2019-07-01

Application and Evaluation of Full-Field Surrogate Models in Engineering Design Space Exploration

Christopher Murray Thelin
Brigham Young University

Application and Evaluation of Full-Field Surrogate Models in Engineering Design Space Exploration

Christopher Murray Thelin

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

John L. Salmon, Chair
Steve E. Gorrell
S. Andrew Ning

Department of Mechanical Engineering

Brigham Young University

ABSTRACT

Application and Evaluation of Full-Field Surrogate Models
in Engineering Design Space Exploration

Christopher Murray Thelin
Department of Mechanical Engineering, BYU
Master of Science

When designing an engineering part, better decisions are made by exploring the entire space of design variations. This design space exploration (DSE) may be accomplished manually or via optimization. In engineering, evaluating a design during DSE often consists of running expensive simulations, such as finite element analysis (FEA) in order to understand the structural response to design changes. The computational cost of these simulations can make thorough DSE infeasible, and only a relatively small subset of the designs are explored.

Surrogate models have been used to make cheap predictions of certain simulation results. Commonly, these models only predict single values (SV) that are meant to represent an entire part's response, such as a maximum stress or average displacement. However, these single values cannot return a complete prediction of the detailed nodal results of these simulations. Recently, surrogate models have been developed that can predict the full field (FF) of nodal responses. These FF surrogate models have the potential to make thorough and detailed DSE much more feasible and introduce further design benefits. However, these FF surrogate models have not yet been applied to real engineering activities or been demonstrated in DSE contexts, nor have they been directly compared with SV surrogate models in terms of accuracy and benefits.

This thesis seeks to build confidence in FF surrogate models for engineering work by applying FF surrogate models to real DSE and engineering activities and exploring their comparative benefits with SV surrogate models. A user experiment which explores the effects of FF surrogate models in simple DSE activities helps to validate previous claims that FF surrogate models can enable interactive DSE. FF surrogate models are used to create Goodman diagrams for fatigue analysis, and found to be more accurate than SV surrogate models in predicting fatigue risk. Mode shapes are predicted and the accuracy of mode comparison predictions are found to require a larger amount of training samples when the data is highly nonlinear than do SV surrogate models. Finally, FF surrogate models enable spatially-defined objectives and constraints in optimization routines that efficiently search a design space and improve designs.

The studies in this work present many unique FF-enabled design benefits for real engineering work. These include predicting a complete (rather than a summary) response, enabling interactive DSE of complex simulations, new three-dimensional visualizations of analysis results, and increased accuracy.

Keywords: surrogate models, design space exploration, finite element analysis, fatigue life, modal analysis, optimization

ACKNOWLEDGMENTS

"But I never could have done it," he objected, "without everyone else's help."

-Norton Juster, *The Phantom Tollbooth*

I would like to thank Evan Selin, Christopher Ruoti, and Joseph Calogero at Pratt & Whitney for their regular guidance over the past three years. Pratt & Whitney is responsible for funding this research, which has changed my life; I am humbly grateful for this opportunity. Dr. Steven Gorrell has helped me weekly to refine my ideas and focus them on real engineering applications; he taught me about the mechanics of jet engines, and every paper I have written is better for his input. I would like to thank Dr. John Salmon for believing in my ability to make an important contribution and for doubting my self-doubts when things were rough. His enthusiastic curiosity and ability to see stories hidden amongst data points has taught me to look both deeply and broadly at the work I do.

The people you are around daily have a profound impact on who you become; I am profoundly grateful for the members of the CAD/XDL lab. Through countless questions, conversations, second-checks, and sharing our various interests, this work would not have been possible (or as much fun) without them. I especially would like to thank Greg Bird for his excellent and steady help with the Goodman diagram research; Landon Wright for his contagious excitement for new ideas and regularly helping see a step further than I could on my own; and of course Spencer Bunnell for being my partner in developing the full field surrogate modeling method, patiently helping me overcome obstacles, and approaching every challenge with optimism.

Finally, I want to thank my companion, best friend, and wife Rachel. She has been a constant source of encouragement and love in all my endeavors, from digging dinosaur bones to designing handcarts for the Apocalypse to studying jet engines. She has sacrificed for this research in as many ways as I have. Like in all our adventures, I am humbled and thrilled that Rachel has chosen to be by my side through it all.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

NOMENCLATURE

D	Response pattern for a new part design
DOE	Design of experiments
DSE	Design space exploration
FEA	Finite element analysis
FF	A full-field surrogate model (predicts a set of nodal values)
E	Magnitude error
K	Kelvin temperature units
K	Thousands
kPa	Kilopascals
lbs	Pound-force units
LHS	Latin Hypercube Sampling
m	Node that hosts the maximum response
MAC	Modal assurance criterion
MD	Magnitude difference
MPa	Megapascals
N	The number of nodes in a part
psi	Pounds-per-square inch
RBF	Radial basis function
RGB	Red Green Blue color system
$RMSE$	Root mean square error
S_e	Endurance strength
S_{ut}	Ultimate strength
SV	A single-value surrogate model (predicts one value)
SQP	Sequential quadratic programming
t	Target node
T	Target response pattern
%E	Percent error
%G	Percent Goodman
$\%G_{max}$	Percent Goodman, calculated at the max alternating and max steady stress
$\%G_{max,nodal}$	Maximum percent Goodman value from a set of $\%G_{nodal}$ values
$\%G_{nodal}$	Percent Goodman, calculated at each node
α	Scaling factor
δ	Displacement/distance
Δ	Magnitude displacement
λ	Constraint limit
σ	Von mises stress
σ_a	Alternating von mises stress
σ_{max}	Maximum von mises stress on a part
σ_{nodal}	Von mises stress on a single node of a part
σ_s	Steady von mises stress

CHAPTER 1. INTRODUCTION

On April 17, 2018, the left engine on Southwest Flight 1380 from LaGuardia to Dallas experienced an uncontained engine failure. Debris from the engine tore through the wing and fuselage. Tragically, the incident killed a passenger. Subsequent investigations showed that one of the fan blades had broken at the root, with evidence indicating that metal fatigue had been the cause of failure. A similar incident occurred in 2016 with a fan blade also failing due to metal fatigue.

While the vast majority of engine parts do not fail catastrophically, the fatigue experienced in these incidents indicate that the current design process is not always sufficiently thorough. Vibration-related failures account for approximately 60% of aero-engine failures, making the ability to accurately and easily predict failure due to fatigue an essential part of the jet engine design process [1]. Moreover, as companies push designs for greater engine performance and efficiency, it will become increasingly difficult to design parts that simultaneously meet aerodynamic, structural, and weight demands [2]. This requires more extreme designs with smaller margins of allowable error. For example, Pratt & Whitney has designed the geared turbofan (GTF) which produces up to 16% more thrust [3]. However, because the fan blades must be simultaneously light and strong, there is very little room for error. Margins of safety, part tolerances, and weight allowances decrease. The extreme performance demands make the parts difficult to design, as well as to repair. When parts are incorrectly manufactured, they must be scrapped and represent a significant source of loss for these companies. Without additional tools for design, the likelihood of parts failing will grow. Specifically, companies that are striving to push the current boundaries but keep passengers safe need efficient and simpler methods for finding designs that achieve higher performance but maintain sufficient strength.

In order to find the best designs possible, it is important to thoroughly explore the design space. The design space is the collection of all possible design variations. Each design in the design

space is uniquely defined by the value of its design parameters. When used with a parameterized model, design space exploration (DSE) consists of varying these parameters and examining the new designs. The new designs are used in engineering simulations to obtain results such as stress, displacements, etc. The results of these simulations are considered the response to changes in the design parameters.

Finite element analysis (FEA) is the most common form of simulation used in structural engineering design [4]. These simulations are invaluable, but are computationally expensive (in terms of time and computing resources). Although computing power has increased over time, simulations have also become more complex in order to meet the demands of a competitive aerospace market [5]. Consequently, even with modern computing power, the quantity and complexity of aerospace simulations that need to be performed continue to make their computational cost a limiting factor for their use during the design process.

Surrogate models are a cheap and efficient solution to the computational expense of full simulations. Surrogate models (also known as response surface models, metamodels, regression models, or emulators [6]) create simplified mathematical relationships between inputs and outputs of a system. The models can then use this relationship to predict the results of a simulation in response to new input values. These surrogate models make it possible to evaluate many design variations during DSE or optimization without the prohibitive cost of full simulations [7–12]. Their convenience and speed when properly trained makes surrogate modeling a valuable part of engineering design. While every surrogate has some degree of error, based on the particular method and training samples used, generally some amount of bias error is acceptable when conducting design space exploration early in a design process.

Most surrogate models map a set of inputs to a single output value. These types of surrogate models will be referred to as single-value (SV) surrogate models. They cannot predict the full response of an entire system, but instead predict a single value that represents the system (see Fig. 1.1 (top)). Often, the single values are used to summarize the quality of a design, such as the maximum stress, aerodynamic coefficients, or the weight of the part. By only predicting single values, the information SV surrogate models provide in an optimization routine is limited [7]. Single values provide little insight into the detailed differences between two complex objects [13], such as the nodal results between FEA models. While values like the maximum stress and weight

are important indicators of model quality, a good design is ideally determined by patterns and relationships between the nodal results across the entire part.

In recent years, methods have been developed for using surrogate models to predict the behavior of every node in an FEA simulation [14–16]. They generally take geometric parameter values as inputs and predict the complete FEA results for a part. These methods give a much more detailed and complete prediction of structural FEA results than SV surrogate model methods, and enable a three-dimensional visualization of the structural results mapped onto the part geometry. These surrogate models can be solved very quickly, allowing the predicted FEA response to changes in parameter values to be shown in real time, where real time is defined as updated responses appearing in under a second [14]. While they necessarily rely on more training data than SV surrogate models (i.e., every node's results instead of a single result for each training sample), they also allow detailed full-field predictions.

Because this surrogate modeling method predicts an entire field of responses from an FEA simulation, these will be referred to as full-field (FF) surrogate models. This work uses the terminology of SV and FF surrogate modeling methods in order to provide a distinction between these different types of surrogate models. The difference between SV and FF surrogate models is illustrated in Fig. 1.1: while SV surrogate models only predict a single value, FF surrogate models predict values for each node on a part.

If industry can confidently adopt the use of FF surrogate modeling methods into these existing analyses, then these methods can begin to make a larger impact [17]. However, while FF surrogate models show promise in terms of speed and accuracy [14–16], they have currently only been used to predict steady stress contours. No studies have been published that test the claims that FF surrogate models can enhance DSE activities and real engineering analyses. Steady stresses are important in engineering design, but more complex analyses require fatigue life and modal data. FF surrogates have not been applied to these types of data and analyses. The benefits of FF surrogate models in these applications has been mostly speculative. In order to build confidence in the abilities of FF surrogate models for real engineering design and DSE, they must be tested in more complex analyses. Their accuracy and abilities compared to SV surrogate models in these contexts must be explored. In this way, FF surrogate models can be further established as viable alternatives or additions to SV surrogate models in engineering applications.

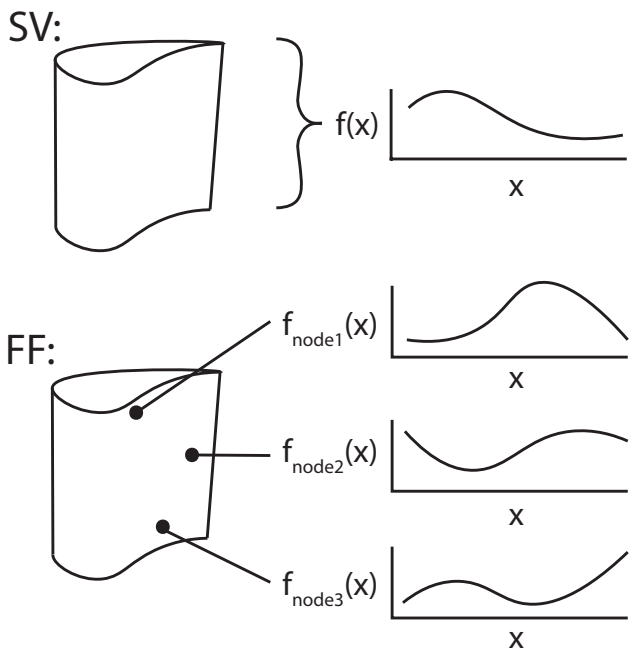

Figure 1.1: (top) Single-value (SV) surrogate models only predict one value that represents the response of the entire part across the design space. (bottom) Full-field (FF) surrogate models predict the response of each node on the part across the design space.

The research objective of this thesis is to evaluate the accuracy and benefits of FF surrogate models when used in real engineering applications and design space exploration. This general objective is addressed by the following research questions:

RQ1 Can FF surrogate models be used reliably in real engineering and design space exploration applications?

RQ2 How do FF surrogate models compare with SV surrogate models in these activities, especially in terms of accuracy and ability?

RQ3 What additional design benefits are gained by using FF surrogate models in the design process?

These questions aim to promote confidence among engineers concerning the usefulness of FF surrogate models for real engineering design, and reveal the associated benefits and shortcomings of the method.

1.1 Thesis Organization

The central chapters of this thesis (Chapters 2 - 5) present work that has been submitted for publication by the author in a variety of academic engineering journals pertaining to FF surrogate models in real engineering applications. In each, the various methods and examples are demonstrated with a jet engine compressor blade model (the Purdue blade [18]). Although the methods are generalized and could apply to any three-dimensional geometry and FEA model, the compressor blade is used throughout for consistency and to take advantage of the expertise of the academic and industry participants.

Chapter 2 contains a study that expands on previous work by testing the claims that FF surrogate models enable enhanced interactive DSE experiences and demonstrates visualization techniques that make DSE with FF surrogate models more efficient. Because the responses of the entire design are obtained at low cost, it becomes possible to quickly compare different designs across the design space; however, the cognitive load of making accurate comparisons is high. A difference model visualization was created that updates dynamically as design parameters are varied. The study describes the procedure and results for a set of user experiments that test how human subjects perform when conducting design space exploration activities with and without this visualization. The results show that the interactive exploration and visualizations enabled by FF surrogate models do provide enhanced results for some design comparison activities. The interactive difference model for comparing FEA results is presented as a benefit gained by using FF surrogate models in the design process. The majority of this chapter has been accepted for publication in the journal *Information Visualization*.

Chapter 3 investigates the effect of calculating values based on two distinct sets of FF surrogate model predictions. FF surrogate models are used to predict steady and alternating stresses. These stresses are mapped onto a Goodman (or Haigh [19]) diagram, a common engineering tool for predicting the fatigue life and risk of a structural part under cyclical loading [20, 21]. The predicted stresses are used to calculate various measures of fatigue safety on the Goodman diagram. The accuracy of these values are then compared to those obtained via SV surrogate models. It is determined that FF surrogate models predict the information with comparable-to-improved accuracy compared to SV surrogate models, making them valuable in this type of analysis. Various benefits associated with using the full nodal predictions are explored which allow designers

to avoid overconservative design and more intuitively understand the interaction between design parameters and fatigue response. A dynamically updating Goodman diagram is presented for interactive DSE of fatigue life data. The content of this chapter has been submitted for publication to the *Journal of Mechanical Design* and is under review as of July 2019.

Chapter 4 investigates cases where a calculation combines all the nodal results of a FF prediction and all the nodal results of an actual simulation. FF surrogate models are used to predict and compare modal displacements, or mode shapes. The modal assurance criterion (MAC) equation is the most common method of measuring similarity between two mode shapes from different designs [10, 22, 23]. The accuracy of the MAC predictions from FF and SV surrogate modeling methods is studied across different modes and training sample sizes. It is concluded that, when the data is highly nonlinear across a design space, the FF surrogate models struggle to predict results with the accuracy of SV surrogate models, and require an extremely high amount of samples. The considerations and methods associated with using FF surrogate models for modal data are detailed. If properly trained, various benefits related to flexibility and full visualizations are made uniquely possible by FF surrogate models. The content of this chapter has been submitted for publication to the journal of *Mechanical Systems and Signal Processing* and is under review as of July 2019.

Chapter 5 applies FF surrogate models in optimization routines for more efficient DSE. Because the FF surrogate models make predictions for each node on a three-dimensional model, each value in the prediction of a design has some geometric or spatial relationship to the others. These values enable spatially-defined constraints and objectives that allow a designer to better control the optimization of the design with respect to the three-dimensional response pattern across the part. The FF-enabled techniques are described and demonstrated with various examples. This study illustrates further ways that FF surrogate models can offer enhanced design benefits, especially in design space exploration and optimization. This content will be submitted to the journal of *Structural and Multidisciplinary Optimization* in August 2019.

Table 1.1: Summary of chapter differences

Chapter	RQ1: Application	RQ2: vs SV?	RQ3: Benefits	Predicted data	General characteristics
2: Difference model	DSE	N	3D comparative visualizations Interactive DSE	Steady stresses	1 set of predicted nodal data for testing DSE enhancements
3: Goodman	DSE Fatigue analysis	Y	Interactive fatigue DSE Less conservative predictions 3D fatigue visualization	Steady stresses Alt. stresses	Nodal calculations using 2 sets of predicted nodal data for fatigue analysis
4: MAC	DSE Modal analysis	Y	Flexible modal comparisons (high cost of training data)	Modal displacements	Single calculation using 1 set of predicted nodal data 1 set of data from simulation for modal comparisons
5: Optimization	DSE Optimization	Y	Spatially-defined optimization Enhanced control over optimization	Steady stresses	1 set of predicted nodal data for DSE using optimization

Table 1.1 displays a summary of the differences between the chapters. The Application column describes which engineering design activities are contained in each chapter in response to the first research question (RQ1). The third column describes whether or not the chapter attempts to determine accuracy and capability comparisons with SV surrogate models, in response to the second research question (RQ2). The Benefits column gives brief descriptions of the benefits FF surrogate models add to the given engineering application in response to the third research question (RQ3). The specific kind of nodal data that was predicted in each chapter is listed in the fifth column. Finally, the sixth column shows a generalized way of describing the differences between each chapter's research.

Chapter 6 describes general conclusions from these four studies. These are followed by a brief description of future work pertaining to this study. Relevant information from the individual chapters is presented in the appendices.

CHAPTER 2. DIFFERENCE MODEL

Existing research has demonstrated that FF surrogate modeling is possible, and has alluded that it can enable enhanced design space exploration and visualization [14]. This enhancement is based on the premise that as more information is available, better conclusions can be drawn about the subject [24]. However, a concentrated study on FF surrogate modeling in design space exploration activities has not previously been published. Also, the ability to explore a design space at interactive rates presents various challenges to the user. With the complete set of nodal results available for every design rapidly displayed, the cognitive load of comparing one design to another in meaningful ways becomes very high.

This chapter contains a study that has been accepted for publication in the journal *Information Visualization*. It describes the basic method of creating FF surrogate models for interactive design space exploration. It also answers the challenge of making meaningful design comparisons by applying a difference modeling technique to the three-dimensional FEA result predictions. This visualization clarifies the differences between two designs. The results of a user experiment are described that uses the FF surrogate models for an interactive design space exploration of steady stress contours on a jet engine compressor blade, and measures the effectiveness of using the difference model visualization for design comparisons. The comparative visualization techniques are shown to help improve perception speed and accuracy when making various design comparisons between two designs. This enables more efficient design space exploration as engineers compare large numbers of complex objects, such as nodal compressor blade predicted by FF surrogate models in an interactive design space exploration.

While this chapter does not make comparisons with SV surrogate models, it does apply FF surrogate models to real engineering activities (RQ1) and explore benefits of using FF surrogate models in the design process (RQ3).

2.1 Introduction

In engineering, design space exploration is an important process for finding optimal designs [25]. The design space is the collection of all possible design variations. When used with a parameterized model, the space may be explored by varying input parameters and running simulations at critical points. Results are then compared to each other with the goal of discovering the interactions between parameters and outputs [26]. The structural results of a design (strength, stiffness, etc.) are affected by the geometry, or the values of the design parameters. Engineers can vary parameters until they achieve a set of parameter values that produces optimal structural results (e.g., lower stresses, minimize displacement, or shift stress contours to a more desirable configuration) [1].

Design space exploration involving structural results can be inhibited by the time required for finite element analysis (FEA) and other analyses to solve. FEA is a simulation method which approximates the shape of a model with a mesh of points called nodes, applies loads and boundary conditions to the model, and then solves for results such as structural stress and displacements [4]. This process is meant to simulate how a part will respond to real-world conditions. These results are produced for each individual node in the model. The speed of FEA is largely dependent on model complexity [14], and models used in engineering disciplines such as aerospace can have hundreds of thousands of nodes and anywhere from 5 to 50 parameters [27, 28]. The expense in time and computer resources can cause designers to consider fewer designs during design space exploration and limits the ability of simulations to facilitate exploration of design alternatives [26].

Surrogate modeling has recently been applied to this problem and allows designers to quickly emulate and visualize structural FEA results across a design space [14–16]. Surrogate models create a relationship between input data and output data. Instead of computationally expensive simulations, these relationships can predict responses very quickly with low cost [12]. The surrogate model can be created, or trained, with the design parameters to a FEA model as inputs and the response, such as the stress, at each node of a FEA model. This allows the surrogate model to predict the complete stress response of the FEA model for new input parameter values, and indicate the strength of a model without needing to solve new computationally expensive simulations. These results can be obtained very quickly. With this process, designers may adjust the input parameters to a model and see the predicted structural results displayed in a three-dimensional

visualization immediately. Bunnell et al. [14] showed that this method can be applied to large, complex models such as those used for turbomachinery compressor blade design.

These methods make thorough exploration of the design space for a part much more feasible; however, even with the ability to visualize simulated FEA results in real-time, better tools are needed to make it easier to visually compare one design to another [26]. With a large number of possible designs in the design space as well as a vast number of points to compare between three-dimensional structural results in each design, recalling the myriad differences becomes quite difficult. When changes between two three-dimensional objects are small or subtle, the human visual system has limited perception [29]. Designing visualizations to aid the user in making such comparisons is a challenging, yet important, endeavor for discovery in science and engineering [30]. This is especially true for comparisons involving spatial three-dimensional data, where spatial means the data has a meaningful three-dimensional length, width, height, etc.

According to the taxonomy for comparative visualization developed by Gleicher et al. [31, 32] and extended for spatial three-dimensional comparisons by Kim et al. [30], a side-by-side comparison of results (juxtaposition) is useful and common. Unfortunately, when objects are large with many points to compare, all the burden is placed upon the user's memory and limits the ability to make effective comparisons [30, 31]. One solution is to use an explicit encoding visualization, such as displaying only the difference between the objects. This type of visualization provides a much more focused comparison of the objects than juxtaposition alone and is useful when the focus of the comparison is the relationship between two objects [30, 31]. Difference images and difference models have been used in many diverse fields for comparison purposes [33–43]. These visualizations often show the degree and location of differences between two objects or models and strip away less essential information.

This research applies a type of difference modeling for visualizing the differences between the structural results of two separate three-dimensional designs. This type of explicit encoding visualization will hereafter be used interchangeably with the phrase "difference model." These differences are visualized to facilitate design comparison, particularly during design space exploration. The difference between the value of each node on one design and the value of each corresponding node on another design is computed using a simple node-by-node relationship [32]. The amount of

change at each node is then displayed on the difference model, allowing the user to easily perceive even slight differences.

It is expected that, when the difference model is displayed alongside the two designs, combining the juxtaposition and explicit encoding methods of visualization in coordinated multiple views, shortcomings in each type of visualization will be overcome [30, 31]. With these visual obstacles removed and differences clearly displayed, users could more quickly and more accurately judge how changing design parameters affects a part design.

Although difference modeling and hybrid comparative visualizations are not new, Kim et al. recently suggested that a valuable branch of research for comparative visualization includes not only exploration of new visualizations but also user studies that help quantify the advantages of various types of visualizations [30]. These user studies should include a variety of tasks and should be used in real data analysis applications.

With these facts in mind, the contributions of this paper are to (1) present an application of using a hybrid juxtaposition and explicit encoding (a difference model) visualization to facilitate comparison between different designs in a design space for surrogate-modeled structural results, and (2) conduct an experiment to help validate the claim that a hybrid visualization, rather than just a juxtaposition visualization, for this type of application will improve speed and reduce error in making meaningful comparisons between parts. The experiment tests the hypothesis for a variety of tasks and information, including search and quantitative estimation, as suggested by Kim et al. [30]. This research presents these results, which show evidence that the difference model does improve performance in some types of design tasks, while for others it produces no statistically significant advantages.

The scope of the application and experiments in this research is limited to showing the difference between the structural results of three-dimensional compressor blade finite element meshes emulated by surrogate models, though this method could apply to a wide variety of engineering part design spaces. The implementation presented uses a reference, or baseline, design as one of the objects to be compared. Because the experiments in this research used random treatment assignment (i.e., participants randomly were assigned a certain visualization method to test), statistical inferences about the causal effect of the difference model may be drawn from the results; however, because the test subjects were self-selected volunteers (rather than by using random sam-

pling methods of selection), general inferences about the larger engineering population are beyond the scope of this research. This study provides useful data that supports conclusions about the effects of the visualization methods in this study, but further testing is required to make conclusions about the effect across the population of all engineers.

This paper will proceed as follows: the surrogate modeling method used to create the data for this research will be briefly covered, followed by an overview of how difference images and difference models have been used in various fields for error visualization and other purposes. In Section 3, the specific methods for creating the difference model will be discussed. In Section 4, the experiment design will be explained and results for each task will be presented. Finally, conclusions are presented that show there is indeed evidence that the difference model improves speed and accuracy for some comparison tasks and that the difference model can be a useful tool in engineering design comparison.

2.2 Related Work

This section will give a brief overview of FEA and attempts to predict responses from finite element analyses with surrogate modeling. This will include a discussion of how surrogate models are created. Other work that covers different comparative visualization techniques fro complex objects will also be presented.

2.2.1 Surrogate-modeled FEA

Finite element analysis is a common engineering tool for obtaining the response of a part to a set of loading conditions, or forces [4]. FEA programs, such as ANSYS, take a two or three dimensional model of a part and convert it into a mesh of distinct points, or nodes. The program determines the response to the loads by solving equations at each node of the mesh. These responses could include values such as the amount of physical displacement or stresses at each node caused by the loads on the part. The response values are assigned a color and mapped onto the model's geometry to create colored contours representing the response across the entire part (see the example in Figure 2.1).

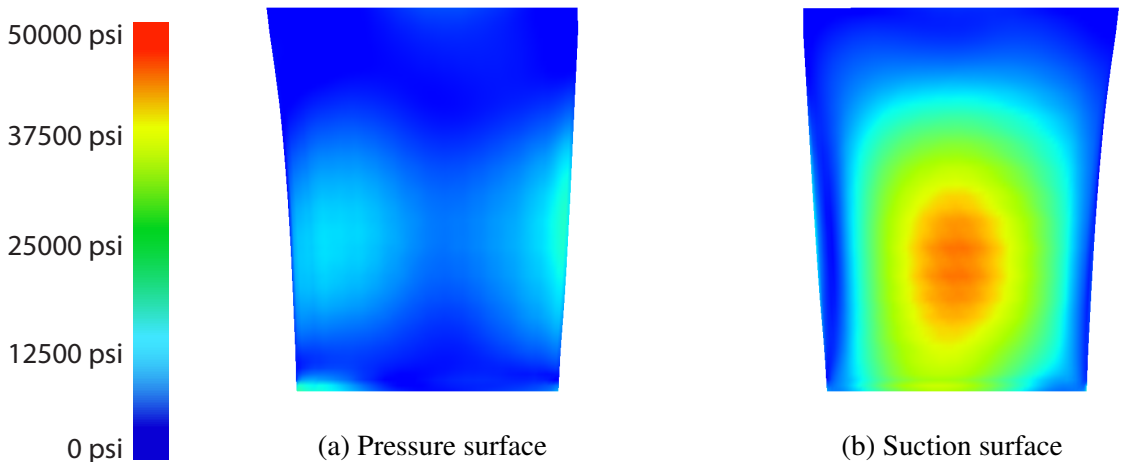

(a) Pressure surface (b) Suction surface

Figure 2.1: An example of FEA results on a compressor blade shape

Surrogate modeling uses training data to create a relationship between inputs and outputs to a system. The resultant relationship can then be used to predict the system's response to a new set of input values in a fraction of the time it takes to solve the original system [6, 12]. In recent years, surrogate models have been used to emulate the stress and displacement response of each node in a finite element mesh [14–16]. A variety of designs are created by using a design of experiments (DOE) to find a set of designs that adequately fill the design space. For a model with n input parameters, each design is represented by an n-dimensional vector [26]. Once the designs are generated, FEA is performed for each design.

The input parameter values for each design constitute the inputs to the surrogate models, and the FEA result values constitute the outputs. The surrogate models are used to model the relationship between these and can then take in new parameter values to predict the response across the part. These predicted results are assigned a color value based on the chosen color scale and mapped onto a reconstructed visualization of the finite element mesh. This process allows a user to change the parameters of a model and obtain visualized results up to 96% more quickly than waiting for the full FEA process [14], allowing a user to conduct rapid design space exploration with instant structural feedback. Being able to quickly visualize alternate designs in a design space without the need to set up many independent models and wait for simulations to solve enables faster and more creative exploration.

2.2.2 Comparative Visualization

Gleicher et al. developed various classifications of visualization techniques for facilitating comparison [31]. These classifications have been used and expanded by others [30, 44], but the basic categories have proven useful:

- Juxtaposition: a side-by-side comparison of two objects
- Superposition: two or more objects shown in the same space (also called "Overlay")
- Explicit encoding: computes a relationship between objects and represents the relationship, not the original models, in space

Each has inherent strengths and weaknesses. Juxtaposition shows each object in its entirety but places the burden of comparison on the user's memory and perception as they look back and forth between objects. This may be helped by linked camera views, etc. [30, 44]. Superposition locates the parts together in space but suffers from occlusion. Explicit encoding directly shows a desired relationship, such as the differences, which makes the user's task quite easy, but results are decontextualized from the original objects [30, 31].

When framing a comparative visualization problem, one must consider various challenges, as described by Gleicher in a later study [32]:

1. Number of items to compare
2. The size/complexity of the objects to be compared
3. The size/complexity of the relationships

The most common form of comparison is between two different objects; comparison of three or more can be incredibly difficult. Objects being compared can be large with many simple small differences, small with very complicated differences, or any combination of these. Specifically, when an object has only small, subtle differences, the challenge for the user to perceive and correctly understand these differences is greater. Finally, the relationships between the objects may be simple or complex. Simple relationships include those where objects to be compared have corresponding elements that may be checked element-by-element [32]. When dealing with comparisons between

two three-dimensional shapes, the challenge of determining relationships becomes much more difficult if there are no clear correspondences between points on one object to the other [32, 45].

Research has suggested that hybrid methods are advantageous when comparing displays of three-dimensional data [30]. These can include coordinated multiple views or tightly integrated combinations of various methods. These combinations help combine strengths and overcome weaknesses. When combining juxtaposition and explicit encoding, the explicit encoding makes the relationship between the juxtaposed views clear, and the juxtaposed views provide context for the explicit encoding [31].

2.2.3 Comparative Visualization in Scientific Studies

There are many examples of how explicit encodings (or difference models and difference images) assist researchers to make scientific comparisons in many fields. In biological studies, these methods have been used to indicate increasing or decreasing accuracy of plant reflectance models [36], exposing toxic leaf chemicals [39], automatically detect cells in the root of a plant [41], and illustrating concentrations of ethanol vapor on a sensor [40]. In medical fields, difference imaging can identify possible breast cancer in patients by taking two images of a patient in different positions and comparing them [38]. In computer vision, a type of difference imaging is used for motion detection by examining pixel differences between two consecutive frame images in a video sequence [42] or calculating the magnitude and direction of change with optical flow vectors [43].

A type of comparative visualization common to scientific fields is error visualization. This involves calculating the difference between an approximation and a ground truth in order to determine the error between the models. Explicit encodings may be used to map error values directly onto a three-dimensional geometry, and allow a user to gain quick intuition about the reliability of the model with respect to its geometric location and features.

A specific example of the use of difference modeling with three-dimensional models is the program *Metro*. Here, a geometric mesh is loaded and then simplified for faster computation [46]. The differences in position between each node of the simplified mesh and of the original mesh are computed, creating an error value for every node on the mesh. For example, a simplified mesh that closely resembles the original mesh will have low error at all nodes. These error values are represented with a hue and then displayed on the surface of the simplified mesh to indicate the

distance between where the node is and where it should be. This allows the viewer to see, at a glance, where the simplified mesh geometry differs from the original and the degree of difference. The result is visually similar to the stress contours found in typical FEA results, but the colors represent varying degrees of signed error instead of stress. Using this visualization, engineers may more easily make judgments about the quality of the simplified mesh.

Similarly, in a 2010 study concerning ocean floor mapping, models of varying fidelity were created and then compared. The distance between each surface location on two different models could be represented as a signed error (positive or negative height differences), which was then converted to a range of colors and displayed on the geometry. This allowed the researchers to determine whether the lower density meshes were reliable enough to use for further research purposes.

Finally, a combination of juxtaposition and explicit encoding has been used in some computational fluid dynamics (CFD) analyses to show the error between different models of a specific flow scenario (e.g., a high vs a low resolution model) [34, 35]. The researcher may compare the two models in their full representations via juxtaposition and then use the difference model, shown to the side of the juxtaposed models, to clearly understand which locations have the most change in accuracy.

While these examples do not use difference modeling in the context of design space exploration or structural results on three-dimensional models, they do illustrate how using explicit encodings can help expose otherwise hidden results in a study. These studies use visualizations, but do not attempt to analyze the visualization choices used. This research aims to not only apply these methods to a unique type of aerospace analysis, but also examine the comparative visualization benefits gained from the chosen visualizations.

2.3 Method

In this section, a method for applying a hybrid visualization to nodal results and the details of its implementation is described.

2.3.1 Visualization Considerations

This research uses a combination of juxtaposition and explicit encoding. The juxtaposition represents the simplest way to compare two objects: side by side, with little to no extra insight to relationships between the two objects. The explicit encoding directly computes and conveys the differences between the two objects as positive or negative. This method is ideal for representing whether stresses at specific locations have increased or decreased for different designs in a design space due to changes in the parameters. By including both the juxtaposition and explicit encoding, the visualization presents the original models to the engineer for examination, clearly shows the differences, and does not suffer from decontextualization.

To address the challenges of comparative visualization set forth by Gleicher [32], it was determined that, for this application, only two designs need be compared at a time. The entire design space could be summarized by statistical methods, but insight into the significance of differences is gained from looking at the full individual models [47]. The focus here is on direct structural comparison, and looking at two sets of results at a time is sufficient.

In regards to object and relationship complexity, the three-dimensional objects being compared may be quite large (each finite element model may be made up of thousands to hundreds of thousands of nodes) and the parts of the objects being compared (the results at each node) are comparatively very small. This increases the difficulty for the user to perceive such small areas of change. However, Gleicher noted that when the relationships between parts are simple, such as checking a set of data in an element-by-element fashion, the challenges are reduced [32].

Each different design in the design space is composed of different geometric parameter values, so pairs of objects being compared in this application will have different shapes. Because of this, it becomes harder to find ways to directly relate one point to another between the two objects [45]. For this purpose, this application makes use of mesh morphing during the data generation phase. Mesh morphing adjusts the nodes in an existing finite element mesh to a new shape, rather than create a new mesh [48]. As the geometry changes, the relative positions and numbering of the nodes in the finite element mesh are preserved. Thus, a consistent number and relative positioning of nodes exists for each design in the design space, and any two designs have an exact set of corresponding nodes and results to be compared. Thus, the relationships between nodal values on the objects are simple as per Gleicher's definition.

2.3.2 Basic Setup

The juxtaposed views consist of two renderings of designs in the design space and their structural results. When displayed, they look like the output of a FEA simulation. The camera views are linked so that both designs are always in the same orientation. This has been done to help the viewer make direct comparisons between both visualizations [30, 44].

The difference model, or explicit encoding visualization, is made by calculating the nodal differences between two designs. The calculated nodal differences are assigned a color based on the magnitude of the difference. This visualization maps these colors onto a third geometry. This produces a separate, three-dimensional representation of the geometry that, instead of displaying the results of a single design, shows the differences between the results of two designs. For the purposes of this study, the geometry used for this rendering is the geometry of the first object to be compared. Though this does not make the results entirely independent of the first object [31], it does provide a clearer, more meaningful representation of the data than applying it to some other geometry outside the comparison. The juxtaposed views are required for analysis of the geometric differences, while the explicit encoding shows the mapped result differences.

The difference model uses the results for each node in the compared designs to calculate difference values. These results could be stresses, displacements, temperatures, or any set of nodal values across a mesh. In this paper, the displayed results will be von Mises stresses [49], which are useful for determining problematic stress locations. Because each design in the design space is built from the same parametric finite element mesh, there exists a result value for each node in every possible design. Thus, no matter which two designs are compared, there also exists a difference value for every node in the mesh.

To render these difference values to the difference model geometry, they must be assigned a color that aids the designer in perceiving the changes that have occurred. The signs of these differences can be determined by calculating whether results at each node on the second design being compared have increased or decreased in relation to the first design. For example, the result values that have increased from the first to the second design can be represented with varying degrees of warm colors, e.g. orange and red. Values that have decreased can use varying degrees of cool colors, e.g. teal and blue. Values that have not significantly changed can be colored a neutral color, such as white. The threshold for significance can be chosen by the user. An example

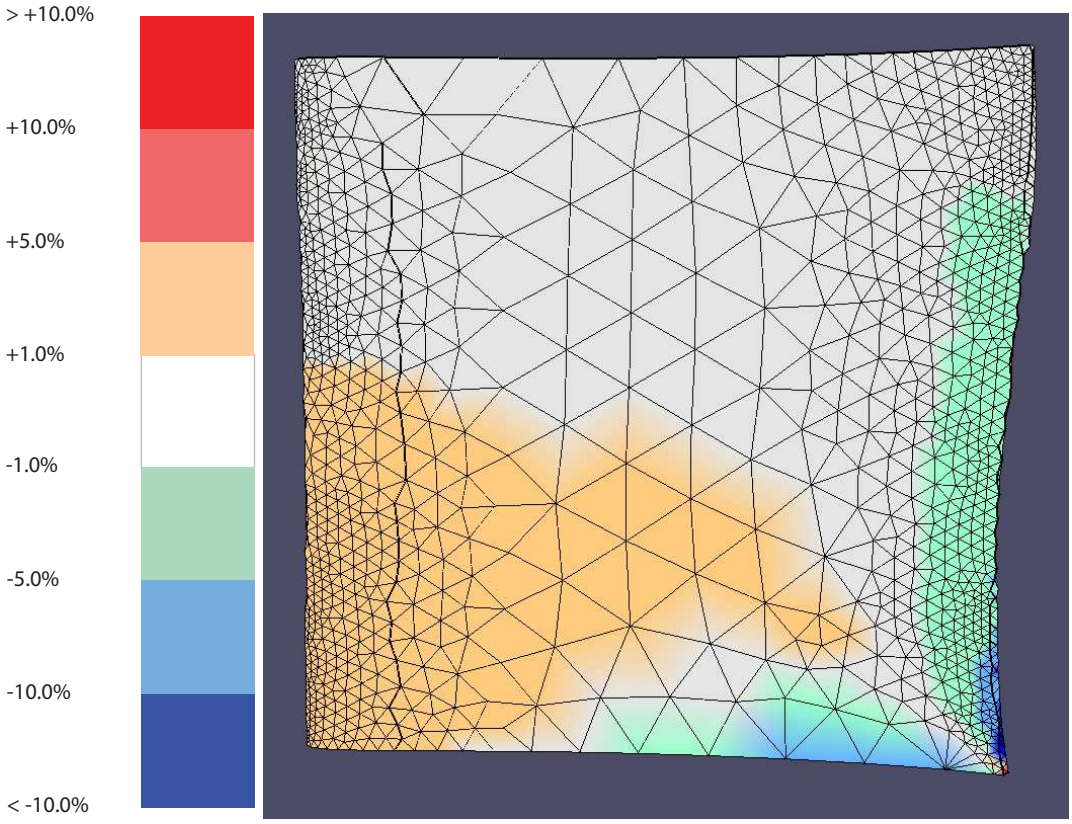

Figure 2.2: The difference model's color scale and the difference model

of these colors is shown in Figure 2.2. The result is a "double-ended" color scale which indicates whether results have increased or decreased by varying degrees, with warm and cool colors at the extremes and colors tending from bold to neutral as the differences decrease in magnitude.

The difference model communicates the differences between the two designs more precisely and more completely than a direct visual comparison of the designs [30, 31]. From observing the two original designs alone, many subtle—yet significant—differences are still difficult to perceive or mentally judge. The difference model can highlight exactly how much each node has changed, where changes have occurred, and quickly provide the designer with a visual map of how design changes affect the structural results between the two designs.

2.3.3 Implementation Details

In this implementation, a simple parametric geometry similar to a turbomachinery compressor blade design is used. The geometry used in this study will be referred to as a "blade" for

simplicity. It is controlled by three parameters: the height, the chord at the root, and the chord at the tip. The root refers to the base and the tip refers to the top face. The leading edge faces the direction of motion, and the trailing edge faces away. The chord refers to the distance from the trailing edge to the leading edge. Unless otherwise specified, the trailing edge will be located on the right hand side of all figures in this paper, and the tip will be oriented towards the top of the figures. Each of these three parameters have a baseline value of 3.0 in (7.62 cm) and can vary by 10% (0.3 in) in either direction, creating a range of 2.7 in to 3.3 in (6.86 cm to 8.38 cm). The design space is three-dimensional, with each axis representing a different parameter between the values of 2.7 and 3.3.

This geometry and the associated nodal stress responses were obtained by creating surrogate models using the process outlined by Bunnell et al. [14]. A design of experiments was used to obtain a set of designs that would suitably fill the design space. In ANSYS, these designs were given loading conditions of a fixed constraint of the root face and a load on the tip's edge. These were solved for each design, and the results for each node were used to train radial basis surrogate models.

At this point, the three parameters' values may be changed by the user, and the new nodal stresses are calculated by the surrogate models fast enough for the visualization to update without the user noticing any delay. Because the surrogate models were trained to respond to parameter values from 2.7 to 3.3, parameter values outside this range are not guaranteed to produce accurate results.

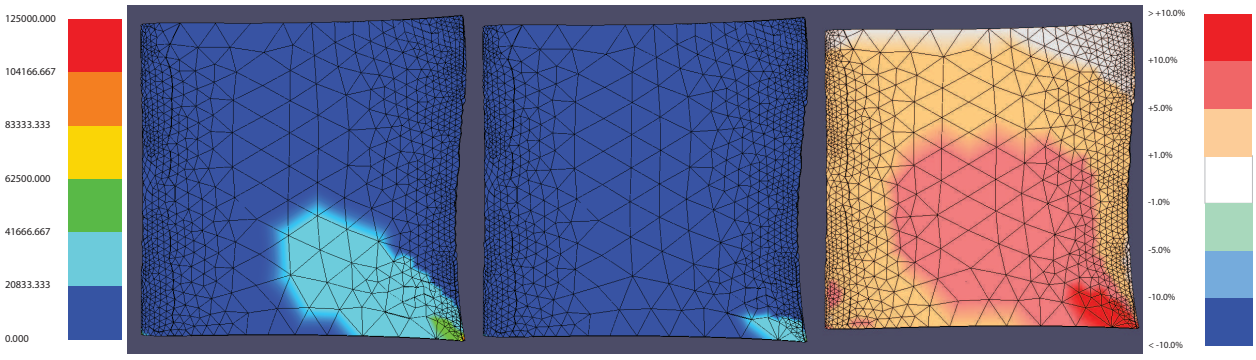

Figure 2.3: From left to right: the shared stress color scale, the new design, the nominal design, the difference model, and the difference model color scale.

The first design to be compared is the nominal case; that is, the solution when all three input parameters are set to the baseline value of 3.0 and is fixed. Hereafter, this will be referred to as the "nominal design." The second design is any new solution, obtained by giving the surrogate models a new set of input parameter values and rendering the new stresses for the model. Hereafter, this will be referred to as a "new design." The nominal design's stresses remain constant, but the new design's stresses update as users input new parameter values. This allows any new design to be compared to the nominal design.

A graphical user interface (GUI) was created to facilitate design comparison and consists of a rendering of the new, the nominal, and the difference models simultaneously on the screen using OpenGL canvases. All models use linked camera views for convenient comparison. Controls are provided for changing the values of the various parameters.

When mapping data values (such as results from an FEA process) to color values, one of the most fundamental choices is what color scale should be used [50–52]. This choice has great impact on all interpretations made about the visualized data. If each design uses its own color scale, mapped to the maximum and minimum value of that particular design, determining differences and comparing patterns between the two designs becomes problematic since the color mappings are not directly comparable [51]. Rather, because this research is focused on visualizing designs from an entire design space, a shared color scale is used based upon the maximum and minimum stress values that occur in the entire design space; that is, out of all possible designs in the design space, the maximum stress and the minimum stress that occur are used to create the range of possible stresses in the design space.

Figure 2.3 shows the canvases and color scales from the GUI. Even though the shared color scale helps make direct comparison between the two designs possible, it can still be difficult to discern the location and magnitude of differences between the two designs. Both designs have a large area of nodes with stress values that fall in the dark-blue range (0 to 20833 psi), but it is not clear whether or not the stresses have changed within this region, or by how much. It is entirely possible that a node's stress may have changed from one design to the next but remains in the same color range as before.

The difference model clarifies these changes. The third OpenGL canvas in the GUI is used to display the same geometry as the other two canvases, but here the primary data displayed are

the differences between the two designs. These differences are calculated using the stress values between each node of the two designs.

The difference calculations could be done in many ways, depending on the application. The simplest method is by a straight comparison, or the absolute difference (stress of Node X in the new design - stress of Node X in the nominal design), with the greatest increase in stress shown as a dark red coloring and the greatest decrease in stress as a dark blue. However, because every pair of designs compared will have variations in how different their nodal stresses are, this method exhibits the same pitfalls as local color scales do in the previous discussion.

This research shows the percent differences between nodes rather than the absolute differences. The percent difference method uses the shared scale for the design space that was previously determined. Once the amount a node has changed from one design to the next, or absolute difference, has been calculated, it is determined what percentage of the shared scale that change comprises. This can be described by Equation 2.1:

$$\Delta\sigma_x = \frac{\sigma_{x,new} - \sigma_{x,nominal}}{\sigma_{range}} * 100\% \tag{2.1}$$

where $\sigma_{x,new}$ and $\sigma_{x,nominal}$ refer to the stress at the same specific Node X in both the first and the second designs, and σ_{range} refers to

$$\sigma_{range} = max[d_1, d_2, ...d_n] - max[d_1, d_2, ...d_n] \tag{2.2}$$

with $[d_1, d_2, ...d_n]$ being the stresses at every design in the design space. This range encapsulates the highest and lowest stresses possible across the entire design space. For the purposes of this research, this was estimated by taking the highest and lowest stresses present in the training data for the surrogate models.

For example, if a particular node were to increase by 12,500 psi from the first to the second design, and the shared scale has a range of 0 to 125,000 psi, then this particular node's difference makes up 10% of the shared scale. An increase of 10% will be colored dark red, while a decrease of 10% will be colored dark blue. For the purposes of this study, any changes between -1% and 1% were chosen to be insignificant and thus are colored neutral white. The values of -1% and 1% are arbitrary and may be adjusted based on the needs of the analysis. This neutral range, as

well as any of the ranges in the color scale, may be adjusted according to the needs of the user. Differences between 1% and 10% are colored in hues that exist between the neutral and extreme colors, depending on if they are increasing or decreasing.

The color scales used in this study were designed to reflect traditional FEA color schemes. These colors, however, were not necessarily designed for people with color-deficient visual impairments [53]. The color schemes could be adjusted in future work to meet these requirements.

Figure 2.3 shows that the new design produces the greatest changes in the bottom right corner, with moderate changes in the center, and small changes along the trailing edge, tip, leading edge, and most of the root. There are also two small areas of moderate change visible near the bottom left corner. No significant change has occurred along the edge of the tip. Some of these details could possibly be inferred from a direct visual comparison of the two designs but now are presented clearly with far less effort and uncertainty.

2.4 User Experiments

An experiment was designed and conducted to validate the claims that adding an explicit encoding of the differences to juxtaposed views of structural results can improve a designer's speed and accuracy when performing certain design comparison tasks. Tasks include comparing general and specific changes, examining groups of elements and single elements on the objects, and responding to specific and more open-ended questions. Speed and accuracy are measured for each type of task.

A request for volunteers was sent via an email from the Brigham Young University Department of Mechanical Engineering to students in the mechanical engineering program - thus, the volunteers used in this study were all undergraduate or graduate mechanical engineering students. Volunteers were randomly assigned to one of two different groups: one group had access to the difference model, and the other group did not. Each group had 14 participants, for a total of 28 participants. The volunteers had no prior knowledge that there were two treatment groups.

The volunteers were self-selected, instead of randomly chosen, therefore this study cannot lead to statistical population inferences about the general population of engineers. However, because the treatment groups were assigned randomly to all volunteers, statistical causal inferences about the effect of using the difference model on user performance may confidently be made in

this study. That is, while this experiment's design does not allow for conclusions to be drawn about groups of engineers beyond those tested, it does allow conclusions regarding the influence of the difference model on performance.

2.4.1 Experimental Setup

Each volunteer was given the same introduction to the test. The test administrator reviewed basic concepts of stress analysis, airfoil parameters, and design space exploration with each volunteer. Those in the group without the difference model had only the new and nominal models (juxtaposed views) rendered in the GUI. For those in the group with the difference model, the concept of the difference model and associated color scale were also explained, and the GUI in their program included the difference model rendering with its associated color scale (hybrid views, see Figure 2.4).

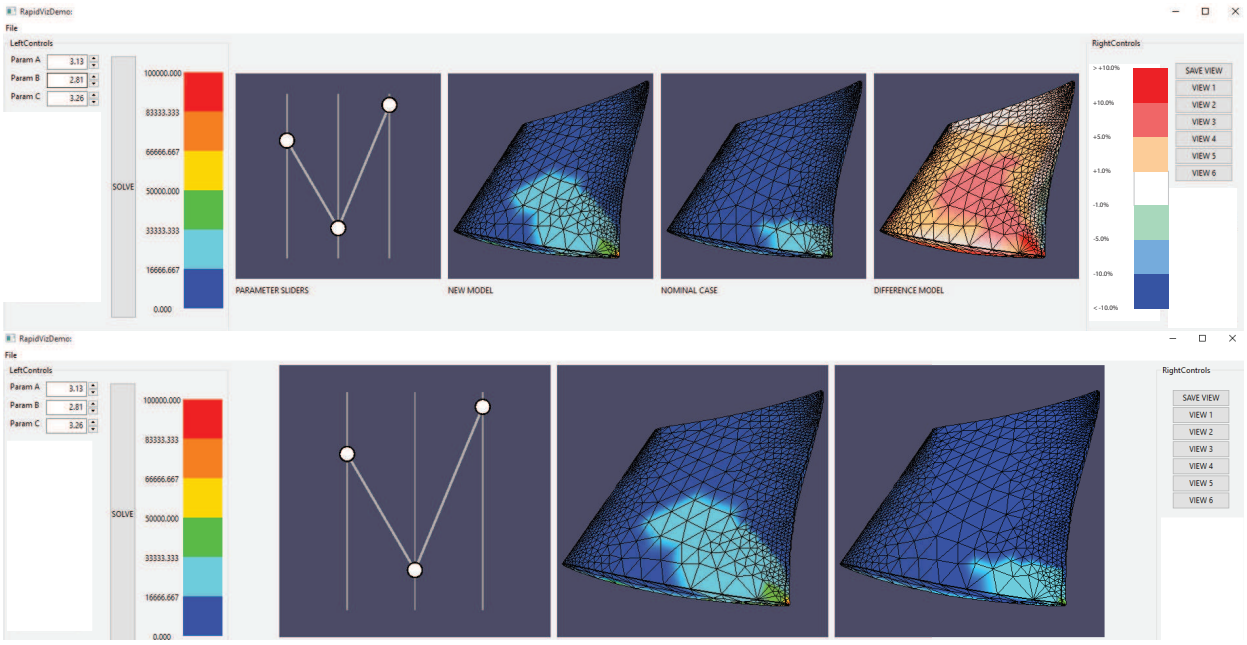

Figure 2.4: The User Test GUIs, with (top) and without (bottom) the difference model. From left to right: parameter value displays, stress color scale, sliders to control the parameter values, the new design, nominal design, and, for those with the difference model, the difference model and difference model color scale.

All participants were given the ability to input different values for the new model's parameter values. For tasks requiring precise inputs, values could be entered into text boxes, while for tasks requiring more exploration of the design space, slider bars could be adjusted. They had full ability to zoom, rotate, and translate the parts, as well as the ability to save different "views" or camera orientations of their own choosing for convenience as they looked at key areas of the blade.

The engine blade used in the study was the same simple parametric test blade described in the Methods section. For simplicity, the participants were not told what the parameter names were, but were given the pseudonyms Parameter A, Parameter B, and Parameter C. The questions were administered and the answers collected via Google Forms. Screen activity was captured visually with oCam software. The participants were given four tasks in order to assess the effect of using the difference model visualization. The order and the research objective of each task are detailed below:

1. **Task A:** An introductory task to familiarize participants with the test environment and controls. No data was collected during this phase.

2. **Task B:** A series of comparisons are presented to the participants along with three questions. The questions were designed to evaluate how the difference model enhances perception of specific kinds of differences between the designs:

   (a) **Q1:** Evaluates perception of general level of difference across the entire part

   (b) **Q2:** Evaluates perception of general level of difference across a specific area or feature of the part

   (c) **Q3:** Evaluates perception of a specific value difference for a specific location on the part

3. **Task C:** Participants are asked to creatively find new designs that differ from the nominal design according to specific criteria. This task was designed to evaluate how the difference model enhances not only perception but also creative design exploration.

2.4.2 Task A

This task familiarized the user with the program by asking them to enter new design points and prompting the user to try to notice changes in the two design points being displayed (the new design and a static nominal design). No data was gathered in this phase. The purpose was to help users with the learning curve associated with the program and served as a tutorial. While it did not introduce the user to every type of task they would be asked to perform, this introductory task helped them understand how to navigate and make sense of the program.

2.4.3 Task B

In Task B, participants were asked to answer three questions about each of eight different designs, one at a time. The objective was to analyze the accuracy and time spent on each set of questions. These times were measured from when each question was presented to the participant to the moment they clicked the button to proceed to the next question.

For each of the eight designs, the participant was asked to answer the following set of questions:

1. Question 1 (Q1): ”Using your best judgment, about what percentage of the surface of the NEW blade has INCREASED in stress?”

2. Question 2 (Q2): “Using your best judgment, about what percentage of the surface on/close to the TRAILING EDGE of the NEW blade has DECREASED in stress?”

3. Question 3 (Q3): “The max stress on the NOMINAL blade is 44724 psi. Is the max stress of the NEW blade greater than the max stress of the NOMINAL blade by 40,000 psi or more?”

For Q1 and Q2, multiple choice answers were provided, and for Q3 the user could answer “yes” or “no.”

Task B Setup

Preliminary testing revealed that the users spent much longer on the first few designs as they learned how to perform the requested tasks. After three question sets, the times and accuracy

became more consistent. Over the entire group of participants, the average of the first three question set times was 165 seconds, while the average of the remaining question sets was 85 seconds. This time discrepancy could be explained by users continuing their familiarization of the interface from Task A or taking some initial time to understand the problem. Consequently, it was decided that the initial three designs presented to the user should not be measured so as to not bias the data. In order to have five measured designs, then, eight total designs presented to the user: three unmeasured sets to accommodate a learning curve (Designs 1-3) followed by five measured sets (Designs 4-8).

The design values for question sets 4-6 were given by a random number generator, producing values from 2.7 to 3.3 for all three parameters in each design. The answers and designs for these questions were not known to the testers beforehand, so they represent a better test of the program features (as they could not be contrived to favor the difference model in any way).

The last two designs (7-8) were identical for every participant. This allows for direct comparison of speed and accuracy on the last two designs.

Task B Hypotheses

Since Q1 relies on perceiving increases in stress, it was hypothesized that using a hybrid visualization with the difference model would reduce the time spent on the question as well as reduce the error. It was designed to test perception of changes across an entire part.

Q2 had a similar hypothesis to Q1: that using a hybrid visualization with the difference model would reduce the time spent on the question as well as reduce the error. However, this question was designed to test perception of changes on a specific feature or area of a part.

Q3 deals with determining stress value changes between the two designs for a specific location. Because the difference model visualized percent differences without providing any additional information in the context of the actual stress values, it was hypothesized that using a hybrid visualization with the difference model would not improve the time spent on the question or the accuracy of answers over just the juxtaposed views. It was designed to test a comparative task involving specific values for a very small location on the objects being compared.

Task B Evaluation Method

In OpenGL, the surface of the model is broken into triangular faces. Each corner of the triangle has a specified RGB color, and the color across the face is a gradient between these corner colors. Thus each corner is "responsible" for coloring $\frac{1}{3}$ of the surface. Using this relationship, it was determined that, if a node was increasing or decreasing, it would produce a change on $\frac{1}{3}$ of the faces associated with that node.

For Q1, the nodes that were increasing or decreasing were identified, and then the surface area each node was associated with was totaled. This was used to produce the actual percentage of the blade that was increasing or decreasing. Because users selected a multiple choice answer, the actual values were rounded to the closest available choice presented to the users during the experiment, and the error between the two was calculated as $|response - actual| = error$. For example, if the actual percentage of the surface increasing in stress was 33.46%, and the answer given was 20%, then 33.46% would be rounded to 30% (also an available choice), and the user's error would be 10%.

For Q2, all the nodes on the trailing edge were identified (see the black region in Figure 2.5). Using a similar method to Q1, the percentage of the surface area of the trailing edge was calculated. Again, the actual value was rounded to the closest available choice presented to the users during the experiment, and the error between the two was calculated in the same way as in Q1.

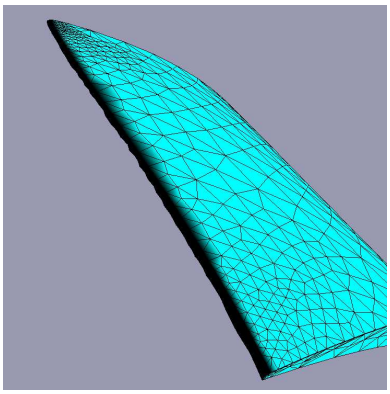

Figure 2.5: Trailing edge nodes discussed in Task B Q2 shown here in black

For Q3, the location of the max stress was found on the new and nominal models. It was then determined if the difference between the two was above 40,000 psi. Answers were given as "yes" or "no," so each answer received a score of 1 if correct, and 0 if incorrect. Thus, instead of average error of responses, this question measures the average accuracy of correct responses.

For each question, the time was measured in seconds from the moment the task was displayed to when the user clicked the "Next" button to progress to the next question. The time for each question as well as the total time spent on the question set were measured.

Since data was only gathered for the last five designs, this means that there were 15 speed values and 15 accuracy values per participant in Task B. With 28 participants, this produced 840 data points.

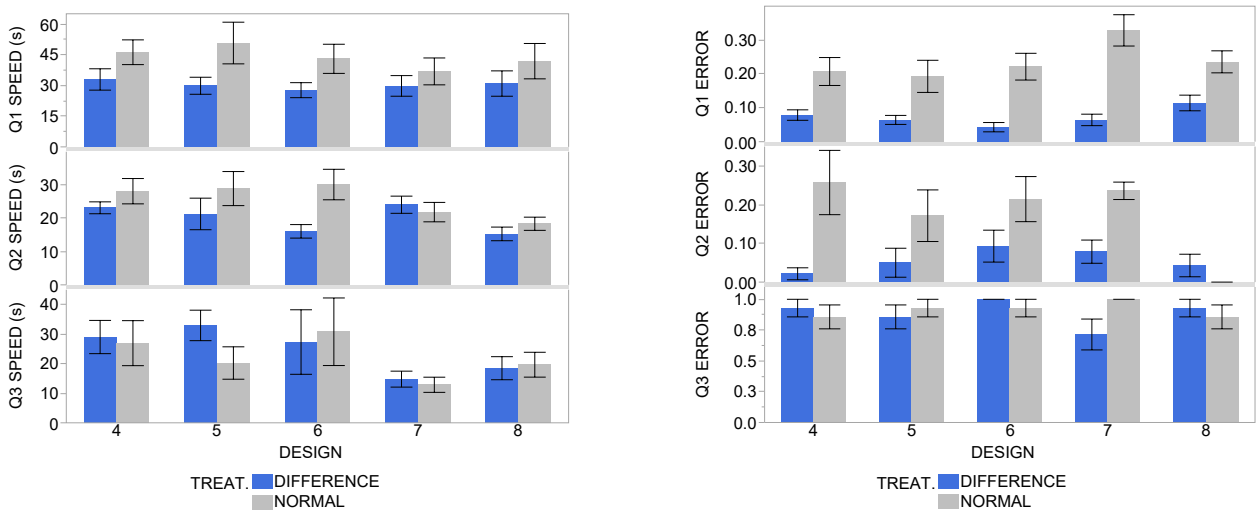

(a) Speeds

(b) Errors

Figure 2.6: Speeds and errors for Q1, Q2, and Q3 for each design in Task B. Each error bar is constructed using 1 standard error from the mean.

Task B Q1 Results

The speed and error for Q1 was calculated for the last five designs for each participant. Table 2.1 presents the mean values for each treatment group. Figure 2.6 shows a comparison of these values for each design presented to the participants.

Those using the difference model performed the task about 14 s faster than those not using the difference model (30.23 s vs. 43.73 s), and the average error from the correct answer was only 7.23%, about 16% lower than those without the difference model. By running standard two-sample t-tests, evidence was found that both differences were significant (p-values < 0.05). Thus, there is evidence that the null hypothesis can be rejected for both the speed and error, which suggests a correlation between using the difference model in a hybrid visualization and reduced time and reduced error for Q1. This statistical evidence supports the hypothesis for Q1.

Table 2.1: Task B Q1 Results

TREATMENT	AVG. SPEED	AVG. ERROR
DIFFERENCE	30.23 s	7.23%
NORMAL	43.73 s	23.71%
p-value	**0.0014**	**0.0001**

For reference, on Design 7, those using the difference model had 26.5% lower error, at about 7 s faster, on average (with only error being statistically significant). On Design 8, those using the difference model had 12% lower error, at about 11 s faster, on average (with only error being statistically significant).

Task B Q2 Results

The speed and error for Q2 was calculated for the last five designs for each participant. Table 2.2 shows the mean of these values for each treatment group. Figure 2.6 shows a comparison of these values for each design presented to the participants.

Those using the difference model answered the question about 5 s faster than those without the difference model (19.98 s vs. 25.49 s), and the average error from the correct answer was only 5.71%, about 12% lower than those without the difference model. Two-sample t-tests revealed that both speed and error, when used with the difference model, were significantly different (p-values $<$ 0.05) than those without the difference model. Thus, there is evidence that the null hypothesis may be rejected for both the speed and error, which suggests a correlation between using the difference model in a hybrid visualization and reduced time and reduced error for Q2. This statistical evidence supports the hypothesis for Q2.

A possible limitation of the data from Q2 is that there was ambiguity of what counted as the "trailing edge" in the question statement. While there was a fixed area used in evaluating the error, the participants had a less specific definition of the area given to them. Participants could have had different understandings about which nodes pertained to the trailing edge and which pertained to the blade faces and may have included more area than intended in their estimation of the trailing edge. This could affect how long participants spent on this question, as well as how accurate their answers were.

Table 2.2: Task B Q2 Results

TREATMENT	AVG. SPEED	AVG. ERROR
DIFFERENCE	19.98 s	5.71%
NORMAL	25.49 s	17.57%
p-value	**0.0146**	**0.0001**

For reference, on Design 7, those using the difference model had 15.7% lower error, at about 2 s slower, on average (with only the error being statistically significant). On design 8, those using the difference model had 4% higher error, at about 3 s faster, on average (neither being statistically significant).

It should be noted that while some designs could produce abnormal results (see Figure 2.6 for Design 7, Q2 Speed, and Design 8, Q2 Error), most of the values follow a predictable trend. The error at Design 8 is near 0.00 for those using the difference model, and is 0.00 for those without the difference model because it was easy to tell that Design 8 did not decrease anywhere, thus both groups did quite well at making a correct judgment.

Task B Q3 Results

The speed and error for Q3 was calculated for the last five designs for each participant. Table 2.3 shows the mean of these values for each treatment group. Figure 2.6 shows a comparison of these values for each design presented to the participants.

Those using the difference model performed the task 2.5 s slower on average than those without the difference model (24.49 s vs. 22.09 s), and the average correct responses given was about 2% lower. The two-sample t-tests revealed that there is no evidence that the differences are

statistically significant (p-values > 0.05), thus the null hypothesis may not be rejected, which does not suggest a correlation between using the difference model and the error and time spent on Q3. This task had little to do with making comparisons between designs, so this statistical evidence supports the hypothesis for Q3 that using the difference model in a hybrid visualization would not improve the results for this question.

The high degree of overlap in the standard error bars for Q3 in Figure 2.6 reflect the findings that there was little difference between the values for those with the difference model and those without.

Table 2.3: Task B Q3 Results

TREATMENT	AVG. SPEED	AVG. ACCURACY
DIFFERENCE	24.49 s	88.57%
NORMAL	22.09 s	91.43%
p-value	0.5722	0.5764

For reference, on Design 7, those using the difference model had 29% lower accuracy, at about 2 s slower, on average (with accuracy being statistically significant). On Design 8, those using the difference model had 7% lower error, at about 1 s faster, on average (neither being statistically significant).

2.4.4 Task C

In Task C, participants were asked to find three separate new designs (combinations of parameter values) that produced at least one location that decreases 10% (10,000 psi) *and*, concurrently, at least one location that increases at least 10% (10,000 psi) from the nominal design. The purpose was to have the participants use the tools to find designs on their own that met specific conditions.

Task C Hypothesis

It was hypothesized that those using the difference model in a hybrid visualization would spend less time on the task and recommend more designs on average that successfully fulfilled

both conditions. This task was designed to test a more open-ended comparative question involving interactive design space exploration as well as perception of trends in change across an entire part.

Task C Evaluation Method

The three different designs recommended by the participants were entered into the program. In the program, a simple function identified if there were any nodes on the new model that had increased or decreased 10% from the nominal model. If there were at least one increasing and one decreasing node that met the conditions, that design was given a score of 1. If it did not meet both conditions, then it received a score of 0. For each participant, the scores were averaged together, e.g., if two of the three designs were correct, the average score would be 0.66.

The time was measured in seconds from the moment the task was displayed to when the user clicked the "Done" button. The responses were measured for speed of the overall task and accuracy.

Task C Results

Table 2.4 shows the average correct designs provided in Task C for each treatment group. Those using the difference model performed the speed on average 91 s faster than those without the difference model (303.62 s vs. 395.00 s), and provided 31% more designs that met both conditions. The two-sample t-tests showed little evidence of there being a statistically significant difference for the speeds (p-value > 0.05), but showed more convincing evidence that the average correct answers given were significant (p-value < 0.05). Thus, there is little evidence that the null hypothesis may be rejected in regards to the speed, which does not suggest a correlation between using the difference model in a hybrid visualization and reduced speeds for Task C. However, there is more convincing evidence that the null hypothesis may be rejected for the average correct answers given, which does suggest a correlation between using the difference model in a hybrid visualization and higher average correct answers. This statistical evidence supports the hypothesis for Task C in regards to accuracy but not for speed.

Table 2.4: Task C Results

TREATMENT	AVG. SPEED	AVG. CORRECT
DIFFERENCE	303.62 s	97.64%
NORMAL	395.00 s	66.79%
p-value	0.1410	**0.0001**

2.5 Discussion

Overall, the results of the experiments show that, for many types of design activities, there is evidence that the difference model does improve the speed and accuracy of users' abilities to perceive differences between the structural results of two different FEA models. However, the type and purpose of the comparison that a user is making affects how much including the difference model in the visualization will assist the user in their efforts.

The results of Task B Q1 suggest that the difference model can improve the speed and accuracy with which a user can assess the general level of difference across an entire part. For comparison tasks where general differences are important, visualizations with a difference model can help users accomplish the task faster and with greater success than visualizations without a difference model (e.g., "Is the new design generally *higher* or *lower* than the original design?", "Is more than half of the new design *higher* (or *lower*) than the original design?" or "How much of the new design has *higher* (or *lower*) values than the original design?").

The results of Task B Q2 similarly suggest that the difference model can improve the speed and accuracy with which a user can assess the general level of difference on a specific area or feature of a part. Visualizations with a difference model can help users make comparisons about differences across an area of a part more quickly and with greater success than visualizations without a difference model (e.g., "Does the top edge of the new design generally have *higher* or *lower* values than the original design?", "Is more than half of the front face of the new design *higher* (or *lower*) than the front face of the original design?").

Confirming the hypotheses for Task B Q1 and Q2 suggests that, for future visualization designs when comparing two intricate designs or objects, like these three-dimensional compressor blades, a difference model will help the users perceive not only the general trends between the two objects, but also the trends over specific areas. Users can estimate the type of changes and the level of change with greater success and speed.

The results of Task B Q3 differ from the previous two by suggesting that the difference model does not improve the speed or accuracy of the user's responses when making comparisons that involve specific values at specific locations on the two designs. This was an expected result; it was included in this study in order to confirm the limitations of the difference model in aiding perception. For tasks where users compare specific values instead of trends across the part or areas of the part (e.g., "Which locations of the new design have a value higher than X?", or "What is the lowest value between these two designs?), there is no evidence that the difference model improve the users' performance. While the difference model will likely not negatively impact the speed and accuracy of the users' perception, it is not worth the cost of adding a difference model to a visualization if these types of comparisons are the only ones being made with the tool.

Finally, the results of Task C suggest that, when a user is asked to creatively explore a design space and find designs that satisfy certain comparison-based criteria, the difference model helps improve the success but not the speed with which users can accomplish the task. These types of tasks are more open-ended than the previous comparison tasks in that they there are many possible answers that will satisfy the question. The open-ended nature of this activity shows that the difference model helps users in both directed and creative comparison tasks. This is more in line with the objectives of design space exploration (e.g., "Find a design that generally has *higher* or *lower* values than the original design." or "Find a design where at least half of the front face is lower than the original design"). To take advantage of the benefits of the difference model in creative design space exploration, the types of objectives should be similar to the types of comparisons examined in Task B Q1 and Q2. While users may not necessarily complete these tasks faster with the difference model, they will have an easier time in correctly identifying target designs.

2.6 Conclusions

This research presented a method for enhancing design comparison in a design space for three-dimensional designs with structural results. This method allows the difference in the nodal results to be calculated and shown to the user with an explicit encoding visualization, which clarifies the location and magnitude of differences between the two designs. When combined in a hybrid visualization with a juxtaposition of the two designs being compared, the context of the difference relationships is preserved and allows the user to make better use of the information.

This was illustrated by comparing the stress responses of turbomachinery compressor blades to different geometric parameter inputs. Stresses that increase from the first to the second design are colored in warm colors, while stresses that decrease are colored in cool colors.

The difference model gives users a way to perceive the differences between the structural results of two designs more quickly than by simple visual inspection of the two designs. It also allows users to understand the magnitude of the differences with greater accuracy than making mental calculations between the two juxtaposed views of the designs. The user can thus understand key aspects of a design comparison more quickly and more accurately, enabling more effective design space exploration.

An experiment was designed to validate the claims that using a difference model combined with juxtaposed views of the designs would improve speed and accuracy for certain design comparison tasks. This study helps provide justification for useful visualizations in future applications of surrogate modeling of three-dimensional structural results for design space exploration

Based on the experiments' results, there is evidence that including the difference model allows a user to perform tasks such as detecting changes across the entire part more quickly and more accurately. Including the difference model also helps users detect changes on a specific feature or area more accurately and more quickly. Finally, there is evidence that including the difference model can help a user suggest designs that meet specifications in an open-ended design comparison task more accurately than without the difference model, but not necessarily more quickly.

When tasks focus on a specific value rather than a decontextualized relationship (or difference), then the experiments did not yield evidence that the hybrid visualization produces advantages in speed and accuracy. This is because the difference model is designed to show the difference between two designs' values rather than the actual values at that point, and thus gives no advantage in this activity.

Further research and experimentation should be done to determine how users utilize the difference model to perform tasks. More focused tests could be done with industry experts (instead of student volunteers) to see how the tool helps with more concrete design tasks that are performed in a professional setting. Further research should also be done in how user perception is affected by displaying result values other than stresses on a part, such as using the difference model to display

differences in displacement, temperature, or vibration. While the principles ought to remain the same for any type of results, perhaps other useful adaptions could be developed to enhance perception in specific applications. Useful future studies could determine how the usefulness of these methods change as the geometry of the comparison objects becomes more complex.

CHAPTER 3. GOODMAN

In jet engine design, fatigue life is an important consideration. Fatigue life analyses involve the results of both structural and modal analyses. Because these simulations are expensive, they are difficult to incorporate into rapid and thorough design space exploration. Therefore, this chapter predicts both steady and alternating stresses with FF surrogate models for a jet engine compressor blade in answer to RQ1. These results are then mapped onto a Goodman diagram for analysis, and used to predict the fatigue risk of the part. This Goodman diagram is similar to those used in industry, except that it can dynamically update as design parameters are changed.

The study fulfills RQ2 by comparing the results on the Goodman diagram with those predicted by SV surrogate models. First, the accuracies are compared, and it is found that FF surrogate models predict the values with a greater degree of accuracy. Then, the various design benefits (RQ3) of predicting nodal responses are described. These include the ability to make judgements with more detailed information, mapping results from the Goodman diagram onto actual three-dimensional geometry, and the ability to make less conservative predictions. With the interactive speeds of the FF surrogate models, the results of the Goodman diagram can be used in real-time design space exploration.

This study of FF surrogate models differs from the other chapters by using the results from two separate FF surrogate models, each predicting a different kind of data i.e., instead of just using FF surrogate models of the steady stresses, here FF surrogate models of the steady and alternating stresses are used together. These two predictions are used in a calculation to achieve a new predicted result at every node. This combination and calculation of new nodal results is not uncommon, and must be tested in order to confidently recommend FF surrogate models for this type of engineering activity. The comparison between FF and SV surrogate models is done based on the predictions of a maximum stress value. Although FF surrogate models can predict much

more information, this maximum calculated value provides a good point of comparison for the two surrogate models.

This study has been submitted for publication to the *Journal of Mechanical Design* and is under review as of July 2019.

3.1 Introduction

In aerospace design, parts often fail due to fatigue from cyclical loadings [20, 54]. Failures related to vibration account for approximately 60% of aero-engine failures, making the ability to accurately and easily predict failure due to fatigue an essential part of the jet engine design process [1]. As aerospace products become more complex, there is an increasing need for more detailed predictions of risk and failure at earlier design stages [55–57]. If this risk can be addressed earlier, time and money that would otherwise be spent on redesign activities can be avoided.

Fatigue life analyses and the prior analyses they depend on can be computationally expensive, which can affect early design practices. Conducting a fatigue life analysis requires results from both structural and modal analyses [1]. For this reason, fatigue life analysis must be considered after these analyses have already been performed. The costs for obtaining highly detailed results in these analyses can cause limits in speed and quality for iterative processes in early design space exploration. To compensate, often fewer iterations are explored or lower fidelity models are used [14]. The level of effort to adjust and properly set up these simulations and analyses often also makes them less useful tools for free exploration of design variations in real time [26].

The Goodman diagram (otherwise known as a Haigh diagram [19]) is a common tool used for predicting the fatigue life of a structural part under cyclical loading [20, 21]. An example of a Goodman diagram is show in Fig. 3.1. The diagram consists of a material limit line (usually the modified-Goodman line) and plotted points determined by the steady and alternating stresses on the part [20]. A point's location relative to the material limit line provides an indication of the expected fatigue life of the part [58–60] Many different measures exist for quantifying the risk of fatigue failure of a part based on a point's proximity to the material limit line. Some of these will be explained in Section 3.2.2, including a specific value called the percent Goodman (%G).

There are various methods for representing a part on the Goodman diagram. Commonly, the entire part is represented with a single point. In the authors' experience, one conservative

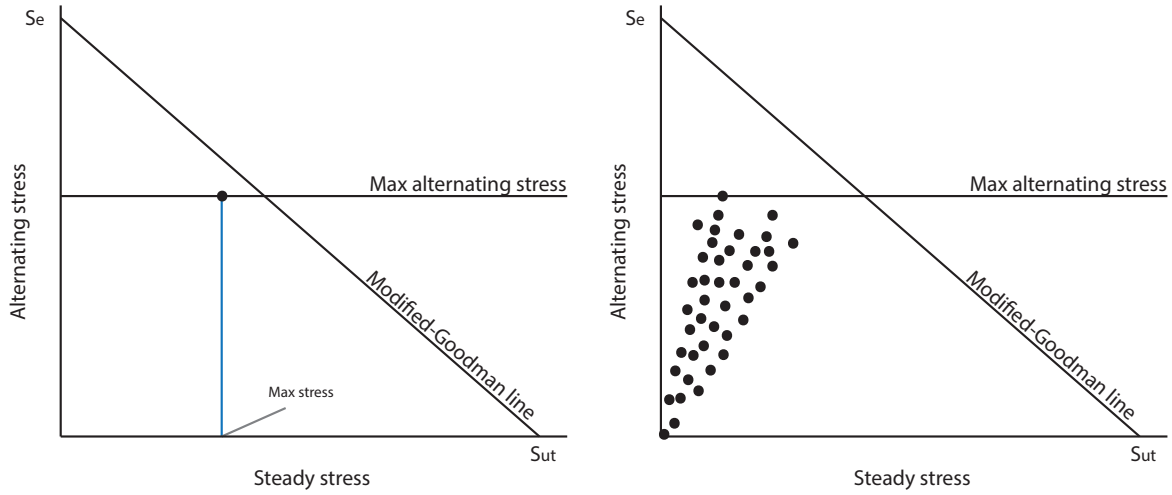

(a) Single point (b) Multiple points

Figure 3.1: A single point (3.1a) and multiple points (3.1b) on a typical Goodman diagram

method involves plotting a point using the maximum steady stress and maximum alternating stress of the part (Fig. 3.1a). This produces a conservative estimate of the part's fatigue life. Conservative methods such as this have been sufficient in the past; however, as performance requirements become more demanding, more complete understanding of the exact blade response is necessary [2]. A more detailed method for representing a part uses a collection of steady and alternating stress values from different locations on the part (Fig. 3.1b). This provides a more complete understanding of the fatigue response of the entire part, but also requires a greater knowledge of the part and more computational resources [61–65].

To overcome the high computation time needed for expensive simulations, many aerospace designers employ surrogate models to quickly predict needed information [7, 8, 66]. Surrogate models (also known as response surface models, metamodels, regression models, or emulators [6]) create simplified mathematical relationships between inputs and outputs of a system. Using new inputs, the surrogate models can then predict results in a fraction of the time it takes to perform the original calculations [6]. In the design process, properly trained surrogate models provide fast, reliable, and computationally cheap predictions of certain values. These are often used in design space exploration or optimization routines in order to quickly determine the best initial designs without the prohibitive cost of full simulations [7–9].

Most surrogate models have been limited to mapping a set of inputs to a single output value [7]. These types of surrogate models will be referred to as single-value (SV) surrogate models. SV surrogate models have been used to predict many values used in fatigue analysis and Goodman diagrams, such as the maximum stress on a part or the stress at a specific location of interest on a finite element model [7, 67]. These SV surrogate models make it possible to quickly update a single point on a Goodman diagram and make general fatigue life assessments without expensive simulations. However, because these surrogate models are only used to predict a single-value, they cannot give a more detailed prediction of the steady stresses and alternating stresses of a part [7] or their full representation on a Goodman diagram.

In recent years, methods for using surrogate models to quickly predict the behavior of every node in a FEA simulation have been developed [14–16]. They generally take geometric or loading parameter values as inputs and predict FEA results for the entire part. Results can be predicted in real time, i.e., under a second for updates to appear [14]. These results can be used to dynamically update three-dimensional visualizations of the nodal results mapped onto the part geometry as input parameters are varied. Thus it becomes possible to predict much more detailed structural FEA results than with SV surrogate model methods. Because this use of surrogate models predicts an entire field of responses from a FEA simulation, these will be referred to as full-field (FF) surrogate models.

This research uses FF surrogate models to predict the results of structural and modal FEA results for compressor blades in a jet engine. For demonstration purposes, the transonic Purdue blade will be used [18]. In addition to being able to predict and visualize geometric changes and stress responses on the part, these FF surrogate models provide real-time updates of information that can be used to plot each node's location on a Goodman diagram. Thus, this research will present a new dynamically updated Goodman diagram that shows a highly detailed fatigue response to changes in geometric inputs for a given compressor blade. This dynamic Goodman diagram will be discussed in the context of early stage design exploration. By obtaining the results in real-time, this Goodman diagram can replace otherwise expensive simulations and avoid the associated time-delay impacts on the design process.

The use of FF surrogate models to update every node's location on the Goodman diagram will be compared to using a SV surrogate model to update a single conservative point on the

Goodman diagram at the maximum stress location. Specifically, the accuracy with which each type of surrogate model predicts the %G values will be examined with reference to a full FEA solution. It is suggested that using every node's results on the Goodman diagram to evaluate the fatigue life of a part (instead of using a single point at a maximum stress location) reduces conservatism and makes it easier to find better designs.

With the accuracy of the FF surrogate method established for predicting %G values, a number of design benefits are presented for design exploration and fatigue analysis. Using the FF surrogate models allows designers to quickly predict and gain intuition about fatigue analysis results, including %G values and other Goodman-based measures. The predicted results can be used to create visualized results, such as a dynamic Goodman diagram that updates in response to parameter changes in real time. Extending previous FF-enabled visualizations, the %G values for each node may be mapped onto the three-dimensional visualization of the part. Because the information is mapped directly onto three-dimensional geometry, the visualization can be used to indicate actual locations of potential fatigue failure. The speed at which information is available to the designer leads to faster intuition development. New result predictions and visualization updates occur in direct response to parameter changes, making the relationship between the design changes and the fatigue analysis results easier to perceive. Finally, this method makes it possible for more detailed fatigue analysis results of new designs to be presented simultaneously with structural and modal analysis results without the need for expensive simulations and delays.

The main contributions of this paper are the use of FF surrogate models to predict steady and alternating stresses (and therefore calculate the %G values), a study on the accuracy of these FF surrogate models as compared to the accuracy of SV surrogate models, and the presentation of the associated design benefits with this method. The interaction between the changing inputs and the dynamically-updated visualizations in real time provides a key improvement over static methods. In addition, this work introduces the terminology of SV and FF surrogate model methods in order to provide a distinction between these different uses of surrogate models.

3.2 Previous Work

3.2.1 Surrogate Models of FEA Results

Finite element analysis is a common engineering tool for obtaining the structural response of a part to a set of boundary conditions. FEA programs, such as ANSYS, take a two or three dimensional model of a part and convert it into a finite element mesh made up of many interconnected nodes. The program determines the response to the loads by solving equations at each node of the mesh. These responses could include values such as the amount of physical displacement or stresses caused by the loads on the part.

SV surrogate models have been used to predict specific values obtained from FEA results, such as the maximum stress of a part, the average displacement of a part, or the response of a specific "monitor" node at some location of interest [7, 67]. The inputs to such a surrogate model could be geometric parameter values or loading parameters used in the FEA simulation. Instead of running a new simulation each time a design change is made, a properly trained SV surrogate model can provide the desired output value very quickly without expensive calculations. These types of surrogate models have been used to predict results in many parts of a jet engine, from the compressor and turbine blade vibrations to aerodynamic coefficients experienced in various nozzle types [8, 11, 66, 67]. It has been noted that when these single values results have a relationship to other points on a part, they can be more difficult to accurately predict with SV surrogate models [7].

In recent years, FF surrogate models have been used to predict the stress and displacement response of each node in a FEA model [14–16]. Both Heap et al. and Bunnell et al. were able to approach "real-time" design exploration of finite element models by representing the nodes of the model with surrogate models. Bunnell et al. used a unique surrogate for every node's stress, x location, y location, and z location in order to predict the stress and geometric response to changes in geometry. Both Schulz et al. and Bunnell et al. were able to predict the results for an entire three-dimensional finite element model.

To properly train FF surrogate models, it is necessary to first create a set of training data. To do this, a variety of designs are created by using a design of experiments (DOE) to find a set of designs that adequately fill the design space. For a model with n input parameters, each design instance is represented by an n-dimensional vector [26]. Once the designs in the DOE

are identified, FEA is performed for each design and the nodal results are collected. The input parameters for the designs and the FEA results are the training data used to train the surrogate models: the input parameter values for each design constitute the inputs to the surrogate models, and the nodal FEA result values constitute the outputs. There exists one unique surrogate model for every node. The surrogate models are trained on these inputs and outputs, and develop a mathematical function to describe the relationship between them. Once trained, the surrogate models can then take in new parameter values to predict the full nodal response of the part.

This process allows a designer to change the parameters of a model and obtain visualized results. The prediction and visualization occurs more quickly than waiting for the full FEA process [14]. The speed of this method is dependent on the type of analysis performed and the complexity of the model, with node count and the number of input parameters playing a significant role in training and solving time. Under the right conditions, the time for calculating and visualizing these results is under a second, which may be considered as a real-time update [14]. With near-instant structural results, a designer can conduct more complete design space exploration and quickly visualize variations of the design without running independent simulations for every new model.

3.2.2 The Goodman Diagram

The Goodman diagram uses steady and alternating stress values to help determine the fatigue life of a part. These stress values determine the position of points on the diagram. The steady, or mean, stress of the part is plotted with respect to the independent axis, and the alternating, or vibratory, stress is plotted with respect to the dependent axis (see Fig. 3.1) [68]. A material limit is represented on the diagram by the modified-Goodman line, typically drawn from the ultimate strength (S_{ut}, on the independent axis) to the endurance strength (S_e, on the dependent axis) [60]. A point's location relative to the modified-Goodman line provides an indication of the expected fatigue life of the part. The equations and processes for creating a Goodman diagram are well documented in the literature [20, 49, 60].

Acceptable designs are those points in the region bounded by the modified-Goodman line and the max alternating stress of the part [59]. Designs that fall into this region are expected to have infinite life (generally 10^7 cycles or greater) [58–60, 68]. Points that fall above this line are determined to be at risk of failing in fatigue [60].

The Goodman diagram has been used for fatigue analysis in a wide variety of applications, including jet engines [19–21, 54, 64, 68, 69], subway frames [61, 62], coronary stents [63, 65], and repair of underground copper cables [70] or aging railroad bridges [60]. The points on the diagram may be obtained from finite element analysis [61–63, 65], physical testing [21, 59, 71], or a combination of these.

The steady and alternating stresses may be obtained from measured test values or from structural and modal FEA simulations. Generally, the mean structural centripetal stresses in rotating components of a jet engine may be calculated with relative certainty because they are closely related to the rotational speeds of the engine, which are relatively straightforward to measure [19, 59, 72]. The alternating stresses, however, tend to be more difficult to accurately calculate because they depend on the vibratory characteristics of the engine, which are more varied and difficult to measure [19].

In design, modal analysis is performed to obtain mode shapes and natural frequencies of the component [68, 73]. An estimate of the mode's stresses may then be obtained from these mode shapes [67, 68]. Because modal analysis results are based on solving a finite element eigenvalue problem, the mode shapes are unscaled. Often, finite element solvers will scale the mode shapes (and the resultant mode stresses) either to unit magnitude or normalize them by the mass. In practice, the mode's stresses are often preliminarily scaled based on previous design analyses, where the maximum value is set to a known max alternating stress, independent of the mean stress [20, 59].

In absence of a known max alternating stress, then an assumed max alternating stress may be used. A designer may choose a value for the max alternating stress based on experience, experimental results, or by determining what additional stress would cause failure when added to the already present steady stress [19, 20, 59].

Many efforts have been made over the last 20 years to reevaluate and improve the use of Goodman diagrams. Some studies [20, 59, 68] have acknowledged that the Goodman diagram neglects certain aspects of design, such as prestress or the effect of damage during operation on fatigue life, and adjustments that address these concerns have been made. Others have attempted to display the information on a Goodman diagram in ways that better convey information to the designer, such as using three-dimensional Goodman diagrams or non-linear material failure

lines [20, 59, 68, 74]. A common observation is that, when using certain assumptions or presentation methods, the Goodman diagram may produce results that are overconservative or obscure important design considerations [59, 75]. Some researchers have found that certain assumptions or methods, such as using minimal material capabilities or neglecting the difference between prestress and mean stress, can lead to overconservative results on the Goodman diagram.

Methods of Evaluation

Rather than evaluating fatigue life solely based on which side of the modified-Goodman line a point is located, it may be important to have a quantitative measure of how close a point is to the line. For this reason, there are a variety of methods for evaluating the proximity of a point to the modified-Goodman line. These values quantify the amount of risk for a given point.

Factor of Safety: Like in many engineering applications, using a factor of safety in the Goodman diagram is a useful way to embed security against physical flaws or unforeseen events into the design. Common examples of the factor of safety in the Goodman diagram include using a fixed fraction of the steady and alternating stress [20], a fixed value of alternating stress [20], or some fraction of the distance of a point to the Goodman line [20, 59, 63, 70, 75].

Generally, the factor of safety is determined along the load line (drawn from the origin to the Goodman line through the point), given by Eqn. 3.1 [75], where n is the factor of safety, σ_a is the alternating stress, and σ_s is the mean stress. This effectively applies the factor of safety to both the steady and the alternating stresses.

$$\frac{1}{n} = \frac{\sigma_a}{S_e} + \frac{\sigma_s}{S_{ut}} \tag{3.1}$$

This value may also be given by a ratio along the load line, or Eqn. 3.2 [70], where AB is the distance from the origin to the point, and AC is the distance from the origin to the Goodman line through the point. This is illustrated in Fig. 3.2.

$$\frac{1}{n} = \frac{AC}{AB} \tag{3.2}$$

In some situations, when one of the types of stress (steady or alternating) is more certain than the other, the factor of safety can be applied only to the uncertain term [59]. Because predic-

tions of the mean stress in a jet engine are fairly certain while alternating stresses are not, a factor of safety can be applied to the alternating stress only. Thus a vertical factor of safety may be used in some cases [20], using Eqn. 3.3. This can be illustrated in Fig. 3.2 (see also Fig. 2 in [59]).

$$\frac{1}{n} = \frac{ED}{EB} \tag{3.3}$$

Percent Goodman: The percent Goodman is another way to measure the vertical distance of a point to the modified-Goodman line, and is calculated by Eqn. 3.4, where %G is the percent Goodman value. It is actually the inverse of the vertical factor of safety (Eqn. 3.3), as seen in Fig. 3.2. It describes the ratio of a point's alternating stress over the value of the Goodman line at the steady stress of that point. In other words, the percent Goodman calculates the height of the point as a percentage of the Goodman line for a specific steady stress value.

$$\%G = (\frac{1}{n})^{-1} = \frac{EB}{ED} \tag{3.4}$$

This type of inverse factor of safety is helpful when reporting the risk of a part in reference to the Goodman diagram. Rather than a factor of safety, it provides a percentage, where values greater than 1 are above the Goodman line and at risk of fatigue failure [49, 63]. This helps facilitate easier comparisons of values during design space exploration. For this reason it will be the chosen method of evaluation for this research. %G $< 100\%$ indicates a point within the modified-Goodman line, %G $= 100\%$ indicates a point directly on the modified-Goodman line, and %G $> 100\%$ indicates a point over the modified-Goodman line.

Methods of Part Representation

As stated previously, the Goodman diagram can be used to represent a part with a single point or many points. A common single-point method is to evaluate the Goodman diagram using Eqn. 3.5, where $\sigma_{s,max}$ is the max mean stress and $\sigma_{a,max}$ is the max alternating stress. An example of this value is shown in Fig. 3.3 with a grey square. Any %G values obtained by using both the max alternating and max steady stress (from Eqn. 3.5) are hereafter referred to as $\%G_{max}$.

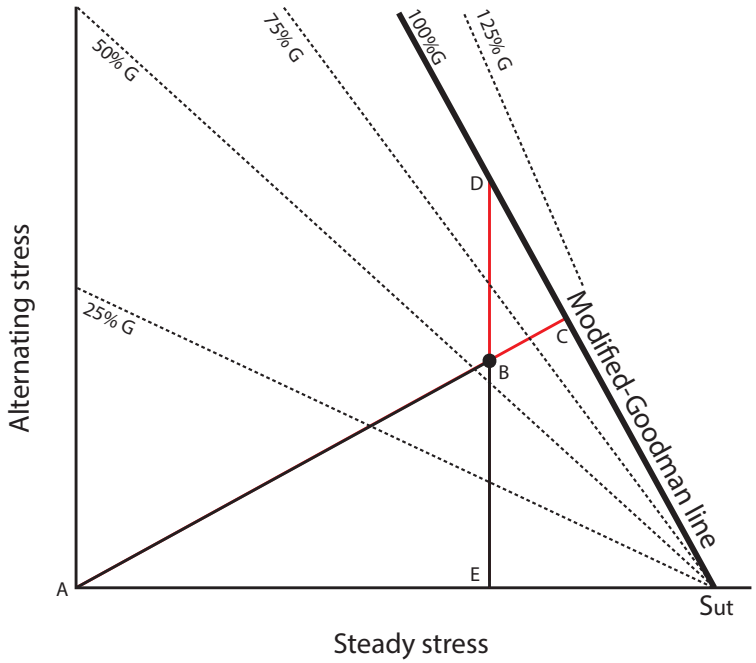

Figure 3.2: Methods of evaluating a point's proximity to the modified-Goodman line. The point (B) is at about 55% for the percent Goodman (Eqn. 3.4) and about 90% for an inverse diagonal factor of safety

$$\%G_{max} = \frac{\sigma_{a,max}}{S_e - (\frac{S_e}{S_{ult}})\sigma_{s,max}} \tag{3.5}$$

When stress data from many different locations on a part are available, this collection of points may be represented on a Goodman diagram together [61–65]. This method has been used with the factor of safety method described in Eqn. 3.2 for each node in an FEA model [63,65]. The values obtained give an indication of the risk of fatigue failure at different nodal locations across the model [61,62,65]. Representing every node's location on the Goodman diagram has been used in a wide range of applications, including FEA models of compressor blades [64].

By representing each node's position on a Goodman diagram, the percent Goodman may then be evaluated for every point on the diagram using Eqn. 3.6, where $\sigma_{s,nodal}$ is a node's mean stress and $\sigma_{a,nodal}$ is a node's alternating stress. Using Eqn. 3.6 for every point gives the designer a more complete understanding of how the whole part will respond in regards to fatigue life. The %G value for any node in this collection will be referred to as a $\%G_{nodal}$ value.

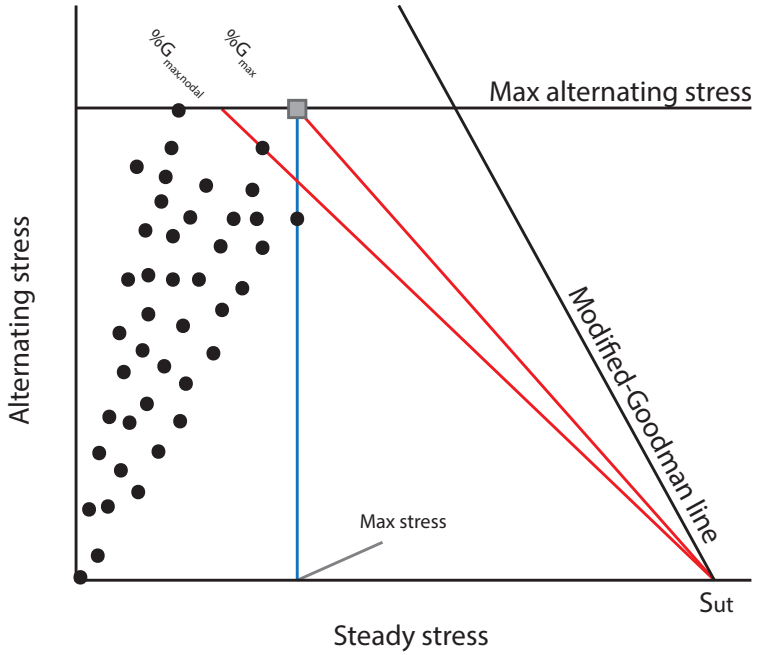

Figure 3.3: The $\%G_{max,nodal}$ and the $\%G_{max}$ shown on a Goodman diagram. The $\%G_{max,nodal}$ is lower than the $\%G_{max}$ in this example.

$$\%G_{nodal} = \frac{\sigma_{a,nodal}}{S_e - (\frac{S_e}{S_{ult}})\sigma_{s,nodal}} \tag{3.6}$$

A value of interest in this set of $\%G_{nodal}$ values is the maximum %G value. This will mark the node (and location) on the part that is at most risk of failing in fatigue. It is found using Eqn. 3.7. Every node's percent Goodman value will be calculated and represented on the Goodman diagram, but this highest value provides a less conservative summary value for the risk of a part than the $\%G_{max}$, and will be referred to as $\%G_{max,nodal}$. The presentation of the $\%G_{max,nodal}$ is compared to the $\%G_{max}$ value in Fig. 3.3.

$$\%G_{max,nodal} = max(\frac{\sigma_{a,nodal}}{S_e - (\frac{S_e}{S_{ult}})\sigma_{s,nodal}}) \tag{3.7}$$

Both the $\%G_{max}$ and $\%G_{nodal}$ values give %G values on the Goodman diagram, but are obtained using different data and have different merits. Because the $\%G_{max}$ is calculated using the maximum possible stress values, it is the most conservative estimate of the fatigue life. In some cases, this may be preferred, but it also may result in overdesign [69]. While safety is paramount

in aerospace designs, even slight excesses in weight and cost have significant ramifications for companies. The $\%G_{max,nodal}$ provides a less conservative estimate of the fatigue life. This value tends to be lower than the $\%G_{max}$, indicating that the $\%G_{max}$ usually predicts a part is closer to being at risk of failing in fatigue than it really is. This can lead to design compensations that may be unnecessary or rejection of designs that are actually safe. Beyond providing the $\%G_{max,nodal}$ value, using the full collection of $\%G_{nodal}$ values can display the complete "shape" of the part's data on the Goodman diagram, which may reveal further insights to the true fatigue life of the part. While calculating the $\%G_{max,nodal}$ requires more detailed knowledge of the part's stress responses, using $\%G_{nodal}$ values also enable further benefits during the design process, as noted in Section 3.6. In contrast, the $\%G_{max}$ can only provide a single representative value.

3.2.3 Rendering Goodman Diagram Values onto Three-dimensional Geometry

Mapping values from the Goodman diagram onto three-dimensional part geometry has been used in many applications. Certain failure evaluation values from a Goodman diagram were used to map risk of failure to different locations of wheel trucks in railroad FEA models [61, 62]. Several studies have analyzed the strength of coronary stents by displaying the position of every node's stress reaction on a Goodman diagram, and using a factor of safety evaluation for each point [63, 65]. One study then mapped these factor of safety values directly onto the three-dimensional geometry [63]. Day et al. did a similar mapping for the Goodman factor of safety onto every node in a FEA compressor blade model [64]. These visualizations helped to indicate locations that could be modified so the part would have a more desirable fatigue response.

These visualizations are useful, but require computationally expensive analyses. When modifications and design changes are made to the model, new analyses must be performed in order to acquire the information needed to update the visualization. Previous studies have not employed surrogate models for these nodal results and geometry-based visualization, and thus have been tied to these expensive simulations.

3.3 Method

This section will discuss the manner in which surrogate models can be trained to predict steady and alternating stress values. The results of these surrogate model predictions can then be plotted on a dynamic Goodman diagram. These points are then used to calculate %G values. The various calculations are presented, followed by the workflow for creating the surrogate models. The accuracy of FF and SV surrogate models with regards to the calculated %G results is presented in the following section.

3.3.1 Surrogate Models

As seen previously, the $\%G_{max}$ depends on the max steady stress and a max alternating stress. In the absence of a known value for the max alternating stress, a value may be assumed (see Section 3.2.2). The $\%G_{max}$ value may be calculated using a SV surrogate model that predicts the maximum steady stress, hereafter referred to as a SV $\%G_{max}$ value (Eqn. 3.8). The SV $\%G_{max}$ calculation depends only on the SV predicted maximum steady stress ($\sigma_{s,max(SV)}$).

$$SV\%G_{max} = \frac{\sigma_{a,max}}{S_e - (\frac{S_e}{S_{ult}})\sigma_{s,max(SV)}} \tag{3.8}$$

A set of FF surrogate models may also be used to predict the maximum steady stress of a specific design. This is found by choosing the maximum value in the set of predicted nodal steady stress values ($\sigma_{s,max(FF)}$). Like the SV $\%G_{max}$ value, the max stress as predicted by the FF surrogate models may also be used to calculate the $\%G_{max}$. This will hereafter be referred to as the FF $\%G_{max}$ value (Eqn. 3.9).

$$FF\%G_{max} = \frac{\sigma_{a,max}}{S_e - (\frac{S_e}{S_{ult}})\sigma_{s,max(FF)}} \tag{3.9}$$

On the other hand, the $\%G_{nodal}$ depends on the nodal steady stress results and the nodal alternating stress results (scaled from zero to the chosen max alternating stress) for an entire part. While a SV surrogate model cannot produce sufficient information for this calculation, the FF surrogate models can predict both the steady and the alternating stress results. The $\%G_{nodal}$ values may be calculated using one set of FF surrogate models for the steady stress and another set of

FF surrogate models for the alternating stress (Eqn. 3.10), creating a FF $\%G_{nodal}$ value for every node. The maximum %G value obtained with the FF surrogate model predictions will hereafter be referred to as the FF $\%G_{max,nodal}$ value (Eqn. 3.11).

$$FF\%G_{nodal} = \frac{\sigma_{a,FFem}}{S_e - (\frac{S_e}{S_{ult}})\sigma_{s,FFem}} \tag{3.10}$$

$$FF\%G_{max,nodal} = max(\frac{\sigma_{a,FFem}}{S_e - (\frac{S_e}{S_{ult}})\sigma_{s,FFem}}) \tag{3.11}$$

Depending on the kind of %G calculation the designer would like to use, a SV surrogate model can be used to predict the maximum steady stress or a FF surrogate model set can be used to predict the steady and alternating stresses for each node. These values are plotted as points on a dynamic Goodman diagram. One of the %G calculations discussed in this section can then be used to evaluate the points created by the surrogate model predictions.

3.3.2 Workflow

This research uses much of the same basic workflow as described by Bunnell et al. for training FF surrogate models [14]. The interested reader is referred to [14] for more complete details; however, the process is briefly presented here to help in understanding the details of this research. This method expands on the previous method by adding modal analyses and creating surrogate models to predict alternating stresses instead of only steady stresses.

Surrogate models require a training set of data that includes inputs and associated outputs. This training set may be constructed using a design of experiments (DOE) with n dimensions, where n is the number of desired input parameters. Each collection of parameter values constitutes a unique design instance. The suggested number of designs in the DOE can depend on criteria such as computational resources, complexity of the models, and number of dimensions.

The parameter values in each design instance constitute the inputs for a training point in the training set of data. These values are used to update the geometry of a parametric model. Rather than create a new finite element mesh for every new design, a mesh-morphing process is applied to the baseline design's mesh. This allows all the designs in the training set to have a common

set of relative node numbers and locations between all designs [14]. The mesh for each design is used in both a static and a modal FEA simulation. The results (steady stress, alternating stress, and geometric information for each node) are saved as the output for the training point. If geometric information is desired, either the deformed or undeformed geometry could be predicted.

This training data is then used to train the surrogate models. Two separate surrogate model sets are created. A steady stress FF surrogate model set is trained using the design parameters as inputs and steady stress for every node in each design as the outputs. Similarly, an alternating stress FF surrogate model set is trained using the design parameters as inputs and the alternating stress for every node in each design as the outputs. If a three-dimensional prediction of the shape of the part is desired, three additional FF surrogate model sets may be trained to predict, respectively, the X, Y, and Z coordinates of each node in response to the parameter changes. If only a SV surrogate model is to be used, a max stress SV surrogate model may be trained using the design parameters as inputs and only the maximum steady stress in each design as the output.

Once the surrogate models are trained, the surrogate models can accept a set of parameter values as a design instance and return predicted results. In this way, for a particular design, the steady stress and alternating stress are predicted with the FF surrogate model sets for every node. The alternating stresses are then scaled against a chosen max alternating stress.

These results can be used to dynamically update a Goodman diagram. Using the steady and alternating stress values from the FF surrogate models, each node is plotted as a point on the Goodman diagram. Equation 3.10 uses these predicted results to calculate the $\%G_{nodal}$ value for every node. The %G values can be assigned colors indicating proximity to the modified-Goodman line, where, for example, blue is 0% and red is 100%. In a similar manner, a SV prediction of the max steady stress can be used to plot a point at the max alternating stress and max steady stress location. This point may be used with Equation 3.8 to calculate the SV $\%G_{max}$ value.

Additionally, if the FF surrogate models are also trained to predict the X, Y, and Z locations, a three-dimensional visualization of the part geometry may be created. As in [14], the steady stress results for each node may be mapped to this geometry. This is also true for the alternating stress values. When the %G values have been calculated using Eqn. 3.10, then these also could be mapped onto the geometry. While these values are all indicated on the dynamic Goodman diagram

visualization, mapping these values to the geometry clearly shows the geometric location of the values and their patterns across the part.

When there are a large number of nodes, generating these visualizations can be more computationally expensive than the actual prediction of the results. In these cases, the time it takes to generate the visualization can be improved by visualizing only a representative portion of the nodal results, such as rendering only a random sample of the nodes on the Goodman diagram or only the surface nodes on the three-dimensional model.

3.4 Accuracy Comparison

The purpose of this accuracy study is to establish the relative accuracy of using FF surrogate models to calculate the %G values used in this method. The proposed design benefits of this research come from the prediction and use of the $\%G_{nodal}$ values. To build trust in the FF surrogate models, some idea of their accuracy must be determined.

Both SV and FF surrogate models are used in this study for comparison purposes. Because the SV surrogate model can only predict a single point, a SV surrogate model cannot be used to predict and or visualize nodal results for an entire part. Instead, some single %G value must be used to compare the FF and SV surrogate models, and Goodman diagrams will be used to help illustrate the relationships. There is no way to directly compare the FF $\%G_{nodal}$ values, which depend on both steady and alternating stresses, to a SV surrogate model because the SV surrogate model cannot predict a large collection of nodes. However, since the maximum stress can be predicted with both the SV and FF surrogate models, this study will focus on the comparison between the SV $\%G_{max}$ and the FF $\%G_{max}$, which depends only on the max stress.

These two values are both found by predicting a max steady stress, and are calculated in the same way, making them good objects of comparison. Their only difference is that the SV $\%G_{max}$ employs a single SV surrogate model that predicts the max steady stress, while the FF $\%G_{max}$ uses a FF surrogate model set to predict the steady stress results for every node on the model, and then uses the maximum value in the $\%G_{max}$ calculation. Though not directly applicable to a visualization, the $\%G_{max}$ is a good representation of a common single summary value for indicating an entire part's fatigue risk. By establishing the accuracy of the FF surrogate models for the $\%G_{max}$,

some degree of confidence can then be given to justifying the use of the FF surrogate models and FF $\%G_{nodal}$ results.

The error of a surrogate model refers to its ability to accurately predict the kind of data with which it was trained. In this research, the surrogate models predict the results of a structural and a modal FEA simulation. The surrogate models may be adjusted to more accurately represent a FEA simulation, and a FEA simulation may be adjusted to more accurately represent reality. While using more accurate FEA simulations is ideal, the objective of this research focuses on validation of surrogate model accuracy, not validation of a finite element model [12]. Thus, the accuracy of the surrogate models is evaluated by comparing the surrogate models' results and the FEA results. Using FEA results of validation samples as a way to evaluate accuracy of values predicted by surrogate models is consistent with previous work [10, 11, 76]. ANSYS 17.0 was the FEA solver used in this research.

3.4.1 Model

This test will use a FEA model of a jet engine compressor blade. The FEA model is based on the transonic Purdue blade [18]. The Purdue blade was developed for research purposes and is a general representation of a compressor blade. It has been parameterized to use the parameters and bounds in Tab. 3.1. The model consists of various airfoil profiles, or airfoil shapes, with connecting surfaces. It has been parameterized with the five parameters and bounds in Tab. 3.1. These parameters change the geometry of the model by adjusting the profiles at the root and tip of the blade. The Angle parameter adjusts the difference in the angle between the profile at the root and the profile at the tip. The two Chord parameters adjust the chord length at the root and at the tip. The Lean parameter offsets the tip profile in a direction perpendicular to the root profile's chord, while the Sweep parameter offsets the tip profile in the same direction as the root profile's chord. This parameterization and associated controls are discussed in [14].

This application of the Purdue blade is a "blade-alone" model, omitting the fillet around the bottom of a real compressor blade where it would connect to a disk or dovetail base. When using a "blade-alone" model in structural FEA, an artificial stress concentrator can develop around the base where the boundary conditions are set. The bottom 4 elements of the blade model have been neglected in order to remove artificial high stresses at the base.

Table 3.1: Parameter bounds

Parameter	Lower Bound	Upper Bound
Angle	-20 deg	20 deg
Chord@Root	20.32 mm (0.8 in)	30.48 mm (1.2 in)
Chord@Tip	20.32 mm (0.8 in)	30.48 mm (1.2 in)
Lean	-7.62 mm (-0.3 in)	7.62 mm (0.3 in)
Sweep	-7.62 mm (-0.3 in)	7.62 mm (0.3 in)

A mesh convergence study was performed in order to find the point where increasing the node count did not significantly alter the structural results. However, because the objective in this study was to calculate and visualize results in real time, a node count of 250,000 nodes was chosen as a good balance between model accuracy and model complexity, which affects the calculation and visualization speeds. While the most converged mesh was not used, the mesh convergence study helped to identify a much more converged mesh to use in these examples and fulfill the competing objectives of speed and accuracy. While using 250,000 nodes does slow the visualization performance slightly, the updates still occur in under one second.

3.4.2 Data and Assumptions

In order to create and then test the surrogate models, two data sets were created for both the SV and FF surrogate models: a training data set and a test data set. These will be briefly described in this section.

Training Data

In Section 3.3.2, it was stated that a training set of design parameters (inputs) and FEA results (outputs) must be generated to train the surrogate models. Preliminary studies showed that, with five parameters, 500 different points would be sufficient to sample the design space. An optimized Latin Hypercube was chosen for a DOE because it is a space-filling design [12, 77], and was used to generate 500 different combinations of the five parameters. A single set of these parameter values defines a unique design. Each of these designs were used to update the parametric compressor blade model. This updated model was used to conduct a structural and a modal analysis

in ANSYS, and nodal data was extracted for each design. This training data was used to create surrogate models to predict the behavior of each node.

For both SV surrogate models and FF surrogate model sets, radial basis functions (RBF) were used. While other functions could be used to build the surrogate models, the authors chose to use RBFs based on previous research and experience [14].

Test Data

In order to test the accuracy of the surrogate models, a new set of designs must be created. The surrogate models are more accurate when predicting results in response to inputs that are near the designs in the training set, so using the training set as test designs would unfairly represent the surrogate model accuracy across the design space. A new DOE and set of data was created with 500 different designs using the same process as before. These designs are distinct from those designs in the training set, but still evenly sample the design space for testing purposes.

In order to test the accuracy of the surrogate models, a separate set of designs must be created. The surrogate models tend to be more accurate when predicting information that is close to information used in the training data set [12]. If the surrogate models are used to predict the same designs used in the training data set, the results may appear artificially accurate. Thus the test designs must be distinct from those designs in the training data set [67, 77]. The objective is to find points for comparing the response of the surrogate models to the actual data across the full span of the design space.

A new DOE and data set were created with 500 different designs using the same process as before. The %G values and max stress values for each design in this data set were found. Because they were calculated with data from ANSYS, and were not predicted directly or calculated from predicted data, they will be referred to as ANSYS values (ANSYS %G and ANSYS $\sigma_{s,max}$). These ANSYS %G and ANSYS $\sigma_{s,max}$ values formed a test data set, and were used for validation against values predicted by the surrogate models [10, 76].

Assumptions

Additionally, to create the Goodman diagram, certain properties were calculated or chosen. The blade is assumed to be made of Titanium 6Al-4V [1]. The ultimate strength for this material is 1182108.7 kPa (170000 psi), and the endurance strength was calculated to be 317158.8 kPa (46000 psi) using Marin factors [49]. This makes the assumptions that the blade has been machined and the operating temperature in the compressor is around 638.9 K (1150 R). This temperature is assumed to roughly simulate an intermediate compressor stage between the inlet and the combustor. This ultimate strength and endurance strength are within the ranges used by other studies on this material [20]. The maximum alternating stress was assumed to be 82737.1 kPa (12000 psi), which is in an acceptable range for alternating stresses for compressor blades.

3.4.3 Test Procedure

The following process was used for each design in the test data set:

1. Using the ANSYS data for each design in the test data set:
   (a) Obtain the max stress (ANSYS $\sigma_{s,max}$)
   (b) Using the ANSYS $\sigma_{s,max}$, calculate the ANSYS $\%G_{max}$
   (c) Using the steady and alternating stresses at each node, calculate ANSYS $\%G_{max,nodal}$
2. Using the FF surrogate model sets:
   (a) Predict the steady stresses
      i. Find the max steady stress (FF $\sigma_{s,max}$) in the set of predicted steady stresses
   (b) Predict the alternating stresses
      i. Scale the predicted alternating stress values from 0 to a max alternating stress
   (c) Using the FF $\sigma_{s,max}$, calculate the FF $\%G_{max}$ value
   (d) Using the steady and alternating stress values at each node, calculate the FF $\%G_{nodal}$ values
      i. Find FF $\%G_{max,nodal}$ value in the set of predicted FF $\%G_{nodal}$ values

(e) Calculate the error between:

i. the FF $\sigma_{s,max}$ and the ANSYS $\sigma_{s,max}$ values

ii. the FF $\%G_{max}$ and the ANSYS $\%G_{max}$ values

iii. the FF $\%G_{max,nodal}$ and the ANSYS $\%G_{max,nodal}$ values

(f) Plot all data points on the dynamic Goodman diagram

3. Using the SV surrogate models:

(a) Predict the max steady stress (SV $\sigma_{s,max}$)

(b) Using the SV $\sigma_{s,max}$, calculate the SV $\%G_{max}$ value

(c) Calculate the error between:

i. the SV $\sigma_{s,max}$ and the ANSYS $\sigma_{s,max}$ values

ii. the SV $\%G_{max}$ and the ANSYS $\%G_{max}$ values

(d) Plot the data point on the dynamic Goodman diagram

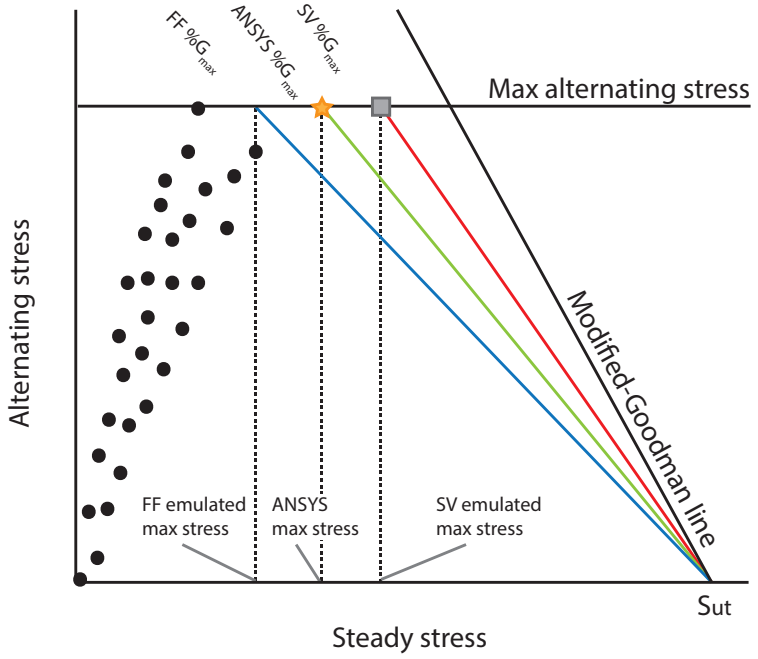

Figure 3.4: Notional comparison of $\%G_{max}$ values as obtained from the max stress from ANSYS, a SV surrogate model, and a FF surrogate model set

Figure 3.4 illustrates a notional comparison between the SV %G$_{max}$, FF %G$_{max}$, and %G$_{max}$ from ANSYS on the Goodman diagram. These points are only examples, and were not generated using actual data. The black points represent FF predicted nodal values and the yellow star represents the max steady stress from ANSYS. The %G$_{max}$ is displayed as calculated from the max steady stress from ANSYS, the FF surrogate models, and the SV surrogate model. This illustrates how both types of surrogate models may predict different values than ANSYS. It also demonstrates how the various values are obtained.

3.4.4 Results

The percent error of the max stress, %G$_{max}$, and max %G$_{nodal}$ is determined using the Eqn. 3.12, where X is either a %G or a max stress value.

$$\%E = \left| \frac{X_{em} - X_{ANSYS}}{X_{ANSYS}} \right| * 100\% \tag{3.12}$$

The magnitude difference was calculated using Eqn. 3.13. These are simple differences in magnitude of the various values. They are included to give additional context for the percent error values. When comparing two values with small magnitudes, a percent error may be high, but the magnitude difference can show that the actual difference may not be very significant in practical terms.

$$MD = |X_{ANSYS} - X_{em}| \tag{3.13}$$

These values were obtained for each of the 500 designs in the test DOE. The average of the percent error and magnitude difference for the %G$_{max}$ and max steady stress values are presented together in Tab. 3.2. The improvement gained refers to how much lower the FF %G$_{max}$ is than the SV %G$_{max}$ and follows the equation $1 - (FF/SV)$. Both the percent error and magnitude difference indicate that, for this data set, the FF surrogate model set provides data that allows for a more accurate calculation of the %G$_{max}$ than the SV surrogate models do. The magnitude difference shows that, on average, the FF %G$_{max}$ is different from ANSYS by a value of 0.4 while the SV %G$_{max}$ is different by a value of 0.6.

This is partly explained by the accuracy results for predicting the maximum steady stress. The $\%G_{max}$ depends on the max stress value. Because the max stress is more accurately predicted using the FF surrogate models, it would be beneficial to use the FF predicted max stress for this calculation instead of the SV predicted max stress. Figure 3.4 helps visualize these values. The data shows that FF $\%G_{max}$ is closer than the SV $\%G_{max}$ is to the ANSYS $\%G_{max}$ value. In this data set, the FF surrogate models predict the max steady stress with a mean percent error that is 53.54% lower than the mean percent error for the SV surrogate model's max steady stress. The average magnitude difference also shows that, on average, the FF surrogate models predict a closer value to ANSYS than the SV surrogate model for max steady stress. It follows that calculating the $\%G_{max}$ using the FF surrogate models from Eqn. 3.9 will be more accurate than the SV counterpart in Eqn. 3.8. Thus, even for applications where the more conservative $\%G_{max}$ value is desirable, this data indicates that the FF surrogate model set would produce much more accurate values than the simple SV surrogate model would.

Table 3.2: FF vs SV $\%G_{max}$ error

Value	**Avg. Perc. Error**	**Avg. Mag. Diff.**
FF $\%G_{max}$	0.904%	0.4%
SV $\%G_{max}$	1.504%	0.6%
Improved	**39.91%**	**27.97%**
FF $\sigma_{s,max}$	2.30%	6.7 MPa (967.2 psi)
SV $\sigma_{s,max}$	4.94%	12.3 MPa (1780.0 psi)
Improved	**53.54%**	**45.66%**

3.4.5 Results for FF $\%G_{max,nodal}$

With some confidence gained concerning the accuracy of the FF surrogate model set, the FF $\%G_{max,nodal}$ can now be examined. Because the SV surrogate model only predicts the maximum steady stress, there is no one-to-one comparison between FF and SV values for $\%G_{max,nodal}$. However, the FF $\%G_{max,nodal}$ can still be compared to the $\%G_{max,nodal}$ as obtained by the test data from ANSYS. These results are presented in Tab. 3.3. The average percent error is still low (1.337%). This provides a frame of reference as the design benefits are outlined in the next section.

Table 3.3: FF $\%G_{max,nodal}$ vs ANSYS error

Value	Avg. Perc. Error	Avg. Mag. Diff.
FF $\%G_{max,nodal}$	1.337%	0.5

3.5 Using $\%G_{nodal}$ Instead of $\%G_{max}$

The $\%G_{max}$ values compared in the last section have some benefits that may make them a more desirable choice over using a collection of $\%G_{nodal}$ values for every node on a part. The $\%G_{max}$ is much simpler to find, and requires far less computation. When used with a SV surrogate model, the training data required is also simpler to obtain, and requires less memory. Because it is only one value, it can conveniently be used to summarize the fatigue risk of an entire part.

However, summarizing an entire part by the stress at one location or some other single value method neglects information about the rest of the part that could be useful in making design choices. No information is conveyed concerning the distribution of risk in different locations. Important differences in designs can also be obscured. In particular, the $\%G_{max}$ may be considered overconservative in some cases. This can result in over-designing for fatigue, raising costs and weight unnecessarily [69]. In the design of parts for jet engine turbomachinery, it is very desirable to minimize both costs and weight if structurally feasible.

Using the information across a part, such as the entire collection of $\%G_{nodal}$ values, can help mitigate some of these limitations. This section will illustrate the advantages of specifically using the $\%G_{nodal}$ values instead of the $\%G_{max}$.

Using only the $\%G_{max}$ values would show the response of a single point. The $\%G_{max}$ provides little information about the full response of the part. If the $\%G_{max}$ increases between two different designs, it might be concluded that the entire part is at more risk of fatigue failure in the new design. However, the shape that comes from using every node's position on a Goodman diagram communicates whether or not a high percentage of points have high risk of fatigue failure.

Figure 3.5 shows the nodal information mapped onto the Goodman diagram for two different designs of a compressor blade. The second design was obtained by adjusting the Lean parameter. The change in data on the Goodman diagram occurred as a fluid real-time response as the parameter was changed. Observing the behavior in this type of design activity leads to better intuition about the effect of variations in a single parameter on the entire part. In Fig. 3.5a, the

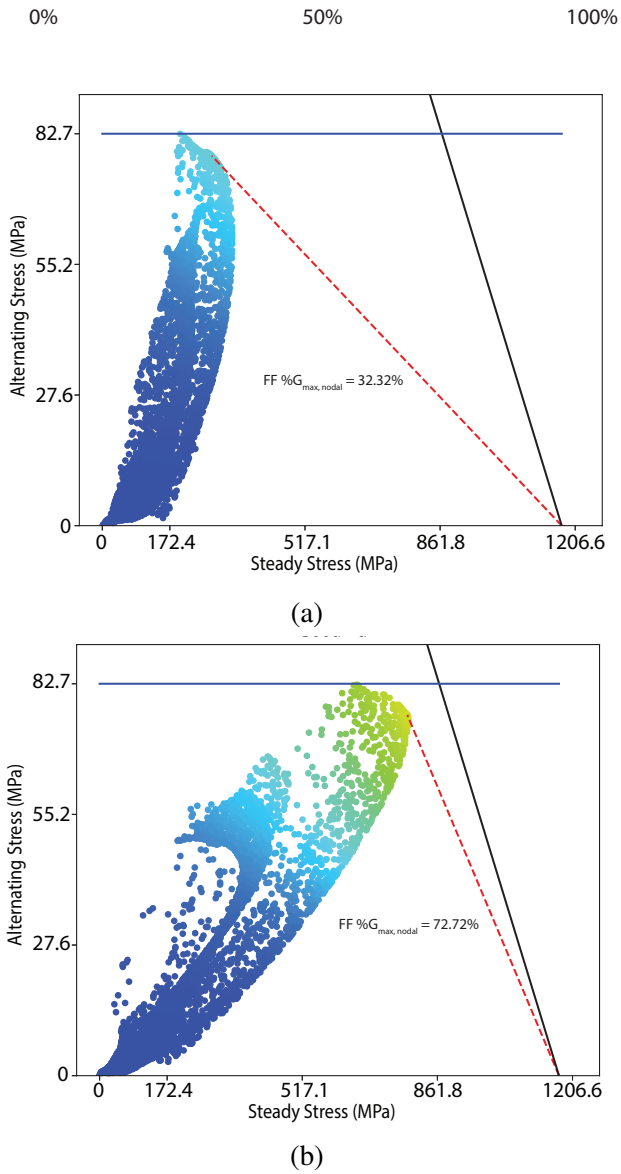

Figure 3.5: The $\%G_{nodal}$ results before and after changing the Lean parameter from 0.00 mm (3.5a) to 6.86 mm (0.27 in) (3.5b). The node with the $\%G_{max,nodal}$ is indicated by the dashed line.

$\%G_{max,nodal}$ is 32.32%, and the $\%G_{nodal}$ values for the rest of the nodes show that there is a fairly even distribution from 0% to 32.32%. In Fig. 3.5b, the $\%G_{max,nodal}$ is much higher (72.72%). However, the rest of the $\%G_{nodal}$ values show that only a small portion of the nodes are very high; the majority are about 30% or lower. The shape of the distribution of points also indicates that there are two major groups: one "branch" is low, while the other has the higher values. The %G values for every node reveal distinctions and patterns that are not visible when neglecting the response at every node.

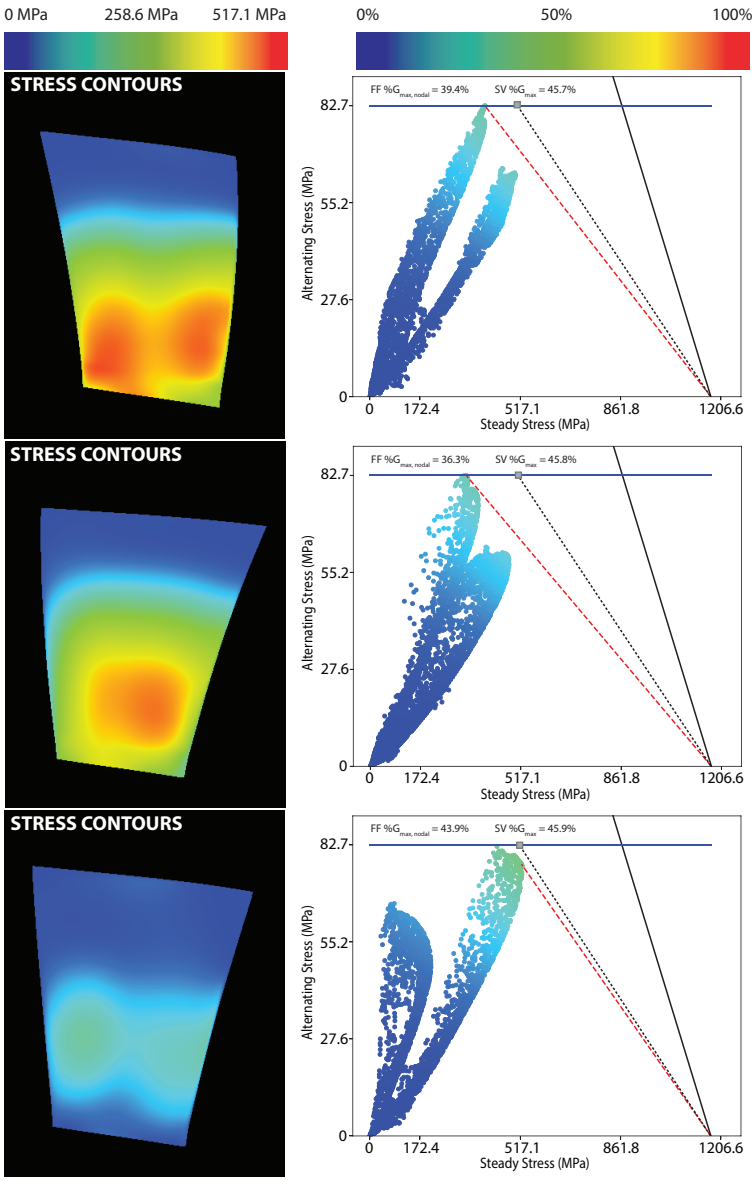

Figure 3.6: Three different blade designs with almost identical SV $\%G_{max}$ values but different geometries, FF $\%G_{nodal}$ values, and shapes on the Goodman diagram. The geometry views (left) show steady stress results while the Goodman diagrams (right) show the $\%G_{max}$ values (gray square with a black line) and the $\%G_{nodal}$ values (colored points, with a red line indicating the $\%G_{max,nodal}$ point). Top = Design 1, Middle = Design 2, Bottom = Design 3.

Using only the $\%G_{max}$ also obscures the differences between designs. There may be similar $\%G_{max}$ values, but they do not convey inequalities of various designs. Consider the case presented in Fig. 3.6. Three different designs are shown, with the main geometric variations occurring in the Lean and Sweep parameters. The Goodman diagrams are presented alongside three-dimensional visualizations of the geometry which have the steady stress values at each node mapped onto the

surface. The three different designs produce very similar $\%G_{max}$ values: the $\%G_{max}$ for each is about 45%, shown with a gray square on the maximum alternating stress line. With only the $\%G_{max}$ values, there is nothing to distinguish these designs from one another. However, the responses for every node on the Goodman diagrams show three different $\%G_{max,nodal}$ values, ranging from 36% to 44% (see Tab. 3.4). In addition, the shape of the data changes greatly from one design to the next. Some designs have nodes that have a mostly linear relationship between the steady and alternating stress, while others have groups of nodes with low steady stress but high alternating stress.

Table 3.4: The $\%G_{max}$ and $\%G_{max,nodal}$ for each of the three designs

Design	$\%G_{max}$	$\%G_{max,nodal}$
1	45.7%	39%
2	45.8%	36%
3	45.9%	44%

The three-dimensional visualization of the steady stress contours helps provide context to the information in the Goodman diagrams. This visualization is made possible by obtaining the steady stress and nodal coordinate values from either an FEA solution or FF surrogate model set predictions, and reveals further differences between the three designs. The max stress is similar between the designs, but the three-dimensional visualizations reveal that the node with the max stress is located in different areas of the part. The part also has very different geometries and stress distributions. Using the $\%G_{max}$ value as the only criteria, these three designs appear to be nearly identical, and no indication of the differences is given. Instead of just a single value, using the full set of nodal values allows a designer to evaluate a design based on many criteria at once, including geometrically visualized data.

These cases illustrate how using the $\%G_{nodal}$ values can help to reduce unnecessary conservatism. The $\%G_{max}$ value is a type of "worst-case" scenario, and neglects the true variations between designs. Instead of merely deeming a design as acceptable by a single desired %G value, using all the $\%G_{nodal}$ values gives a much more complete view of the fatigue life of the entire part, and can reveal many differences. Moreover, the $\%G_{max,nodal}$ values are often shown to be far smaller than the $\%G_{max}$ values, indicating that the $\%G_{max}$ may be overly conservative.

3.6 Design Benefits

The previous section gave an indication of the accuracy of using FF surrogate model sets to predict Goodman response data. The benefits of using the $\%G_{nodal}$ instead of the $\%G_{max}$ were also discussed. This section builds on those findings and suggests the design benefits that come from using the following:

1. FF surrogate models to predict the alternating and steady stresses at every node
2. Calculating the %G values from these predicted stress values
3. The $\%G_{nodal}$ values as a method of evaluating risk of fatigue failure rather than conservative summary methods like $\%G_{max}$

These design benefits include:

1. A dynamic Goodman diagram
2. Faster response times
3. Better intuition about parameter and response relationships
4. Understanding about which geometric locations are at most risk of fatigue failure
5. An ability to combine fatigue analysis with other early design process activities

The data in the examples for this section were obtained using a 25K node model of the Purdue blade, rather than the 250K node model used in the previous section. Though predicting and visualizing results for the 250K node model occurs in real time (under one second per update), the researchers used a model with a smaller node count for convenience.

3.6.1 Dynamic Goodman Diagram

By using two FF surrogate model sets - one to predict the steady stress at every node and one to predict the alternating stress at every node - it becomes possible to create a dynamic Goodman diagram. The predicted points can be plotted on the Goodman diagram, and as variations

are made to the inputs of the surrogate models, the collection of points on the Goodman diagram will update in real time. These points can be colored by using the steady and alternating stress predictions to calculate the $\%G_{nodal}$ values at each node (Eqn. 3.10), and give an evaluation of the fatigue life of each node in the model.

The ability to control and vary the parameter values allows the designer to adjust the inputs to the FF surrogate model sets, and thus see the response of the data on the Goodman diagram update. This control increases the usefulness of this application for design space exploration activities.

3.6.2 Faster Response Time

Without a rapid method of predicting this information, a designer must run a new analysis in a finite element solver to obtain the full response of a part to changes in the design. This requires an expense of time and computer resources (a computational cost). For a typical part, the time to solve is affected by factors such as node count, model complexity, and computing power, and can take anywhere from a few minutes to days [14]. For obtaining Goodman response information, such as the %G values, this requires waiting for a structural and a modal analysis to complete, calculating the %G values from the steady and alternating stress results at each node, and mapping the updated %G values onto either a Goodman diagram or the geometry.

With the FF surrogate model sets, the response to the parameter changes occurs without any appreciable delay. This includes a prediction of the full structural, modal, and the calculated full Goodman response. These results and the associated visualizations occur as direct responses to the design changes. This shortens the time needed to generate this information during design space exploration activities when evaluating the validity of a design. It also makes it possible to continuously change parameters during design exploration and see the responses fluidly updated.

3.6.3 Intuition About Parameter/Response Relationships

Creating a Goodman diagram without surrogate models requires a designer to perform separate static and modal analyses for each new design. This makes it difficult to gain intuition about how variations in the design parameters affect the responses at different locations on the part. The

calculation and visualization of the Goodman diagram's results to the designer are often several layers removed from the moment the actual design changes occur; this involves adjusting the model environment and boundary conditions in FEA packages and then assembling results [26]. These tasks are largely unrelated to the design changes and, coupled with the delay in response, can make creative design exploration less fluid and feasible [26]. To save time, only certain representative designs are evaluated instead of the entire range of designs across the design space.

With a dynamic Goodman diagram, designers may vary design parameters and observe the responses immediately. A designer can thus consider information regarding the max steady stress, the $\%G_{nodal}$ values, each node's steady and alternating stress, and the shape of the plotted data on the Goodman diagram all as direct responses to the changes in parameter values. In this kind of application, the process of making changes and observing responses is much clearer and simpler to use than setting up separate analyses and waiting for each variation's results. There is a clear link between the changed parameters and the resultant behavior, making it easier to develop intuition and understanding about the part. It also is possible to explore more of the design space at very low cost. This speed, the clarity of parameter/response relationships, the completeness of available information, and the freedom to explore the design space rapidly makes it far easier to gain intuition about which parts of the design space will yield the best results. Observing the behavior in this type of design activity leads to better intuition about the effect on the entire part of variations in specific parameters.

This type of information can be obtained without the FF surrogate model sets by using a finite element solver, but would be subject to longer wait times. A structural and modal analysis would have to be solved for each design. It would also take more experience and iterations to determine candidate designs. The surrogate models allow the design space to be quickly explored and identify designs of interest. If a design is deemed desirable by the designer, further small variations to the design parameters can be made to find the optimal blade in that neighborhood of the design space without the time loss of running a full analysis. This could be useful when considering the outcome of an optimization routine. The surrogates could show the recommended optimal design, and the designers could then adjust the parameters and generate visualizations to explore the other variations in the vicinity of the solution. This could help them build confidence

in the solution, or possibly reveal important factors the optimization routine was not set up to consider.

3.6.4 Geometric Locations of Failure

By using the $\%G_{nodal}$ values on the Goodman diagram, other visualizations become possible. The colors indicating the $\%G_{nodal}$ values of every node on the Goodman diagram may be mapped onto the three-dimensional geometry of the part. This lets a designer connect each node's percent Goodman value with a physical location on the part. Thus, a designer may see that a part is at risk of failing in fatigue, but can also determine which specific locations are most likely to be failing in fatigue. This visualization becomes essentially a different way of viewing and analyzing the information on a Goodman diagram that plots the results of all nodes on the part.

Studies using similar three-dimensional visualizations were mentioned previously; however, these visualizations were not created with data predicted by FF surrogate models. Depending on the number of nodes predicted, this method can generate visualizations for new designs in less than a second, or essentially in real time. Under these conditions, these visualizations update fluidly in response to parameter changes and present a smooth exploration of designs throughout the design space. This fluidity can provide much greater intuition about behavior than discrete visualizations of different designs in the design space.

In Fig. 3.7, the three-dimensional geometry visualizations are shown with the $\%G_{nodal}$ values mapped onto their surfaces. For reference, they are shown next to corresponding visualizations of the $\%G_{nodal}$ values on the dynamic Goodman diagrams. A designer can find a point on the Goodman diagram and use its %G color to identify where it is on the blade surface. Just like the Goodman diagram and three-dimensional mapping of the steady stress updated in real time to parameter changes in the last section, here the Goodman diagram and the three-dimensional mapping of the %G values also update in real time.

While examining this information only on a Goodman diagram can indicate risk of fatigue failure for each node, mapping and viewing that same information on the geometry provides more information. Not only do the values on the geometry communicate the distribution of %G values, but also show exactly where the %G values occur on the part. A designer can use the %G color to determine how close a geometric location is to failing in fatigue. In these examples, the highest

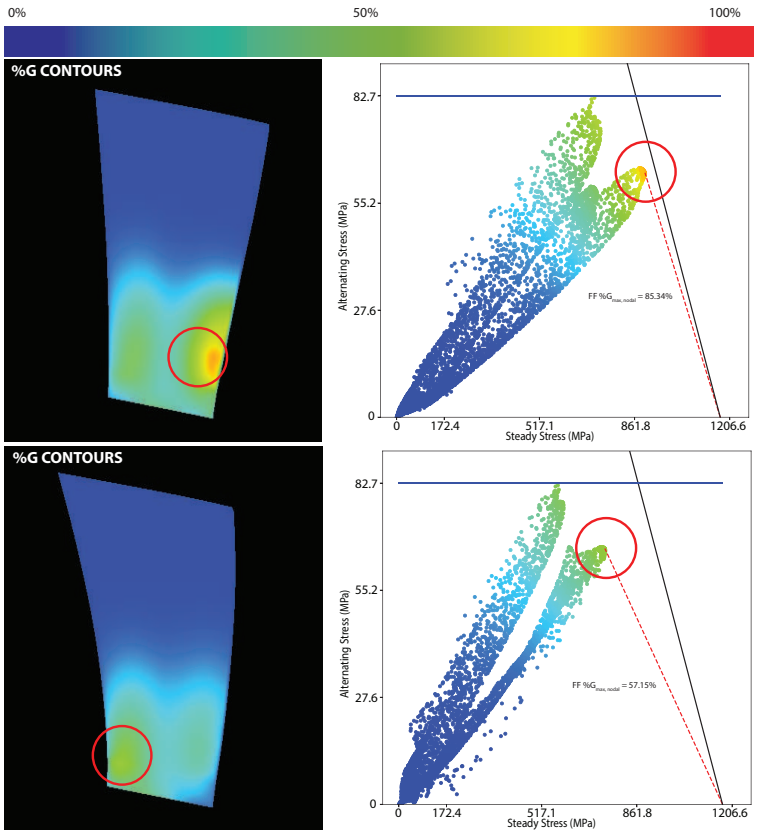

Figure 3.7: The data on the Goodman diagram can be mapped onto three-dimensional geometry to provide insight into locations of potential failure. The %G values and colors on the Goodman diagrams (right) are directly mapped onto the corresponding nodes on the blade surfaces (left). The colorscale used on the Goodman diagram is now shared with the three-dimensional geometry (top).

%G values occur closer to the root (bottom) of the blade, while the lowest values are found at the tip (top) of the blade where the vibratory and steady stresses are smaller.

The relationship between this type of visualization and the accompanying Goodman diagram can inform the design process in two directions; each visualization gives context and insight into the other. Mapping the %G values onto the geometry allows designers to understand which locations on the part are at the most risk of fatigue failure. On the other hand, seeing the colors on the geometry allows a designer to understand possible geometric influences for why the points behave a certain way on the Goodman diagram.

With only a SV surrogate or the $\%G_{max}$, there is no information given to the designer about how to improve the design. All that is provided is a single measure of the fatigue risk of the part. In contrast, the FF surrogate model sets allow a designer to understand the entire part's fatigue risk

as well as see which physical locations have the highest %G. This understanding can reveal what changes need to be made.

Consider the case where the nodes of a part are represented as points on a Goodman diagram, and all but a few have low %G values. The few points with higher %G values are not above 1.0, therefore they are not definitively failing in fatigue, but they are perhaps at greater risk than the rest of the part. Rather than accepting or rejecting the design on this basis, a designer can look at these points mapped onto the geometry. With the aid of FF surrogate model sets, this visualization can respond to changes in parameter values. If the points with higher fatigue risk are in a critical location, the designer may decide to adjust the design to reduce this risk. However, if they occur at a location that is known to be less sensitive, then the designer may choose to accept the design.

A designer may also use the geometric location of those high %G value points to determine why they are high. By pairing this type of visualization with real-time FF surrogate model predictions, the values on both the Goodman diagram and the three-dimensional visualization can be updated as parameters are adjusted. By visualizing predicted results on a Goodman diagram and the three-dimensional geometry, the effect of making parameter adjustments becomes clearer. This allows a designer to make changes to the geometry and affect the distribution of %G values across the part. This type of exploration may reveal certain parameters that can reduce the high %G values, or even how to shift them to a less critical location on the part. With the %G values mapped onto the geometry, a designer may also be able to use past engineering experience to know which parameters could be shifted to lower the value or change the location of the high risk areas. For example, if the high %G values are located at the root of a jet engine compressor blade, a designer may see that increasing the chord at the root could reduce the alternating stress experienced by that region and thus reduce the risk of fatigue for that area. This visualization makes it easier to determine geometric causes for risk of fatigue failure at each node than it would be by simply looking at these points a Goodman diagram. The low computational cost of predicting these results with FF surrogate models also make it easier for designers to explore and adjust the design.

An additional visualization is possible in which the steady stress, alternating stress, and %G contours are each rendered on three separate geometries in juxtaposed views. An example of this presentation is shown in Fig. 3.8. When the three visualizations are shown with the Goodman diagram, even more information is simultaneously available to the designer. Judgments about

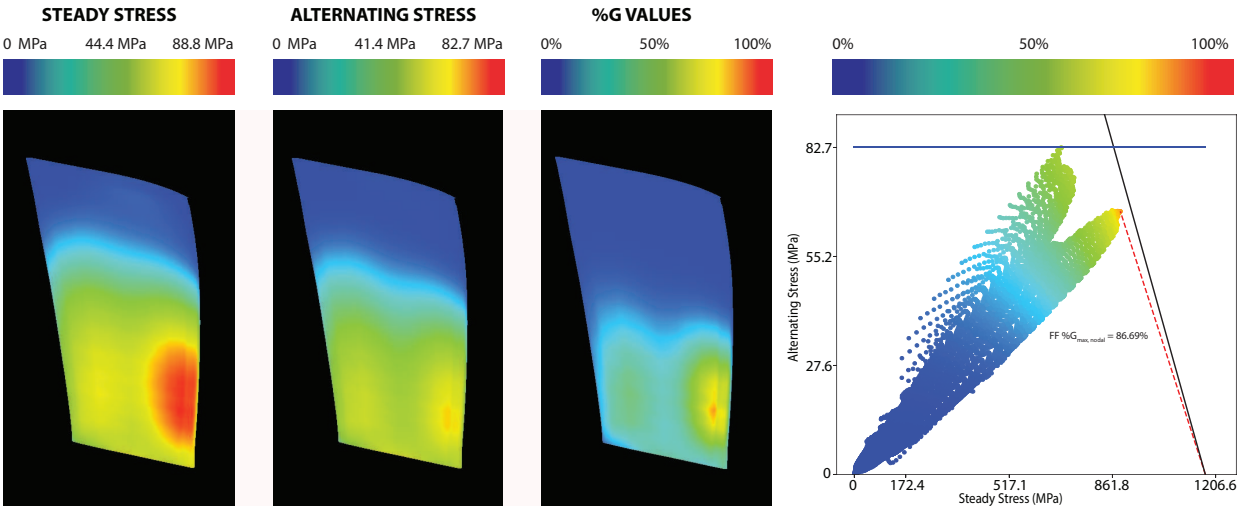

Figure 3.8: With FF surrogate models, it is possible to display a real-time updated view with 3-dimensional visualizations of the steady stress, alternating stress, and %G contours alongside the Goodman diagram

steady, alternating, and %G values at specific locations on the blade can be made together across different designs in the design space. Conveniently, this presents all relevant information in one interface for the designer: the Goodman diagram, the geometry, the steady and vibe stress for each node on the blade, as well as the %G values. Unlike traditional methods, all of this information is instantly responsive to any changes to the input design parameters.

This once again helps the designer avoid unnecessary conservatism. Instead of guessing where to add material or how to change the shape of the blade in order to improve the fatigue life of the part, now it is clear exactly what area needs to be improved. These improvements could target fatigue life improvement for the entire part or just specific areas of interest. Even slight reductions in weight or efficiency can have large financial savings in the jet engine industry. Thus being able to directly see the location of needed improvement allows designers to make better and more efficient judgments about how to improve the %G values on the Goodman diagram while still saving costs.

3.6.5 Early Design Process Benefits

FF surrogate model predictions of the nodal steady and alternating stresses allow fatigue analyses to be updated and presented alongside structural and modal analysis. Without these pre-

dictions, fatigue analysis must occur after choices concerning geometry and structural design are made [1]. If a design passes the criteria set for the structural and modal analyses, it is then evaluated with fatigue analysis. The fatigue analysis must occur after the previous analyses because the Goodman diagram and associated information depend on their results. Designs that fail at the fatigue analysis stage lead to further iterations of these processes. These iterations are repeated until suitable candidate designs are chosen.

With FF surrogate models and the dynamic Goodman diagram, these analyses may be considered simultaneously. The early design exploration can consider geometry, steady stress, alternating stress, and fatigue life assessments together without waiting for expensive simulations, thus freeing the designer to explore and observe many more design variations as needed. FF surrogate models link all these responses to design space exploration. Where previously fatigue analysis was removed from this type of exploration, here it is linked as a direct result of these design changes. Insights about how fatigue analysis responds to changes in parameters may now be formed more easily.

Having this type of information available at earlier design stages can reduce design and analysis costs downstream, and can better prevent failure [56, 57]. Designers can be guided away from designs that fail in fatigue without needing to perform as many design iterations. Fatigue information can be considered for a larger number of design alternatives instead of only a few selections [56].

3.7 Conclusions

A study about the accuracy of FF surrogate model sets compared to SV surrogate models was conducted in the context of information on a Goodman diagram. This study has shown that FF surrogate models can predict certain values, such as the maximum steady stress on a FEA model, more accurately than a traditional SV surrogate model. Calculations based on these predictions, such as the $\%G_{max}$, are also shown to be more accurate when using FF surrogate model sets than when using SV surrogate models. While FF surrogate models require more work to set up, they can provide a more accurate representation of the data they predict.

It has also been shown that using nodal information on a Goodman diagram can reduce conservatism in evaluating fatigue life for an FEA model. When using a single point for evaluation,

such as the conservative $\%G_{max}$ value, only one value is obtained. Using every node's position for evaluation, such as the $\%G_{nodal}$ method, reveals variations in design and provides a much truer understanding of the risk of fatigue failure across the entire part.

While a SV surrogate model can only predict one value, e.g. a SV $\%G_{max}$, FF surrogate models can predict both single values and complete sets of nodal information, e.g. FF $\%G_{max}$ and FF $\%G_{nodal}$. It can be noted that, in this particular study, the surrogate models were used to predict stresses and the %G values were calculated off of these values. Simply predicting the %G values for a part is less useful than also predicting the steady and alternating stresses that contributed to those values. By predicting the stresses, the nodes' locations may be plotted dynamically on a Goodman diagram. This also allows for the contours on the three-dimensional geometry to be dynamically switched between steady stress, alternating stress, and %G values.

The design benefits of using FF surrogate models to predict the nodal results in these analyses have been described as:

1. A dynamic Goodman diagram
2. Faster response time for results
3. Gaining intuition about parameter/response relationships
4. Predicting the geometric locations of failure
5. Combining steps used in early design process iterations

In this study, FF surrogate models have been applied to fatigue analysis of jet engine compressor blades, which enabled a dynamic Goodman diagram that can respond to parameter changes in real time. This study focused on the benefits of using these FF surrogate models with the dynamic Goodman diagram, but similar methods could be applied to many other analyses and processes in engineering design. They have the potential to allow more thorough design exploration, shorten iterative design activities, and provide more accurate predictions than traditional uses of surrogate models. Experimenting with these applications and their implications will be of great benefit in establishing the usefulness of this method in industry and real-world design problems.

CHAPTER 4. MODAL ASSURANCE CRITERION

In addition to fatigue life predictions, modal analyses are used during design space exploration to study how a part's geometry affects the vibratory displacement patterns (or mode shapes). However, these mode shapes can be difficult to classify, and require comparison with known mode shapes. Mode shapes are compared to one another with an equation called the modal assurance criterion (MAC).

This chapter uses FF surrogate models to predict the modal displacements at every node on the part. The full collection of displacement predictions is then used in the MAC equation with a set of modal displacements from a reference mode shape attained in ANSYS (RQ1). The accuracy of these predicted MAC values is compared with those obtained with SV surrogate models (RQ2). It is found that FF surrogate models require a very high number of training samples to achieve comparable accuracy to SV surrogate models for certain mode conditions. This is due to the highly nonlinear nature of the data being predicted. While there are some benefits from using FF surrogate models in this application, SV surrogate models remain a more feasible option (RQ3).

This study differs from the other chapters by predicting a displacement result (instead of a stress result) at every node. Rather than examining predicted nodal results or calculating a new result at every node from predicted data (like the $\%G_{nodal}$ values), this study combines all the results from a static simulation with a single calculation (the MAC equation). The comparison between FF and SV surrogate models in this study is not based on a single node's results, but between a type of summary of all the nodal results. This kind of result makes this comparison unique from the other studies in this research.

This research has been submitted for publication to the journal of *Mechanical Systems and Signal Processing* and is under review as of July 2019.

4.1 Introduction

Modal analysis is a finite element analysis (FEA) process that allows aerospace engineers to predict the modal responses for a part under specific conditions. Modal analyses calculate the mode shapes, or relative displacements across the part, which occur at the natural frequencies. These mode shapes help engineers predict the ways in which a part will vibrate in operation. Understanding these vibrations and patterns is critical for designing safe and reliable parts.

The modal assurance criterion (MAC) measures the similarity between two mode shapes from different designs [10, 22, 23]. Often, the MAC equation is used to compare each mode, d, in a design to each mode, r, in a reference design with [22, 23, 78]. This reference design may be a baseline design, a design that has been previously manufactured and tested, or simply another proposed design that is being evaluated. Modal comparisons between specific modes are denoted as d:r. The values of d and r are ordered according to the natural frequency values at which the modes occur (i.e., d=1 or r=1 indicates the mode associated with the lowest natural frequency in the design).

When comparing different designs with the MAC equation, the computational expense of conducting a wide variety of modal analyses may make a thorough exploration of a design space infeasible [12,67,79–81]. The computational cost of modal analyses depends on model complexity as well as the number of mode shapes to be obtained. Although computing power has increased, so has the need for more complex models [5]. Often in design space exploration fewer designs are explored or lower fidelity models are used to compensate for the computational expenses [12, 14]. This could result in sub-optimal designs that are not necessarily the best choice from the design space. The level of effort to adjust and properly set up each modal analysis often also makes them less ideal tools for free real-time exploration of design variations [26].

To overcome the high computation time needed for expensive simulations, many aerospace designers employ surrogate models to quickly predict needed information [5,7,8,12,66]. Surrogate models (also known as response surface models, metamodels, regression models, or emulators [6]) create simplified mathematical relationships between inputs and outputs of a system. Using new inputs, the surrogate models can then predict results in a fraction of the time it takes to perform the original calculations [6, 12, 77]. These surrogate models are often used in design space exploration or optimization routines in order to evaluate many design variations without the prohibitive cost of

full simulations [7–12]. Denimal et al. [10] suggested that surrogate model methods ought to be developed for better design space exploration and emulation of modal results.

Most surrogate models map a set of inputs to a single output value, and thus are limited in the information they can predict [7]. They cannot predict the response of an entire system, but instead predict a single value that represents the system (see Fig. 4.1 (top)). These types of surrogate models will be referred to as single-value (SV) surrogate models.

Applied to MAC values, SV surrogate models can be used to predict MAC values for a specific mode-to-mode (d:r) comparison. If properly trained, these SV surrogate models make it possible to quickly update the MAC value in response to different input parameters without needing to run a modal analysis for every new design. However, because they are trained for only a specific type of modal comparison, they are limited in their utility (e.g., a SV surrogate model trained to predict the MAC values in a d:r comparison cannot predict the MAC values in a d+1:r or a d:r+1 comparison). They also can only predict the MAC values for comparisons with one reference design. If comparisons with a different reference design are desired, then a new set of SV surrogate models need to be trained on MAC values performed with the new reference design. This results in SV surrogate models that only predict MAC values specific to a single comparison and reference design. These limitations make SV surrogate models less adaptable to unforeseen changes in design space exploration.

In recent years, methods for using surrogate models to predict the behavior of every node in an FEA simulation in real time have been developed [14–16]. They generally take geometric parameter values as inputs and predict FEA results for the entire part. These methods allow a much more detailed and complete prediction of structural FEA results than SV surrogate model methods, and enable a three-dimensional visualization of the structural results mapped onto the part geometry. These surrogate models can be solved very quickly, allowing the predicted FEA response to changes in parameter values to be shown in real time, where real time is defined as updated responses appearing in under a second [14]. Because this use of surrogate models predicts an entire field of responses from a FEA simulation, these will be referred to as full-field (FF) surrogate models. This work uses the terminology of SV and FF surrogate model methods in order to provide a distinction between these different uses of surrogate models. The difference between

SV and FF surrogate models is illustrated in Fig. 4.1. While SV surrogate models only predict a single value, FF surrogate models predict values for each node on a part.

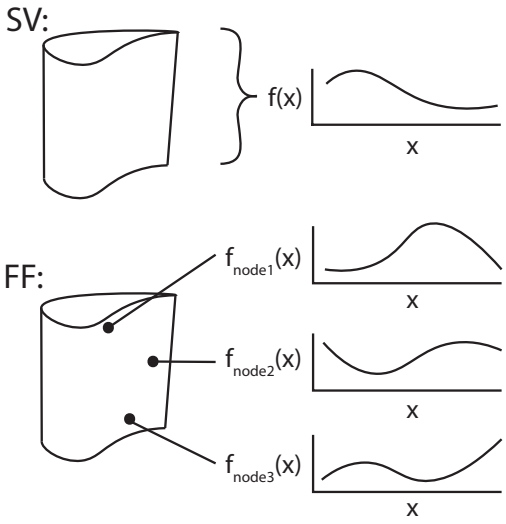

Figure 4.1: (top) Single-value (SV) surrogate models only predict one value that represents the response of the entire part across the design space. (bottom) Full-field (FF) surrogate models predict the response of each node on the part across the design space.

With the possibility of predicting the complete nodal response across a part in real time, FF surrogate models can be applied to predicting a particular mode shape (d), or the displacements for every node in the mode shape. This predicted mode shape can then be used to calculate the MAC value with another mode shape. The resultant MAC calculations would not be limited to a specific modal comparison or a specific reference design (d:r, where r is any mode shape from any reference design). These FF surrogate models provide a much more flexible tool to use in design space exploration.

This research uses both SV and FF surrogate models to predict the MAC values for compressor blades in a jet engine. The transonic Purdue compressor blade [18] will be used for demonstration purposes. The main contributions of this paper are the introduction of using FF surrogate models as a method for predicting MAC values, and a comparison of the relative accuracy and utility of both SV and FF surrogate models for this application. Accuracy is established by comparing predicted responses to actual FEA responses. This is accompanied by a discussion on considerations that should be made when using these methods. It is found that FF surrogate models may be

used, but require a much larger number of samples than SV surrogate models to create reasonably accurate predictions of the MAC values across the design space. Although this study uses a jet engine compressor blade as the test object, the methods described are not limited to this specific geometry.

Related topics and research associated with this study will be presented in the next section, followed by a detailed explanation of the method for creating SV and FF surrogate models to predict MAC values. Section 4.4 will describe the testing process for the accuracy study, and the results will be presented in Section 4.5. This will be followed by a discussion of possible causes, factors, and benefits of the two methods in Section 4.6 and general conclusions in Section 4.7.

4.2 Background

This research uses two forms of surrogate modeling to predict modal assurance criterion values. This section outlines related concepts and research that pertain to the modal assurance criterion and how surrogate models have been used previously to predict data from FEA simulations.

4.2.1 The Modal Assurance Criterion

The modal assurance criterion (MAC) computes the similarity between two mode shape vectors [10, 22, 23]. The mode shape vectors ϕ consist of the modal displacements δ in X, Y, and Z ($\delta_{x,i}$, $\delta_{y,i}$, $\delta_{z,i}$) at each node i of the models [10, 23]. The structure of these vectors is shown in Eqn. 4.1, where N is the total number of nodes in the model.

$$\phi = \begin{bmatrix} \delta_{x,1} & \delta_{x,1} & \dots & \delta_{x,N-1} & \delta_{x,N} \\ \delta_{y,1} & \delta_{y,1} & \dots & \delta_{y,N-1} & \delta_{y,N} \\ \delta_{z,1} & \delta_{z,1} & \dots & \delta_{z,N-1} & \delta_{z,N} \end{bmatrix} \tag{4.1}$$

The MAC value is calculated with Eqn. 4.2 [23], where $\{\phi_D^{(d)}\}_i$ is the i^{th} node of the mode shape vector for Mode d in the first design, D, and $\{\phi_R^{(r)}\}_i$ is the i^{th} node of the mode shape vector for Mode r in the second or reference design, R. This returns a value between 0.0 (no correlation) and 1.0 (exact correlation) [23]. In this way, the MAC value indicates the degree of similarity

between two mode shapes. Generally, a value of 0.9 or above is accepted as an indication that the modes are of the same type [12].

$$MAC_{d:r} = \frac{\left|\sum_{i=1}^{N} \left\{\phi_D^{(d)}\right\}_i \left\{\phi_R^{(r)}\right\}_i\right|^2}{\left(\sum_{i=1}^{N} \left\{\phi_D^{(d)}\right\}_i \left\{\phi_D^{(d)}\right\}_i\right)\left(\sum_{i=1}^{N} \left\{\phi_R^{(r)}\right\}_i \left\{\phi_R^{(r)}\right\}_i\right)} \tag{4.2}$$

The MAC value can be a useful tool in identifying a mode shape [23, 78]. When compared against a known mode, a high MAC value identifies a mode as of the same type. For example, if a first bending mode shape occurs in both the proposed and the reference design, then the MAC value between them will be close to 1.0.

Mode switching is a problem that must be addressed when performing modal comparisons [12,23,82]. Mode shapes are not guaranteed to occur in the same order for every design [10,12,82]. Thus, for some modes, the order in which the particular mode shapes occur may be different.

Fig. 4.2 shows an example of mode shapes that could occur at the fourth and fifth natural frequencies (i.e., Mode 4 and Mode 5). At different points in the design space, the frequency values change and eventually switch places: the mode shape that occurs at the fourth natural frequency at one point in the design space occurs at the fifth natural frequency in another, and vice versa. This behavior is commonly exhibited only in higher modes; the first few modes are often similar across the entire design space, while higher modes tend to switch places more often [83]. For meaningful modal comparisons to be made, this type of behavior must be identified [12, 23, 82].

The MAC value can be used as a way to detect this mode switching [12]. When all designs in a design space are compared against a common reference design and its mode shapes, then the MAC values can indicate high correlation between different modes for different designs. If there is no mode switching present, it is expected that the MAC values for primary comparisons (i.e., d:r, where $d = r$) will be high (0.9 or above) and the MAC values for all secondary comparisons (i.e., d:r, where $d \neq r$) will be low (closer to 0.0) [23, 81]. For designs that have modes that experience mode switching, however, there will be some primary comparisons that have low MAC values and secondary comparisons that have high MAC values.

Fig. 4.3 presents an example of the MAC values between five design modes and five reference modes, where mode switching is present. These results are displayed on a traditional MAC table to simultaneously present all modal comparisons between the new design and the reference

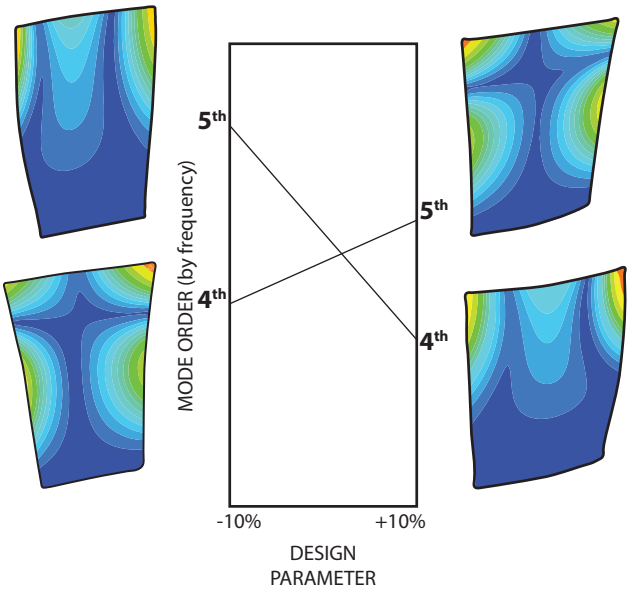

Figure 4.2: A typical example of mode switching as a design parameter is varied.

design. Primary comparisons are highly correlated for Modes 1-3. Modes 4 and 5 have high correlations with opposite modes, as a result of mode switching.

4.2.2 Surrogate Models of FEA Results

To properly train surrogate models with FEA data, it is necessary to first create a set of training data [5, 10, 79, 83, 84]. To do this, a variety of designs are created by using a design of experiments (DOE) to find a set of designs that adequately fill the design space. For a model with n input parameters, each design instance is represented by an n-dimensional vector [26]. These parameters values often control the geometry or conditions of a parameterized FEA model. The parameter values for each design are considered the inputs for the surrogate model.

Once the designs in the DOE are identified, FEA is performed for each design and the results are collected [5]. FEA produces results such as stress or modal displacements for every node on a model. The FEA results are considered the outputs for the surrogate model. The surrogate models are trained on these inputs and outputs, and develop a mathematical function to describe the relationship between them. Because the training depends on samples of the actual data taken from the design space, this is known as a sample-based method [10, 79]. Once trained, the surrogate models can then take in parameter values for new designs to predict the full response of the

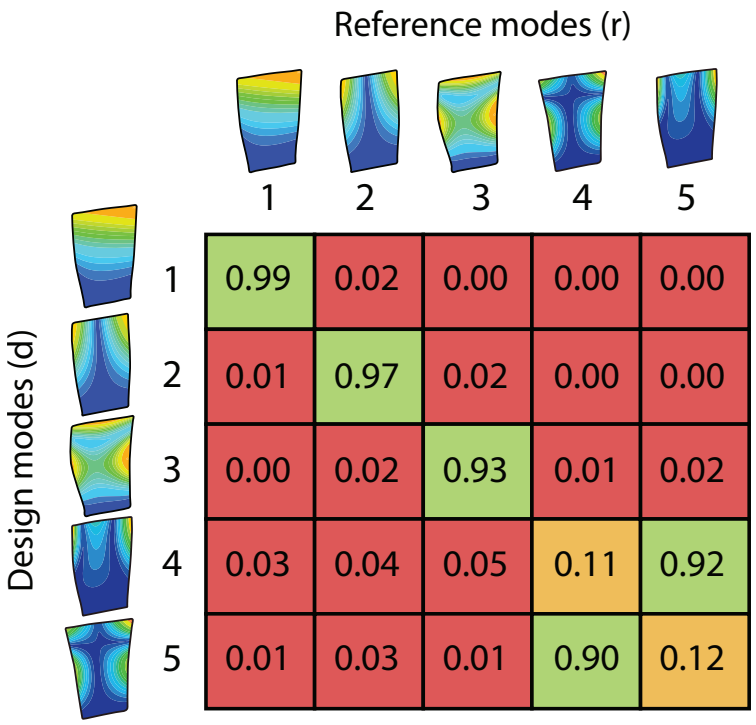

Figure 4.3: An example set of MAC values for all comparisons when mode switching is present.

part. Instead of executing a new simulation each time a design change is made, a properly trained surrogate model can provide the desired output value very quickly without expensive calculations.

Because SV surrogate models can only represent one value at a time, a SV surrogate model can only predict specific values from a complete finite element analysis (FEA) results. Such values may include the maximum stress of a part, the average displacement of a part, or the response of a specific "monitor" node at some location of interest [7, 67]. These types of surrogate models have been used to predict results in many parts of a jet engine, from the compressor and turbine blade vibrations to aerodynamic coefficients experienced in various nozzle types [8, 11, 66, 67]. Geller et al. [7] noted that when single values are related to a geometric location, such as maximum stresses or displacements, they can be more difficult to accurately predict with SV surrogate models.

SV surrogate models have been used to predict MAC values and other modal results. Qin et al. [83] constructed SV surrogate models to predict modal frequency and MAC values in order to update an existing FEA model to more accurately reflect experimental modal results. Bae et al. [81] used surrogate models to predict modal displacements for integrated blade rotors in jet engine turbomachinery. Nobari et al. [12] used surrogate models to map the relationship across a

design space between design inputs and modal frequencies for a complex FEA model in order to make recommendations for design improvement.

Surrogate model sets have been used to predict the stress and displacement response of each node in a FEA model [14–16]. Both Heap et al. and Bunnell et al. were able to approach "real-time" design exploration of finite element models by representing the nodes with surrogate models [14, 16]. Bunnell et al. used a unique surrogate for every node's stress, X location, Y location, and Z location in order to predict the stress and geometric response to changes in geometry. These collections of surrogate models represent an entire field of results from a FEA simulation, and are called FF surrogate model sets. Both Schulz et al. and Bunnell et al. were able to predict the results for an entire three-dimensional finite element model [14, 15].

The FF surrogate modeling method does not share the same limitations of the SV surrogate models. Because the training data for FF surrogate model sets consist of the response at every node, the FF surrogate model sets can predict the full response of the part. With these predicted results, additional benefits are possible. Several implementations of the FF surrogate model sets map the predicted nodal results onto a reconstructed visualization of the FEA model [14, 15]. By contrast, the SV surrogate models cannot give full responses nor enable full visualizations of the response.

Surrogate model accuracy depends largely on the number of training samples used [5]. It has been shown that this is true for both SV surrogate models [11, 12, 85] and FF surrogate model sets [14]. Bunnell et al. established that the accuracy of FF surrogate model sets when predicting steady stress values decreased greatly as the number of training samples used approached zero. Conversely, as the number of training samples increased, they approached 5% normalized root mean square error (NRMSE) [14]. When the training data for surrogate models depends on computationally expensive FEA simulations, it is desirable to get acceptable accuracy with a minimum amount of training data [5]. Acceptable accuracy varies depending on the application and field.

Although the accuracy of FF surrogate models has been demonstrated for steady stress data, it has not yet been established for modal data. Modal data in various forms has been found to be challenging to emulate properly with surrogate models [2, 81, 86]. When using SV surrogate models to emulate mode switching, Bae et al. found that more training samples are needed to predict the data with sufficient accuracy [81]. This is because the modal data is highly nonlinear,

especially when mode switching behavior is present. Because FF surrogate model sets predict each node's behavior across the design space, it is possible that they handle the nonlinearity of modal data differently than do SV surrogate models. This research establishes that FF surrogate modeling methods do not provide an advantage over SV surrogate modeling methods when predicting highly nonlinear modal data.

4.3 Method

This section will discuss the basic principles for creating both SV and FF surrogate models for predicting MAC values. A more specific application will be presented in Section 4.4. The process for creating training data for FF surrogate models uses much of the same basic workflow as described by Bunnell et al., and the interested reader is referred to [14] (especially Fig. 3) for more complete details. This research expands on the previous method by using FF surrogate models to predict modal displacement FEA results instead of steady stress FEA. It also performs calculations (MAC values) on the predicted results.

The training data set for both the SV and FF surrogate models may be constructed using a DOE with n dimensions, where n is the number of desired input parameters. Each collection of parameter values constitutes a unique design instance. The suggested number of designs in the DOE can depend on criteria such as computational resources, model complexity, and number of dimensions.

These parameter values are used to update the geometry of a parametric model. Rather than create a new finite element mesh for every new design, a mesh-morphing process is applied to the baseline design's mesh [2, 14]. This allows all the designs in the training data set to have a common set of relative node numbers and locations [14]. The behavior of a specific node can be determined across all designs in the training data set. The mesh for each design is used in a modal FEA simulation, and the modal displacement results are stored.

These modal displacements are then scaled. Modal analyses produce mode shapes that often are assigned an arbitrary scaling. Although the MAC calculation is invariant to scaling, the surrogate models are not. By using a common scaling on the training data, the surrogate models will predict a more even range of data and are more likely to be accurate [5]. This is done by

scaling each sample design's modal displacements by the largest magnitude displacement [78, 87]. This process is described in the following steps:

1. Nodal displacements at each node are found for a particular design D and a particular mode d:

$$\phi_D^d = \begin{bmatrix} \delta_{x,1} & \delta_{x,1} & \dots & \delta_{x,N-1} & \delta_{x,N} \\ \delta_{y,1} & \delta_{y,1} & \dots & \delta_{y,N-1} & \delta_{y,N} \\ \delta_{z,1} & \delta_{z,1} & \dots & \delta_{z,N-1} & \delta_{z,N} \end{bmatrix} \tag{4.3}$$

2. Magnitude displacements are calculated for each node, i:

$$\Delta_i = \sqrt{\delta_{x,i}^2 + \delta_{y,i}^2 + \delta_{z,i}^2} \tag{4.4}$$

3. The maximum magnitude displacement is found:

$$\Delta_{max} = max(\Delta_1, \Delta_2, ..., \Delta_{N-1}, \Delta_N) \tag{4.5}$$

4. All displacements are scaled by the maximum magnitude displacement:

$$\frac{\phi_D^d}{\Delta_{max}} = \begin{bmatrix} \frac{\delta_{x,1}}{\Delta_{max}} & \dots & \frac{\delta_{x,N}}{\Delta_{max}} \\ \frac{\delta_{y,1}}{\Delta_{max}} & \dots & \frac{\delta_{y,N}}{\Delta_{max}} \\ \frac{\delta_{z,1}}{\Delta_{max}} & \dots & \frac{\delta_{z,N}}{\Delta_{max}} \end{bmatrix} \tag{4.6}$$

4.3.1 SV Surrogate Models

The process for training a SV surrogate model for a specific modal comparison (d:r) is shown in Fig. 4.4 (top). A reference design is chosen a priori, and the reference mode shapes are obtained. For each design in the training DOE, a specific mode shape d is chosen, and the

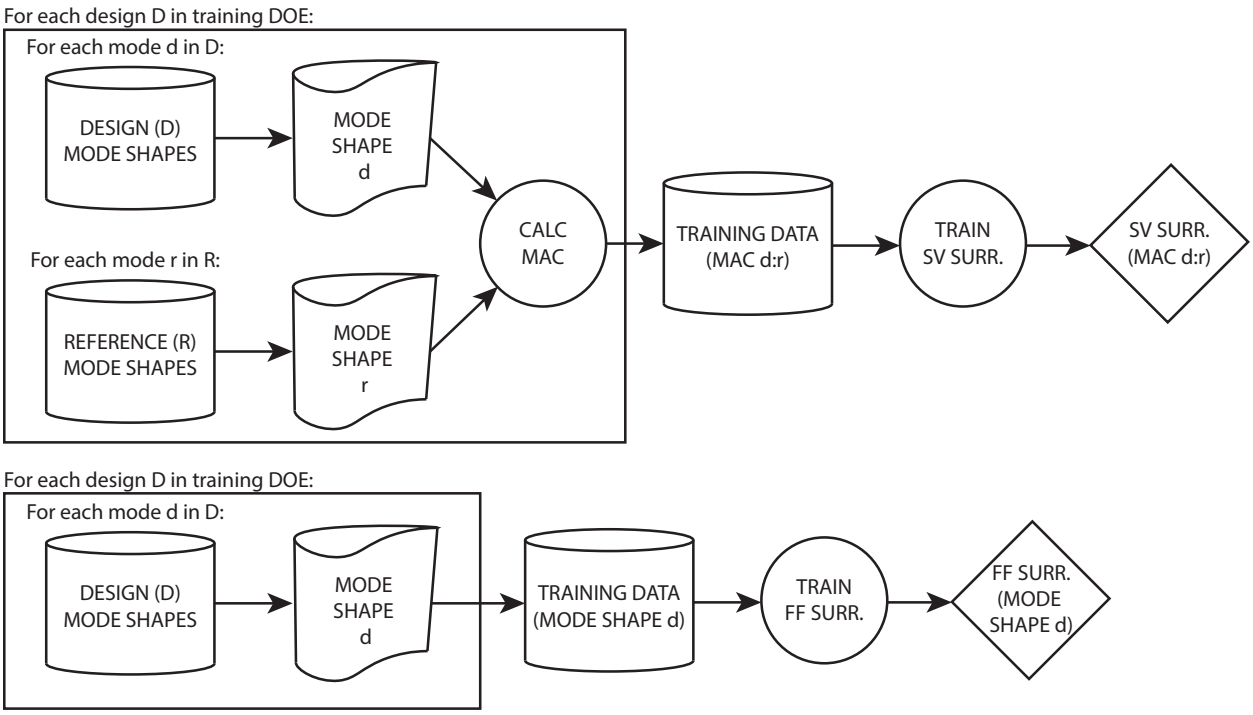

Figure 4.4: The procedure for creating the SV (top) and FF (bottom) surrogate models

MAC values between design mode d and reference mode r are calculated. The collection of MAC values for the comparison d:r constitute the training data for the SV surrogate model. Once the SV surrogate model is trained, the surrogate model can accept a set of parameter values as a design instance and return a prediction of the MAC value for comparison d:r (see Fig. 4.5 (top)). This predicted MAC value will be referred to as a MAC_{SV} value.

Because each SV surrogate model can only predict a specific modal comparison, multiple SV surrogates are needed to predict all modal comparisons. The MAC value is calculated for every modal comparison between the reference design and each design in the training data set. For d design modes, and r reference modes, there will be dr MAC values for every design in the training data set. A different SV surrogate model is trained on the MAC values for each modal comparison. Thus, for dr MAC values, dr SV surrogate models are needed. Together, these SV surrogate models can predict dr different MAC_{SV} values.

4.3.2 FF Surrogate Models

The FF surrogate models are created differently. Fig. 4.4 (bottom) illustrates this process. The scaled modal displacements for a particular mode d constitute the training data for a FF surrogate model set. For d modes, there will be d FF surrogate model sets. Once the surrogate models are trained, each FF surrogate model set can accept new parameter values as a design instance and predict the nodal displacements of a mode shape.

These predicted modal displacements, or mode shapes, can then be used to calculate MAC values. A MAC value is calculated with a predicted mode shape for the new design and a mode shape from a reference design, as described in Fig. 4.5 (middle). These MAC values will be referred to as MAC_{FF} values. Because the MAC_{FF} values are not predicted directly, the results of a single FF surrogate model set are not limited to a single modal comparison as they are with SV surrogate model sets. For a reference design with r modes, a FF surrogate model set trained to predict a particular mode shape can be used to calculate r MAC_{FF} values. If there are d FF surrogate model sets (one for each of d modes), then there can be dr MAC_{FF} values calculated.

By predicting the mode shapes and not specific MAC values, the FF surrogate models are also not limited to comparisons with a single reference design. MAC_{FF} values can be calculated between the predicted mode shape and the mode shapes of any number of reference designs. The FF surrogate model sets allow the user much more freedom to quickly obtain a wide variety of modal comparisons with a single prediction. This enables much more flexibility and extends the utility of a FF surrogate model set even if criteria for making comparisons changes during design space exploration.

The process outlined in this paper trains SV surrogate models to predict specific modal comparisons and FF surrogate model sets to predict specific mode shapes. The modes are not identified prior to training; the surrogate models predict the modes according to the original order (e.g., a FF surrogate model set that was trained on data for Mode 4 will predict the mode shape that occurs at the fourth natural frequency across the design space). If mode switching occurs, the surrogate models should reflect this behavior.

4.4 Accuracy Comparison

This section describes the conditions and process for evaluating accuracy of the surrogate models in this study. The purpose is to test and compare the abilities of SV and FF surrogate models to predict MAC values. This study provides general factors to consider when selecting a surrogate model method for predicting the MAC values in design space exploration, and suggests ways to address error in the models.

The error of a surrogate model refers to its ability to accurately predict the kind of data with which it was trained. In this research, the surrogate models predict the results of a modal FEA simulation. Surrogate models may be adjusted to more accurately represent a FEA simulation by increasing the number of training samples, and a FEA simulation may be adjusted to more accurately represent reality by increasing the fidelity of the model, adding parameters, and other means. While using more accurate FEA simulations is ideal, the objective of this research focuses on validation of surrogate model accuracy, not validation of a finite element model [12]. Thus, the accuracy of the surrogate models is evaluated by comparing the surrogate models' results and the FEA results. Using FEA results of validation samples as a way to evaluate accuracy of values predicted by surrogate models is consistent with previous work [10, 11, 76]. ANSYS 17.0 was the FEA solver used in this research.

4.4.1 Model

This test will use the transonic Purdue blade model of a jet engine compressor [18], which is a "blade-alone" model. This particular model has been meshed with 25000 nodes. The Purdue blade was developed for research purposes and is a general representation of a compressor blade. The model consists of various airfoil profiles, or airfoil shapes, with connecting surfaces. It has been parameterized with the five parameters and bounds in Tab. 4.1. These parameters change the geometry of the model by adjusting the profiles at the root and tip of the blade. The Angle parameter adjusts the difference in the angle between the profile at the root and the profile at the tip. The two Chord parameters adjust the chord length at the root and at the tip. The Lean parameter offsets the tip profile in a direction perpendicular to the root profile's chord, while the

Sweep parameter offsets the tip profile in the same direction as the root profile's chord. This parameterization and associated controls are discussed in [14].

Table 4.1: Parameter bounds

Parameter	Lower Bound	Upper Bound
$Angle$	-20 deg	20 deg
$Chord_{root}$	20.32 mm (0.8 in)	30.48 mm (1.2 in)
$Chord_{tip}$	20.32 mm (0.8 in)	30.48 mm (1.2 in)
$Lean$	-7.62 mm (-0.3 in)	7.62 mm (0.3 in)
$Sweep$	-7.62 mm (-0.3 in)	7.62 mm (0.3 in)

The reference design used in this study is the baseline design; i.e., it is the design where each parameter is at the center of its range ($Angle = 0$ deg, $Chord_{root} = 25.4$ mm (1.0 in), $Chord_{tip} = 25.4$ mm (1.0 in), $Lean = 0$ mm (0 in), $Sweep = 0$ mm (0 in)). Only the surface nodes are used in order to reduce computational costs. In the MAC calculation, it is appropriate to not include every node of a mode shape as long as there are the same number of corresponding nodes in both mode shapes in the comparison. Some methods include using only a representative area of the part, a critical area where displacements are deemed more significant than other areas [12], or only the surface nodes. The nodes with the highest, or most significant, displacements for each mode will occur on the surface of the part; therefore only using the results on the surface nodes is an acceptable choice for this study.

4.4.2 Data

In order to create and then test the surrogate models, two data sets were created for both the SV and FF surrogate models: a training data set and a test data set. These will be briefly described in this section.

Training Data

As described in Section 4.3, a training set of design parameters (inputs) and FEA results (outputs) must be generated to train the surrogate models. An optimized Latin Hypercube was chosen because it is a space-filling design [5, 12, 77, 83]. Because the number of samples needed

for accurate surrogate models in this context is unknown, this study used four training data sets with different numbers of training samples: 100, 250, 500, 800, and 3500. Each of these samples was used to update the parametric compressor blade model. The updated model was used to conduct a modal analysis in ANSYS, and nodal displacement data was extracted for each mode of each sample.

The MAC values for the comparison between every mode in the training data set and every mode of the reference design were calculated and used to train dr SV surrogate models. The SV surrogate models predict the MAC values across the design space for a particular mode comparison. This was repeated for each training data set in order to evaluate accuracy with different numbers of training samples (SV_{100}, SV_{250}, SV_{500}, SV_{800}, and SV_{3500}).

The displacements for each mode shape in the training data set were used to train d FF surrogate models. The FF surrogate model sets predict the nodal displacements across the design space for each mode, and MAC values are calculated with these predicted results. This was again repeated for each training data set in order to evaluate accuracy with different numbers of training samples (FF_{100}, FF_{250}, FF_{500}, FF_{800}, and FF_{3500}).

For both SV surrogate models and FF surrogate model sets, radial basis functions (RBF) were used. While other functions could be used to build the surrogate models, the authors chose to use RBFs based on previous research and experience [14]. The RBF is an appropriate choice for sparse training data in design spaces with many parameters [14, 16]. The RBF has also been recommended as a good choice of surrogate model function for capturing highly nonlinear behavior [76].

Test Data

In order to test the accuracy of the surrogate models, a separate set of designs must be created. The surrogate models tend to be more accurate when predicting information that is close to information used in the training data set [12]. If the surrogate models are used to predict the same designs used in the training data set, the results may appear artificially accurate. Thus the test designs must be distinct from those designs in the training data set [67, 77]. The objective is to find points for comparing the response of the surrogate models to the actual data across the full span of the design space.

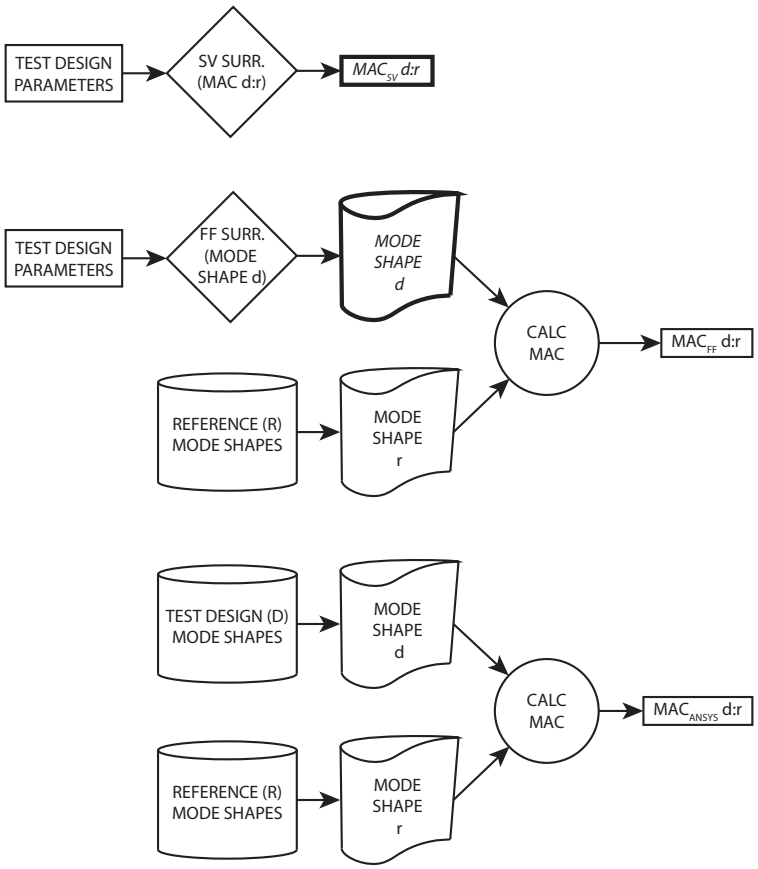

Figure 4.5: The process used in the test procedure to obtain the MAC_{SV} (top), MAC_{FF} (middle) and MAC_{ANSYS} (bottom) values. Shapes in bold are directly predicted by a surrogate model (MAC_{SV} for SV surrogate models and mode shape d for FF surrogate models)

A new DOE and data set were created with 500 different designs using the same process as before. The MAC values for each mode shape in this data and each mode of the chosen reference design were calculated. Because they were calculated with data from ANSYS, and were not predicted directly or calculated from predicted data, they will be referred to as MAC_{ANSYS} values. These MAC values formed a test data set, and were used for validation against values predicted by the surrogate models [10, 76].

The number of designs in the test data set was also independent of the number of designs in the training data set. Given that the same parameters and parameter bounds are used in different surrogate models, the same test data set can be used to test accuracy across surrogate models with any number of training samples. Thus, for this study, the test data set with 500 designs was used to test the accuracy of the 100, 250, 500, and 800 sample surrogate models.

4.4.3 Test Procedure

The following process was used for each design in the test data set. It is represented graphically in Fig. 4.5.

1. Using the ANSYS data from the test data set:
   (a) Obtain the modal displacements for each mode in each design of the test data set
   (b) For each mode d of the test design:
      i. Calculate the MAC_{ANSYS} values between the displacements for the test design mode d and each mode r in the reference design
2. Obtain the predicted MAC values for the test design with the surrogate models
   (a) Using the FF surrogate model sets:
      i. For each mode d of the test design:
         A. Obtain the predicted full field of modal displacements with the appropriate FF surrogate model set
         B. Calculate the MAC_{FF} value between the predicted displacements for the test design mode and each mode r in the reference design
      ii. Calculate the error between the MAC_{FF} values and the MAC_{ANSYS} values (E_{FF})
   (b) Using the SV surrogate models:
      i. For each mode comparison d:r of the test design modes d and reference design modes r:
         A. Use the appropriate SV surrogate model for the comparison d:r to predict the MAC_{SV} value between the test design and the reference design
      ii. Calculate the error between the MAC_{SV} values and the MAC_{ANSYS} values (E_{SV})

4.5 Results

This section describes the results from the accuracy study. Error between predicted MAC values (MAC_{SV}, MAC_{FF}) and MAC_{ANSYS} values is measured by the magnitude difference. These

are simple differences in magnitude of the various values, and are calculated by Eqn. 4.7 and Eqn. 4.8. The magnitude error was chosen over percent error because, with small values, the percent error equation can produce large error and convey misleading information about the true difference. Since all MAC values are between 0.0 and 1.0, the magnitude error is an appropriate choice. These magnitude error values were obtained for all mode comparisons for each of the 500 designs in the test DOE; that is, there are 25 MAC values for each design in the test DOE, and thus 25 E_{SV} values for the SV surrogate models and 25 E_{FF} for the FF surrogate model sets for each design.

$$E_{SV} = |MAC_{ANSYS} - MAC_{SV}| \tag{4.7}$$

$$E_{FF} = |MAC_{ANSYS} - MAC_{FF}| \tag{4.8}$$

Fig. 4.6 shows all of the magnitude error results for the 500 sample surrogate models (SV_{500}, FF_{500}). The MAC values from the surrogate models ($MAC_{SV,500}$, $MAC_{FF,500}$) are compared to the ANSYS MAC values (MAC_{ANSYS}) for every mode comparison of every design in the test data set. With perfect accuracy (E_{SV}=E_{FF}=0.0), the data would appear as a straight, diagonal line between the origin and (1.0, 1.0). Because of error in the surrogate models, there is deviation in the data. The points are colored by the level of magnitude error.

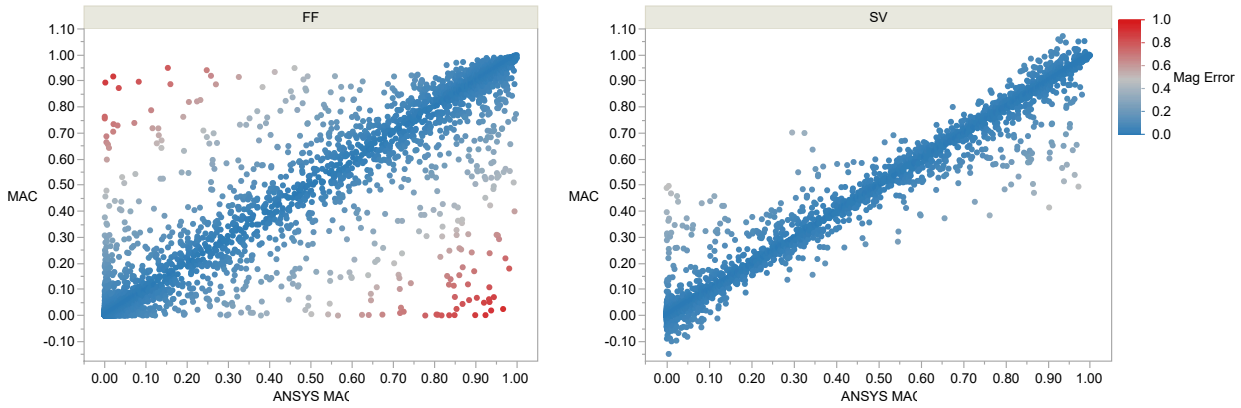

Figure 4.6: The MAC_{FF} (left) and MAC_{SV} (right) MAC values from the FF_{500} and SV_{500} surrogate models for all test designs compared to the MAC_{ANSYS} values. Coloring reflects the magnitude error values, with red being higher error and blue being lower error.

This image shows that, overall, FF_{500} has greater magnitude errors than SV_{500}. There are certain designs and mode comparisons where FF_{500} predicts the opposite values as the MAC_{ANSYS} ($MAC_{FF,500} = 1.0$ but MAC_{ANSYS}=0.0). While SV_{500} is more accurate, it still has a large degree of error.

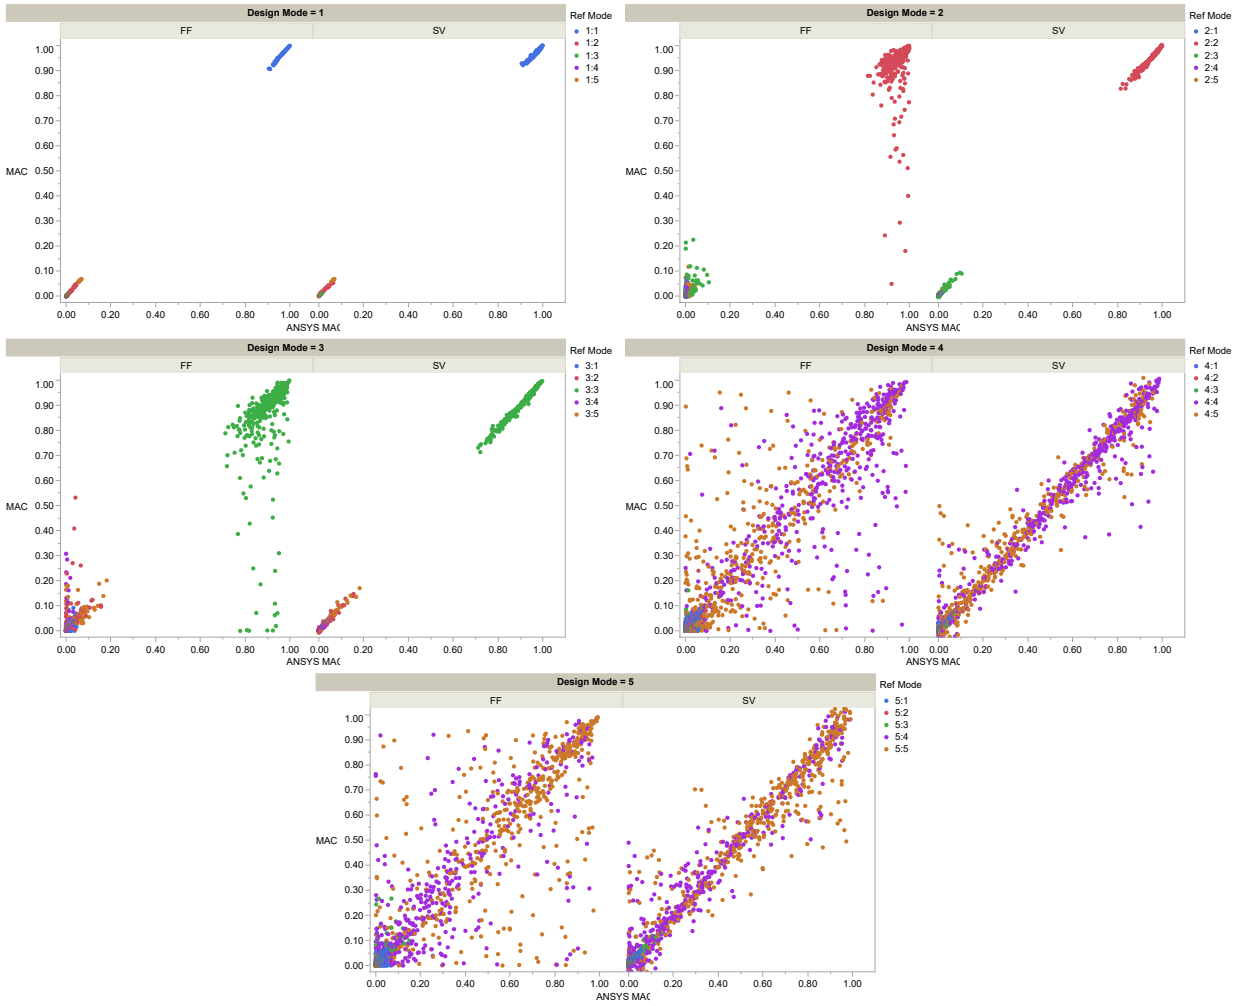

Figure 4.7: The predicted FF and SV MAC values from the SV_{500} and FF_{500} surrogate models for all test designs compared to the MAC_{ANSYS} values separated by each mode. Points are colored by the reference mode used in each MAC comparison.

Fig. 4.7 shows the same data as Fig. 4.6, but it is separated by the design modes into subplots. The comparisons with each design mode and the reference modes are represented by different colored points. The figures show how well the SV and FF surrogate models predict the

MAC_{ANSYS} values. For some comparisons, the MAC_{ANSYS} values are closer to 0.0, while others are closer to 1.0.

While the SV surrogate models still have less severe errors for the 500 sample training data set, this figure demonstrates how the surrogate models' accuracy depends on the particular design modes and reference modes involved in the MAC comparisons. For Modes 1-3, no mode switching behavior is present: the MAC_{ANSYS} values are either close to 0.0 (no correlation) or close to 1.0 (high correlation). As expected, the high correlations occur for primary comparisons while the low correlations occur for all other secondary comparisons. The SV surrogate models for the most part represent the same behavior as the MAC_{ANSYS} values for all comparisons in Modes 1-3 ($0.00 \leq E_{SV} \leq 0.04$, $E_{SV,avg} = 0.0009$). The FF surrogate models, however, have more error. While the FF surrogate models predict MAC values close to the true MAC_{ANSYS} values for all comparisons with design mode 1, the next two plots show how the MAC_{FF} values for the specific comparisons 2:2 and 3:3 have significant deviations from the true values ($0.00 \leq E_{FF} \leq 0.90$, $E_{FF,avg} = 0.05$). This tendency towards higher error in the FF surrogate model increases with each mode.

On Fig. 4.7, the MAC_{ANSYS} values for comparisons between Modes 4 and 5 do not occur only near 0.0 or 1.0; thus, in this design space, this is an indication that mode switching is occurring. For example, at various places in the design space the MAC_{ANSYS} values for 4:4 are close to 1.0 and 4:5 are close to 0.0, while at other places the MAC_{ANSYS} values for 4:4 are closer to 0.0 and 4:5 are closer to 1.0. The same behavior is seen between 5:4 and 5:5. There are also a range of designs in the design space where these modes are in the process of switching places (e.g., both 4:4 and 4:5 are closer to 0.5 than either 0.0 or 1.0).

The MAC_{SV} values for Modes 4 and 5 represent the same general behavior as the MAC_{ANSYS} values, but have greater deviations than previous modes ($0.00 \leq E_{SV} \leq 0.50$ for design modes 4 and 5, $E_{SV,avg} = 0.019$). The MAC_{FF} values have a much greater range of error ($0.00 \leq E_{FF} \leq 0.94$ for design modes 4 and 5, $E_{FF,avg} = 0.05$). This primarily occurs in comparisons between Modes 4 and 5 ($E_{FF,avg} = 0.10$ for 4:4 and $E_{FF,avg} = 0.12$ for 5:5). There are places in the design space where, again, MAC values predicted by the FF surrogate models are exactly opposite those found by ANSYS. While both SV and FF surrogate models have more designs with large magnitude errors at the higher modes, the SV surrogate models do so to a much lower degree.

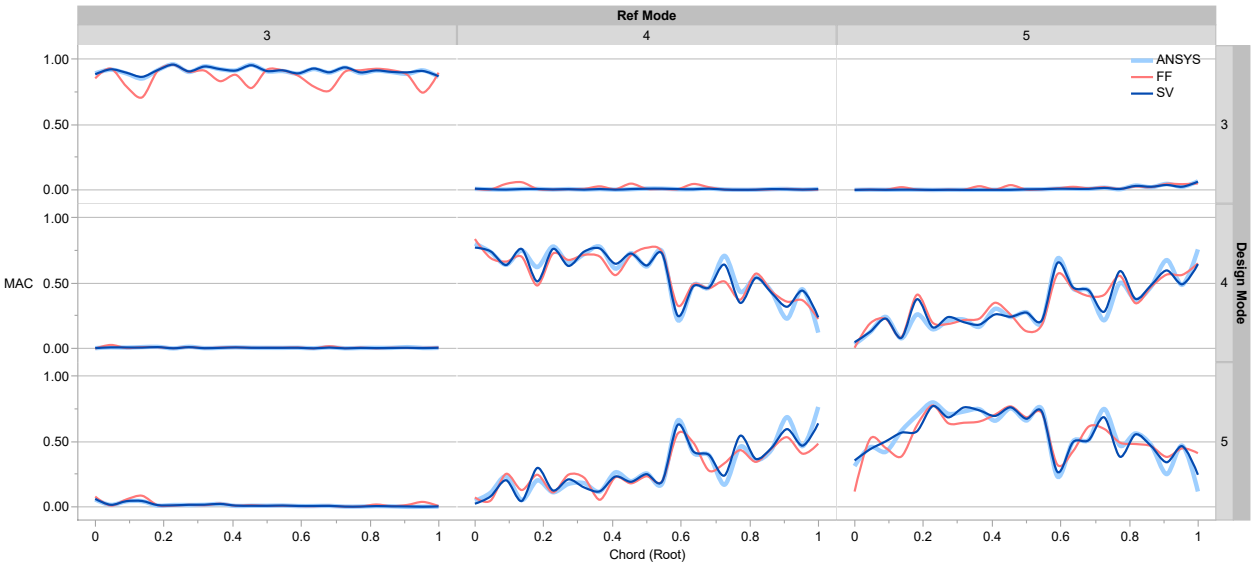

Figure 4.8: The moving average trends of the MAC_{ANSYS}, MAC_{FF}, and MAC_{SV} values across the design parameter $Chord_{root}$. The FF and SV lines indicate how closely the predicted MAC values follow the trends of the MAC_{ANSYS} values. Mode switching is present in the comparisons between Modes 4 and 5.

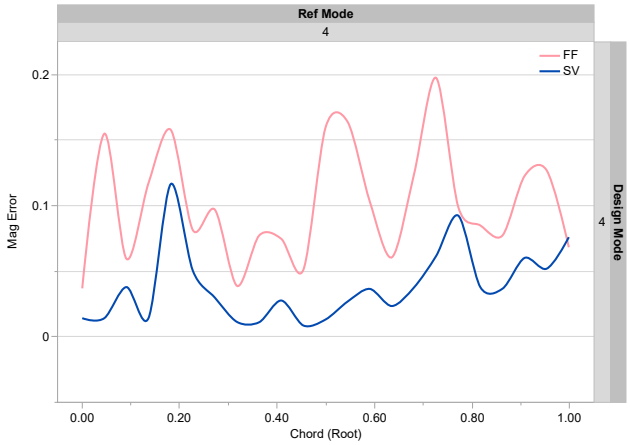

Figure 4.9: The trends of E_{FF} and E_{SV} for F_{500} and SV_{500} across the design parameter $Chord_{root}$ for comparison 4:4. The SV surrogate models have lower error at almost every part of the design space.

4.5.1 Fit to Nonlinear Data

These plots establish that, with 500 samples, the MAC values as calculated from mode shapes predicted by the FF surrogate models have a higher degree of error than the MAC values predicted by the SV surrogate models. This is especially true for modes where mode switching

is present. To understand why, it is helpful to look at how the MAC values change throughout the design space. For simplicity, Fig. 4.8 shows just the root chord parameter, $Chord_{root}$, and the interaction between the last three modes (Modes 3-5) in both the new design and the reference design. The moving averages of the MAC values for each modal comparison are represented as separate lines. This figure gives an indication of how well, on average, the SV and FF surrogate models predict the actual ANSYS data.

The top row of Fig. 4.8 shows the modal comparisons involving Mode 3 from the test DOE and Modes 3, 4, and 5 from the reference design. Because there is no mode switching in Mode 3, the comparison of 3:3 has high values close to 1.0 (high correlation) and the comparisons of 3:4 and 3:5 have low values close to 0.0. This is true for the MAC_{ANSYS}, MAC_{FF}, and MAC_{SV} values across every value of $Chord_{root}$. However, the MAC_{FF} values for comparison 3:3 have much more variability than the MAC_{SV} values.

In the next two design modes, mode switching is present. If no mode switching were present, comparisons 4:4 and 5:5 would be close to 1.0, like comparison 3:3. However, the data shows that, as $Chord_{root}$ changes, so does the correlation between Modes 4 and 5 in the test designs and Mode 4 and 5 in the reference design. When the $Chord_{root}$ values are low, the MAC values are high for comparison 4:4 (high correlation) and are low for comparison 4:5 (low correlation). However, as $Chord_{root}$ increases, Mode 4 of the test design becomes less correlated with Mode 4 of the reference design and more highly correlated with Mode 5 of the reference design. Similar behavior is observed in Mode 5 of the test designs.

While this is the general trend, the MAC_{ANSYS} values show that the data is very erratic. At different points, the modes cross back and forth (e.g., when $Chord_{root}$ has values between 0.5 and 0.7). This is partially because the points are also varied across all five geometric parameters, not only $Chord_{root}$, and it is difficult to represent in two dimensions. However, the mode ordering for mode switching across a design space can be inherently nonlinear. Examining the moving average lines for MAC_{FF} and MAC_{SV}, there are certain areas where the surrogate models fail to capture the behavior of the MAC_{ANSYS} data. For example, at $Chord_{root} = 0.60$, both SV and FF surrogate models struggle to predict the MAC values for comparisons 4:4, 4:5, 5:4, and 5:5; however, the moving average lines show that the FF surrogate models predict these values with less accuracy than the SV surrogate models. There are regions where the FF surrogate models predict very

poorly ($Chord_{root} = 0.5$ for design mode 4 and $Chord_{root} = 0.7$ for design mode 5). Overall, the SV surrogate models predict the MAC_{ANSYS} values much more closely than the FF surrogate models.

As a closer illustration of this data, the values of E_{SV} and E_{FF} for comparison 4:4 as $Chord_{root}$ varies are shown in Fig. 4.9. These plots illustrate the difference in accuracy between the FF and SV surrogate models. The 4:4 SV surrogate model has less magnitude error than the FF surrogate model set at almost all locations.

Fig. 4.8 illustrates that the MAC data being predicted becomes highly erratic and nonlinear when mode switching is present. This is one possible reason why FF_{500} has less accuracy than SV_{500}. The moving average lines in these plots also show that there is possible underfitting in predicting the behavior of the MAC_{ANSYS} values. To better represent highly nonlinear data, using a greater number of training samples is a typical strategy for increasing the accuracy [12, 81].

4.5.2 Number of Training Samples

To observe the effect of the number of training samples on the surrogate model accuracy, the analyses shown above for the 500 sample surrogate models were repeated for the SV and FF surrogate models created with 100, 250, 500, 800, and 3500 sample training samples.

To represent the accuracy of the SV and FF surrogate models between different numbers of training samples, the root mean square error (RMSE) of the MAC values at each comparison was measured. The RMSE is a common metric for measuring the accuracy of a surrogate model over a design space [5, 12, 77]. This allows a direct comparison between the SV and FF surrogate models for the results of each of the 25 different modal comparisons. The behavior of these RMSE values was tracked as the number of training samples used in the surrogate models was varied.

The changes in RMSE for each modal comparison are shown in Fig. 4.10. Each subplot displays the comparisons with a different design mode (e.g., the first plot of Fig. 4.10 shows comparisons with design mode 1: 1:1, 1:2, 1:3, 1:4, and 1:5). The dashed lines show the RMSE for the MAC_{SV} values and the solid lines show the RMSE for the MAC_{FF} values over the different numbers of training samples.

These figures indicate that, for all comparisons, increasing the number of samples reduces the RMSE. For Mode 1, both FF and SV have extremely low RMSE, even with only 100 training

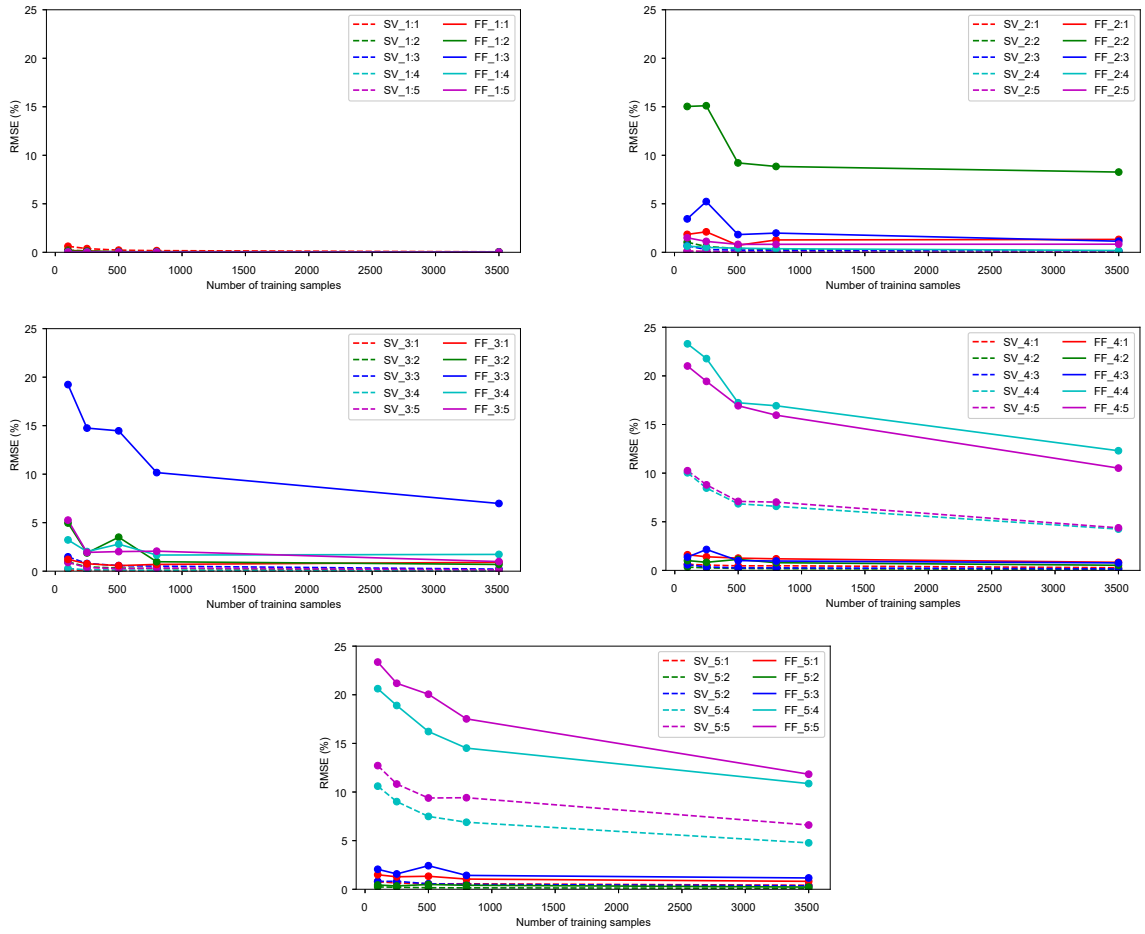

Figure 4.10: The effect of different numbers of training samples on the RMSE of the surrogate model predictions of each comparison. Each plot shows comparisons with a different design mode. The RMSE of the MAC_{SV} values are represented with dashed lines, and MAC_{FF} values are represented with solid lines.

samples. For Mode 2, the 2:2 MAC_{FF} values had a 15.1% RMSE. With 500 samples, this decreased to 9.2% RMSE, and with 3500 training samples it reached 8.3% RMSE. All other comparisons using the FF surrogate models reached under 2.0% RMSE with 500 samples or higher. For the SV surrogate models, all comparisons for Mode 2 had below 2.0% RMSE for any number of training samples tested. Mode 3 had similar behavior, with FF surrogate models reaching 7.0% RMSE comparison 3:3 after 3500 samples, and all other comparisons using the FF surrogate models reaching under 2.0% RMSE with 800 samples or higher. Again, all comparisons for Mode 3 had below 2.0% RMSE for any number of training samples tested.

Modes 4 and 5 illustrate how comparisons that exhibit mode switching affect the number of samples needed. While non-switching comparisons do not require very many training samples for fairly low RMSE, comparisons between the design and reference Modes 4 and 5 (4:4, 4:5, 5:4, and 5:5) require many more samples than previous modes. This is true for both FF and SV surrogate models, although the SV surrogate models required less training samples than the FF surrogate models. Even at 3500 samples, the FF surrogate models still have between 10.0-12.0% RMSE. By contrast, these are about the same RMSE ranges that can be expected of the SV surrogate models with only 100 samples. The SV surrogate models have 6.0-9.0% RMSE for these modes at only 800 training samples. With 3500 samples, this goes down to 4.0-6.0%.

These plots demonstrate how, if the mode comparisons that are to be predicted are non-principal comparisons (e.g., 1:2, 2:5, 5:3, etc), a relatively low number of samples is required to create a reliably accurate design space. However, for principal comparisons (e.g., 1:1, 2:2, 3:3, etc), the FF surrogate models require a much greater number of samples to achieve this accuracy than do SV surrogate models. Further, when mode switching is present, both SV and FF surrogate models struggle to predict the behavior accurately. Increasing the number of training samples does indeed help the problem; however, SV surrogate models require far fewer samples than do FF surrogate models. The FF surrogate models still have relatively high error even after using a near-prohibitively high number of training samples, indicating that even more samples should be used.

4.6 Discussion

The results of the previous section show that SV surrogate models are more accurate than FF surrogate model sets when considering the same parameters and training sample size. Both methods were affected by the erratic shape of the data due to mode switching, as well as the increasing complexity of higher modes.

These results support the conclusions of previous research. Just as previous methods had more challenges in using SV surrogate models to emulate mode shapes that experience mode switching behavior than other mode shapes [81, 82], this appears to be true for FF surrogate models as well. As Bae et al. [81] observed, this is likely due to the highly nonlinear behavior that accompanies mode switching.

It has been recommended that, for SV surrogate models, higher numbers of training samples can improve predictions of this kind of nonlinear behavior [81]. This research confirms that this is true for both the SV and FF surrogate methods used in this study. It also adds the conclusion that FF surrogate models require up to 3400 more training samples to achieve the same level of accuracy as SV surrogate models in making these predictions.

4.6.1 SV vs FF Surrogate Model Accuracy

Although this does not appear to be the case in this study, FF surrogate models can be more accurate than SV surrogate models in some applications. Consider the case of predicting the maximum steady stress of the same part used in Section 4.4. A SV surrogate model is created that is trained on the maximum steady stress of 500 training samples, and directly predicts the maximum steady stress. A FF surrogate model set is created that predicts the steady stress response of every node on the part across the design space. The maximum steady stress is then identified from this set of predicted nodal steady stresses.

Tab. 4.2 shows the results of this case. The improvement gained refers to how much lower the error for the FF max stress is than the error for the SV max stress, and follows the equation $improvement = 1 - (FF/SV)$. In this data set, the FF surrogate models predict the max steady stress with a mean percent error that is 53.54% lower than the mean percent error for the SV surrogate model's max steady stress. The average magnitude difference also shows that, on average, the FF surrogate models predict a closer value to ANSYS than the SV surrogate model for max steady stress.

Table 4.2: FF vs SV max steady stress error

Value	**Avg. %Error**	**Avg. Mag. Error**
FF $\sigma_{s,max}$	2.30%	6.7 MPa (967.2 psi)
SV $\sigma_{s,max}$	4.94%	12.3 MPa (1779.9 psi)
Improvement	**53.54%**	**45.66%**

The question of why the FF surrogate models perform better for predicting the maximum steady stress than the MAC values in this study is worth considering. The maximum steady stress

depends on the accurate prediction of each individual node, and the maximum steady stress is the predicted value of a single node. Fig. 4.11 presents a simple example of how, when two nodes are present, a SV surrogate model may inadvertently create a smooth function that misses the detail present in the system. When using a FF surrogate model, the behavior of the nodes is more fully captured, and thus the FF surrogate model is better at predicting the results of the max stress. Though there is some error present in the prediction of an individual node's results, it has less effect than the error in the SV surrogate model's interpolation. This same principle works when using higher numbers of nodes, such as the results of Tab. 4.2.

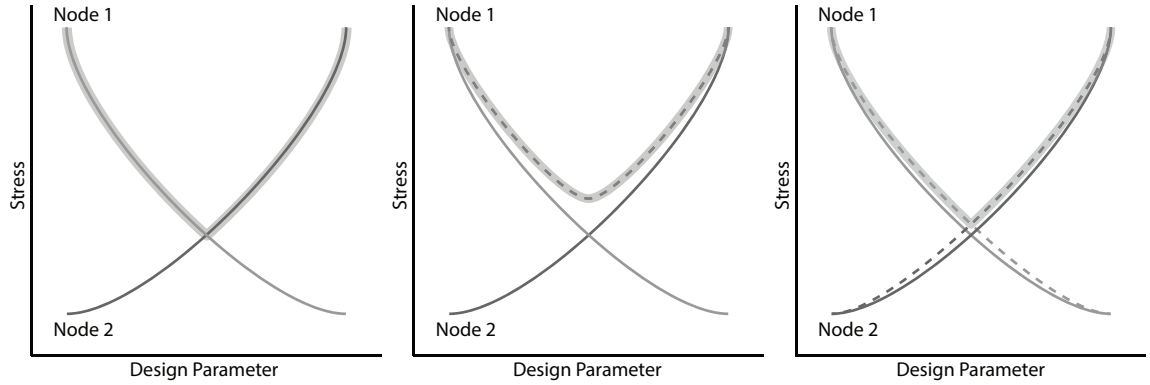

Figure 4.11: An example of how a FF surrogate model set predicts maximum stress better than a SV surrogate model on a 2-node model. (left) the behavior of each node (solid lines), with the true maximum stress behavior (shaded); (middle) the predicted maximum stress behavior (shaded) by a SV surrogate model (dashed); (right) the predicted behavior of each node by a FF surrogate model set (dashed) with the predicted maximum stress behavior (shaded). TODO: add labels.

However, unlike predicting a single nodal response like the maximum stress, the MAC calculation combines the nodal responses across the entire part into one value. Unlike a SV surrogate model, the FF surrogate model has error at each individual node. When using a calculation that combines all the nodal responses, these small errors are likely compounded, causing greater error in the calculated value. The highly nonlinear nature of the modal data contributes to this.

4.6.2 Improving Accuracy

The accuracy of the surrogate models depends greatly on the number of training samples [10, 81, 85]. The results shown in Fig. 4.10 reflects this. As the number of training samples are

increased, the RMSE decreases. As a general rule, as complexity and variance in the data increases, more training samples are required. This is affected by factors such as the types of parameters used in the design space as well as how much variance exists in the data being predicted. When mode switching occurs in the mode shapes being predicted, the data becomes highly nonlinear, and the number of training samples required increases dramatically.

While these trends were observed in both SV and FF surrogate models, it occurred to a much lesser degree with the SV surrogate models than with the FF surrogate models. Because each training sample requires a full modal analysis, the cost of creating a surrogate model begins to be outweighed by the cost of creating training data if the number of training samples is extremely high. The higher numbers of training samples required may make FF surrogate models less advantageous.

Because the detection and identification of the switched modes would be done prior to training, it would negate the purpose of using the surrogate models to detect mode switching, but could still preserve other benefits. The surrogate models would still be able to show the range of variation across the design space for specific, identified modes. Evaluating this method could be a subject for future work.

4.6.3 Benefits

While FF surrogate models were found to be less accurate than SV surrogate models when predicting MAC values, there are associated benefits that may make either of these options desirable, provided that enough samples are used for sufficient accuracy. Both methods (SV and FF) can be used in real-time design space exploration to produce the results of all modal comparisons as a direct and real-time response to changes in input parameters. As inputs are modified, this becomes a way to detect mode switching and perform mode identification; high MAC values indicate that a certain proposed design has a positive correlation with a reference design, while low MAC values indicate poor correlation. When traversing the design space, as mode switching occurs, the predicted MAC values for the relevant modes will tend towards 0.5, and then switch places. A designer can easily explore and determine which areas of the design space exhibit this behavior. These MAC results could be presented in a classic MAC table, such as Fig. 4.3, that dynamically updates as parameters are changed.

In addition, the FF surrogate model sets enable a more general approach. Because they predict the modal displacements for every node of a proposed design's mode shape, these displacements can be used to calculate the MAC value against any mode and any reference designs. In other words, the reference design and reference modes can be chosen and changed at runtime. This allows designers to have more flexibility in the design process, without needing to train further data. By contrast, the SV surrogate models are trained on specific MAC values for specific modal comparisons with a specific reference design, and cannot be changed at runtime.

It is also possible for FF surrogate models to display the predicted modal displacements on a three-dimensional representation of the part. In a similar manner to the methods shown by Bunnell et al. [14], these predicted displacements can be shown as contours on the visualization, and update in response to new parameter inputs in real time. This allows a designer to quickly explore the design space and gain intuition about the behavior of the mode shapes with instant visual feedback. This speed depends on the fidelity (node count) of the model as well as the number of parameters in the design space [14].

While the FF surrogate models allow for more flexibility and visualizations, they require far more training samples than do SV surrogate models to produce reliable results. The user must consider whether or not the associated benefits are worth the resources required to create the training data. If not, the SV surrogate models provide a good alternative. With far fewer training samples required, creating these surrogate models is less computationally expensive. They allow a user to explore the changes in MAC values across the design space very quickly, as well as detect mode switching behavior, in the same way as the FF surrogate models. Though they cannot adapt to show comparisons with new reference designs, less resources are needed to create new training data and train new SV surrogate models.

If SV surrogate models are chosen, one potential pitfall is that there is no guard against predicting values outside the range of acceptable MAC values: sometimes they predict values below 0.0 and above 1.0. This can be seen in the SV MAC values in both Fig. 4.6 and Fig. 4.7. This behavior seems to occur in the areas where the surrogate model has most difficulty matching the erratic nature of the real system, such as in cases of mode switching. It is not present in FF surrogate models because the MAC calculations are performed after predictions are made, and the MAC equation does not produce results outside of that range. However, the SV surrogates,

although trained on proper MAC values, can produce a regression fit to nonlinear data that causes the predicted values to vary outside the range.

4.7 Conclusions

This research analyzed the use of surrogate models to predict the MAC values between two different mode shapes, specifically for a jet engine compressor blade. Although this study uses a jet engine compressor blade for its study, the methods described are not limited to this specific geometry.

Previously, SV surrogate models had been used to predict these MAC values; this research applied FF surrogate models to the same problem, and compared the results of the SV and FF surrogate model methods. It was proposed that using surrogate models to predict the MAC values for various comparisons between mode shapes across a design space and a reference mode would allow a user to interactively identify modes and detect mode switching in the design space.

It was found that the accuracy of both methods is affected by the number of the parameters in the design space and the number of training samples used to train the surrogate models. The number of training samples required is greatly increased when trying to predict mode switching behavior. While this is true for both SV and FF surrogate model methods, the FF surrogate models require a much larger number of training samples to achieve the same level of accuracy as the SV surrogate models.

The benefits of using the FF surrogate model method were described. The usefulness of these benefits varies in accordance with the level of accuracy. When using FF surrogate models to predict mode shapes, the MAC value can be calculated for a comparison with any mode and any reference design; the comparison is not static like it is in the SV surrogate modeling methods. In addition, because the displacement response is predicted for every node on the part, a three-dimensional representation can be displayed that updates in response to new parameter inputs in real time, extending the ability of the user to explore and evaluate the design space. Finally, both SV and FF surrogate models could allow the predicted MAC values to be displayed on a MAC table that can update in response to new parameter inputs in real time.

CHAPTER 5. OPTIMIZATION

Because a design space is often large and complex, optimization is an important tool for conducting design space exploration and finding improved designs. The methods of the previous chapters all could benefit from optimization: finding designs with specific stress criteria on different regions of part (e.g., the user study from Chapter 2), determining designs that have acceptable levels of fatigue risk (e.g., below a certain %G value on the Goodman diagram in Chapter 3), or designs that have mode shapes that match a desired pattern (e.g., the MAC equation in Chapter 4).

This chapter explores the benefits of using FF surrogate models in optimization routines, with a focus on design space exploration and design improvement (RQ1, RQ3). Because the FF surrogate models make nodal predictions, every node's result can be used in the optimization. These results also all inherently pertain to a geometric location; thus, unlike SV surrogate models, the results of a FF surrogate model prediction can be used in an optimization to define objectives and constraints based on location. These spatially-defined objectives and constraints offer unique methods for controlling the optimization and searching for specific results. Although the accuracy is not compared with SV surrogate models, the differences in possible optimization techniques between FF and SV surrogate models is discussed (RQ2).

While these methods could apply to any of the preceding methods and applications in this thesis, the techniques are demonstrated with steady stress only for simplicity. The research objective of determining if FF surrogate models can be used in real engineering applications (RQ1) was answered in the previous two chapters by predicting a specific kind of engineering data; in this chapter, optimization is treated as a worthy engineering application on its own. Thus, the complexity of combining the other results is beyond the scope of the study. This chapter's work does expand on several of the design space exploration activities used in the experiments of Chapter 2.

This study is being cleared by export control at the industry research sponsor and will be submitted to the journal of *Structural and Multidisciplinary Optimization* in August 2019.

5.1 Introduction

Engineers conduct design space exploration during the early phases of design in order to explore potential variations of structural parts. The design space is the collection of all possible design variations. Each design in the design space is uniquely defined by the value of its design parameters. When used with a parameterized model, the space may be explored by varying these parameters, running simulations at critical points, and examining the results (stress, displacements, etc.).

Optimization routines provide an algorithmic and efficient manner by which to explore a design space and obtain an "optimum" design. The design parameters are generally the inputs to the optimization, and the responses to these design changes provide a way for evaluating the design. These optimization routines often rely on callable functions in order to obtain the response of the designs to the design parameter changes. These functions allow the optimization to iterate over a design space by adjusting the design parameters and determining if the responses satisfy pre-defined constraints and objectives [80].

While analytical functions are preferred, they are rarely available for complex behavior; finite element analysis (FEA) is used instead to simulate the behavior. Since these FEA results are obtained by running computationally expensive simulations, including them in an optimization routine can severely slow the process and make thorough design space exploration infeasible [11, 12, 67, 77, 79–81, 88–90]. When optimization routines are used to explore a design space, these same limitations make efficient searches difficult to achieve. Although computing power has increased over time, the growing need for more complex models makes the computational cost of FEA simulations a continuing problem [5]. Optimization routines using these expensive analyses must be carefully designed to avoid wasted resources or long convergence times.

To overcome the high computation time needed for expensive simulations, many aerospace designers employ surrogate models to quickly predict desired results [5, 7, 8, 12, 66]. Surrogate models (also known as response surface models, metamodels, regression models, or emulators [6]) create simplified mathematical relationships between inputs and outputs of a system. Using new

input values, the surrogate models can then predict results in a fraction of the time it takes to perform the original calculations [6, 12, 77]. These surrogate models are often used in design space exploration or optimization routines in order to evaluate many design variations without the prohibitive cost of full simulations [7–12, 77]. They provide cheap approximations of the expensive simulations in the form of a callable function. This makes including certain results from FEA simulations in an optimization routine feasible.

Most surrogate models map a set of inputs to a single output value. These types of surrogate models will be referred to as single-value (SV) surrogate models. They cannot predict the response of an entire system, but instead predict a single value that represents the system (see Fig. 5.1 (top)). Optimization routines often employ these surrogate models in order to obtain cheap estimates of the system's outputs [11, 17, 77, 91]. SV surrogate models are widely used; however, by only predicting single values, the information SV surrogate models provide in an optimization routine is limited [7]. Single values provide little insight into the detailed differences between two complex objects [13]. While values like the maximum stress and weight are important indicators of model quality, a good design is ideally determined by patterns and relationships between the nodal results across the entire part.

In recent years, methods have been developed for using surrogate models to predict the behavior of every node in an FEA simulation [14–16, 92–94]. They generally take geometric parameter values as inputs and predict FEA results for the entire part. Like SV surrogate models, these surrogate models may be treated as a callable function. Previous work has shown that, under certain conditions, these surrogate models can predict the results in real time, allowing interactive design space exploration with a complete prediction and visualization of the results [14, 92–94]. The nodal results provide a much more detailed and complete prediction of the structural FEA response than do SV surrogate modeling methods. It has also been shown that, in some applications, these surrogate models can predict results with greater accuracy than the SV surrogate models [14, 94].

Because this surrogate modeling method predicts an entire field of responses from a FEA simulation, these will be referred to as full-field (FF) surrogate models. This work uses the terminology of SV and FF surrogate modeling methods in order to provide a distinction between these different uses of surrogate models. The difference between SV and FF surrogate models is

illustrated in Fig. 5.1: while SV surrogate models only predict a single value, FF surrogate models predict values for each node on a part's FEA mesh.

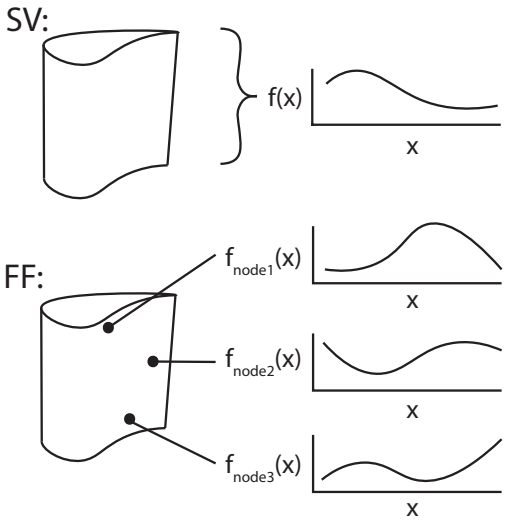

Figure 5.1: (top) Single-value (SV) surrogate models only predict one value that represents the response of the entire part across the design space. (bottom) Full-field (FF) surrogate models predict the response of each node on the part across the design space.

The ability to quickly obtain a complete prediction of the FEA nodal results in the form of a callable function suggests that FF surrogate models would be well suited for use in optimization routines. While previous work has demonstrated how FF surrogate models could be used in various design space exploration tools and activities, this research explores new FF-enabled techniques for improving a design and efficiently searching a design space with optimization routines. Specifically, these techniques are used to control the optimization in spatially-defined ways that are not possible unless the complete nodal response is available. The Purdue model of a jet engine compressor blade is used in the examples, but these methods apply to any three-dimensional structural FEA model.

A brief review of previous work concerning SV and FF surrogate models will be given in Section 5.2. This will be followed by a description of how the surrogate models and data used in this study were created in Section 5.3. Section 5.4 presents case studies that demonstrate how FF surrogate models enable enhanced optimization techniques for design space exploration. Finally, the conclusions will be summarized in Section 5.5.

5.2 Background

Finite element analysis (FEA) is a common engineering tool used for obtaining the structural response to a set of boundary conditions for a given part. FEA programs, such as ANSYS, take a two or three dimensional model of a part and convert it into a finite element mesh made up of many interconnected nodes. The program determines the response to the boundary conditions by solving equations at each node of the mesh. These responses could include values such as the amount of physical displacement or stresses caused by the loads on the part.

SV surrogate models have been used to predict specific values obtained from FEA results, such as the maximum stress of a part, the average nodal displacement of a part, or the response of a specific "monitor" node at some location of interest [7, 67]. These types of surrogate models have also been used to predict many kinds of results in a jet engine, from vibrations in compressor and turbine bladed to aerodynamic coefficients experienced in various nozzle types [8, 11, 66, 67]. Specifically, SV surrogate models have been used to predict single-valued FEA results inside of optimization routines. This has been described as "meta-model based design optimization" [95]. These predicted FEA values have included maximum stresses [7], displacements [7], aerodynamic coefficients [11, 96], physical properties [91], and natural frequencies [12, 83, 86]. These applications demonstrate that single values can be sufficient for some cases, but these SV surrogate models do not describe the full response of a part with as much detail as a full simulation.

FF surrogate model sets have been used to predict the full stress and displacement response of each node in a FEA model [14–16]. Both Heap et al. and Bunnell et al. were able to approach "real-time" design exploration of finite element models by representing the nodes with surrogate models [14, 16]. Bunnell et al. used a unique surrogate for every node's stress, X location, Y location, and Z location in order to predict the stress and geometric response to changes in geometry. These collections of surrogate models represent an entire field of results from a FEA simulation. Both Schulz et al. and Bunnell et al. were able to predict the results for an entire three-dimensional finite element model [14, 15].

The FF surrogate modeling method does not share the same limitations of the SV surrogate models. While SV surrogate models can only predict a single value, FF surrogate models can predict the full response of the part. This allows for a more detailed and complete understanding of the response. Geller et al. [7] noted that when single values are related to a geometric location, such

as maximum stresses or displacements, they can be more difficult to accurately predict with SV surrogate models. FF surrogates can produce better predictions because the values and geometry are coupled, unlike with SV surrogate models [7]. Additional benefits are then possible, such as mapping the predicted nodal responses onto a reconstructed three-dimensional visualization of the FEA model [14, 15].

While surrogate models are convenient and make fast predictions, their results contain a degree of error. The number of training samples, quality of training data, and type of surrogate modeling method affects the amount of error in the predictions. When the training data for surrogate models depends on computationally expensive FEA simulations, an important goal is to achieve acceptable levels of accuracy with a minimum amount of training data [5]. Acceptable accuracy varies depending on the application and field.

It has been shown that the relationship between surrogate accuracy and number of training samples is valid for both SV surrogate models [11, 12, 85] and FF surrogate model sets [14]. Bunnell et al. established that FF surrogate models predicting steady stress values approached 5% normalized root mean square error (NRMSE) as the number of training samples increased [14]. Thelin et al. demonstrated that with 500 training samples, FF surrogate models predict the maximum steady stress with 53% less error than SV surrogate models [94]. For highly nonlinear modal data, however, SV surrogate models have much lower error than FF surrogate models [93].

FF surrogate models are somewhat more difficult to create than SV surrogate models. Each training sample consists of nodal results rather than a single value. For large numbers of training samples and node counts, the amount of training data increases dramatically. With the increase in training data, they also take longer to train than SV surrogate models. When used with optimization, the introduction of nodal data can lead to slower solve times, more function calls, and more complexity.

However, the ability to predict the full response and make reasonably accurate steady stress predictions make FF surrogate models a desirable alternative to SV surrogate models in optimization routines. While SV surrogate models are simpler to create, the information they provide to an optimization routine is limited. The nodal predictions of FF surrogate models allow objective and constraint functions to have more detailed control over the design. This research explores these FF-enabled optimization techniques.

5.3 Method

This section will discuss the basic principles for creating FF surrogate models of nodal steady stress values. Fig. 5.2 illustrates the steps in this process. This study uses much of the same basic workflow as described by Bunnell et al., and the interested reader is referred to [14] for more complete details.

To train surrogate models with FEA data, it is necessary to first create a set of training data [5, 10, 79, 83, 84]. The training data set for FF surrogate models may be constructed using a design of experiments (DOE). The DOE finds a set of designs that adequately fill the design space. The suggested number of designs in the DOE can depend on criteria such as computational resources, model complexity, and number of dimensions. For a model with n input parameters, each design instance is defined by an n-dimensional vector [26]. These parameters values often control the geometry or conditions of a parameterized FEA model. The parameter values for each design are considered the inputs for the surrogate model.

These parameter values are used to update the geometry of a parametric model, such as a parameterized three-dimensional model. Rather than create a new finite element mesh for every new geometry, a mesh-morphing process is applied to the baseline design's mesh. This allows all the designs in the training data set to have a common set of node numbers and relative locations [14]. With these common meshes, there exists a sample of the behavior for each specific node across all of the training data.

The mesh for each design is used in a structural FEA simulation, and the steady stress results are stored [5]. FEA produces results such as stress or modal displacements for every node on a model. The FEA results are considered the outputs for the surrogate model.

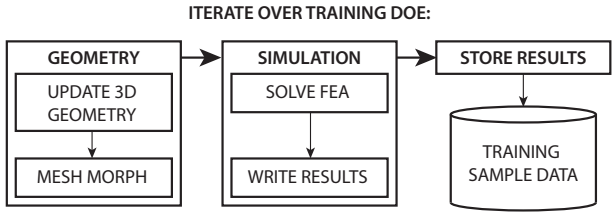

Figure 5.2: The basic workflow for creating FF training data.

The surrogate models are trained on these inputs and outputs, and develop a mathematical function to describe the relationship between them. Because the training depends on samples of the actual data taken from the design space, this is known as a sample-based method [10, 79]. Once trained, the surrogate models can then receive parameter values for new designs and predict the full response of the part. Instead of executing a new simulation each time a design change is made, a properly trained surrogate model can provide the desired output value very quickly without expensive calculations. This process is shown in Fig. 5.3.

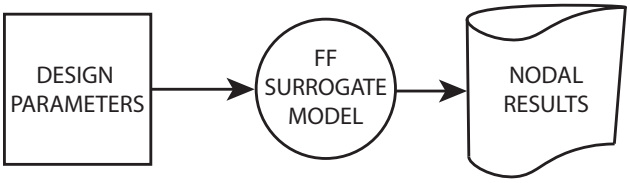

Figure 5.3: FF surrogate models accept new parameter values and predict a design's nodal results

Although this study focuses on the FF predictions of steady stress values for FF surrogate models, any nodal results could be used. Thelin et al. used FF surrogate models to predict alternating stresses [94] and modal displacement values [93].

5.3.1 SV Surrogate Models

The techniques presented in this research use FF surrogate models, but can be combined with SV surrogate models for additional constraints or objectives. There are some kinds of data that simply make more sense to predict as single values, such as a part's weight value. The examples in this paper highlight unique FF surrogate model applications, but do use a SV surrogate model to predict weight for an objective function in Section 5.4.2. Thus, it is worth noting here the general process for creating a SV surrogate model.

To create a SV surrogate model that predicts a weight value, the weight is calculated for each sample in the training data set. The surrogate model is trained with these weight values instead of the complete nodal responses. New input parameters yield a prediction of the weight only; the SV predictions provide no further insight to the part's behavior.

Table 5.1: Bounds for blade design parameters

Parameter	Lower Bound	Upper Bound
Angle	-20 deg	20 deg
$Chord_{root}$	20.32 mm (0.8 in)	30.48 mm (1.2 in)
$Chord_{tip}$	20.32 mm (0.8 in)	30.48 mm (1.2 in)
Lean	-7.62 mm (-0.3 in)	7.62 mm (0.3 in)
Sweep	-7.62 mm (-0.3 in)	7.62 mm (0.3 in)

5.3.2 Model

This study uses the transonic Purdue blade model of a jet engine compressor [18], which is a "blade-alone" model. The Purdue blade was developed for research purposes and is a general representation of a three-stage compressor blade. Further details about the Purdue research blade can be found in [18]. This particular model has been meshed with 25000 nodes. The model consists of various airfoil cross-sections with connecting surfaces. It has been parameterized with the five parameters and bounds in Tab. 5.1. These parameters change the geometry of the model by adjusting the profiles at the root and tip of the blade. The *Angle* parameter adjusts the difference in the angle between the profile at the root and the profile at the tip. The two *Chord* parameters adjust the chord length at the root and at the tip. The *Lean* parameter offsets the tip profile in a direction perpendicular to the root profile's chord, while the *Sweep* parameter offsets the tip profile in the same direction as the root profile's chord. This parameterization and associated controls are discussed in [14] with more detail. Only the surface nodes are used in this study in order to reduce computational costs.

5.3.3 Training Data

As described in Section 5.3, a training set of design parameters (inputs) and FEA results (outputs) must be generated to train the surrogate models. An optimized Latin Hypercube was chosen because it is a space-filling design [5, 12, 77, 83]. Based on preliminary studies of the surrogate model accuracy, 500 samples were generated. Each of these samples was used to update the parametric compressor blade model. The updated models were used to conduct structural analyses in ANSYS, and nodal stress data was extracted for each sample.

5.3.4 SQP

This study uses a sequential quadratic programming (SQP) optimization method called sequential least squares quadratic programming [97]. SQP navigates in a bounded design space and subjects the design to constraints [98]. The boundaries used in this study are the same as those used in the DOE [77] (see Tab. 5.1). Using an optimization method involving constraints is particularly helpful in the design of a jet engine compressor blade because the part is often subject to various stress requirements that keep it under the maximum allowable stress. The SQP method can also handle functions of several variables such as the input parameters used by the surrogate models for each new designs [98]. These features make it an appropriate choice for the type of optimization that this study will use.

The SQP method is one of the most efficient optimization algorithms and often gives more accurate results than other algorithms [97]. It is a gradient-based method, which converges quickly and is well suited for large design spaces [99]. However, it does not enforce feasibility of the constraints at each step. Thus the SQP optimization is only guaranteed to be feasible at the end of the search. This does provide a danger for searches that terminate early, as they are not guaranteed to be feasible designs or even better than the starting point [97]. Any designs that terminated early were rejected in this study as a consequence.

5.4 FF-enabled Optimization Techniques

The following cases explore and demonstrate optimization techniques enabled by FF surrogate models. Because FF surrogate models give a complete three-dimensional prediction of the response, it is possible to use spatially-defined constraints and objectives in an optimization routine. The FF surrogate models can be used to create constraints that apply to regions of nodes, or used in objective functions that use the location of a node or the full response pattern as objects to minimize or maximize. A summary of these techniques is presented in Tab. 5.2.

The techniques will make use of the following generic optimization problem:

Min $f(\mathbf{X})$ or **Max** $f(\mathbf{X})$

s.t.

$c(\mathbf{X})$

where $f(\mathbf{X})$ is the objective function, $c(\mathbf{X})$ is a constraint function, and $\mathbf{X}$ is set of current design parameters. The optimization techniques will either minimize or maximize $f(\mathbf{X})$, depending on the objective. The specific constraint functions used will vary from application to application.

Table 5.2: FF-enabled optimization techniques

Method	Description
Regional constraints	Constrain a subset of nodes to be under a limit
Location-based objectives	Finds designs where a nodal response is nearer to/farther away from a target node
Response pattern matching	Finds designs with a similar response pattern

5.4.1 Nodal Constraints

FF surrogate models allow optimization routines to use predicted information for every node. Instead of only constraining a predicted global value, such as maximum stress, constraints can be applied individually.

For an entire part, constraints can either be uniformly applied or vary across the nodes. Constraints that are uniform, c_u, across the entire part can be implemented in one of two ways. The first predicts the results for each node, and then finds the maximum value in the predicted set. This value is subjected to the constraint, λ, as indicated in Eqn. 5.1:

$$c_{u,max}(\mathbf{X}) : max([\phi_1, \phi_2, \cdots, \phi_n]) \leq \lambda \tag{5.1}$$

where ϕ_i is the response value of an individual node and λ is the global constraint. Although the constraint is applied to the maximum value only, the maximum value is obtained by predicting the

results of all the nodes. This produces a more accurate prediction than using a SV surrogate model prediction of the maximum stress only, as noted by Thelin et al. [93, 94].

The second method predicts the results for each node in the same manner, but then subjects each individual node's response to the constraint, λ, as indicated in Eqn. 5.2:

$$c_{u,i}(\mathbf{X}) : \phi_i \leq \lambda \tag{5.2}$$

where ϕ_i is the response value of an individual node and λ is the global constraint. While these methods are mathematically similar, they may cause different behavior in optimization routines. The first (max) method creates one constraint, while the second (nodal) method creates n constraints for n nodes. With very large numbers of constraints, the optimizer requires more constraint function evaluations in order to determine if the current iteration violates or satisfies the constraints.

Applying constraints to individual nodes may cause the optimizer to behave differently, but it also enables more interesting constraint capabilities. Because constraints can be applied to individual nodes, they also can be varied from node to node, c_v. This allows the optimization to apply a different limit, λ_i, to each node as indicated in Eqn. 5.3:

$$c_{v,i}(\mathbf{X}) : \phi_i \leq \lambda_i \tag{5.3}$$

where ϕ_i is the response value and λ_i is the constraint of an individual node, i. These limits can vary from node to node, or can vary between subsets of nodes. The next section presents an example of applying nodal constraints to different subsets of nodes.

5.4.2 Regional Constraints

FF surrogate models allow constraints to be applied to each node individually. Rather than constrain every node in a mesh, it may be desirable to constrain only certain nodes. Thus, constraints can be applied to the nodes in different subsets, or regions, of a part. This allows the designer to have more control over which designs are found within the design space by the optimization routine. Beck et al. examined the effects and risks of localized responses on as-manufactured blades, and used surrogate models to predict this behavior [2]. Regional constraint

methods such as those presented here could be used to constrain localized, or regional, responses in specific ways.

Implementation

Eqn. 5.2 shows the general case of using varied constraints for individual nodes. For the regional constraint method, $c_{region,i}$, a very similar concept is used, but is applied only to a subgroup of nodes. These constraints are made by assigning a limit to every node in a particular region as shown in Eqn. 5.4:

$$c_{region,i}(\mathbf{X}) : \phi_{region,i} \leq \lambda_i \tag{5.4}$$

where $\phi_{region,i}$ is the response at node i of a region and λ_i is the constraint on node i. For n nodes in the region, there will be n nodal constraints. The ability to constrain nodes individually allows for much more nuanced control over the blade's stress pattern. Additional regional or global constraints could be applied for further control in the optimization.

Application Examples

To demonstrate this method, two regions were defined on the Purdue blade. Fig. 5.4 shows these regions with the relevant nodes highlighted. The stress of each node in the regions used was constrained to be equal to or less than a given limit.

Minimize $f(\mathbf{X})$: weight

s.t.

$\forall$ nodes, i, in region

$c_{region,i}(\mathbf{X}) : \sigma_{region,i} \leq \lambda_i$

The following examples show the results of using FF steady stress surrogate model predictions to constrain the stress response for the nodes in a particular region. The FF steady stress surrogate model set was used to predict the stresses at each node on the blade. The responses for nodes in the specified region were constrained, and all other nodal responses were left unconstrained.

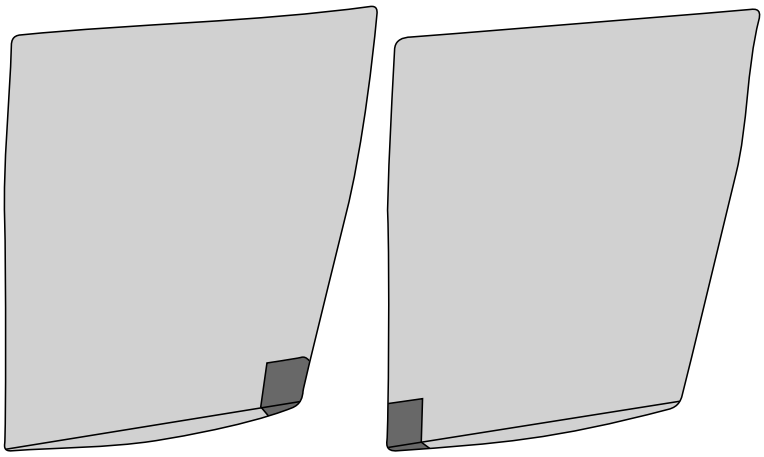

(a) Trailing edge root corner (b) Leading edge root corner

Figure 5.4: Nodal regions on the blade

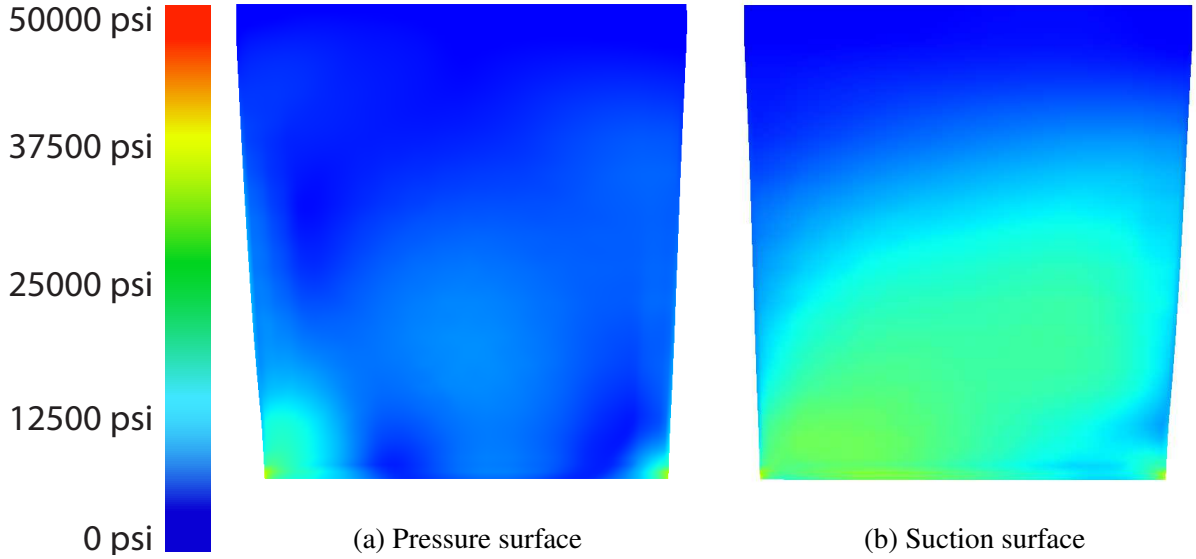

(a) Pressure surface (b) Suction surface

Figure 5.5: Baseline design.

The focus of this example is on the FF regional constraints; however, the optimization cannot search for a design without a specified objective function. In order to simplify the problem and clarify the effect of the regional stress constraint, the objective function for this example seeks to minimize the weight of the part. A SV surrogate model was chosen for weight rather than a FF surrogate model since weight is not a nodal value.

Each optimization started from the baseline design shown in Fig. 5.5, which is considered the center of the design space. At the baseline design, the global maximum stress is 36590 psi. The maximum stress in the trailing edge region is 36590 psi and the maximum stress in the leading edge region is 35367 psi. The weight at the baseline design is 0.108 lbs.

Fig. 5.6 shows the result of minimizing the weight while constraining the nodes in the corner of the root and the trailing edge to be under 25000 psi. The optimum design has a global maximum stress of 45640 psi, and a regional maximum stress of 18257 psi, which satisfies the constraint of 25000 psi. The weight was reduced to 0.0746 lbs. In comparison with the baseline design shown in Fig. 5.5, the global maximum stress is now located in the center of the suction side of the blade.

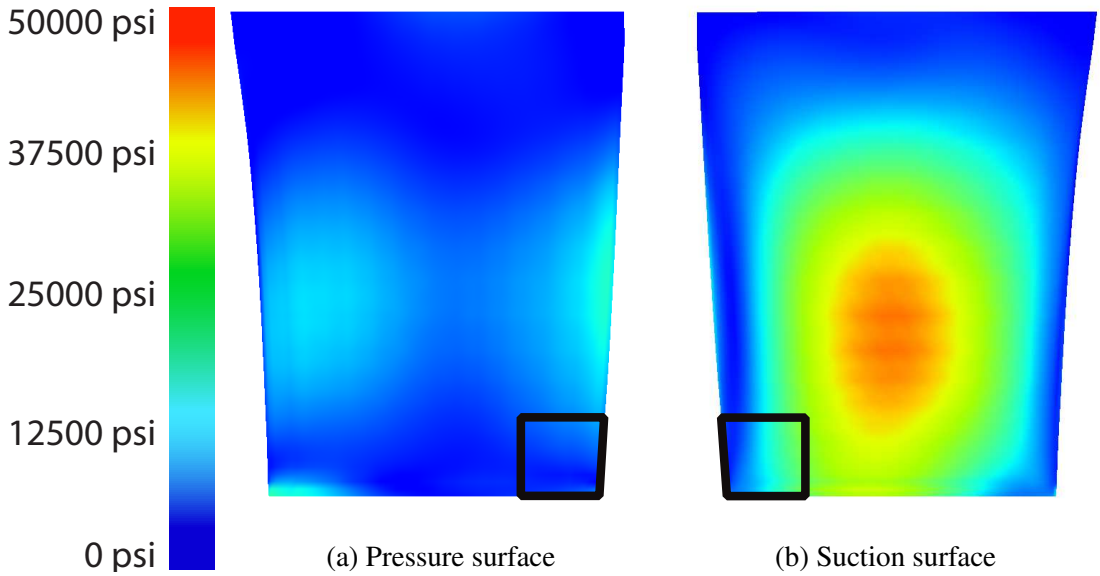

(a) Pressure surface (b) Suction surface

Figure 5.6: The optimization found a design that minimizes the weight and constrains the regional maximum stress in the corner of the root and the trailing edge. Compare to Fig. 5.5.

Similarly, Fig. 5.7 shows the optimum design when the nodes at the corner of the root and the leading edge were constrained to be under 25000 psi. The optimum design has a global maximum stress of 35612 psi, and a regional maximum stress of 24970 psi, which satisfies the constraint of 25000 psi. The global maximum stress is not located in the constrained leading edge corner, but instead in the corner of the root and the trailing edge. The weight was reduced to 0.0697 lbs.

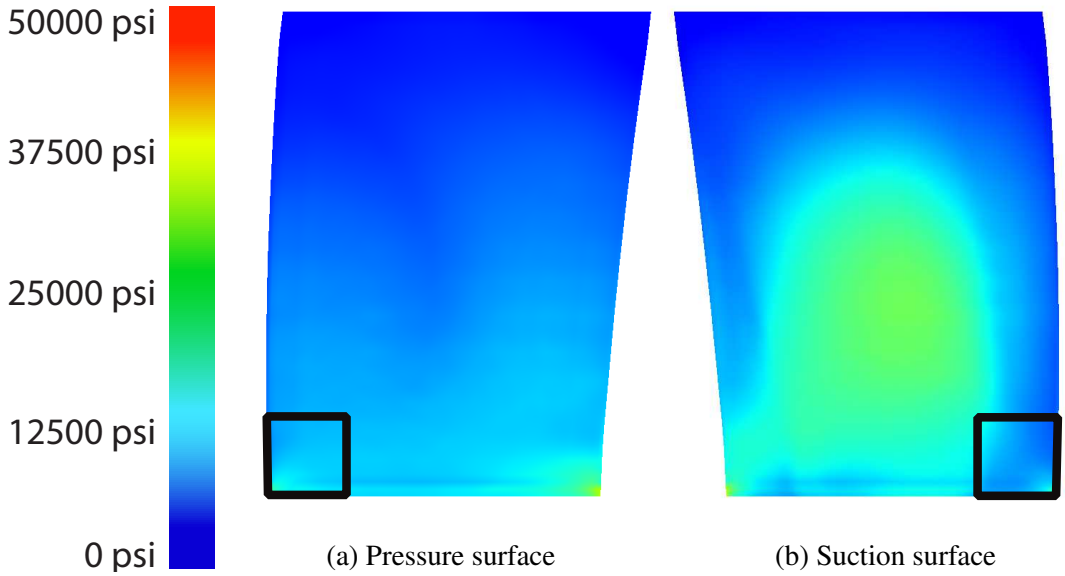

(a) Pressure surface (b) Suction surface

Figure 5.7: The optimization found a design that minimizes the weight and constrains the regional maximum stress in the corner of the root and the leading edge. Compare to Fig. 5.5.

Different Starting Points

The particular design reached by the optimization can depend on the starting point used. If the constraints are too demanding, the optimization's search path may be limited to local minima and not have the freedom to find the true global optimum. Thus, it is possible that the optimum design found by the optimization may not be the true global optimum, and the optimization results may vary. It is important to try starting the optimization from many different starting points in order to confidently identify the global optimum.

Although this principle applies to all of the techniques in this study, it is only illustrated in detail with this example for brevity. In this example, the SV weight prediction was minimized and the nodes of the corner of the root and the trailing edge were constrained under 25000 psi. The optimization was then tested with 100 different starting points generated with an optimized Latin Hypercube.

Of the 100 starting points, 29 resulted in optimization paths that failed to converge within the maximum allowable number of iterations. Some imposed constraints will create infeasible designs. The optimization cannot create new results; it may only be used as a design exploration tool to find existing results. If there are no feasible designs, the optimization will fail. However,

71 starting points did result in successful optimizations; the failures were a result of poor starting points rather than the absence of feasible designs in the design space. The failed optimization runs were searching for a feasible design while trying to minimize the weight prediction. The maximum stress in the trailing edge root corner for each failed design was above the constraint of 25000 psi. Some starting points did not provide as clear of an optimal gradient and caused the search path to remain in infeasible regions of the design space longer than the successful designs. Thus, while feasible designs exist, some starting points can cause optimization routines to fail.

The average predicted regional maximum stress in the corner nodes for these successful designs was 22500 psi. Fig. 5.8 shows the FF predicted and ANSYS regional maximum stress values for the found optimum designs. The predicted values ranged from 17178 psi to 24999 psi, all of which fulfill the constraint of 25000 psi. Though the objective and constraint functions did not vary, the found optimum designs reached different results due to the different starting points. While the majority of starting points resulted in values that were closer to the constraint of 25000 psi, some achieved much lower maximum stresses in the region. The objective searched for the minimum weight; the stress in the region was only a constraint. However, the different starting points greatly affected the stress in this corner. Only one starting point allowed the optimizer to find the lowest maximum stress in the trailing edge root corner. This illustrates the value of using different starting points in order to establish confidence that the optimization has found the desired result.

When checked with ANSYS simulations, the average actual regional maximum stress was found to be 21880 psi. Thus, the average percent error of the predicted regional maximum stresses was 4.11%. Fig. 5.8 shows the range of values obtained in ANSYS for these optimum designs. The histograms clarify the relationship between the FF predicted and ANSYS values. In these results, the ANSYS values for some points are actually higher than the FF predictions, and violate the 25000 psi constraint. However, the ANSYS values for other points are actually lower than the prediction, as in the case of the lowest point found. Thus, depending on the design, the FF predicted results could be better or worse than the actual results. The FF surrogate models used in this study had some error, as do all surrogate models. The errors in the models depend on the particular method in which the surrogate models were created. Greater accuracy can be obtained by training the surrogate models with more training points or by other methods [93].

Generally, it is important to verify surrogate model predictions by checking the results with the full simulation (such as in ANSYS). To allow for this error, a designer may want to use a lower constraint than the desired constraint. For example, in order to avoid finding designs that have feasible predictions but are actually infeasible, a designer in this study could reduce the constraint to 24000. This helps ensure that, even if the ANSYS values are higher than the predictions, they do not violate the desired constraint.

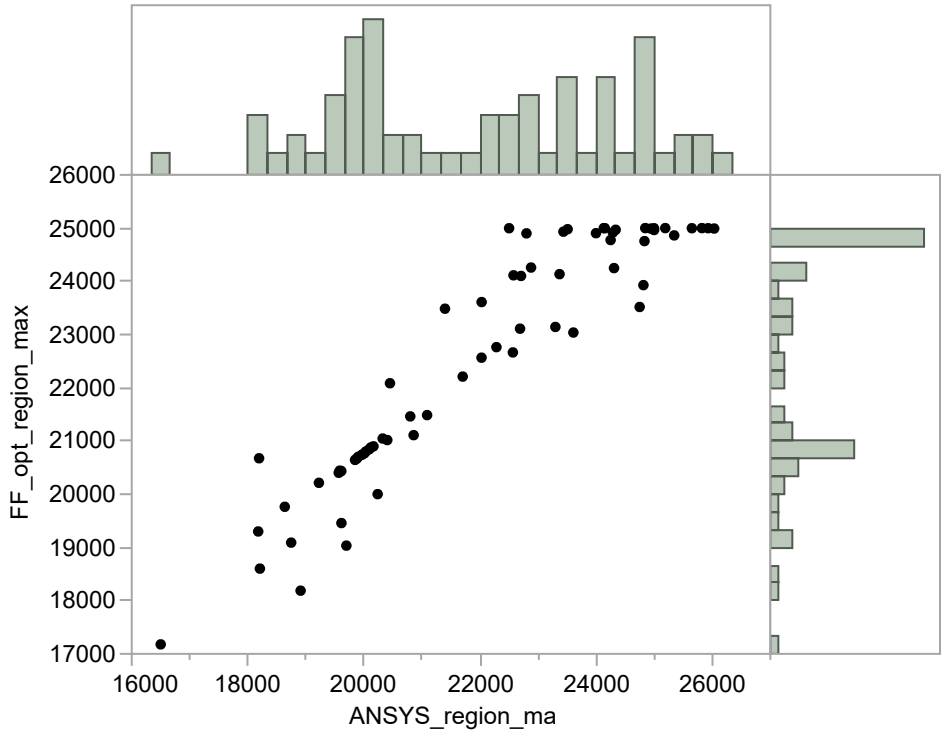

Figure 5.8: The maximum steady stress values in the trailing edge corner region for 71 successful optimum designs. Using different starting points in the optimization can produce different optimal and results.

5.4.3 Location-based Objectives

The location of certain responses on a part can be a critical factor in accepting or rejecting a design. These specific responses could include the global maximum stress, a maximum displacement, or some other unique response that is found at a particular node. For example, when the maximum stress of a compressor blade is located at the root, it can signal a structural risk. This is

especially true if the location of the maximum alternating stress is in the same location, leading to compounding stresses and a higher risk of eventual fatigue failure at that point [69]. In the case of mode shapes (the pattern of displacements due to vibrations), high displacements in certain locations could also present risks. Thus, it may be desirable to search for designs in the design space based on the location of the maximum responses.

The FF surrogate models predict the response of each node on a part. Because each node inherently has an individual relative location on the part, the FF surrogate models also can be used to predict the relative location of each response. Bunnell et al. [14] predicted the steady stresses, but also the Cartesian coordinates of each node in space to visualize the geometry of each new design. Using these features, FF surrogate models may be used in an optimization to control the location of specific responses.

Implementation

A target node, t, can be used to control the location of the response. This is the node where the designer would prefer that the response would occur. The objective function predicts the full nodal response of the part. The node where the global maximum response is actually located, m, is identified for each new iteration of the optimization. The global maximum response can be a maximum stress, displacement, or other result value as long as it is located at a single node. Then, the predicted Cartesian coordinates of the target node (e.g., t_x, t_y, and t_z) and the node with the global maximum response (e.g., m_x, m_y, and m_z) are obtained. The Euclidean distance δ between the nodes is calculated with Eqn. 5.5:

$$\delta = \sqrt{\sum_{k=1}^{3} (t_k - m_k)^2} \tag{5.5}$$

where d is the number of dimensions. This distance is returned to the optimizer as the result of the objective function. By minimizing this distance, designs are found in which the node with the global maximum response approaches the location of the desired target node.

The previous section described how using nodal constraints over a region can help to control the location of the global maximum stress. Here, the nodal location of the global maximum stress will be controlled with the objective function and a target node.

Application Examples

To demonstrate the method, a target node was chosen. Fig. 5.9 (top) shows the baseline design, where the global maximum stress is indicated with a black point and the target node is indicated with a grey point. The optimization used the method described above with the FF steady stress surrogate model set to minimize the distance between the target node and the node with the global maximum stress. The optimization found a design that minimized the distance between the points as shown in Fig. 5.9 (bottom). The location of the node with the global maximum stress is now nearly at the location of the target node.

Min $f(\mathbf{X}) : \sqrt{\sum_{k=1}^{3}(t_k - m_k)^2}$ (Eqn. 5.5)

The node with the global maximum response can also be driven away from the target node. In this application, the exact location of the global maximum response is less important than ensuring it does not occur at the target node. The objective function uses exactly the same process, except the optimizer maximizes the distance δ between the nodes. Fig. 5.10 (top) shows the baseline design again, where the global maximum stress is indicated with a black point and the target node is indicated with a grey point. In this example, the two nodes begin close to one another. Fig. 5.10 (bottom) shows the optimum design, where the distance between these two points has been maximized.

Max $f(\mathbf{X}) : \sqrt{\sum_{k=1}^{3}(t_k - m_k)^2}$ (Eqn. 5.5)

Maximizing the distance between two response locations is especially useful when considering both steady and alternating stresses. Consider the example in Figure 5.11(a). This shows the steady and alternating stresses for a starting design, as well as the percent Goodman (%G) values mapped onto the geometry and the Goodman diagram. The nodes at which the maximum steady stress and the maximum alternating stress occur are both very close together, and the %G values

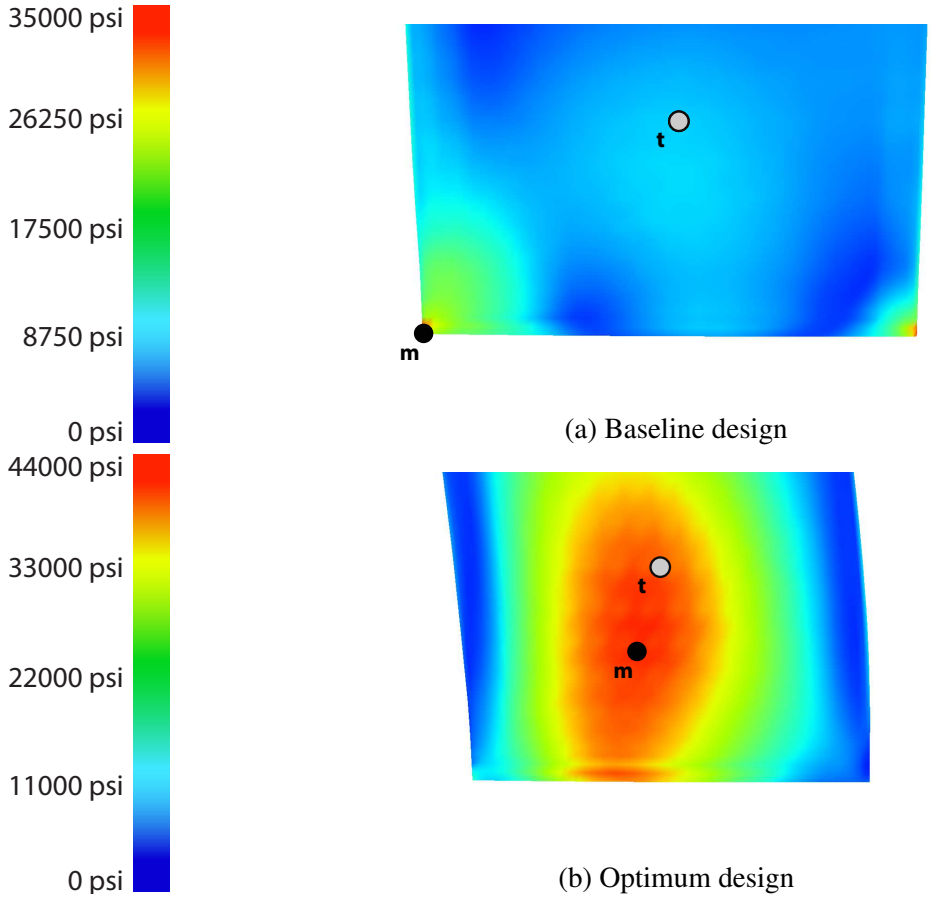

(a) Baseline design

(b) Optimum design

Figure 5.9: The optimization finds a design that minimizes the distance between the maximum stress (m) and a target node (t). Views depict the bottom half of the suction side of the blades. Colorscales are not common in order to clarify the max stress locations.

are relatively high in this area ($\%G_{max,nodal}$ = 56%). More information about the %G values is given in Chapter 3. Figure 5.11(b) shows the optimum design after maximizing the distance between these two responses. Now, the maximum steady stress is no longer in the same area as the maximum alternating stress. This has helped the %G values in the area to be reduced by about half ($\%G_{max,nodal}$ = 27%). Thus, this optimization strategy can be used as an effective method for reducing the risk of fatigue failure.

Considerations

Again, there may be designs in the design space that are slightly more optimal than the design found by the optimization. For the example in Fig. 5.9, the design space did contain designs

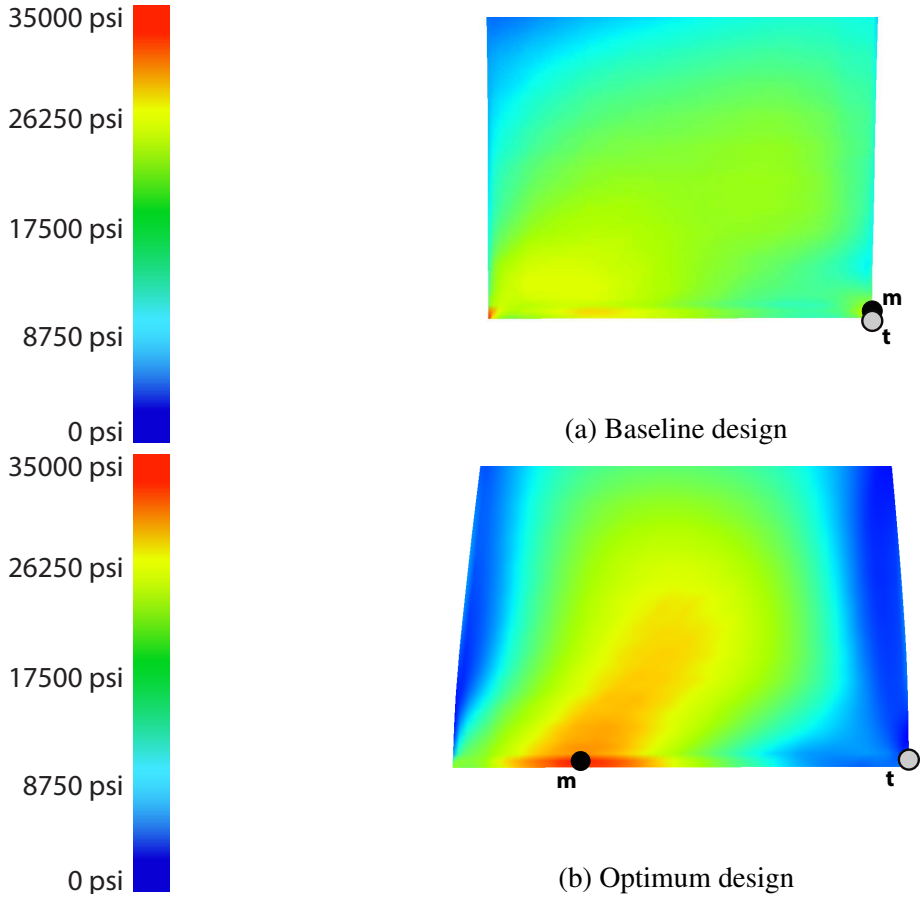

(a) Baseline design

(b) Optimum design

Figure 5.10: The optimization also can find a design that maximizes the distance between the maximum stress (m) and a target node (t). Colorscales are not common in order to clarify the max stress locations.

in which the node with the global maximum stress was exactly at the target node. Using different starting points could allow the optimization to find these closer matches. Because the FF surrogate models can be used to predict results quickly, a designer may want to interactively explore the design space around the optimum design in order to determine that closer matches do not exist. This type of interactive exploration is described in [14] and [92–94].

This method requires consideration of the general part geometry. Because the Euclidean distance between two nodes does not provide any information about what face the nodes are on, the optimization may find that the minimum distance between the nodes is through the part. This is especially problematic if the part is very narrow with respect to one or two dimensions, such as the thin compressor blades used in these examples. For example, when minimizing the distance as

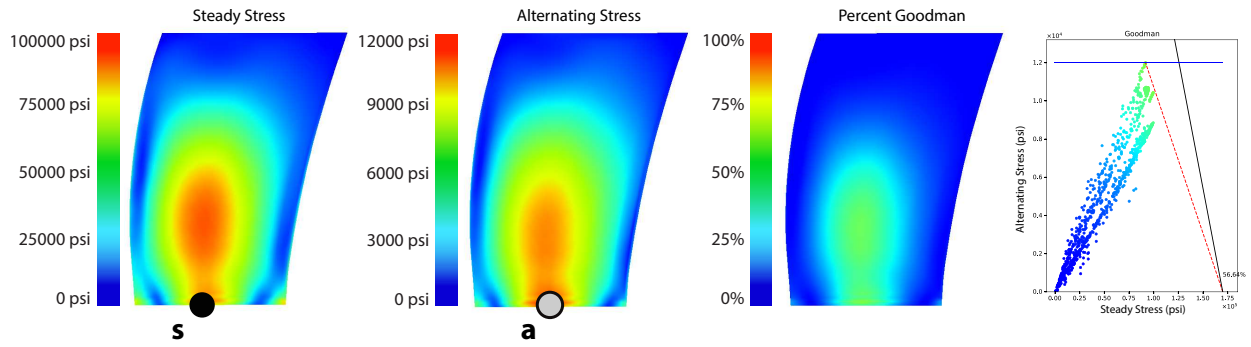

(a) Starting design

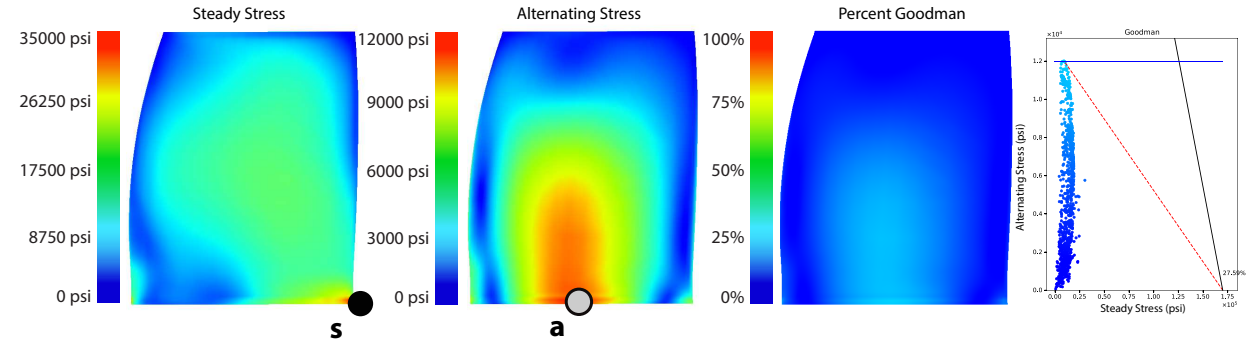

(b) Optimum design

Figure 5.11: Maximizing the distance between the maximum steady stress (s) and the maximum alternating stress (a) can reduce the fatigue risk on the part. Steady and alternating stresses are shown next to the percent Goodman values mapped onto the part and the Goodman diagram. Colorscales are not common in order to clarify the max stress locations and different types of data.

depicted in Fig. 5.9, using the unaltered distance formula in Eqn. 5.5 actually yields the optimum design in Fig. 5.12.

The problem can be accounted for by scaling the coordinate values to have similar magnitudes. Eqn. 5.6 shows the general method for scaling the coordinates in the distance formula, where α are the scaling factors for each coordinate. In these three-dimensional examples, since the compressor blade model is very narrow with respect to the Y direction compared to the X and Z coordinates, the coordinates were all scaled to have similar magnitudes. This allowed the optimizer to treat each dimension more equally, and obtain the results presented in Fig. 5.9.

$$\delta = \sqrt{\sum_{k=1}^{3} \alpha_k (t_k - m_k)^2} \tag{5.6}$$

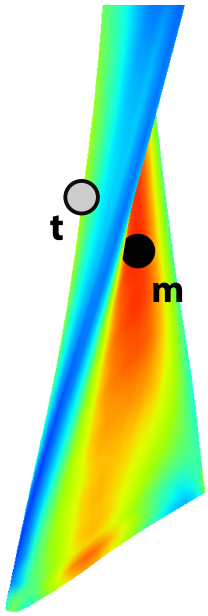

Figure 5.12: Without scaling, the minimum distance found could be through the part.

5.4.4 Response Pattern Matching

The previous two methods have shown how nodal constraints or nodal objectives can be used in optimization to have control over the response at specific regions or nodal locations. However, for certain types of design exploration, it may be important to have control over the entire nodal response pattern.

For example, a specific stress pattern may have been previously analyzed and studied, but the current design's global maximum stress may exceed the desired allowable stress. A designer's goal would be to find a design that has as similar a stress pattern as possible, but with a reduced maximum stress.

Designers are also sometimes given a specific stress pattern and asked to find the geometric parameters that would create such a response. In this case, the goal of design space exploration is to find a design that produces a matching stress pattern.

Manually searching the design space to find a design with these responses is time-consuming and prone to human error. An optimization routine that can use the entire response as an objective provides a more efficient method for finding the desired design as well as its parameter values.

Implementation

A design is chosen before optimizing, and the response pattern saved. This set of nodal responses becomes a target pattern that the optimization must try to match. During each iteration of the optimization, the FF surrogate models are used to predict the full nodal response of each new design. The nodal responses of the target pattern and the current stress pattern are scaled from 0.0 to 1.0. Then, the nodal responses are assembled into two vectors, each of the form shown in Eqn. 5.7.

$$\phi = \begin{bmatrix} \sigma_1 & \sigma_1 & \dots & \sigma_{n-1} & \sigma_n \end{bmatrix} \tag{5.7}$$

The modal assurance criterion (MAC) is an equation which finds the degree of similarity between two vectors [22, 23]. Usually, it is used to compare mode shape vectors for mode shape identification. The MAC equation may also be used to find the similarity between the two stress pattern vectors. This application is shown in Eqn. 5.8, where $\{\phi_T\}_i$ is the i^{th} node of the target pattern vector, and $\{\phi_D\}_i$ is the i^{th} node of the current design's pattern vector. The MAC equation returns values between 0.0 to 1.0, where values close to 0.0 indicate poor correlation and values close to 1.0 indicate high correlation [23].

$$\text{MAC}_{T:D} = \frac{|\sum_{i=1}^{n} \{\phi_T\}_i \{\phi_D\}_i|^2}{(\sum_{i=1}^{n} \{\phi_T\}_i \{\phi_T\}_i)(\sum_{i=1}^{n} \{\phi_D\}_i \{\phi_D\}_i)} \tag{5.8}$$

The starting point of the optimization should not be the design which generated the target pattern; if the starting point and the target design are the same, the MAC value will already be maximized at 1.0 (a perfect correlation). The objective function predicts the stress contours and computes the MAC value. The optimizer then maximizes the MAC value. Constraints can be added to this objective to help the optimization to improve the design.

Max $f(\mathbf{X})$: $\text{MAC}_{T:D}$ value (Eqn. 5.8)

s.t.

$c_{u,max}(\mathbf{X})$: $\max([\sigma_1, \sigma_2, \cdots, \sigma_n]) \leq \lambda$ (Eqn. 5.1)

Application Examples

For this example, the chosen response pattern is the steady stresses at each node shown in Fig. 5.13 (top). The global maximum steady stress of this pattern is 52130 psi. The goal of the optimization is to find a design with a similar pattern but lower stresses. Thus the optimization seeks designs that have high MAC values but constrains the global maximum stress to be below 40000 psi. The optimum design found by the optimization is shown in Fig. 5.13 (bottom). The geometry of the optimum design is very similar, but the chord at the tip is slightly smaller than the starting design. The stress pattern is nearly identical with MAC$=0.996$. The greatest difference between the target design's and optimum design's patterns is that the global maximum stress at the optimum design is 39999.9 psi. This optimization improved the design while maintaining the desired stress pattern.

Because there is not always a design in the design space that fits the criteria, the optimum design is not always a perfect match. Fig. 5.14 shows a different pattern matching example where the pattern of the optimum design does not match the target pattern as closely. The MAC value is 0.940, indicating that the match is still very good, but not perfect. This is most apparent in the high stress areas on the front of the blade. The target design's pattern has two distinct regions of high stress, while the optimum design's pattern has only one. Many other details of the two patterns do match, such as the stresses around the corners and the general distribution of stresses. The geometry difference between the starting and optimum designs are also larger than the previous example. The optimum design has 0.05 in of positive *Sweep*, 12 deg positive *Angle*, and 0.84 in $Chord_{root}$ while the target pattern's design had 0.29 in negative *Sweep*, -10 deg negative *Angle*, and 1.14 in $Chord_{root}$. However, the pattern found by the optimization does satisfy the criteria and is considered a success. If maintaining similar geometry is also an important consideration to the designer, then further constraints or penalties could be added to control for geometric variation.

5.5 Conclusions

This research demonstrated new techniques for optimization of three-dimensional FEA results using full-field (FF) surrogate models. FF surrogate models predict the complete nodal response of FEA results very quickly without the need for computationally expensive simulations

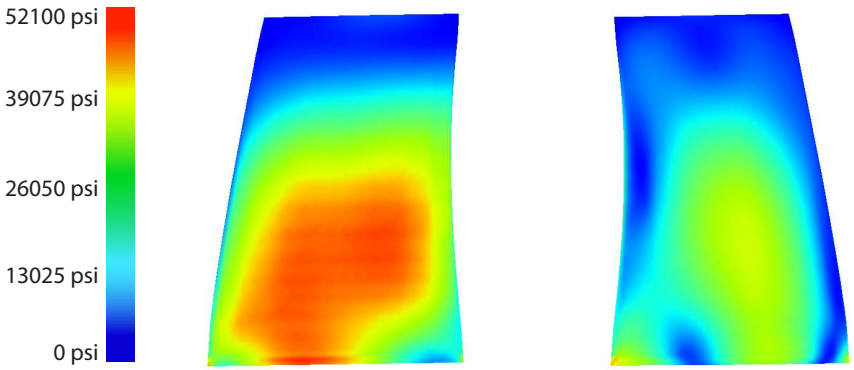

(a) Target pattern

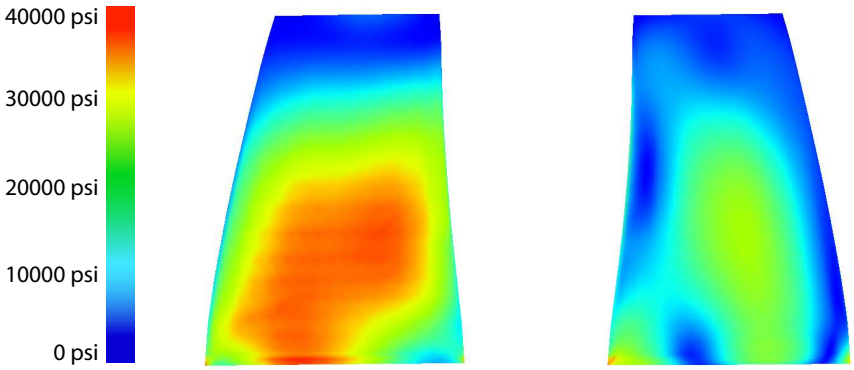

(b) Optimum pattern

Figure 5.13: Using the MAC equation, the optimization found a similar stress pattern with stresses that met the constraint. Colorscales are not common in order to highlight the relative stress patterns.

at run time. This allows optimization routines to use much more detailed information about a structural part in constraint and objective functions.

These methods extend the ability of optimization routines to have control over FEA results at a nodal level. Rather than rely on cheap but limited predictions of single values, the optimization can control regions, nodal locations, and complete patterns of responses as it traverses a design space. This enables designers to be much more specific when searching a design space with optimization, but also to more easily improve designs in spatially-defined ways.

Because FF surrogate models can predict FEA results for each node, they can be used to apply individual constraints to the nodes. By defining regions of interest on a part, FF surrogate models make it possible to constrain specific regions differently than others. It was demonstrated

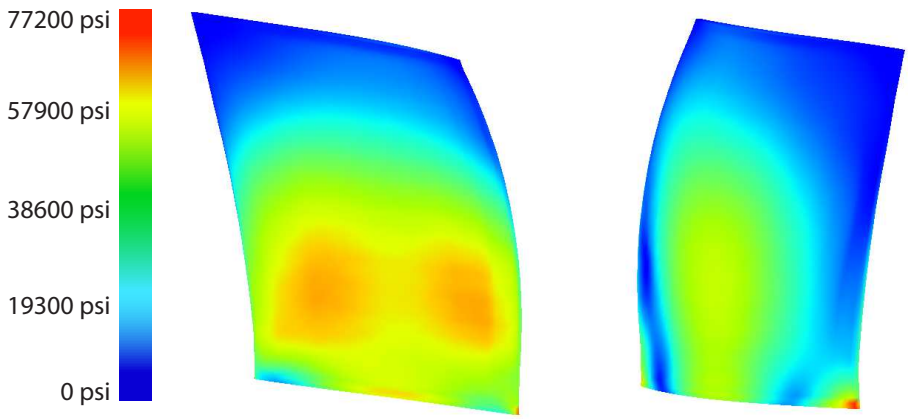

(a) Target pattern

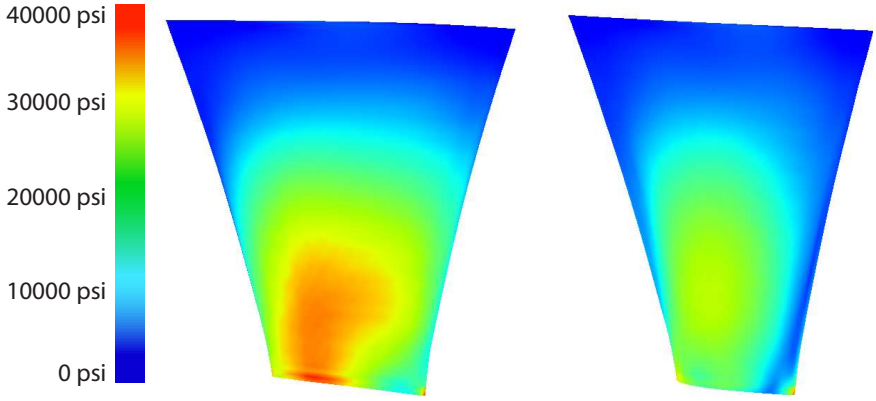

(b) Optimum pattern

Figure 5.14: The design space may not always have a design that perfectly fits the criteria. Colorscales are not common in order to highlight the relative stress patterns.

that the optimization successfully found designs in the design space that minimized the objective while respecting the regional constraints.

Objective functions that minimized a distance formula could find designs that either drew a particular response on the part's surface closer to a target location. The inverse case of driving a response away from a particular location was also shown. These methods were successful, but do require some tuning for the best results.

Finally, a method was described for searching a design space for designs with similar response patterns. This employed the modal assurance criterion as a means for determining similarity

between the three-dimensional responses. When combined with other constraints, designs may be improved while maintaining the same basic response patterns.

The examples in this research focused on design space exploration activities for the Purdue model of a jet engine compressor blade, but the methods could be used in other applications. The compressor blade model could be replaced with any three-dimensional structural FEA model. Further exploration would need to be made in order to make conclusions about the efficacy of the methods for the results and specific surrogate models used. Also, although this paper focused on early design exploration activities, using FF surrogate models for the optimization of three-dimensional geometry and nodal results could apply to other activities such as updating models to match experimental data.

CHAPTER 6. CONCLUSION

This chapter presents general conclusions pertaining to the preceding studies and results, as well as a brief discussion about future work related to this research. While the publications found in Chapters 2 - 5 contain conclusion sections that pertain to those individual studies, the conclusions in this chapter combines the results in the context of the overall research goals. This research set out to answer three main questions:

RQ1 Can full-field (FF) surrogate models be used reliably in real engineering and design space exploration applications?

RQ2 How do FF surrogate models compare with single value (SV) surrogate models in these activities, especially in terms of accuracy and ability?

RQ3 What additional design benefits are gained by using FF surrogate models in the design process?

This research concludes that FF surrogate models can be used in real engineering analyses for certain types of data where SV surrogate models are typically used. Accurate predictions of the steady and alternating stresses allow FF surrogate models to be used in creating dynamic Goodman diagrams for predicting fatigue life. The steady and alternating stresses, as well as Goodman-related measures that are calculated with those stresses, can be predicted with better accuracy than SV surrogate models. However, highly nonlinear data, such as mode shapes, are more prone to error when predicted with FF surrogate models rather than SV surrogate models. For these types of data, FF surrogate models are only reliable with an extremely large amount of training samples. This makes SV surrogate models a better choice. FF surrogate models also perform well in optimization routines for design space exploration.

FF surrogate models are found to present unique design benefits when used for design space exploration activities. The ability to predict the response at every node in real time enables an

interactive design space exploration experience, where predicted three-dimensional results appear immediately as direct responses to changes in the design parameters. This leads to greater intuition and understanding about the behavior of a part through the design space. The benefits gained are specific to the applications; generally, presenting all the nodal data for a part allows a designer to quickly evaluate design variations as a whole instead of through abstracted single values that represent the results.

The three-dimensional result predictions also provide opportunities for new visualizations that enhance the design process. The difference model simplifies the complex differences between two designs, and is shown to make a significant difference in the speed and accuracy of certain design comparisons. Every node's prediction can be plotted onto a Goodman diagram instead of a single conservative point. Values from the Goodman diagram can be plotted directly onto the part's three-dimensional visualization and indicate which locations on the part are at most risk of failing in fatigue. The results of spatially-defined optimization criteria can be displayed to the user.

6.1 Future Work

While FF surrogate models have been demonstrated and tested in real engineering applications, their potential is still just being uncovered. This thesis presents research has helped discover domains where FF surrogate models are helpful, as well as others where they are not. Comparisons with current SV surrogate methods illustrate the advantages or disadvantages gained by using FF surrogate models. However, the methods and examples, while using real analysis methods and example geometry, have been focused on general principles instead of a specific implementation. Further studies ought to apply the principles and methods of this research to the design of a specific part and provide concrete examples of their benefits. As the findings of this research are applied and shown to be truly beneficial, confidence in these methods will increase to the point that they may be incorporated into current design practices.

This research examined the engineering activities of design space exploration, fatigue analysis with the Goodman diagram, mode shape comparisons with the modal assurance criterion, and optimization. There are many other types of engineering analyses and activities in which FF surrogate models could be used. Fields such as heat transfer, fluid dynamics, and material studies could potentially benefit from FF surrogate modeling. The principles presented in this thesis should ap-

ply to the use of FF surrogate models in any of these applications, but like the individual studies in this research, each new application of FF surrogate models will likely reveal specific details regarding implementation, prediction quality, and added benefits.

The examples presented in this thesis applied FF surrogate models to jet engine design and a specific jet engine compressor blade model. The methods are generic to any geometry, and thus further research should demonstrate the usefulness of these methods in other contexts. Due to the cost of creating training data and setting up the framework, these other geometries and design applications ought to have the same challenges of computational cost and result complexity as does jet engine design. This will provide motivation for the applications and provide useful knowledge to those fields. Potential applications could include simulations of crash testing in the automotive industry, spaceflight design, and manufacturing equipment design.

Highly nonlinear data, such as mode shapes, were found to limit the utility of FF surrogate models; unless a prohibitively large amount of training data is used, the FF surrogate models suffer from high inaccuracy. Future studies should examine strategies for overcoming these obstacles. Preprocessing of the training data could help.

These applications were examined independently of each other. Future work could synthesize the methods in this study. An optimization that includes the percent Goodman values would need to ensure that similar mode shapes are being used from one design to the next; therefore, the modal assurance criterion would have to be used in order to properly determine if the correct mode shape is predicted. This level of synthesis is beyond the goals of this research, and would require further work to overcome the problems with modal data found in Chapter 4. However, this would provide a useful contribution to promoting industry confidence in this work.

REFERENCES

[1] Amoo, L. M., 2013. "On the design and structural analysis of jet engine fan blade structures." *Progress in Aerospace Sciences,* **60**, pp. 1–11. 1, 10, 40, 59, 74

[2] Beck, J., Brown, J., Kaszynski, A., Carper, E., and Gillaugh, D., 2019. "Geometric mistuning reduced order model development utilizing bayesian surrogate models for component mode calculations." In *Proceeds of ASME Turbo Expo 2019*. 1, 41, 84, 85, 118

[3] Pratt and Whitney GTF engine Accessed: 2019-07-01. 1

[4] Hughes, T. J., 2012. *The finite element method: linear static and dynamic finite element analysis*. Courier Corporation. 2, 10, 13

[5] Becker, W., Oakley, J., Surace, C., Gili, P., Rowson, J., and Worden, K., 2012. "Bayesian sensitivity analysis of a nonlinear finite element model." *Mechanical Systems and Signal Processing,* **32**, pp. 18–31. 2, 77, 82, 84, 85, 90, 99, 108, 112, 113, 115

[6] Han, Z.-H., Zhang, Y., Song, C.-X., and Zhang, K.-S., 2017. "Weighted gradient-enhanced kriging for high-dimensional surrogate modeling and design optimization." *AIAA Journal*(12), Dec, pp. 4330–4346. 2, 14, 41, 77, 108, 109

[7] Geller, M., Schemmann, C., and Kluck, N., 2017. "Optimization of the operation characteristic of a highly stressed centrifugal compressor impeller using automated optimization and metamodeling methods." ASME Paper GT2017-63262. 2, 41, 42, 44, 77, 78, 83, 108, 109, 111, 112

[8] Zhang, M., Gou, W., Li, L., Yang, F., and Yue, Z., 2017. "Multidisciplinary design and multi-objective optimization on guide fins of twin-web disk using kriging surrogate model." *Structural and Multidisciplinary Optimization,* **55**(1), pp. 361–373. 2, 41, 44, 77, 78, 83, 108, 109, 111

[9] Farias Filho, U. P., Antunes, A. R., Bastos, S. M., and Lyra, P. R., 2015. "Minimization of vortex induced vibrations using surrogate based optimization." *Structural and Multidisciplinary Optimization,* **52**(4), pp. 717–735. 2, 41, 78, 109

[10] Denimal, E., Nechak, L., Sinou, J.-J., and Nacivet, S., 2016. "Kriging surrogate models for predicting the complex eigenvalues of mechanical systems subjected to friction-induced vibration." *Shock and Vibration,* **2016**. 2, 6, 56, 58, 77, 78, 80, 81, 82, 89, 92, 103, 109, 113, 114

[11] Goulos, I., Otter, J., Stankowski, T., MacManus, D., Grech, N., and Sheaf, C., 2016. "Aerodynamic design of separate-jet exhausts for future civil aero-enginespart ii: Design space exploration, surrogate modeling, and optimization." *Journal of Engineering for Gas Turbines and Power,* **138**(8), p. 081202. 2, 44, 56, 78, 83, 84, 89, 108, 109, 111, 112

[12] Nobari, A., Ouyang, H., and Bannister, P., 2015. "Uncertainty quantification of squeal instability via surrogate modelling." *Mechanical Systems and Signal Processing,* **60-61**, pp. 887 – 908. 2, 10, 14, 56, 57, 58, 77, 78, 81, 83, 84, 89, 90, 91, 99, 108, 109, 111, 112, 115

[13] Alexander, E., and Gleicher, M., 2015. "Task-driven comparison of topic models." *IEEE transactions on visualization and computer graphics,* **22**(1), pp. 320–329. 2, 109

[14] Bunnell, S., Thelin, C., Gorrell, S., Salmon, J., Ruoti, C., and Hepworth, A., 2018. "Rapid visualization of compressor blade finite element models using surrogate modeling." In *ASME Turbo Expo 2018: Turbomachinery Technical Conference and Exposition*, no. GT2018-77188, American Society of Mechanical Engineers, pp. V07AT30A011–V07AT30A011. 3, 9, 10, 11, 14, 21, 40, 42, 44, 45, 53, 54, 56, 58, 68, 77, 78, 84, 85, 90, 91, 105, 109, 111, 112, 113, 115, 125, 128

[15] Schulz, A., Xu, J., Zhu, B., Zheng, C., Grinspun, E., and Matusik, W., 2017. "Interactive design space exploration and optimization for cad models." *ACM Transactions on Graphics (TOG),* **36**(4), July, p. 157. 3, 10, 14, 42, 44, 78, 84, 109, 111, 112

[16] Heap, R. C., Hepworth, A. I., and Jensen, C. G., 2015. "Real-time visualization of finite element models using surrogate modeling methods." *Journal of Computing and Information Science in Engineering,* **15**(1), Mar, p. 011007. 3, 10, 14, 42, 44, 78, 84, 91, 109, 111

[17] Viana, F. A., Simpson, T. W., Balabanov, V., and Toropov, V., 2014. "Special section on multidisciplinary design optimization: metamodeling in multidisciplinary design optimization: how far have we really come?." *AIAA Journal,* **52**(4), pp. 670–690. 3, 109

[18] Cross, C. J., 1998. "Turbomachine airfoil vibration control utilizing active and passive piezoelectric elements.". 5, 42, 56, 79, 89, 115

[19] Nicholas, T., and Maxwell, D., 2003. "Mean stress effects on the high cycle fatigue limit stress in ti-6al-4v." In *Fatigue & Fracture Mechanics: 33rd Volume*. ASTM International, pp. 476–492. 5, 40, 46

[20] Nicholas, T., and Zuiker, J., 1996. "On the use of the goodman diagram for high cycle fatigue design." *International Journal of Fracture,* **80**(2-3), pp. 219–235. 5, 40, 45, 46, 47, 48, 59

[21] George, T. J., Shen, M. H., Nicholas, T., Cross, C. J., et al., 2006. "A new multiaxial fatigue testing method for variable-amplitude loading and stress ratio." *Journal of Engineering for Gas Turbines and Power,* **128**(4), p. 857. 5, 40, 46

[22] Selin, E. D., 2012. "Application of parametric nurbs geometry to mode shape identification and the modal assurance criterion.". 6, 77, 80, 131

[23] Pastor, M., Binda, M., and Harabik, T., 2012. "Modal assurance criterion." *Procedia Engineering,* **48**, pp. 543 – 548 Modelling of Mechanical and Mechatronics Systems. 6, 77, 80, 81, 131

[24] Britt, B., Chure, D., Engelmann, G., Dalla Vecchia, F., Scheetz, R., Meek, S., Thelin, C., and Chambers, M., 2015. "A new, large, non-pterodactyloid pterosaur from a late triassic interdunal desert environment within the eolian nugget sandstone of northeastern utah." In *75th Annual Meeting Society of Vertebrate Paleontology, Dallas*, p. 97. 9

[25] Platt, C., Pratt, T., and Brown, K., 1982. Structural tailoring of engine blades (staebl) NASA CR-167949. 10

[26] Coffey, D., Lin, C., Erdman, A. G., and Keefe, D. F., 2013. "Design by dragging: An interface for creative forward and inverse design with simulation ensembles." *IEEE Transactions on Visualization and Computer Graphics,* **19**(12), Dec, pp. 2783–2791. 10, 11, 14, 40, 44, 69, 77, 82, 113

[27] Rygiel, P., Obrocki, W., and Sieniawski, J., 2017. "Numerical vibration analysis of turbine engine compressor blades depending on geometry and position of the damage." *Advances in Manufacturing Science and Technology,* **41**(1). 10

[28] Kou, H., Lin, J.-s., Zhang, J.-h., and Fu, X., 2017. "Dynamic and fatigue compressor blade characteristics during fluid-structure interaction: Part I blade modelling and vibration analysis." *Engineering Failure Analysis,* **76**, Jun, pp. 80–98. 10

[29] Nader, G., Wang, K., Hétroy-Wheeler, F., and Dupont, F., 2016. "Just noticeable distortion profile for flat-shaded 3d mesh surfaces." *IEEE transactions on visualization and computer graphics,* **22**(11), pp. 2423–2436. 11

[30] Kim, K., Carlis, J. V., and Keefe, D. F., 2017. "Comparison techniques utilized in spatial 3d and 4d data visualizations: A survey and future directions." *Computers & Graphics,* **67**, pp. 138–147. 11, 12, 15, 16, 19, 20

[31] Gleicher, M., Albers, D., Walker, R., Jusufi, I., Hansen, C. D., and Roberts, J. C., 2011. "Visual comparison for information visualization." *Information Visualization,* **10**(4), pp. 289–309. 11, 12, 15, 16, 19, 20

[32] Gleicher, M., 2018. "Considerations for visualizing comparison." *IEEE transactions on visualization and computer graphics,* **24**(1), pp. 413–423. 11, 15, 16, 18

[33] Johnson-Roberson, M., Pizarro, O., Williams, S. B., and Mahon, I., 2010. "Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys." *Journal of Field Robotics,* **27**(1), pp. 21–51. 11

[34] Blanc, T. J., Gorrell, S. E., Jones, M. R., and Duque, E. P., 2013. "Analysis and compression of time-accurate turbomachinery simulations using proper orthogonal decomposition." In *49th AIAA/ASME/SAE/SEE Joint Propulsion Conference 2013*, p. 3620. 11, 17

[35] Shen, Q., Pang, A., and Uselton, S., 1998. "Data level comparison of wind tunnel and computational fluid dynamics data." In *Proc. VIS*, IEEE Computer Society Press, pp. 415–418. 11, 17

[36] Behmann, J., Mahlein, A.-K., Paulus, S., Dupuis, J., Kuhlmann, H., Oerke, E.-C., and Plümer, L., 2016. "Generation and application of hyperspectral 3d plant models: methods and challenges." *Machine Vision and Applications,* **27**(5), July, pp. 611–624. 11, 16

[37] Stalling, D., Westerhoff, M., Hege, H.-C., et al., 2005. "Amira: A highly interactive system for visual data analysis." *The visualization handbook,* **38**, pp. 749–67. 11

[38] Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L., Leach, M. O., and Hawkes, D. J., 1999. "Nonrigid registration using free-form deformations: application to breast mr images." *IEEE transactions on medical imaging,* **18**(8), Aug, pp. 712–721. 11, 16

[39] Cordero, C., Zebelo, S. A., Gnavi, G., Griglione, A., Bicchi, C., Maffei, M. E., and Rubiolo, P., 2012. "HS-SPME-GC× GC-qMS volatile metabolite profiling of *Chrysolina herbacea* frass and *Mentha* spp. leaves." *Analytical and bioanalytical chemistry,* **402**(5), Feb, pp. 1941–1952. 11, 16

[40] Park, S., Seo, D., Kim, Y., and Lee, K., 2010. "Organic vapor detection using a color-difference image technique for distributed Bragg reflector structured porous silicon." *Sensors and Actuators B: Chemical,* **147**(2), Jun, pp. 775–779. 11, 16

[41] Chopin, J., Laga, H., Huang, C. Y., Heuer, S., and Miklavcic, S. J., 2015. "Rootanalyzer: a cross-section image analysis tool for automated characterization of root cells and tissues." *PLoS ONE,* **10**(9), Sept, p. e0137655. 11, 16

[42] Lu, N., Wang, J., Wu, Q., and Yang, L., 2008. "An improved motion detection method for real-time surveillance." *IAENG International Journal of Computer Science,* **35**(1), pp. 119–128. 11, 16

[43] Aggarwal, J., and Nandhakumar, N., 1988. "On the computation of motion from sequences of images - a review." *Proc. IEEE,* **76**(8), Aug, pp. 917–935. 11, 16

[44] Javed, W., and Elmqvist, N., 2012. "Exploring the design space of composite visualization." In *Pacific Visualization Symposium (PacificVis), 2012 IEEE*, IEEE, pp. 1–8. 15, 19

[45] Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., and Guibas, L., 2012. "Functional maps: a flexible representation of maps between shapes." *ACM Transactions on Graphics (TOG),* **31**(4), p. 30. 16, 18

[46] Cignoni, P., Rocchini, C., and Scopigno, R., 1998. "Metro: Measuring error on simplified surfaces." *Computer Graphics Forum,* **17**(2), Nov, pp. 167–174. 16

[47] Alexander, E., and Gleicher, M., 2016. "Task-driven comparison of topic models." *IEEE transactions on visualization and computer graphics,* **22**(1), pp. 320–329. 18

[48] Vurputoor, R., Mukherjee, N., Cabello, J., and Hancock, M., 2008. "A mesh morphing technique for geometrically dissimilar tessellated surfaces." In *Proceedings of the 16th International Meshing Roundtable*, Springer, pp. 315–334. 18

[49] Shigley, J. E., 2011. *Shigley's mechanical engineering design*. Tata McGraw-Hill Education. 19, 45, 48, 59

[50] Bergman, L. D., Rogowitz, B. E., and Treinish, L. A., 1995. "A rule-based tool for assisting colormap selection." In *Proc. VIS*, IEEE Computer Society, pp. 118–125. 22

[51] Tominski, C., Fuchs, G., and Schumann, H., 2008. "Task-driven color coding." In *Information Visualization*, IEEE, pp. 373–380. 22

[52] Rheingans, P., 2000. "Task-based color scale design." In *Proc. SPIE*, Vol. 3905, SPIE, pp. 35–43. 22

[53] Light, A., and Bartlein, P. J., 2004. "The end of the rainbow? color schemes for improved data graphics." *Eos, Transactions American Geophysical Union,* **85**(40), Oct, pp. 385–391. 24

[54] Suarez, E. L., Duffy, M. J., Gamache, R. N., Morris, R., and Hess, A. J., 2004. "Jet engine life prediction systems integrated with prognostics health management." In *Aerospace Conference, 2004. Proceedings. 2004 IEEE*, Vol. 6, IEEE, pp. 3596–3602. 40, 46

[55] Dong, P., 2005. "A robust structural stress method for fatigue analysis of offshore/marine structures." *Journal of offshore mechanics and Arctic engineering,* **127**(1), pp. 68–74. 40

[56] OHalloran, B. M., Stone, R. B., and Tumer, I. Y., 2013. "Developing new design requirements to reduce failures in early complex systems design." In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, American Society of Mechanical Engineers, pp. V005T06A037–V005T06A037. 40, 74

[57] Lough, K. G., Stone, R., and Tumer, I. Y., 2009. "The risk in early design method." *Journal of Engineering Design,* **20**(2), pp. 155–173. 40, 74

[58] Patil, S., Patil, S., and Karuppanan, S., 2013. "Modal and fatigue analysis of a camshaft using fea." *International Journal of Applied Engineering Research,* **8**(14), pp. 1685–1694. 40, 45

[59] Nicholas, T., 1999. "Critical issues in high cycle fatigue." *International Journal of Fatigue,* **21**, pp. S221–S231. 40, 45, 46, 47, 48

[60] Ghafoori, E., Motavalli, M., Nussbaumer, A., Herwig, A., Prinz, G., and Fontana, M., 2015. "Design criterion for fatigue strengthening of riveted beams in a 120-year-old railway metallic bridge using pre-stressed cfrp plates." *Composites Part B: Engineering,* **68**, pp. 1–13. 40, 45, 46

[61] Kim, J.-S., Yoon, H.-J., Lee, S.-H., Lee, W.-G., and Shin, K.-B., 2011. "Durability evaluation of the composite bogie frame under different shapes and loading conditions." *ICCM18, Korea, August June*, pp. 22–26. 41, 46, 49, 51

[62] Kim, J.-S., 2006. "Fatigue assessment of tilting bogie frame for korean tilting train: Analysis and static tests." *Engineering Failure Analysis,* **13**(8), pp. 1326–1337. 41, 46, 49, 51

[63] Morlacchi, S., Pennati, G., Petrini, L., Dubini, G., and Migliavacca, F., 2014. "Influence of plaque calcifications on coronary stent fracture: a numerical fatigue life analysis including cardiac wall movement." *Journal of biomechanics,* **47**(4), pp. 899–907. 41, 46, 47, 48, 49, 51

[64] Day, W. D., Fiebiger, S. W., and Patel, H. N., 2012. "Parametric evaluation of compressor blade blending." In *ASME Turbo Expo 2012: Turbine Technical Conference and Exposition*, American Society of Mechanical Engineers, pp. 1099–1110. 41, 46, 49, 51

[65] Perry, M., Oktay, S., and Muskivitch, J., 2002. "Finite element analysis and fatigue of stents." *Minimally Invasive Therapy & Allied Technologies,* **11**(4), pp. 165–171. 41, 46, 49, 51

[66] Houck, L., Sewell, D., Burke, M., and Vogel, G., 2015. "A fully coupled aero, thermal, and structural lifetime model for root cause failure analysis and robust redesign of an industrial f class gas turbine blade." ASME Paper GT2015-42505. 41, 44, 77, 83, 108, 111

[67] Hüls, M., Panning-von Scheidt, L., and Wallaschek, J., 2019. "Influence of geometric design parameters onto vibratory response and high-cycle fatigue safety for turbine blades with friction damper." *Journal of Engineering for Gas Turbines and Power,* **141**(4), p. 041022. 42, 44, 46, 58, 77, 83, 91, 108, 111

[68] Shen, M.-H. H., 1999. "Reliability assessment of high cycle fatigue design of gas turbine blades using the probabilistic goodman diagram." *International journal of fatigue,* **21**(7), pp. 699–708. 45, 46, 47

[69] Ni, K., 2013. "Blade stress estimation during multiple vibratory modes." In *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, p. 1772. 46, 50, 63, 125

[70] Foresta, M., Gnanasambandam, S., Weston, D., Li, F., Pan, J., and Le Blanc, M., 2018. "Corrosion fatigue of phosphor bronze reinforcing tapes on underground power transmission cables-failure analysis." *Engineering Failure Analysis,* **91**, pp. 507–515. 46, 47

[71] George, T. J., Seidt, J., Shen, M.-H. H., Nicholas, T., and Cross, C. J., 2004. "Development of a novel vibration-based fatigue testing methodology." *International Journal of Fatigue,* **26**(5), pp. 477–486. 46

[72] Maxwell, D. C., and Nicholas, T., 1999. "A rapid method for generation of a haigh diagram for high cycle fatigue." In *Fatigue and Fracture Mechanics: 29th Volume*. ASTM International. 46

[73] Muthu, S. E., Dileep, S., Kumar, S. S., and Girish, D., 2013. "Life prediction of directionally solidified air cooled hpt gas turbine blade used in a supersonic aircraft using fem." In *ASME 2013 Gas Turbine India Conference*, American Society of Mechanical Engineers, pp. V001T05A005–V001T05A005. 46

[74] Kwofie, S., 2001. "An exponential stress function for predicting fatigue strength and life due to mean stresses." *International Journal of fatigue,* **23**(9), pp. 829–836. 47

[75] Jelaska, D., 2011. "On the goodman's fatigue safety factor." *International Journal of Advanced Engineering,* **5**(1), p. 27. 47

[76] Clark, D. L., and Bae, H.-r., 2017. "Non-deterministic kriging framework for responses with mixed uncertainty." In *19th AIAA Non-Deterministic Approaches Conference*, p. 0593. 56, 58, 89, 91, 92

[77] Wang, J., Wang, C., and Zhao, J., 2017. "Frequency response function-based model updating using kriging model." *Mechanical Systems and Signal Processing,* **87**, pp. 218–228. 57, 58, 77, 90, 91, 99, 108, 109, 115, 116

[78] Allemang, R., 2003. "The modal assurance criteriontwenty years of use and abuse." *Sound and vibration,* **37**, pp. 14–23. 77, 81, 86

[79] Chikhaoui, K., Bouhaddi, N., Kacem, N., Guedri, M., and Soula, M., 2016. "Low cost metamodel for robust design of periodic nonlinear coupled micro-systems." *MATEC Web Conf.,* **83**, p. 05004. 77, 82, 108, 113, 114

[80] Hamza, K., and Saitou, K., 2004. "Crash mode analysis of vehicle structures based on equivalent mechanism approximations." In *Proceedings of the Fifth International Symposium on Tools and Methods of Competitive Engineering*, pp. 13–17. 77, 108

[81] Bae, H., Boyd, I. M., Carper, E., and Brown, J. M., 2019. "Non-deterministic emulator for mistuned bladed rotor responses with multi-fidelity modeling approach." In *AIAA Scitech 2019 Forum*, p. 2003. 77, 81, 83, 84, 99, 101, 102, 103, 108

[82] Yadav, V., Venkataraman, S., and Bland, S., 2016. "Mode-tracking in surrogate-based inverse identification of rotor blade geometry using campbell diagram." In *54th AIAA Aerospace Sciences Meeting*, p. 1642. 81, 101

[83] Qin, S., Zhang, Y., Zhou, Y.-L., and Kang, J., 2018. "Dynamic model updating for bridge structures using the kriging model and pso algorithm ensemble with higher vibration modes." *Sensors,* **18**(6). 81, 82, 83, 90, 111, 113, 115

[84] Mehmani, A., Chowdhury, S., and Messac, A., 2015. "Predictive quantification of surrogate model fidelity based on modal variations with sample density." *Structural and Multidisciplinary Optimization,* **52**(2), pp. 353–373. 82, 113

[85] Lamecki, A., Kozakowski, P., and Mrozowski, M., 2004. "Cad-model construction based on adaptive radial basis functions interpolation technique." In *15th International Conference on Microwaves, Radar and Wireless Communications (IEEE Cat. No.04EX824)*, Vol. 3, pp. 799–802 Vol.3. 84, 103, 112

[86] Lew, T., Spencer, A., Scarpa, F., Worden, K., Rutherford, A., and Hemez, F., 2006. "Identification of response surface models using genetic programming." *Mechanical Systems and Signal Processing,* **20**(8), pp. 1819–1831. 84, 111

[87] Gres, S., Döhler, M., Andersen, P., and Mevel, L., 2018. "Variance computation of the Modal Assurance Criterion." In *ISMA 2018 - 28th Conference on Noise and Vibration Engineering*, pp. 1–12. 86

[88] Liu, H., Xu, S., and Wang, X., 2016. "Sampling strategies and metamodeling techniques for engineering design: comparison and application." In *ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition*, American Society of Mechanical Engineers, pp. V02CT45A019–V02CT45A019. 108

[89] Chaudhuri, A., Lam, R., and Willcox, K., 2017. "Multifidelity uncertainty propagation via adaptive surrogates in coupled multidisciplinary systems." *AIAA Journal*, pp. 1–15. 108

[90] Wang, X., and Shi, L., 2014. "A new metamodel method using gaussian process based bias function for vehicle crashworthiness design." *International Journal of Crashworthiness,* **19**(3), pp. 311–321. 108

[91] Liao, X., Li, Q., Yang, X., Zhang, W., and Li, W., 2008. "Multiobjective optimization for crash safety design of vehicles using stepwise regression model." *Structural and multidisciplinary optimization,* **35**(6), pp. 561–569. 109, 111

[92] Thelin, C., Salmon, J., Bunnell, S., and Gorrell, S., In press. Difference modeling for design space exploration and comparison of 3D structural simulation results. 109, 128

[93] Thelin, C., Salmon, J., Bunnell, S., Gorrell, S., Bird, G., Routi, C., Selin, E., and Calogero, J., Manuscript submitted for publication. Evaluation of surrogate-modeling methods to predict the modal assurance criterion. 109, 112, 114, 118, 123, 128

[94] Thelin, C., Salmon, J., Bunnell, S., Gorrell, S., Bird, G., Routi, C., Selin, E., and Calogero, J., Manuscript submitted for publication. Real-time emulation of goodman diagram for FEA design space exploration. 109, 112, 114, 118, 128

[95] Hamman, R. A., 2015. "The analysis and optimization of an axial compressor." PhD thesis, Stellenbosch: Stellenbosch University. 111

[96] Immonen, E., 2017. "2d shape optimization under proximity constraints by cfd and response surface methodology." *Applied Mathematical Modelling,* **41**, pp. 508–529. 111

[97] Parkinson, A. R., Balling, R., and Hedengren, J. D., 2013. "Optimization methods for engineering design." *Brigham Young University,* **5**, p. 11. 116

[98] Jones, E., Oliphant, T., Peterson, P., et al., 2001–. SciPy: Open source scientific tools for Python [Online; accessed ¡today¿]. 116

[99] Mueller, L., and Verstraete, T., 2017. "Cad integrated multipoint adjoint-based optimization of a turbocharger radial turbine." *International Journal of Turbomachinery, Propulsion and Power,* **2**(3), p. 14. 116