemes exist in the literature, many of these schemes have some limitations. There is an opportunity to design user-friendly systems that are built with an understanding of designers’ needs, and that reflect the considerations of design tools discussed in prior research. Regardless of the exact visualization employed, it is expected that effective data visualization of a range of possible orientations will enable the designer to refine their preference and select a build orientation that more fully satisfies their needs. However, this expectation has not been evaluated yet.

Other

Although tools focused on manufacturability and orientation selection are of the most relevance to this thesis, other software tools have been described in the literature that also evaluate geometry to predict or optimize some AM objective. Although it is outside the domain of traditional mechanical engineering design, the proceedings of various Association for Computing Machinery conferences contains a large number of papers focused on AM. By and large, the goal of AM-focused research published in ACM proceedings seems to be

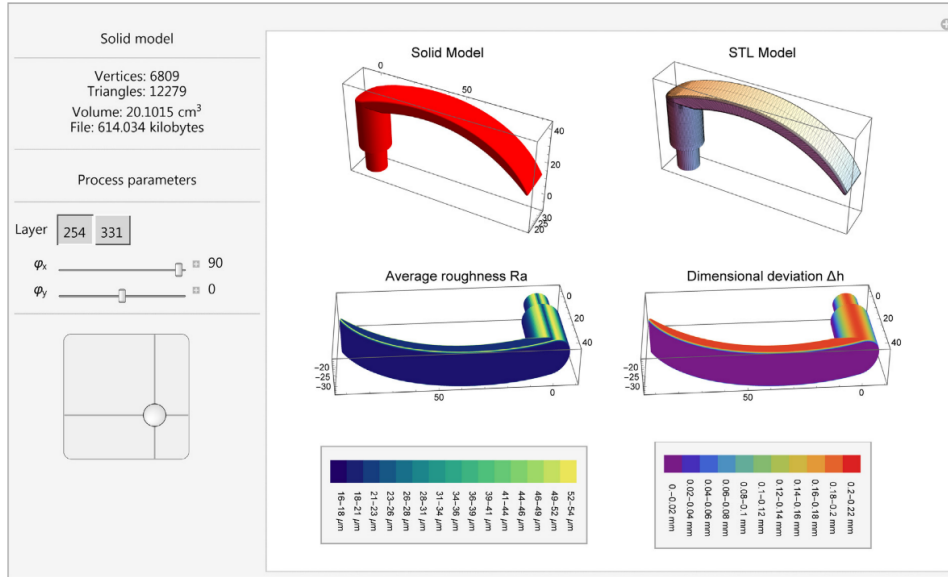


Figure 2.13: A GUI that allows the user to change the build orientation and view the expected surface roughness and dimensional accuracy of the part at that orientation. Reprinted from Additive Manufacturing, 12, A. Boschetto et al., “Integration of FDM surface quality modeling with process design,” p. 339, Copyright 2016, with permission from Elsevier [142].

to improve the appearance of non-functional parts. A few articles are summarized here to illustrate the range of research in ACM. Prévost et al. developed an algorithm to ensure decorative objects can be displayed after printing without toppling over [143]. Luo et al. developed an algorithm to determine how to partition large parts to be printed as separate sub-components [144]. Zhang et al. used user preferences to determine the least critical faces to add support material to ensure visually important faces were not marred during support removal [102]. Stava et al. morphed the geometry of AM parts to ensure drooping did not occur [145]. Most AM-focused research in ACM proceedings describes tools to perform automatic or semi-automatic shaping of non-functional, decorative parts. Part geometry is changed to ensure a better appearance with no regard for functional requirements (presumably because the parts are assumed to have little to no functional use). It is expected that, while some insight can be drawn from this research regarding good human-computer interaction, it is less relevant for mechanical engineering design research, where geometry and functionality are inextricably linked.

Another area of computer-based DFAM tools is the automatic generation of new parts to be created using AM. Goguelin et al. [146] developed a data visualization dashboard based on parallel coordinates that allow a designer to evaluate a set of automatically-created chair designs to be manufactured using AM. Topology optimization, which seeks to use algorithms to determine where to add or remove material within a preliminary part geometry outline, in order to withstand design loads with minimal material, also falls into the automatic design generation category. This category of tools is outside the scope of this thesis, which is

focused more on assisting designers in performing DFAM, rather than fully automating the design process. It is interesting, however, to note that several recent papers have begun to stress the importance of “feasible” topology optimization, which factors in some printability constraints like self-supporting faces [71, 147]. This trend speaks to the growing recognition in the field that understanding the limitations of AM is necessary to create cost-effective, feasible designs.

2.3 Summary

The number of design tools that have been created to support DFAM demonstrates that there is considerable interest in this area. However, most of these tools are limited in the metrics they evaluate and their practical utility to designers. Another limitation is that geometry analysis DFAM tools have generally not been evaluated to assess their effectiveness and how they impact the design process.

Two areas where geometry analysis in support of DFAM can be helpful are assessing manufacturability and predicting achievable quality. However, existing tools that assess manufacturability have a narrow focus on one or two types of manufacturability problems and do not address overall manufacturability. Also, while there is an understanding of sources of geometric inaccuracy in AM, this knowledge has not been used to build tools that accurately predict achievable geometric accuracy.

Many existing tools developed by engineering researchers are limited because they are generally not developed with a focus on promoting designer interaction or learning. Some tools, like those described by researchers in ACM CHI, are implemented with best practices from the field of human-computer interaction, but these tools do not adequately address metrics relevant to mechanical engineering design, like geometric accuracy. Because human factors and data visualization impact the effectiveness of software tools, there is a need to bring more human-computer interaction knowledge into the design of engineering design-support tools that use metrics designers care about.

Algorithms and frameworks developed by academic researchers are typically described in articles but are not implemented into tools that are shared so that other researchers or the general public can use them. Most commercial tools and academic algorithms have not been evaluated independently in the literature. Because many of the tools are difficult or impossible to replicate unless the source code is shared, we cannot build on or even evaluate many existing tools. Are the predictions accurate? Do designers actually benefit from using the tools? How does using software tools impact the design process? These are difficult questions to assess without a tool to use as a testbed. By developing, validating, and evaluating two testbed DFAM software tools, this thesis seeks to fill in gaps in our understanding of what characteristics enable effective DFAM software-based tools.

Chapter 3

Geometry analysis for design for additive manufacturing feedback

Previous chapters outlined existing work in the area of design for manufacturing (DFM) and design for additive manufacturing (DFAM), with a particular focus on software-based systems. The focus of this chapter is describing the novel tools that we developed. The development of these tools served two purposes: to address unmet needs of designers regarding predicting manufacturing outcomes of additive manufacturing (AM); and to test the basic hypothesis of this thesis.

3.1 Introduction

In additive manufacturing, process parameters can have a significant influence on the quality of the produced part, making it difficult to understand if a part can be manufactured successfully or with acceptable quality. Novices, unfamiliar with additive manufacturing, try to print parts with poor design, resulting in failed prints. One recent study found that close to 20% of parts students tried to print in a university makerspace failed due to poor design [116]. A manufacturability analysis system could help students avoid failed prints. Such a system could also be used by more experienced designers early in the design process, to help compare the difficulty of manufacturing different concepts. Later in the design process, designers need to develop tolerance schemes, limiting the amount of dimensional and geometric variation on different features on their parts to ensure that the functionality of their part is as acceptable. At this stage, design for manufacturing feedback should address what tolerances on key features are actually achievable, a feature that is not adequately addressed in most DFM tools [24]. Another problem is that, because AM is a new technology, designers are uncertain as to the level of tolerances that were achievable by different AM processes [22]. To address these needs of different designers at different stages in the design process, two software-based tools were developed: one that evaluates the manufacturability of a part to help prompt geometry or orientation changes to improve manufacturability, and another that performs a more sophisticated analysis of features to determine what geometric

tolerances are achievable for those specific features.

The first system is called Will It Print. Will It Print evaluates part geometry to check for compliance with DFAM guidelines. Specifically, the tool uses an STL file as input and analyzes the part geometry to evaluate the part for potential problems regarding warping, tipping, surface roughness, small features, and overhanging features. This tool helps designers evaluate the AM-manufacturability of their part and provides them with general suggestions for geometry and orientation changes to improve the ease-of-manufacturing of their design.

We also present a tolerance allocation system, which we call Tolerance Allocation for Additive Manufacturing (TAAM). Tolerance allocation is the process of setting the exact value of individual tolerances based on manufacturing cost or capability [105]. TAAM can rapidly analyze part geometry and predict achievable geometric tolerances (namely parallelism, perpendicularity, angularity, and cylindricity), based on layer thickness and build direction. TAAM can analyze multiple distinct features and their corresponding tolerances and datums to identify the build directions where all specified tolerances can be achieved. This tool can be used to select an optimal build direction and to analyze whether specified tolerances are manufacturable using additive manufacturing. In the development of TAAM, we use definitions and concepts defined by ASME Y14.5-2009 to quantify geometric deviations on additively-manufactured parts.

The theory and development between the two different tools are described in this chapter. Will It Print is presented first. Both tools use the same coordinate system, which is described in the next section.

3.2 Coordinate system

For both tools, we use a spherical coordinate system for our analysis, oriented to the normal vector, \mathbf{n}_f , of the face of interest. Following Arni and Gupta [90], the polar axis is aligned with \mathbf{n}_f ; the polar angle, θ and the azimuthal angle, ϕ , describe the orientation of the build vector, as seen in Fig. 3.1. The build vector, \mathbf{B} , is the direction normal to the build platform of the AM machine, orthogonal to the flat layers that are deposited.




Figure 3.1: Spherical coordinate system, adapted from [90]. \mathbf{B} is the build direction.

The polar angle is an important parameter because it controls the stair-step error and determines if support material is needed. As described in Chapter 2, for certain AM processes

such as FDM, support material is needed to support overhanging faces. Typically, support material is added when the angle, θ , between \mathbf{n}_f and \mathbf{B} is larger than some threshold angle, θ_{sup} , which is set by the slicer program. For FDM processes, θ_{sup} is often set at 135° .

3.3 Manufacturability assessment - Will It Print

In an effort to address manufacturability problems in parts designed by novices, we developed a MATLAB tool that we titled “Will It Print.” Novices tend to employ a trial-and-error approach to design, implementing a design before completing much or any analysis on their design, whereas more experienced designers will perform more analysis before implementation [48]. The framework of a tool to help novices assess manufacturability would need to be compatible with this trial-and-error approach when addressing common errors. Will It Print improves upon the trial-and-error process we observed students to take while printing. They would print, realize that their chosen orientation was not satisfactory, change geometry or orientation, and then reprint. We wanted to reduce the cycle time and eliminate wasted material from failed parts by giving them manufacturability feedback by flagging poor orientation choices and other common errors *before* they physically printed anything.

In order to determine what manufacturability guidelines to evaluate, we sought to categorize common problems with printing. We reviewed commonly cited AM-manufacturability guidelines (summarized in Chapter 2). Additionally, we reviewed parts that had been printed by students in our university makerspace who were enrolled in a sophomore-level introduction to manufacturing and tolerancing course. We evaluated common issues described by the students. Most common issues related to a combination of orientation and geometry, such as poor surface finish due to the stair-step effect or overhanging faces (Fig. 3.2).

Based on the evaluation of existing guidelines and observations about student errors, we decided to address the following DFAM guidelines:

- Tall parts with a small base should be avoided because they tend to tip.
- Small features should be avoided because they tend to fail to print or print with poor quality.
- Surfaces should generally be oriented normal or parallel to the build plate, or else they will have poor surface roughness.
- Long faces on the build platform should be avoided because they tend to warp.
- Support material should be avoided because it mars faces and is difficult to remove.

Most of these guidelines are impacted both by the part geometry and the build orientation selected, and so we wanted to develop a tool that evaluated compliance with the guidelines for the designer’s specific part while allowing the designer to explore different build orientations.



Figure 3.2: Students described a variety of issues relating to orientation, geometry, and printer settings. The reported issues are grouped into categories, and the number of students reporting each issue is listed.

In order to communicate whether the part geometry and orientation were not in compliance with a guideline, we decided to implement a system of warnings. The designer uploads the part geometry information in the form of an STL file, Will It Print conducts a series of analyses to evaluate the part and orientation for compliance with the guidelines mentioned above. As the designer clicks on the button in the main GUI for each individual analysis, a visual warning is issued if the part and orientation were not compliant. The designer can access the results of each analysis, one by one. For example, the designer can check to see if Will It Print highlights any part features as being too small. Then, the designer can move on to the tipping analysis to determine if the geometry and current build orientation result in a tipping warning. If Will It Print issues a warning, they can use the tool to rotate their part and evaluate the new build orientation. The warnings, which are graphically plotted in a GUI (more details of which will be described later), quickly communicate compliance information to the designer. By evaluating each guideline separately, on demand from the designer, we avoid inundating the designer with too much information. This approach of evaluating compliance at each new orientation the designer is interested in mimics the typical

workflow of novices to some extent: they decide on some combination of part geometry and orientation, print it, and see if their part was successfully printed. If not, then they try a new combination.

The algorithms used to calculate warning for out-of-compliance geometry and orientation are described in detail in the following subsections. Warnings that are calculated on a per-facet or per-feature basis are denoted ξ , and warnings calculated on a per-part basis are denoted κ . In the next chapter, an example design scenario will be presented to illustrate the typical workflow of Will It Print, and to compare Will It Print to other, existing solutions. The focus of this chapter is to provide the technical background of the tool itself. The algorithms used to evaluate compliance with the chosen guidelines (tipping, small features, surface roughness, warping, and support material) will be discussed.

3.3.1 Tipping

Some students had parts fail to print that were relatively tall, with a small base (i.e., the part area in the first layer, printed directly on the build platform). The parts would tip over or peel off of the build platform at some point in the print, and the print would fail, resulting in a mess of filament (Fig. 3.3).




Figure 3.3: The small base cannot hold the part in place when the applied moment due to tipping forces becomes larger.

Based on our observations, we grouped these errors into a common category, which we called tipping or wobbling. Tipping or wobbling errors are caused when the base of a part is not large enough to resist its own weight or the reaction forces of the extruder, and it becomes either partially or totally dislodged from the build plate. To avoid tipping, we implemented a check to determine if the combination of part orientation and geometry was such that the base was large enough to prevent especially tall parts from tipping over.

Tipping does not generally occur when the part has a large area on the build plate. Taller parts were observed to experience tipping more often because the larger distance from their center of mass to the build plate means that they are more likely to tip from even

small displacements. We approximated the risk of tipping by an experimentally determined threshold based on the ratio of the height of the part to the base area. The area on the base is calculated in a several step process.

We use ray shooting (Fig. 3.4), shooting rays through a grid below the lowest z -height in the current orientation of the STL. Ray shooting allows us to find intersections with the part facets, indicating where the part is located in space at the current orientation. We use this method to identify the base area of the part, with a tolerance of 0.05 mm to account for rounding errors in positions of the facet vertices. Each individual, disconnected area, a_i , making up the total base area, can be found using the built-in MATLAB function `polyarea`. If support material is present, the supported area, a_j^s , where support material is attached to the build plate, should also be included in the sum. The area where support is attached to the base plate is also found via ray shooting at the part's current orientation, checking for which facets on the part have a normal vector where the θ is greater than θ_{sup} . To ensure only the facets that are supported by support material that touches the build plate are included, we only consider the facet with the lowest intersection point from each ray.




Figure 3.4: Rays shooting (rays shown in yellow) to find each base area, a_i , and supported area, a_j^s .

While the base area of the part is completely filled in with filament, the supported area is only partially filled with support material. To account for this fact, the sum of the support area is multiplied by the support density, η_s , which varies between 0 and 1. The exact value of the support density term depends on the slicing program used to create the support structure. The total area in contact with the base plate, A , can then be found, where n_f is the number of disconnected areas touching the build platform and n_{sf} is the number of areas requiring support area:

$$A = \sum_{i=1}^{n_f} a_i + \eta_s \sum_{j=1}^{n_{sf}} a_j^s. \quad (3.1)$$

The total height of the part’s center of mass, measured normal to the build platform, H_c , is calculated from STL vertices, rotated to the current part orientation. The tipping warning, κ_{tip} , is turned on when the ratio of the area on the base plate, A , to the height of the center of mass, H_c , was less than some experimentally determined threshold, T_w :

$$\kappa_{tip} = \begin{cases} 1, & \frac{A}{H_c} < T_w \\ 0, & \frac{A}{H_c} \geq T_w. \end{cases} \quad (3.2)$$

We set T_w to be equal to 1.5 based on experimental measurements. This value was set by determining the ratio of $\frac{A}{H}$ from several student prints that failed due to tipping. T_w can be updated for other printers and with more extensive experimental study. Using this approach, Will It Print will highlight the base area of parts where $\frac{A}{H}$ is less than T_w to inform the designer that the part is at risk of tipping. If the ratio is greater than the threshold, the designer is told the part is not at risk of tipping. An example of this flagging is shown in Fig. 3.5 for the part shown in Fig. 3.3 when it is at the orientation in the actual build orientation used. The designer would be warned that this is a risky orientation for the geometry because it would be at risk of tipping, matching the physical results we saw.




Figure 3.5: For the example part, Will It Prints issues a tipping/wobbling warning for the orientation where tipping was observed.

3.3.2 Small features

As summarized in Chapter 2, DFAM guidelines often include instructions to avoid very small features. These features will either fail to print because they will not be detected during

the slicing process, or they will not print well because the positioning system on the printer is not capable of making sharp enough movements or depositing filament to the desired accuracy, or they may break off afterward. One example of this behavior is seen in Fig. 3.6, which shows a part with many small features that failed to print because they were below the minimum feature resolution.




Figure 3.6: A model representing the Eiffel Tower, with missing latticework above each level and inaccurate geometry near the top of the tower. Image courtesy of Yilin Sun.

To evaluate whether small features existed for a given part, we employed the method described by Telea and Jalba [29], similar to the approach taken by Ghiasian et al. [126]. This method, which uses a top-hat transform to identify small features on a voxelized representation of the user’s part, will be briefly summarized here.

The STL representing the part geometry is transformed into a binary voxel model using the polygonal mesh voxelization of Aitkenhead [148]. The voxelization process produces sets of voxels representing the part, Ω , and the complement of Ω , i.e., the non-occupied space around the part, $\bar{\Omega} = \mathbb{R}^3 \setminus \Omega$.

Thin features are identified as regions of the part with a local thickness smaller than a threshold, τ , which is set to the minimum recommended feature size for the relevant printing technology. The resolution of the voxelization needs to be smaller than τ to ensure that thin features are adequately detected. Thin regions, Θ , of the voxelized part, Ω , are identified using a set dilation and erosion. In order to identify these thin regions, we define a 3D unit ball B , where the Euclidean distance is less than or equal to 1, following Telea and Jalba [29]:

$$B \equiv \{ \mathbf{x} \in \mathbb{R}^3 \mid \|\mathbf{x}\|_2 \leq 1 \}. \quad (3.3)$$

This unit ball can be scaled by some factor, s , and used to find the multi-scale dilation δ_B and erosion ϵ_B of Ω at scale:

$$\begin{aligned}\delta_B(\Omega, s) &\equiv \Omega \oplus sB \\ \epsilon_B(\Omega, s) &\equiv \Omega \ominus sB.\end{aligned}\tag{3.4}$$

Thin regions, Θ , of the voxelized part, Ω , are identified using a multi-scale top-hat transform, setting the scale to the threshold, $s = \tau$:

$$\Theta(\Omega, \tau) = \Omega \setminus \delta_B(\epsilon_B(\Omega, \tau), \tau).\tag{3.5}$$

Part features can generally be divided into positive (i.e., walls, extrusions) and negative (i.e., holes, extruded cuts) features. The voxel set, Θ , represents thin positive features smaller than the printer resolution (e.g., sharp exterior corners, thin walls, and small extruded features). To identify negative features smaller than the printer resolution, Φ (e.g., sharp interior corners, shallow hole features, and small holes), this process is repeated, reversing the roles of $\bar{\Omega}$ and Ω . MATLAB’s Euclidean distance transform function `bwdist` is used to identify both Θ and Φ . Next, voxel classification (based on voxel connectivity) and removal of one-voxel-thick “interface shells” that have no relevance for printability, are completed, as described by Telea and Jalba [29].

Although Telea and Jalba calculate and display metrics meant to convey the importance of each voxel in Θ and Φ , we choose instead to display all of Θ and Φ to the designer and to allow them to interpret the importance of the feature themselves. This decision was made to avoid having to interpret the relative importance of different features, which is difficult, and can vary for different designers, parts, and part functions. For each individual voxel, $\mathbf{x} \in \Omega$, the warnings for positive features, $\xi_{small,+}$, and negative features, $\xi_{small,-}$, are calculated as:

$$\xi_{small,+} = \begin{cases} 0, & \mathbf{x} \notin \Theta \\ 1, & \mathbf{x} \in \Theta \end{cases}\tag{3.6}$$

$$\xi_{small,-} = \begin{cases} 0, & \mathbf{x} \notin \Phi \\ 1, & \mathbf{x} \in \Phi. \end{cases}\tag{3.7}$$

To communicate the information to the designer, the part is displayed, and all voxels where $\xi_{small,-} = 1$ and $\xi_{small,+} = 1$ are highlighted (using different colors). If for all voxels, $\xi_{small,-} = 0$ or $\xi_{small,+} = 0$, the initial part is displayed without any highlighting. An example of the Will It Print output for small features is shown in Fig. 3.7, with voxels where $\xi_{small,+} = 1$ shown in green and labeled “eroded” and voxels where $\xi_{small,-} = 1$ shown in blue and labeled “filled in”. Material that is not smaller than the minimum printer resolution is shown in gray. In this example, only the base and level platform features are larger than the resolution. This prediction matches the physical results shown in Fig. 3.6, where almost every feature has a very rough appearance or was missing altogether because of their small size.

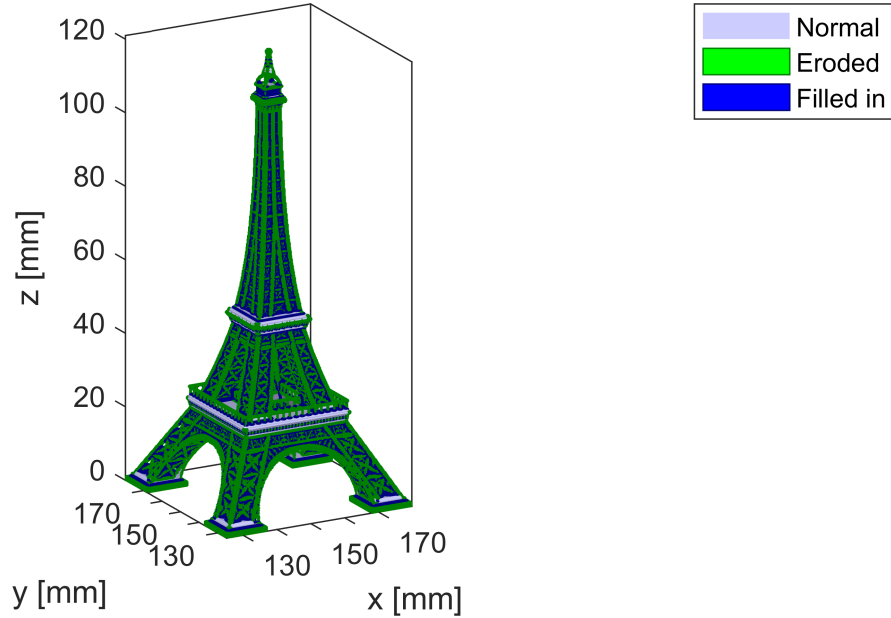


Figure 3.7: Will It Print identifies small positive features on the Eiffel tower model where $\xi_{small,+} = 1$ in green and small negative features where $\xi_{small,-} = 1$ in blue.

3.3.3 Surface roughness

Although designers in industry typically use GD&T to specify allowable geometric deviations on their part, novices are sometimes unfamiliar with GD&T. Also, tolerance schemes are usually developed late in the design process, after completion of detailed design. Because Will It Print was intended as a tool for novices or designers in early stages of design, before tolerance schemes were established, it is not advisable to quantify geometric error using GD&T. Surface roughness is an alternative measure to quantify inaccuracy caused by the stair-step error. Quantifying the stair-step error using surface roughness has two main benefits: it is familiar to designers, and it can capture stair-step error adequately because measurements of surface roughness are on the same order of magnitude as the stair-step error. Figure 3.8 shows a part with regions of relatively high surface roughness, where the stair-step error is large.

There are several different options for communicating what faces will have large surface roughness to a designer. One simple approach is to highlight any face where the stair-step error and therefore, the surface roughness, will be high. A high surface roughness flag, ξ_{surf} , is turned on for facets with high surface roughness (demonstrated in prior experiments, as seen in Fig. 2.6) when the angle between the facet normal and the build direction, θ , is in




Figure 3.8: The stair-step effect is clearly visible on curved surfaces (orange circles) where the surface normal is almost but not fully parallel to the build direction (shown in red).

the ranges:

$$\xi_{surf} = \begin{cases} 0, & 0 < \theta < \frac{\pi}{4} \\ 1, & \frac{\pi}{4} \leq \theta \leq \frac{3\pi}{4} \\ 0, & \frac{3\pi}{4} < \theta < \pi. \end{cases} \quad (3.8)$$

We implemented this approach in the version of Will It Print used in this thesis. Facets with high surface roughness are highlighted, and facets that are not predicted to have high surface roughness are not highlighted. For example, the same part shown in Fig. 3.8 was uploaded into the Will It Print tool, and the output for surface roughness detection is shown in Fig. 3.9.

An alternative approach is to estimate the actual surface roughness according to an empirical model. For example, surface roughness can be estimated according to equations derived from Thrimurthulu et al. [133], which matches the experimental trends seen by Boschetto et al. [96]. This model calculates surface roughness on a face based on the angle between the build direction and the face normal and the layer thickness. Once the surface roughness for every facet is estimated, the system can highlight any facets where the surface roughness on that facet is above some threshold, or similar to prior researchers, [131, 142], we can represent the magnitude of the surface roughness as a color on the surface.




Figure 3.9: Faces where surface roughness is expected to be high are highlighted.

3.3.4 Warping

Long parts tend to curl or warp when printed with a long face on the build plate. Figure 3.10 shows a relatively long part printed on a Type A Series 1 FDM printer. The amount of warpage is substantial enough to be obviously visible, with the bottom face curling up at the outside edges. For other printers, deviation due to warpage is much less substantial, in part due to improved adhesion between the part and the build platform. However, warping has been observed on parts printed on a wide variety of printers [80, 84, 149–151].

Warping is caused by the thermal gradients present during the printing process. In FDM, the filament is heated before being extruded and deposited on the build plate. The temperature of the build plate and the build chamber are typically below the temperature of the filament. As the extruded material cools to the ambient temperature, it experiences thermal contraction. However, the extruded material also adheres to the build plate (or to previous layers of material), which constrains contraction, and so stresses develop within the part. These thermal stresses cause deformation of the part, either during printing or after removal from the build plate. The deformation caused by thermal stresses, which we refer to as warpage, tends to be most evident on large, flat parts, but can be seen in a variety of parts.

Research into the relationship between part geometry and deviation due to warpage is still relatively young. However, a consensus in existing research is that the size of the part is a crucial factor [80, 149, 150]. Armillotta et al. [80] found that the length (i.e., the largest dimension) of a rectangular part had the strongest effect on the total amount of warpage deviation. Based on this research, we use the largest dimension of the part area on the build plate to estimate the likelihood of warping.




Figure 3.10: A truncheon part displaying warping, with the right side (orange circle) several milimeters above the center.

The area contacting the baseplate is found, following the procedure described in the Tipping section (including both supported area and base area). Each individual connected area, a_i and, if support material is present, the supported area, a_j^s , are identified. Because this area can be complex, it is necessary to find a way to estimate the “length” of the area, even when the area is not a rectangle. We estimate the largest dimension of the part by finding the dimensions of the axis-aligned bounding box of each connected area on the build plate, using `regionprops` in MATLAB.

Once the largest dimension for each area, Λ , is found, it is compared to an experimentally determined threshold, Λ_w . A warping flag, κ_{warp} , is turned on when:

$$\kappa_{warp} = \begin{cases} 0, & \Lambda < \Lambda_w \\ 1, & \Lambda \geq \Lambda_w. \end{cases} \quad (3.9)$$

If there are separate, disconnected areas on the build plate, their size is evaluated individually. The warping flag is turned on if any area’s largest dimension exceeds the threshold. For the example truncheon part, the largest dimension, including support material, is larger than Λ_w and so Will It Print highlights the base area and issues a warping warning (Fig. 3.11).

3.3.5 Support material

In FDM, support material can be difficult to remove fully. A user may find it impossible to remove all support material from internal cavities with limited accessibility from outside the parts, such as deep pockets. Also, during manual removal of support material, the part surface can be nicked and scarred (Fig. 3.12i), or for small or delicate features, the part feature may accidentally be completely removed from the part. Also, even with support material, overhanging features printed with FDM tend to droop slightly (Fig. 3.12ii).

We should warn the user when support material is needed and advise them about these possible repercussions. To do this, a support flag, ξ_{supp} , is turned on when support material




Figure 3.11: Will It Print highlights the base area of the part, indicating that the part is expected to warp.

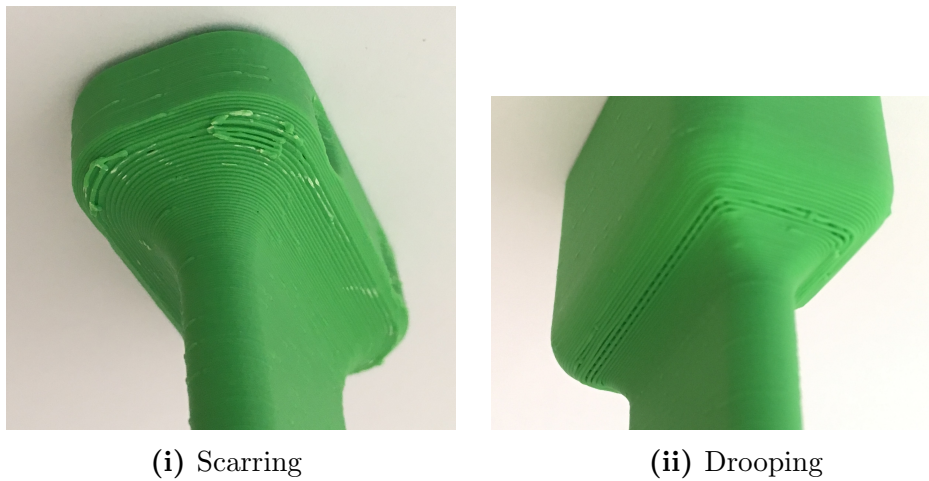


Figure 3.12: Example images showing issues associated with support material, with (i) scarring visible in white, surrounded by excess support material and (ii) drooping filament.

will be needed. Whether or not support material is needed is determined by θ for each facet in the current build orientation and comparing it to the support angle, θ_{sup} :

$$\xi_{supp} = \begin{cases} 0, & 0 \leq \theta < \theta_{sup} \\ 1, & \theta_{sup} \leq \theta \leq \pi. \end{cases} \quad (3.10)$$

All facets requiring support material are flagged. If the facet does not require support material, it is not highlighted. As an example, the same part shown in Fig. 3.12i was uploaded

into Will It Print. Faces that need to be supported are highlighted in yellow (Fig. 3.13), matching the physical results shown in 3.12i.




Figure 3.13: Overhanging faces where support material is needed are highlighted.

In future work, several capabilities should be added. First, we do not take into account the size of the overhanging face in Will It Print, but other design guidelines indicate that smaller overhanging faces are often self-supporting and therefore less of a concern. Second, we currently only highlight faces that are overhanging and do not highlight areas on the part *below* the supported face, even though support material will mar that surface as well. In the future, Will It Print should be updated to use ray tracing to highlight overhanging faces as well as part faces below overhanging faces.

3.4 Implementation of Will It Print

The interaction of a designer with the Will It Print tool is summarized in Fig. 3.14. Required inputs include machine parameter information, such as support material cutoff angles and an STL file representing the geometry.

Once the designer inputs the geometry and any other necessary parameters, the part is voxelized and small feature warnings are calculated. Because small features do not change with orientation, it is only necessary to perform this calculation once. The part is plotted in the initial orientation of the STL file. The designer can reorient the part, and view warnings related to all orientation-dependent guidelines (such as tipping and surface roughness). If the designer chooses, they can reorient the part (the loop seen in Fig. 3.14), and the warnings will be recalculated based on the new orientation.

At any time, the designer can choose to make changes to the part geometry. However, because we implemented the tool as a MATLAB-based GUI, the designer must make geom-

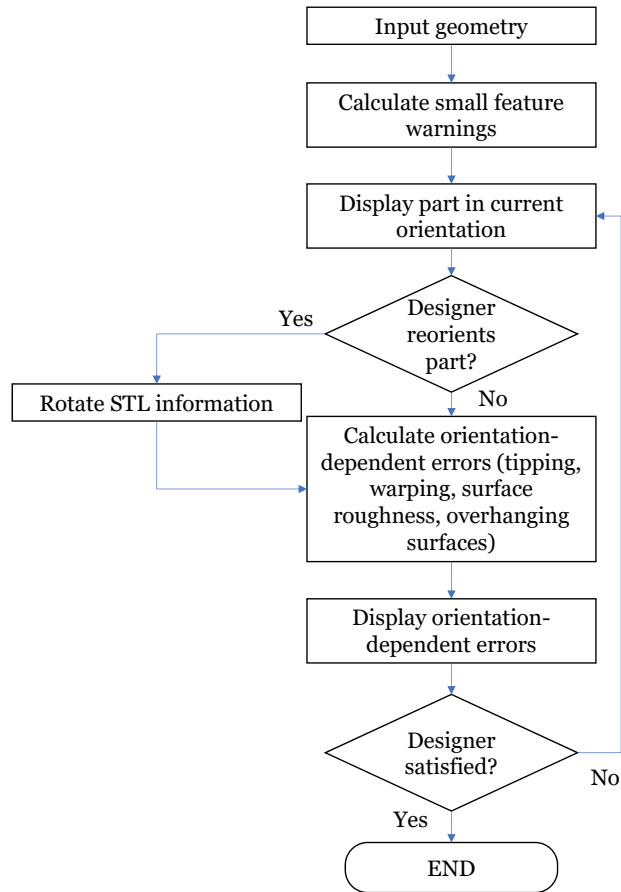


Figure 3.14: Process flow for user interaction with Will It Print.

entry changes in a separate CAD package, and reimport the geometry into Will It Print to evaluate the impact of their changes. In the future, Will It Print could be adapted as an add-in for common CAD packages.

In the current implementation of Will It Print, the thresholds used for issuing warnings about different guideline violations are set for FDM, but they can easily be customized to fit a different printer or AM process. Different printers and processes may require the evaluation of different guidelines, which could be added into the Will It Print tool.

GUI

We created a MATLAB-based GUI where the designer can perform all of the interactions outlined in Fig. 3.14. The GUI also has “Info” buttons next to each guideline, allowing the user to access a brief explanation of the guideline and general suggestions for how to change their part (Fig. 3.15). The information is displayed in a new window. An example of the

information window is shown in the next chapter. The information contained in each info window is listed in Appendix C, as rows in a worksheet that was adapted from the GUI.

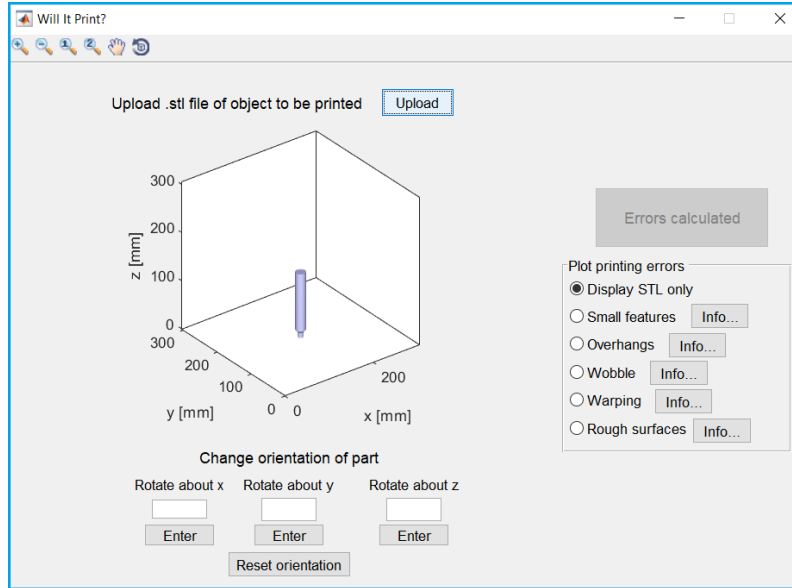


Figure 3.15: Will It Print GUI.

The part geometry is represented by plotting the STL faces using MATLAB’s built-in `patch` function to create filled polygons. Although the geometry is voxelized, plotting of each individual voxel is cumbersome and slow. The only result that is currently displayed as individual voxels is small features, which are plotted on top of the STL patches. Every other warning is plotted by changing the color of the patch in MATLAB.

The source code for this tool will be made available at the author’s GitHub site.

3.5 Analysis of achievable tolerances - TAAM

Now the focus of this chapter shifts to the other tool that we developed. The analysis of achievable tolerances was developed into a tool that we called “Tolerance Allocation for Additive Manufacturing,” (TAAM). Two main attributes make TAAM unique: its mathematical analysis of geometric accuracy and achievable tolerances, and its framework to enable designers to consider tradeoffs between cost and accuracy during the tolerance allocation process. Both attributes will be described in detail in this section, starting with the mathematical analysis.

As mentioned in the introduction for this chapter, TAAM was developed to help with the tolerance allocation process, where designers set the values associated with individual tolerances. Specifically, TAAM deals with orientation tolerances (i.e., parallelism, perpendicularity, and angularity), as well as flatness and cylindricity. Note: TAAM analyzes orientation

tolerances only on planar and cylindrical surfaces and planar datum surfaces. A designer specifies tolerances to control the amount of geometric variation on particular features that are key to the function of a design.

Flatness and orientation tolerance zones for planar features are defined by two parallel planes within which the toleranced feature must be contained. For orientation tolerances, the tolerance zone is oriented relative to one or more datums. Orientation error is defined here as the distance between two parallel planes, oriented by the datum, that contain every point on the toleranced surface. This error value can be compared to the tolerance value to determine if the as-manufactured feature is within tolerance. We denote this error as ϵ_{para} for parallelism error, ϵ_{perp} for perpendicularity, ϵ_{ang} for angularity, ϵ_{flat} for flatness, and ϵ_{cyl} for cylindricity. An example part with several specified tolerances is shown in Fig. 3.16, along with the error measurements for an example as-manufactured part.

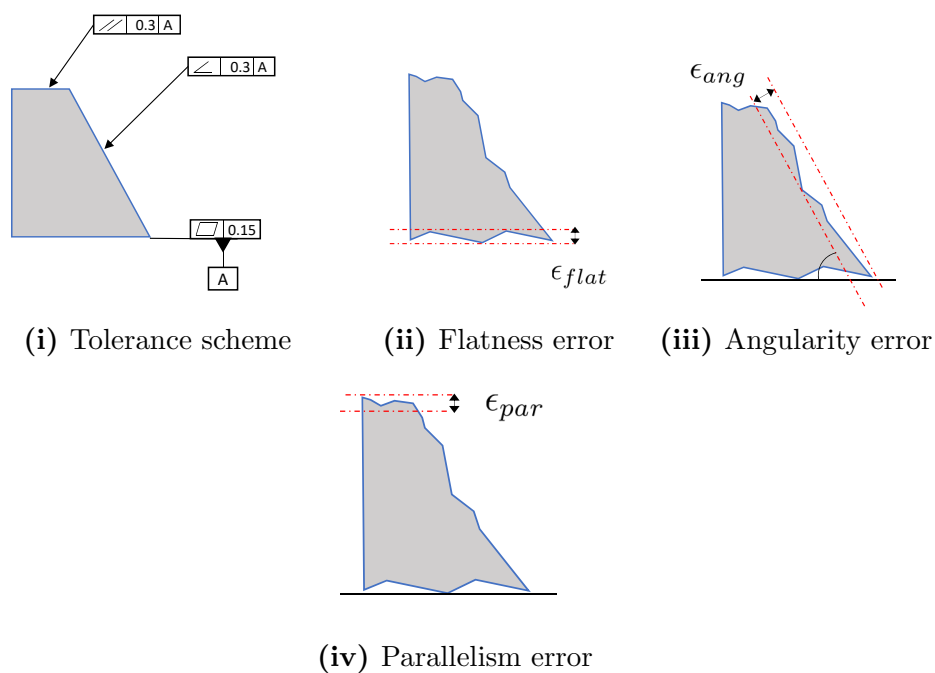


Figure 3.16: Example of a (i) tolerance scheme and a (ii-iv) demonstration of how the error associated with each tolerance would be found, as the distance between two parallel planes. For orientation errors, the planes are oriented relative to a datum.

The basic workflow of TAAM is that a designer uploads geometry information and their initial tolerance allocation into the tool. The tool evaluates the geometry and estimates the amount of error associated with each initially-specified tolerance, informing the designer whether their initially-specified tolerances were actually achievable given the printer settings and part geometry. If not all of the initially-specified tolerances are achievable, the tool helps the designer explore the trade-offs between manufacturing cost and quality to set new

tolerance values for each of their tolerances. The tool also helps the designer to select a build orientation that best satisfies their tolerances.

The mathematical analysis portion of this chapter will discuss the development of the mathematical model we use to estimate dimensional accuracy corresponding to tolerances on features of AM-printed parts. This section starts with describing the mathematical models that will be developed for orientation, flatness, cylindricity errors associated with tolerances of the same name. Next, the impact of various factors (such as layer thickness, layer deviations, and support material) will be discussed in detail. After this discussion, the refined mathematical model that we implement in TAAM will be presented. Then the visualization and multi-objective optimization used in TAAM are described.

3.5.1 Mathematical analysis of the impact of orientation, layer thickness, and layer deviations on geometric accuracy

Previous studies [85,96] have shown that orientation and layer thickness significantly influence the quality of the part itself. The effect these two factors have on geometry is more predictable than that of other process parameters, and, for this reason, we focus on these factors in our mathematical analysis. We have also incorporated analysis of some simple errors (related to deviations on each individual layer) into our mathematical framework, but other process factors can affect the quality of the produced part, such as large part size and material defects. Errors caused by process factors are difficult to predict because they tend to be specific to a given AM process, material, and even particular machine. The associated calculations can also be time intensive, making them impractical to calculate in near-real-time. For these reasons, we do not attempt to include them in this general framework. In this section, we develop mathematical models relating achievable tolerances (namely orientation, flatness, and cylindricity tolerances) to part orientation, layer thickness, and layer deviations.

Orientation tolerances

Most of the analysis that is summarized builds on Arni and Gupta's [90] work to come up with similar equations for orientation tolerances. The reader is referred to their paper for an in-depth description of how orientation and layer thickness impacts achievable flatness.

Errors associated with orientation tolerances depend on the size of the part itself. The length of the datum face and feature face in a particular direction, L_d and L_f respectively, can be calculated for $0 \leq \phi < 2\pi$ using the procedure that was described by Arni and Gupta [90] in the context of flatness error. Note that both the datum face and feature face lengths are therefore functions of ϕ even though, for notational convenience, we will refer to them as L_f and L_d . Length calculations are done separately for both the toleranced feature and the datum feature.

In order to derive equations to predict orientation tolerance error, a more complex coordinate system was needed. For parallelism, the error for each feature face can be described

using a single spherical coordinate system, oriented based on \mathbf{n}_f , because \mathbf{n}_f is parallel to the datum face normal, \mathbf{n}_d . For angularity and perpendicularity, \mathbf{n}_f and \mathbf{n}_d are offset from each other by some angle; therefore, two local spherical coordinate systems are used. A second spherical coordinate system like that applied to the feature face (as shown in Fig. 3.1) is aligned to \mathbf{n}_d . The polar angle between \mathbf{B} and \mathbf{n}_d , defined as θ_d , is shown in Fig. 3.17. For parallelism, $\theta_d = \theta$ because \mathbf{n}_d and \mathbf{n}_f are parallel. Because angularity and perpendicularity require two coordinate systems while parallelism only requires one, we will define a set of mathematical models specifically for angularity and perpendicularity and another mathematical model for parallelism.

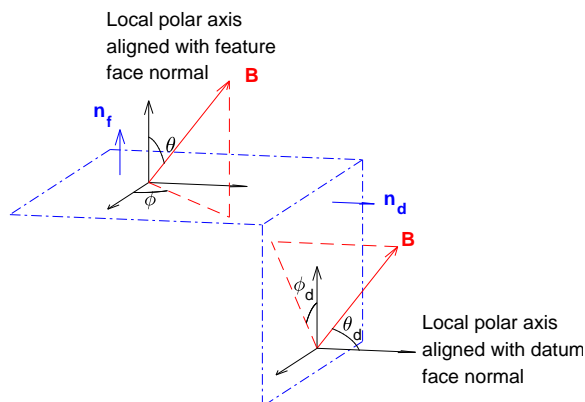


Figure 3.17: The local coordinate systems for the datum and a perpendicular feature.

To incorporate geometric deviations of individual layers, we adopt the method of Arni and Gupta [90], and model a zone of finite thickness along the z - and xy - directions of each layer. The value of this deviation zone, δ_{xy} (parallel to \mathbf{B}) and δ_z (perpendicular to \mathbf{B}), can be experimentally determined for a particular process and machine. The deviations are caused by a variety of factors, such as shrinkage and machine precision. For simplicity's sake, we model layer-level deviations as rectangular. To further simplify our analysis, we make an assumption regarding the error on the datum feature. Because a datum simulator used in orientation error measurement would only contact high points of the datum feature face, it is assumed that the layer-level deviation does not greatly affect datum plane orientation. Thus, the layer-level deviations are only modeled on the feature plane. This assumption will be evaluated in future work.

The first tolerance analyzed is parallelism. We divide the equations that capture parallelism's error into three regimes. The error is equal to zero when the build vector is aligned or nearly aligned with \mathbf{n}_f and \mathbf{n}_d . Both surfaces are approximated by one layer in this case, so the as-manufactured surfaces would both be parallel flat planes. The error in this regime, Regime 1, is equal to δ_z for parallelism, as shown in Fig. 3.18i.

The next regime, Regime 2, occurs when the datum face is approximated by at least two layers, and the feature face is approximated by one layer. As the angle between \mathbf{B} and \mathbf{n}_f increases, the projection of the datum face on \mathbf{B} will become large enough that the datum

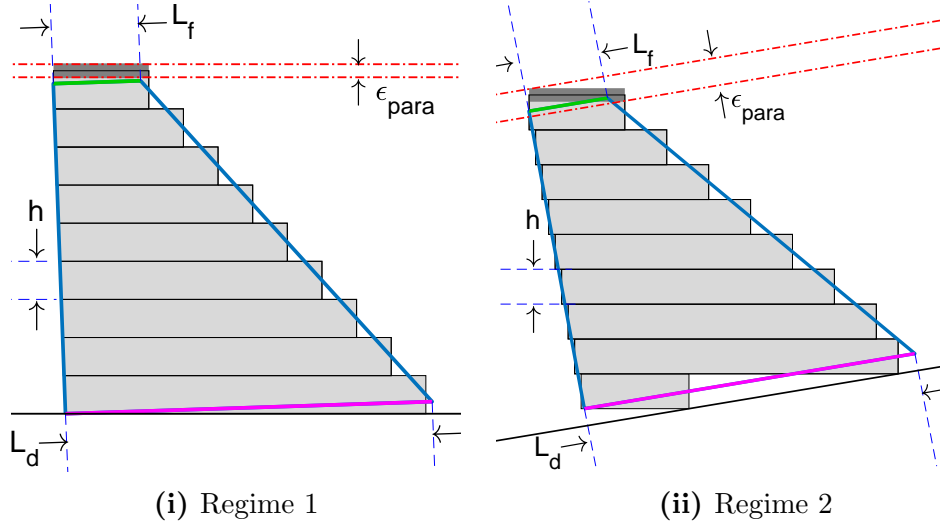


Figure 3.18: Parallelism error, ϵ_{para} . in (i) Regime 1 and (ii) Regime 2 is shown as the distance between two red planes, oriented by the datum plane (black line). The outline of the perfect theoretical geometry is also shown, with the datum face highlighted in pink and the feature face displayed in green. δ_z is shown in dark gray.

face is approximated by two layers. This transition occurs when the projection, $L_d \sin \theta$, is equal to $2h$ (where L_d is the length of the datum evaluated at a given ϕ). The polar angle at this transition, $\theta_{cr,1}$, can be calculated using Eq. 3.11. The error in this regime depends on the length of the feature, L_f , evaluated at a given ϕ , as seen in Fig. 3.18ii.

$$\begin{aligned} \theta_{cr,1}|\phi &= \arcsin\left(\frac{2h}{L_d}\right) \\ \theta_{cr,2}|\phi &= \arcsin\left(\frac{h}{L_f}\right) \end{aligned} \quad (3.11)$$

Equation 3.11 shows two distinct critical angles. The need for two angles is due to the different roles of the datum feature and tolerated feature. For the datum feature, two layers were needed to establish a datum plane parallel to the theoretically perfect feature. For the approximated tolerated feature, the layered geometry does not define a plane. Rather, it must be contained by two planes whose orientation is determined by the datum plane. As long as the projection of the surface onto \mathbf{B} is greater than h , the surface will be approximated by at least one plane, and the error follows the behavior of Regime 3. This regime switch occurs when θ is greater than $\theta_{cr,2}$, as defined in Eq. 3.11. If $\theta_{cr,2}$ is smaller than $\theta_{cr,1}$, Regime 1 dominates and Regime 2 will not be present. After this critical angle, the Regime 3 error is equal to $(h + \delta_z) |\cos \theta| + \delta_{xy} \sin \theta$. Regime 1 and 2 reoccur as θ approaches π . The errors and their corresponding ranges are shown in Tab. 3.1.

Angularity and perpendicularity follow a similar pattern to that of parallelism, but because of the offset between \mathbf{n}_d and \mathbf{n}_f , there are now two separate coordinate systems to

Table 3.1: Parallelism error in various ranges, evaluated at a given value of ϕ

Regime	Range	ϵ_{par}
1	$0 \leq \theta < \theta_{cr,1}$ $\pi - \theta_{cr,1} \leq \theta < \pi$	δ_z
2	$\theta_{cr,1} \leq \theta < \theta_{cr,2}$ $\pi - \theta_{cr,2} \leq \theta < \pi - \theta_{cr,1}$	$L_f \sin \theta + \delta_z \cos \theta $
3	$\theta_{cr,2} \leq \theta < \pi - \theta_{cr,2}$	$(h + \delta_z) \cos \theta + \delta_{xy} \sin \theta$

consider. Regime 1 occurs when \mathbf{B} is aligned or nearly aligned with \mathbf{n}_d . The behavior in this regime is identical to that in Regime 1 for parallelism, where one layer is used to approximate the slanted datum surface, resulting in an offset between the datum plane and the perfect intended face. The offset between \mathbf{n}_d and \mathbf{n}_f results in a different error than that for parallelism in Regime 1. As shown in Fig. 3.19i, the angle misalignment causes an error approximately equal to $|L_f \sin \theta_d|$, while the layer approximation of the feature face results in an error equal to $(h + \delta_z) |\cos \theta| + \delta_{xy} \sin \theta$.

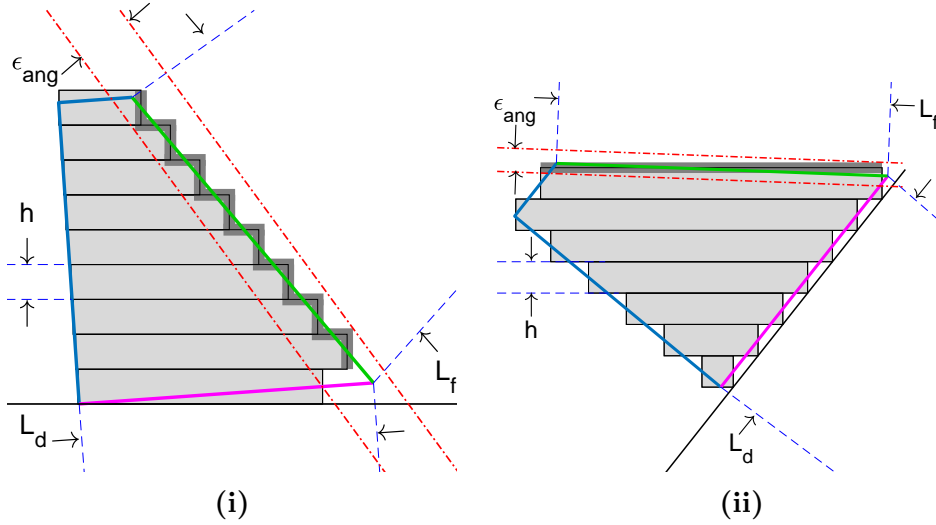


Figure 3.19: ϵ_{ang} in (i) Regime 1 and (ii) Regime 2. δ_{xy} and δ_z are shown in dark gray.

Regime 2 will occur when θ_d is greater than $\theta_{cr,1}$ but θ is less than $\theta_{cr,2}$. In this regime, \mathbf{n}_f is nearly parallel to \mathbf{B} . The slanted feature face will be approximated by only one layer and the error is approximately equal to $|L_f \sin \theta| + \delta_z |\cos \theta|$, as shown in Fig. 3.19ii. Both Regime 1 and 2 reoccur as θ approaches π . Regime 3 is the same as that for parallelism, where both features are approximated by several layers. The angularity or perpendicularity error as a function of the polar angle of the feature face, θ , and the polar angle of the datum face, θ_d , is summarized in Tab. 3.2.

Table 3.2: Angularity or perpendicularity error, evaluated at a given value of ϕ .

Reg.	θ_d Range	θ Range	ϵ_{par} or ϵ_{ang}
1	$0 \leq \theta_d < \theta_{cr,1}$ $\theta_d \geq \pi - \theta_{cr,1}$	$0 \leq \theta \leq \pi$	$L_f \sin \theta_d + (h + \delta_z) \cos \theta + \delta_{xy} \sin \theta$
2	$\theta_d \geq \theta_{cr,1}$	$0 \leq \theta < \theta_{cr,2}$ $\theta \geq \pi - \theta_{cr,2}$	$L_f \sin \theta + \delta_z \cos \theta $
3	$\theta_d \geq \theta_{cr,1}$	$\theta_{cr,2} \leq \theta < \pi - \theta_{cr,2}$	$(h + \delta_z) \cos \theta + \delta_{xy} \sin \theta$

Flatness

Flatness error is simpler than orientation tolerances because flatness callouts do not refer to a datum feature. No Regime 2 is needed. This mathematical model for flatness is described in detail in [90].

Table 3.3: Flatness error in various ranges, evaluated at a given value of ϕ .

Regime	Range	ϵ_{flat}
1	$0 \leq \theta < \theta_{cr,2}$ $\pi - \theta_{cr,2} \leq \theta < \pi$	δ_z
3	$\theta_{cr,2} \leq \theta < \pi - \theta_{cr,2}$	$(h + \delta_z) \cos \theta + \delta_{xy} \sin \theta$

Cylindricity

Cylindricity can be viewed as an extension of flatness tolerances, applied on a cylindrical surface rather than a planar surface. Cylindricity error, ϵ_{cyl} is the radial distance between two concentric cylinders that enclose all of the as-manufactured part's surface (Fig. 3.20).

We define a new angle, θ_{ax} , as the angle between the axis of the cylinder and the build direction. The model for flatness is adapted to accommodate this new angle, as seen in Tab. 3.4.

Table 3.4: Cylindricity error in various ranges, evaluated at a given value of ϕ .

Regime	Range	ϵ_{cyl}
1	$0 \leq \theta_{ax} < \theta_{cr,2}$ $\pi - \theta_{cr,2} \leq \theta_{ax} < \pi$	δ_{xy}
3	$\theta_{cr,2} \leq \theta_{ax} < \pi - \theta_{cr,2}$	$(h + \delta_z) \sin \theta_{ax} + \delta_{xy} \cos \theta_{ax} $

Other tolerances

There are additional GD&T callouts that were not analyzed in this work, including circularity, straightness, profile, location, and runout. However, other tolerances are likely to have a




Figure 3.20: A cross-section of a cylinder showing ϵ_{cyl} ($\theta_{ax} = 10^\circ$).

similar dependence on the layer thickness, the orientation of the feature face and datum face (if applicable), and various other errors. Adding additional tolerances is an area of future development.

Discussion of mathematical models

The mathematical models for orientation errors were created to accurately capture interactions between datum and feature face size and the number of layers used to approximate those faces. However, after developing these mathematical models, it became clear that the impact of the size of the datum and feature faces are most important for very small faces. In general, when the part faces are relatively large, the critical angles $\theta_{cr,1}$ and $\theta_{cr,2}$ are small, and so the regions associated with Regime 1 and 2 are small. The dominant error regime for all orientation tolerances is Regime 3.

For the rest of this thesis, we assume that the datum face and feature face are about 100 times much larger than the layer thickness in every direction, and therefore, $\theta_{cr,1}$ and $\theta_{cr,2}$ are less than one degree. Then, only the equations for Regime 3 and a simplified version of Regime 2 need to be used, if the error is calculated in one-degree increments (the approach we use). This assumption is not unrealistic: assuming a typical layer thickness of 0.15 mm, this assumption requires the datum and toleranced feature faces to both be larger than 17.25 mm in all directions. This assumption allows us to use the same equations for ϵ_{ang} , ϵ_{par} , ϵ_{perp} , and ϵ_{flat} .

3.5.2 Other factors impacting achievable tolerances

Other factors influence the expected geometric accuracy of AM-printed parts, including support material and the profile shape of individual layers. The impact of these factors, although less important than the impact of orientation and layer deviations, is still important to consider. These factors will be briefly discussed.

Support material

As summarized in Chapter 2, the removal of support material can cause scarring on overhanging surfaces to which the support material is attached. This phenomenon is also visible in Fig. 3.12i. To factor the impact of scarring due to support material removal into our model, we introduced a support material geometric deviation, δ_s , which is added to δ_z if a feature face needs support. Whether a feature face needs support is when the angle between the build vector and the face normal vector, θ , is larger than θ_{sup} . This cut-off angle θ_{sup} is process dependent, as is the value of δ_s .

Layer profile

To keep the model general for different AM processes, each individual layer is assumed to have a rectangular cross-section with a constant thickness. However, as summarized in Chapter 2, experimental studies have shown that individual layers are not perfectly rectangular [95–98]. We can customize the equations to reflect the profile of each layer better.

For FDM, the profile of each layer tends to be similar to an elongated half-circle (Fig. 2.5). To incorporate this profile shape into our model, we can model each layer as a half-circle (Fig. 3.21). While this is not a perfect representation of the shape of each layer, it is more accurate than a rectangular shape. In the next subsection, we update our previously derived equation to reflect this more accurate shape for FDM, as well as to include the effect of support material.

3.5.3 Refined error equations for FDM

The equations described in Tab. 3.4 and Tab. 3.3 can be updated to reflect the half-circle profile of FDM and the effect of support material (Equations 3.13 and 3.12). For the sake of simplicity, for the rest of this thesis, we will assume that all tolerances faces are adequately large compared to the layer thickness and therefore only the reduced set of equations is needed to describe orientation errors. These assumptions are reflected in the equations for




Figure 3.21: Flatness error, ϵ_{flat} , for a simple trapezoid. (The outline of the theoretical perfect geometry is shown in blue.)

FDM:

$$\epsilon_{flat}, \epsilon_{ang}, \epsilon_{perp}, \epsilon_{par} = \begin{cases} \delta_z & \theta = 0 \\ \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) |\cos \theta| + \delta_{xy} \sin \theta & 0 < \theta < \theta_{sup} \\ \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) |\cos \theta| + \delta_{xy} \sin \theta + \delta_s & \theta_{sup} \leq \theta < \pi \\ \delta_z + \delta_s & \theta = \pi \end{cases} \quad (3.12)$$

$$\epsilon_{cyl} = \begin{cases} \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) \sin \theta_{ax} + \delta_{xy} |\cos \theta_{ax}| & 0 \leq \theta_{ax} < \theta_{sup} - \frac{\pi}{2} \\ \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) \sin \theta_{ax} + \delta_{xy} |\cos \theta_{ax}| + \delta_s & \theta_{sup} - \frac{\pi}{2} \leq \theta_{ax} \leq \frac{3\pi}{2} - \theta_{sup} \\ \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) \sin \theta_{ax} + \delta_{xy} |\cos \theta_{ax}| & \frac{3\pi}{2} - \theta_{sup} < \theta_{ax} \leq \pi \end{cases} \quad (3.13)$$

3.5.4 Visualization of errors corresponding to tolerances

To illustrate the visualizations used in the system, we will introduce a simple example. The geometry of a jet engine bracket, provided by GE for the GrabCAD additive manufacturing redesign challenge [152], was evaluated. We used the original geometry provided by GE (Fig. 3.22), however, because the original design geometric tolerances were unspecified, all tolerances were inferred. The bottom face was chosen as a datum feature. Geometric tolerances were applied to three features, as shown in Fig. 3.22. We analyzed the achievable tolerances with layer thickness set to 0.15 mm. The value of δ_{xy} was set to 0.12mm and δ_z was set to 0.05mm, based on experimental information [97]. δ_s was set to 0.05mm and θ_{sup} was set to 135° .




Figure 3.22: GE redesign bracket [152] with tolerances added.

The error for each tolerance is calculated separately. The build vector orientation is varied from $0 \leq \phi < 2\pi$ and $0 \leq \theta \leq \pi$, and the error for each tolerance is calculated based on the equations summarized in Eq. 3.12 and 3.13. The computation time of this error is short as only simple trigonometry functions need to be evaluated.

We use color to indicate the magnitude of the error. This error map is similar to a visibility map [153], where the surface of the sphere is related to orientations of a vector, relative to a part. To make the color scaling easy to understand and universally readable, we choose the Viridis colormap for its perceptual uniformity, its readability in grayscale, and interpretability by people with various forms of colorblindness [154]. The error is plotted on a unit sphere, whose axes are aligned with the part axes. This visualization is shown in Fig. 3.23, showing 200 possible orientations, each represented as a colored dot.




Figure 3.23: Flatness error at different build orientations are shown as colored dots, plotted on a unit sphere around the coordinate axes of the STL file.

By calculating errors at a larger range of orientations (i.e., 10,000 orientations), we can create a denser visualization for the error corresponding to each callout (Fig. 3.24). These maps show the designer regions of orientations that would result in low error for one or more of their callouts. Angularity error is minimized if the build vector is exactly aligned with the feature normal, or if the build vector is orthogonal to the feature normal. This behavior is

also true for the perpendicularity error. The flatness error is oriented about the base plane. Flatness error is minimized if the build vector resides in the xy -plane or if it is parallel to the z -axis. The increased error due to support material removal on the feature face is visible at the bottom of the angularity error plot and the top of the flatness error plot.




Figure 3.24: Estimated error as a function of build vector orientation for example part.

While the error maps are helpful on their own, we also generate a plot that combines the information shown in each error map into a single visualization, showing how many tolerance callouts are met at each orientation. Feasible regions are found by identifying build vector orientations where all specified tolerances are met. We also plot all the orientations that do not currently satisfy the tolerances, using color to represent the number of tolerances satisfied, because a designer may decide to relax or change some of their initially-specified tolerances. Figure 3.25 shows a unit sphere containing all possible orientations of the build vector. Green regions denote orientations of the build vector that result in all specified tolerances on the example part being satisfied. If \mathbf{B} was aligned with the positive x - or y -axis of the example part (see coordinate system shown in Fig. 3.22), then all three specified tolerances would be satisfied, which would not be the case with the original build orientation (positive z -axis). Region A contains only a small range of orientations, which is why it is displayed as a small dot rather than a large area. Region B is much larger, giving a wide range of feasible orientations.

To help a designer choose a single specific orientation from this range of acceptable tolerances, we developed a framework that can visualize and help sort through trade-offs. This framework is discussed in detail in the subsequent section. Another benefit of the framework is helping a designer relax certain tolerances if they originally specified a tolerance scheme that resulted in no feasible regions (i.e., all tolerances could not be simultaneously




Figure 3.25: Feasibility regions for example part.

satisfied at a single orientation). This relaxation process will be described in the next two sections.

3.5.5 Tolerance allocation tool details

As summarized in Chapter 2, previous approaches to orientation optimization are limited because of how they deal with multi-objective optimization for AM. Typically, the relative importance of the different objectives (e.g., minimize error, minimize support volume, minimize build time) is assigned by the researcher or provided by the designer as scalar weights, which is not intuitive and can be complicated by the need to combine objectives with widely varying units. An approach called physical programming eliminates the need for scalar weighting of each objective [155]. Instead of a meaningless scalar weight, it asks the designer to break each objective into ranges of different relative desirability, which is more intuitive. Because we seek to create an interactive design tool, rather than a non-interactive orientation optimization algorithm, it is crucial that the input required from the designers is easy to understand.

Physical programming has been used as the basis of an interactive system [156]. Previous orientation optimization systems for AM are not interactive and assume that the designer’s preferences are well informed and inflexible, which is not consistent with the iterative, evolving nature of the design process. An interactive system that allows designers to iterate and refine their tolerance allocation while learning about the manufacturability of their design is still needed.

Physical programming has been applied to many different engineering problems in the past, as summarized by Ilgin and Gupta [157], but has not to our knowledge been utilized for tolerance allocation. As summarized by Chase [158], tolerance allocation often uses a theoretical, general equation to connect part quality and manufacturing cost. New methods

are needed for tolerance allocation for AM because it is possible to improve part quality by changing the orientation of the part without necessarily increasing the build time or cost.

The tool described here, in addition to enabling visual exploration of the data set, allows a designer to refine their preference structure interactively, and then uses an algorithm to present them with a few “best candidates.” This strategy is in line with the findings of Brill et al. [55] who found that designers can best formulate their preferences when they are presented a few diverse solutions.

The goal of our trade-off and tolerance relaxation/refinement tool is to elucidate the relationships between achievable geometric tolerances and other AM process parameters. This section presents a summary of the physical programming ranking used by our tool. Rather than asking designers to sort through thousands of data points to find the best option, we turn their preferences into minimization goals and constraints, which reduces the need for unnecessary user interaction.

Printing cost

In addition to the geometric errors, designers considering AM to produce their parts must also consider manufacturing cost. Predictive models have been developed to optimize part orientation to minimize cost [28, 134, 135]. In order to provide quick, real-time estimates of cost at thousands of orientations, we use cost estimates partially based on simplified estimates of build time. Following Armstrong, Barclift, and Simpson [128], with modifications to consider the volume of support material needed, the total cost to produce a part, \bar{C} , is calculated as:

$$\bar{C} = \rho C_m [\eta_s(V_s) + \eta_p(V_p)] + \bar{t}C_t \quad (3.14)$$

where ρ is the material density, C_m is material cost, η_s and η_p are support and infill density, respectively, V_s and V_p are the support material volume and part volume, respectively, \bar{t} is the total manufacturing time, and C_t is the cost per time charged for the manufacturing process. The total manufacturing time, \bar{t} , can be calculated as (adapted from [128]):

$$\bar{t} = \kappa \left(\frac{\eta_s(V_s) + \eta_p(V_p)}{vwh} \right) + \frac{H'}{h}(t_r) \quad (3.15)$$

where κ is a complexity factor, v is the extrusion rate, w is the width of the filament, h is the layer thickness, H' is the height of the STL file’s bounding box in orientation \mathbf{B} , and t_r is the additional time required to reset and move to the next layer. The total part volume, V_p , is calculated [159] and does not change with orientation. The support volume, V_s , does depend on orientation and is calculated as the product of the area and the average height of each supported facet. Currently, we do not determine if there is another part feature below the supported feature that would reduce the volume of support material needed, but in future implementations, we will use ray shooting to check for additional features and use this information to more accurately calculate support volume.

Rather than using the total estimated cost, \bar{C} , directly in our optimization, we instead calculate a cost parameter, λ . This approach is taken so we can provide generally applicable physical programming bounds in the next section. For a given orientation, λ is calculated as:

$$\lambda = \frac{\bar{C} - \min(\bar{C})}{\max(\bar{C}) - \min(\bar{C})} \quad (3.16)$$

where \bar{C} is the cost corresponding to that orientation, and $\max(\bar{C})$ and $\min(\bar{C})$ are the maximum and minimum costs for all orientations, respectively.

Physical programming

Messac introduced the basic approach of physical programming [155]. For our approach, we used linear physical programming [160]. In physical programming, instead of asking a designer to set weights to turn a multi-objective problem into a single-objective problem, the designer instead sets meaningful boundaries between different values of each objective, and these boundaries are used to establish preference class functions that are used to solve the multi-objective problem.

Physical programming includes several preference types, called preference class functions. If the goal is to minimize, the class type is 1; if the goal is to maximize, the class type is 2; and if the goal is to achieve a particular range, the class type is 3. There are soft and hard versions of each class, denoted S and H respectively, that can be used to characterize different levels of attainment by the designer. For each class function, a certain desirability is assigned to different ranges of objective values. For a soft class function in which the goal is to minimize the value of the objective, denoted 1S, the range of objective values is divided by boundaries, t_{is} where $s = 1, 2, \dots, 5$ and i is equal to the index of the objective in question. These boundaries separate what range of objective values are highly desirable (HD), desirable (D), tolerable (T), undesirable (U), highly undesirable (HU), and unacceptable. An example of a 1S class function is shown in Fig. 3.26.




Figure 3.26: Preference class function. Green dotted line indicates boundary of acceptable error, which we set equal to the initially specified tolerance, Δ .

Because a designer’s goal will always be to minimize error, the class function 1S is an obvious choice to represent each tolerance error minimization objective. The shape of this class function also closely resembles the Taguchi loss function (Fig. 3.27), which describes the increased cost of quality loss associated with imperfection and error, especially with regards to design specification and tolerances [161]. Although the 1S class function is a good initial first guess at the designer’s preferences for all tolerances, there may be errors that the designer would not consider relaxing. The designer has the option to make these tolerances mandatory, in which case they are represented by a 1H preference class function. We have also chosen to represent the goal of minimizing cost, as represented by minimizing λ , using a 1S class function.




Figure 3.27: Taguchi quadratic quality loss function, after [162], where Δ represents the tolerance.

Rather than asking designers to set their initial preference boundaries for each tolerance objective, we instead automatically position the boundaries of the class function relative to the initially specified tolerance, Δ . These class boundaries are summarized in Tab. 3.5 in the row for optional tolerances, ϵ_{op} . These initial assumptions will be evaluated in future work to determine if they match the initial preferences of designers.

After being presented with the assumed boundary positions, the designer is given the opportunity to adjust the boundaries, if desired. For designers with vague ideas regarding which tolerances are truly critical, no changes to the boundaries are needed. However, the designer has the option to change the position of the bounds if the initial guess does not match her preferences. The designer can choose what value of error defines the transition between the desirable region, tolerable region, undesirable region, and so on. The designer also has the option to change each tolerance to mandatory, ϵ_{mnd} , as described above. This phase of refining the preference structure further justifies the use of physical programming: the designer can more easily answer what range of error is tolerable to them than determining what scalar weight to assign to each of many tolerances.

Once class function types are determined and boundaries are calculated, weights can be calculated for each soft function, defining the shape of the class function for each objective, following [160]. Using these weights, \tilde{w} , and by calculating the deviational variables, denoted d_{is} , the overall minimization problem can be formulated. (Only a subset of the full physical programming implementation used by Messac [160] is described here since there

Table 3.5: Preference class function parameters where t_{is} separate the different preference regions.

	Class	HD	D	T	U	HU
	Boundary	t_{i1}	t_{i2}	t_{i3}	t_{i4}	t_{i5}
ϵ_{op}	1S	0.3Δ	0.8Δ	Δ	1.3Δ	2Δ
λ	1S	5%	30%	75%	90%	100%
		Acceptable				
ϵ_{mnd}	1H	0	Δ			

are no maximization, value, or range goals needed for our application.) Our main physical programming objective for class 1S and 1H functions is summarized as:

$$\min_{d_{is}} J = \sum_{i=1}^{n_{sc}} \left[\sum_{s=2}^5 \tilde{w}_{is} d_{is} \right] \quad (3.17)$$

subject to

$$\begin{aligned} \mu_i - d_{is} &\leq t_{i(s-1)}; \quad d_{is} \geq 0; \quad \mu_i \leq t_{i5} \quad (1S) \\ &\text{or} \\ \mu_i &\leq t_{j,max} \quad (1H) \end{aligned} \quad (3.18)$$

where n_{sc} is the number of soft class objectives, n_{hc} is the number of hard class objectives, $i = 1, 2, \dots, n_{sc}$, $j = 1, 2, \dots, n_{hc}$, $s = 2, \dots, 5$, x is the design variable vector, and $\mu_i = \mu_i(x)$ are the design objectives.

In a typical physical programming problem, Eq. 3.17 would be optimized using a commercial optimization code, resulting in a single, Pareto-optimal set of objective values. (Pareto-optimal refers to a state where it is impossible to improve one objective value without making at least one other objective value worse off). For our problem, because the error calculations in Eq. 3.12 and Eq. 3.13 have a quick computation time, it is trivial to generate thousands of solutions to use for data visualization. Designers can learn more about the trade-offs between objectives by examining the range of possible orientations and the errors associated with them, exploring how each objective is related, and determining the range of errors associated with each objective. Our data visualizations (described in the next section) allow designers to learn from a wide range of solutions without needing to examine each orientation individually. This visualization and exploration can help designers to refine preferences to better identify a most preferred orientation. Interactive physical programming has been explored before [156], but the process can be somewhat tedious since only one Pareto-optimal point is generated at a time. Because of the quick computation time of our problem, hundreds of Pareto-optimal points can be generated.

3.6 Implementation of tolerance allocation tool - TAAM

TAAM requires as input: geometry information for the feature face and datum face, normal vectors of the face and datum, the type of tolerance on each face, and the associated tolerance value. This geometry information can be manually found from the STL file or exported from a CAD system from the CAD model itself. The calculations performed for the example in this paper use geometry information directly from the CAD model and do not account for errors introduced by coarse STL resolution.

The interaction with the designer begins with entering geometry and tolerance information. Then, orientations are sampled (randomly distributed points on a unit sphere that are chosen using the method of Marsaglia [163] as well as orientations parallel to each tolerated face’s normal vector) and the error associated with each specified tolerance is calculated. We use all generated orientations for plotting, including non-Pareto-optimal points because it results in a more cohesive visualization, but the orientations are then filtered to remove dominated, non-Pareto-optimal points. Only the non-dominated points are ranked according to the physical programming weights, and the single most-preferred orientation is found and presented to the designer. This process is summarized in Fig. 3.28.

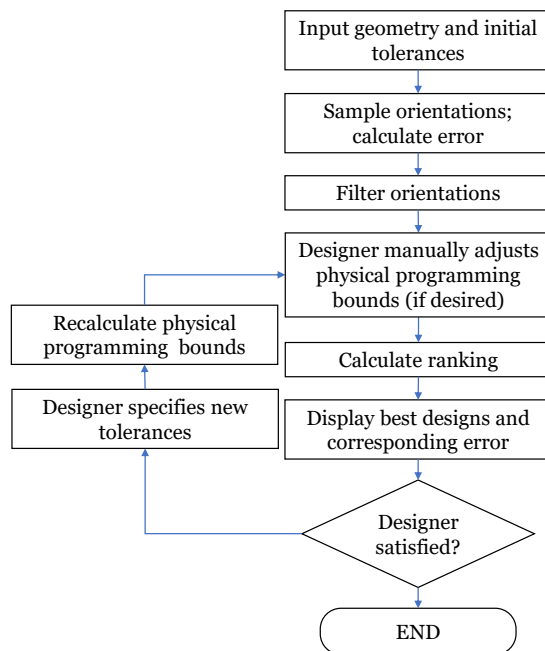


Figure 3.28: Process flow for interaction with TAAM.

If the designer is satisfied with the first orientation selected for them and does not need to refine or relax tolerances, the interaction ends. If the first orientation is not satisfactory, the designer can use a GUI we developed to relax the specified tolerances selectively. The

designer is presented with several data visualizations, illustrated in previous sections, to help them analyze the data. Once the designer is satisfied with the relaxed tolerances as input into the GUI, the physical programming bounds and weights are recalculated, and the process can restart (see the loop in Fig. 3.28). Because all the preference ranges are calculated from the specified tolerances, the new tolerance will change all the preference ranges for that updated tolerance. Additionally, the designer can opt to change some of the tolerances to mandatory (class function 1H) or to change the 1S boundaries.

This tolerance-exploration and -relaxation process minimizes initial input required from designers by making informed guesses about preferences regarding error and cost. After quickly generating candidate orientations and presenting these orientations, the system helps designers to explore what-if scenarios by guiding them through selectively relaxing or refining tolerances. Once designers have explored the data enough that they are satisfied that they cannot achieve a better result, the physical programming ranking identifies the orientation that best meets their needs. This exploration process will be illustrated in the next chapter with an example design scenario. Designers can then indicate that the part must be built in this orientation, as detailed in ASME Y14.46 [109]. The source code for this tool will be made available at the author's GitHub site.

3.7 Conclusion

TAAM is able to calculate orientation, flatness, and cylindricity errors for multiple features quickly and can identify build directions that result in the satisfaction of all tolerances. Mathematical analysis of the tolerance error enables efficient prediction of error due to the stair-stepped approximation characteristic of AM. The system is material-general and can be utilized for any layer-based process with fixed layer thickness. A tool that predicts tolerances are mutually achievable for a part can improve the tolerance allocation process. Further, this tool can elucidate relationships between layer thickness, orientation, and geometric errors, helping designers understand and optimize tradeoffs between quality and cost before the detail design of a part is finalized.

Similarly, Will It Print assess part geometry to determine compliance with a set of DFAM guidelines. The guidelines that it evaluates were chosen through an evaluation of commonly cited guidelines from academic research and 3D printer manufacturers, as well as an analysis of common student printing problems. By evaluating compliance with the guidelines using algorithms, the tool avoids relying on possibly faulty designer analyses. This approach also frees designers from route tasks so they can focus on more creative or analytical aspects of the design. Will It Print can help reduce cycle time and failed prints during the design of a new part.

Both of the tools developed in this chapter use geometry analysis to support designers' analysis of parts by providing them with information about manufacturing outcomes. These types of design tools can help bridge the gap between designers and manufacturers, allowing designers to understand and optimize manufacturing constraints before their parts are

manufactured. In the next chapter, both tools will be illustrated in more detail when they are applied to a design scenario.

Chapter 4

Example application of geometry analysis tools

The previous chapter detailed the development and technical details of the design for additive manufacturing (DFAM) systems we have developed. The purpose of this chapter is to demonstrate, through an example design scenario, the practicality of the systems. We seek to assess if the tools introduced earlier can be set up to be efficient, both in calculation time and in the effort level required by the designer. By detailing the workflow and computation time of our systems, Tolerance Allocation for Additive Manufacturing (TAAM) and Will It Print, we hope to demonstrate the viability of using geometry analysis to support DFAM. We also will use the example design scenario to evaluate the performance of our tools compared with existing DFAM systems.

4.1 Introduction

Although both TAAM and Will It Print generally address the manufacturability of AM parts, they were designed for different phases of the design process. Will It Print helps evaluate major manufacturability problems relatively early in the design process, so that a designer can mitigate the problems with geometry changes. TAAM helps with assessing whether a certain level of geometric accuracy, specified by a set of geometric tolerances, can be achieved by a given AM process. The process of assessing tolerances, which is when a designer would use TAAM, occurs later in the design process, once the preliminary detail design has been completed.

As summarized in earlier chapters, Will It Print was created to help designers by automatically assessing the manufacturability of a part. For AM, manufacturability information is typically given as a list of guidelines, but a few tools that predict manufacturability do exist, both in academia [28,30,126,127] and as commercially available software packages such as Meshmixer. We will use the example design scenario to demonstrate a typical workflow of using Will It Print. Comparisons to other existing tools will also be drawn. Will It Print has additional manufacturability assessments that are not available in existing tools and also

includes richer feedback about each category of manufacturability, so we expect it to better support designers, especially if they are novices.

The goal of TAAM is to allow a designer to predict geometric error and use that information to choose a build orientation and to change their tolerance allocation if necessary. Researchers have previously developed systems that perform a similar function, typically in the form of an orientation optimization algorithm. As summarized in Chapter 2, there have been many orientation optimization algorithms that have been proposed for AM, with a wide variety of optimization objectives. Several algorithms that have been described in literature focus on minimizing cost, printing time, or support volume [129, 130]. Others focus on minimizing geometric error on the part, using metrics like average volumetric error [134, 135], surface roughness [131–133], or error associated with geometric tolerances [90, 93, 136–138]. Some commercial tools also exist: Meshmixer performs a weighted sum optimization of the objective functions related to support volume, strength, and support area. For this example design scenario, we will focus on two orientation optimization algorithms that have the most in common with TAAM, in that they simultaneously try to optimize for metrics representing part cost and geometric error.

We will use our example design scenario to evaluate similarities and differences between the workflows of TAAM and other approaches (evaluating the amount of interactivity and utility of different visualizations) and to compare the recommended final build orientation of the part. The level of interactivity, which we posit is important for designer learning, varies widely, but most systems are not interactive. In addition, most systems simplify the geometric error by creating an average metric, while TAAM calculates error corresponding to geometric tolerances applied on individual features, which is the system designers typically use. We hope our example design scenario will highlight some of the improved practicality of the approach we used with TAAM.

By using an example design scenario, we can demonstrate the ease-of-use of both of our systems, their visualizations, and compare these attributes to other similar systems. The example design scenario will be summarized in the next section. Then, the results and discussion of the design scenario will be presented. We will present results for Will It Print first, as it would be used earlier in the design process.

4.2 Materials and methods

Because these tools are interactive, we want to show a typical workflow, and so we will describe actions taken and choices made by a hypothetical “designer.” No user testing was undertaken in this chapter. User testing will be described in Chapter 6. The workflows and output for the two tools are different by design because they have different target audiences. We attempt to unify the two tools by presenting how they might be used at separate times during a new product development cycle (Fig. 4.1).

For our example design scenario, we will describe a fictional new product development for a simple product. The example part used in the design scenario is a triple flagpole bracket

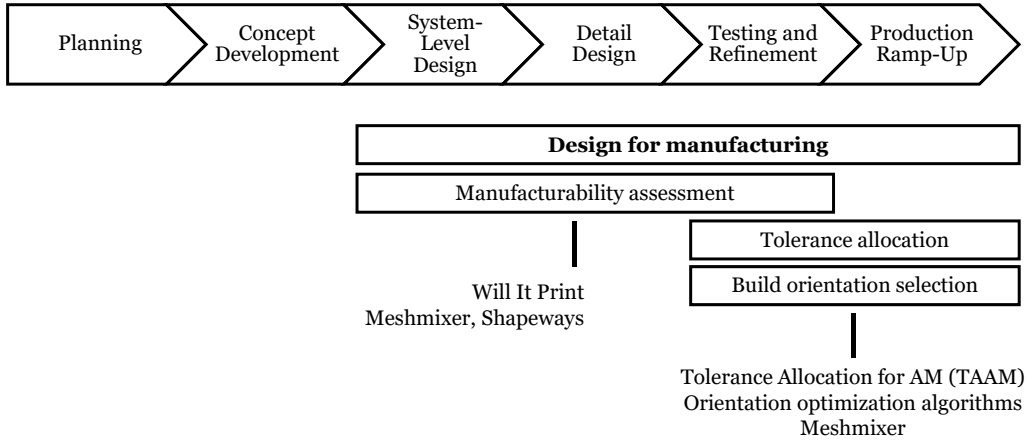


Figure 4.1: Our tools fit into different stages of the product development process (process flow adapted from [164]).

used to mount three flagpoles on the corner of a building, which was chosen because it has both planar and cylindrical features, not all of which are orthogonal to each other. (Selecting a build orientation for parts with features that are all orthogonal to one another tends to be more straightforward, and is less informative for demonstrating the tools.)

Different designs for the flagpole bracket were created, varying the type, complexity, and size of the part features. We created several designs in order to replicate the process of evaluating different design concepts, early in the detail design stage. A designer could use Will It Print to assess the manufacturability of all designs. The tool could be used to identify geometry changes that could be made to improve the manufacturability of the best design. The geometries of the different designs we generated for this design scenario are shown in Fig. 4.2. A designer could also use Will It Print to explore the impact of orientation on manufacturability and consider design changes to make their design easy to print in many different orientations, if desired.




Figure 4.2: Geometry of the different design alternatives evaluated using Will It Print.

After geometry improvements are made, later in the development process, the designer would have a semi-finalized design and would develop a preliminary tolerance allocation scheme. In our example design scenario, the theoretical designer chooses Design C for further development. Before approving their design for production, the designer may want to ensure that their tolerance allocation scheme is achievable, and may want to select an orientation to ensure optimal cost and quality of the manufactured product. As discussed in Chapter 3, TAAM uses geometric tolerances to quantify allowable geometric error. To replicate a possible tolerance scheme that a designer might specify, a flatness tolerance was assigned to one face of the bracket that mounts on a building. A perpendicularity tolerance and three cylindricity tolerances were also defined. Figure 4.3 shows the geometry and geometric tolerances used as input in TAAM.




Figure 4.3: Example flagpole bracket part with specified geometric tolerances.

This materials and methods section first describes the implementation of Will It Print and describes the process for benchmarking Will It Print against other tools. Then, the implementation of TAAM and the process for benchmarking TAAM against similar tools is described. One aspect of benchmarking was comparing computation times. We measured computation time for all methods tested in MATLAB using `tic` and `toc`. For commercial tools not implemented in MATLAB, we similarly estimated computation time as the amount of time elapsed between uploading the part and seeing the results of the geometry analysis.

4.2.1 Benchmarking performance of Will It Print

There are few software tools that perform extensive manufacturability analysis for AM. Some academic tools have been developed [28,30,126,127], but they are not typically shared or made available online. A few commercial tools are available as well. The academic and commercial tools have different capabilities to assess different categories of manufacturability. The capabilities are summarized in Tab. 4.1. Two commercial tools, Meshmixer (Autodesk, San Rafael, CA) and Shapeways 3D Tools (Shapeways, New York City, NY), which is a service provided as part of Shapeways' 3D part ordering process, were used to compare the predictions of Will It Print.

Shapeways 3D Tools has only a few manufacturability analyses. It checks to make sure the part is smaller than the maximum size allowed for a given machine. It also checks to

Table 4.1: Comparing capabilities of manufacturability analysis tools (1. Small features; 2. Warping; 3. Tipping; 4. Surface roughness; 5. Overhanging faces).

Tool	Capability					Other capabilities
	1	2	3	4	5	
Will It Print	X	X	X	X	X	
Shapeways 3D Tools	X					Maximum part size
Meshmixer	X		X		X	Strength
Tedia & Williams [28]	X				X	Build time estimate
Ranjan, Samant, & Anand [30]	X			X	X	Overall manufacturability metric
Ghiasian et al. [126]	X				X	Build time estimate

see if the feature sizes on a part are larger than the minimum wall thickness. The minimum wall thickness used in Shapeways is 0.7 mm (a value that reflects settings for SLS, while Meshmixer and Will It Print are set up to reflect settings for FDM). In Shapeways 3D Tools, the minimum wall thickness threshold is not adjustable, so we purposely designed the flagpole designs with some geometry smaller than 0.7 mm to compare the predictions of the tools. We could not find an explanation of how Shapeways 3D Tools computes feature size.

Meshmixer has more extensive manufacturability analyses than Shapeways 3D Tools. Meshmixer performs a “stability” analysis to assess if the part was in danger of tipping, similar to Will It Print’s tipping analysis. For its stability analysis, Meshmixer calculates if the object’s center of mass falls within the convex hull of contact points on the build plate, to assess if the object will be stable under its own weight. Before checking the stability of a part in Meshmixer, we used Meshmixer’s built-in support material generator to create support structures for overhanging faces. Evaluating tipping considering support material is the approach Will It Print takes, and so we wanted our comparison to Meshmixer to be consistent. For “thickness” analysis, which is the name Meshmixer uses for its evaluation of small features, the minimum thickness was set to 2 mm. Meshmixer’s thickness analysis is performed, according to the Meshmixer software, “by shooting a number of random rays from each vertex into the interior of the surface, and measuring the distance to the closest ray intersections. The average of these distances is our measure of thickness.”

To evaluate the manufacturability tools, we tested models of varying complexity, and with different features (Fig. 4.2), which we will describe briefly here. Each model also had different STL file resolution, in order to assess the impact of STL resolution on the accuracy of the manufacturability tools’ analyses. For Design A, the thinnest features were the cylindrical walls, which were 2 mm thick. The resolution of the STL file for Design A was made very fine, with 7120 facets. The bounding box size of Design A was 148 mm \times 148 mm \times 108 mm. Design B had several overhanging features. The bounding box size of Design B was 126 mm \times 126 mm \times 96 mm. The thinnest features were the cylindrical walls, which were 2 mm thick. The resolution of the Design B STL file was in between Design A and Design B, with 2882 facets. Design C was the simplest design. Design C featured cylindrical features

that had a wall thickness of 0.65 mm and small countersunk holes with diameters equal to 1.4 mm. The STL resolution of Design C was very low, with 1066 facets. The bounding box size of Design C was 155 mm \times 155 mm \times 110 mm.

Will It Print and Meshmixer are meant to be used interactively, with the user choosing orientations based on intuition and tool feedback. However, to systematically investigate differences between the tools, we also used Meshmixer and Will It Print to evaluate five specific orientations. Two orthogonal orientations were chosen ($\mathbf{B} = [0, 0, 1]$, and $\mathbf{B} = [0, 1, 0]$), as well as two orientations that were randomly generated ($\mathbf{B} = [0.28, -0.85, -0.44]$ and $\mathbf{B} = [-0.79, 0.167, 0.59]$). Because the only capability of Shapeways 3D Tools that we assessed, small features, does not depend on orientation, Shapeways 3D Tools was not used for this part of the evaluation.

Similarities and differences between tools will be discussed in detail in the Results and Discussion section. The next section describes the implementation of TAAM and the process for benchmarking it against other tools.

Will It Print implementation

Chapter 3 contains detail of the algorithms and MATLAB implementation used to assess manufacturability, but we will summarize the specific settings used in the tool here. Assuming the part was going to be prototyped using fused deposition modeling (FDM), the minimum feature size was set at 2 mm, the angle at which overhanging feature warnings were enabled was set at 135°, and the length associated with warping was set to 80 mm. The support density, η_s , was set to 15%. These values were determined from best practices cited in industry guidelines and from experimentation using the Type A Series 1 Pro printer.

4.2.2 Benchmarking performance of TAAM

In order to evaluate the performance of TAAM, we compare its outputs to the recommendations of two orientation optimization schemes that have been proposed in literature. Specifically, we will compare our approach with orientation optimization schemes that optimize for surface roughness, based on the work of Thrimurthulu et al. [133], and simplified geometric tolerances, based on the work of Das et al. [138]. A summary of the implementation of these two approaches is included below, but for more technical detail, the reader is referred to the original papers that described these approaches.

Average surface roughness optimization

The first approach used for comparison, based on the optimization scheme proposed by Thrimurthulu, Pandey, and Reddy [133] is selecting an orientation to optimize the average surface roughness on the part. This approach will be referred to as “average roughness optimization,” or ARO, in subsequent sections.

The surface roughness is evaluated at each facet of the STL using an experimentally-derived function relating face orientation to surface roughness, $Ra(\theta, \phi)$. An average surface roughness for the entire part is calculated from the surface roughness (Ra_i) and area (A_i) of the i^{th} STL facet, for the total number of facets, n_f :

$$Ra_{av}(\theta, \phi) = \frac{\sum_{i=1}^{n_f} Ra_i A_i}{\sum_{i=1}^{n_f} A_i}. \quad (4.1)$$

The minimization problem from [133] is summarized as:

$$\begin{aligned} \min_{\theta, \phi} O = & \quad \omega Ra_{av} + (1 - \omega) \bar{t} \\ \text{s.t.} \quad & \quad 0 \leq \theta \leq \pi \\ & \quad 0 \leq \phi < 2\pi. \end{aligned} \quad (4.2)$$

We calculated printing time, $\bar{t}(\theta, \phi)$, according to Eq. 3.15. The designer sets the scalar weights applied to the two objectives by setting ω to some value between 0 and 1. The designer is generally not given any guidance or advice about setting these parameters. In the original paper describing this approach [133], it is assumed that the designer knows the appropriate scalar weights to apply. A genetic algorithm [165] is used to solve the optimization problem described in Eq. 4.2, using 50 generations.

Weighted penalties of geometric tolerances

As another comparison to the orientation selection employed by TAAM, we implemented the approach of Das et al. [138], which is based on that of Paul and Anand [166]. The approach described in these two papers uses a weighted penalty function to minimize support volume and error associated with different tolerances simultaneously. This approach will be referred to as “penalized tolerance optimization” or PTO in subsequent sections.

These papers calculate geometric error associated with different tolerances using the following equations, where h is the layer thickness, θ_{ax} is the angle between the cylinder axis and the build vector, and θ is the angle between a face’s normal vector and the build direction:

$$\begin{aligned} \epsilon_{cyl} &= h \sin \theta_{ax} \\ \epsilon_{flat}, \epsilon_{perp}, \epsilon_{par} &= h |\cos \theta|. \end{aligned} \quad (4.3)$$

Penalties, $p(\theta)$, for these errors are defined according to the amount of error associated with each tolerance. In addition to minimizing geometric tolerances, another objective is minimizing support volume. The maximum support volume for all orientations is found, and a normalized support volume, $V_{S-norm}(\theta, \phi)$, for each orientation relative to that maximum is calculated, using the same method described in Chapter 3 for λ . Briefly, Das et al. define

the minimization problem as:

$$\begin{aligned}
 \min_{\theta, \phi} E &= \sum_{i=1}^{n_{cyl}} p_i \omega_i + \sum_{j=1}^{n_{flat, perp, par}} p_j \omega_j + \omega_s V_{S-norm} \\
 \text{s.t.} \quad &0 \leq \theta \leq \pi \\
 &0 \leq \phi < 2\pi
 \end{aligned} \tag{4.4}$$

where $n_{flat, perp, par}$ and n_{cyl} are the number of flatness, perpendicularity, or parallelism callouts, and the number of cylindricity callouts respectively, and the weights (ω_i, ω_j , and ω_s) are user-assigned scalar weights for the planar callouts, cylindrical callouts, and support material. Following their approach [138], the minimization problem is solved using the `fmincon` routine available in MATLAB.

TAAM implementation

Our TAAM uses physical programming to set boundaries of desirability for each objective, rather than using scalar weights or weighted penalties. For the physical programming calculations, the preference region boundaries, t_{i1} to t_{i5} , were set based on guidelines described in Chapter 3, with certain amounts of error being deemed highly desirable (HD), desirable (D), tolerable (T), undesirable (UD), and highly undesirable (HU) (Tab. 4.2). These boundaries are also applied to the part cost parameter, λ . The objective function is given in Eq. 3.17. The class for each objective was set to 1S, indicating a soft class function in which the goal is to minimize the value of the objective (for more information, see Chapter 3).

Table 4.2: Example part physical programming parameters. For full explanation of notation used in this table, see Chapter 3.

		HD	D	T	U	HU
	Class	t_{i1}	t_{i2}	t_{i3}	t_{i4}	t_{i5}
ϵ_{flat}	1S	0.045	0.135	0.165	0.195	0.3
ϵ_{perp}	1S	0.075	0.225	0.275	0.325	0.5
ϵ_{cyl}	1S	0.06	0.18	0.22	0.26	0.4
λ	1S	5%	50%	80%	90%	100%

Normal vectors of the toleranced features and STL information were input into the system. The layer-level deviations, δ_{xy} , δ_z , and δ_s were set to 0.10, 0.05, and 0.05 mm, respectively while the layer thickness, h , was set to 0.20 mm. The support angle, θ_{sup} was set to 135°. These values were chosen to reflect layer-level error and layer thickness for an FDM process.

4.3 Results and Discussion - Will It Print

The results and discussion in this chapter are divided into two sections. The first section describes the output of the Will It Print tool, comparing it to other available design tools. The second section describes the output of the TAAM tool and compares it to orientation optimization approaches.

4.3.1 Results - Will It Print output

To demonstrate the general workflow of using the Will It Print tool, we will illustrate the full workflow for one part design in particular: Design C. The workflow for assessing any of the designs is similar, so we will demonstrate the process for one design to avoid being repetitive.

The manufacturability analysis of Will It Print is divided into several categories, each focusing on a particular design for additive manufacturing guideline. We will present the results of the analysis for each of those guideline categories: small features, warping, tipping, surface roughness, and overhanging faces. The designer views the results for each guideline one at a time.

Initial analysis

The first guideline listed in the GUI evaluates small features. (A screenshot of the GUI is presented in the previous chapter in Fig. 3.15). Once the user clicks the “small features” option, Will It Print evaluates the part geometry. Will It Print accurately indicates that the cylindrical features were thinner than the recommended 2 mm, and also highlights the small holes (Fig. 4.4). The designer could look at the information window for further details regarding minimum size recommendations (Fig. 4.5). The info guide suggests increasing the size of all features to be greater than 2 mm in all directions.

Moving on, the next guideline listed in the GUI is overhanging features. Here, the designer can see that the majority of the part is not overhanging in this orientation, indicating that not much support material is needed. However, a few facets inside the cylindrical features are highlighted in yellow, indicating they are overhanging (Fig. 4.6i). It would be difficult to remove support material from inside these internal features fully, and so the designer should identify an alternative orientation.

The next guideline category is tipping, which is also referred to as wobbling. Because the base of the part is relatively large, and the center of mass is not very tall, tipping is not likely at this orientation (Fig. 4.6ii).

The next guideline category is warping. Warping was predicted (Fig. 4.6iii), indicating that the part may curl or bow slightly. The part is relatively large, and so the maximum dimension of the base area’s bounding box is above the maximum recommended size.

The final guideline is surface roughness. There are some facets on the part with a normal such that high surface roughness is predicted at the top of each cylinder (Fig. 4.6iv). Most

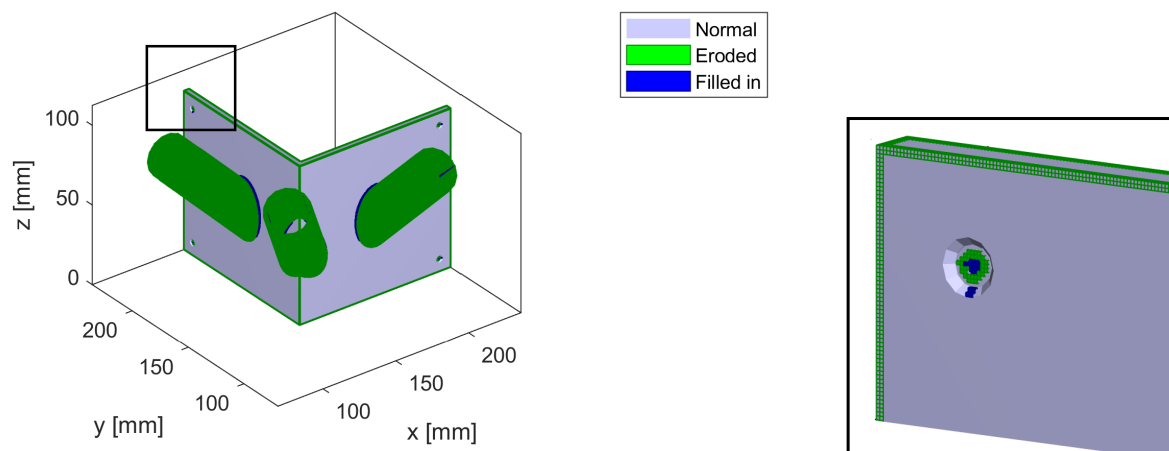


Figure 4.4: Will It Print analysis indicates that the thickness of the cylindrical features and the size of the small holes are smaller than the recommended size (close-up of a hole is shown in the inset). As described in Chapter 3, “eroded” indicates positive features that are too small and “filled in” indicates negative features that are too small.

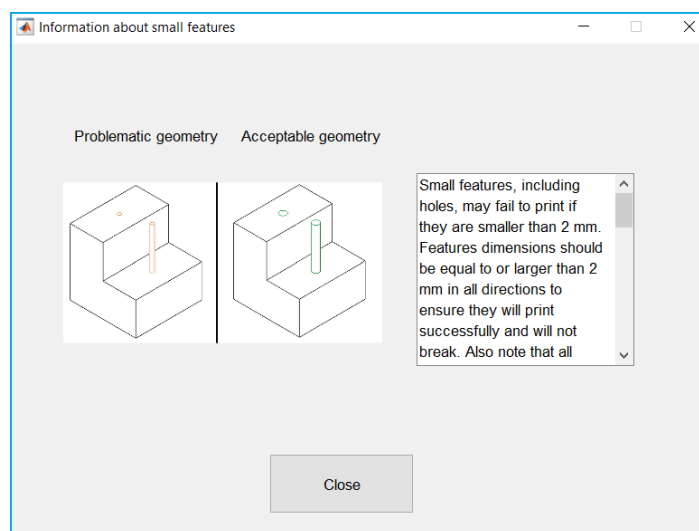


Figure 4.5: Info guide for small feature analysis with an explanation and suggestions for improvement.

of the facets on the non-cylindrical surfaces are oriented 90° from the build direction, and so the surface roughness is low on those faces.

Design changes: Small features

To address the warnings from Will It Print, the designer could decide to make some geometry changes. Based on the feedback described above for small features (Fig. 4.4), the designer

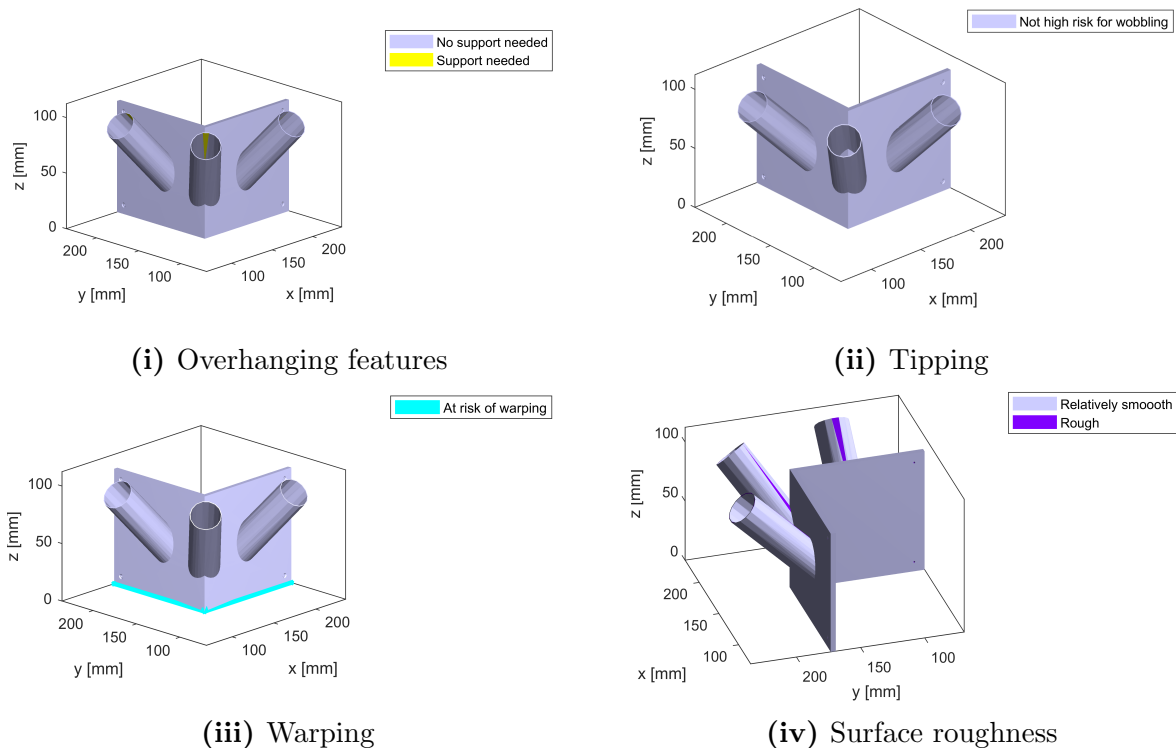


Figure 4.6: Will It Print output for Design C, showing results for surface roughness, overhanging features, and tipping for $\mathbf{B} = [0, 0, 1]$.

may decide that the hole features can be added in a post-processing operation, and so the hole features can be removed from the geometry sent to the 3D printer. To address the thin cylindrical walls, which are currently 0.65 mm thick, the designer could thicken the walls to more than 2 mm. The designer removes the holes and thickens the walls to 3 mm in an external CAD program, creates an STL for the new geometry, and uploads the revised STL into the tool. When re-evaluating the part for small features, the designer can see that the highlighting has been largely eliminated (Fig. 4.7), showing that the geometry changes adequately addressed the problem.

Design changes: Exploring orientations

After making the geometry changes to address the too-small features, the designer may choose to try to eliminate the other warnings regarding surface roughness, overhanging features, and warping by changing the orientation of the part. The designer might choose an unusual orientation, such as $\mathbf{B} = [0.5, 0.71, 0.71]$, to try to minimize the area on the build plate to eliminate warping problems. The error displays are updated in near real-time (Fig. 4.8). Although the surface roughness is now low on key features, a tipping warning is issued.




Figure 4.7: Will It Print indicates cylindrical features are acceptable when thickened.

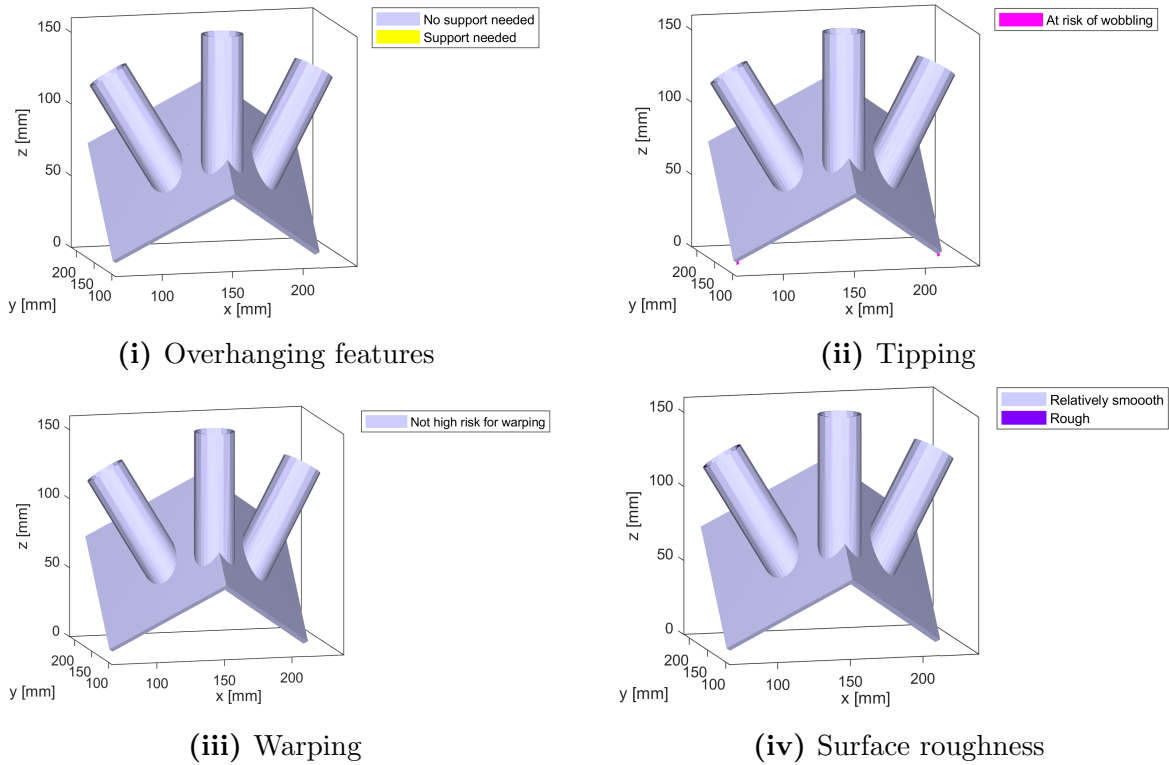


Figure 4.8: Will It Print output for the modified version of Design C, showing that no geometry in the current view is highlighted for surface roughness, overhanging features, and warping, but a warning is issued for tipping when $\mathbf{B} = [0.5, 0.71, 0.71]$.

To fix the problem of not enough surface area on the build plate that is causing the tipping issue, the designer could rotate the part (new build vector $\mathbf{B} = [-0.06, -0.06, -0.99]$) so that the top face of the part would require support, which also moves the center of mass down.

At this new orientation, no tipping warning is issued (Fig. 4.9). Additionally, no warping warning is issued. Features with overhanging and surface roughness warnings are not critical to the function of the part. The designer would likely be satisfied with the manufacturability of this design, after implementing these geometry and orientation changes, and would have a sense of promising build directions to explore as the design process moves towards production.




Figure 4.9: Will It Print analysis indicates that the modified version of Design C is no longer at risk of tipping when $\mathbf{B} = [-0.06, -0.06, -0.99]$.

At this point, the designer could move on to evaluating Design A and B, to see if they required any geometry or orientation changes as well. We will not show the evaluation of Design A and B for the sake of brevity, but all three designs are used to compare and contrast Will It Print to other existing tools.

4.3.2 Comparison of Will It Print to other tools

For comparison of the predictions of the Will It Print Tool to other design for AM tools, the output of Meshmixer and Shapeways 3D Tools is analyzed. All three designs and multiple orientations were tested with each tool.

Small features

Small features do not depend on orientation, so the tools were compared with all designs at the orientation $\mathbf{B} = [0, 0, 1]$. We compared the original version of Design C, not the version that was created with geometry and orientation changes described in the previous section.

In Will It Print, all of the cylindrical features were highlighted for by the small feature analysis Designs A, B, and C (Fig. 4.10). In all designs, sharp corners are also highlighted. As mentioned earlier, the small holes of Design C are also flagged.

Meshmixer highlights all of the cylindrical faces and none of the surrounding walls structure for Design A (Fig. 4.11i), which had a fine STL resolution. Meshmixer uses red highlighting to indicate features that are too small. For Design B, which has a lower STL resolution

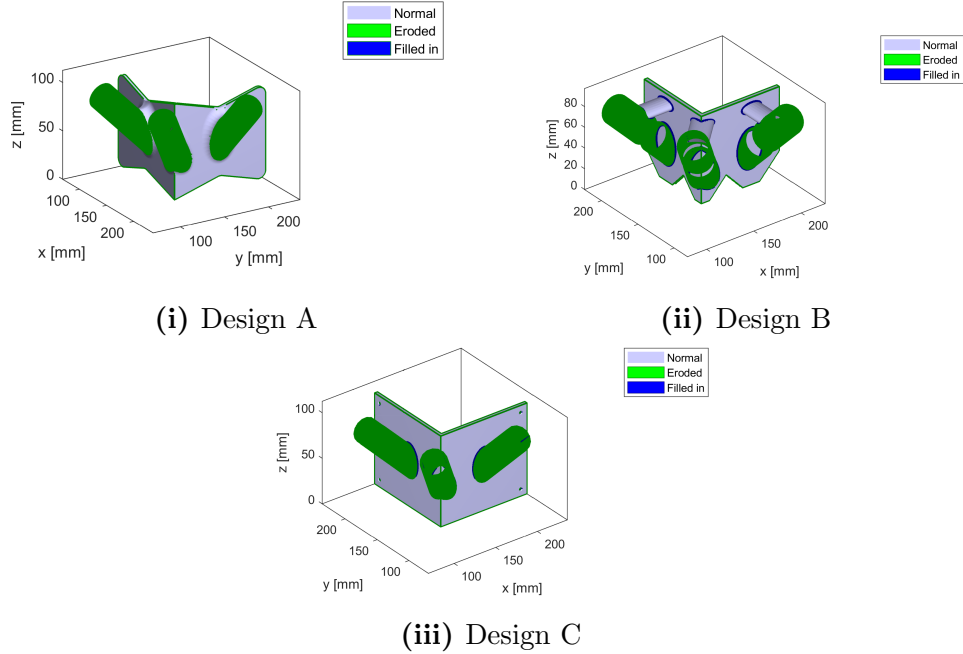


Figure 4.10: Will It Print identification of too-small features for (i) Design A, (ii) Design B, and (iii) Design C, showing that the cylindrical features and sharp corners are below the recommended thickness.

and more complex geometry, Meshmixer highlights both the cylindrical features and portions of the wall behind them (Fig. 4.11ii). For Design C, most of the part is highlighted but the small holes are not highlighted.

Shapeways 3D Tools did not highlight any portion of the geometry as being below the recommended size for Design A, (Fig. 4.12i). Similarly, for Design B, Shapeways 3D Tools did not highlight any portion of the part as being below the recommended size, (Fig. 4.12ii), but the whole part was highlighted in green. It is unclear why Shapeways 3D Tools employed different visualizations for Design A and C. It is possible that this difference is related to a bug in their system. For Design C, Shapeways 3D Tools identified portions of the very thin cylindrical walls as being below the recommended thickness (Fig. 4.12). The walls of the CAD file used to create the STL were uniformly 0.65 mm thick, thinner than the recommended thickness from Shapeways 3D Tools, 0.7 mm, but only isolated portions of the surfaces were highlighted in Shapeways 3D Tools. Because Shapeways 3D Tools is not open-source, we are unsure what method they use to calculate thickness and why it would give these erroneous results. The small holes in Design C are not highlighted, indicating that Shapeways 3D Tools does not assess small negative features.

In general, the small feature detection results were similar for Meshmixer and Will It Print (Tab. 4.3). All three tools highlighted at least portions of the cylinders that were only 0.65 mm in thickness. However, Meshmixer also highlighted surrounding facets near




Figure 4.11: Meshmixer identification of too-small features for (i) Design A, (ii) Design B, and (iii) Design C, showing that the cylindrical features and some surrounding areas are highlighted in red, indicating they are too small.

the cylindrical faces. Will It Print also highlighted sharp, 90° corners, which neither of the other tools highlighted. The analysis of Shapeways did not have good agreement with the other two tools. All tools have some flagging that could be improved, but the flagging errors have varying severity. Both Will It Print and Meshmixer could highlight too much: Will It Print highlights sharp corners that designers may not care about and Meshmixer highlights portions of planar faces that are actually adequately large. Shapeways 3D Tools, on the other hand, fails to highlight portions of the cylinder that are below the threshold, an issue that is more severe.

Computation times

A summary of the different computation times of the comparison tools and Will It Print is shown in Tab. 4.4. The computation time for Meshmixer seems to scale with the number of facets in the model, with finer STL resolution resulting in longer computation time. However, the computation time of Meshmixer was extremely quick compared with the other tools. The computation time of Will It Print is impacted by the number of voxels and, therefore, the size of the part. Shapeways 3D tools took several minutes to compute. The computation times for Will It Print were in between those of the other tools.

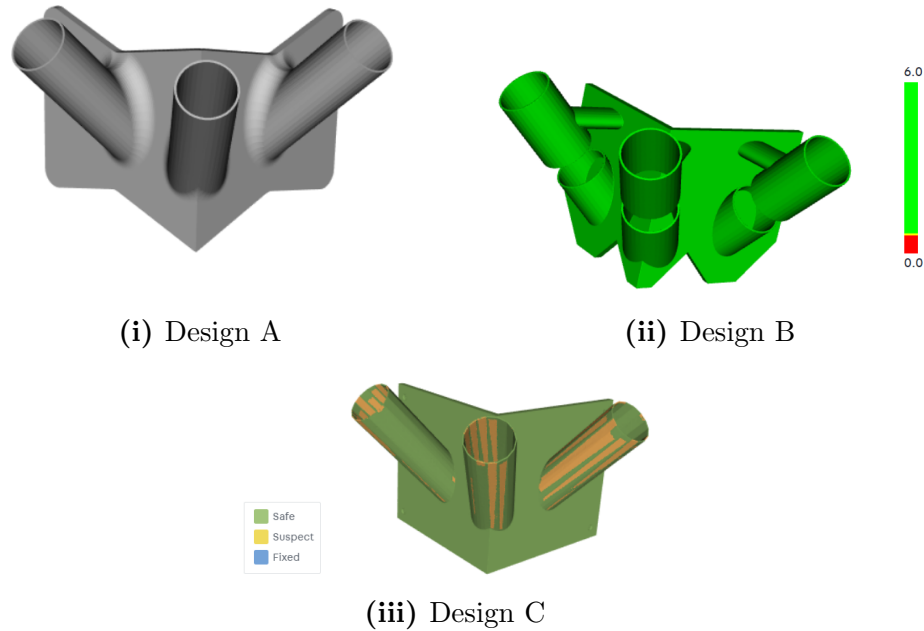


Figure 4.12: Shapeways 3D Tools identification of too-small features for (i) Design A, (ii) Design B, and (iii) Design C, showing that only parts of the cylindrical features in Design C are flagged as being too small.

Table 4.3: Comparison of features flagged with small feature error of the tools

	Model		
	Design A	Design B	Design C
Will It Print	Sharp corners; cylindrical faces; holes	Sharp corners; cylindrical faces	Sharp corners; cylindrical faces
Shapeways 3D Tools	None	None	Portions of cylindrical faces
Meshmixer	Cylindrical faces	Cylindrical faces; portions of planar faces	Cylindrical faces; portions of planar faces

Exploring orientations

Many of the manufacturability problems (e.g., overhanging features, surface roughness, tipping, and warping) depend on the build orientation of the part. Because the tools have different functionality (i.e., both Shapeways or Meshmixer have fewer analysis capabilities than Will It Print), it is difficult to compare all orientation-dependent manufacturability

Table 4.4: Comparing computation times for small feature detection using the different test parts.

	Computation time for part [sec]		
	Design A	Design B	Design C
# of facets	7120	2882	1066
Will It Print	95.8	162.5	121.6
Shapeways 3D Tools	≈ 600	≈ 600	≈ 600
Meshmixer	2.3	<1	<1

analyses fully. However, both Will It Print and Meshmixer compute warnings for tipping and overhanging features, and so we will compare those two errors for different orientations.




Figure 4.13: Meshmixer shows the part is not at risk of tipping (communicated by the green color of the ball marking the center of mass) when $\mathbf{B} = [0, 1, 0]$.

All designs were evaluated for warping and tipping at $\mathbf{B} = [0, 0, 1]$, $[0, 1, 0]$, $[0.28, -0.85, -0.44]$, and $[-0.79, 0.167, 0.59]$, as described in the methods section of this chapter. An example of the results of Meshmixer is shown in Fig. 4.13. In Tab. 4.5, we summarize overhanging warnings outside or inside the cylindrical features, as well as tipping warnings for the whole part. Meshmixer identified areas of the part as being *nearly* overhanging but did not flag them as overhanging in Design A and Design C at orientation $\mathbf{B} = [0, 0, 1]$, while Will It Print flagged facets at that orientation as being overhanging for Design C and did not flag facets at that orientation for Design A. Will It Print did not identify any of the tested combinations of orientation and geometry as being problematic for tipping, while Meshmixer only identified one combination (Design B, orientation $\mathbf{B} = [0.28, -0.85, -0.44]$). The difference between predictions is likely due to the way the support material was generated in Meshmixer, which created only a few support structures under overhanging faces, whereas Will It Print assumes a more solid, grid-like support structure will be added under all overhanging faces. Generally, there was good agreement between the predictions of both

overhanging features and tipping by the two tools.

Table 4.5: Comparing orientation-dependent errors of the tools at different orientations, where O denotes an overhanging warning on the cylindrical features and T denotes a tipping warning for the whole part.

Model		Warnings at orientation			
		[0, 0, 1]	[0, 1, 0]	[0.28, -0.85, -0.44]	[-0.79, 0.17, 0.59]
Design A	Will It Print	None	O	O	O
	Meshmixer	None	O	O	O
Design B	Will It Print	O	O	O	O
	Meshmixer	O	O	O;T	O
Design C	Will It Print	O	O	O	O
	Meshmixer	None	O	O	O

4.3.3 Discussion

Generally, there were some similarities between Will It Print and the tools we used for comparison. Will It Print has analysis capabilities that are also included in Meshmixer (i.e., identification of overhanging faces, small features, and tipping) and Shapeways (identification of small features). However, Will It Print has types of analysis capabilities that are not available in either tool, namely evaluation of warping and surface roughness. Will It Print also was able to identify regions that were smaller than the recommended minimum feature size more accurately than the other tools.

Because Meshmixer’s approach to thickness measurement relies on shooting rays from each vertex, thick regions close to thin regions are also highlighted as being below the recommended size. This problem is not visible when the resolution of the STL file is extremely high, as it was for Design A, but would be visible for most normal or rough STL files, as we saw in Design B and C. Extraneous highlighting could lead the designer to modify geometry unnecessarily. Similarly, Shapeways 3D Tools only highlighted part of the too-small geometry (possibly also because of the rough STL resolution), which could also lead the designer to make unnecessary changes. Will It Print was most able to highlight features at or below the minimum feature size without highlighting extraneous faces and was robust to STL resolution.

Will It Print was also unique compared to Shapeways 3D Tools and Meshmixer in that it flagged several features that the other tools did not: 90° corners, and small negative features (i.e., small holes). The flagging of sharp corners is a result of the multi-scale top-hat transform process used by Will It Print to identify small features. When parts are printed, sharp corners do print slightly rounded because the printer cannot make infinitely sharp turns

while depositing material. Will It Print conveys this limitation to designers, while the other tools do not. However, we are unsure of the importance of this rounded-corners effect on designers, and the importance of this visualization will need to be assessed further. However, it does seem critical to communicate to designers about too-small negative features like holes, because these are more likely to be crucial for the physical performance and aesthetic look of parts. Will It Print can identify and communicate this important information to designers, while the other tools do not.

The tools also differed in how they were set up for user interactions. One example of this is that Shapeways 3D Tools had a hard-coded threshold for small feature size, while in Meshmixer and Will It Print, this level can be easily adjusted. While hard-coding the threshold may be easier to implement, it does not give the designer the flexibility to use the tool for different machines or AM processes, which have different constraints regarding the small feature threshold.

Computational speed was another difference between the tools. For small features detection, the ray shooting method used by Meshmixer was much faster than the voxelization approach used by Will It Print or the approach used by Shapeways 3D Tools. However, Will It Print was also able to identify small features in a relatively small amount of time, typically under 2 minutes. There appears to be a trade-off between accuracy and speed with the computational approach employed by Meshmixer and Will It Print: the more accurate small feature detection of Will It Print takes more time. In future iterations of Will It Print, we could ask the designer about the criticality of small feature detection, and decrease the computational time by either coarsening the voxel resolution or employing a faster but less-accurate approach like Meshmixer.

Another important consideration in user experience is the visualizations that the tool employs. The visualizations of all three tools tended to be somewhat similar in that they use color to flag portions of the geometry. Meshmixer uses their visualizations to show extra information (such as geometry that is almost but not quite at the threshold for an error). Will It Print could adapt this approach to communicate more nuanced information about compliance with guidelines. For example, rather than using one color to highlight rough surfaces, the color could be scaled according to the estimated surface roughness. However, there may be a point of diminishing returns where dense visualizations clutter the GUI and overwhelm the designer with information.

Information regarding the manufacturability assessments performed by each tool is another consideration in the user experience. Because we targeted Will It Print to be used by novices, with relatively little AM experience, we were careful to include educational feedback for novices in the form of graphical examples and descriptions of each manufacturability guideline we assessed. Similarly, Shapeways 3D Tools gives some helpful background information about the minimum thickness of walls. Meshmixer, on the other hand, gives only basic information about manufacturability, and instead focuses on how to interpret their visualizations. Meshmixer does not give suggestions on how to correct any manufacturability problems. We expect that the guideline information we convey in Will It Print will make the tool more useful for novices who are likely not aware of manufacturability guidelines.

However, all of these decisions regarding user experience should be the subject of more study and user testing.

In summary, Will It Print provides additional analysis capabilities not provided by existing tools while also providing enhanced accuracy and user experience. Designers should be able to use our tool to assess the manufacturability of their design and use the feedback and information provided about the guidelines to make improvements to their design geometry manually. Will It Print also helps designers understand the impact of orientation on AM outcomes, so that they can begin to select a build orientation for their part or to create a design that is relatively insensitive to build orientation. Later in the design process, they can use our other tool, TAAM, to further investigate manufacturing outcomes and build orientation. TAAM will be the subject of the next section.

4.4 Results and Discussion - TAAM

This section will illustrate a typical workflow for TAAM. TAAM is more targeted at experienced designers, who understand and utilize GD&T. TAAM is also targeted to be used later in the design process, once the preliminary detail design is complete, and a rough tolerance scheme has been developed. After a designer explored the manufacturability of a design using Will It Print and developed a detail design with a tolerance scheme, TAAM can be used to ensure the tolerance allocation is appropriate and achievable. Another difference between Will It Print and TAAM is that a designer must manually try out different orientations, while TAAM evaluates many orientations for the designer and suggests promising orientations. These features of TAAM will be illustrated in the following section. We analyze a modified version of Design C using TAAM and other similar tools.

4.4.1 TAAM Output

In TAAM, the geometric error corresponding to each tolerance was calculated using the method described by Budinoff and McMains [2] and summarized in Chapter 3. The data visualizations described in Chapter 3 are used here, with the error associated with each orientation represented as a color, plotted on the surface of a unit sphere. For a full explanation of the data visualization, the reader is referred to Chapter 3.

The error associated with the flatness tolerance, perpendicularity tolerance, and middle cylinder cylindricity tolerance is shown in Fig. 4.14 as a function of the orientation. The other cylindricity errors are not plotted here to limit redundancy. The coordinate axes in these figures correspond to the coordinate axes shown in Fig 4.3. These plots illustrate the error functions described earlier (i.e., Eq. 3.12 and Eq. 3.13). These maps are made by evaluating the error at each orientation and plotting the error as a dot. The color of each dot represents the magnitude of the error at that orientation. The slight fuzziness of the coloring of these plots is due to the discrete nature of the points. Regions with increased

error due to support material being needed are visible in the left, right, and sides of the three plots for flatness, perpendicularity, and cylindricity respectively.




Figure 4.14: Geometric errors of the example part.

In this example, the initially specified tolerances were not all simultaneously achievable at any particular tolerance, as shown in Fig. 4.15. The color scaling of this sphere, developed by Rinaldi [140], shows orientations where all tolerances are met as green, and orientations where some tolerances are not met as shades of red. Even though there is no orientation where all tolerances are simultaneously satisfied (visualized as no green region on the sphere), using the physical programming ranking of the Pareto-optimal points, we can show the designer a particular orientation that best meets her needs. For this example, that orientation was calculated to be $\mathbf{B} = [0.19, 0.47, 0.86]$ (shown as a yellow square in Fig. 4.15 and Fig. 4.16). The object rotated to this orientation is shown in Fig. 4.19i.




Figure 4.15: No orientation satisfies all five initial tolerances. The pareto-optimal orientation $\mathbf{B} = [0.19, 0.47, 0.86]$ is marked with a yellow square.

The errors at this orientation are summarized in Tab. 4.6, along with the error at the initial orientation shown in Fig. 4.3, $\mathbf{B} = [0, 0, 1]$, the orientation of the part used by




Figure 4.16: Example part cost parameter. The Pareto-optimal orientation $[0.19, 0.47, 0.86]$ is marked with a yellow square.

the designer during geometric modeling. Comparing the naive, initial orientation to the orientation output by the program, $[0.19, 0.47, 0.86]$, the errors associated with the flatness and perpendicularity tolerances have increased, but both are in the tolerable range. The cylindricity error has decreased for all cylinders. At this orientation, the insides of the cylindrical bosses no longer need support, which decreases the error on those features. The cost parameter has increased only slightly.

Table 4.6: Errors at naive and optimal orientations suggested by TAAM.

	\mathbf{B}	ϵ_{flat}	ϵ_{perp}	ϵ_{cyl1-3}	λ
Naive orientation	$[0, 0, 1]$	0.10	0.10	0.25, 0.28, 0.25	0%
First optimal orientation	$[0.19, 0.47, 0.86]$	0.15	0.21	0.24, 0.18, 0.18	1%
Second optimal orientation	$[-0.10, -0.34, -0.94]$	0.13	0.18	0.24, 0.21, 0.20	2%

At this point, the designer may choose to relax the tolerances slightly so that the error associated with this “best” orientation falls within the newly specified tolerances. However, the designer may want to understand better how the relaxation of particular tolerances affects the number of feasible orientations. Figure 4.17 shows the GUI TAAM provides for this purpose, with a set of relaxed tolerances for the example problem.

After exploring the effect of relaxing different tolerances, the hypothetical designer decides to relax the cylindrical tolerances to 0.25 mm each, does not switch any tolerances to the 1H class, and does not modify the boundaries of the 1S class functions. Based on these inputs, the feasibility plot is regenerated (Fig. 4.18). A yellow square at the very bottom of the sphere in Fig. 4.18 represents the new preferred orientation, $\mathbf{B} = [-0.10, -0.34, -0.94]$, output by the program, based on the revised tolerances. At this new preferred orientation, there is slightly increased error of the cylindrical tolerances, but the flatness and perpendic-

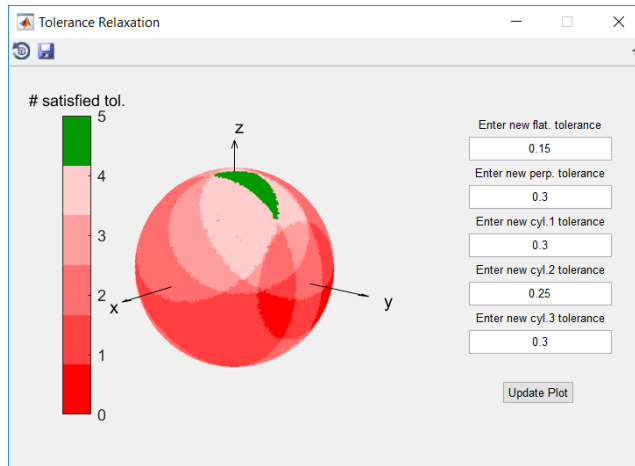


Figure 4.17: Tolerance relaxation GUI.

ularity error has been improved over the previous orientation (summarized in Tab. 4.6). The part is displayed in this new preferred orientation in Fig. 4.19ii. At this new orientation, all five errors are within the new tolerances, as opposed to the initial, naive orientation, where only the flatness and perpendicularity error satisfied the original tolerance allocation.




Figure 4.18: Feasible orientations are visible with revised tolerances. The Pareto-optimal orientation $[-0.10, -0.34, -0.94]$ is marked with a yellow square (the back side of the sphere is shown).

The visualizations used in the figures shown in this section used 40,000 data points. If desired, the number of points generated can be reduced, reducing the computation time, but the visualization becomes sparser and potentially less informative (Fig. 4.20).




Figure 4.19: Example part at the pareto-optimal orientations.

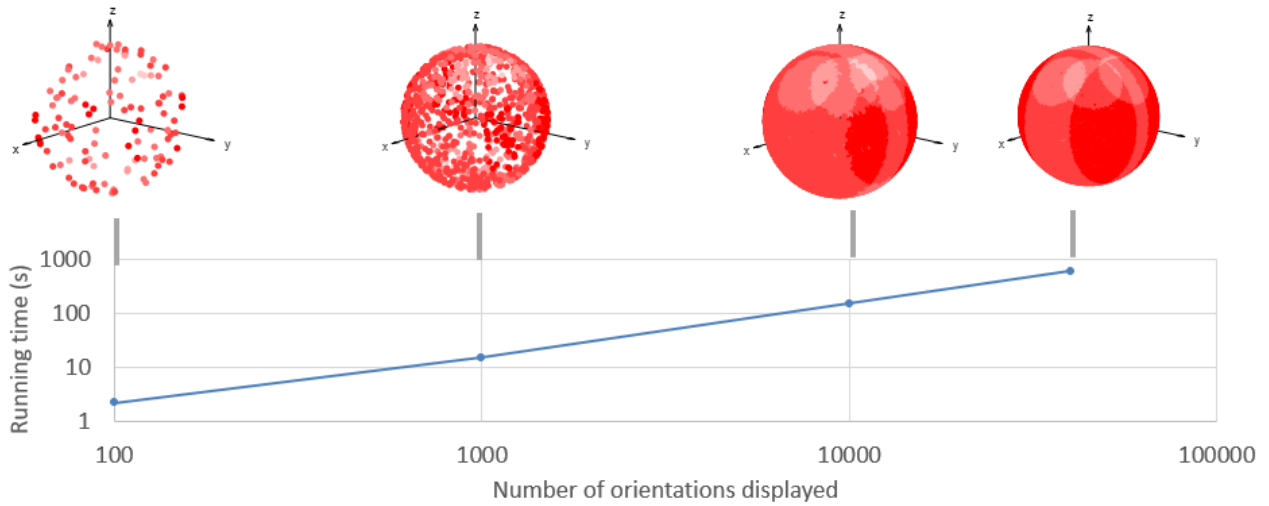


Figure 4.20: As the number of orientations at which error and part cost are calculated increases, the total computational time increases, so a balance between dense visualization and speed must be identified.

4.4.2 Comparison of TAAM to other orientation optimization tools

Now that we have described the workflow and output of TAAM, we will compare TAAM to two other approaches described in previous literature: ARO and PTO.

ARO

For the more typical approach, average roughness optimization (ARO), the designer needs choose scalar weights for the two objectives (surface roughness and print time). The hypothetical designer initially set the weight for the two objectives to be equal ($\omega = 0.5$). The optimal orientation with these weights was the initial orientation, with $\mathbf{B}=[0, 0, 1]$. The designer decides to try to improve the surface roughness specifically on the cylinders, but because the ARO approach only calculates an average roughness over the whole part, the only option to achieve this goal is to increase the scalar weight for the average roughness. The designer changes the weights to $\omega_1=0.9$ and $\omega_2=0.1$. These weights resulted in this same optimal orientation, with no improvement in surface roughness. When the weights were changed to $\omega_1=1$ and $\omega_2=0$, the optimal orientation was $\mathbf{B}=[0.34, 0.34, 0.88]$. This orientation will result in slightly improved surface roughness on the cylinders over the default orientation.

Table 4.7: Comparing average roughness optimization outcomes.

Weights	Ra_{avg} [μm]	\bar{t} [sec]	O	Suggested \mathbf{B}
$\omega_1 = 0.5, \omega_2 = 0.5$	15.6	222	118.7	$[0, 0, 1]$
$\omega_1 = 0.9, \omega_2 = 0.1$	15.6	222	36.2	$[0, 0, 1]$
$\omega_1 = 1, \omega_2 = 0$	15.4	234	15.4	$[0.34, 0.34, 0.88]$

PTO

Like the ARO approach, the penalized tolerance optimization (PTO) approach requires the designer to set scalar weights for all included objectives. The minimization approach used in the Das et al. model also requires an initial point at which to begin the minimization search. Using the weights $\omega_s = 0.6$, $\omega_{cyl} = 0.2$, $\omega_{flat} = 0.1$, $\omega_{perp} = 0.1$, and an initial point ($\theta = 90^\circ$, $\phi = 270^\circ$), the build vector $[-0.39, -0.39, -0.84]$ was identified by the optimization algorithm as minimizing the objective function. Using the same weights but initial point ($\phi = 90^\circ$, $\theta = 90^\circ$), $\mathbf{B} = [0.38, 0.40, 0.83]$ was identified as the optimal orientation. Using a range of other initial points did not result in any additional optimal orientations, so the designer assumed that there are two local minima for the objective function using these weights. The penalties, p , normalized volume, V_{s-norm} , and objective function value, E , were found for each optimal point identified (Tab. 4.8).

The hypothetical designer decides that the flatness and perpendicular tolerances are more important than the cylindricity tolerances and decides to adjust the weights, to $\omega_s = 0.3$, $\omega_{cyl} = 0.2$, $\omega_{flat} = 0.3$, $\omega_{perp} = 0.2$. Again, two different orientations are suggested, depending on the starting point chosen. Objective function values, which are the equal at these orientations, are also shown in Tab. 4.8. The designer could select one of the two suggested orientations. They could also evaluate the error associated with each tolerance callout at this orientation, and relax or tighten tolerances as necessary. However, the PTO

approach is based on an optimization algorithm, which would need to be rerun each time the tolerances were adjusted.

Table 4.8: Penalized tolerance optimization optimal orientations for a given set of weights, with objective function values calculated at each orientation.

Weights	Initial	Penalties	Objec. value		Suggested \mathbf{B}
	θ, ϕ	$p_{flat}, p_{perp}, p_{cyl1-3}$	V_{s-norm}	E	
$\omega_s = 0.6,$ $\omega_{cyl} = 0.2,$ $\omega_{flat} = 0.1,$ $\omega_{perp} = 0.1$	90,270	0.15, 0.15, 0.25, 0.05, 0.25	0.02	0.14	$[-0.39, -0.39, -0.84]$
	90,90	0.15, 0.16, 0.26, 0.04, 0.24	0.003	0.14	$[0.38, 0.40, 0.83]$
$\omega_s = 0.3,$ $\omega_{cyl} = 0.2,$ $\omega_{flat} = 0.3,$ $\omega_{perp} = 0.2$	180,270	0.09, 0.12, 0.30, 0.10, 0.23	0.02	0.18	$[-0.29, -0.35, -0.89]$
	0,10	0.09, 0.12, 0.30, 0.10, 0.23	0.02	0.18	$[0.30, 0.35, 0.89]$

TAAM and the other approaches we implemented have different objectives, visualizations (or lack thereof), and approaches to identify optimal points, but they all attempt to find an orientation that optimizes cost (as measured by print time, support volume, or total print cost) while also optimizing quality (as measured by surface roughness or errors associated with tolerances). Tab. 4.9 summarizes all three approaches and the optimal build vector identified by the tool, after adjusting boundaries or weights. The table lists computation times, showing that the PTO approach is the quickest. However, the computation time for the PTO approach excludes calculating the maximum support volume at a wide range of orientations (a value needed for calculating the support volume penalty), which is pre-calculated and can take several minutes.


However, it should also be noted that the computation times for the ARO and PTO approaches only result in giving the designer one suggested optimal point, while the approach we take in TAAM also provides the designer with the cost and error information at many other possible orientations. The final orientation selected by the designer using each of these approaches is plotted in Fig. 4.21.

4.4.3 Discussion


There were several similarities between TAAM and the tools we used for comparison. TAAM performs a similar function to orientation optimization schemes that have been developed previously, in that it suggests orientations to minimize geometric error and printing cost simultaneously. The orientations recommended by TAAM were similar to those recommended

Table 4.9: Comparing performance and predictions of orientation optimization approaches.


Tool	Objectives	Comp. time	Modeled sources of geometric error	Suggested \mathbf{B}
TAAM	Geometric tolerances; printing cost	122.2 s	Stair-step effect; machine precision; layer profile; support material	$[-0.10, -0.34, -0.94]$
Average roughness optimization (ARO)	Average surface roughness; printing time	108.0 s	Stair-step effect; layer profile; support material	$[0.34, 0.34, 0.88]$
Penalized tolerance optimization (PTO)	Geometric tolerances; volume of material	3.1 s	Stair-step effect	$[-0.29, -0.35, -0.89]$



(i) TAAM orientation, $\mathbf{B}=[-0.10, -0.34, -0.94]$



(ii) ARO orientation, $\mathbf{B}=[0.34, 0.34, 0.88]$



(iii) PTO orientation, $\mathbf{B}=[-0.29, -0.35, -0.89]$

Figure 4.21: Example part at optimal orientations identified using different algorithms and designer inputs.

by both tools we compared it with, which makes sense because each was optimizing for similar goals.

The approach taken by PTO is similar to the approach we take with TAAM, but it does not take into account several key sources of geometric error: positioning error and support material removal marring the surface. In PTO, the effect of support material is only considered in how it affects the total print volume, so if a designer is not concerned with cost or printing time, that weight can be set to zero, and an orientation with lots of support material would be recommended, despite the fact that support material removal has been found to damage part surfaces [102]. The predictions of the PTO approach provide a helpful

starting point for assessing the achievability of tolerances, but they are less accurate than our estimates, based on the above limitations.

Another difference is that the PTO approach uses a minimization scheme to identify a single, optimal point. This scheme, however, does not give the designer enough information to understand what tolerances they could relax, and by what amount each tolerance needs to be relaxed (information we provide using our spherical data visualizations). Another limitation of the minimization technique used in PTO is that it cannot identify global minima, only local minima, so depending on the initial point chosen by the designer, a more suitable orientation may exist but may not be identified.

A benefit of TAAM and the PTO approach, when compared with ARO, is that they analyze geometric tolerances. Geometric tolerances are how designers specify allowable geometric error, so it makes sense to predict and optimize geometric errors using the same language designers speak: geometric dimensioning and tolerancing (GD&T). Anecdotally, through discussions with industry representatives at conferences, it seems that tolerance allocation is a problem that needs further support. Industry representatives have commented that this tool is helpful and necessary. While surface roughness is certainly an important consideration, it may be better to combine it with other ways of quantifying geometric error, like GD&T, rather than using it as the only metric.

Physical programming, used in the TAAM tool, has advantages over the weighted sum-based approach, used in the PTO and ARO approaches. Although the objective function used in the ARO approach takes into account both print time and average surface roughness, it tends to weigh the print time objective more, because that objective has a larger unweighted magnitude. Unless the weights are scaled according to the magnitude of the different objectives (which also have different units) this approach does not consider the two objectives on equal footing. The PTO approach uses a weighted sum objective function that scales each objective so that it can only vary between 0 and 1, which avoids this problem. However, the setting of scalar weights can be somewhat arbitrary (i.e., how can the designer determine if the quality on one face is a 0.25 or a 0.3?). The physical programming boundaries are more intuitive than scalar weights.

There are potential issues with using an average error metric for the whole part, as is done in the ARO. As described in the results, it is impossible to prioritize quality on a particular face when only using a metric that averages quality over the whole part. Another limitation of this approach is that calculating an average over all features means that large faces, even if they are less important to the part function, are given more weight than small features. Designers set different tolerances on different features because non-functional features need much looser tolerances than functional features. A design-support tool should reflect these needs by allowing designers to set different levels for each feature, as is done in TAAM and by the PTO approach.

Compared to existing tools, our tool offers more helpful features. TAAM has additional functions in helping designers understand the range of achievable tolerances and exploring which tolerances they may want to relax. Rather than subjective or guess-and-check approaches for adjusting scalar weights and choosing an orientation, TAAM allows a designer

to visualize each objective for all possible orientations at once, which aids in faster evaluation. These visualizations and the interactive nature of the tool make trade-offs between geometric accuracy and production time explicit. Our visualization allows for quick analysis showing not only the feasible orientations but also the *almost* feasible orientations, where some but not all specified tolerances are met. It is hoped that this visualization, along with the interactive GUI, will help designers quickly explore the relationships between each objective.

There are other options for presenting designers with visual analyses of manufacturability information. Rinaldi has previously used a 2D projection of the orientations [140]. We plan to explore other forms of data visualization, including parallel coordinates, in the future. In order to examine the effectiveness of the data visualizations used in TAAM, we hope to test the proposed tool with actual designers. We assume that designers, when presented with manufacturability information, will be willing to selectively relax less critical tolerances in the manner described in this chapter. Because the tool can allow the designer to examine ranges of options, rather than immediately selecting one orientation based on assumed preferences as previous tools have done, we hypothesize that the designer will feel more satisfied with the final, optimized orientation.

Another benefit of TAAM is its generality: it can be applied to many different AM processes, incorporating different sources of geometric error, machine-specific deviations, layer thickness, and support generation cut-off angles into the error calculations. Experimental verification, which will be discussed in the next chapter, will help refine the mathematical analyses.

An additional benefit is TAAM's computational speed. Error calculations can be computed quickly at all orientations, enabling designers to interactively explore trade-offs and tolerance allocation during the highly iterative design stage. Other optimization tools for AM process planning are not designed to be interactive, and require the designer to rerun the optimization each time they decide to adjust scalar weights. Our approach with TAAM is to calculate error at many orientations, giving the designer a sense of the range of options available to them. The initial computation of the error is in near real-time (calculating error at 10,000 orientations for this example took 1.5 seconds on a laptop) and the relaxation GUI can update the feasibility sphere in real time, because the same orientations and error calculations are re-used and only changes to plotting color are made, after comparing the already-calculated-error to the revised tolerances. Calculation of cost is slower due to the calculation of support volume, but this calculation is only performed once and is not recalculated during tolerance relaxation. The computation time could be reduced using GPU techniques [167].

There is a trade-off between dense, informative, visualizations, and computation time. Our approach of calculating error at many orientations is computationally expensive, though, and takes more time than a simple minimization like that used in PTO. There may be situations where a designer wants a single, optimal orientation, without wanting to explore other options or to relax tolerances, and a minimization resulting in only one suggested orientation may be preferable in that situation. However, we expect that most designers are

willing to relax some tolerances slightly in order to achieve a cheaper manufacturing process, and our tool is much more informative and helpful for that process.

TAAM also requires minimal initial input from the designer, making educated guesses about the designer's goals. The designer can explore feasible solutions and then refine these assumed preferences. The time the designer spends interacting with the tool is spent productively refining the preference structure and exploring the data, not tediously defining preferences or guessing at objective weights.

If the goal of developing new algorithms and analyses to help designers understand manufacturing constraints is to create systems that designers will want to use, it is important to understand designers needs and to speak their language. TAAM has more features, include more sources of geometric error, and quantify geometric error in a way that is meaningful to mechanical designers.

4.5 Conclusion

In this chapter, we have presented an example design scenario illustrating the use of Will It Print and TAAM. Will It Print has additional capabilities not found in other existing tools. The comprehensive analysis capabilities of Will It Print, which were developed through analysis of common errors described by students, will help novices using the tool identify avoid printing failures. Unlike any widely distributed tool that the authors are aware of, Will It Print can identify areas that will have high surface roughness, assess the risk of warping, and identify small negative features like holes. Additionally, Will It Print more accurately identifies small features than other available tools. In the design scenario, Will It Print identified several manufacturability problems with the simple part used in this example. Will It Print was designed to be helpful for novices seeking to apply design for manufacturing principles to their designs because it helps them evaluate compliance and provides them with enough information about the guidelines that they can make educated changes to their part or its selected build orientation.

TAAM allows designers to predict achievable tolerances for parts produced using AM. Designers can explore which orientations enable them to meet all specified tolerances while also considering the printing time and associated cost. A ranking system based on meaningful, easy-to-understand interpretations of designer preferences was implemented to minimize demands on the designer's time and patience. Compact data visualizations were presented to quickly convey a large amount of information regarding objectives and their interconnectivity. Additionally, an interactive, iterative GUI for tolerance relaxation was presented so designers can learn and shape their preferences based on what-if scenarios of tolerance refinement and relaxation. Designers must analyze and optimize often conflicting objectives and requirements while designing products and preparing for production. As part of the DFM process, designers seek to minimize cost while maximizing the quality of the parts that are produced. In order to do this, the trade-offs of quality and cost must be carefully analyzed. As AM becomes an increasingly viable option for production parts, it is necessary

to develop tools specifically for AM to assist designers in trade-off analysis and subsequent decision making.

This design scenario illustrated some overlap between Will It Print and TAAM. The final orientation selected by the theoretical designer using Will It Print is very similar to the orientation selected using the orientation optimization built into TAAM. This similarity makes some intuitive sense because both tools evaluate some measurement of the stair-step error, and both tools suggest the user stay away from very tall parts (Will It Print does so because of the risk of tipping, while TAAM does so because of the higher cost of printing). The tools were purposely developed to suit the needs of different users at different times in the design process, and for this reason, they are not combined into a single tool. However, the similarity between the tools highlights opportunities to adapt the functionality of one tool into the other. For example, adding an automatic assessment of manufacturability constraints like warping and tipping into TAAM rather than assuming the designer has already assessed manufacturability could further its usefulness. We will further evaluate opportunities to improve the capabilities of both tools in future work.

In summary, this chapter demonstrated that our geometry analysis tools are practical to implement, with relatively low computation times, a high level of interactivity, and features that are not currently offered by other existing systems. We expect that these systems will be helpful for designers in understanding complex relationships between geometry and manufacturing outcomes.

Chapter 5

Experimental verification of geometry analysis

The focus of the previous chapters has been on detailing the development and implementation of the geometry analysis tools for design-for-additive-manufacturing feedback. The development of two tools was discussed in detail: a tolerance allocation tool for additive manufacturing (TAAM) and Will It Print, a tool that helps assess manufacturability. This chapter focuses on experimentally validating the predictions of our systems.

5.1 Introduction

This chapter seeks to validate the predictions of the tools developed in this thesis and identify areas to improve. For TAAM, the prediction that we seek to validate is the mathematical model used to predict error associated with different geometric tolerances (e.g., $\epsilon_{flat}, \epsilon_{cyl}$). For Will It Print, we seek to validate that the thresholds for compliance with manufacturability criteria accurately capture transitions from acceptable print quality to unacceptable print quality.

Accurately predicting what level of geometric tolerance is achievable on different part features, as is our goal in TAAM, should fulfill an unmet need for designers. To this point, little has been published in terms of experimental evaluation of achievable geometric tolerances in AM. Some researchers [21, 85, 86, 110, 111] have published experimental results of achievable geometric tolerances of test parts printed using different AM processes, but they did not evaluate the sources of geometric error or seek to build a predictive model. Given the relative lack of literature describing even approximate ranges of achievable geometric tolerances for AM processes, it is not surprising that few tools exist to predict quality. Outside of TAAM and similar tools discussed in Chapter 3, there are currently few models available to predict how different process parameters may impact the accuracy of features on a part produced using AM. Some empirical models for part quality have been developed entirely from experimental data [168, 169]. However, the extent to which these models can be applied more generally is unclear because they are based on experiments conducted using

a single process and machine. A benefit of the approach we have taken with TAAM is that, because it is based on first principles, it is general and can be adapted to a wide range of AM machines, different materials, and even different AM processes. However, experimentation is needed to confirm that the important causes of geometric error are included in the model and that the model predictions accurately reflect reality.

As discussed in Chapter 2, there are many sources of geometric error in parts made by additive manufacturing. The accuracy of different features on a part are affected by process parameters (such as the orientation and layer thickness), the precision and repeatability of the machine itself, and by factors beyond the direct control of the user, like shrinkage [21, 85, 92]. It is unclear which sources of geometric error are most impactful for different kinds of geometric tolerances. Because TAAM was developed to analyze cylindricity, flatness, and orientation tolerances, we hope to understand better what drives geometric errors related to those tolerances specifically.

Like tolerance prediction, manufacturability criteria and design guidelines for AM have been developed but are not frequently built into any predictive or analytic system. Because the Will It Print system uses a mixed theoretical and empirical approach to assessing the manufacturability of parts (as summarized in Chapter 3), it is already built on actual observations. However, further experimentation is needed to validate the predictions of the tool and ensure that the predictions are widely applicable, beyond the parts and observations that were used to develop the tool.

Generally, there has been some prior work evaluating experimental error, and some prior work developing mathematical models to predict geometric error, but these two bodies of work are disconnected. The developers of predictive models have not generally attempted to experimentally validate their models, and experimentalists have not tried to validate existing models. We undertake experimental testing of our geometry analysis to close this gap between theory and practice. We seek to evaluate whether mathematical models and manufacturability heuristics actually capture the physical results observed in FDM printing, and to evaluate the relative impact of different sources of error.

This chapter describes the experiments we conducted to validate the geometry analysis tools described in Chapter 3 and Chapter 4, namely TAAM and Will It Print. Although the materials and methods section in this current chapter applies to all of the experiments we conducted, the results section is divided into two studies. Study 1 validates the manufacturability predictions of the Will It Print tool. Study 2 focuses on the validation of the quality predictions of the TAAM tool and the analysis of general sources of geometric error.

5.2 Materials and methods

Details regarding our experimentation procedure will be provided in the following paragraphs, but a rough overview of the process is as follows. An experimental part was chosen as well as a geometric dimensioning and tolerancing (GD&T) scheme, which defined the type and magnitude of geometric error allowed on relevant features. The part was printed on a

3D printer at a particular build orientation. Then, the part was removed from the build plate, and all support material was manually removed. Finally, the part was scanned, and the point cloud representing the as-printed geometry was processed through the metrology software, PolyWorks Inspector, to quantify geometric error on the part.

5.2.1 AM equipment

All of the experimentation discussed here used FDM machines because these machines were the most widely available on the University of California, Berkeley and the University of Arizona campuses, where all experimentation was conducted. Type A Series 1 Pro FDM printers using polylactic acid (PLA) filament with a heated print bed were used for earlier experimentation and for manufacturability tests. Additional parts were printed in PLA on a LulzBot TAZ 5 and an Ultimaker 3. A Stratasys Dimension 1200ES SST and Zortrax M200, both of which use acrylonitrile butadiene styrene (ABS), were also used. The specifications of these printers are given in Tab. 5.1.

Table 5.1: Specifications of 3D printers and materials used in experiments.

Printer	Precision [μm]		Layer height [μm]	Material
	<i>xy</i> -directions	<i>z</i> -direction		
Type A Series 1 Pro [170]	6.67	6.25	50 - 250	PLA
Zortrax M200 [171]	1.5		90 - 390	ABS
Stratasys Dimension 1200ES SST			254 - 330	ABS
LulzBot TAZ 5 [172]	50		75 - 500	PLA
Ultimaker 3 [173]	12.5	2.5	20 - 600	PLA

5.2.2 Parts printed

Several different parts, summarized in Tab. 5.2, were printed as part of the experimentation, as part of Study 1 and Study 2. Some of these parts (Fig. 5.1) were printed to validate the predictions of the Will It Print tool as part of Study 1. Because the main focus of Study 1 was validating the most novel predictions of Will It Print, namely warping and small feature detection, our experimental parts were chosen mainly to test those two predictions. A flat rectangular prism was created to test warping. This part was printed with a thickness of either 3.5 mm or 15 mm, and the corners were 90° corners or filleted with a radius of 5 mm.


Several tapered cylinder parts were printed with varying diameters to test tipping. A part, based on that of the small hole benchmark part described in [121], was used to measure the minimum negative feature size (note: the holes with a diameter of 1 mm, which falls in the middle of the range of holes tested, were accidentally excluded from the part we printed).

Table 5.2: Description of parts printed.

Part name	Size [mm]	Material	Measurements	Printer	Layer height [mm]
GE bracket	40×55×110	ABS	Flatness; Orientation	Zortrax M200	0.2
Flagpole bracket	110×110×85	ABS	Flatness; Cylindricity	Stratasys Dimension	0.254
Heptagonal prism	50×40×45	ABS	Flatness; Orientation	Zortrax M200	0.2
Flat rectangular prisms	140×40×3.5; 140×40×15; 170×40×3.5; 170×40×15	PLA	Deviation due to warpage	Ultimaker 3; LulzBot	0.2
Tapered cylinder part	135×(10-30)	PLA	Tipping threshold	Type A S1Pro	0.4
Hole feature benchmark part	90×75×13	PLA	Small feature threshold	Type A S1Pro	0.4


For Study 2, several different parts were printed. A few of these parts have been introduced in previous chapters. The General Electric (GE) bracket was introduced in Chapter 3. The flagpole bracket was introduced in Chapter 4. We wanted to assess general trends associated with various orientations and achievable tolerances. To achieve this goal, we added additional callouts on the GE bracket and flagpole bracket and measured the error associated with these additional callouts. Additionally, a heptagonal prism was designed to test the TAAM’s performance for a wider range of build orientations. Figure 5.2 shows images of these parts, and Tab. 5.2 summarizes which printers they were printed on. Detailed drawings with key dimensions are included in Appendix A.

As discussed in Chapter 3, a crucial component of TAAM’s predictions of achievable tolerances is estimating the geometric error caused by the stair-step effect, which depends on the build orientation of the part. In order to investigate the effect of orientation on



(i) Flat rectangular prism (ii) Tapered cylinder part (iii) Hole feature benchmark part

Figure 5.1: Parts used in Study 1 to validate Will It Print.



(i) GE bracket

(ii) Flagpole bracket

(iii) Heptagonal prism

(iv) Angled cylinders part

Figure 5.2: Parts used in Study 2 to validate TAAM.

achievable quality and test the hypothesis that orientation would be the primary driver of geometric error, the parts used in Study 1 were printed at different orientations. The GE bracket was printed at three orientations (Fig. 5.3). The support material (visible underneath the overhanging features) was not removed at the time that these pictures were taken. The flagpole bracket was printed at two orientations (Fig. 5.4). The heptagonal prism was printed at four orientations (Fig. 5.5).




Figure 5.3: GE bracket printed at different orientations.




Figure 5.4: Flagpole bracket printed at different orientations (support material has been removed).

5.2.3 TAAM measurement procedure

All experimental parts for Study 2, focusing on TAAM, were scanned using a ROMER Absolute Arm 7525 SEI produced by Hexagon Manufacturing Intelligence, which has a minimum point spacing of 0.014mm and a theoretical scanning accuracy of 0.063 mm [174]. These scans generated point clouds of the as-printed geometry, allowing us to quantify the differences in shape and size between the nominal (perfect) geometry defined in the CAD model and the actual printed part. Then, we can use metrology software to determine the error corresponding to different GD&T callouts we defined on part features. Each large planar feature on the GE bracket, flagpole bracket, and heptagonal prism was inspected for flatness as well as orientation error (relative to several different datums). Cylinders on the flagpole




Figure 5.5: Heptagonal prisms at different build orientations (from left to right: $\mathbf{B} = [1, 0, 0]$, $\mathbf{B} = [0.87, 0, 0.5]$, $\mathbf{B} = [0.5, 0, 0.87]$, $\mathbf{B} = [0, 0, 1]$).

bracket and angled cylinders part were inspected for cylindricity error.

We tried two different workflows with the point clouds: importing the point cloud into SolidWorks via the Geomagic add-in, creating a mesh, and exporting that mesh into InnovMetric PolyWorks Inspector metrology software for error calculation; and importing the point cloud directly into PolyWorks Inspector for error calculation. These two different workflows produced nearly identical results ($R^2=.96$ for flatness and orientation measurements on the two heptagonal prisms built at orientations $\mathbf{B} = [0, 0, 1]$ and $\mathbf{B} = [0.5, 0, 0.87]$). The more straightforward workflow of importing directly into PolyWorks Inspector was used for all remaining scans, and those are the results reported here.

An example of the flatness measurement process is shown in Fig. 5.6. Flatness and orientation tolerances, as discussed in Chapter 3, are measured by enclosing the entire toleranced surface between two parallel planes (for orientation, those planes are aligned relative to a datum plane) and measuring the distance between the planes. Similarly, cylindricity is measured by finding the distance between two concentric cylinders that encompass the part's entire surface.

5.2.4 TAAM implementation and comparison

In preliminary testing with FDM, we found that the deviation of a flat face oriented parallel to the build platform ($\theta = 0$) was dominated by the depressions between roads of deposited filament. The magnitude of the deviation caused by the depressions between roads was larger than any small layer-level deviation caused by shrinkage or positioning error represented by δ_z in the mathematical models described in Chapter 3. To account for this finding, below we update Eq. 3.12 for orientation and flatness errors so that, when $\theta = 0$, the corresponding error is equal to $\delta_{z,0}$, an experimentally determined value. When $\theta < 1^\circ$, it is likely that the geometric error would also equal $\delta_{z,0}$, but we did not evaluate error at very small θ . The updated mathematical models that we attempt to validate in this chapter are reproduced




Figure 5.6: The experimental setup involves a laser scanner (black) that scans the part surface (blue) to find points (orange) and then calculates error corresponding to geometric tolerances, such as ϵ_{flat} , ϵ_{par} , and ϵ_{cyl} .

below:

$$\epsilon_{flat}, \epsilon_{ang}, \epsilon_{perp}, \epsilon_{par} = \begin{cases} \delta_{z,0} & \theta = 0 \\ \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) |\cos \theta| + \delta_{xy} \sin \theta & 0 < \theta < \theta_{sup} \\ \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) |\cos \theta| + \delta_{xy} \sin \theta + \delta_s & \theta_{sup} \leq \theta < \pi \\ \delta_z + \delta_s & \theta = \pi \end{cases} \quad (5.1)$$

$$\epsilon_{cyl} = \begin{cases} \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) \sin \theta_{ax} + \delta_{xy} |\cos \theta_{ax}| & 0 \leq \theta_{ax} < \theta_{sup} - \frac{\pi}{2} \\ \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) \sin \theta_{ax} + \delta_{xy} |\cos \theta_{ax}| + \delta_s & \theta_{sup} - \frac{\pi}{2} \leq \theta_{ax} \leq \frac{3\pi}{2} - \theta_{sup} \\ \frac{h}{2} + \left(\frac{h}{2} + \delta_z\right) \sin \theta_{ax} + \delta_{xy} |\cos \theta_{ax}| & \frac{3\pi}{2} - \theta_{sup} < \theta_{ax} \leq \pi \end{cases} \quad (5.2)$$

where h is the layer thickness, δ_s is the support material deviation, δ_{xy} and δ_z are the xy - and z -direction geometric deviations of each layer, θ is the angle between a face's normal vector and the build direction, θ_{ax} is the angle between a cylinder's axis and the build direction, and θ_{sup} is the angle at which support material is added. The following settings were for these mathematical models, based on achieving a good fit the experimental data, $\delta_{xy} = 0.02$, $\delta_z = 0.02$, and $\delta_s = 0.3$, and $\delta_{z,0}$ equal to 0.2 mm.

We also compared our predictions to the mathematical model proposed by Das et al. [138], given in Eq. 4.3. This mathematical model approximates each layer as a rectangle and does not include the impact of support material, or layer-level deviations. This mathematical model will be referred to as the “simple model.”

5.3 Results - Study 1: Will It Print

The results are organized into two subsections: one for Study 1 and one for Study 2. Study 1, which is described in this current subsection, focuses on predicting combinations of part geometry and build orientation that impact the manufacturability of a part and are likely to lead to build failures. We will discuss only the predictions for warping, tipping, and small features, because the impact of overhanging surfaces in need of support and surface roughness has been studied elsewhere, as summarized in Chapter 2.

5.3.1 Warping

As discussed in Chapter 2, warping can be present on AM-printed parts, especially large parts. While a small amount of warpage is often acceptable to designers, large amounts of warpage can cause repeated build failures. Warping tends to be most heavily impacted by the largest dimension of the part (i.e., the length for rectangular objects) [80].

In order to validate the predictions of Will It Print regarding warping, and to identify additional factors that should be included in the prediction algorithms in the future, a series of experiments were conducted. For testing warping, we printed relatively simple, thin rectangular prisms. A rectangular shape was chosen to limit the impact of extraneous geometry factors on warping, and because the flat rectangular prism shape has been used extensively in other experimental studies examining warping [80, 84, 149, 151]. Our first warping experiments validated the current predictions implemented in Will It Print by validating the warping threshold. Then we evaluated the relative impact of different part dimensions and geometry features. Finally, we investigated how the type of machine used impacted warpage. Each of these topics will be summarized in the following subsections.

Warping threshold

In Will It Print, we used a length of 80 mm as the threshold Λ_w above which the system would issue a warning saying that warping was likely. This threshold was based on some preliminary experimental testing using the Type A Series 1 Pro. Although the threshold of acceptable flatness error due to warping is subjective, the effect of warping should be low enough that the part can finish printing successfully and should not be visibly misshapen. We wanted to verify that the threshold we set was appropriate.

Several of the flat rectangular prisms (Fig. 5.1i) were printed on the Type A Printer. On Type A printer, 5 out of 7 attempts to print the thin rectangular prism with a length of 140 mm failed. The part would warp substantially, with the corners lifting off of the build plate, and eventually, the part would detach before printing finished. After observing the repeated failures, we limited testing on the Type A to shorter prisms, with a length of 100 mm. However, the warpage was still substantial, with visible curling of the prism (Fig. 5.7), and flatness error on the top face of 0.8 mm. This level of deviation and high failure rate confirmed that a roughly 80 mm threshold was appropriate for Type A printers.




Figure 5.7: A “successful” print of the flat rectangular prism on the Type A Series 1 Pro, showing extreme deviation due to warping at the corners of the part.

We decided to continue testing of warpage using other printers that we observed to have less warpage, namely the Lulzbot and the Ultimaker. Our first set of experiments, conducted using the Lulzbot, were conducted to confirm that length was the best predictor of the risk of warping, to assess if fillets help reduce warpage, and to assess if thin parts were at significantly higher risk of warping than thicker parts.

Impact of dimensions and geometry

In order to study the impact of different dimensions, several different versions of the flat rectangular prism geometry were printed on the LulzBot. The width of all prisms was 40 mm. The length and thickness both had two levels: 140 mm and 170 mm for the length; and 3.5 mm and 15 mm for the thickness. To test the impact of fillets (filleting sharp edges is frequently recommended to reduce warpage in online 3D printing forums), half of the parts were printed with fillets of 5 mm on the edges and the other half were printed with 90° corners. We performed a full factorial analysis with eight combinations of factors. Only one replicate of each combination was performed due to time and cost constraints. In general, the magnitude of warpage (Fig. 5.8) was less significant than on the parts printed on the Type A printers.



Figure 5.8: Flat rectangular prism with fillets printed on the LulzBot printer, showing a small amount of warpage, with center of part deformed so that it is higher than the sides.

The rectangular prisms were inspected after printing to find the flatness error on the bottom face. The flatness error was used as a measure of the amount of warpage on the part. A three-way analysis of variance (ANOVA) was completed to assess the relative impact of fillets, length, and thickness on the amount of warpage observed on the parts. No significant interaction effects were found when two-way interactions were included in the ANOVA, and

so the ANOVA was re-run to exclude these interactions. The results of this ANOVA are shown in Tab. 5.3.

Table 5.3: Analysis of variance evaluating relative impact of different dimensions and geometry.

Source	DOF	SS	F	<i>p</i> -value
Length	1	0.03864	2.04	.226
Fillets	1	0.00259	0.14	.7299
Thickness	1	0.22378	11.84	.0263
Error	4	0.07561		
Total	7	0.34062		

The only significant main effect was found to be the thickness of the part. In previous studies [80, 84, 149, 151], only thin parts were tested, but with the Type A, we noticed that warping was present on both thick and thin parts. Although warping was present on thick parts in our study, the warping was less significant than on thin parts with the same geometry. Length, although its effect was not significant given our small sample size, seemed to be related to an increase in warpage. Fillets were not found to affect the amount of warpage significantly. These effects are shown graphically in Fig. 5.9.




Figure 5.9: Main effects plot for the first analysis of variance, showing that thickness and length have largest impact on flatness error for the LulzBot printer.

Variations between different machines

As summarized in the last subsection, we observed much less warpage on parts printed on the LulzBot than on parts printed on the Type A printer. To more quantitatively study the impact of type of printer on the amount of warpage observed, we re-printed the four flat rectangular prism specimens with a length of 140 mm on an Ultimaker printer. The results of the three-way analysis of variance, examining the effects of machine type, fillets, and thickness are summarized in Tab. 5.4.

Table 5.4: Analysis of variance evaluating relative impact of machine type

Source	DOFs	SS	F	<i>p</i> -value
Machine	1	0.07239	1313.2	.0176
Fillets	1	0.01758	318.88	.0356
Thickness	1	0.0005	9	.2048
Machine*Fillets	1	0.0005	9	.2048
Machine*Thickness	1	0.07277	1320.11	.0175
Fillets*Thickness	1	0.0027	49	.0903
Error	1	0.00006		
Total	7	0.16649		

In this ANOVA, the two-way interaction between machine and thickness was significant, which makes it difficult to interpret the main effects of machine, fillets, and thickness. Interaction plots are shown in Fig. 5.10. For the LulzBot, less warpage was observed for the thicker parts, but the opposite was true for the Ultimaker. In general, the Ultimaker had even less warpage than was seen on the LulzBot. In a similar result to the first ANOVA, we did not identify that fillets helped reduce warpage on either machine. The implications of these findings will be discussed in the discussion for Study 1.

5.3.2 Tipping

In order to validate the threshold used by Will It Print to flag combinations of geometry and orientation that are at risk of tipping, several parts were printed in risky orientations using the Type A. The threshold for the tipping warning, κ_{tip} , was set to 1.5, based on experimentation using the Type A printer. To validate this threshold, we printed several other parts on the Type A printer. Most of the parts we printed were a variation of the tapered cylinder part, which was designed to have a small area in contact with the build plate and a tall center of mass. Most of the attempts to print this part ended in print failure, with the part becoming dislodged from the printer platform early in the printing process (Fig. 5.11), but some printed successfully (Fig. 5.12).

A summary of some parts that were printed and compared with the predictions for Will It Print is given in Tab. 5.5. The risk of warping was calculated from the ratio of the area

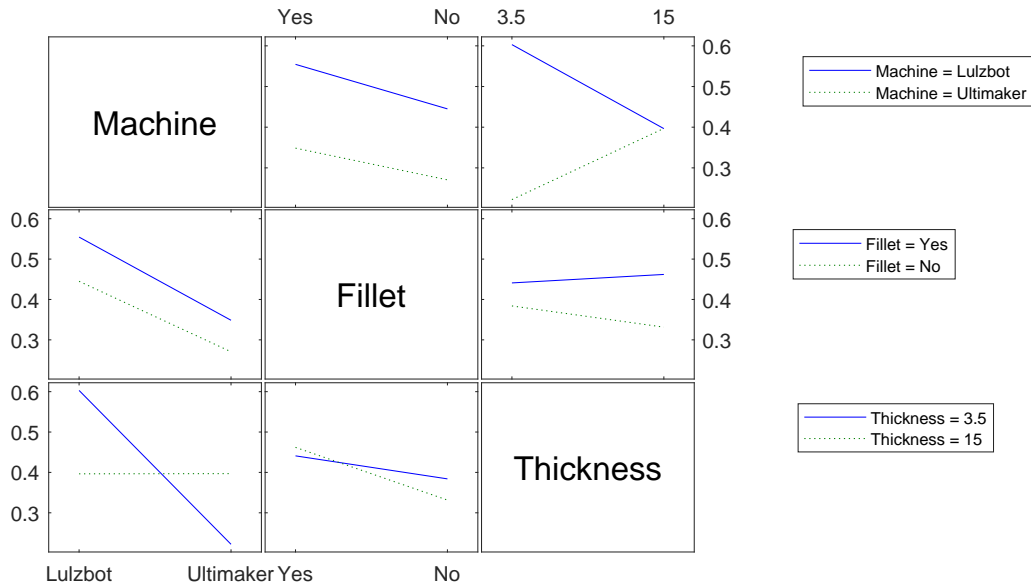


Figure 5.10: Interaction plot for the second analysis of variance, showing that many variables interact.



Figure 5.11: Three attempts to print versions of the tapered cylinder part that became dislodged from printer platform, causing the print to fail.

on the base plate, A , to the height of the center of mass, H_c . If this ratio was less than T_w , Will It Print issues a tipping warning.

Parts that had a large ratio of the area on the base plate to the height of the center of mass, $\frac{A}{H_c}$, printed successfully. Parts with small $\frac{A}{H_c}$ tended to fail to print: 3 of the 5 cylinders and tapered cylinders printed failed by toppling over and peeling off the base plate, near the very beginning of the print process. However, some parts with $\frac{A}{H_c} < 1.5$ printed successfully,



(i) Cylinder part (130 mm \times 10 mm) (ii) Tapered cylinder part (140 mm \times 10-20 mm)

Figure 5.12: Two attempts to print versions of the tapered cylinder part that printed successfully.

Table 5.5: Summary of tipping experiments.

Part name	Size [mm]	Printing failed?	$\frac{A}{H_c}$ [mm]	Tipping warning?
Cylinder part	130 \times 10	N	1.23	Y
Tapered cylinder part	140 \times 10-20	N	1.06	Y
Tapered cylinder part	140 \times 10-30	Y	1.03	Y
Tapered cylinder part	140 \times 10-30	Y	1.03	Y
Tapered cylinder part	140 \times 15-30	Y	2.37	N
Small hole benchmark part	90 \times 75 \times 13	N	17.8	N
Flagpole bracket	110 \times 110 \times 85	N	15.8	N
Cat figurine	55 \times 30 \times 30	N	34.7	N

namely the first two parts listed in Tab. 5.5. It is likely, however, that these parts might fail to print, if they were printed many times, if there was more vibration from the printer. False positives (i.e., warning the designer that the geometry is at risk of tipping even if it may print successfully) are better than false negatives. Because of this fact, it may be advisable to increase the value of T_w used for the Type A printers to 2.5. With this value, a warning would be issued for the largest tapered cylinder part that was printed, which failed to print

despite no warning being issued from Will It Print. Another possible way to improve the tipping warnings is by calculating warping risk in a different manner. This possibility will be addressed in the discussion section.

We leave as future work experimentation on other printers and with various print settings. Anecdotally, adhesion does appear to be impacted by the printer and settings used (i.e., if a brim or raft is used, and the size of that brim or raft).

5.3.3 Small features

To verify the ability of Will It Print to flag any geometry that is below a threshold of printability, we used a benchmark part developed by Seepersad et al. [121]. This part was developed to evaluate the minimum hole size a printer can successfully print. To verify Will It Print, we printed the benchmark part and compared physical observations with the geometry flagged by Will It Print as being too small to print successfully.

Will It Print flags all features that are smaller than approximately 2 mm in any dimension, which includes features that are so small that the slicer program will not detect it (and thus it will not be printed), and geometry that is large enough to print but will likely break or have a poor appearance after printing. Our decision to set up Will It Print to flag geometry smaller than 2 mm was based on an analysis of DFAM guidelines: as summarized by Booth et al. [116], it is generally advisable to have features larger than 1-2 mm for FDM.

In order to verify Will It Print's predictions, the small hole benchmark part was printed at two orientations: horizontally, with the large flat side printed on the build plate; and vertically, with the short side printed on the build plate. Photos of these two versions of the part printed are shown in Fig. 5.13.




Figure 5.13: Small hole feature benchmark part printed at different orientations.

The small feature analysis of Will It Print for the small hole benchmark part is shown in

Fig. 5.14. Will It Print indicates that the first 9 holes are acceptably large, but the remaining 11 holes are likely to be filled in (shown as blue highlighting inside the holes). The thinnest portion of the part, on the far right, is highlighted in green as being too small. From these results, we would expect the first 9 holes to print successfully, the smallest 11 holes to be printed poorly or to fail to print, and the thinnest region on the right to be very fragile or to print smaller than expected.




Figure 5.14: Will It Print analysis for the small feature benchmark part, showing small positive features in green and small negative features in blue. (See previous tables for plate thickness.)

We inspected the holes in the printed parts using the pass/fail criteria given in [121]. The results are summarized in Tab. 5.6 and Tab. 5.7 for the benchmark part printed horizontally and vertically, respectively.

Holes larger than 2 mm printed successfully for all plate thicknesses and both part orientations. For the part printed vertically, holes larger than 1 mm printed for most plate thicknesses but did not print well on the thinnest plate thickness. Holes in the range of 2-0.4 mm tended to be slightly or completely filled in, which limits their functionality and detracts from their appearance. Holes smaller than 0.4 mm failed to print altogether for both orientations and did not appear in the toolpaths planned by Cura (the slicing software used to generate the toolpaths).

Will It Print also highlighted sharp interior and exterior corners, as discussed in Chapter 3, because 90° corners have very small thickness at the extreme edge of the corner. This highlighting is shown in Fig. 5.15.

The highlighting of sharp corners matches the physical behavior of FDM machines, which cannot make perfectly sharp 90° corners due to the fixed filament size and imprecision of the positioning system. Sharp, 90° corners were observed to be slightly rounded (Fig. 5.16). The rounding of sharp corners was especially prominent for features printed parallel to the

#	Hole diameter [mm]	Plate thickness [mm]						
		12.7	9.327	7.152	5.253	3.755	1.877	0.939
1	4	Green	Green	Green	Green	Green	Green	Green
2	3.75	Green	Green	Green	Green	Green	Green	Green
3	3.5	Green	Green	Green	Green	Green	Green	Green
4	3.25	Green	Green	Green	Green	Green	Green	Green
5	3	Green	Green	Green	Green	Green	Green	Green
6	2.75	Green	Green	Green	Green	Green	Green	Green
7	2.5	Green	Green	Green	Green	Green	Green	Green
8	2.25	Green	Green	Green	Green	Green	Green	Green
9	2	Yellow	Yellow	Yellow	Green	Green	Green	Green
10	1.75	Red	Red	Red	Red	Yellow	Green	Green
11	1.5	Red	Red	Red	Red	Red	Yellow	Green
12	1.3	Red	Red	Red	Red	Red	Yellow	Green
13	1.1	Red	Red	Red	Red	Red	Red	Yellow
14	0.8	Red	Red	Red	Red	Red	Red	Yellow
15	0.6	Red	Red	Red	Red	Red	Red	Red
16	0.5	Red	Red	Red	Red	Red	Red	Red
17	0.4	Red	Red	Red	Red	Red	Red	Red
18	0.3	Red	Red	Red	Red	Red	Red	Red
19	0.25	Red	Red	Red	Red	Red	Red	Red
20	0.125	Red	Red	Red	Red	Red	Red	Red

Table 5.6: Hole success for horizontally printed small hole benchmark part where red denotes failure, yellow denotes neutral, and green denotes a pass

#	Hole diameter [mm]	Plate thickness [mm]						
		12.7	9.327	7.152	5.253	3.755	1.877	0.939
1	4	Green	Green	Green	Green	Green	Green	Green
2	3.75	Green	Green	Green	Green	Green	Green	Green
3	3.5	Green	Green	Green	Green	Green	Green	Green
4	3.25	Green	Green	Green	Green	Green	Green	Green
5	3	Green	Green	Green	Green	Green	Green	Green
6	2.75	Green	Green	Green	Green	Green	Green	Green
7	2.5	Green	Green	Green	Green	Green	Green	Green
8	2.25	Green	Green	Green	Green	Green	Green	Green
9	2	Green	Green	Green	Green	Green	Green	Green
10	1.75	Green	Green	Green	Green	Green	Green	Yellow
11	1.5	Green	Green	Green	Green	Green	Green	Red
12	1.3	Green	Green	Green	Green	Green	Green	Red
13	1.1	Green	Green	Green	Green	Green	Green	Red
14	0.8	Yellow	Yellow	Yellow	Yellow	Yellow	Yellow	Red
15	0.6	Red	Red	Red	Red	Red	Red	Red
16	0.5	Red	Red	Red	Red	Red	Red	Red
17	0.4	Red	Red	Red	Red	Red	Red	Red
18	0.3	Red	Red	Red	Red	Red	Red	Red
19	0.25	Red	Red	Red	Red	Red	Red	Red
20	0.125	Red	Red	Red	Red	Red	Red	Red

Table 5.7: Hole success for vertically printed small hole benchmark where red denotes failure, yellow denotes neutral, and green denotes a pass




Figure 5.15: Another view of the Will It Print analysis for the same small feature benchmark part, showing that sharp exterior and interior corners will be rounded. (Note that rounding can result from erosion or being filled in.)

build plate.




Figure 5.16: Rounding of sharp interior and exterior corners is visible.

The varying thicknesses of the small hole benchmark part also allowed us to assess the impact of differing levels of positive feature size. The portion of the part with the smallest thickness (0.939 mm) was very flexible and delicate for the versions of the benchmark part printed in both orientations. The flexibility was especially pronounced in the part printed horizontally. The area with the next largest thickness (1.877 mm) was somewhat flexible and delicate, especially on the part printed horizontally. The flagpole part was printed with cylinders with a thickness of 1 mm. These walls printed successfully but were somewhat delicate. Based on these observations, the thickness of positive features should be above approximately 2 mm. This value agrees with other reports that estimate the minimum recommended wall thickness [116,175].

We did not extensively study the impact of orientation on minimum feature size. However, other sources indicate that the minimum feature size for both positive and negative features depends slightly on orientation, but the effect of orientation is more evident for very small features [175]. If the designer ensures all features are above 2 mm, they need not be overly concerned about the impact of orientation on small features in their design.

5.4 Discussion - Study 1: Will It Print

In summary, we printed several parts that were deemed to be near the threshold for having poor manufacturability in order to validate the predictions of Will It Print. For warping, several very long parts that would tend to warp were printed. For tipping, several tall parts with a small base were printed. For small features, a part with very small features was printed. In general, this experimental study confirmed the appropriateness of the general thresholds that had been set in Will It Print, corresponding to values for the Type A printers. However, our experiments demonstrated that there is some variability in the printing outcomes. Additionally, we discovered some areas that should be addressed in future work.

In the experiments dealing with warping, we confirmed that extreme warping was observed on parts printed on the Type A printer with a length of above 80 mm, confirming the threshold in Will It Print was appropriate. This threshold needs to be customized on different machines. A threshold of roughly 150 mm would likely be appropriate for the Ultimaker and LulzBot printers, which we observed to generally have less warping. Adding fillets to sharp corners did not appear to reduce the amount of warpage seen on parts, despite the fact that this was recommended as a strategy to reduce warping on online forums and lists of guidelines for FDM. In general, our results indicated a need for additional experimentation on the impact of various geometric factors on warping, such as a larger range of thickness and more complicated shapes. Additionally, the printer parameters used to print the parts should be considered. We observed that the adhesion type selected (i.e., raft, brim, none) affects the amount of warping, as does the infill pattern. While using length as the primary predictor of warping seems appropriate, based on our preliminary study and other research [80], additional experimental research could lead to more nuanced predictions.

The experiments focused on tipping showed that the ratio of the area to the height of the center of mass was generally an acceptable indicator of risk of tipping. However, some unflagged geometry did tip, and some geometry that was flagged did not tip. It may be advisable to increase the tipping threshold, T_w , up to 2.5, rather than the current value of 1.5. We hope to evaluate alternative methods of estimating tipping risk in future work. For example, the method used in Meshmixer, using the convex hull of points on the build plate, holds promise. Another area for future research is to explore how geometry effects adhesion to the build plate. The shape of the contact area (i.e., the area on the base plate) can impact adhesion between two surfaces, so certain geometry, such as pointed corners, may reduce adhesion [176], which could increase the risk of tipping.

The small feature predictions matched observations of physical parts well. There ap-

peared to be an interaction between orientation and the success of small features whose size was close to the threshold. This interaction was not modeled in Will It Print. The current threshold is conservative, however, alerting designers that their part may be too small, so it including this interaction in future iterations may not be necessary.

Several additional findings apply to all manufacturability guidelines. For all guidelines, Will It Print requires thresholds of acceptability. We set these thresholds through a varied set of experiments, combined with referencing some design guidelines listed in academia and by manufacturers of 3D printers. While this approach worked for our purposes, and our experimental study confirmed that the thresholds we selected were appropriate, this experimental setting of thresholds is not highly standardized and could be difficult for others to adopt. A more standardized alternative would be designing and sharing a set of standard test parts that could be printed on any 3D printer, accompanied with a guide for assessing the quality of the printed test parts and setting the Will It Print thresholds based on that quality.

Another finding that applies to many of the manufacturability criteria is the need to consider more interactions between process parameters and the manufacturability criteria. Experienced designers can mitigate the impact of risky geometry or design choices (such as printing very large parts so that their long face is printed on the build plate, which tends to lead to warping) by selecting particular sets of process parameters (such as including a large brim). Will It Print can be enhanced by offering suggestions for process parameter settings that are customized to the particular risks identified. For example, if the designer wants to print a part in an orientation where tipping may occur, Will It Print can suggest that they use brim of a certain width to secure the part to the base plate more firmly.

These findings can help improve Will It Print in future iterations. However, we found that the current predictions of Will It Print matched physical printing results relatively well. All parts that Will It Print flagged for warping either failed to print or exhibited visible warpage at the corners. Parts that were flagged to be at risk of tipping either failed to print, or printed, but looked precarious and unstable. Small features flagged by Will It Print were either: printed filled in, had poor quality, or were very fragile and at risk of breaking. Experimental verification of Will It Print indicates that it can be used to accurately flag geometry that can have poor manufacturability.

5.5 Results - Study 2: TAAM

Study 2 focused on the accuracy of the mathematical model used in TAAM to predict achievable tolerances. By gathering experimental data on tolerances, like flatness, orientation (perpendicularity, parallelism, and angularity), and cylindricity, we hoped to understand the effect of build orientation on achievable quality. We also aimed to identify other factors that substantially impacted achievable quality.

5.5.1 Flatness

As summarized in Tab. 5.2, both the heptagonal prism and the GE bracket were used to evaluate these predictions for flatness. We evaluated the flatness on every face, except the face used to support the part during scanning. From these parts, we were able to obtain a wide range of data for faces printed at different orientations.

As discussed in Chapter 2 and shown in Eq. 3.12, TAAM predicts that flatness error on a particular face would vary with the angle, θ , between that face's normal vector and the build direction. The error is also influenced by the layer thickness, h , and precision in the xy - and z - directions, δ_{xy} , and δ_z , respectively. Based on Eq. 3.12, we would expect relatively low flatness error on faces with their normal vector oriented 0° , 90° , or 180° from the build vector. The experimental data demonstrates that the magnitudes of measured and predicted error are generally similar (Fig. 5.17). Also, the data generally follows the trend predicted by TAAM, with higher flatness error between approximately 140° and 179° , and lower flatness error at 0° , 90° , and 180° .




Figure 5.17: Measured flatness error falls close to the trend predicted by TAAM as a function of θ , the angle between the normal vector of the face with the flatness callout and the build vector, while the trend for the simple model is lower than the measured data. (For measurements from different faces oriented such that they have the same angle to the build vector, the mean flatness error is shown, along with error bars indicating the standard error of the repeated measurements.)

For further comparison of the predicted and measured errors, we plotted geometric error on the part faces of the GE bracket part printed at the orientation corresponding to $\mathbf{B} =$

[1,0,0]. For the measured error, the geometric error is represented by deviation from the nominal CAD geometry, with positive and negative deviations plotted on opposite ends of the color spectrum (Figure 5.18). For predicted error, we plot the flatness error predicted by TAAM for each face (Figure 5.19). Flatness error can only be positive, so the color bar goes from 0 to high flatness error. These two plots represent different metrics of geometric error, but we can qualitatively assess some differences and similarities.




Figure 5.18: Romer arm scan of the GE bracket printed in Orientation 2, with color indicating the total deviation from the nominal (perfect) geometry. Region 1 indicates higher flatness error to the right of angled bosses. Region 2 indicates higher error near the semi-circular edges of the part, especially on the left. Region 3 indicates low, relatively constant error across the bottom of the part. Region 3a shows two vertical areas of increased negative deviation. Region 4 indicates relatively low, constant error across all faces on the top of the part.

Higher deviations are visible near where the support structure was removed (Region 1), although this area extends farther in the measured deviations than was predicted. The back face and front, angled faces have similarly low levels of deviation, as we would predict given that their face normals are all equal to or close to 90° from the build vector (Regions 3 and 4). Larger deviations are visible around the curved corners of the parts (Region 1).

As seen in the predicted scans, we expect the two triangular faces adjacent to the holes visible in the top view to have higher flatness error than most of the other faces visible in that view, due to increased stair-step effect. However, these small differences in the stair-step effect at different orientations are not visible in the scanned data. This limitation will be analyzed in the discussion section.

A feature of note in the scan data is Region 3a, made up of two parallel blue regions (smaller than nominal geometry) visible on the backside of the part (Fig. 5.18). These two small linear indents, parallel to the build platform, are also visible based on a visual inspection of the part itself, as seen in Fig. 5.20. These indents appear near the layers where the printing of the angled features on the part began.




Figure 5.19: Predicted error of the GE bracket printed in Orientation 2, with color indicating flatness error in millimeters. Region 1 indicates higher flatness error to the right of angled bosses. Region 2 indicates higher error near the semi-circular edges of the part, especially on the left. Region 3 indicates low, relatively constant error across the bottom of the part. Region 4 indicates relatively low, constant error across all faces on the top of the part.



Figure 5.20: GE bracket printed at Orientation 2. Support material is visible under the holes and on the backside of the part, supporting the angled features.

Another interesting observation was the occurrence of unexpectedly high flatness error on faces of parts that were printed on the build platform. Figure 5.21 shows an example of this phenomenon, comparing the scans of the bottom face of the GE bracket when that face was printed on the build platform and perpendicular to the build platform. The deformation visible in the left scan is symmetrical and longer in the long dimension of the part. Deformation is highest at the corners on the left scan. In the right scan, deformation is low in magnitude and mostly constant across the entire face. This observation, which we attribute

to warping, will be analyzed in the discussion section.




Figure 5.21: When printed with the large base on the build platform, significant warping was visible (left), but when printed on its short side, little warping was visible (right). Color indicates deviation from nominal geometry.

5.5.2 Cylindricity

Our evaluation of model performance on predicting cylindricity errors is based on measurements taken with two parts: the flagpole bracket used in Chapter 4 and the angled cylinders part. The flagpole bracket was printed in one of the optimized directions (Orientation 1) suggested in Chapter 4 as well as in its initial orientation (Orientation 2) on a Stratasys printer. The parts printed in those orientations can be seen in Fig. 5.4. The angled cylinders part was printed in one orientation using the Zortrax printer. Its many cylindrical features enabled several measurements. Cylindricity measurements were made on the outside cylindrical faces for each part and the results of those measurements are plotted in Fig. 5.22, along with the predictions of TAAM and the simple model.

Figure 5.22 shows the impact of layer thickness, with the prints that used a larger layer thickness showing higher cylindricity error. The measured data falls close to the trend predicted by TAAM, while the simple model predicts much lower error than what was observed.

Visual observation of the flagpole parts confirmed the measured results. There is a slight improvement in the quality of the cylinders in the first, optimized orientation due to the reduction in the stair-step effect, compared with the second, naive orientation. However, other, larger flaws and bulges are visible on the cylindrical surfaces of both orientations. These shifts, which are visible as discontinuities in the surface, are visible just above the colored lines in Fig. 5.23. These bulges occur when the shape of the contour printed changes significantly between layers.




Figure 5.22: Measured cylindricity error falls close to predicted trend for both Dataset 1 (Zortrax, $h=0.2$ mm) and Dataset 2 (Stratasys, $h=0.254$ mm). For measurements from different cylinders oriented such that they have the same angle to the build vector (θ_{ax}), the mean cylindricity error is shown, along with error bars indicating the standard error of the repeated measurements.

5.5.3 Orientation (angularity, parallelism, perpendicularity)

Using different faces as datum features (see Appendix A for details), many repeated angularity, perpendicularity, and parallelism measurements were taken of the heptagonal prisms. Figure 5.5 shows the prisms, printed at four different orientations (each prism’s orientation is rotated 30° from the previous). Additionally, a perpendicularity and angularity measurement were taken from the GE brackets at each of the three printed build orientations.

The data from the prisms and the brackets is shown together in Fig. 5.24. The predicted trend for orientation tolerance data with respect to the angle between a face’s normal and the build vector, which is identical to that of flatness (neglecting the effect of small feature faces), is also plotted for comparison. The experimental error is significantly larger than the error predicted by both TAAM and the simple model. Measurements of orientation tolerance error for any given θ also tend to vary widely, depending on which datum they are referenced to, which can be seen in the large height of the error bars for almost every data point.




Figure 5.23: Changes in the contour shape at slices 1 and 2 lead to distinct bulges on the printed part (slice contours were created using Slic3r).

5.6 Discussion - Study 2: TAAM

This discussion begins with an analysis of the effect of the independent variables that were examined in Study 2, followed by a discussion of the impact of extraneous variables, which were not explicitly examined in Study 2, but appeared to impact the results. Finally, we will summarize and interpret the overall significance of the results of Study 2.

5.6.1 Effect of independent variables

Part orientation

Part orientation has been shown to have a clear effect on surface roughness of parts printed using AM and FDM specifically, so we expected orientation to similarly affect geometric tolerance errors like flatness, orientation, and cylindricity. We did see some effect, as reflected in Figs. 5.17, 5.24, and 5.22. However, our measurements did not have the necessary resolution to distinguish between small changes in geometric error on the scale of surface roughness.

Support material

For the Zortrax, Type A, Ultimaker, and LulzBot prints, removal of the support material was performed manually with a utility knife. Manual removal of support material was found to leave visible scratches and gouges on the surface of the part. Also, laser scan data tended




Figure 5.24: Measured orientation error was highly variable and did not fall close to the trend predicted by TAAM or the simple model as a function of θ , the angle between the normal vector of the face with the flatness callout and the build vector. (For measurements from different faces oriented such that they have the same angle to the build vector, the mean flatness error is shown, along with error bars indicating the standard error of the repeated measurements.)

to show higher geometric error in areas where support material was removed. In addition to higher error on the faces that were being supported, we also observed higher error on the adjacent faces, where the support material was touching. For dissolvable supports (like those used in the Stratasys Dimension), this additional error was less prominent. The support term, δ_s , should be experimentally determined for each printer used.

Layer shape and deviation

The primary differences between the simple model and the model used in TAAM are the inclusion of the effect of support material, layer-level deviations, and the shape of each layer in the mathematical model used in TAAM. The support material term was found to better match the observed trends in the experimental data, as was discussed in the previous paragraph. Similarly, the impact of layer shape and layer-level deviations, which increased the error predicted by TAAM for $\theta < \theta_{sup}$, were found to better match experimental data. Inclusion of these factors increased the accuracy of TAAM.

5.6.2 Effect of extraneous variables

Warping

In several of the parts used for our experimental verification, the parts showed signs of large scale deformation due to warping. In the Type A printer, this typically manifested itself in failed builds, where the part would become dislodged from the build platform, and the print would fail. However, even in the more reliable printers used, warping was visible on faces of parts that were built directly on the build platform.

For large faces, the deformation of warping was large enough that it impacted the tolerance measurements. In the case of the GE bracket shown in Fig. 5.21, the flatness of the large face printed on the build platform was 0.38 mm, but when a smaller face of the same geometry was printed on the build platform, its flatness was 0.119 mm. It appears that warping, which was not included in the TAAM predictive model, can have a significant impact on the achievable flatness of a face.

Datum quality

The TAAM predictive model assumes that datum planes are established on the protruding edges of layers and that the edges are positioned such that the datum plane is established parallel to, although possibly offset from, the nominal, perfect datum feature. Based on this assumption, we would not expect the datum feature itself to cause errors with a tolerance that is defined relative to that datum feature. However, our experiments showed that this was not a valid assumption. The quality of the datum feature that was used to establish a datum plane did influence the measured error associated with tolerances defined relative to that datum. Figure 5.25 shows that the orientation error tends to increase as the flatness error of the datum face increases, making it difficult to see a clear relationship between orientation error and θ .

Analyzing data from Yang et al. [21], who printed and inspected parts printed using SLS (a Sinterstation 2500+ machine), there are several similar trends regarding datum quality. For any given face, the flatness on that face is less than the orientation callouts. Also, datums with lower straightness errors correspond to lower errors for orientation tolerances referenced to those faces. However, only a few orientations were actually built in the Yang et al. study [21], so more study is needed to confirm these results for SLS.

If the experimental data is restricted to orientation callouts defined relative to relatively smooth datum features (flatness error less than 0.35 mm), the data does follow the trend predicted by TAAM, but the magnitude is generally lower than predicted. If we increase each TAAM prediction by a certain constant amount, which we will denote δ_{orient} and set equal to 0.2 mm, the measured data closely aligns with the new TAAM predictions (Fig. 5.26).

Simply adding an experimentally determined value, like we do here with δ_{orient} , is a simple solution to account for the difference we observed between measured and predicted orientation error. However, this does not account for the relationship we observed between




Figure 5.25: Measured orientation error versus θ , the angle between the normal vector of the face with the orientation callout and the build vector, with the marker size indicating the flatness error on the datum face. Orientation error tends to be very high for faces referenced to datums with high flatness error.

orientation error and flatness error on the datum face. More research is needed to analyze if our approach with δ_{orient} is appropriate or if another method would be superior.

Change in contour

In several different printers, we noticed that when there was an abrupt change in the contour being traced by the print head, there would be a visible shift in the deposited filament before and after the layer occurred. This phenomenon was remarked upon the discussion of the achievable tolerances of the flagpole, bracket earlier. Defects caused by changes in contour is also visible in the two circumferential indents on the backside of the GE bracket in Fig. 5.20.

From observing these and other parts, it appears that the flaws are caused by abrupt changes between the contour or toolpath of two subsequent layers. We observed this flaw on several machines, namely the Stratasys Dimension, the Zortrax M200, and the Type A printers. The effect of abrupt changes in contour on the geometric error of a part was not included in the model, but it appears to have a potentially significant effect on the amount of geometric error observed. Although abrupt changes in contour may not cause additional geometric error in all printers, its effect should be included in models like TAAM for printers




Figure 5.26: When limited to callouts with relatively smooth datums, measured orientation error follows the trend predicted by the modified TAAM equation, while the simple model predicts much lower error. (For measurements from different faces oriented such that they have the same angle to the build vector, the mean flatness error is shown, along with error bars indicating the standard error of the repeated measurements.)

that do demonstrate this phenomenon. The impact of wall thickness on this phenomenon should also be evaluated.

5.6.3 Overall evaluation of TAAM performance

In summary, the predicted error of the TAAM tool was similar in magnitude to the measured error. Because the magnitude of the predicted error matches the measured error, it appears TAAM is accurately capturing some sources of error (namely stair-step error, machine precision, and the profile shape of each layer). However, it was difficult to observe the effect of these small, (<0.2 mm scale) deviations in the laser scan data. Additional factors that were not included in the model were observed to contribute to errors in the printed parts, including warping, datum quality, and abrupt changes in contour between layers.

It was difficult to distinguish between faces with varying levels of stair-step effect due to different orientations. This is likely caused by the way the laser system works — its sampling of points may be such that it is difficult to scan each small layers, or perhaps

the layer deviations are smoothed over when the point cloud is processed and optimized by the software. By using a metrology system with higher resolution, or by printing parts with larger layer thickness, we may be able to measure the dependence of error on face orientation with greater precision.

After future experimental research using improved measurement techniques, we should understand better the impacts of the extraneous factors we do not currently include in the TAAM mathematical model. The mathematical model should be updated accordingly. For example, the prediction of orientation errors, governed by Eq. 5.1, does not account for the dependence between datum face quality and orientation error. The approach we implemented here, adding a new term, δ_{orien} , should be analyzed to ensure it is appropriate. The impact of warping and changes in contours should also be evaluated and possibly be included in the mathematical model.

As AM technology improves, AM machine manufacturers will no doubt improve their designs to eliminate some of the sources of geometric error. This process has already begun, with lower layer thicknesses and improved machine precision. By basing TAAM's mathematical model on first principles, it should be relatively easy to customize it for different machines. Additionally, TAAM can aid AM machine manufacturers in their development efforts for new machines: by understanding the different sources of error and their relative impact, manufacturers can address the most significant sources of error in new technology. Based on our results, the stair-step error caused by large layers, support material scarring, warping, and bulging caused by changes in contour between layers should all be addressed.

In its current form, TAAM successfully predicts approximate flatness, cylindricity, and orientation errors at many orientations and for different layer thicknesses. The accuracy of TAAM is significantly higher than the simple model proposed in previous research [93, 138]. The predictions of TAAM can help designers estimate the level of geometric accuracy that is possible on different part features, which is helpful during the tolerance allocation process.

5.7 Conclusion

This chapter focused on experimentally verifying the predictions of the two tools developed in this thesis: Will It Print and TAAM. Will It Print was successfully able to issue warnings that correlated to experimental results, regarding failed and poor-quality prints due to excessive warping, parts tipping during printing, and small features. Similarly, TAAM predicted the amount of error associated with geometric tolerances with reasonable accuracy, on a variety of part features and several different FDM printers. These tools can help designers estimate the manufacturing outcomes for their parts, exploring the impact of build orientation, before printing their part.

Experimental testing of mathematical models for geometric error has typically been omitted in prior work. Through the experimental testing we undertook, we identified phenomena and sources geometric error that were not predicted in prior mathematical models. Because of the dearth of prior experimental testing of mathematical models, researchers have failed

to identify these gaps between theory and practice up until now. This chapter contributes to our understanding of the impact of different sources of error and helps direct the focus of future work on improving our predictions of AM outcomes.

Future work should focus on improving our predictive abilities of both tools through more investigation of factors that the tools do not currently consider. Another avenue for future research is estimating the uncertainty of the predictions, which could enable us to adapt the tools to give designers a sense of how reliable the predictions are or are not. Perhaps by adopting a more probabilistic approach, with a range of expected values for errors, or a range of expected outcomes for manufacturability, designers can feel more confident in the tools' predictions.

Chapter 6

Verification of utility of geometry analysis tools

Previous chapters of this thesis detailed the practical aspects of implementing geometry analysis in a design for additive manufacturing (DFAM) software system, and the current level of accuracy associated with our system. This chapter will detail efforts to examine the usability, or ease of use, of our DFAM software.

6.1 Introduction

Because we want our software to be useful and easy to use for designers, it is important to evaluate the utility of our tool and the experience of users who interact with it. In order to understand utility and user experience, a semi-controlled experiment with human subjects was conducted, which is detailed in this chapter. Due to practical constraints such as access to industrial designers, we choose to evaluate only the Will It Print tool. However, findings from this study can be useful for DFM tools in general.

Similar to traditional manufacturing processes like milling and injection molding, there are certain limitations of AM processes, such as size, minimum resolution, and the necessity for supports on overhanging features. As summarized in Chapter 2, these guidelines are frequently shared with designers in the form of lists or tables with small example pictures showing acceptable or unacceptable geometry, requiring designers to memorize the lists or to refer back to them frequently to ensure their design does not violate any guidelines. Another option for sharing these guidelines, which has not been widely implemented in academia or industry, is to use geometric analysis to automatically analyze a part's geometry to determine if the guidelines are being followed. Geometric analysis has several potential advantages. It is likely that geometry analysis can more accurately flag problems on the geometry, because it does not rely on humans to perform calculations or recognize certain geometric features. Also, it can likely offset some of the mental effort required of the designer by performing route calculations, freeing the designer to focus on more creative or analytical parts of the problem.

Our user testing was planned to evaluate the effectiveness of such a geometric analysis software tool in helping designers understand geometric constraints and limitations associated with additive manufacturing, as compared with a worksheet that lists design guidelines. Our hypotheses are as follows:

H1: The software-based manufacturability feedback system is more effective at improving designs than the worksheet.

H2: The software system requires less mental effort to use than the worksheet.

H3: Designers' confidence in manufacturability is improved more with the software system than with the worksheet.

H4: Designers spend less time interacting with the worksheet than with the software tool.

6.2 Study method

Data was collected from students enrolled in the fall 2018 semester of Introduction to Manufacturing and Tolerancing, a sophomore-level course at the University of California, Berkeley. During the semester, all students used the worksheet and software tool as part of two different homework assignments. This chapter describes the homework assignment that occurred first, eight weeks into the semester.

The study task was to complete the homework assignment using either the software or worksheet tool (details of both tools, as well as a description of the study task, will be described in subsequent sections). Briefly, the study task was assess a pencil holder with an engraved logo based on DFAM guidelines, to redesign the pencil holder to improve its compliance with DFAM guidelines, and to choose a build orientation for the part. The participants were asked to use either the software tool or the worksheet tool while completing the study task described in this chapter.

Assignment into experimental groups took the form of stratified randomization. Our goal was to assign half of the students to the software tool (experimental) group and half to the worksheet tool (control) group for the study, randomizing to control for the confounding effects of which sections students enrolled in (i.e., the tendency of certain students to enroll in a lab section meeting at 9 AM versus 2 PM). There were four lab sections. In each lab section, students were randomly assigned to teams of four to six students. (They were allowed to adjust the team membership, though few did.) Team assignments were made for a class project not relevant to this study (students completed the study activity individually), but for convenience, we assigned all members of a team to the same treatment group. Each team was collectively assigned to the control or experimental group, with the goals of ensuring a roughly equal number of total students and female students in the control and experimental groups, and a roughly equal number of students in each lab section assigned to the control and experimental groups.

Completion of the study task took place outside of class with no fixed time limit. Although the participants for this experiment were all engineering students, we hypothesize that our findings would generally apply to designers, especially novice designers, and the

Maker population in general.

6.2.1 Study participants

Out of 83 students in the class, 58 (10 female and 48 male) consented to participate in this study. We will refer to these students as the study participants. Before the class module about additive manufacturing began, participants were asked to respond to the question “What is your personal experience with additive manufacturing (colloquially, 3D printing)?”, rating their experience as 1 (“I have never heard of it”), 2 (“I am familiar with it, but haven’t 3D printed anything myself”), 3 (“I have 3D printed one or two things”), 4 (“I have 3D printed many things”), or 5 (“I own a 3D printer”). The average experience level was 2.84.

6.2.2 Description of software and worksheet manufacturability tools

In order to test these hypotheses, two versions of a testbed DFAM tool were developed: a software-based and worksheet-based version. Specifically, we chose to evaluate the DFAM manufacturability tool described in Section 3.3, although many of the findings described in this current chapter can be generalized to other DFAM or DFM systems. The development of this tool and examples illustrating its use can be found in Chapter 3 and Chapter 4.

Some of the issues that we chose to analyze, such as warping and support material mar-
ring the surface of the print, are not just a function of the part geometry. These issues are also affected by printing parameters and specifications of the printer used. However, the vast majority of students enrolled in the class in which the study was conducted used Type A Series 1 Pro FDM printers for their projects. In our university makerspace, students are generally instructed to use the makerspaces’ predetermined printing parameters. Therefore, this chapter assumes that process parameter changes will not mitigate geometry/orientation issues and that the designer must make geometry and orientation changes to ensure the best manufacturing outcomes. Our tools also do not include the effect of orientation on total cost, as students do not pay for filament or print time on the Type A printers in our university makerspace.

The workflow of the worksheet is straightforward and similar to that of someone using the “Design for Additive Manufacturing Worksheet” [116]. The worksheet was a two-page PDF with generic pictorial examples of acceptable and potentially problematic geometry, and suggestions for redesigning a part to avoid those problems. An example of a description of the guideline for warping shown is shown in Fig. 6.1. A designer would read the worksheet, and mentally assess their part geometry to determine if it was in keeping with the guidelines.

The main difference between the worksheet tool and the software tool was that the worksheet tool explained DFAM guidelines and required designers to analyze their part themselves to determine how and if their geometry violated any of the described guidelines,




Figure 6.1: Sample of guideline description used in both versions of tool.

whereas the software performed geometry analysis specific to the uploaded part geometry and highlighted areas of the geometry with each manufacturability issue for the participant. If a designer changed the part geometry, they could re-upload the geometry into the tool to assess if the problems were resolved, whereas a designer would have to rely on their own analysis alone to evaluate redesigns using the worksheet. Similarly, if the orientation was changed, the analysis and plotting were automatically updated in the software tool, but a designer using the worksheet would have to mentally assess if the new orientation resulted in any guideline violations. As mentioned earlier, the presentation of the design guidelines was nearly identical in both versions, in keeping with best practices from game-based learning [177].

6.2.3 Design problem

The study task was designed to test the study hypotheses. Participants were asked to evaluate an angled pencil holder with an engraved logo, like that shown in Fig. 6.2, for its manufacturability using a hobbyist FDM printer. The pencil holder was a remix of a design found on Thingiverse [178]. Participants were asked to analyze the part using either the software or worksheet version of the tool (depending on whether they were assigned into the experimental or control group). Participants were instructed to use the tool to identify any problems that could interfere with the successful printing of the part. Then, they were asked to redesign the pencil holder, with geometry changes to mitigate the DFAM guideline violations of the original design. Participants were given a few design restrictions (they were told that the logo must remain the same size on the front of the pencil holder and that the front face and logo must be relatively smooth for aesthetic reasons) and general size constraints. Participants were also asked to choose a build orientation for their redesigned part. Students submitted a visual representation of their redesign (either CAD or engineering drawings were permitted) and a written description and justification of all redesigns they




Figure 6.2: Study participants were asked to evaluate the manufacturability of a pencil holder with an engraved logo, shown here with the class number.

had made. The full assignment description is included in Appendix B.

In order to ensure that the redesign problem was challenging and to avoid having a single obvious orientation that would eliminate all potential manufacturability issues, the part was designed such that there were several conflicting problems with the geometry. One of the problems was relatively straightforward though potentially difficult to notice: there were two very thin walls that would be prone to breaking. The other problems largely depended on orientation. The geometry featured many long, thin faces with sharp corners that could be prone to warping if printed on the build plate. These faces were not parallel, and so it was difficult to find a single orientation that simultaneously resulted in good surface finish on all faces that also avoided requiring support material in internal cavities, which would be difficult to remove.

Participants using the software version of the tool were given instructions for how to open and load the part STL file in the tool. All participants were provided with a dimensioned engineering drawing of the part. Both tools were designed to be self-explanatory, with built-in instructions guiding the user through the use of the tool. No training was provided to either group. We wanted to see which tool could be quickly learned and put into use by novices.

6.2.4 Study measurements

One of the main goals of this study was to determine if the automated geometry analysis provided by the software tool made participants more effective at completing the objectives of the design task than participants who only used the worksheet, which could not provide any geometry analysis. To determine effectiveness, we analyzed the redesigns submitted by each participant and counted the number of features on each participant's redesigned part that violated either a design guideline given in the problem statement or a DFAM guideline that was described by the worksheet and software tools. This metric is referred to as the number of design or manufacturing guideline violations. The redesigns were analyzed by a single evaluator (the author of this dissertation). The evaluator was aware of which group

each participant was in during evaluation, which may have introduced some biases.

In addition to understanding designer effectiveness, we sought to understand how difficult the task was for the participants in both the experimental and control group. To evaluate task load while redesigning the part, we asked all participants to rate the difficulty of the task according to the NASA Task Load Index (NASA-TLX), with indices such as mental demand and frustration [179]. We report each of the individual NASA-RTLX components and an unweighted average of the NASA-RTLX component ratings. Participants also responded to questions about the ease of use of the tool they were assigned to use, based on the System Usability Scale (SUS) [180]. SUS is a flexible assessment scale that can be used to evaluate the usability of a wide variety of tools and systems. Participants were also asked to record and report how much time the study task took them to complete, and to report how confident they felt that their redesigned part would print successfully. Confidence was reported on a 5-point scale, ranging from 1 (“Not at all confident”) to 5 (“Very Confident”). All reporting was done online after they completed the study task.

6.2.5 Statistical tests

Two-sample t -tests were used to determine differences between the software and worksheet groups. The effect size of the differences between groups is quantified using Cohen’s d . Spearman’s rank correlation coefficient, ρ , was used to quantify correlations between study measurements. For 2x2 contingency tables, a two-tailed Fisher’s exact test was used. A significance level of .05 was used throughout this study. All statistical tests were implemented in MATLAB.

6.3 Results

From each of the 58 participants, we collected the following: survey data that included responses on the SUS and NASA-TLX, geometry of their redesign, and a written explanation of all redesigns they had made.

Several participants were excluded from further analysis after their redesigns were analyzed. One participant from the software group and one participant from the worksheet group did not complete the design task. Six participants (all male) who were assigned to use the software tool were found have completed the design problem without consulting the software tool. Participants were judged to have not consulted the tool if they made no reference to any of the guidelines described by the tool and included no geometry changes related to any of the guidelines. One participant (female) who was assigned to use the worksheet was also excluded from the study for the same reason. The excluded participants who completed the design task reported completing the task in an average of 56 minutes, which is substantially less than the average of students who completed the task per the instructions (118 minutes).

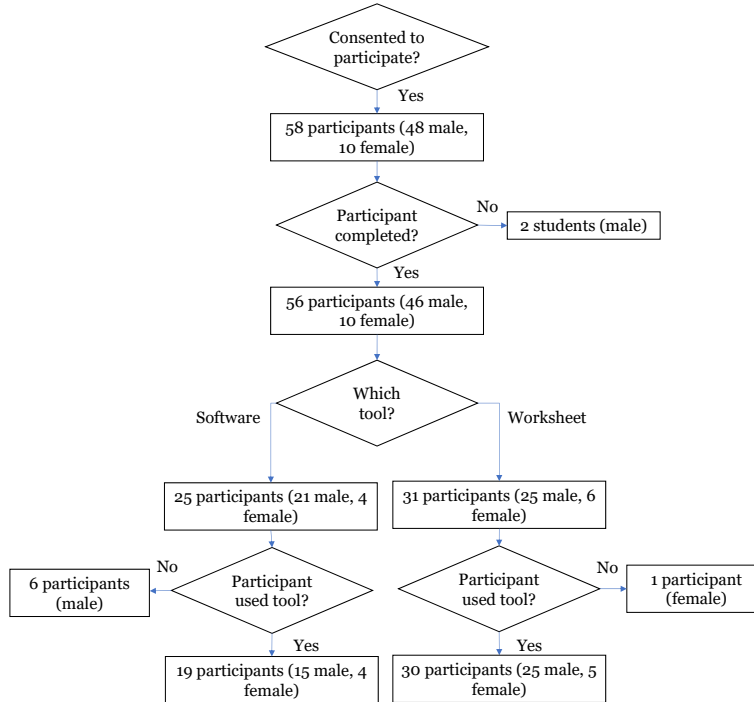


Figure 6.3: Summary of process for selecting study participants.

After these participants were excluded, our samples consisted of 4 female and 15 male participants who used the software tool, and 5 female and 25 male participants who used the worksheet tool. In total, data from 9 female and 40 male participants were included in this study. This breakdown of participants is summarized in Fig. 6.3.

6.3.1 Evaluating study hypotheses

A summary of the study measurements is presented in Tab. 6.1, along with the effect size of the difference and the p -value calculated from a two-sample t -test. Each hypothesis presented in the introduction will be analyzed separately here. In the subsequent sections, we will elaborate on these results and present additional findings, outside of these hypotheses.

H1: The hypothesis that the software manufacturability tool is more effective at improving designs was supported. Participants who used the software tool came up with redesigns that contained fewer design or manufacturing guideline violations.

H2: The hypothesis that the software tool requires less mental effort to use was partially supported. Based on the mental demand question of the NASA-TLX, software participants reported lower mental demand. The effect size for this result was 0.44 (medium to large effect), but in part because of our relatively small sample size, the difference was not statistically significant at a .05 significance level. Software participants also reported lower task

Table 6.1: Summary of study measurements.

Variable	Metric	Average score (SD)		Effect size	<i>p</i> -value
		Worksheet	Software		
Workload	Mental demand (0-100)	57.0 (19.3)	47.9 (22.6)	0.43	.14
	Physical demand (0-100)	34.7 (23.1)	28.7 (27.3)	0.24	.42
	Temporal demand (0-100)	47.7 (24.6)	49.7 (18.3)	0.10	.75
	Performance (0-100)	37.3 (28.7)	33.9 (22.9)	0.13	.67
	Effort (0-100)	68.7 (19.9)	65.3 (21.1)	0.17	.57
	Frustration (0-100)	51.2 (26.0)	56.3 (26.4)	0.20	.51
	NASA-TLX average (0-100)	49.4 (15.5)	47.0 (15.1)	0.16	.59
Confidence	5-point scale	4.2 (0.6)	4.1 (0.7)	0.19	.52
Usability	SUS scale (0-100)	68.3 (19.5)	61.3 (20.4)	0.35	.23
Performance	# of guideline violations	1.9 (1.1)	1.1 (0.7)	0.92	.004
	Time taken	119.7 (108.4)	116.1 (61.2)	0.04	.89

load associated with performance, effort, and physical demand, but reported slightly higher task load associated with temporal demand and frustration than worksheet participants.

H3: The hypothesis that designers' confidence in the manufacturability of their part is higher with the software tool than with the worksheet was not supported. In fact, participants who used the worksheet reported higher confidence in their designs, though this difference was small in magnitude and not statistically significant. It appears that the type of tool used did not strongly influence the participants' confidence.

H4: The hypothesis that designers spend less time interacting with the worksheet than with the software system was not supported. Participants who used the worksheet spent similar amounts of time to those who used the software, although the variation in the time of worksheet participants was much larger. For both groups, there was significant variation between participants in the amount of time spent on the design task, ranging from 27 minutes to 10 hours. Because the completion time was self-reported, it is not a perfect measure of the actual time spent. However, it was clear from reviewing the participants' redesigns and descriptions that the dedication of participants varied. Some participants took the assignment extremely seriously and spent hours optimizing their design, while others were more cursory, settling for the first combination of orientation and design changes that satisfied most of the guidelines. This variation makes it difficult to distinguish the effect that the type of tool may have on completion time.




Figure 6.4: A typical design and orientation for worksheet participants featuring orthogonal walls, and using the same, default orientation shown in the drawings and CAD model given to participants.

6.3.2 Differences in designer effectiveness

As summarized in the discussion of H1, our results indicate that the software tool was more effective at helping participants identify design/manufacturing violations than the worksheet. This is shown graphically in Fig. 6.6. Beyond a simple count of the number of violations, we identified several interesting trends in different participants' redesigns.

Common design changes made by participants using the worksheet were to increase the thickness of all small features and to change all sloped feature to be perpendicular to the build plate (an example of such changes is shown in Fig. 6.4). We hypothesize that this is due, in part, to the fact that the dimensions and angles of these features were not directly given on the engineering drawing or in the CAD file provided to participants. Participants using the worksheet would have to use CAD measuring tools or calculate thickness and angles from the other given dimensions. It seems that many participants decided to make changes conservatively, rather than trying to calculate any non-obvious dimensions.

There may be some effect on creativity and design space exploration with the software tool. Of the 30 participants who used the worksheet, 24 suggested the part be printed in its default orientation, seen in Fig. 6.2, versus 8 of 19 software participants. This result approaches statistical significance ($p=.12$ from a two-tailed Fisher's exact test). There was more variety in the geometry changes and orientations suggested by software participants (an example of which is shown in Fig. 6.5). This finding may be related to design fixation, a connection that will be explored in the discussion section.

Two software participants and six worksheet participants reported importing the original part or a redesigned part into Cura, a widely available 3D printing software package, in order to check for overhanging features. It is interesting that these two software participants chose to use a separate package to check for overhanging features even though the software




Figure 6.5: Some software participants came up with more creative combinations of build orientation and geometry changes that allowed them to preserve more of the original shape, an example of which is shown here.

tool they were asked to use had this same capability. This preference for Cura probably is an indication that the usability and visualization used by the tool need improvement. However, the frequent usage of Cura further highlights the desire for geometry analysis tools. Participants seem to prefer the automatic calculation of overhanging angles, rather than calculating anything by hand.

The software tool would have flagged issues on participants' redesigns, but most participants (16 out of 19 software participants) did not report re-uploading their parts into the tool. Some participants may have not re-uploaded because they were uncomfortable using SolidWorks to edit the part. However, five participants included revised CAD files, indicating they went through the trouble of editing the SolidWorks file, but then did not report re-uploading the file into the software tool, so lack of familiarity with SolidWorks cannot be the sole explanation. Another explanation is motivation. Some participants seemed to settle for a certain level of performance (i.e., how well their chosen orientation optimizes manufacturability, and the corresponding grade earned on the design task) that they deem adequate while avoiding tasks they considered unnecessary. A similar finding was reported in [116], where few students consulted manufacturability guidelines before printing unless they were required to do so. The observation of motivation modifying performance may be unique to academic settings, but it seems likely that different designers in industry have differing levels of persistence in improving their design. Although motivation and self-efficacy (an individual's judgment of their ability to accomplish a task) regarding the engineering design process tend to increase with designer experience, there is variance in the level of self-efficacy and motivation reported by practicing engineers [181]. Additionally, this finding highlights the need to lower the effort level required to make changes to the part geometry and have the tool reanalyze the part, possibly by making this tool an add-in to a preexisting CAD tool.

From the three participants who changed their part in SolidWorks and reported re-uploading it into the tool, one participant had no violations, and two participants had one

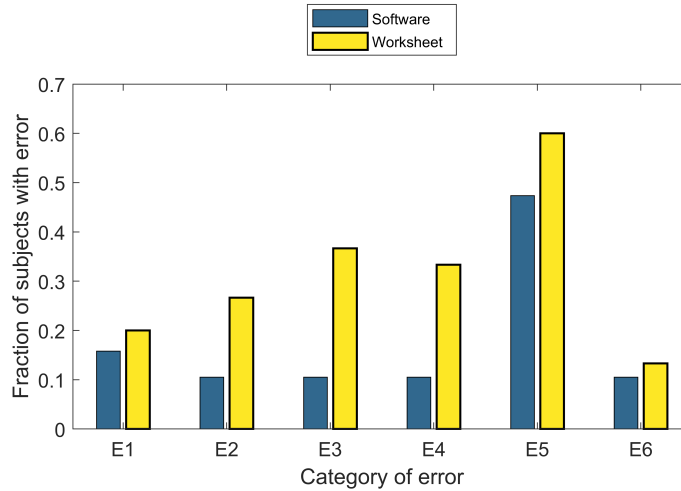


Figure 6.6: A smaller percentage of software participants had designs with design or manufacturing guideline violations. The types of DFAM violations are coded as: E1–Very long, thin face on build plate, which is likely to warp; E2–Thin wall behind logo; E3–Thin wall between pencil compartments; E4–Thin wall at base; E5–Support material that is likely to mar the surface is needed inside logo; E6–Support material that will be difficult to remove is needed inside compartments.

violation each (both of their redesigns required support material inside the logo). Neither of these two participants appeared to notice their design violation. This lapse is likely due in part to the poor visibility of the highlighting of the faces inside the logo. This issue will be discussed in a subsequent section.

6.3.3 Effect of prior experience

Not surprisingly, experience and task load are negatively correlated. Analyzing the total sample (both treatments), experience and task load had a negative correlation of $\rho=-.40$, which is significant, $p=.005$. Although participants of the software tool reported slightly lower mental load and average task load than those who used the worksheet (as summarized in the discussion of Hypothesis 2), participants who used both tools and who were experienced with 3D printing tended to report lower task load than novices. This trend can be seen graphically in Fig. 6.7.

After completing the study task, in a subsequent homework, the participants were asked to use the other version of the tool to complete a similar redesign problem and then were asked which version of the tool they preferred. The software tool was more preferred by non-experienced participants who had not used a 3D printer before beginning this class than experienced participants (64% versus 58%, respectively).

Although there was a significant variation in the amount of time spent on the design problem, more experienced participants tended to spend slightly less time. There is a slight




Figure 6.7: Experienced participants tended to have lower workload using both tools. Experience levels range from low (1) to high (5), and are described in Sec. 6.2.1.

negative correlation between time and experience for both groups combined ($\rho=-.20$, $p=.17$), which can be seen in Fig. 6.8. Although all participants were asked to use the tool to help them to analyze the part and to decided what changes to make, some more experienced participants described only cursory use of either the software or worksheet tool. Most of their changes seemed to be based on intuition and experience with prior prints. Although most implemented a few of the tool’s suggestions, they spent less time interacting with the tool and seemed to trust their intuition more than the tool. This lack of trust in the tool also led, in a few cases, to poor performance because they did not carefully look at the tool suggestions.

The level of prior AM experience was not strongly correlated to confidence in the success of the print ($\rho=.14$, $p=.33$) or the number of design violations made ($\rho=-.10$, $p=.48$). It appears that DFAM tool usage may help equalize the performance and confidence of novices, relative to experienced designers.

Interestingly, confidence was positively correlated with the number of design violations, ($\rho=.29$, $p=.04$), meaning that participants who were very confident the part would print well tended to have more design violations. This is seen in Fig. 6.9. (The notations n_w and n_s indicate sample sizes for worksheet participants and software participants, respectively.) The trend of high confidence with a large number of guideline violations is more prevalent for worksheet participants than software participants, indicating that the software tool may be better at preventing overconfidence in the success of a print. There are several possible explanations for this trend. It is possible that as participants discovered and corrected more violations, it made them worried that even more violations existed, lowering their confidence. This theory may also explain the lower confidence of software participants compared with worksheet participants, as summarized in Tab. 6.1 — they were more aware of the possible




Figure 6.8: Experienced participants tended to use slightly less time to complete the assignment. An outlier who reported using 600 minutes (a worksheet participant with experience level 3) is not included.

errors on the part and therefore less confident that the part would print successfully.

There are several possible explanations of this trend. It is possible that as participants discovered and corrected more violations, it made them worried that even more violations existed, lowering their confidence. This may explain the lower confidence of software participants compared with worksheet participants, as summarized in Tab. 6.1 — they were more aware of the possible errors on the part and therefore less confident that the part would print successfully.

6.3.4 Participant comments

The feedback on the tools was generally positive. 61% of all participants preferred the software tool after using both tools. A chi-square test of goodness-of-fit was performed, but the null hypothesis that the tool versions were equally preferred was not rejected, $X^2(1, N=41)=1.56, p=.21$.

The participants were given the opportunity to provide comments about both tools. Participants who provided comments tended to be those who had used a 3D printer before this class ($p=.08$ from a two-tailed Fisher’s exact test). Several comments included suggestions or criticisms on the visualizations used in the software tool. Others critiqued the content of the tools, like Participant 22 (experience level 4) who commented, “Personally, these tools, but particularly the [software] tool, is exceedingly worried about features that I am confident any competent printer will print just fine,” and Participant 13 (experience level 4), who said, “Once you’ve internalized the techniques for FDM printing, the tools are essentially useless. After running a print or two, you should have enough experience to reason through




Figure 6.9: Participants with low confidence in a successful print tended to have fewer design/manufacturing guideline violations (Confidence 3, “Neutral”; Confidence 4, “Confident”; Confidence 5, “Very confident”).

the things presented by the tools without referencing them at all.”

Several participants commented that the tools were useful for different reasons, like Participant 5 (experience level 4), who commented, “Good in conjunction, [the software] tool points out things, but worksheet does a better job explaining them,” and Participant 39 (experience level 3), who said, “[The software tool is] not designed for the best user experience but it’s cool and easier than troubleshooting [without] a computer.”

Our general takeaway from participant comments is best summed up by Participant 29 (experience level 5), who said, “The worksheet tool is full of tips that are easy to memorize for somebody who prints regularly. The [software] tool is more useful for finding small mistakes/flaws in design that could bypass a designer.”

6.3.5 Evaluation of the software tool

Based on participants’ comments and analysis of participants’ redesigns, we observed several areas for improvement of the software tool. The software tool was implemented in MATLAB, using MATLAB’s default zooming and rotating features. Participants found the zooming and rotating difficult, so it may be best to move the tool to a different platform with a more intuitive view rotation. The difficulty of view interaction and movement may have been one of the primary reasons for the lower usability rating and the higher reported task load associated with frustration reported by software participants, compared to worksheet participants (see Tab. 6.1).

It is also desirable to combine this type of analysis with popular CAD or slicing software because this would make the workflow from redesign to printing less cumbersome. Including

geometry analysis in a slicing program would also enable us to include some interactions with process parameters that the system does not currently consider (e.g., chamber and build plate temperature, and layer thickness)

Participants also found it difficult to interpret which features would be rounded (e.g., corners) versus features that were so small that they were below the resolution of the printer and were at risk of either failing to print or breaking, since these features were identified with the same color. This identification needs to be improved in future versions.

Although participants did not comment on this point, it is clear from their performance that the flagging of guideline violations should be improved. Improved flagging is especially necessary for overhanging faces. Currently, the system identifies problem geometry only when the designer clicks on each guideline category. In the current version of the tool, the STL faces are highlighted if they are overhanging, but if they are hidden from view at a given rotation (as they often were inside the logo), participants missed the highlighting. For each of the guidelines, a warning or flag of some kind needs to be added to the main menu to ensure users know that a guideline violation (like an overhanging feature) exists, so they know they need to find where it is located. A summary of violations could potentially reduce the occurrence of perfunctory use of the tool by overconfident designers who do not realize their design has problems that need to be addressed.

One potential feature that could be added is allowing the user to select key faces that the user cares about, and then suppressing surface finish or overhanging face warnings on the faces that the user did not select. (A similar approach using visual saliency is used in some orientation optimization schemes, like that of Zhang et al. [102]). This feature would reduce the number of warnings and the amount of visual clutter on the interface, possibly reducing the user's mental load.

We could also add estimates of build time and cost. Several participants described choosing an orientation based partially on print time or cost, even though these criteria were not included in the design problem. Although cost may not be a concern of our participants, who do not pay for filament or print time on our university's basic makerspace 3D printers, it is clearly of interest, and could easily be added.

Based on the comments we received, it appears that novices, in particular, found most of the features helpful for learning, but more experienced participants wanted to use the tool only for a quick check to confirm their analysis that was mostly based on prior experience. A more pared-down version of the tool, with fewer built-in descriptions of the guidelines and suggestions for redesign, could be more valuable to experienced designers, especially with improved notifications informing the designer that potential DFAM problems exist.

The software tool workflow currently requires more initial effort from the designer than the worksheet, because it requires the designer to access a computer, run the tool in MATLAB, and become familiar with the GUI. The extra effort required seemed to be a barrier to participants, as evidenced by the number of participants excluded from the study (six software participants who did not use the software tool versus one worksheet participant who did not use the worksheet). Additionally, Participant 5 (experience level 4) commented that he disliked having to install the image processing toolbox in MATLAB to use our software

tool. The inconvenience of setting up a DFM software tool decreases with repeated use and could be mitigated in industry with access to IT support.

6.4 Discussion

Analysis of our results indicates several useful findings and highlights many areas of future research. Although the usability of the worksheet was only moderately higher than the software tool, it is clear from participant comments that the usability of the software tool impacted participants' performance and perceptions. The impact of usability is also indicated in the higher task load associated with frustration for the software tool, which was one of the only two task load components where the worksheet outperformed the software tool. The need for a certain level of usability is likely a drawback for implementing software tools in either an educational or industry setting: the usability of the tool must be improved enough during the development process that designers are willing to use the tool repeatedly. Based on the fact that the majority of the participants in this study preferred the software tool, it appears that the usability of the current version of the tool is past this threshold, but there is still room for improvement.

We observed more variety in the build orientation chosen by participants using the software tool, whereas the difficulty in calculating which features would be overhanging seemed to encourage worksheet participants to prefer a basic orientation with orthogonal features. Using software tools could help designers consider a wider range of possible orientations and geometry than self-assessment alone, which could help design creativity. This observation may be related to findings reported by Abdelall, Frank, and Stone [182], which suggest that designers fixate less on non-producible features when they interact with DFM software. Further study on how CAD-based DFM tools could help ease design fixation, a frequent problem among designers [183], is needed.

Our findings, echoing those of Knijnenburg and Willemseen [50], indicate that a designer's prior experience, trusting propensity, and persistence played a role in user interaction and redesign efficacy. The DFAM tools, especially the software version, enabled similar redesign performance of experienced and inexperienced participants, but less experienced participants reported higher task load. Given the large range in the time reported by participants to complete the redesign task, as discussed in the evaluation of H4, it appears that a designer's persistence can play a role in their performance, with more persistent participants spending hours to fine-tune their designs. Our observations might be influenced by the fact that participants self-reported completion time and summarized their redesign process, which could lead to us missing trends regarding participant behavior if participants' reports were incomplete or inaccurate. However, we excluded participants who did not fully participate in the study activity, and trends emerged despite these limitations.

Based on participant comments, different participants trusted the recommendations of the system to varying degrees, with some preferring to rely on their prior experiences rather than the recommendations of the tool. Participants with high confidence in a successful

print had more guideline violations than their less confident peers (Fig. 6.9). This trend may indicate that some participants trusted their own judgment over the system and were not receptive to its analysis. To some degree, overconfidence could be mitigated with improvements in how the design guideline violations are presented to the designer. Hopefully, by first presenting the designer with a summary of the number of guideline violations on their part, we can engage even confident designers to explore the tool and improve their design. Another option to improve usage of this tool in an educational setting is to require students to actually print their redesigned part to reinforce trust in the tool's predictions and stress the importance of following DFAM guidelines. Future study is needed to better understand how to mitigate differences between designers.

Ahmed, Wallace, and Blessing [48] found that novice designers implemented and then evaluated design ideas, whereas experienced designers tended to ask more critical and evaluating questions, relying on their prior experiences to decide if a change was worth making, before implementing it. Based on these findings, our tool in its current form would be especially useful for novices, who prefer to repeatedly implement a design and evaluate it, learning from their mistakes as they go. A more pared-down version, with fewer design suggestions and background information, could be created for experts. Experts, who analyze the feasibility of their design based on prior experience, could use the system to check for subtle details they may have missed (such as hard-to-notice small or overhanging features, or features with geometry that is close to a cut off for size or orientation). Customization based on experience level is also in line with previous findings that suggest experts prefer systems that allow them significant control and enable them to leverage their expertise [50], and that novices and experts prefer different levels of detail [184]. These findings are likely to apply to other manufacturing processes as well, and developers of other DFM software should take heed.

Although the software tool helps identify these hard-to-notice geometry problems and seems to reduce mental load because of this, neither form of the tool truly addresses one of the main demands on working memory that a design problem such as this presents: resolving many conflicting objectives at once. The difficulty in balancing multiple objectives is a reason that optimization of build orientation in AM is so popular — it is impossible to try out every combination of geometry and build orientation. This type of tool is likely a good educational tool, and useful for quick checks for one-off parts, but could be used in tandem with design and orientation optimization in the future.

The geometric analysis of the software tool led to better redesigns with fewer design and manufacturing guideline violations, but several software participants failed to retest their redesigned parts. Lowering the effort required to test redesigns in the tool (e.g., creating a CAD add-in) could further improve performance. Our results indicate promise in the effectiveness of geometry analysis for improving designs. If our results and those of similar studies [42] are indicative of broader trends, DFM software can help designers evaluate the manufacturability of their designs more effectively, which has the potential to make concurrent engineering design easier.

6.5 Conclusions

In conclusion, we found that a software-based DFAM tool was more effective at eliminating DFAM guideline violations than asking participants to read and apply guidelines listed on a worksheet. Prior AM experience was correlated with lower task load and slightly faster performance during the experimental design task. Further development of the software tool should improve its usability and visualization. Our results highlight a need to understand further how designer attributes can affect performance using DFM tools, and how the design of such tools can help ensure good user experience and effectiveness for a wide range of designers. Using geometry analysis to power software-based DFM tools is a promising area to improve educational outcomes and concurrent engineering practices in industry.

In reference to the overarching hypothesis of this thesis, it appears that although software usability considerations can impact user experience negatively, geometry analysis does show great promise in improving the effectiveness of design for additive manufacturing feedback. Most subjects preferred using the version of tool that provided them with real-time, part-specific geometry analysis. This testing indicates that, especially once improvements have been implemented, software tools that perform geometry analysis will help designers learn manufacturing principles and evaluate parts for manufacturability problems.

Chapter 7

Conclusion

This final chapter will review the previous chapters of this thesis and interpret the implications of the key findings. Areas of future work will be discussed. The contributions of the thesis will be summarized.

7.1 Review of research objectives

This thesis evaluated the practicality, accuracy, and utility of geometric analysis in support of design for additive manufacturing (DFAM). Chapter 2 contextualized the research problem this thesis addressed. The design process for a new product or part is challenging, requiring designers to make countless interconnected decisions about the geometry, functionality, and price point of their part. Currently, designers struggle to predict how decisions regarding their product's geometry during the design stage can ultimately impact manufacturing cost and quality. Additive manufacturing (AM), a relatively new but promising technology, is being used increasingly for production parts, and so there is a need to help designers make informed trade-off decisions regarding manufacturing outcomes. Research has indicated that manufacturability and geometric accuracy of parts made using AM are impacted by a variety of factors, including part orientation, machine type, and layer thickness. By assessing the geometry of the part to be printed and the AM process parameters, we can automatically provide designers with design for additive manufacturing (DFAM) feedback. This feedback enables them to make changes to the geometry of their part or their planned manufacturing process before finalizing their design.

As summarized in Chapter 2, while some DFAM tools have been evaluated for more traditional technologies like casting or machining, few tools exist to support AM. The tools that do exist that have been developed in academia tend to have low usability, which makes it difficult to test the effectiveness of the geometry analysis itself and limits the adoption of these tools in industry. Although our survey of related work indicates that geometry analysis in support of DFAM has the potential to be practical, accurate, and effective, these characteristics have not been clearly demonstrated previously. To satisfy the unmet needs of designers and to explore the effectiveness of geometric manufacturability analysis for AM,

we developed two MATLAB-based tools, meant to be used by designers at different design stages.

Chapter 3 presented the background and development of these two tools: “Will It Print” and “Tolerance Allocation for Additive Manufacturing” (TAAM). Will It Print was developed to meet designer needs relatively early in the design process when designers are evaluating preliminary concepts. Will It Print was developed to help designers, especially those unfamiliar with AM, to evaluate their preliminary design concepts for compliance with common DFAM guidelines by evaluating their part geometry. Designers could explore the impact of geometry and orientation changes to their part, with analysis being updated in near-real-time. TAAM was designed to be used later in the design process when designers have both a detailed design and a preliminary tolerance allocation scheme. TAAM presents designers with visualizations evaluating all possible build orientations, and the achievable tolerances on their part at that build orientation, and helps a designer choose a single, “best” build orientation and refined tolerance scheme.

In the next chapter, Chapter 4, we demonstrated the use of geometry analysis tools with an example design scenario. The design scenario highlighted the practicality of the approaches we employ in the tools, showing how they can be used at different stages of the design process. We also used the example design scenario to compare both tools to any existing alternatives. Will It Print was compared to two commercial tools, Shapeways 3D Tools and Meshmixer. Will It Print had a more comprehensive analysis of manufacturability, evaluating compliance with more guidelines than Meshmixer or Shapeways. Will It Print also served as an improvement over existing tools, because it included more instructive information and suggestions for changes (useful for novices seeking to improve the manufacturability of their designs), and had a more accurate assessment of small features.

TAAM was compared with orientation optimization schemes that have been proposed in prior literature. Benefits of the approach we employed in TAAM include a more intuitive method for establishing preferences for optimization, and that more possible alternative orientations were displayed to the designer, in keeping with suggestions provided for other design-support tools [54, 55]. However, the computation time was longer. This chapter demonstrated that geometry analysis is practical for DFAM, but highlighted the need for developers to be intentional about decisions regarding the software architecture and user experience. Software systems need to achieve a balance of ease of use and the level of abstraction of the data that is presented in order to be useful. Our design scenario illustrated that geometry analysis tools should seamlessly fit into common design processes, rather than assuming designers will adapt to drastic new workflows or processes, an attribute that is also highlighted by Shukor and Axinte [39].

In Chapter 5, we sought to test the accuracy of the predictions regarding manufacturability and achievable geometric tolerances. Some research has been conducted to develop ranges of achievable geometric tolerances for parts manufactured using AM. However, this research is empirical and cannot be easily transferred to different machines. Moreover, there is no easy way to interpret how the achievable tolerances will be impacted by different process parameters (i.e., layer thicknesses, orientation, and so on). The mathematical model used

in TAAM, however, is designed to predict achievable tolerances based on orientation, layer-level inaccuracy, the shape of each layer, layer thickness, and presence of support material. Model testing indicated that TAAM is able to correctly predict the order of magnitude of error corresponding to flatness and cylindricity.

Moreover, the experimental relationship between achievable flatness and cylindricity for different orientations generally followed our predicted trend. However, our results indicated that the prediction of orientation errors needs to be improved. The quality of the datum face (i.e., the presence of stair-step error and warping on the datum face) needs to be factored into the mathematical model. The predictions used in TAAM were closer to the measured values than simpler models that only included layer thickness and orientation (such as the model used by Das et al. [138]). Similarly, the Will It Print tool was able to capture the risk of warping, tipping, and small features adequately, but could not convey the relative magnitude of the risk, and sometimes missed features close to thresholds of acceptability. The accuracy of the predictions of our tools is an improvement over what is currently available as it reflects the impact of some important sources of error, but the predictions could be improved by modeling additional sources of error. Our results support the suggestion made by Fornasini and Schmidt [185] that process parameters need to be considered in order to provide designers with the most accurate DFAM feedback. Geometry feedback and design guidelines should reflect the machine chosen and the process parameters that will be used.

Chapter 6 detailed a user study with a sample of 48 undergraduate students, which was conducted to confirm that geometry analysis tools can improve the efficacy of designers. The students completed a redesign task where they suggested geometry and orientation changes for a simple part. Roughly half of the students completed the task using the second software tool developed in this thesis, Will It Print, while the rest formed a control group that used written DFAM guidelines. The results of the study indicated that the use of geometry analysis was more effective than using written guidelines. Interestingly, some designer characteristics such as their prior experience impacted their use of the tool. The user study indicated that geometry analysis for DFAM can be effective, but the usability of the tool and certain designer characteristics can impact user experience and design efficacy. Designer characteristics, like experience and trusting propensity, have been shown to influence other design activities [48,50], and our results confirm that they are influential on DFAM activities as well.

In summary, the three different research thrusts generally supported the research hypothesis of this thesis. Geometry analysis for DFAM can be practical to implement into tools that designers find useful. An experimental study showed that the algorithms used to predict manufacturability and quality can provide feedback that is indicative of actual phenomena observed in AM, but refinement is needed to improve accuracy further. Finally, the user study showed the effectiveness of geometry analysis in a design scenario, helping designers to improve their designs more than analyzing the part without a tool.

7.2 Contribution

This thesis had several key contributions. Specifically, the contributions of this thesis are:

- An evaluation of the state of the art of DFAM feedback regarding geometric accuracy and manufacturability criteria, focusing on existing software-based DFAM tools.
- The development and dissemination of two DFAM tools, Will It Print and TAAM, that can be used by designers or researchers.
- A demonstration of geometry analysis for DFAM, illustrating that geometry analysis can provide useful predictions of AM outcomes with reasonable computation times.
- Experimental evaluation of the predictions of the two DFAM tools, highlighting areas for improving accuracy.
- Specific recommendations derived from user testing to improve the tools developed in this thesis as well as other, similar DFAM tools.
- Insights into designer characteristics that should be considered during the design of software-based design-for-manufacturing tools.

7.3 Future research

Future research is needed on the design process, DFAM feedback, and manufacturability analysis systems. Some specific suggestions for particular areas of future research are discussed below.

Process parameters not included in the current model could impact manufacturing outcomes. Parameters relating to the build method, infill method, support and infill density, build chamber temperature, adhesion type, and so on, may have an impact on manufacturing outcomes like deviations, warping, and tipping. Our understanding of the impact of these parameters could be aided by the collection and analysis of data sets describing the printing process. By collecting and analyzing part geometry, a wide set of process parameters, and machine specifications simultaneously, researchers may identify new relationships that improve the accuracy of current geometry analysis predictions.

One possible solution to enable a better understanding of the relative impact of different process parameters is a repository of printing data that is shared between researchers. Hobbyists and researchers could upload descriptions of their parts, describing manufacturability problems, or measuring accuracy on different features, along with the g-code and machine specifications used to create the part. Using data analytics, insights could be gained regarding settings or strategies to get good manufacturing outcomes for certain parts (e.g., large parts, parts at risk of tipping). In addition to collecting this information in a collaborative

online repository, more information could be scraped from existing online discussion boards for 3D printing.

Another important area of future research focuses on the human-computer interaction aspect of this thesis. There is great potential to bring in insights from the human-computer interaction field into mechanical design research and the development of design-support tools. One area in particular that needs attention is how to encourage more adoption of design-support tools by designers. We saw some resistance to adopting Will It Print in our user study. In general, encouraging the adoption of design research in industry is an on-going problem. Previous research in the field of human-computer interaction can serve as a jumping off point to further investigate what could cause low adoption of DFAM and design for manufacturing (DFM) tools. Users tend to prefer data visualizations that they are familiar with, even if those familiar visualizations are not as effective as alternatives [186]. Emotional and attitudinal factors in a software system can enhance acceptance and productivity [187]. Possible strategies to explore to promote adoption are improving trust by “making the algorithms of the automation simpler or by revealing their operation more clearly” [187] and involving the targeted user group in the development process of new tools [186]. Future research should seek to understand to what extent familiarity and trust impact usage of DFM and DFAM tools. Further, we should assess the applicability of strategies derived from the field of human-computer interaction to improve adaption of DFM and DFAM software tools.

While this thesis added to our knowledge about how designers consider manufacturing constraints, there is still much that is not known regarding the design-for-manufacturing process. The timing of DFM feedback in the design process could impact the efficacy of the feedback. Research could evaluate the efficacy of early, preliminary manufacturability feedback on reducing cost while also evaluating if the feedback constrained creativity and caused designers to ignore promising but complex concepts. Another area for further research is the impact of computation time on DFM decisions. It is likely that designers will respond better to real-time DFM feedback so they can explore a wide range of options quickly, as has been seen in other design-support tools [58].

As summarized by Ulrich and Pearson [188], manufacturing and design decisions are inexorably linked in determining product cost. However, if a designer can explore these trade-offs in a design tool, how would this knowledge impact designer decisions? By researching how designers currently consider manufacturing constraints during the design process, we can enable the development of better tools to support designer exploration of design and manufacturing trade-offs. Perhaps with improved manufacturing cost modeling, design and manufacturing companies can make more strategic decisions over investment strategies (i.e., whether to cut costs via redesign or moving production to a location with lower labor costs). Another promising area for development is creating software tools that enable designers to explore the suitability of different processes for manufacturing a part by estimating manufactured cost and expected quality based on geometry analysis, as proposed by Kim and Simpson [189].


7.4 Conclusion


As AM becomes a more viable option for manufacturing production parts, it will become increasingly important to help designers understand AM and how its strengths and limitations can affect the performance and design of their parts. Geometry analysis can help enable quick analysis of trade-offs between manufacturing cost and quality for AM, but its efficacy and feasibility had not been studied extensively in prior work. This thesis investigated the use of geometry analysis tools to support DFAM activities through the development and testing of two novel software tools. Three research thrusts were conducted: comparing the tools developed in this thesis with existing, alternative tools; experimental verification of the predictions of the tools; and user testing to evaluate the efficacy of the use of one of the tools on improving the manufacturability of an example part.






DFAM as a research field will likely continue to grow in the coming years, as AM technology improves and designers seek to develop products that take advantage of the freedom of form that AM enables. Design-support software for DFAM can help support this growth. Given that this thesis concluded that geometry analysis in support of manufacturability assessment for AM was feasible, accurate, and effective, we feel that geometry analysis is an especially promising area for development. Dual DFAM software systems (i.e., systems that help designers to consider both the benefits and limitations of AM) should be the focus of future research, using geometry analysis to evaluate manufacturability while still supporting design exploration and creativity. Interdisciplinary thinking will become increasingly important in developing the next generation of design-support tools for AM. By combining insights from the fields of design research, human-computer interaction, data visualization, in addition to manufacturing, we can develop more effective tools that are more widely adopted in industry. Geometry analysis systems can reach peak effectiveness when developers and researchers seek to *support*, rather than replace, humans in the design process.

Appendix A


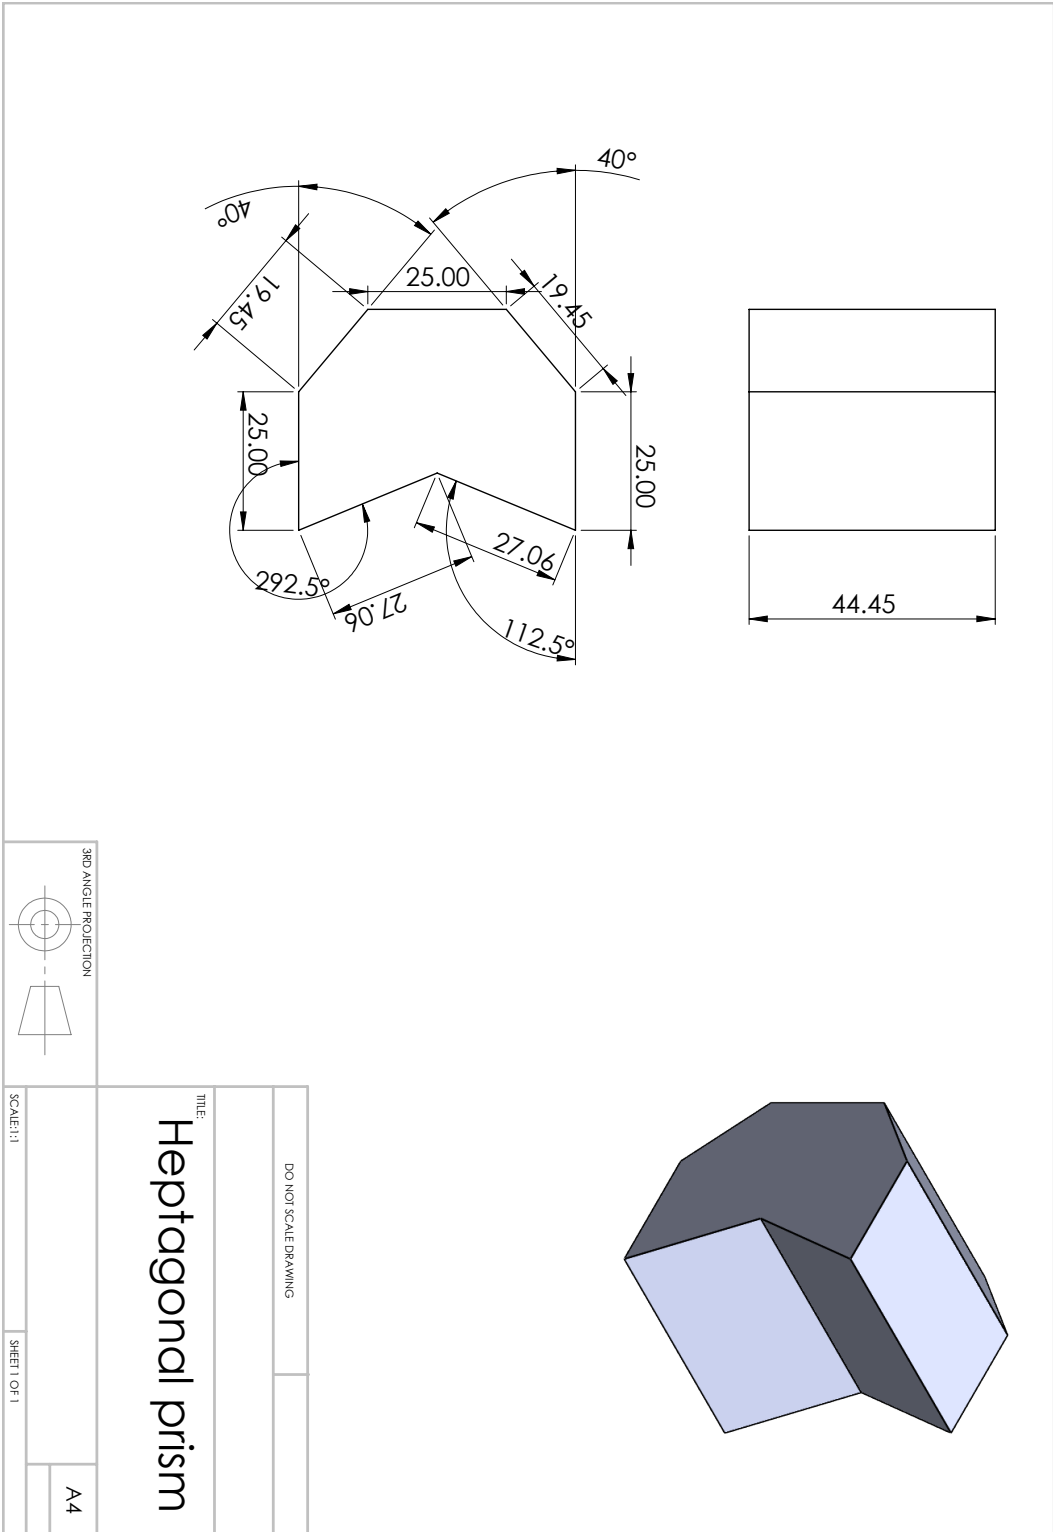
Experimental parts







DO NOT SCALE DRAWINGS	
TITLE: Flagpole bracket	
SCALE: 1:3	SHEET 1 OF 1
A4	



DO NOT SCALE DRAWINGS	
TITLE: Heptagonal prism	
SCALE: 1:1	SHEET 1 OF 1
A4	


Appendix B

Usability tool assignment

The assignments that were given to students using the Will It Print MATLAB tool, the worksheet, and the engineering drawing provided with both assignments are included in this appendix.

Problem 2

Download the [pencil_holder.stl](#) file from the embedded link². The STL format is the industry standard for describing 3D geometries that are to be printed. It consists of a list of triangular surface elements, which together form the surface of the object. You are tasked with identifying how to change the geometry of this part and choose a build orientation to avoid major printing errors when 3D printing the part. Download the [WillItPrint.zip](#) file from the embedded link and follow the unzipping instructions in this [tutorial](#). WillItPrint is a tool which will help you analyze the STL file to identify geometry that will need to be changed and choose a build orientation. The tool will highlight different regions of the part geometry that may cause printing errors (e.g. walls being too thin), and undesired characteristics (e.g. surfaces having a rough finish). Some of these errors depend on part orientation, so the tool enables you to explore the effect of changing the orientation of the part to choose a build direction that minimizes errors. You will need to run the tool using MATLAB, but no prior MATLAB knowledge is required. MATLAB is installed in the CAD labs (10 Jacobs and 1171 Etcheverry) if you don't have it installed on your personal computer. The [tutorial](#) will show you how to use the tool. Email the tool developer at hdb@berkeley.edu if you run into issues with the tool.



Start timing yourself on how long the following portion of the question takes you to complete. Decide what changes you would make to the part geometry to avoid printing errors, and choose a build orientation for printing the part. You will be annotating the multi-view drawing [pencil_holder.dwg.pdf](#) with your suggested changes. You can suggest changes to any part of the geometry with the following limitations: 1) the logo must be visible and remain the same size, although some rounding of the edges is acceptable; 2) the front face and logo must be relatively smooth for aesthetic reasons; 3) the height of the part should be between 50 and 70mm; and 3) there must be one hole that can fit a 18 by 30mm object and four holes that can fit a 12 by 12mm object, all with a depth of no less than 55 mm. This is an open-ended design problem without a single “correct” solution.


- a) Make a list of all changes you would make to the model's geometry. For each change, describe exactly how you would change the model geometry, and explain why you would make those changes, making references to the information contained in the tool. (You can include sketches, etc. if it makes it easier to explain the changes.) If you used any other tools to analyze the part other than the WillItPrint MATLAB tool, please mention them here.
- b) Print out and annotate the multi-view drawing of the part to show your proposed changes to the geometry and dimensions and upload a scan or picture of your annotated drawing with your homework. You can also use Adobe Acrobat to annotate the pdf electronically, and upload the edited pdf, if you prefer. (Alternatively, you can directly edit the [Solidworks file](#) and upload it to bCourses, if you are comfortable with Solidworks.)

² This design is a remix of another design from Thingiverse: <https://www.thingiverse.com/make:218614>

- c) Describe and sketch the orientation you would use to print the part. Be sure to clearly describe or show how the key faces are oriented, relative to the build plate. Why did you choose this orientation? What other orientations did you consider and why was this orientation better?
- d) Write down the amount of time Problem 2 took you to complete. You will also need it later to state how much time you spent on this problem in your HW submission quiz, where we will ask about it in addition to asking about the total time HW3 took you.

Problem 2




Download the [pencil holder dwg.pdf](#) file from the embedded link². You are tasked with identifying how to change the geometry of this part and choose a build orientation to avoid major printing errors when 3D printing the part. Please download the [WillItPrintWorksheet.pdf](#) file from the embedded link. This worksheet is a tool that will help you identify geometry that will need to be changed, and choose a good build orientation. The tool will help you identify different regions of the part geometry that may cause printing errors (e.g. walls being too thin), and undesired characteristics (e.g. surfaces having a rough finish). Some of these errors depend on part orientation, so the tool guides you through exploring the effect of changing the orientation of the part to choose a build direction that minimizes errors.



Start timing yourself on how long the following portion of the question takes you to complete. Analyze the part geometry shown on the drawing, decide what changes you would make to the part geometry to avoid printing errors, and choose a build orientation for printing the part. You can suggest changes to any part of the geometry with the following limitations: 1) the logo must be visible and remain the same size, although some rounding of the edges is acceptable; 2) the front face and logo must be relatively smooth for aesthetic reasons; 3) the height of the part should be between 50 and 70mm; and 3) there must be one hole that can fit a 18 by 30mm object and four holes that can fit a 12 by 12mm object, all with a depth of no less than 55 mm. This is an open-ended design problem without a single “correct” solution.

- Make a list of all changes you would make to the model’s geometry. For each change, describe exactly how you would change the model geometry, and explain why you would make those changes, making references to the information contained in the tool. (You can include sketches, etc. if it makes it easier to explain the changes.) If you used any other tools to analyze the part other than the WillItPrint worksheet, please mention them here.
- Print out and annotate the multi-view drawing of the part to show your proposed changes to the geometry and dimensions and upload a scan or picture of your annotated drawing with your homework. You can also use Adobe Acrobat to annotate the pdf electronically, and upload the edited pdf, if you prefer. (Alternatively, you can directly edit the [Solidworks file](#) and upload it to bCourses, if you are comfortable with Solidworks.)
- Describe and sketch the orientation you would use to print the part. Be sure to clearly describe or show how the key faces are oriented, relative to the build plate. Why did you choose this orientation? What other orientations did you consider and why was this orientation better?
- Write down the amount of time Problem 2 took you to complete. You will also need it later to state how much time you spent on this problem in your HW submission quiz, where we will ask about it in addition to asking about the total time HW3 took you.

² This design is a remix of another design from Thingiverse: <https://www.thingiverse.com/make:218614>



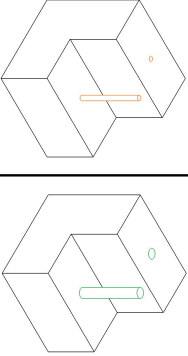
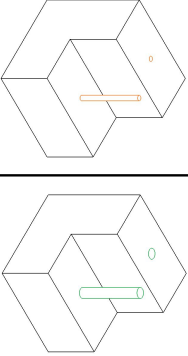
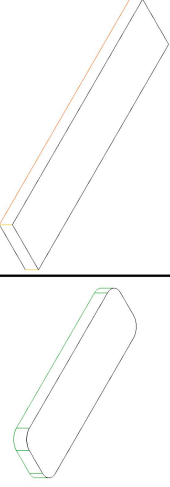
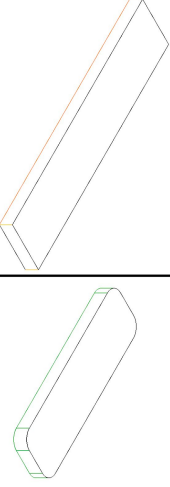
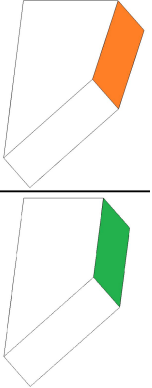
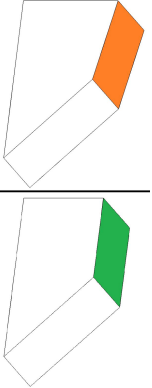
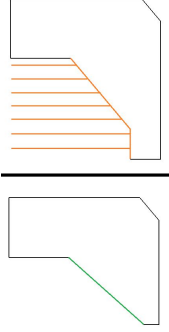
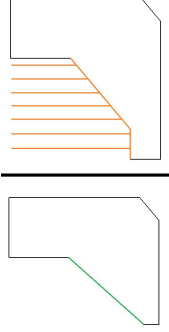
Appendix C

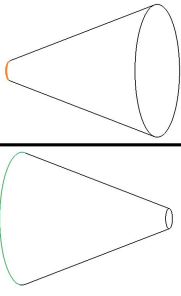
DFAM worksheet

The DFAM worksheet used as a control in the usability study described in Chapter 6 is included on the next two pages.

Introduction:

This worksheet should serve as a helpful tool for identifying geometry that might cause printing errors when making a part using fused deposition modeling (often called 3D printing). Read the descriptions of different types of problematic geometry, and then follow the tool's instructions on the next page for identifying geometry that needs to be changed and choosing an orientation to build your part in that best meets your needs.

Category	Problematic geometry	Acceptable geometry	Description
Small features			<p>Small features, including holes, may fail to print if they are smaller than 2 mm. Features dimensions should be equal to or larger than 2 mm in all directions to ensure they will print successfully and will not break.</p> <p>Also note that all sharp corners in your geometry may print slightly rounded. You can avoid this by avoiding sharp corners.</p>
Warping			<p>Long areas on build plate tend to curl or warp, peeling off of the build plate and potentially ruining your print. Reduce the length of the area touching the build plate by re-orienting the part or modifying the part geometry.</p> <p>Rounding sharp corners of the area on build plate can also help reduce warping.</p>
Rough surfaces			<p>If a face is almost but not exactly horizontal, the printed surface will be rough, with each printed layer clearly distinguishable from the next. This can be problematic when it occurs on features where a smooth finish is important for aesthetic or functional reasons. Re-orient the part or change the angle of the feature. Surfaces that are exactly horizontal, exactly vertical, or within 45deg from vertical will be smoother.</p>
Overhangs			<p>Surfaces that are overhanging (i.e. are oriented downward) need to be supported with extra material, called support material, that will need to be removed after printing. The base of your part, contacting the build plate, will also be printed with support material attached if you choose the "Rat" setting when printing your part. Removing support material can be difficult or impossible, especially from small cavities and from small/thin features that might break during removal. Also, removing support material can damage the surface of the part. To avoid support material on a key feature, re-orient the part or change the angle of the feature. Upward facing surfaces or downward facing surfaces that are more than 45deg from horizontal do not need support.</p>

Wobbling		<p>If there is only a small area on the build plate, parts can suffer from vibration issues, especially if they are tall. Printer movement can cause the part to wobble, possibly becoming detached from the build plate. If support material is present, it increases the area of material touching the build plate and helps stabilize the part. If the sum of the support material and part area touching the build plate is small, it is best to re-orient the part or change the geometry in order to minimize part height and to ensure the area on the build plate is large enough to provide a stable base.</p>
----------	--	---

Instructions:

1. Check to see if any of your part features are smaller than 2mm in any dimension. If there are any, the geometry needs to be changed to ensure every dimension is 2mm or larger.
2. Check if there are any sharp corners. If so, determine if some small rounding of those sharp corners is acceptable. If this is not acceptable, consider modifying the geometry to avoid sharp corners or choose another manufacturing process.
3. Choose a build orientation. Then check each of the following:
 - a. Check your part for sharp corners contacting the build plate. Consider rounding them to avoid warping.
 - b. Check your part for long areas contacting the build plate. If any dimension of the area is longer than 80 mm, consider shortening the geometry. If the geometry cannot be shortened, re-orient the part and return to Step 3.
 - c. Check your part for key features that are nearly but not exactly horizontal (parallel to the build plate). If it is important that this feature be smooth, consider changing the angle of the feature. Surfaces that are exactly horizontal, exactly vertical, or within 45deg from vertical will be smoother. If the geometry cannot be changed, re-orient the part and return to Step 3.
 - d. Check your part for key features that are overhanging (pointing downwards with an angle of less than 45deg from horizontal). This feature will need extra support material to be printed. If it will be difficult to remove the support material, or if you need a smooth surface on that feature, consider changing the angle of the feature to be upward facing or downward facing with an angle of more than 45deg from horizontal. If you don't want to change the geometry, re-orient the part and return to Step 3.
 - e. Check to see if the area of your part contacting the build plate, including any support material that will be required, is enough to provide a stable base for your part. This is especially important if your part is tall in this orientation. You should know how much support material your part will have in your chosen orientation from Step 4d. Consider changing the geometry to minimize the part height and maximizing the base height. If you don't want to change the geometry, re-orient the part and return to Step 3.
4. If you've found an orientation and a geometry that have passed the above checks, congrats! Your part should print successfully.

Bibliography

- [1] Budinoff, H., McMains, S., and Rinaldi, A., 2018, “An interactive manufacturability analysis and tolerance allocation tool for additive manufacturing,” *Proc. ASME IDETC/CIE Conf.*, Quebec City, Canada, August 26–29, ASME Paper No. DETC2018-86344, American Society of Mechanical Engineers.
- [2] Budinoff, H. and McMains, S., 2018, “Prediction and visualization of achievable orientation tolerances for additive manufacturing,” *Procedia CIRP*, **75**, pp. 81–86.
- [3] The World Bank, 2019, “Manufacturing, value added (% of GDP),” accessed 06/01/2019, <https://data.worldbank.org/indicator/nv.ind.manf.zs>
- [4] Bureau of Economic Analysis, 2019, “Gross Domestic Product by Industry: Fourth Quarter and Annual 2018,” Tech. rep., Bureau of Economic Analysis, <https://www.bea.gov/newsreleases/industry/gdpindustry/2018/pdf/gdpind417.pdf>
- [5] Scott, R. E., 2015, “The Manufacturing Footprint and the Importance of U.S. Manufacturing Jobs,” Tech. rep., Economic Policy Institute, <http://www.epi.org/publication/the-manufacturing-footprint-and-the-importance-of-u-s-manufacturing-jobs/>
- [6] National Science Board, 2016, “Science & Engineering Indicators,” Tech. rep., <https://nsf.gov/statistics/2016/nsb20161/uploads/1/nsb20161.pdf>
- [7] Barton, J. A., Love, D. M., and Taylor, G. D., 2001, “Design determines 70% of cost? A review of implications for design evaluation,” *J. Eng. Des.*, **12**(1), pp. 47–58.
- [8] Clive, D., Agogino, A., Eris, O., Frey, D. D., Leifer, L. J., Dym, C. L., Agogino, A., Eris, O., Frey, D. D., and Leifer, L. J., 2005, “Engineering Design Thinking, Teaching, and Learning,” *J. Eng. Educ.*, **94**(1), pp. 103–120.
- [9] Andersson, F., Hagqvist, A., Sundin, E., and Björkman, M., 2014, “Design for manufacturing of composite structures for commercial aircraft-The development of a DFM strategy at SAAB aerostructures,” *Procedia CIRP*, **17**, pp. 362–367.
- [10] Jarratt, T. A., Eckert, C. M., Caldwell, N. H., and Clarkson, P. J., 2011, “Engineering change: An overview and perspective on the literature,” *Res. Eng. Des.*, **22**(2), pp. 103–124.
- [11] Hartley, J. R., 1998, *Concurrent Engineering: Shortening Lead Times, Raising Quality, and Lowering Costs*, CRC Press.
- [12] Youssef, M. A., 1995, “Design for manufacturability and time-to-market: Part 2: Some empirical findings,” *Int. J. Oper. Prod. Manag.*, **15**(1), pp. 6–23.

- [13] Dowlatshahi, S., 1994, “A Comparison of Approaches to Concurrent Engineering,” *Int. J. Adv. Manuf. Technol.*, **9**, pp. 106–113.
- [14] Youssef, M. A., 1994, “Design for manufacturability and timetomarket: Part 1: Theoretical Foundations,” *Int. J. Oper. Prod. Manag.*, **14**(12), pp. 6–21.
- [15] Gupta, S. K., Regli, W. C., Das, D., and Nau, D. S., 1997, “Automated Manufacturability Analysis : A Survey,” *Res. Eng. Des.*, **9**(3), pp. 168–190.
- [16] Chiu, M.-C. and Okudan, G. E., 2010, “Evolution of Design for X Tools Applicable to Design Stages: A Literature Review,” *Proc. ASME IDETC/CIE Conf.*, Montreal, Canada, August 15–18, ASME Paper No. DETC 2010-29091.
- [17] Thompson, M. K., Moroni, G., Vaneker, T., Fadel, G., Campbell, R. I., Gibson, I., Bernard, A., Schulz, J., Graf, P., Ahuja, B., and Martina, F., 2016, “Design for Additive Manufacturing: Trends, opportunities, considerations, and constraints,” *CIRP Ann. - Manuf. Technol.*, **65**(2), pp. 737–760.
- [18] Hague, R., Campbell, R., and Dickens, P., 2003, “Implications on design of rapid manufacturing,” *Proc. Inst. Mech. Eng. Part C J. Mech. Eng. Sci.*, **217**(1), pp. 25–30.
- [19] Doubrovski, Z., Verlinden, J. C., and Geraedts, J. M. P., 2011, “Optimal design for additive manufacturing: Opportunities and challenges,” *Proc. ASME IDETC/CIE Conf.*, Washington, DC, August 28–31, ASME Paper No. DETC2011-48131.
- [20] Gao, W., Zhang, Y., Ramanujan, D., Ramani, K., Chen, Y., Williams, C. B., Wang, C. C., Shin, Y. C., Zhang, S., and Zavattieri, P. D., 2015, “The status, challenges, and future of additive manufacturing in engineering,” *Comput. Des.*, **69**, pp. 65–89.
- [21] Yang, L. and Anam, A., 2014, “An investigation of standard test part design for additive manufacturing,” *Proc. Int. Solid Free. Fabr. Symp.*, Austin, TX, pp. 901–922.
- [22] Munguía, J., Ciurana, J. D., and Riba, C., 2008, “Pursuing successful rapid manufacturing: a users’ best-practices approach,” *Rapid Prototyp. J.*, **14**(3), pp. 173–179.
- [23] Seepersad, C. C., Allison, J., and Sharpe, C., 2017, “The need for effective design guides in additive manufacturing,” *Proc. 21st Int. Conf. Eng. Des.* (Vol. 5), pp. 309–316.
- [24] Conradson, S., Barford, L., Fisher, W., Weinstein, M., and Wilker, J., 1988, “Manufacturability tools: an engineer’s use and needs,” *Fifth IEEE/CHMT Int. Electron. Manuf. Technol. Symp.*, pp. 155–158.
- [25] Ameta, G., Lipman, R., Moylan, S., and Witherell, P., 2015, “Investigating the Role of Geometric Dimensioning and Tolerancing in Additive Manufacturing,” *J. Mech. Des.*, **137**(11).

- [26] Nelaturi, S. and Shapiro, V., 2015, “Representation and analysis of additively manufactured parts,” *Comput. Aided Des.*, **67-68**, pp. 13–23.
- [27] Nelaturi, S., Kim, W., and Kurtoglu, T., 2015, “Manufacturability Feedback and Model Correction for Additive Manufacturing,” *J. Manuf. Sci. Eng.*, **137**(2).
- [28] Tedia, S. and Williams, C. B., 2016, “Manufacturability analysis tool for additive manufacturing using voxel-based geometric modeling,” *Proc. Solid Free. Fabr. Symp.*, Austin, TX.
- [29] Telea, A. and Jalba, A., 2011, “Voxel-based assessment of printability of 3D shapes,” *Int. Symp. Math. Morphol. Its Appl. to Signal Image Process.*, pp. 393–404, doi:10.1007/978-3-642-21569-8_34.
- [30] Ranjan, R., Samant, R., and Anand, S., 2015, “Design for Manufacturability in Additive Manufacturing Using a Graph Based Approach,” *Proc. ASME 2015 Int. Manuf. Sci. Eng. Conf.*, Charlotte, NC.
- [31] Subramaniam, B. L. and Ulrich, K. T., 1998, “Producibility Analysis Using Metrics Based on Physical Process Models,” *Res. Eng. Des.*, **10**(4), pp. 210–225.
- [32] Prasad, S., Zacharia, T., and Babu, J., 2014, “Design for Manufacturing (DFM) approach for Productivity Improvement in Medical Equipment Manufacturing,” *Int. J. Emerg. Technol. Adv. Eng.*, **4**(4), pp. 79–85.
- [33] Marion, T. J., Friar, J. H., and Simpson, T. W., 2012, “New product development practices and early-stage firms: Two in-depth case studies,” *J. Prod. Innov. Manag.*, **29**(4), pp. 639–654.
- [34] Bralla, J., 2000, *Design for manufacturability handbook*, 2nd ed., McGraw-Hill.
- [35] de Sam Lazaro, A., Engquist, D. T., and Edwards, D. B., 1993, “An Intelligent Design for Manufacturability System for Sheet-metal Parts,” *Concurr. Eng.*, **1**(2), pp. 117–123.
- [36] Venkatachalam, A. R., Mellichamp, J. M., and Miller, D. M., 1993, “A knowledge-based approach to design for manufacturability,” *J. Intell. Manuf.*, **4**(5), pp. 355–366.
- [37] Gadh, R., Herbert, D., Kott, A., and Kollar, C., 1991, “Feature-based design for manufacturability critique in concurrent engineering,” *Comput. Coop. Prod. Dev. MIT-JSME Work.*, D. Sriram, R. Logcher, and S. Fukuda, eds., Springer-Verlag, pp. 393–410, doi:10.1007/BFb0014288.
- [38] Hirschi, N. and Frey, D., 2002, “Cognition and complexity: An experiment on the effect of coupling in parameter design,” *Res. Eng. Des.*, **13**(3), pp. 123–131.

- [39] Shukor, S. A. and Axinte, D. A., 2009, “Manufacturability analysis system: Issues and future trends,” *Int. J. Prod. Res.*, **47**(5), pp. 1369–1390.
- [40] Kim, W., 2015, “A Framework for Set-based Manufacturing Analysis and Visual Feedback,” Doctoral dissertation, Pennsylvania State University.
- [41] Ullman, D. G., 2002, “Toward the ideal mechanical engineering design support system,” *Res. Eng. Des.*, **13**(2), pp. 55–64.
- [42] Barnawal, P., Dorneich, M. C., Frank, M. C., and Peters, F., 2017, “Evaluation of Design Feedback Modality in Design for Manufacturability,” *J. Mech. Des.*, **139**(9).
- [43] Riggs, B., Poli, C., and Woolf, B., 1998, “A Multimedia Application for Teaching Design for Manufacturing,” *J. Eng. Educ.*, **87**(1), pp. 63–69.
- [44] Lynn, R., 2017, “Enhancing awareness of additive and subtractive manufacturability with virtualized voxel-based simulations,” Master’s thesis, Georgia Institute of Technology.
- [45] Chandrasegaran, S. K., Ramani, K., Sriram, R. D., Horváth, I., Bernard, A., Harik, R. F., and Gao, W., 2013, “The evolution, challenges, and future of knowledge representation in product design systems,” *Comput. Aided Des.*, **45**(2), pp. 204–228.
- [46] Dillon, A. and Sweeney, M., 1988, “The application of cognitive psychology to CAD.” *People Comput. IV*, D. Jones and R. Winder, eds., Cambridge University Press, pp. 477–488.
- [47] King, N. and Majchrzak, A., 1996, “Concurrent engineering tools: are the human issues being ignored?” *IEEE Trans. Eng. Manag.*, **43**(2), pp. 189–201.
- [48] Ahmed, S., Wallace, K. M., and Blessing, L. T., 2003, “Understanding the differences between how novice and experienced designers approach design tasks,” *Res. Eng. Des.*, **14**, pp. 1–11.
- [49] Yu, B. Y., Honda, T., Sharqawy, M., and Yang, M., 2016, “Human behavior and domain knowledge in parameter design of complex systems,” *Des. Stud.*, **45**, pp. 242–267.
- [50] Knijnenburg, B. P. and Willemsen, M. C., 2011, “Each to His Own : How Different Users Call for Different Interaction Methods in Recommender Systems,” *Proc. 5th ACM Conf. Recomm. Syst. - RecSys '11*, Chicago, IL, pp. 141–148.
- [51] Randall, T., Terwiesch, C., and Ulrich, K. T., 2007, “User Design of Customized Products,” *Mark. Sci.*, **26**(2), pp. 268–280.

- [52] Simpson, T. W. and Martins, J. R. R. A., 2011, “Multidisciplinary Design Optimization for Complex Engineered Systems: Report From a National Science Foundation Workshop,” *J. Mech. Des.*, **133**(10).
- [53] Aloysius, J. A., Davis, F. D., Wilson, D. D., Taylor, A. R., and Kottemann, J. E., 2006, “User acceptance of multi-criteria decision support systems: The impact of preference elicitation techniques,” *Eur. J. Oper. Res.*, **169**(1), pp. 273–285.
- [54] Balling, R., 1999, “Design by shopping: A new paradigm?” *Proc. Third World Congr. Struct. Multidiscip. Optim.*, I. S. f. S. Optimization and Multidisciplinary, eds., Berlin.
- [55] Brill, E. D., Flach, J. M., Hopkins, L. D., and Ranjithan, S., 1990, “MGA: A Decision Support System for Complex, Incompletely Defined Problems,” *IEEE Trans. Syst. Man Cybern.*, **20**(4), pp. 745–757.
- [56] Mueller, S., 2016, “Interacting with Personal Fabrication Devices,” Doctoral dissertation, University of Potsdam.
- [57] Shneiderman, B., 1983, “Direct manipulation: A step beyond programming languages,” *Computer*, **16**(8), pp. 57 – 69.
- [58] Ligetti, C., Simpson, T. W., Frecker, M., Barton, R. R., and Stump, G., 2003, “Assessing the Impact of Graphical Design Interfaces on Design Efficiency and Effectiveness,” *J. Comput. Inf. Sci. Eng.*, **3**(2), p. 144.
- [59] Abi Akle, A., Minel, S., and Yannou, B., 2017, “Information visualization for selection in Design by Shopping,” *Res. Eng. Des.*, **28**(1), pp. 99–117.
- [60] Stump, G., Yukish, M., Martin, J., and Simpson, T., 2004, “The ARL Trade Space Visualizer: An Engineering Decision-Making Tool,” *10th AIAA/ISSMO Multidiscip. Anal. Optim. Conf.*, Albany, New York.
- [61] Kollat, J. B. and Reed, P., 2007, “A framework for Visually Interactive Decision-making and Design using Evolutionary Multi-objective Optimization (VIDEO),” *Environ. Model. Softw.*, **22**(12), pp. 1691–1704.
- [62] Yan, X., Qiao, M., Li, J., Simpson, T. W., Stump, G. M., and Zhang, X., 2012, “A work-centered visual analytics model to support engineering design with interactive visualization and data-mining,” *Proc. Annu. Hawaii Int. Conf. Syst. Sci.*, IEEE, pp. 1845–1854.
- [63] Daskilewicz, M. J. and German, B. J., 2012, “Rave: A Computational Framework to Facilitate Research in Design Decision Support,” *J. Comput. Inf. Sci. Eng.*, **12**(2).
- [64] Tweedie, L., Spence, B., Dawkes, H., and Su, H., 1995, “Influence explorer,” *Conf. Hum. Factors Comput. Syst. - Proc.*, **2**, pp. 129–130.

- [65] Seletos, T. M., 2016, “An Analysis of the Effectiveness of a Multi- Disciplinary Decision Support System on System- Level Decision Making,” Master’s thesis, Brigham Young University.
- [66] Laverne, F., Segonds, F., Anwer, N., and Le Coq, M., 2015, “Assembly Based Methods to Support Product Innovation in Design for Additive Manufacturing: An Exploratory Case Study,” *J. Mech. Des.*, **137**(12).
- [67] Rosen, D., 2014, “Design for Additive Manufacturing: Past, Present, and Future Directions,” *J. Mech. Des.*, **136**(9).
- [68] Valjak, F., Bojčetić, N., and Lukić, M., 2018, “Design for additive manufacturing: Mapping of product functions,” *Proc. Int. Des. Conf. Des.*, **3**, pp. 1369–1380.
- [69] Laverne, F., Segonds, F., D’Antonio, G., and Le Coq, M., 2017, “Enriching design with X through tailored additive manufacturing knowledge: a methodological proposal,” *Int. J. Interact. Des. Manuf.*, **11**(2), pp. 279–288.
- [70] Blösch-Paidosh, A. and Shea, K., 2019, “Design Heuristics for Additive Manufacturing Validated Through a User Study,” *J. Mech. Des.*, **141**(4).
- [71] Tang, Y. and Zhao, Y. F., 2016, “A survey of the design methods for additive manufacturing to improve functional performance,” *Rapid Prototyp. J.*, **22**(3), pp. 569–590.
- [72] Lieneke, T., Adam, G. A. O., Leuders, S., Knoop, F., Josupeit, S., Delfs, P., Funke, N., and Zimmer, D., 2015, “Systematical determination of tolerances for additive manufacturing by measuring linear dimensions,” *Proc. Solid Free. Fabr. Symp.*, Austin, TX, pp. 371–384.
- [73] Bähr, F. and Westkämper, E., 2018, “Correlations between Influencing Parameters and Quality Properties of Components Produced by Fused Deposition Modeling,” *Procedia CIRP*, **72**, pp. 1214–1219.
- [74] Turner, B. N. and Gold, S. A., 2015, “A review of melt extrusion additive manufacturing processes: II. Materials, dimensional accuracy, and surface roughness,” *Rapid Prototyp. J.*, **21**(3), pp. 250–261.
- [75] Mahesh, M., Wong, Y. S., Fuh, J. Y., and Loh, H. T., 2004, “Benchmarking for comparative evaluation of RP systems and processes,” *Rapid Prototyp. J.*, **10**(2), pp. 123–135.
- [76] Mercelis, P. and Kruth, J. P., 2006, “Residual stresses in selective laser sintering and selective laser melting,” *Rapid Prototyp. J.*, **12**(5), pp. 254–265.

- [77] Wu, A. S., Brown, D. W., Kumar, M., Gallegos, G. F., and King, W. E., 2014, “An Experimental Investigation into Additive Manufacturing-Induced Residual Stresses in 316L Stainless Steel,” *Metall. Mater. Trans. A Phys. Metall. Mater. Sci.*, **45A**, pp. 6260–6270.
- [78] Bouthillier, J. L., 2016, “The Importance of High Quality 3D Printer Filament,” accessed April 8, 2019, bootsindustries.com/the-importance-of-high-quality-3d-printer-filament/
- [79] Wittbrodt, B. and Pearce, J. M., 2015, “The effects of PLA color on material properties of 3-D printed components,” *Addit. Manuf.*, **8**, pp. 110–116.
- [80] Armillotta, A., Bellotti, M., and Cavallaro, M., 2018, “Warping of FDM parts: Experimental tests and analytic model,” *Robot. Comput. Integr. Manuf.*, **50**, pp. 140–152.
- [81] Comb, J., Priedeman, W. R., and Turley, P. W., 1994, “FDM technology process improvements,” *Proc. Solid Free. Fabr. Symp.*, Austin, TX, pp. 42–49.
- [82] Boschetto, A. and Bottini, L., 2014, “Accuracy prediction in fused deposition modeling,” *Int. J. Adv. Manuf. Technol.*, **73**, pp. 913–928.
- [83] Sood, A. K., Ohdar, R. K., and Mahapatra, S. S., 2009, “Improving dimensional accuracy of Fused Deposition Modelling processed part using grey Taguchi method,” *Mater. Des.*, **30**(10), pp. 4243–4252.
- [84] Zhang, Y. and Chou, K., 2008, “A parametric study of part distortions in fused deposition modelling using three-dimensional finite element analysis,” *Proc. Inst. Mech. Eng. Part B J. Eng. Manuf.*, **222**(8), pp. 959–967.
- [85] Mahmood, S., Talamona, D., Goh, K. L., and Qureshi, A. J., 2016, “Fast Deviation Simulation for Fused Deposition Modeling Process,” *Procedia CIRP*, **43**, pp. 327–332.
- [86] Mahmood, S., Qureshi, A. J., and Talamona, D., 2018, “Taguchi based process optimization for dimension and tolerance control for fused deposition modelling,” *Addit. Manuf.*, **21**, pp. 183–190.
- [87] Pennington, R. C., Hoekstra, N. L., and Newcomer, J. L., 2005, “Significant factors in the dimensional accuracy of fused deposition modelling,” *Proc. Inst. Mech. Eng. Part E J. Process Mech. Eng.*, **219**(1), pp. 89–92.
- [88] Cattenone, A., Morganti, S., Alaimo, G., and Auricchio, F., 2018, “Finite Element Analysis of Additive Manufacturing Based on Fused Deposition Modeling: Distortions Prediction and Comparison With Experimental Data,” *J. Manuf. Sci. Eng.*, **141**(1).
- [89] Minetola, P., Iuliano, L., and Marchiandi, G., 2016, “Benchmarking of FDM Machines through Part Quality Using IT Grades,” *Procedia CIRP*, **41**, pp. 1027–1032.

- [90] Arni, R. and Gupta, S. K., 2001, “Manufacturability Analysis of Flatness Tolerances in Solid Freeform Fabrication,” *J. Mech. Des.*, **123**(1), pp. 148–156.
- [91] Lieneke, T., Denzer, V., Adam, G. A., and Zimmer, D., 2016, “Dimensional Tolerances for Additive Manufacturing: Experimental Investigation for Fused Deposition Modeling,” *Procedia CIRP*, **43**, pp. 286–291.
- [92] Hernandez, D., 2015, “Factors Affecting Dimensional Precision of Consumer 3D Printing,” *Int. J. Aviat. Aeronaut. Aerosp.*, **2**(4).
- [93] Paul, R. and Anand, S., 2011, “Optimal part orientation in Rapid Manufacturing process for achieving geometric tolerances,” *J. Manuf. Syst.*, **30**(4), pp. 214–222.
- [94] Campbell, R. I., Martorelli, M., and Lee, H. S., 2002, “Surface roughness visualisation for rapid prototyping models,” *CAD Comput. Aided Des.*, **34**(10), pp. 717–725.
- [95] Serdeczny, M. P., Comminal, R., Pedersen, D. B., and Spangenberg, J., 2018, “Experimental validation of a numerical model for the strand shape in material extrusion additive manufacturing,” *Addit. Manuf.*, **24**, pp. 145–153.
- [96] Boschetto, A., Giordano, V., and Veniali, F., 2012, “Modelling micro geometrical profiles in fused deposition process,” *Int. J. Adv. Manuf. Technol.*, **61**(9-12), pp. 945–956.
- [97] Boschetto, A., Bottini, L., and Veniali, F., 2017, “Roughness modeling of AlSi10Mg parts fabricated by selective laser melting,” *J. Mater. Process. Technol.*, **241**, pp. 154–163.
- [98] Reeves, P. E. and Cobb, R. C., 1997, “Reducing the surface deviation of stereolithography using in-process techniques,” *Rapid Prototyp. J.*, **3**(4), pp. 20–31.
- [99] Ahn, D., Kweon, J.-h. H., Kwon, S., Song, J., and Lee, S., 2009, “Representation of surface roughness in fused deposition modeling,” *J. Mater. Process. Technol.*, **209**(15-16), pp. 5593–5600.
- [100] Strano, G., Hao, L., Everson, R. M., and Evans, K. E., 2013, “Surface roughness analysis, modelling and prediction in selective laser melting,” *J. Mater. Process. Technol.*, **213**(4), pp. 589–597.
- [101] Boschetto, A., Giordano, V., and Veniali, F., 2013, “3D roughness profile model in fused deposition modelling,” *Rapid Prototyp. J.*, **19**(4), pp. 240–252.
- [102] Zhang, X., Le, X., Panotopoulou, A., Whiting, E., and Wang, C. C. L., 2015, “Perceptual models of preference in 3D printing direction,” *ACM Trans. Graph.*, **34**(6).

- [103] Jiang, J., Lou, J., and Hu, G., 2019, “Effect of support on printed properties in fused deposition modelling processes,” *Virtual Phys. Prototyp.*, **14**.
- [104] Bacchewar, P. B., Singhal, S. K., and Pandey, P. M., 2007, “Statistical modelling and optimization of surface roughness in the selective laser sintering process,” *Proc. Inst. Mech. Eng. Part B J. Eng. Manuf.*, **221**(1), pp. 35–52.
- [105] Hong, Y. S. and Chang, T. C., 2002, “A comprehensive review of tolerancing research,” *Int. J. Prod. Res.*, **40**(11), pp. 2425–2459.
- [106] ElMaraghy, H. A., ed., 1998, *Geometric Design Tolerancing: Theories, Standards and Applications*, Springer.
- [107] ASME, 2009, “Dimensioning and Tolerancing Y14.5,” American Society of Mechanical Engineers.
- [108] ISO1101, 2017, “Geometrical product specifications (GPS),” International Organization for Standardization.
- [109] ASME, 2017, “Product Definition for Additive Manufacturing Y14.46,” American Society of Mechanical Engineers.
- [110] Moylan, S., Slotwinski, J., Cooke, A., Jurrens, K., and Dommez, M. A., 2012, “Proposal for a Standardized Test Artifact for Additive,” *Proc. Solid Free. Fabr. Symp.*, Austin, TX, pp. 902–920.
- [111] Hanumaiah, N. and Ravi, B., 2007, “Rapid tooling form accuracy estimation using region elimination adaptive search based sampling technique,” *Rapid Prototyp. J.*, **13**(3), pp. 182–190.
- [112] Nancharaiah, T., Raju, D. R., Raju, V. R., Ranga Raju, D., Ramachandra Raju, V., Raju, D. R., Raju, V. R., Ranga Raju, D., and Ramachandra Raju, V., 2010, “An experimental investigation on surface quality and dimensional accuracy of FDM components,” *Int. J. Emerg. Technol.*, **1**(2), pp. 106–111.
- [113] Gausemeier, J., Echterhoff, N., Kokoschka, M., and Wall, M., 2011, “Thinking ahead the Future of Additive Manufacturing Analysis of Promising Industries,” Tech. rep., DMRC, University of Paderborn, Patterborn.
- [114] Mani, M., Witherell, P., and Jee, J., 2017, “Design rules for additive manufacturing: A Categorisation,” *Proc. ASME IDETC/CIE Conf.*, Cleveland, OH, August 6–9, ASME Paper No. DETC2017-68446.
- [115] Brockotter, R., “Key design considerations for 3D Printing,” accessed 7/15/2018, <https://www.3dhubs.com/knowledge-base/key-design-considerations-3d-printing>

- [116] Booth, J. W., Alperovich, J., Chawla, P., Ma, J., Reid, T., and Ramani, K., 2017, “The Design for Additive Manufacturing Worksheet,” *J. Mech. Des.*, **139**(10).
- [117] Kumke, M., Watschke, H., and Vietor, T., 2016, “A new methodological framework for design for additive manufacturing,” *Virtual Phys. Prototyp.*, **11**(1), pp. 3–19.
- [118] Adam, G. A. and Zimmer, D., 2015, “On design for additive manufacturing: Evaluating geometrical limitations,” *Rapid Prototyp. J.*, **21**(6), pp. 662–670.
- [119] Thomas, D., 2009, “The Development of Design Rules for Selective Laser Melting,” Doctoral dissertation, University of Wales Institute, Cardiff.
- [120] Teitelbaum, G., 2009, “Proposed build guidelines for use in fused deposition modeling to reduce build time and material volume,” Master’s thesis, University of Maryland, College Park.
- [121] Seepersad, C. C., Govett, T., Kim, K., Lundin, M., and Pinero, D., 2012, “A Designer’s Guide for Dimensioning and Tolerancing SLS parts,” *Proc. Solid Free. Fabr. Symp.*, Austin, TX, pp. 921–931.
- [122] Meisel, N. and Williams, C., 2015, “An Investigation of Key Design for Additive Manufacturing Constraints in Multimaterial Three-Dimensional Printing,” *J. Mech. Des.*, **137**(11).
- [123] Allison, J., Sharpe, C., and Seepersad, C. C., 2019, “Powder bed fusion metrology for additive manufacturing design guidance,” *Addit. Manuf.*, **25**, pp. 239–251.
- [124] Burton, M. J., 2005, “Design for rapid manufacture: developing an appropriate knowledge transfer tool for industrial designers,” Doctoral dissertation, Loughborough University.
- [125] Samperi, M. T., 2014, “Development of design guidelines for metal additive manufacturing and process selection,” Master’s thesis, The Pennsylvania State University.
- [126] Ghiasian, S. E., Jaiswal, P., Rai, R., and Lewis, K., 2018, “From Conventional to Additive Manufacturing: Determining Component Fabrication Feasibility,” *Proc. ASME IDETC/CIE Conf.*, Quebec City, Canada, August 26–29, ASME Paper No. DETC2018-86238.
- [127] Lynn, R., Saldana, C., Kurfess, T., Reddy, N., Simpson, T., Jablokow, K., Tucker, T., Tedia, S., and Williams, C., 2016, “Toward Rapid Manufacturability Analysis Tools for Engineering Design Education,” *Procedia Manuf.*, **5**, pp. 1183–1196.
- [128] Armstrong, A. P., Barclift, M., and Simpson, T. W., 2017, “Development of a CAD-Integrated Cost Estimator to Support Design for Additive Manufacturing,” *Proc. ASME IDETC/CIE Conf.*, Cleveland, OH, August 6–9, ASME Paper No. DETC2017-68330.

- [129] Barclift, M., Armstrong, A., Simpson, T. W., and Joshi, S. B., 2017, “CAD-Integrated Cost Estimation and Build Orientation Optimization to Support Design for Metal Additive Manufacturing,” *Proc. ASME IDETC/CIE Conf.*, Cleveland, OH, August 6–9, ASME Paper No. DETC2017-68376.
- [130] Luo, Z., Yang, F., Dong, G., Tang, Y., and Zhao, Y., 2016, “Orientation optimization in layer-based additive manufacturing process,” *Proc. ASME IDETC/CIE Conf.*, Charlotte, NC, August 21–24, ASME Paper No. DETC2016-59969.
- [131] Ahn, D., Kim, H., and Lee, S., 2009, “Surface roughness prediction using measured data and interpolation in layered manufacturing,” *J. Mater. Process. Technol.*, **209**(2), pp. 664–671.
- [132] Frank, D. and Fadel, G., 1995, “Expert system-based selection of the preferred direction of build for rapid prototyping processes,” *J. Intell. Manuf.*, **6**, pp. 339–345.
- [133] Thrimurthulu, K., Pandey, P. M., and Reddy, N. V., 2004, “Optimum part deposition orientation in fused deposition modeling,” *Int. J. Mach. Tools Manuf.*, **44**(6), pp. 585–594.
- [134] Xu, F., Loh, H., and Wong, Y., 1999, “Considerations and selection of optimal orientation for different rapid prototyping systems,” *Rapid Prototyp. J.*, **5**(2), pp. 54–60.
- [135] Alexander, P., Allen, S., and Dutta, D., 1998, “Part orientation and build cost determination in layered manufacturing,” *Comput. Aided Des.*, **30**(97), pp. 343–356.
- [136] Moroni, G., Syam, W. P., and Petrò, S., 2014, “Towards early estimation of part accuracy in additive manufacturing,” *Procedia CIRP*, **21**, pp. 300–305.
- [137] West, A. P., Sambu, S. P., and Rosen, D. W., 2001, “Process planning method for improving build performance in stereolithography,” *Comput. Aided Des.*, **33**(1), pp. 65–79.
- [138] Das, P., Chandran, R., Samant, R., and Anand, S., 2015, “Optimum Part Build Orientation in Additive Manufacturing for Minimizing Part Errors and Support Structures,” *Procedia Manuf.*, **1**, pp. 343–354.
- [139] Savage, V. and Cheng, D., 2013, “Support Support : Visualizing Part Orientation and Effects Thereof for 3D Printing,” Tech. rep., University of California, Berkeley, <http://vis.berkeley.edu/courses/cs294-10-fa13/wiki/index.php/FP-ValkyrieSavage>
- [140] Rinaldi, A., 2017, “Design methodology for geometric dimensioning and tolerancing in additively manufactured parts,” Master’s thesis, ETH Zurich.

- [141] Mao, H., Kwok, T. H., Chen, Y., and Wang, C. C., 2019, “Adaptive slicing based on efficient profile analysis,” *Comput. Aided Des.*, **107**, pp. 89–101.
- [142] Boschetto, A., Bottini, L., and Veniali, F., 2016, “Integration of FDM surface quality modeling with process design,” *Addit. Manuf.*, **12**, pp. 334–344.
- [143] Prévost, R., Whiting, E., Lefebvre, S., and Sorkine-Hornung, O., 2013, “Make It Stand: Balancing Shapes for 3D Fabrication,” *ACM Trans. Graph.*, **32**(4), pp. 81:1—81:10.
- [144] Luo, L., Baran, I., Rusinkiewicz, S., and Matusik, W., 2012, “Chopper: partitioning models into 3D-printable parts,” *ACM Trans. Graph.*, **31**(6).
- [145] Stava, O., Vanek, J., Benes, B., Carr, N., and Měch, R., 2012, “Stress Relief: Improving Structural Strength of 3D Printable Objects,” *ACM Trans. Graph.*, **31**(4).
- [146] 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.
- [147] Gaynor, A. T. and Guest, J. K., 2016, “Topology optimization considering overhang constraints: Eliminating sacrificial support material in additive manufacturing through design,” *Struct. Multidiscip. Optim.*, **54**(5), pp. 1157–1172.
- [148] Aitkenhead, A., 2010, “Mesh Voxelisation,” accessed July 20, 2018, <https://www.mathworks.com/matlabcentral/fileexchange/27390-mesh-voxelisation>
- [149] Wang, T. M., Xi, J. T., and Jin, Y., 2007, “A model research for prototype warp deformation in the FDM process,” *Int. J. Adv. Manuf. Technol.*, **33**, pp. 1087–1096.
- [150] Guerrero-De-Mier, A., Espinosa, M. M., and Domínguez, M., 2015, “Bricking: A New Slicing Method to Reduce Warping,” *Procedia Eng.*, **132**, pp. 126–131.
- [151] Alsoofi, M. S. and Elsayed, A. E., 2017, “Warping deformation of desktop 3D printed parts manufactured by open source fused deposition modeling (FDM) system,” *Int. J. Mech. Mechatronics Eng.*, **17**(4), pp. 7–16.
- [152] GrabCAD, 2013, “GE jet engine bracket challenge,” accessed 2017-07-11, <https://grabcad.com/challenges/ge-jet-engine-bracket-challenge>
- [153] Woo, T. C., 1994, “Visibility maps and spherical algorithms,” *Comput. Aided Des.*, **26**(1), pp. 6–16.
- [154] van der Walt, S. and Smith, N., 2015, “mpl colormaps,” accessed 2017-10-01, <http://bids.github.io/colormap/>
- [155] Messac, A., 1996, “Physical programming - Effective optimization for computational design,” *AIAA J.*, **34**(1), pp. 149–158.

- [156] Huang, H. Z., Tian, Z. G., and Gu, Y. K., 2004, “Reliability and Redundancy Apportionment Optimization Using Interactive Physical Programming,” *Int. J. Reliab. Qual. Saf. Eng.*, **11**(03), pp. 213–222.
- [157] Ilgin, M. A. and Gupta, S. M., 2012, “Physical programming: A review of the state of the art,” *Stud. Informatics Control*, **21**(4), pp. 359–366.
- [158] Chase, K., 1999, “Tolerance allocation methods for designers,” Tech. Rep. 99-6, Brigham Young University.
- [159] Suresh, K., 2010, “Volume of a surface triangulation,” accessed 2019-03-20, <https://www.mathworks.com/matlabcentral/fileexchange/26982-volume-of-a-surface-triangulation>
- [160] Messac, A., 2015, *Optimization in Practice with MATLAB®: For Engineering Students and Professionals*, Cambridge University Press, New York, NY.
- [161] Taguchi, G., 1986, *Introduction to Quality Engineering*, UNIPUB/Kraus International.
- [162] Kalpakjian, S. and Schmid, S., 2001, *Manufacturing Engineering and Technology*, 4th ed., Prentice-Hall, Upper Saddle River, NJ.
- [163] Marsaglia, G., 1972, “Choosing a Point from the Surface of a Sphere,” *Ann. Math. Stat.*, **43**(2), pp. 645–646.
- [164] Ulrich, K. T. and Eppinger, S. D., 2008, *Product Design and Development*, 4th ed., McGraw-Hill.
- [165] Blasco, X., 2018, “Basic Genetic Algorithm,” accessed 2019-03-20, <https://www.mathworks.com/matlabcentral/fileexchange/39021-basic-genetic-algorithm>
- [166] Paul, R. and Anand, S., 2015, “Optimization of layered manufacturing process for reducing form errors with minimal support structures,” *J. Manuf. Syst.*, **36**, pp. 231–243.
- [167] Khardekar, R. and McMains, S., 2006, “Fast Layered Manufacturing Support Volume Computation on GPUs,” *Proc. ASME IDETC/CIE Conf.*, Philadelphia, PA, September 10–13, ASME Paper No. DETC2006-99666.
- [168] Lynn, C. M., West, A., and Rosen, D. W., 1998, “A Process Planning Method and Data Format for Achieving Tolerances in Stereolithography,” *Proc. Solid Free. Fabr. Symp.*, Austin, TX, pp. 293–302.
- [169] Lynn-Charney, C. and Rosen, D. W., 2000, “Usage of accuracy models in stereolithography process planning,” *Rapid Prototyp. J.*, **6**(2), pp. 77–87.

- [170] MatterHackers, “Type A Machines Series 1 Pro 3D Printer,” accessed 05/01/2018, <https://www.matterhackers.com/store/printer-kits/type-a-machines-series-1-3d-printer>
- [171] Zortrax, “Zortrax M200,” .
- [172] LulzBot, “LulzBot TAZ 5,” accessed 05/01/2018, <https://www.lulzbot.com/store/printers/lulzbot-taz-5>
- [173] Ultimaker, “Ultimaker 3,” accessed 05/01/2018, <https://ultimaker.com/en/products/ultimaker-3/specifications>
- [174] Hexagon Metrology, 2015, “ROMER Absolute Arm Overview Brochure. Hexagon Metrology.” accessed 05/15/2019, <https://www.hexagonmi.com/{~}/media/HexagonMILegacy/hxrom/romer/general/brochures/ROMERAbsoluteArm{ }overview{ }brochure{ }en.ashx>
- [175] UT Austin, 2017, “Design for AM Knowledge Base: Designer’s Guide for Additive Manufacturing,” accessed 07/01/2019, <http://designforam.me.utexas.edu/category/fdm/fdm-features>
- [176] Popov, V. L., Pohrt, R., and Li, Q., 2017, “Strength of adhesive contacts: Influence of contact geometry and material gradients,” *Friction*, **5**(3), pp. 308–325.
- [177] All, A., Nuñez Castellar, E. P., and Van Looy, J., 2016, “Assessing the effectiveness of digital game-based learning: Best practices,” *Comput. Educ.*, **92-93**, pp. 90–103.
- [178] Drato, 2016, “Office Pencil Holder,” accessed 08/30/2018, <https://www.thingiverse.com/thing:1338488>
- [179] Hart, S. and Staveland, L., 1988, “Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research,” *Adv. Psychol.*, **52**, pp. 139–183.
- [180] Brooke, J., 1996, “SUS-A quick and dirty usability scale,” *Usability Eval. Ind.*, P. W. Jordan, B. Thomas, I. L. McClelland, and B. Weerdmeester, eds., Taylor & Francis, pp. 189–194.
- [181] Carberry, A. R., Lee, H. S., and Ohland, M. W., 2010, “Measuring engineering design self-efficacy,” *J. Eng. Educ.*, **99**(1), pp. 71–79.
- [182] Abdelall, E. S., Frank, M. C., and Stone, R. T., 2018, “Design for Manufacturability-Based Feedback to Mitigate Design Fixation,” *J. Mech. Des.*, **140**(9).
- [183] Ullman, D. G., Stauffer, L. A., and Dietterich, T. G., 1986, “Preliminary Results of an Experimental Study of the Mechanical Design Process,” Tech. rep., Oregon State University.

- [184] Kumke, M., Watschke, H., Hartogh, P., Bavendiek, A. K., and Vietor, T., 2018, “Methods and tools for identifying and leveraging additive manufacturing design potentials,” *Int. J. Interact. Des. Manuf.*, **12**(2), pp. 481–493.
- [185] Fornasini, G. and Schmidt, L. C., 2015, “A Call for FDM Design Rules To Include Road Deposition,” *Proc. Int. Conf. Eng. Des. ICED15*.
- [186] Dasgupta, A., 2017, “Towards Understanding Familiarity Related Cognitive Biases in Visualization Design and Usage,” *DECISIVE 2017 Deal. with Cogn. Biases Vis.*, Phoenix, AZ.
- [187] Lee, J. D. and See, K. A., 2004, “Trust in Automation: Designing for Appropriate Reliance,” *Hum. Factors J. Hum. Factors Ergon. Soc.*, **46**(1), pp. 50–80.
- [188] Ulrich, K. T. and Pearson, S. A., 1993, “Does Product Design Really Determine 80% of Manufacturing Cost?” Working Paper #3601-93.
- [189] Kim, W. and Simpson, T. W., 2013, “Toward Automated Design for Manufacturing Feedback,” *Adv. Prod. Manag. Syst. Sustain. Prod. Serv. Supply Chain.*, V. Prabhu, M. Taisch, and D. Kiritsis, eds., Springer Berlin Heidelberg, pp. 40–47, doi:10.1007/978-3-642-41266-0_5.