Multi-Objective Optimization

for the

Conceptual Design of Structures

By

Nathan C. Brown

B.S.E. in Civil and Environmental Engineering
Princeton University, 2012

Submitted to the Department of Architecture
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Building Technology

at the

Massachusetts Institute of Technology

June 2016

Signature of Author: ____
Department of Architecture
May 6, 2016

Certified by: ____
Caitlin T. Mueller
Assistant Professor of Architecture & Civil and Environmental Engineering

Accepted by: ____
Takehiko Nagakura
Professor of Architecture
Chair, Departmental Committee for Graduate Studies

Thesis Committee

Caitlin T. Mueller
Assistant Professor of Architecture and Civil and Environmental Engineering
Massachusetts Institute of Technology
Thesis Supervisor

John A. Ochsendorf
Class of 1942 Professor of Architecture and Civil and Environmental Engineering
Massachusetts Institute of Technology
Thesis Reader

Multi-Objective Optimization for the Conceptual Design of Structures

By

Nathan C. Brown

Submitted to the Department of Architecture on May 6, 2016
in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Building Technology

Abstract

Using computational tools, fast and accurate predictions of building performance are increasingly possible. In parallel, the expectations of a high-performance building have been rising in contemporary architecture, as designers must synthesize many inputs to arrive at a design that fulfills a wide range of requirements. Despite the clear need for assistance in prioritizing and managing different design objectives, advances in performance analysis have not commonly translated into guidance in early stage design, as the limits of the traditional design process and a separation of disciplines have relegated performance feedback to later phases.

In order to facilitate better design on a holistic level, researchers in related areas have developed multi-objective optimization (MOO), which is a methodology intended for navigating complex design spaces while managing and prioritizing multiple objectives. However, after reviewing existing design optimization research and considering current usage of optimization in AEC practice, a number of clear research questions arise: How can conceptual, architectural design problems be formulated and solved using MOO in a way that generates diverse, high-performing solutions? What is the best way for the designers of buildings and structures to interact with MOO problems? Finally, how does the use of MOO in the conceptual phase affect design possibilities and outcomes?

This thesis addresses these key research questions, along with a number of secondary questions, through a combination of design case studies, tool development, user experience testing, and historical analysis. First, it presents a conceptual framework for implementing MOO within architectural parametric design tools in flexible, interactive way. Next, it shows the outcomes of a conceptual design exercise in which participants are given increasing access to performance feedback. Finally, through the application of MOO to three long span roof case studies, it demonstrates how MOO can lead to diverse, high-performing results that are difficult to generate through other means, before introducing a new way in which multi-objective techniques can be used to analyze historical structures. Together, these contributions encourage more widespread and effective use of multi-objective optimization in conceptual design, leading to better performing buildings and structures without overly constraining creative, innovative designers.

Key words: multi-objective optimization, design space exploration, conceptual design, design tradeoffs, interactive design tools, structural design, embodied and operational energy

Thesis Supervisor: Caitlin T. Mueller
Title: Assistant Professor of Architecture and Civil and Environmental Engineering

Acknowledgements

I am tremendously grateful for the many advisors, colleagues, collaborators, and friends who made this thesis possible. First, I want to thank my thesis advisor, Professor Caitlin Mueller, who was a source of knowledge and creativity throughout the research process. It would be impossible to mention here all that she has taught me. Her combination of energy, enthusiasm, and sharpness, along with her breadth of skills within the topics of structural design and computation, have made research enjoyable while pushing me to think deeply about the challenges of integrating performance into early-stage design, and how to overcome them.

This research was also significantly influenced by Professor John Ochsendorf, who was instrumental in my transition to graduate school and the early development of my research goals. Throughout my last two years, I am appreciative of his mentorship and guidance, as well as the many opportunities he as provided for me to develop as a designer, teacher, and researcher. I am also grateful for the thoughtful feedback, encouragement, and technical assistance of the other Building Technology faculty: John Fernandez, Leon Glicksman, Les Norford, and Christoph Reinhart. Together with the structures faculty, these researchers and educators have cultivated an environment that is supportive, challenging, and highly conducive to forming and answering clear, meaningful research questions. Through their comments during BT Seminar and other meetings, they have freely lent their time and expertise while strengthening this research.

To the other students currently or previously associated with the Digital Structures and Structural Design Lab research groups, including Renaud Danhaive, Pierre Cuvilliers, Noor Khouri, William Plunkett, Catherine De Wolf, Leonardo Todisco, and Daniel Åkkeson, I am appreciative for their feedback on my research and lively discussions about related topics. More broadly, I am thankful for the Building Technology students with focus areas outside of my own, who have helped me understand other aspects of building design and performance that interact with structure. Notably, these include Timur Dogan, David Blum, Carlos Cerezo, Cody Rose, Chris Mackey, Arfa Aijazi, Irmak Turan, and Jaime Farrell. The many other BT students are also responsible for creating a friendly, engaging, and intellectually stimulating environment, which has made working and playing softball with them a pleasure.

Additionally, I am indebted to two undergraduate research assistants, Anthony McHugh and Jonathas Felipe de Oliveira, who made valuable contributions to this research. It must also be noted that the case study used in Chapter 3 was inspired by the conceptual design of the Curragh Racecourse redevelopment in County Kildare, Ireland. I am grateful for the opportunity to collaborate with Grimshaw Architects in their design process for this project. I would also like to thank Carl Rosenberg of Acentech Inc. and Alejandra Menchaca of Payette for their contributions concerning the acoustic and rain protection evaluation methods, respectively.

The many staff within Building Technology and the Department of Architecture have also been very supportive while conducting this research. In particular, Kathleen Ross has been extremely friendly, reliable, and eager to assist with anything I have needed. I am also grateful for CRON, who have been responsive to my needs and quickly solved the computer problems that happen to occur while doing computational research. I also sincerely appreciate the sources of funding for my study, which include the Leon Hyzen Fellowship and the Department of Architecture.

In closing, I am grateful to my parents for raising me and teaching me how to read and make things, and also to the many friends and family members who have given love and support during my time in graduate school. Finally, I would like to thank my wife Sarah for everything she has done for me over the last two years.

Table of Contents

List of Figures

1 Introduction

At the intersection of architecture, structural engineering, and related fields in building design lies a fundamental pursuit: the generation of creative forms that perform well and respond effectively to a variety of design requirements. As each of these fields has advanced, the list of objectives has grown longer and more complex. When putting together a contemporary building, the design team must consider aspects of structural efficiency, environmental impact, architectural quality, constructability, durability, materials, comfort, economics, and other goals. Some design objectives are easily quantified and analyzed, while others are more qualitative due to the role of aesthetics, expression, visitor experience, and related considerations in architecture. Prioritizing and managing interactions between these design goals, while effectively synthesizing each potential input into a high-performance, sustainable design, are of utmost importance to today's designers.

Many tools have been developed to keep pace with the increasing complexity of building design. Through parametric modeling, the ability to run extensive performance simulations, and the increased ease of accessing information about design precedents, there is significant potential for the buildings of today to be structurally intelligent, energy efficient, comfortable, and durable, while still being evocative from an architectural perspective. However, the traditional design process has not substantially changed in recent years despite the vast improvement in the speed and precision to which the performance of buildings can be predicted. Designers often still approach the challenge of synthesizing design inputs in a way that minimizes the potential influence of performance feedback to slight adjustments that keep within the bounds generated by an initially proposed conceptual design. This fragmented design process that isolates each discipline can lead to design inflexibility and other challenges. In order to facilitate better design on a holistic level, researchers must find ways to improve on the current state of the field by integrating performance-based design tools into early-stage design workflows in a way that makes them accessible and useful for assisting design decisions that involve multiple disciplines and aspects of performance.

1.1 Current practices in conceptual design

The typical building design process begins with a phase called conceptual design. In this stage, architects will make general decisions about the overall shape and geometry of a building through synthesis of experience, knowledge of built work, and input from design tools (Dorsey et al. 2007). Conceptual design is the least constrained of the phases, as an architect or other type of designer is free to imagine any form that satisfies the architectural purpose of the proposed building. Throughout the conceptual design process, the series of design decisions an architect makes plays a large role in the eventual success of the building. In order to make these decisions, the designer must prioritize different design goals and weigh the importance of each, while considering how different inputs may affect architectural quality and performance. This prioritization can be done unconsciously or explicitly if a particular aspect of performance dominates—for example, an architect might choose to emphasize a structurally efficient form for a convention center, or the use of daylighting for an elementary school.

However, despite the freedom of the conceptual design phase and the importance of early design decisions, practitioners typically consider only a few conceptual designs before moving on to schematic design, the next phase in the process. According to a survey of a leading design firm concerning their recent projects (Flager, Welle, et al. 2009), an average of fewer than three alternatives was considered during the conceptual design phase, because it can take more than a month to manage design information and translate between architectural and engineering software programs for a single iteration. Long iteration times are especially common for larger projects, where more building complexity creates a web of interrelated design inputs that can be difficult to traverse. As such, the typical design process often becomes linear, where engineering simulation tools are used to confirm design decisions made earlier, rather than to provide feedback with the intent of changing the conceptual design to respond to performance. The status quo, which involves architects creating geometry without performance feedback and then handing off their designs for engineering analysis, severely limits the impact of such analysis on the final design, regardless of how sophisticated the analysis tools may be.

Another significant barrier to the effective use of performance tools in conceptual design is the balance that must be struck between designer preferences and a pure reaction to performance feedback. Although many design outcomes (such as cost or energy usage) can be clearly measured, architects also consider aspects of a building that are not easily analyzed. Architectural outcomes, such as how a building looks or is experienced, are often significant and can dominate the conceptual design phase. In these cases, although performance simulation tools are state of the art in building practice, many architects would prefer to maintain control over the early process rather than be completely subservient to the results of simulations. Consequently, an effective performance feedback loop is one in which simulation tools are able to guide the design, but architects maintain the ability to express preference.

1.2 Need for rapid generation and evaluation of designs

In order for performance tools to be effectively integrated into conceptual design, they must be able to complete at least two main tasks: the generation of viable design alternatives, and the evaluation of each alternative. These tasks must be executed quickly and in an interface that is easy to use and compatible with existing environments preferred by designers. The initial performance evaluation does not need to be perfectly accurate, but should be reliable enough that a better design identified in the conceptual phase would outperform a worse design once the building has been fully actualized. Ideally, the process of generating and evaluating design alternatives balances a number of competing requirements: speed and accuracy of simulation results, freedom of geometric expression and guidance towards high performers, and easy-to-use versus more detailed feedback.

1.3 Use of computers in conceptual design

Computation is clearly the preferred method for generating and evaluating designs, as both generative and analytical software are ubiquitous in the building design field. Computer programs can perform all of the separate steps required for rapid generation and evaluation of conceptual designs, and are doing so with increasing effectiveness following continuous upgrades to computing power and software capability. In some aspects of conceptual design, computers are even be better than humans at discrete tasks—this is certainly true for performance simulations, but may even be true for the process of design synthesis, as computers have been shown to effectively pursue simultaneous design goals, prioritize them, and pick between trade-offs (Cvetkovic & Parmee 2002).

However, despite their success at specific tasks within conceptual design, there are a number of reasons why computers are not being effectively used as an integrated tool for preference- and performance-based design. The first is that it is often difficult to switch between the programs that execute separate tasks (Flager, Adya, et al. 2009). For example, software such as McNeel's *Rhinoceros* and Bentley's *GenerativeComponents* provides designers with the generative capabilities to create complex geometric forms and explore these forms parametrically. Analysis software, including finite element structural modelers and a wide range of energy usage evaluation programs, can take in information about a design and create precise approximations of how it will behave. Although considerable effort has gone into making the generative and analytical programs compatible, a complicated translation of design information often still occurs when they are used together. The complexity of the translation makes the overall process too slow to be effectively used as a support tool in conceptual design. Furthermore, the required inputs and results of analysis programs are often too extensive to be adequately created, processed, and used by a conceptual designer who is only generating the basic geometry, massing, and form of the building.

The second issue that limits the use of computation in conceptual design is that difficulties arise when computers try to mix "hard" measureable performance evaluations with "soft" architectural evaluations. In general, a computer prefers to solve problems that contain a clean set of mathematical relationships and lead to a specific solution. Performance evaluations that primarily use numbers work perfectly well in this context. However, an architect is not looking for a single answer, and often would like to consider different design solutions based on the architectural effects the designs produce. Researchers have attempted to solve this problem by turning design variables into numbers, facilitating interactivity, or through other methods. These techniques for integrating architectural preference into computer programs have had varying degrees of success, but there is no consensus on the best process, leaving it an open question. Despite the difficulties caused by software translation and non-numerical variable requirements, though, potential methods do exist for using computers as both a preference- and performance-driven conceptual design tool.

1.4 Optimization in design

One promising method for integrating performance feedback into the conceptual design phase is optimization. In simplest terms, optimization is a mathematical technique for minimizing an objective function while subject to a particular set of variable bounds and constraints. It is widely used in different fields of applied mathematics and engineering. In a particular building design, many aspects of building performance could be optimized, as long as the design problem can be reduced to a set of design variables that influence an objective. When a single aspect of building performance is clearly most important, the optimization problem is straightforward to compute and represent. However, in many cases designers are interested in more than one objective. Thus, for an architect, multi-objective optimization (MOO) may be the subfield of optimization that most closely resembles the conceptual design process in its entirety. Multi-objective optimization requires a more complicated set of methods for navigating amongst trade-offs between objectives in generated design solutions. If used in a particular way, MOO could be used to improve current conceptual workflows and generate high-performance designs by addressing a variety of important objectives at once.

1.5 Significant quantitative design objectives and their implications

1.5.1 Sustainability

Although different architectural and structural design optimization problems may call for specific quantitative objectives, most modern buildings share a common, overarching design requirement. In the contemporary design and construction communities, sustainability has become increasingly emphasized to the point that it is always present as a fundamental design goal. Broadly speaking, sustainable design in buildings is concerned with minimizing energy consumption throughout the four main stages of a building's lifetime: materials manufacturing, construction, use and maintenance, and end of life (Bayer et al. 2010). When observing broad energy trends, buildings form the single biggest usage category by a wide margin. In the United States in 2011, buildings accounted for 40% of total energy usage, which was 44% more than the transportation sector and 36% more than the industrial sector (U.S. Department of Energy 2011). This trend is similar in other parts of the world, as the International Energy Agency (2013) estimates that buildings account for over one third of all final energy and half of global electricity consumption. The need for a reduction in energy usage and carbon emissions has been well documented, making clear the importance of sustainability as a building design goal.

The most obvious component of building energy consumption is operational energy, which includes the amount of energy consumed by a building to keep it heated, cooled, lit, and otherwise functioning throughout the year (Bayer et al. 2010). Operational energy has many different components, and each of these parts could be optimized separately. For example, an architect could focus on reducing lighting power density through daylighting, or minimizing heating loads by insulating the envelope, or situating the building geometry to use natural ventilation instead of powered fans. However, for conceptual design, it is often more appropriate to use the broad objective of complete operational energy, since it most effectively addresses the need for overall efficiency and includes each of the individual components. In most existing buildings with relatively long lifetimes, operational energy dominates the other three main stages of a buildings' lifetime, and thus features most prominently into lifecycle sustainability calculations. However, as current design trends improve energy efficiency with the goal of low or even zero emissions, the other stages will become increasingly important.

A second major component of building energy consumption comes in the form of embodied energy, which is present in the structure of the building as well as all other materials and construction processes. For large buildings or design problems with a specific spatial need, such as skyscrapers or long span roofs, structural efficiency has a proportionally larger effect on overall sustainability, and considerations of structural form can dictate the conceptual design. Structurally efficient designs reduce material usage and lifecycle costs, and they are generally safer, more durable, and easier to build. This is especially true because structural efficiency is closely tied to the geometry that is being developed in conceptual design, and because shape matters more than material, sizing, and other aspects of structural design that happen in later phases. Structural efficiency in these cases is also significant because it can have a larger relative effect on the cost and constructability than in smaller, more standard building designs. This research will focus on structurally significant design problems, and thus always consider structure as one of the main design objectives. Structural efficiency can be measured by a variety of metrics, including amount of material, stiffness, or strain energy, and each of these metrics can be used in an optimization problem. However, material amount is likely the objective related most directly to overall sustainability.

When considering both embodied energy of the structure and operational energy in an optimization problem, it is possible to compare the two objectives on the same scale. Most measurements of structural efficiency calculate the weight of material used, which can then be converted to embodied energy or carbon emissions using commonly accepted coefficients for a given material and location. This conversion does not perfectly account for all the effects of structural efficiency (such as improved constructability or durability), but it can give a rough approximation of the relative importance of structural and operational efficiency. Once the embodied energies of the structure and other construction materials of a design are calculated, they can be combined with the projected annual operational energy to estimate the total lifecycle usage of the building.

The possibility of converting both objective measurements to energy raises the question of whether or not a conceptual design attempting to optimize structure and operational energy is actually multi-objective. In some cases, it may be advantageous to simply convert both objectives to energy and optimize for lowest total lifecycle usage. However, for the conceptual design of structurally significant buildings, it is important to have the choice to compare structure and energy separately for a number a reasons. The first is designer preference—an architect may weight the two goals differently, and reducing them to one objective eliminates the precision to which this differentiation is possible. Another is that a pure energy comparison may lose accuracy and utility because money, materials, and carbon now are not exactly equivalent to their counterparts in the future, a detail which would get lost in a single objective of lifecycle energy. Finally, since structural efficiency and material reduction have implications for capital costs rather than operational costs, certain clients may emphasize upfront structural costs if they are trying to finance a project.

As a result, it is often important to the designer that different objectives be held separate during the design process, which allows more flexibility for the designer to exercise preference and control different priorities. This keeps many relevant problems in architecture and structural engineering truly multi-objective, despite these different objectives contributing to overall design sustainability. Through thoughtful navigation of these complex problems, the designers of buildings and other large-scale architectural structures have significant opportunity to minimize emissions and positively impact the future built environment in terms of sustainability.

1.5.2 Other architectural criteria

In addition to structure and energy, a number of other broad performance objectives could be integrated into the conceptual design phase. The most obvious one is cost, although information about cost can often be estimated from the other objectives, since relative cost can roughly correspond to structure and energy. Another potential objective is constructability, which is difficult to define but could be useful in conceptual design. Further architectural criteria that are quantifiable but are often problem specific include daylighting, rain protection, glare reduction, acoustics, and thermal comfort. In addition, basic architectural criteria that do not require sophisticated analysis could also be involved in an optimization problem, including usable programmatic area or spatial requirements such as column spacing. Each aspect of building performance mentioned in this section can contribute to the overall success of a design, and some of them either reduce emissions or must be balanced against environmental considerations. However, given the contemporary design climate, these secondary criteria must often be considered in conjunction with overall sustainability.

1.6 Response

This thesis addresses the limitations of current design processes, demand for high-performance contemporary buildings, and potential of optimization techniques by creating strategies and tools for simultaneously integrating performance inputs into the conceptual design process in a way that produces diverse, high-performing designs. It will emphasize structural optimization and design problems where structural considerations are particularly significant, including stadiums and long-span buildings. However, considerations of building operational energy and overall lifecycle energy usage will also be explored extensively, as will interactions between structure and other related inputs and objectives.

1.7 Thesis organization

This thesis is divided into seven chapters. The first chapter serves as the introduction, provides background information on the conceptual design of structures, and details the problems that motivate this thesis. The second chapter presents a broad overview of tools for building design and analysis, with an emphasis on structural and energy software. It concludes with a description of critical unmet needs and corresponding research questions.

Chapter 3 demonstrates and discusses the benefits of live interactivity in multi-disciplinary, multi-objective conceptual design for structures. It also provides a conceptual framework for implementing multi-objective optimization within architectural parametric design tools in a way that gives designers the flexibility to choose between different optimization techniques. In addition, this chapter presents a design case study of a cantilevered stadium roof to show how designers can effectively set up and navigate architectural design spaces using MOO. This research was accepted for presentation at the International Conference for Structures and Architecture (Brown et al. 2016).

Chapter 4 presents the results of a behavioral case study that demonstrates the effect of access to multi-objective performance feedback and optimization tools on design outcomes. It contains a review of literature concerning idea generation in related disciplines, as well as a methodology for collecting and analyzing the responses of designers to computational environments with progressively more performance-based information as part of an educational exercise. An abstract for this research has been accepted by the 2016 conference for the International Association of Shell and Spatial Structures (Brown & Mueller 2016).

Chapter 5 applies multi-objective optimization to a set of long span building designs to investigate tradeoffs between operational energy consumption and structural embodied energy, with a focus on results that are geometrically diverse and architecturally interesting. In doing so, this research contextualizes the importance of the embodied energy of structural material within the general goals of sustainable design, while revealing specific and sometimes unexpected ways in which the initial embodied energy and long term operational energy performance trade off in different climates. This chapter also includes background on previous applications of multi-objective optimization to building design. Overall, these contributions demonstrate how MOO can be applied to geometrically, architecturally interesting conceptual design problems, leading to diverse, high-performing results that are difficult to generate through other means. Research related to this chapter was published in the 2015 proceedings of the conference for the International Association of Shell and Spatial Structures (Brown et al. 2015), as well as in *Automation in Construction* (Tseranidis et al. 2016).

Chapter 6 uses the MOO techniques discussed in previous sections as a tool for historical analysis while studying the relationship between structural and daylighting performance in the Amiens Cathedral. This research demonstrates how multi-objective optimization, design space exploration, and related concepts can be applied analytically as well as imaginatively, while quantitatively establishing a relationship between the two most prominent aspects of performance in Gothic cathedrals. This research was developed along with Professor John Ochsendorf as part of the MIT course 4.444: Analysis of Historical Structures.

Chapter 7 summarizes the specific research contributions of the thesis, discusses potential impact on the field of computational design, and identifies opportunities for future work.

2 Background: Performance in Conceptual Design

This section describes contributions to the development and application of performance-based tools for conceptual design. It begins with a review of separate analysis and design tools, before moving to recent efforts to integrate rapid performance feedback and guidance into the conceptual design process. Since this thesis extensively examines both structural and building energy efficiency as design objectives, existing tools in both fields are mentioned. Other objectives are pursued in later chapters, and discussions of their analysis methods will be provided where appropriate. Similarly, background information and literature reviews more specific to each chapter are given in the beginning of that chapter. As such, this initial background chapter is not exhaustive, but rather gives a broad review of existing tools and research relevant to the conceptual design of structures, before concluding with critical unmet needs in the field and research questions addressing those needs.

2.1 Existing tools for analysis

Most computational structural analysis software programs are based on finite element analysis. These programs require the user to input geometry, materials, boundary conditions, loads, and other properties of the design they wish to analyze. Once the problem is set up, the user is able to solve for stresses, deflections, forces, reactions, and many other aspects of both static and dynamic structural behavior. A quick Internet search reveals the existence of over 60 commonly used software programs, ranging in their use, capability, and availability. Some of these programs are open-source, whereas many are proprietary and commercially available. A few of the more recognizable programs include ETABS, SAP2000, STAAD.Pro, GSA, and LUSAS. Although some programs are more advanced in terms of speed, accuracy, and available functions, they mostly use the same general procedures.

Although some efforts have been made to integrate finite element analysis into architectural design tools, such as Advance Design by GRAITEC, Robot by Autodesk, and Geometry Gym (Mirtschin) for Rhino, analysis programs are on their own not useful for conceptual design. Many are standalone and require a difficult model translation process to obtain accurate results. Most of the programs cannot analyze parametric design representations, help with geometry generation, or go beyond form verification and member sizing, limiting their usefulness in early design.

Energy usage in a building depends on a variety of complex factors, and simulation programs tailored individually to each of these factors are readily available. However, this review will focus on software designed to execute quick simulations of whole building energy usage. Most commonly used whole building energy programs in the United States can be traced to two freely available simulation engines. The first of these engines is DOE-2, which was developed at Lawrence Berkeley National Laboratory (LBNL) with funding by the U.S. Department of Energy. The most popular software interface for the DOE-2 engine is eQUEST (James J. Hirsch & Associates & LBNL). The second and more recent engine developed by LBNL is EnergyPlus, which models heating, cooling, lighting, ventilation, water, and other energy flows in buildings (Crawley et al. 2001). EnergyPlus is a standalone program without a designer-friendly interface, so it is commonly used with OpenStudio or other plugins that make it easier to set up the model and extract results.

Whole building simulation programs face the same obstacles to integration with early-stage design as structural analysis programs, since a model must be set up carefully to ensure accurate results. An added challenge is the time it takes to run a full simulation, which makes interactivity in conceptual design difficult to achieve. The next section will address recent efforts to remedy these problems and employ simulation software as a design aid, rather than just an analysis tool.

2.2 Existing tools for design

Mueller (2014) identifies the limitations of analysis software for use in conceptual design, and concludes that effective design tools must include capabilities for feedback and guidance. Ideally, such tools would be able to rapidly evaluate given designs, but also guide the user by generating alternatives in response to performance evaluations. The following is a brief survey of research efforts that have attempted to address feedback and guidance.

In structural design, there are a number of existing tools that are effective in rapid evaluation of design alternatives. One is graphic statics, which has been around for over a hundred years and relies on graphical methods to calculate forces in axially loaded structures. Graphic statics was commonly used before numerical methods and finite element analysis become ubiquitous in structural engineering, and recently there have been a number of efforts to bring it back into practice, including Allen & Zalewski's (2010) book *Form and Forces*. There has also been a push to implement graphical methods in structural optimization, including Beghini et al. (2013). However, graphical methods require simple, statically determinate structures, and thus lose flexibility to solve a wide range of design problems.

Another class of design tools focuses on speed of evaluation and interactive exploration by employing real-time structural analysis. Early examples of real-time structural analysis programs include Arcade (Martini 2006), which dynamically simulates the behavior of structures, and Dr. Frame3D (Dr. Software 2009), which allows the user to manipulate loading and other model parameters and immediately see how a structure responds. Another standalone real-time analysis program was developed by Clune (2010), which also offers some optimization functionality. Clune's program permits experimentation with different truss designs while obtaining instantaneous information about the structure's cost, stiffness, and weight. Some of these real-time structural analysis programs have simulation engines that were inspired by or are similar to those found in computer graphics, such as the quasi-static or fully dynamic deformation methods evaluated in Boeing & Bräunl (2007) and mentioned in Parker & O'Brien (2009).

A few structural software programs widely used in building analysis offer internal real-time functionality, such as SAP2000's Model-Alive feature. Other researchers have also developed real-time finite element analysis modelers that operate as plug-ins to architectural software, such as React Structures (Autodesk 2016a) for Revit (Autodesk 2016b) and Dynamo (Autodesk). Karamba (Preisinger & Bollinger-Grohmann-Schneider), a plug-in for Grasshopper (Robert McNeel & Associates), is another recent example of this type of program and will be used later in this research. While these real-time programs have obvious benefits, they are still limited to structural problems that are simple enough to be analyzed quickly, and they do not consider other architectural goals that may be significant in conceptual design.

In addition to the software mentioned above, the most cutting edge structural design tools move beyond feedback and trend towards optimization and the generation of increasingly higher performing designs. Due to the direct relevance of these methods to the overall framework of this research, structural optimization will be given its own section.

Most attempts at integrating whole building energy simulation into early stage design focuses on overcoming its two fundamental limitations—complexity and slowness (Dogan & Reinhart 2013). To do this, researchers have built tools that either facilitate the transfer of design information between programs, speed up the simulations, or both. Morbitzer et al. (2001) identifies key obstacles and proposes ways in which a design tool's interface, analysis, and presentation of results could be simplified to an appropriate level for early stage design. Pratt et al. (2012) developed a generalized protocol for automatically translating between architectural models and modeling software. This protocol was prototyped to allow for building energy consumption to be displayed side-by-side with a changing architectural model using SketchUp, Grasshopper, 3ds Max, and EnergyPlus. Dogan & Reinhart (2013) created a plug-in for Grasshopper that allows Rhino models to interact with EnergyPlus. It also presents the development of an automated method for converting architectural building geometry into 'shoebox' models that can be evaluated more rapidly. These improvements and others have made energy simulations more accessible to designers, but these simulations are still primarily focused on evaluating an existing design and have not begun to generate alternatives and guide designers towards better solutions.

2.3 Optimization background

In order to provide designers with both feedback and guidance, design tools can make use of optimization. In its simplest form, optimization is concerned with finding the best solution out of many potential options that are bound by a set of variable ranges and mathematical constraints. Thus, optimization-based conceptual design tools can guide designers towards alternatives that are the best or near-best design possibilities. An optimization problem contains a number of formal parts, which will be described here for clarity. The bolded terminology is used in this thesis, but other common terms are given for reference.

The **objective function** (also called the fitness, cost, value, or payoff function) represents the quantity that must be minimized or maximized in the problem. The main structural objective used in this research is material weight, but could also be stiffness, strain energy, or deflection. An energy objective used is annual operating energy, although lifecycle cost or a subset of energy usage such as heating load could also be employed. Other architectural criteria are used as objective functions in this research as well.

The **design vector** contains all of the design decisions that must be made in the form of numerical **design variables**. These variables can include geometry or any other property of the design that can be represented by numbers or the indices of a list. Some authors use the term **parameter** or decision variable to describe design variables. However, parameter commonly refers to constants or settings within an optimization problem that can be adjusted for separate runs, and thus this thesis will only use the term variable when referring to the elements of a design vector.

The **variable bounds** represent the limits to which the variables can be adjusted. Additional **design constraints** can be added to ensure that solutions satisfy certain conditions, such as a maximum allowable deflection for a structure.

Together, these elements form the **design space** (or decision space), which contains all potential solutions to the problem. The set generated by performance evaluations during optimization is an **objective space** (also called criterion or cost space), which represents the value of the objective function at every possible combination of design variables. The design space and objective space are linked, and depending on the exact concept and visualization being used, researchers employ the terms in slightly different ways. Regardless, in a simple optimization problem the goal is to search the design space to converge on the best possible combination of design variables (Marler & Arora 2004).

Although the basic elements are present in all optimization problems, a variety of optimization algorithms exist, and the algorithms can be applied in many different ways. Optimization techniques relevant to the conceptual design of buildings will be discussed next.

2.3.1 Structural optimization

The advent of structural optimization was catalyzed by the desire of a few scientists and engineers to find the minimal amount of structural material required to carry a given load. Initially, these problems and solutions were analytical, including Galileo's optimization of a cantilevered beam in 1638 and Michell's (1904) method for finding an ideal cantilevered truss configuration to support a point load. Closed-form solutions such as those presented by Galileo and Mitchell are typically not possible for the more complicated load patterns and geometries found in buildings (Beghini et al. 2013). The arrival of capable computers made numerical, iterative optimization methods achievable instead, including early contributions by Schmit (1981). However, most of these early approaches are limited to simple problems or a single structural typology, an issue the field still struggles with today.

In recent years, a number of more advanced structural optimization techniques have been proposed. One approach involves continuum methods for topological optimization, in which material is distributed freely without regard for member size, shape, or connectivity (Céa et al. 2000). However, engineers often have trouble realizing these particular optimization solutions as physical designs, which limits the effectiveness of this method. The remaining structural optimization methods mostly fit into the categories of ground methods or truss optimization. Ground methods are used for member sizing given an established layout, while truss optimization involves moving nodal coordinates as design variables and changing cross sectional areas in response to the node position changes. Currently, many structural optimization researchers are focused on evolutionary algorithms and shape grammars as a way to cross structural typologies and exert aesthetic preferences on an optimization problem; these efforts will be discussed later.

Although some existing approaches to structural optimization are successful and create diverse, high-performing solutions to a wide range of structural problems, they have had limited influence on design practice. This is at least partially due to their failure to represent the practical needs of a conceptual designer on typical buildings. Pure structural optimization may work for a bridge or an open-air long span roof, but buildings have many other performance goals that must be considered early in the process if designers aspire to global success.

2.3.2 Energy optimization

Optimization is also used in building energy design, and it is becoming increasingly common as genetic algorithms have developed and spread. An early example of energy optimization is given by Coley & Schukat (2002), which optimizes lighting and space conditioning systems to minimize overall energy usage while presenting these designs for aesthetic evaluation by architects. Caldas & Norford (2003) note that energy optimization problems can be used to generate designs that minimize capital and operating costs, emissions, or maximize thermal comfort. The main variables that influence these objectives are building geometry, window surface area, amount of thermal mass, type and layout of HVAC systems, and controls for flow of heated fluids. Many additional examples exist, but since the design of energy management systems involves the manipulation of tradeoffs, most of these examples become multi-objective, such as the multi-criteria search between heating, lighting, and cooling performance given by Méndez Echenagucia et al. (2015).

2.4 Multi-objective optimization (MOO) in design

Marler & Arora (2004) give a detailed literature overview of multi-objective optimization, consolidating different terminologies and algorithms used across engineering disciplines. The goal of multi-objective optimization is to accurately model the preferences of the decision-maker, who is often an architect in the case of conceptual design. There are three main ways this can be done:

In ***a priori*** **articulation** of preferences, the designer indicates the relative importance of different design goals before running the optimization algorithm.

In ***a posteriori*** **articulation** of preferences, the designer picks a single solution from a set of mathematically equivalent, often Pareto-optimized solutions after completing the algorithm.

In **progressive articulation** of preferences, the designer continuously provides input while the algorithm runs.

Each of these methods could be employed in conceptual design, as all three provide an opportunity for the designer to express preferences while being guided by performance feedback towards an optimal solution. In architecture, these preferences are often aesthetic and difficult to translate into appropriate design variables, making the second two methods more attractive, but researchers have pursued all three. Specific examples of MOO for architecture are given in later chapters.

2.5 Combining qualitative preferences and performance

Although some existing multi-objective research deals with non-quantifiable variables, many efforts to include designer preferences in an optimization problem come from structural optimization. This is likely because conceptual design primarily deals with overall building geometry and massing, and there is a strong relationship between the aesthetics and structural performance of a form. As a result of this relationship, there are a few recent, useful examples of structural design tools and methods that can help reveal the best ways to create an aesthetic, performance based design process.

Machwe et al. (2005) propose integrating subjective input through interactive evolutionary algorithms where aesthetic preferences are formulated numerically, as designers set the design variables by visually evaluating options. During the first few iterations of the algorithm, a machine-learning agent picks up the designer's preferences and continues running the process. Unfortunately, this requires aesthetic preferences to be based on rules, such as the assumption that symmetry and slenderness are always good and their opposites are always bad, which is not always the case in architecture. Shea et al. (2005) combined a parametric modeler with a structural design program to stimulate the exploration of new shapes, while illustrating the pros and cons of this approach in the design of a stadium roof. Although the authors automate the translation between programs and rapidly iterate, this procedure did not produce diversity on its own until a few arbitrary geometric constraints were added to force more interesting results. Holzer et al. (2008) attempted to create a similar process in a more practice-driven setting to design a stadium roof that is structurally efficient, aesthetically pleasing, and in keeping with the architect's design intent. While quick feedback and iterations made the development of structurally efficient designs more accessible to architects, the authors found that an expert was still required to guide the process towards the best solutions.

A more flexible, accessible design environment is provided by StructureFIT (Mueller & Ochsendorf 2015), which is a browser-based conceptual design tool that allows users to progressively express preference by selecting parent structures for the next iteration of an interactive evolutionary algorithm. Danhaive (2015) developed Stormcloud, which extends the main functionality of StructureFIT into the Grasshopper environment. These tools have been shown to achieve significant diversity in generated designs, while also improving on the performance of traditional geometries in a number of structural applications. However, StructureFIT is purely concerned with truss optimization and does not take into account the relationship between competing objectives. While Stormcloud is agnostic to design variables, objective function, and evaluation method, it still only allows for one quantitative objective. Tseranidis et al. (2016) also make contributions by using surrogate modeling and other techniques to speed up analysis, and Mueller (2014) relies on grammar rules to generate diversity in conceptual design. Both of these ideas could be applied to multi-objective optimization design problems for buildings, but there is a clear need for design tools that offer multi-criteria functionality within a similarly flexible, accessible, interactive framework.

2.6 Critical needs and research questions

In summary, the historical separation of design and analysis software has created a practice environment in which it is difficult to implement performance feedback loops that guide early-stage design. Although optimization techniques can help generate this system of feedback and guidance, the many non-quantifiable design goals of the field of architecture complicates their application. Researchers have nevertheless made significant strides in developing methods for integrating performance into the conceptual design of structures. However, many of these contributions must be improved upon, more thoroughly studied, or extended to multi-objective problems to achieve greater relevance.

After reviewing current practices, noting the limitations of existing analysis procedures, recognizing optimization as a conceptual design tool, and identifying the most important design objectives, a number of clear research questions arise:

- How can multi-objective optimization conceptual design problems be formulated and solved in a way that generates diverse, high-performing solutions?

-What is the most useful way for designers to interact with multi-objective optimization problems?

-How does the use of multi-objective optimization in conceptual design affect design outcomes?

-What are the most important design tradeoffs, and how should they be handled?

-Are there other research applications of multi-objective thinking?

These research questions will be addressed in greater detail by the specific investigations of the following four chapters.

3 Interaction and Visualization in Multi-Objective Optimization

In the current era of architectural and structural design, computational and numerical methods have expanded the shapes and forms designers are able to analyze and build. However, the software that has enabled this freedom is also largely responsible for removing considerations of performance from the early design phase. Since architects are now able to model complex geometries computationally without reference to gravity or other considerations, a typical workflow of architects coming up with conceptual designs before passing them off to engineers and other specialists for analysis has become commonplace. This process, along with the accompanying difficulty in translating between different software models, severely limits the ability of designers to create a performance feedback loop that influences overall building forms. As a result, sophisticated techniques such as structural optimization, which can be used to reduce the cost and embodied energy of a building design, are often not used in the design process.

Furthermore, architects must simultaneously consider a variety of design objectives early on in the process, including ones than can be quantified, such as structural and energy efficiency, as well as others that cannot, such as aesthetic expression. Many of these objectives may trade off with one another, meaning any optimization process would require designer intuition and input to make overall design decisions. In order to facilitate better designs on a holistic level, optimization techniques must help designers navigate a complicated design space by providing functionality for searching through many designs options, prioritizing different objectives, and presenting rapid and reasonably accurate performance feedback. As mentioned in the background section, multi-objective optimization (MOO) is a methodology designed for this purpose, and if used appropriately, it can account for designers' needs and guide them towards high performing solutions in conceptual design.

MOO is commonly applied in fields as diverse as finance and aerospace engineering, and researchers have developed a wealth of methods and algorithms to apply to various problems. Moreover, parametric modeling is already a commonly used tool in many architectural design firms, which lends itself to optimization. Yet for complex problems with high dimensional design and objective spaces, current architectural design tools lack the functionality and accessibility necessary for enabling widespread use of MOO techniques. Many conceptual designers also have little experience with setting up and using multi-objective workflows, since their development has been geared towards more pure engineering problems than expressive architectural design. Specific architectural designs processes or designers themselves might also prefer a wide range of techniques within the broader field of MOO. Thus, there is a need for research contributions that address the accessibility, ease of use, and flexibility of MOO processes for aiding conceptual designers in their search for diverse, creative forms.

3.1 MOO for conceptual architectural design

Many initial contributions towards building accessible interactive optimization tools that consider aesthetics and other non-quantifiable objectives come from the field of structural engineering. Von Buelow (2012) developed ParaGEN, which combines performance feedback from a variety of simulation programs (structural, lighting, acoustical) with aesthetic preferences in design, also through the use of interactive evolutionary algorithms. The fields of mechanical, aerospace, and product design have also contributed significantly to the theory and application of MOO or Multidisciplinary Design Optimization (MDO), which is related. These contributions including many numerical optimization techniques (Vanderplaats 1999), overviews of simplified optimization workflows for engineering (Arora 2004), and the concept of isoperformance (de Weck & Jones 2006).

Other researchers have proposed and tested MOO workflows specifically for architectural design. Asl et al. (2014) establish and test an optimization method for whole building energy performance and daylighting using a parameterized BIM model, illustrating clear trade-offs in window size between usable daylight and thermal performance. Flager et al. (2009) present a methodology and case study results for MOO while considering structure and energy as objectives. Krem et al. (2013) use MOO to demonstrate how the position of a structural core and the shape of the floor plan in a high-rise building affect its energy and structural performance differently depending on the climate and location of the building. Quaglia et al. (2014) generate a Pareto optimal solution set that trades off between structure and energy for the design of origami-inspired, rapidly deployable shelters. Mendez Echenagucia (2013) presents shape optimizations for a number of competing architectural performance objectives, including a study of the acoustic and structural aspects of shell concert halls (Méndez Echenagucia et al. 2014). Although these research contributions show the immense potential of MOO for use in design, few are focused on how conceptual designers can easily and interactively apply MOO to arrive at satisfying design solutions. Many proposed workflows also require a lengthy setup and computationally intense simulations, or complicated information transfers between multiple pieces of software, which may be beyond the scope of some designers.

3.2 Existing tools for MOO

A number of computational tools that can facilitate some MOO techniques have already been integrated into typical conceptual design software. Both evolutionary and gradient-based solvers already exist in parametric software, including Galapagos and Goat (Flöry) for Grasshopper. The Core Studio at Thornton Tomasetti has developed TT Toolbox, which includes a range of different computational tools that may be useful, such as a brute force solver for enumerating the design space and various visualization components. Another plug-in for Rhino, Octopus (Vierlinger), was designed to apply evolutionary principles to parametric design problems with multiple objectives. It incudes features such as searching for tradeoffs, forcing diversity of solutions, changing objectives during a search, and visualizing and exporting results. A lighter version, called octopus.E, allows users more flexibility in picking and choosing only some of these functions. A similar tool, called Optimo (Rahmani), has been implemented for the parametric modeling plug-in Dynamo, which interfaces with many different architecture and engineering software. Users of Optimo can also utilize multi-objective evolutionary algorithms to help explore tradeoffs and generate optimal designs within a design space. As mentioned in the last chapter, Clune (2010) wrote a program that allows for real-time, interactive structural design optimization while pursuing the objectives of weight, compliance, and cost. However, this program is limited to 2D structural applications.

3.3 Research contributions

Some of these existing tools are limited in that they take an all or nothing approach to implementation of MOO, which limits flexibility for designers that prefer particular methods. Many also require sophisticated software knowledge to implement and efficiently move through a design space to select the best options. This greatly reduces the number of designers that can use MOO effectively, or at least increases the time and effort a designer must spend considering performance. Ideally, future implementations will be as integrated as possible with typical conceptual design workflows, while building on the functionality of existing tools to create systematic workflows for applying a number of potentially useful MOO techniques.

In response, this chapter describes a number of multi-objective optimization methods and discusses how conceptual designers can interact with them most effectively, while also considering how specific aspects of each method affect the eventual design solutions it generates. It also proposes a workflow of necessary components for enabling generalizable, interactive multi-objective optimization for architectural applications. To illustrate the potential of interactive MOO for architecture, a design case study of a cantilevered stadium roof is presented. The case study gives visualizations associated with each interaction mode and illustrates how designers can use these methods to formulate a design space and select a best design.

3.4 Designer interaction with optimization

3.4.1 Free exploration with performance feedback

The first method for integrating performance feedback into the conceptual design phase is simply providing performance "scores" onscreen while a designer is interacting with a parametric model. This method depends on the designer to prioritize objectives and weigh tradeoffs with only the most basic computational assistance. However, despite the reliance on user intuition to move towards better performing designs, this method allows for complete design freedom within the parametric model. In an architectural workflow, it only requires evaluation methods for each dimension of performance, rather than additional optimization algorithms, and it most closely replicates a traditional parametric workflow by using sliders or other means of parametric navigation.

3.4.2 Composite functions

A second prominent method for multi-objective optimization is to form a composite objective function of the weighted sum of the objectives. In this method, the weight of each objective corresponds to a preference factor assigned to that particular objective (Deb 2001). Although this greatly simplifies the problem by reducing it to a single-objective optimization and opening up the possibility of a single best solution, a user must still decide on these preference weights. In a conceptual design scenario, a designer would likely have to try a number of different weight combinations to arrive at a satisfying design. This can be done systematically through sampling, or interactively as part of a creative process. Special care must be given to architectural problems in which different objective functions have vastly different units or sensitivities, as certain objectives or combinations of correlating objectives can dominate the problem.

3.4.3 Sampling and Pareto optimal design sets

Broadly speaking, the most common methods for multi-objective optimization require two main steps: finding multiple tradeoff solutions with a wide range of objective values, and choosing one of the obtained solutions using higher level information (Deb 2001). One of the most effective methods for initially finding the range of best possible solutions uses the concept of Pareto optimality. A design is Pareto optimal for a given design space if it is not dominated by another design that performs better in every other dimension of performance. Thus, the Pareto optimal set of a given objective space likely includes the most important options for performance-conscious designers to consider.

There are a number of ways to determine—or at least approximate—the Pareto front of an objective space. The first is to simply brute-force sample the entire design space and use a Pareto-culling function after each design has been evaluated. The sampling can be random, on a grid, or use a more sophisticated algorithm such as the Latin Hypercube (McKay et al. 1979). One advantage of sampling is that it can generate many different designs a particular distance away from the Pareto front, but that an architect may want to consider for non-performance reasons. However, brute-force sampling can also waste considerable computational time while exploring large portions of the design space that are poor performing. The other method is to use an evolutionary algorithm such as the NSGA-II that approximates the Pareto front more quickly by breeding progressively higher-performing generations (Deb et al. 2002).

Another technique that loosely fits in the family of multi-objective optimization sampling is single variable sampling, where a designer changes one particular design variable while holding all others constant, and then visualizes the performance. This technique is useful when there is a preferred design, determined through some immeasurable architectural criteria, and designers wish to understand the importance of each variable to the problem and the direct effect of changing one of these variables from a given starting point. When the range of this single variable is plotted against each objective function, the slope of the objective function can indicate how sensitive it is to that particular variable.

3.4.4 Other methods

Other promising workflows exist, but many of them are a combination of the ones already mentioned. For example, sampling could be coupled with a composite function so that different preference weights are sampled thoroughly, but an individual optimization occurs at every step of the sample. In addition, promising single-objective interactive evolutionary workflows, such as Stormcloud, could be used with a composite function, or have their functionality extended to pursue multiple objectives by randomly distributing the evaluations of a generation across multiple objective functions. Thus, the methods listed in this thesis are not exhaustive, but provide a useful overview for designers considering MOO in conceptual design.

3.4.5 Tools to facilitate these interaction modes

Although various workflows exist for each of these MOO techniques, they often require multiple transitions between software and manual manipulation of data. This section introduces a suite of tools created within Rhinoceros that are designed to work together to facilitate various MOO methods. Figure 1 gives a flowchart describing the connectivity of the developed components. Components 1-3 have been used in various research, teaching, and design exercises, but components 4 and 5 are still in development. Component 1 reads in design variables and their ranges and generates a list of design vectors that represent the design space. Geometry generation and objective functions are design specific, and must be set up individually by the designer. Component 2 records screenshots and performance, while component 3 contains an implementation of the NSGA-II MOO algorithm for use in parametric design (Durillo & Nebro 2011). Component 4 uses surrogate modeling to approximate computationally intense objectives, and component 5 helps organize and visualize overall design data for inspection.

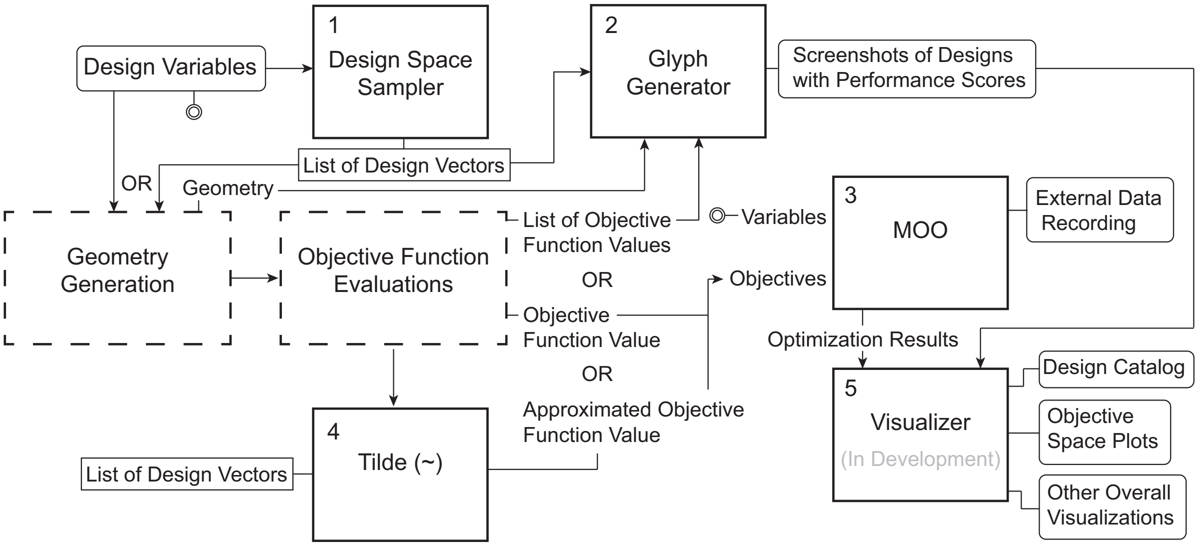

Figure 1: Process flowchart for implementing a MOO workflow

3.5 Case study demonstration: cantilevered stadium roof

3.5.1 Design space formulation

The potential outputs and utility of these interactive optimization methods will be demonstrated through their application to the architectural design problem of the cantilevered stadium roof. The cantilevered stadium form is commonly found at racetracks or other programs that require a covered grandstand. Many examples of creative structural solutions to this problem exist, including the reinforced concrete designs of the Zarzuela Hippodrome and the Miami Marine Stadium. This chapter utilizes a steel truss that sits on two column supports and cantilevers out over the crowd seated below (Figure 2). A parametric model of a cantilevered roof was created using Rhinoceros and Grasshopper. The model is made of a mixture of 2D and 3D geometry. The structural truss represents one section of the roof, corresponding to a single set of columns. The actual design would include a number of these trusses arrayed longitudinally with additional elements spanning between the trusses, but the partial reduction to 2D allows for faster evaluations and visual clarity.

Four design variables drive the parametric logic of the truss: overhang length, backspan length, truss depth ratio denominator, and height of the free edge. These variables were selected because they drive the overall geometry of the roof in a way that allows for design variables to clearly effect performance in multiple architectural objectives. The truss model is projected onto the existing geometry of a crowd seating area and a hospitality suite above, which were previously dictated by the program of the stadium. It is assumed that the surfaces of the roof roughly follow the outer chords of the truss, although a slight offset at the tip location was created to allow for bullnose detail on the edge of the roof.

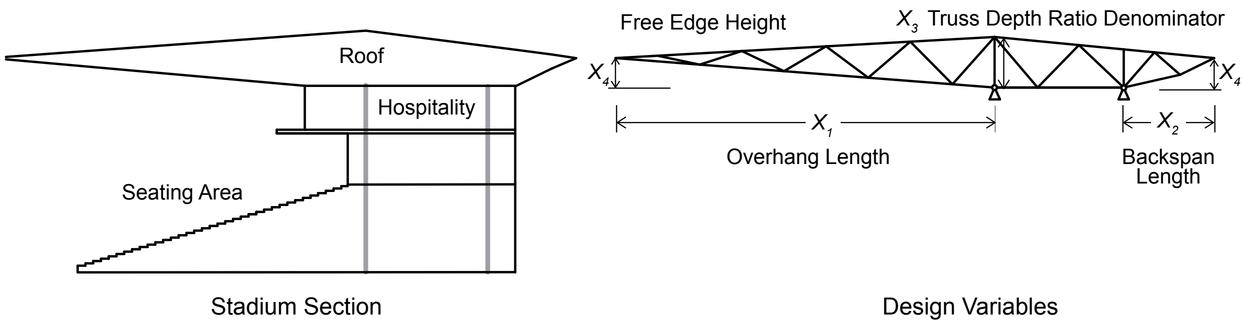

Figure 2: Geometric description of the design space for the cantilevered roof design case study

Design Space Variables		
Symbol	**Variable**	**Range**
x_1	Overhang Length	20 m < x < 30 m
x_2	Backspan Length	17 m < x < 25 m
x_3	Truss Depth Ratio Denominator	5 < x < 9
x_4	Free Edge Height	0 m < x < 4 m

Table 1: Design variables for the case study

3.5.2 Objective function selection

Four different performance objectives were selected for the multi-objective optimization process. The evaluations functions used for each were simplified so that the interaction with the model yields nearly real-time feedback. Although some of the MOO methods described above can be applied to models with longer evaluation times if an automated way of collecting results is established, the methods that require interactive articulation of design preferences require fast evaluations. Table 2 provides the assumptions used for each performance evaluation used in the model. The first objective is structural weight, where a design that uses the least amount of steel is considered to be optimal. To determine this weight, a finite element model of the truss was created using Karamba, which is a plug-in for Grasshopper. Karamba's sizing functionality was used, which reads in the applied loads, runs an analysis, and determines the smallest section size for each element that can provide adequate capacity to resist forces and buckling.

The second objective is rain protection, which is measured as the percentage of the crowd that would not be touched by blowing rain at a given windspeed and orientation. Although a computational fluid dynamics simulation could give a more accurate profile of blowing rain penetration into the crowd based on drop size, the effect of this penetration can be approximated geometrically by drawing a line from the tip of the roof at an angle based on the wind speed and droplet vertical velocity. In addition to rain protection, the shaded area of the crowd was also calculated and considered as an objective function, but this metric was deemed too similar to rain protection and less important in certain climates.

The third objective is sound dispersion. Acoustics in architecture is a complex field, and measurements of acoustic performance can be difficult to reduce to a single geometric objective function, especially since material selection often matters more than shape. However, it was assumed for this design that sound dispersion was desirable, and a flat or concave shape to the underside of the roof would be worse than an open, convex shape. Thus, the angle of the bottom chord of the truss was used as the objective function, where the largest angle is considered the best sound performance.

The final objective involves daylighting in the hospitality suites above the seating area, which is affected by the shape of the roof. A number of different evaluations and tools were considered for this metric, including those that run full annual simulations or calculate a daylight factor for one specific date and time. In the end, it was decided that the portion of sky visible from the hospitality suite, which can be calculated geometrically, was an acceptable approximation. The calculation involved finding the angle between a horizontal line starting at the base of the suite and a line passing from this point through the tip of the roof.

Objective Function Assumptions		
Objective	**Metric**	**Evaluation Method**
Structural Efficiency	Weight of steel (min)	FEM + Member sizing
Rain Protection	Portion of crowd protected (max)	Intersection: rain trajectory and seats
Sound Dispersion	Angle of lower truss chord (max)	Geometric measurement
Portion of Visible Sky	Angle of line: hospitality to roof tip (max)	Geometric measurement

Table 2: Objective function evaluation methods and objective function metrics

3.5.3 Application of optimization methods

Using a combination of existing software, native Grasshopper functionality, and custom tools, each of the potential MOO workflows for design was implemented on the model of the cantilevered stadium roof. The limits of design space for this problem were based on ranges that were considered acceptable from an aesthetic point of view, and the performance outputs were normalized, with the worst possible performance receiving a normalized score of '0', and the best performance receiving a score of '1'.

3.6 Results and visualizations

The most basic method for interacting with multi-objective feedback is setting up a catalogue of generated designs. A small example catalogue for the cantilevered stadium is given in Figure 3, and it contains a glyph of the roof truss and a few typically used performance data visualizations. This type of glyph communicates information rapidly and visually and is appropriate for comparing all sorts of designs in which the variables are primarily geometric. Although this catalogue is presented statically on a page as separate design options, these visuals could be made interactive and dynamic using the tools introduced in Section 3.4.5. In this case, the graphics would update in nearly real time during free parametric exploration, depending on the evaluations, and allow the designer to weight objective importance, test performance sensitivity, and consider aesthetics of the design while using intuition to drive towards a satisfying design.

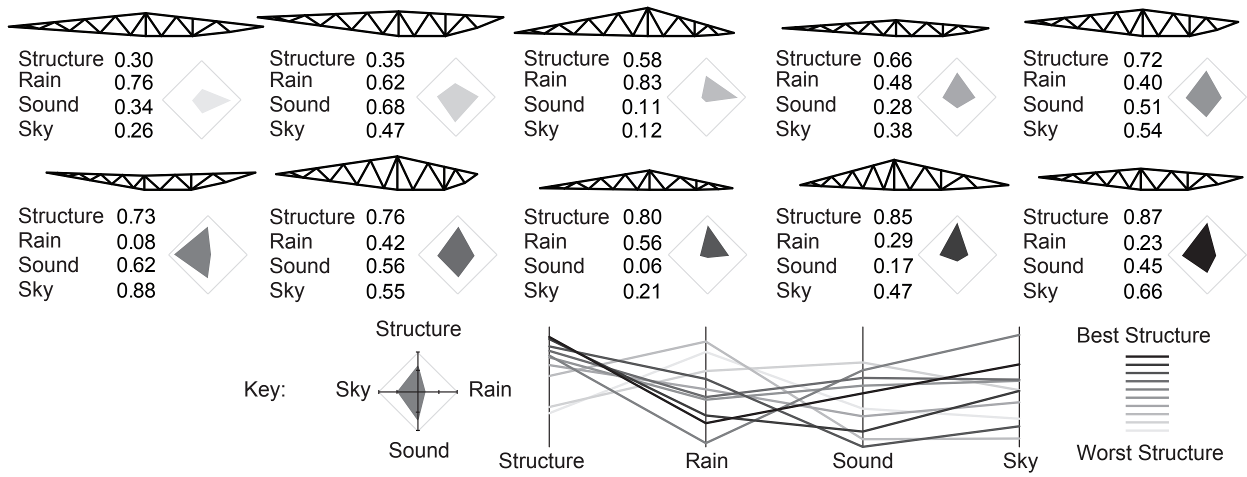

Figure 3: Example of a catalogue from design space sampling

If a composite function is used, the different objective weights could be projected onto standard glyphs and visualizations, which gives additional information to designers testing different preference weight combinations. As illustrated by the cantilevered roof designs in Figure 4, performance in certain dimensions may be easy to obtain without significant weighting, suggesting that there are objectives, such as sound dispersion in this example, that are largely independent and do not have major tradeoffs. Other objectives may dominate the composite function if a significant portion of the design space has high performance in that single dimension, depending on the normalization. Designers might also notice that certain objectives correlate closely with others, and should be eliminated from the problem. Since each of the objectives can be in vastly different ranges and units, a designer can use this composite function MOO optimization technique to test various tradeoffs, correlations, and sensitivities. This process can be done gradually to build intuition about the design problem or systematically to generate more design options, but it ensures that each selected design is optimal for a particular combination of designer preferences, ultimately leading to a successful design.

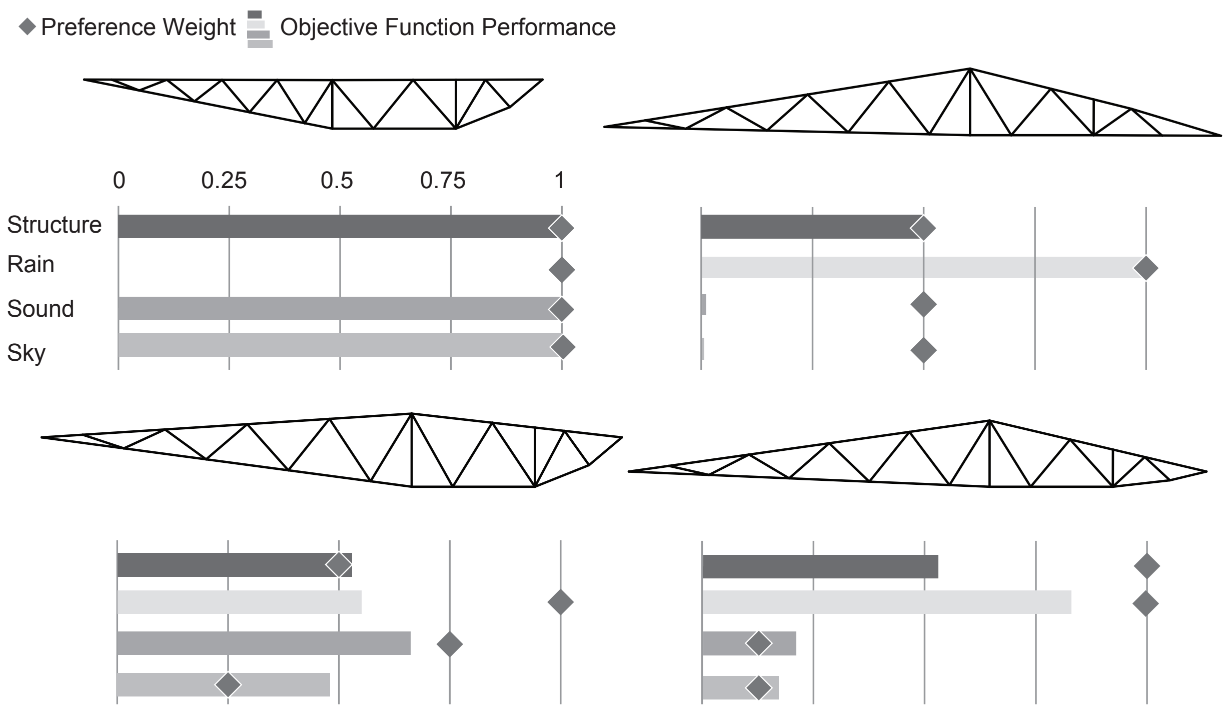

Figure 4: Optimal designs using different objective functions

Another typical visualization, especially useful while exploring a sampled design space or a generated Pareto front, is to plot the objective space with each design represented as a point. When there are more than three objectives, this can only be done using projections to various bi-objective plots, such as the ones shown in Figure 5. In these plots, the designs are colored based on structural score, and the best designs are found on the top right. At this scale, only the trends can be viewed effectively, but within a parametric design space this could be made interactive, where a user could click on one design and have it highlighted on every other plot. Overall, these plots give context to the set of highest performing designs, as well as indicating clear correlations (visible light and sound dispersion), almost linear tradeoffs (rain versus both of these), and curved Pareto fronts (structure and rain protection). By setting up a Pareto front, a conceptual designer is already moving towards high-performing designs, but likely has a variety of visual options to choose from within the set.

In the case where a designer already has a preferred design for aesthetic, geometric, or other reasons, single variable sampling can prove useful. This technique, sometimes called one at a time analysis, can show the direct relationship between design variables and performance sensitivities by essentially visualizing the performance gradient at a particular point. Figure 6 gives an example of single variable sampling from a cantilevered roof design with preferred variables.

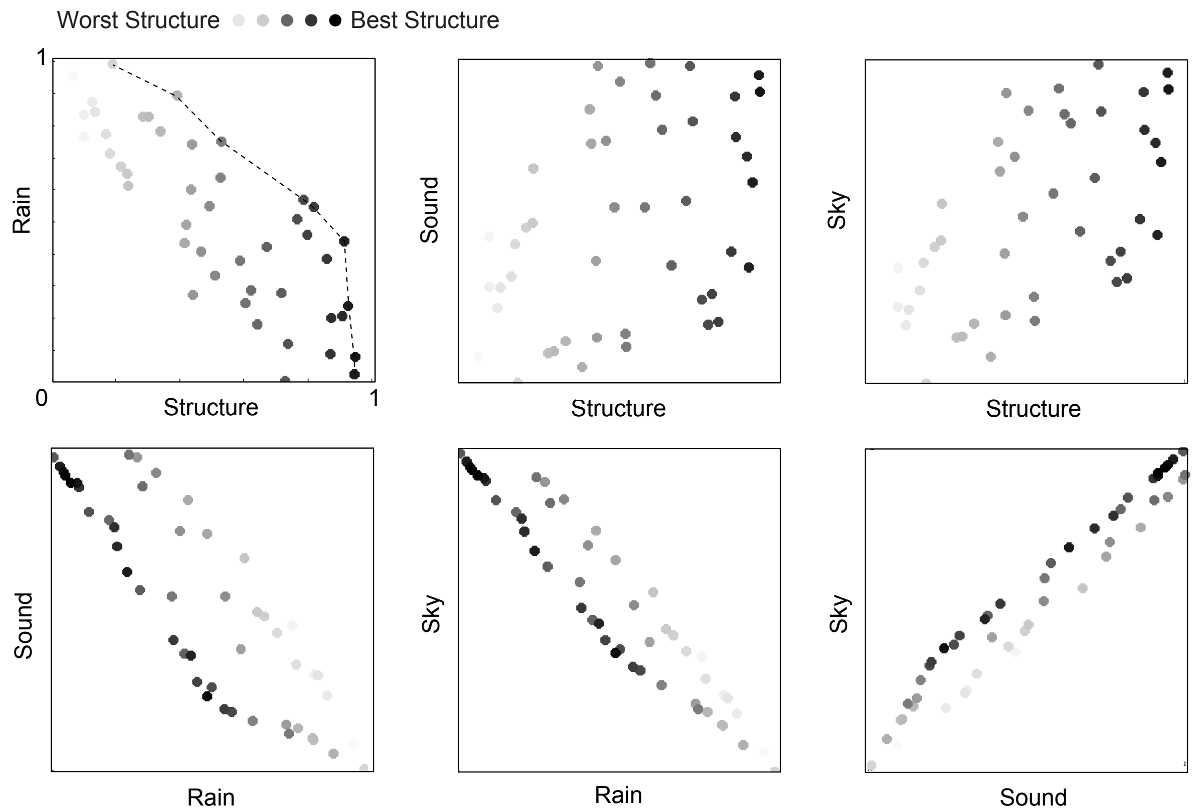

Figure 5: Projections of the Pareto front to 2D bi-objective plots for each objective

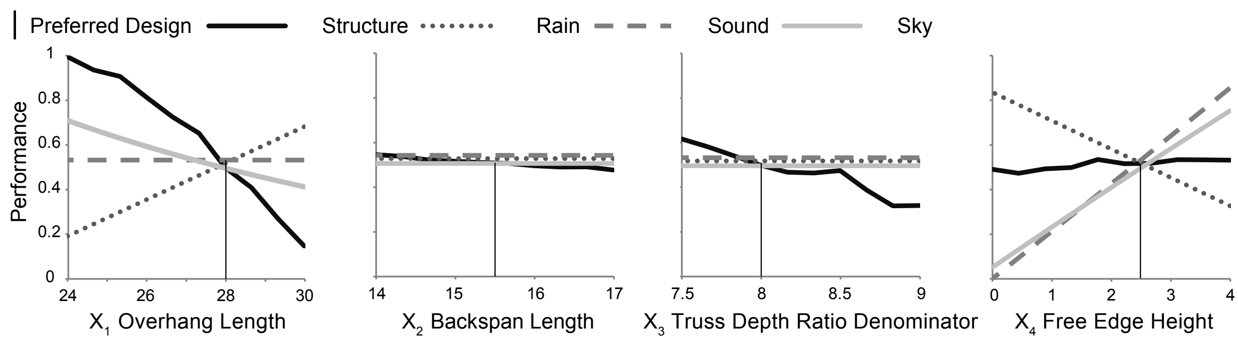

Figure 6: Results of a single variable sampling exercise for a preferred design

3.7 Future work and concluding remarks

Currently, many of these generated visualizations require the export of data from the parametric modeler to another tool. Extensions to the already completed tools, which are being developed by the author and collaborators and are presented in Section 3.4.5, would add this functionality as its own interface within the parametric modeler. Another significant consideration when applying

MOO is computational speed, which can limit interactivity in MOO. In light of this, the author and collaborators are currently developing and testing surrogate modeling tools for evaluations built in parametric modelers, which could approximate longer simulations with a less accurate but much faster evaluation, enabling free play, composite functions, and generally enhancing overall interactivity.

This future work will build on the contributions of this chapter, which has discussed multi-objective, interactive optimization methods that are well within reach of conceptual designers who are already using parametric design tools. This chapter has also presented a number of computational tools developed to facilitate MOO methods, while demonstrating their usefulness through visualizations, which are fundamental to interactive design methods that must consider and balance both quantitative and qualitative design goals. In addition, the case study of the cantilevered stadium roof provides a first concrete example of how designers can effectively formulate and navigate performance-based design spaces. These contributions have the potential to make MOO more accessible and successful in architectural and structural design.

4 The Effect of MOO Performance Feedback on Design Outcomes

In recent years, researchers in the area of computational design for structures have emphasized the development of tools that enable both performance feedback and guidance in early stage design. These contributions respond to limitations of the traditional design process in which design and analysis software are entirely separate, making substantial design iterations that respond to performance difficult and time-consuming to pursue. As a result, engineers and other specialists are often given a rigid design and limited to small adjustments or simply "making it work". Researchers assert that when performance analysis is present in the conceptual phase, which is when a design is most flexible, it has much greater potential to improve the performance of an eventual building or structure. Due to its context of structural design, much of this research is concerned with large, global decisions such as typology, material, or specific geometry that are difficult to change later in the design process, but typically have a large influence on building performance.

This chapter tests the assertion of improved performance by investigating the effect of performance feedback and optimization techniques on architectural design processes through a behavioral case study. In the study, a group of design students are given a structural design problem and provided with a series of increasingly performance-driven computational design tools to complete the task. The designs chosen by the students are then analyzed for quality and diversity while comparing outcomes across different design environments. Overall, this research seeks a better understanding of the link between performance data and design outcomes in early stage design.

4.1 Background: design studies

Design studies, which focuses on developing an understanding of design processes across areas such as engineering, product, architectural, and urban design, is an established academic field with robust supporting literature. An initial contribution towards establishing the field is given by Cross et al. (1981), which argues that design should be regarded as a technology rather than a science, since design and technology include the application of knowledge other than the purely 'scientific' kind. A thorough review of early design theory and methodology in mechanical engineering is provided in Finger & Dixon (1989). More recently, Cross (2004) presents a review paper addressing the nature and role of experts in design, including topics such as empirical studies of design cognition and comparisons between the processes of novice and expert designers.

Within the examination of design processes, behavioral studies that test the relationship between different engineering design environments and idea generation are common. Shah et al. (2000) establish experimental guidelines for evaluating conceptual design strategies, and these guidelines are widely followed. Mckoy et al. (2001) analyze the influence of design representation on the effectiveness of idea representation. Schlecht (2013) tests the impact of prototyping environments on ideation, while Faas et al. (2014) address the question of whether or not designers who are more engaged, as measured by presence and immersive tendency questionnaires, produce better designs. Moreno et al. (2014) even push outside the domain of engineered artifacts by exploring the topic of design by analogy and its application to transactional problems.

Architectural design provides its own developing history of design processes, which have recently shifted from drawings and physical modeling to computational, generative, and performance-based parametric and morphological models (Oxman 2008). The requirements of an architectural design and subsequent decision-making processes are often subjective and particular to the field, which can complicate efforts to evaluating the quality and novelty of designs using traditional means. While specific architectural design process have been studied from a behavioral standpoint (Suwa & Tversky 1997, Goldschmidt 1994), there is a clear need for research contributions that explore how designers interact with newer digital design methodologies.

The recent emphasis on high-performance architecture and the integration of analysis tools into conceptual design workflows demands additional testing on how these feedback mechanisms influence the design process. Many recent contributions to the computational design of structures are based on the assumption that overall design quality can be improved when performance feedback and guidance are part of the conceptual design process (Mueller & Ochsendorf 2015). Some efforts have been made to test this theory, including Arnaud (2013), which analyzes how architects and engineers generate structural designs when given access to an environment that contains both performance feedback and guidance for a flexible, parametric model. However, many architectural and even structural design problems are multi-objective, requiring the synthesis of feedback in multiple dimensions. This chapter proposes a behavioral study in which participants must pursue multiple objectives simultaneously, with the intent of uncovering more generalizable knowledge concerning the relationship between performance input and design quality in conceptual design.

4.2 Methodology

4.2.1 Participants

The study involved 26 undergraduate and graduate students at the Massachusetts Institute of Technology. The vast majority of participants were in their second year of a graduate architectural degree program, while 3 were pursuing degrees in civil and environmental engineering. All of these students were taking an architecture class in building structural systems and were given the design task as an educational exercise. Students in the class were provided the option of participating or requesting that their chosen designs not be included in the aggregate study results. No compensation was provided to the participants. Materials including software files and instructions were distributed online, and participants were able to complete the exercise on their own computers and upload answers at their convenience. Most participants had at least a year worth of formal design training, including two structures courses, and were concurrently taking a comprehensive architectural design studio.

4.2.2 Procedure

Each experiment consisted of three design phases plus a short survey about the overall experience. The phases consisted of the same design problem, but with different conditions that progressively gave participants more access to performance-driven computational tools. In the first phase, students were given a design problem and a corresponding parametric model that defined the potential geometry of the architectural solution. Participants were free to adjust any of the design variables while the corresponding geometry updated in real time. The second phase consisted of the same parametric model and variable sliders, but with performance feedback also updating on the screen along with the model geometry. This feedback included both actual and normalized performance values for a variety of design objectives, as well a simple bar graph to visualize the performance scores. In the third phase, participants were instructed to use a number of optimization tools to guide their design exploration. The students were asked to record their "favorite" design along with 5 desirable alternatives for each phase. After completing the exercise, participants were given survey questions comparing the three different design environments in terms of ease of use, quality of outcomes, likeliness of using in their own workflows, and related topics. Participants were provided with a prepared form to record all of their answers.

4.2.3 Task

The task given in this study was the design of a canopy structure for the outdoor seating area of the restaurant (Figure 7). Due to the desire for a free edge and the ability to anchor into the wall above, the hypothetical client asked for a cable-stayed structure. The main topology of the structure is formed by beams that cantilever out from the wall and are supported by a series of cables, which also anchor into the wall. Within this main geometry, participants were allowed to adjust the anchor point spread, height of cable and beam connections, height and horizontal distance to the canopy tip, number of cables, curvature of the canopy, and the structural material. A full description of the adjustable design variables is given in Table 3.

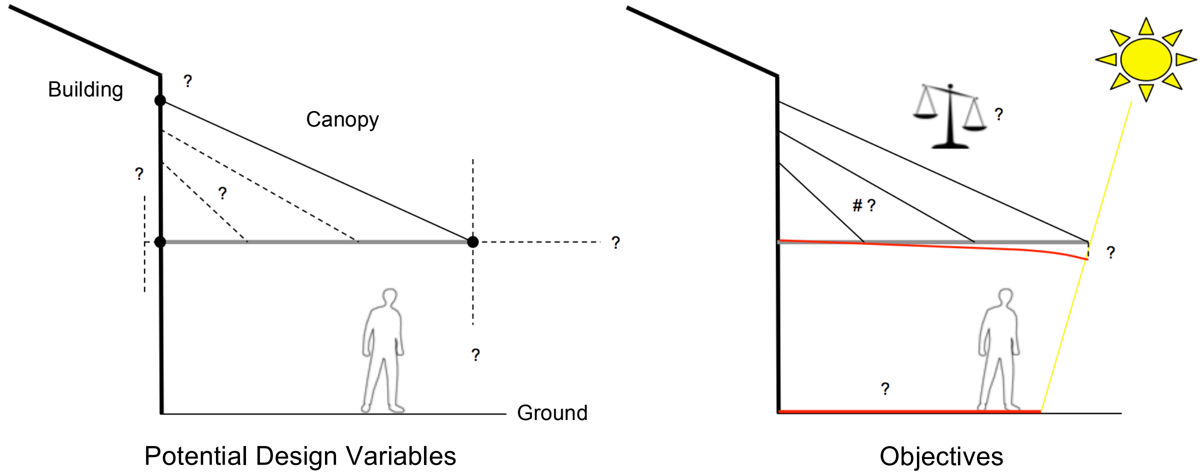

Figure 7: The design problem, a cable-stayed canopy roof, as presented to participants

Design Space Variables		
Symbol	**Variable**	**Range**
x_1	Anchor Point Spread	0 < x < 1
x_2	Height of (Top) Cable Anchor Point	2.44 m < x < 9.14 m
x_3	Height of Canopy Anchor Point	2.13 m < x < 7.62 m
x_4	Length of Canopy	1.52 m < x < 12.19 m
x_5	Height of Canopy Tip	2.13 m < x < 4.57 m
x_6	Number of Cables	1 < x < 10
x_7	Curvature	0 < x < 1
x_8	Material	Steel, Aluminum, Wood, Carbon Fiber

Table 3: Design space variables and bounds for the design task given to participants

As part of the exercise, a number of computational models were built to quantitatively assess the performance of each design. These models calculated the shaded area provided, carbon emissions due to the materials, number of connections, and maximum deflection of the design. The shaded area provided is an architectural metric that would be a priority of the client, and it was calculated geometrically using an assumed static sun angle. The number of connections roughly approximates how difficult the design is to construct. Carbon emissions and maximum deflection are structural metrics and were calculated using finite element analysis. Carbon emissions is linked to material selection and quantity, while maximum deflection indicates response to loads. Table 4 gives additional information about the objectives and their evaluation methods.

The stated goal of the design exercise, which is implicit in most design situations, was to prioritize, navigate, and explore interrelated performance and aesthetic objectives to arrive at a satisfying design solution. Participants were given the freedom to choose which objectives were the most important and encouraged to use their own design sensibilities when judging the expressiveness of the solution. Each objective was normalized so that the best performing design received a score of '1', and every other score is a multiple showing how much worse the design performs than the optimal. Both raw and normalized feedback were given to the participant, while the normalized values were used in the optimization parts of the exercise.

Objective Function Assumptions		
Objective	**Metric**	**Evaluation Method**
Shade	Area shaded by canopy (max, ft^2)	Geometric measurement; assuming 50^o sun angle
Connections	Number of connections (min)	Sum of intersections between cables and canopy/wall
Carbon Emissions	Emissions due to material volume (max, kg CO_2)	FEM + Sizer
Deflection	Maximum deflection in model (min, in)	FEM

Table 4: Quantitative objective functions for the design experiment

4.2.4 Materials

Each participant was provided with a parametric model of the conceptual design space created by the author. The model was developed in Rhinoceros and Grasshopper. The model geometry and the evaluation metrics for shade and connections were built using native components. The structural model was created using Karamba, a plug-in for Grasshopper that interacts with Rhinoceros geometry. A vertical distributed load of 0.60 kN/m was applied to each beam, which corresponds to the tributary area of multiple beams arrayed longitudinally along the wall of the restaurant. To estimate the overall weight of structural material, Karamba's sizing feature was used. This feature checks the allowable axial, bending, and buckling loads for each member and then searches through a structural section library to determine the smallest member that can adequately handle each load before outputting the total weight of the structural system. These material quantities were then multiplied by carbon coefficients calculated in the Inventory of Carbon & Energy (ICE) (Hammond & Jones 2010) database to arrive at embodied emissions.

Although a more detailed analysis involving uplift and other forces could be completed on the structure, this study focuses on conceptual design, and thus the evaluations are greatly simplified so that designs can be explored instantaneously and compared based on relative performance. The purpose was to give participants rapid feedback to build intuition about the design space and ultimately guide exploration while maintaining creative flow. For this reason, square footage and deflection were given in Imperial units, since these units might be understood more quickly by students with experience in U.S. universities and design offices. Units for emissions were given in kg of CO_2, which is consistent with the ICE Database. In addition to the model shown in Figure 8, a 3D version was created to allow for creative exploration of curvature as part of the educational exercise. However, the results of this chapter only include the 2D model.

In phase 1 of the exercise, only the geometry of the design was shown to participants. In phase 2, this same geometry was shown alongside the performance feedback information (Figure 8). In phase 3, students were able to use either Galapagos, an evolutionary solver native to Grasshopper, or one of the derivative-free optimization algorithms contained within Goat, which is an alternative optimization plug-in. These optimization components were connected to a composite function combining each of the objective scores along with an adjustable importance weight for each objective. Although participants were encouraged to explore different weight combinations as part of the exercise, the instructions were to simply use the tools to produce satisfying designs, matching the task of the earlier two phases.

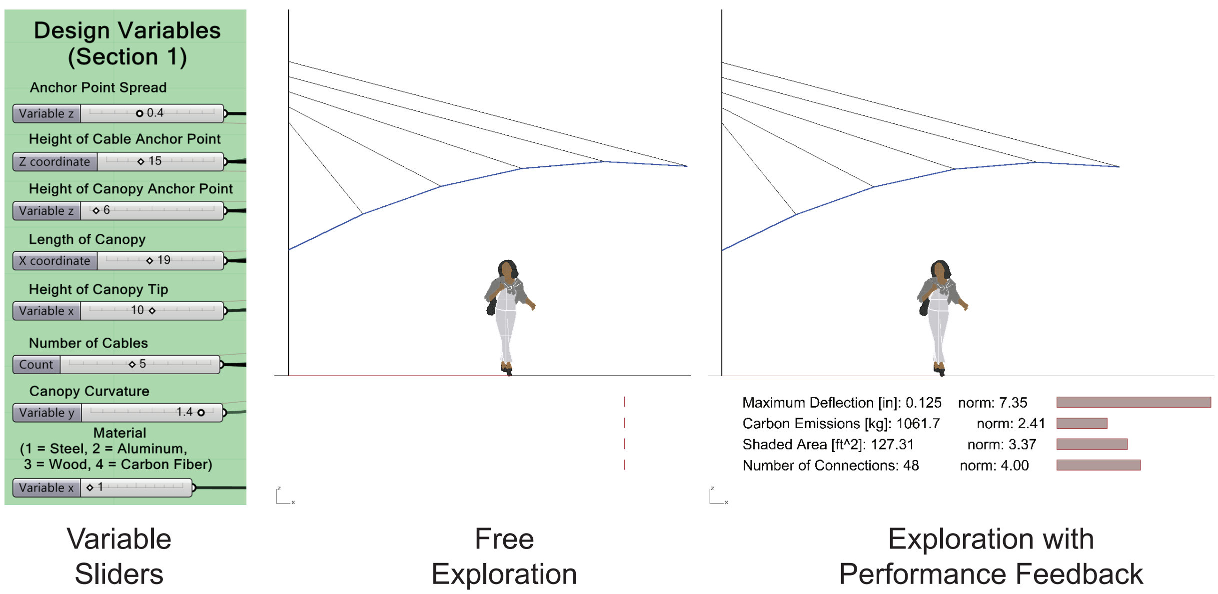

Figure 8: The main parametric model and design environment, showing phases 1 and 2

4.2.5 Design outcome metrics: quality and diversity

The results of this study include both the measured quality of the designs produced by the participants and an evaluation of their novelty and diversity. The quality of the designs refers to their relative performance in each of the four dimensions. Performance data collected throughout the study was separated by both phase and objective. Outliers that did not comply with the original bounds of the problem were removed from each dataset. A number of statistical tests were then completed on the datasets to compare overall performance averages and medians between the three different design environments.

To determine any significant effect the performance feedback and optimization tools had on design quality, a single factor ANOVA test was completed on the three datasets for each performance measurement. This analysis tests the null hypothesis that the averages of each set are statistically equal. If the ANOVA test determined that at least one of the datasets was significantly different, a series of two-tail t-tests were completed on the different sets to find out which of the performance categories were statistically different from the free exploration dataset. Each of these tests used an alpha of 0.05 for a 95% confidence level.

In addition to the quality of examples produced in a conceptual design exercise, researchers are also concerned with creativity, and there are established criteria for assessing this aspect of a particular design environment. These criteria include novelty and variety, which can both be obtained by defining what is not novel, or by identifying key attributes and functionalities, developing a hierarchical rating system, and scoring each design based on this system (Shah et al. 2003). It is also possible to mathematically compute the diversity of a dataset, which measures how different the designs in a set are from one another. This diversity measurement can be completed by computing certain geometric relationships between the different design vectors, such as the radius of the smallest enclosing ball, area or perimeter of a convex hull, or taking an average or total distance to the centroid of all design points in n-dimensional space, where n is the number of design variables.

Novelty, variety, and diversity are all important concepts in this exercise due to the creative requirements of architectural and structural design processes. If a design tool produces high-performing designs but constrains results to visually uninteresting solutions or only a small portion of the design space, it is likely not useful to expressive designers desiring the ability to exert preference. As such, computational methods that provide performance-driven design assistance must be flexible enough to yield diverse design spaces. However, the most common metrics for novelty and variety both require researchers to dictate a scoring system for the design problem, which is unsuitable for the open-ended nature of this study. The author ran a number of different diversity calculations on each dataset, but the metrics did not agree in their ranking and comparison between the different design environments. Consequently, visualizations of the design vectors and example geometries generated for each phase will be presented directly in the results section of this chapter, along with commentary on their meaning.

4.2.6 Participant survey

Upon finishing the exercise, participants were asked to complete a short survey related to their experience. The survey asked participants to rank the design environments (free exploration, performance feedback, optimization tools) on a number of factors: which environment is the easiest to work with, leads to the best design outcome, they would most like to mention to a client or architectural reviewer, and they are most likely to apply in their own design workflows. Analysis of these responses supplements the measurement of design quality and discussion of design diversity.

4.3 Results and discussion

4.3.1 Design quality

The aggregate performance scores for each experiment are given in the box plots in Figure 9. The plots show the median, interquartile range, upper and lower whiskers, and outliers for each dataset. A note is also included if a dataset for performance feedback or optimization tools differed significantly from the dataset for free exploration.

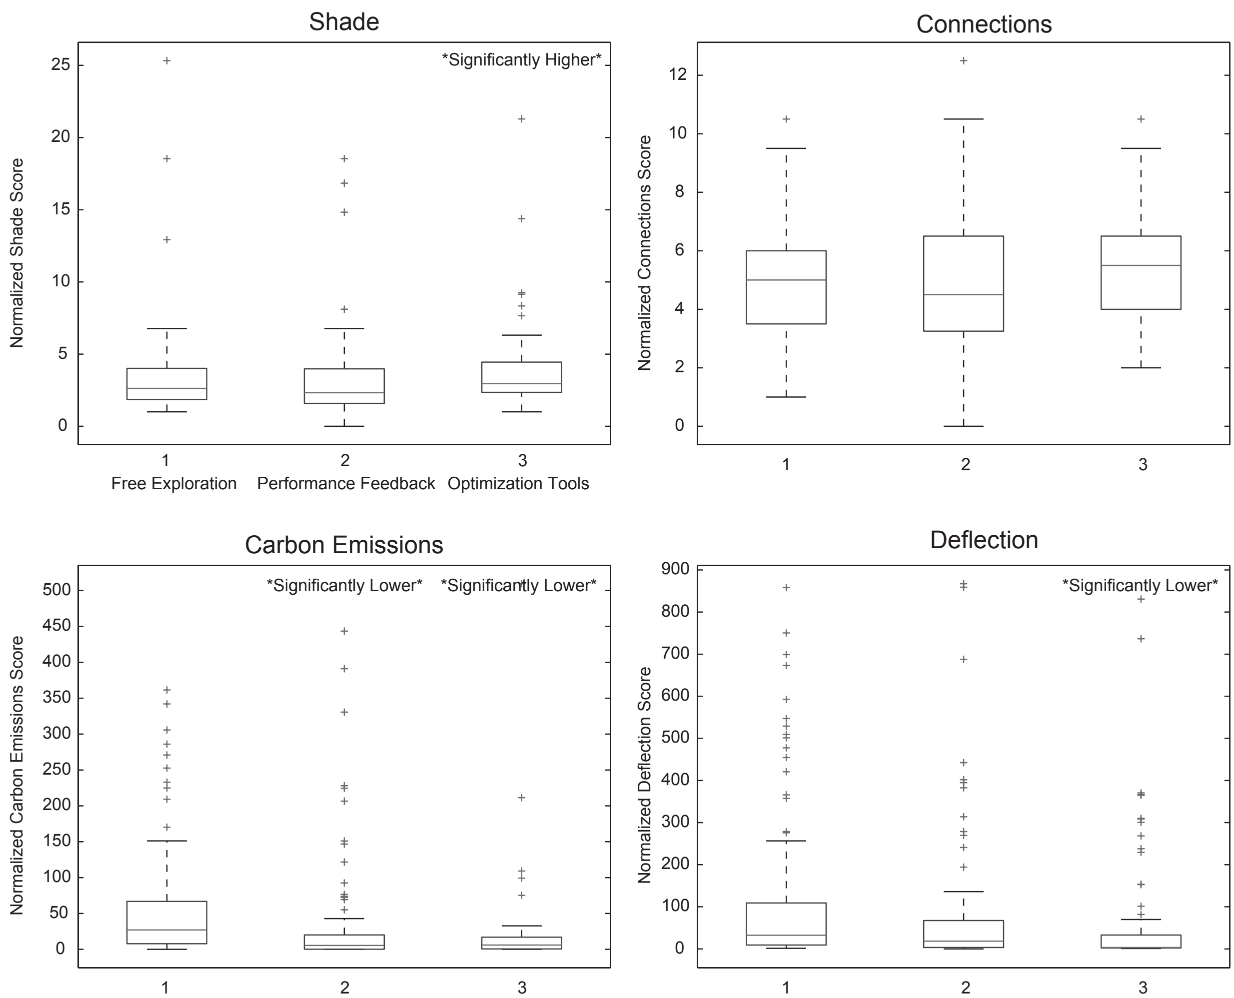

Figure 9: Box plots showing performance scores of designs for different environments

In the shade performance metric, only the optimization tools dataset showed any significant difference, with designs in this set performing substantially worse than in free exploration. There was no statistical difference in the number of connections measured between the design environments. Both performance feedback and optimization tools significantly reduced the estimated embodied carbon in selected designs compared to free exploration. While the median deflection values for both feedback and optimization were lower than in free exploration, only the optimization dataset registered as significantly lower.

The lack of improvement in shaded length and number of connections can be partially attributed to the fact that these objectives are easy to understand visually regardless of design environment. In addition, these performance metrics seem not to have been prioritized as highly as the structural metrics across all three phases. This effect could be a result of the study's educational setting, or of the differences in magnitude between normalized scores for lower performers—carbon and deflection could score hundreds or thousands of times worse than the optimal, while the design space and normalization technique only allowed for much smaller scores in shade and connections.

In contrast, the performance scores of embodied carbon and deflection improved significantly when designers were given additional data. The designs generated while using performance feedback and optimization tools resulted in 43% and 68% lower carbon emissions on average, respectively. Structural performance is less easy to predict intuitively unless designers have advanced education or experience in the field. The results of this study suggest that access to computational tools helps mitigate lack of structural intuition, while guiding participants towards better performing designs. Optimization tools magnified this effect, since participants seemed more willing to balance the four objectives during exploration with feedback than during optimization, which pushed priorities towards the structural criteria at the significant expense of shaded length. The optimization tools also reduced the number of very poorly performing outliers, mostly because optimization is unlikely to pick a design solution with large scores in any category unless that category is completely ignored in the composite function. This effect is noticeable in Figure 10, which shows that for carbon and deflection in phases 2 and 3, almost every participant design is in the lowest bin of the histogram.

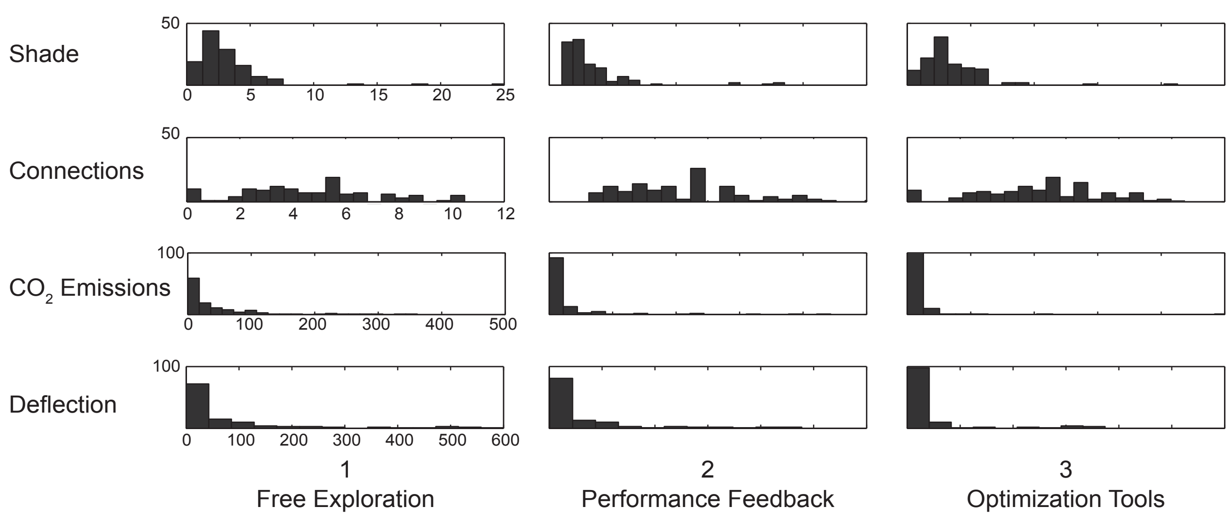

Figure 10: Performance histograms showing the tendency towards eliminating structural outliers

4.3.2 Design diversity

All of the design vectors produced during the exercise are displayed in the parallel coordinate plots in Figure 11. Material selection, which is one of the most important variables relating to structural performance, is noted for each design. A representative sample of designs produced in each environment is also shown geometrically in Figure 12. Each cleaned dataset contained over 110 designs, making it impossible to show every single result. In all three cases, the vector plots leave few gaps, illustrating that participants covered a large portion of the design space and did not produce designs that can be categorized simply. Nevertheless, a number of trends are visible when viewing both the design vector plots and representative examples.

One noticeable difference between environments is that fewer carbon fiber and aluminum solutions are present when either feedback or optimization is utilized. Both of these materials are considered high-performance by many in the architectural community due to their strength and weight, leading to numerous generated solutions during free exploration. However, these materials have energy-intensive manufacturing practices and relatively high carbon coefficients, which becomes evident once performance feedback is enabled. The vector plots also show that the optimization tools favored designs with less extreme values for number of cables and curvature of the canopy, which corresponds to subtle curves and more than one or two cables.

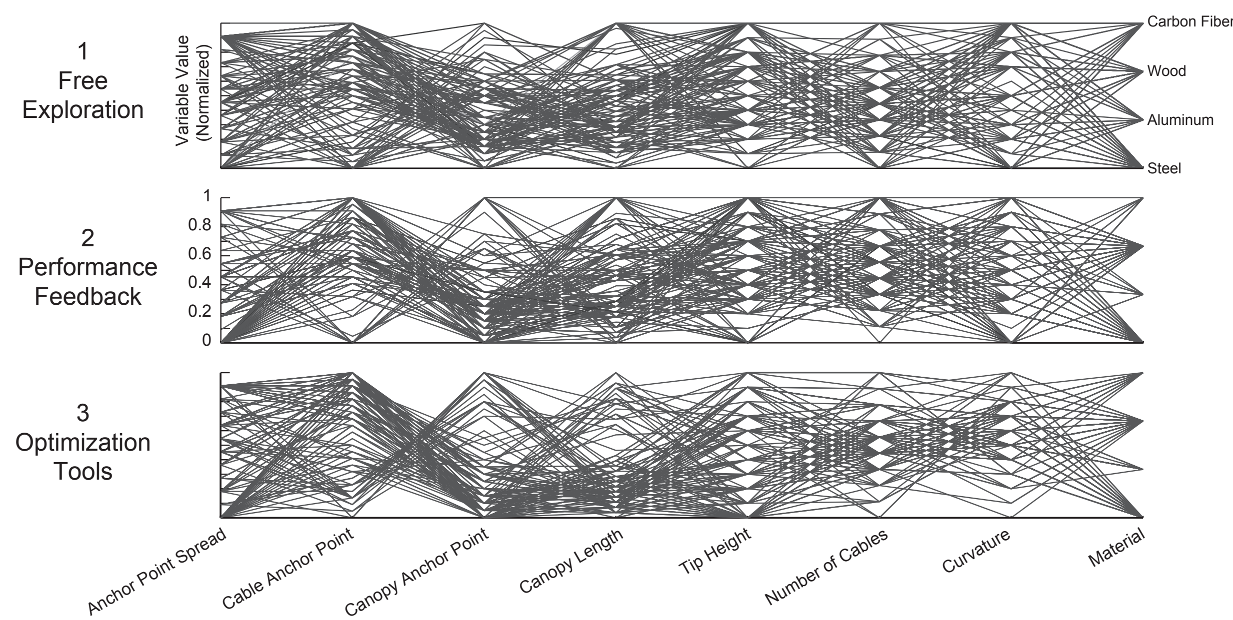

Figure 11: Parallel coordinate plots of the design vectors for each collected dataset

In addition, the optimization tools pushed designs to have lower tip heights, which improves shaded area without a large negative impact on other metrics. This effect is somewhat noticeable for the performance feedback phase, but it is especially obvious when comparing the prevalence of designs with the lowest tip height in free exploration versus optimization tools. Similarly, structural performance feedback and optimization tools encouraged designers to produce geometries with low canopy anchor points and high cable points, as indicated by the high density of lines moving diagonally between the two extreme values. Interaction with performance data also drove canopy anchor heights much lower in general. A number of creative participants realized that setting a canopy anchor point above the cables turned them into compressive struts and yielded a different typology, but these designs required considerably more material to resist buckling and also generally paid a penalty on shaded area due to high tip heights, thus making them undesirable in performance-based design. Despite these noticeable differences, there was still a high degree of geometric diversity present in each design environment. Although performance feedback and optimization both discouraged some areas of the parametric design space, the large variations demonstrated in Figure 12 suggest that designers still had sufficient flexibility to exert preference. The survey questions in the next section were meant to further assess this assertion.

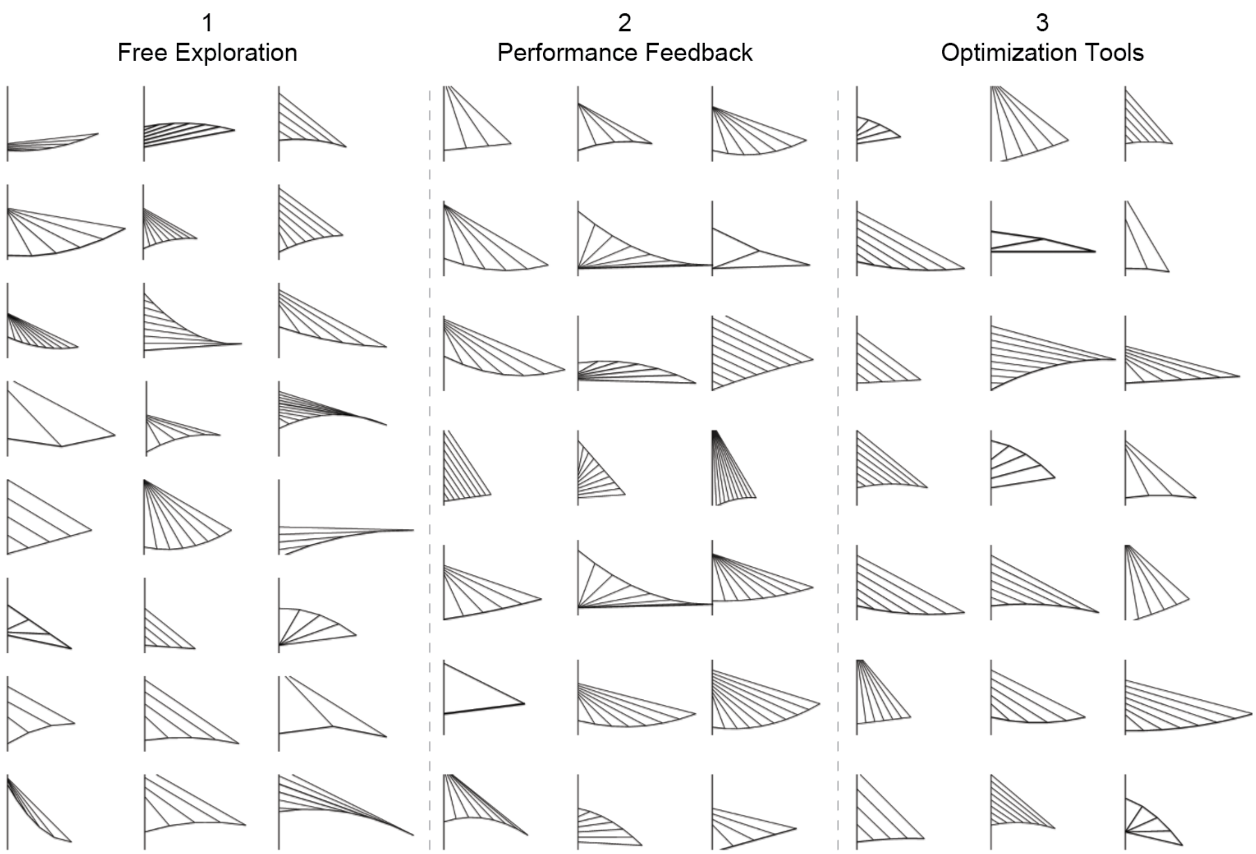

Figure 12: Representative designs generated during each phase of the study

4.3.3 Survey responses

After completing the exercise, students were asked a series of optional questions related to the three different design environments. The results of their answers are given in Table 5. None of the questions led to an overwhelming winner in any of the categories. Surprisingly, the single most common response was that optimization tools were the easiest to manage, perhaps because participants felt that the algorithms were doing part of the work of eliminating bad outcomes. Although participants seemed to feel on average that performance data does improve design quality, respondents were perfectly split on which technique would give them the most confidence when explaining their design to a client or critic. Overall, it is clear that different designers have different preferences, and flexibility is often key when setting up a design environment.

Participant Survey Results			
Question	**Free Exploration**	**Performance Feedback**	**Optimization Tools**
Easiest to work with:	30%	15%	55%
Leads to best design outcome:	24%	38%	38%
Would most like to mention to a client or reviewer:	33%	33%	33%
Most likely to apply in their own workflow:	24%	52%	24%

Table 5: A comparison of participant preferences for the three different design environments

4.4 Summary of contributions

In conclusion, this chapter presents new data concerning the effects of performance feedback and guidance on the conceptual design of structures through an experimental case study. The study, composed of students with formal design training, gave participants the task of designing a cable-stayed canopy roof, along with a parametric model for design exploration and progressive access to performance-based computational tools. The design exercise contained three phases, which each asked the students to record their favorite designs. The first phase was free design exploration without any performance feedback. Next, the students were asked to again find their favorite designs, but with the maximum deflection, embodied carbon, shaded area, and number of connections displayed on screen. Finally, the students were given access to optimization tools that could guide exploration in an automated way. Students were then asked to assess their experience in each of the different environments.

Although there was no consensus between the students on aesthetic preferences, prioritization of design variables and objectives, or favorite design workflow, designs chosen using either real-time performance feedback or a directed optimization process performed significantly better in terms of deflection and emissions than designs chosen through free exploration. The average effect of feedback and guidance on the shaded area and connections metrics, which can be more easily determined visually, was not significant except in the case of shaded area, which was made worse when students used optimization. In terms of design diversity, performance feedback and optimization tools virtually eliminated some poor-performing areas of the design space, while encouraging a number of specific design characteristics, such as low canopy tip heights and a large spread between cable and canopy anchors. However, there was still noticeable variation and diversity within the designs produced using performance-based methods, which indicates that flexibility and creativity were still possible even as considerations of performance clearly guided participants.

Overall, this research determined that the use of performance data in conceptual design can have a noticeably positive effect, a finding that encourages the further development of performance-based computational tools. Future study in this area could include additional case studies, a higher number of participants, and more controlled settings. However, this research is an initial quantitative step in testing and supporting the arguments commonly made for greater integration of rapid, multi-dimensional performance feedback into architectural and structural conceptual design workflows.

5 Structure + Energy in Long Span Building Design

In the conceptual design of buildings, many traditional optimization methods have seen only limited application, despite the emphasis contemporary designers place on building performance. This is largely due to the complex requirements of contemporary architecture, and the fact that human intuition and judgment are still central to the design process. Even as early as the conceptual design phase, architects must simultaneously consider and prioritize a multitude of interrelated design objectives, and while an increasing number of these objectives are quantitatively measureable, many are not. Two of the most important objectives related to building performance are the embodied and operational energy used in a building's materials and operations, respectively. Often, the goals of reducing each of these quantities trade off with one another, as well as with other qualitative design goals, in unexpected ways. This is especially true in building types with substantial structural requirements such as the long span roof, which is the focus of this chapter.

5.1 Background: embodied and operational energy

As mentioned in Chapter 1, a high-performance, sustainable building has been identified as one that minimizes energy consumption throughout the four main stages of a building's lifetime: materials manufacturing, construction, use and maintenance, and end of life. Both the large amount of energy consumed by buildings and the subsequent need for a reduction in this consumption and subsequent carbon emissions have been well documented. In light of this, a conceptual designer could simply convert every aspect of the design to a unit of emissions and run a traditional optimization to find a single solution. However, in practice this would hamper the ability of designers to express preference, and also ignore financial constraints and other architectural complications that influence the development of a real building. As such, architects often find a pure performance optimization approach reductive, overly simplified, and deterministic, which can lead to resistance towards the adoption of optimization methods in design (ACADIA 2015).

Furthermore, determining the correct structural form for long span roofs and other large structures is especially important because structural efficiency depends more on the geometry of a building than on material, sizing, and other building characteristics developed in later stages, and because structural material makes up a more sizable portion of the overall embodied energy. For example, De Wolf (2014) found that the amount of embodied energy within a building's 50-year lifecycle usage can range from 4-22% of the total. When considering carbon emissions instead of energy consumption, the portion of embodied can rise to as high as 80% of total emissions, depending on which exact source is consulted. As shown by Kaethner & Burridge (2012), the material contained in the substructure and superstructure can be responsible for over half of embodied emissions.

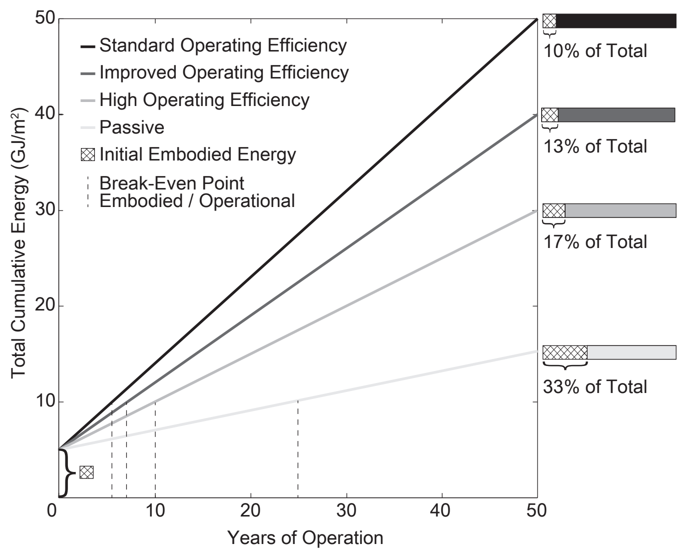

Figure 13: The contributions of embodied and operational energy to the cumulative total energy usage of a typical building. Values from (Sartori & Hestnes 2007)

With increased emphasis on cutting operational energy usage while pushing towards net zero buildings, the embodied energy of future buildings will make up an increasingly larger portion of total energy usage over their lifetimes (Figure 13) (Sartori & Hestnes 2007). Consequently, this chapter focuses on multi-objective optimization with structural efficiency and operational energy efficiency as the two measurable objectives, since a MOO approach gives the designer flexibility but encompasses the most significant quantitative performance goals of contemporary architecture.

5.2 Applications and limitations of multi-objective optimization

Although MOO has demonstrated a greater potential than traditional optimization to assist conceptual designers in generating and deciding between high-performing, early-stage designs, it too has seen only limited use in practice. This lack of application can be attributed to the complicated model translation that must occur between design and analysis software, the often linear process in which members of a design team are only given small latitude to 'optimize' for their own performance goals without reference to other disciplines, and the difficulty of using optimization within a process that includes subjective preferences and design goals that are difficult to formulate numerically. Related fields such as aerospace, mechanical, and other pure engineering design disciplines have been more successful than architecture in overcoming some of these obstacles. The differences in scale, production, and customization of buildings when compared with airplanes or cars have contributed to a building industry that is more fragmented and mostly unexposed to optimization workflows. For a conceptual MOO procedure to become popular with building designers looking for original, expressive forms, researchers must overcome difficulties of non-quantitative objectives and disconnected disciplines while showing how an integrated process can lead to a diverse range of design outcomes that meet a variety of aesthetic preferences.

Many academic researchers have addressed the limitations of multi-optimization for use in conceptual design, but few have studied the strong link between architectural form and different performance metrics simultaneously in a way that demonstrates the significant potential of MOO to influence the leading edge of architectural design. Many studies have been restricted to geometries that are primarily made of rectangular boxes, which are easy to model in terms of energy usage. However, even within the typology of the long span roof, there are a wide variety of architectural forms and corresponding building shapes and structural systems that could be optimized for performance. In addition, researchers have been largely unable to define a clear way for architects to interact with MOO data, which can include performance feedback from multiple engineering disciplines in different units and scales, in a way that leads to good design decisions. In order to have greater impact on innovative and creative architectural practices, it is important to develop methodologies that effectively navigate meaningful tradeoffs and produce design examples that are applicable to a wider range of building geometries.

In response, this chapter demonstrates how MOO can be used to generate geometrically diverse architectural design solutions in different climate regions through three case studies of buildings with long span roofs. Since selecting the right building shape and form in the early stages has a large effect on the overall success of a building with demanding structural and spatial requirements, the case studies focus on these large-scale conceptual design decisions.

The optimization method in this chapter uses simulation to focus simultaneously on the embodied energy found in structural material and the operational energy of the building. The case study results are presented in terms of overall energy requirements, but the embodied and operational components are kept independent since the two are not always equal when time, financing, and other practical realities of construction are taken into account. These results demonstrate the utility of separating structural efficiency (primarily upfront emissions and cost) and operational energy efficiency (emissions and cost over time) in optimization for conceptual design, as well as showing the effect of the two separate objectives on architectural form. Overall, this chapter illustrates how the application of MOO can yield a wide range of expressive, high-performance designs, provides increased awareness of how architects might respond to particular climates while designing long span buildings, and contextualizes structural efficiency within the broad goal of sustainable design.

5.2.1 Optimization for structure or energy

This research builds on a wide body of existing scholarship concerning the integration of visual criteria into optimization algorithms, geometry optimization for building performance, and multi-objective optimization in architectural design. A brief overview of major contributions is given here, beginning with research that focuses exclusively on either structure or energy usage. As mentioned in Chapter 2, Coley & Schukat (2002) optimize lighting and space conditioning systems to minimize overall energy usage and set up a system to present optimal designs for aesthetic evaluation. Caldas & Norford (2003) find that energy optimization problems can be used to generate designs that minimize capital and operating costs, emissions, or maximize thermal comfort. Wang et al. (2005) apply a multi-objective genetic algorithm while considering lifecycle cost and lifecycle energy usage as different objectives. Asl et al. (2014) establish and test an optimization method for whole building energy performance and daylighting. Magnier & Haghighat (2010) use surrogate modeling and genetic algorithms to optimize a building design for thermal comfort and operating energy.

5.2.2 MOO with global building geometry variables

Other researchers have applied MOO to building design with significant geometric design variables. Marks (1997) explores shape optimization for minimizing initial building cost and annual heating cost for a variety of building plans. Khajehpour & Grierson (2002) optimize commercial buildings for minimum capital and operating cost and maximum income, as well as for profitability and safety Khajehpour & Grierson (2003). Suga et al. (2010) generate Pareto-optimal solution sets for window design optimization with energy consumption, cost, uniformity, and draft performance as objectives. Tuhus-Dubrow & Krarti (2010) use a genetic algorithm to optimize for lifecycle cost in a conceptual building, but they also consider the relationship between operational and embodied energy when compared across climates. Diakaki et al. (2008) propose a multi-objective approach to optimizing the window type, wall insulation material, and wall insulation thickness of a building for costs and savings. Fialho et al. (2011) investigate the effects of varying building materials and orientations on construction costs and energy efficiency using a multi-objective optimization algorithm.

5.2.3 MOO for structure and energy

Additional existing research proposes workflows that allow designers to explore different structural shapes while also optimizing for energy or other objectives. Some of this research was also mentioned in Chapter 3, since it has general relevance to multi-objective optimization in architecture. Von Buelow (2012) aims to combine performance feedback from a variety of simulation programs (structural, lighting, acoustical) with aesthetic preferences in design through the use of interactive evolutionary algorithms. By breeding designs that have different fitness functions and storing multiple alternatives in a searchable, graphical database, von Buelow's proposed workflow gives designers the opportunity to pursue multiple design objectives and visualize tradeoffs between them. Flager, Welle, et al. (2009) propose a framework for automating the prototyping, simulation, and analysis tools required to run an optimization, and then use this framework to optimize a classroom design for structural cost and lifecycle energy cost. Krem et al. (2013) demonstrate how the position of a structural core and the shape of the floor plan in a high rise building affects its energy and structural performance differently depending on the climate and location of the building. Quaglia et al. (2014) use parametric modeling and performance simulations to generate an optimal solution set for the design of origami-inspired, rapidly deployable shelters. Méndez Echenagucia (2013) presents shape optimizations for a number of competing performance objectives, including acoustics, structural efficiency, and energy efficiency.

5.3 Unmet needs

This chapter adds to the existing knowledge of multi-objective optimization for structure and energy in a number of ways.

First, it employs MOO on structural typologies and building shapes that are not primarily rectangular, demonstrating that the technique, which involves both structural finite element and energy surface models, can be used on more expressive architectural designs.

While testing generalizable, known architectural problems, this research also generates diverse results that may be structurally, geometrically, and architecturally interesting. In different climates, the optimization results are often surprising or unexpected.

Finally, this chapter compares structure-energy tradeoffs in absolute energy units, which contextualizes structural efficiency and presents a fuller picture of how structural optimization can influence architectural forms even while emphasizing overall design sustainability.

5.4 Methodology

To demonstrate the capabilities of MOO in the conceptual design of long span structures, three case study models were developed and analyzed. The procedure for testing the case studies included the development of a parametric design space, using a MOO algorithm to find the Pareto front of best performing designs, visualizing the results in a way that designers can easily compare between samples, and analyzing the overall data across different forms and contexts (Figure 14). The design lessons learned from the combination of these three case studies could be applied to a variety of architectural programs that contain large, open, conditioned volumes, such as airports, convention centers, or stadiums.

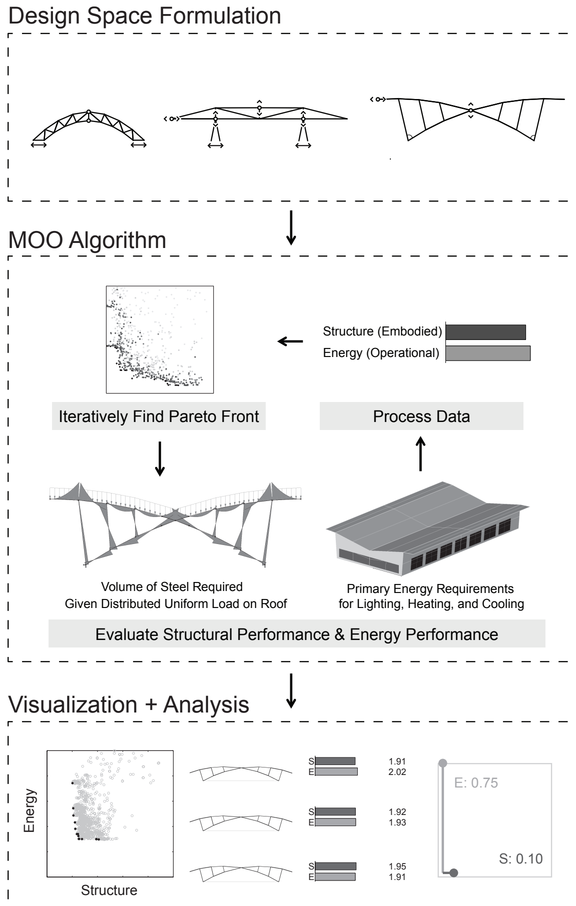

Figure 14: Flowchart describing the computational methodology for generating Pareto optimal solution sets of designs for each case study

5.4.1 Problem selection

While demonstrating the potential of a multi-objective optimization framework for conceptual design, this chapter seeks to explore two generalizable architectural problems that trade off between structure and operational energy. The first problem is the enclosed arch, where a higher arch rise can lead to greater structural efficiency, but also results in greater interior volume and envelope surface area. The significance of this architectural problem is demonstrated by a spirited debate about the Montreal Olympic Stadium (Figure 15) between the engineer Anton Tedesko and architect-engineer Frank Moffet in a 1976 issue of *Civil Engineering*, the journal of the American Society of Civil Engineers (Tedesko 1976). Tedesko, with support in later letters from Professor David Billington of Princeton University, argued that the architecturally arbitrary form of subtly sloped cantilevers required massive section sizes and wasted material, while Moffett commented that the lower slope results in operational energy savings (Sommerschield et al. 1977). Although the argument was never resolved, contemporary optimization techniques and performance simulations can address this tradeoff quantitatively.

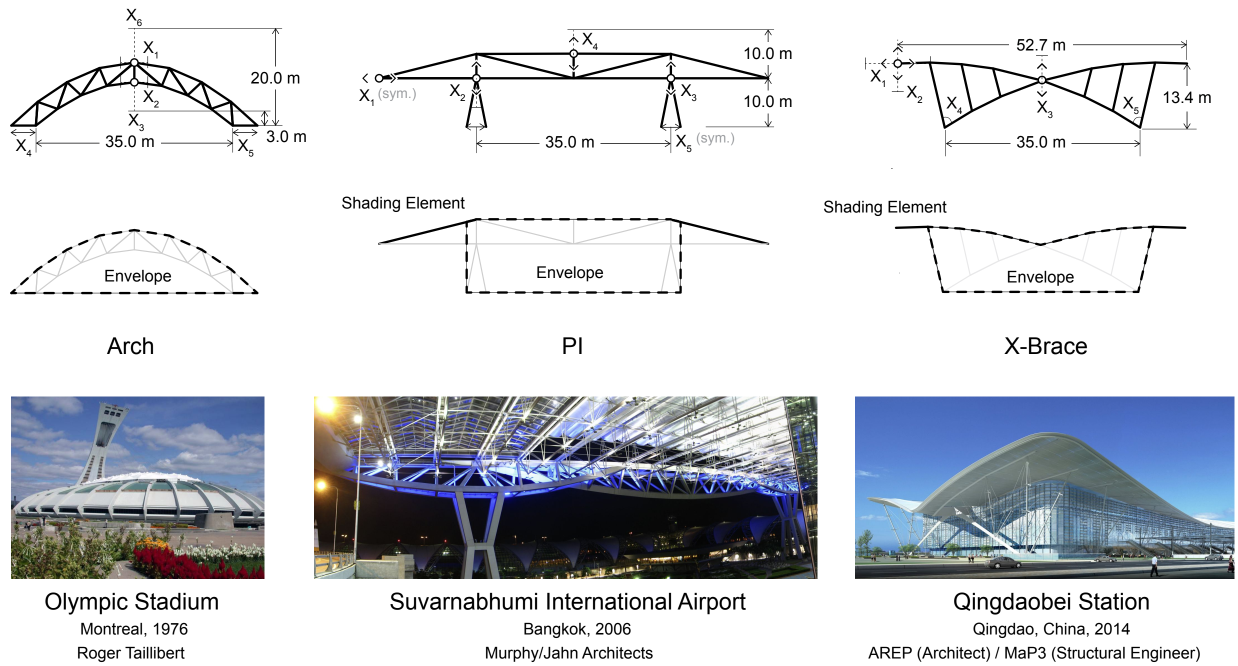

Figure 15: Geometric diagrams of the design space and constructed buildings that correspond to each architectural problem. Images courtesy of (Wikipedia_User:Tolivero 2006), (Wikipedia User: Wikigod 2009), (MaP3 2014)

The second architectural problem is the cantilevered overhang, which can allow for efficient shading strategies in certain climates and orientations, but may require more structural material to build. Roof overhangs are common architectural elements, and have been used as a passive cooling strategy since ancient vernacular dwellings (Zhai & Previtali 2010). Overhangs have many different functions, including protecting an open window from wind and rain or providing a covered entrance in the case of a transit terminal, but their influence on building energy usage can be substantial: a case study by Raeissi & Taheri (1998) shows that an optimized static overhang design lead to a 12.7% reduction of cooling loads with only a marginal increase in heating loads. However, when overhangs are formed by a large cantilevered structure, such as in the Suvarnabhumi Airport, they can substantially affect the amount of structural material required by a design, creating a tradeoff similar to the enclosed arch. These two architectural problems can also be explored simultaneously, as in the example of the Qingdaobei Station in China, a bus terminal completed in 2014 (MaP3 2014).

5.4.2 Design space formulation

The performance implications of the enclosed arch and the cantilevered overhang are explored through three different long span building case studies. The first design is an enclosed, trussed arch, which can vary in terms of truss depth, overall height, and two skew parameters. This design has no large windows, as would be the case for a hangar, warehouse, roller skating rink, or possibly the main span of a sports arena. As a result, the optimization isolates the pure arch rise tradeoff between structure and energy without interference from daylighting effects. Although the geometry of the case study is an extruded arch, a revolved arch or dome, such as the Montreal Olympic Stadium, would show the same tradeoff. The second design is a 'PI' structure, which contains two columns and a spanning truss that cantilevers past the columns on both sides, resembling the Greek letter. Although the parameterization of this design creates a few opportunities for other small performance tradeoffs, it primarily studies how designers might choose an appropriate overhang configuration for optimal daylighting performance. The last case study is the 'x-brace', which consists of a three-hinged arch supporting cantilevering roof beams through a series of vertical struts, inspired by the Qingdaobei Station. The x-brace geometry studies both the arch and overhang tradeoffs simultaneously.

In each of these designs, the main two-dimensional structural system is arrayed longitudinally, with transverse beams spanning in between the arches, PIs, and x-braces to create the enclosed volume. Since all of the case studies have a long main span, there is free interior floor space between the supports, allowing for easy occupant movement and a variety of interior programs. For each of the case studies, the total floor area remains constant at 3,500 m^2. The decision to parameterize different structural and envelope geometries while treating floor area as a given parameter realistically represents the typical architectural process, in which the programmatic area requirements would drive the design. This value, rather than conditioned volume or another metric, would be most important to both the client and architect. As such, the energy usage metrics given in the results section are normalized by floor area.

In the PI and x-brace cases, the enclosure consists of a panelized curtain wall system that mixes regularly spaced transparent glazed panels with spandrel glass panels concealing interior insulation. This configuration captures the effects of geometric changes on both natural daylighting and heat transfer through the building's envelope. The arch enclosure contains no glazing and is entirely insulated. A parametric model of each design was generated in Rhinoceros and Grasshopper based on a number of global geometric parameters, which are given in Table 6. To prepare for evaluation of both structural and energy performance, the model included linear elements representing steel structural members, as well as surfaces representing the exterior envelope of the building. An abstraction of the structural steel geometry is used to represent the overall building shape in the results section; the enclosure boundary approximately follows the line of the main structural elements on the outside of each truss, excluding the overhangs.

Design Space Variables		
Arch		
Symbol	Variable	Range
x_1	Top Skew	-2.5 m < x < 2.5 m
x_2	Bottom Skew	-2.5 m < x < 2.5 m
x_3	Height of Lower Chord	3 m < x < 10 m
x_4	Left Width at Support	0 m < x < 6 m
x_5	Right Width at Support	0 m < x < 6 m
x_6	Truss Depth	0 m < x < 10 m
PI		
x_1	Overhang Length	1 m < x < 15 m
x_2	Left Column Height	4 m < x < 10 m
x_3	Right Column Height	4 m < x < 10 m
x_4	Truss Height	0 m < x < 10 m
x_5	Column Width	0.25 m < x < 1 m
X-Brace		
x_1	Starting Edge Height	12 m < x < 21.3 m
x_2	Starting Overhang Length	7.6 m < x < 12.2 m
x_3	Starting Hinge Height	7.6 m < x < 15.2 m
x_4	Angle Arch A	-0.150 rad < x < 0.150 rad
x_5	Angle Arch B	-0.150 rad < x < 0.150 rad

Table 6: Design space variables for each case study

Four different model locations were selected for simulation to test how optimal building forms change for different climates. These four locations represent the main world climate zones as described in Köppen-Geiger classification: cool, temperate, arid, and tropical (Kottek et al. 2006). The cool location, Boston, has both heating and cooling loads but is heating dominated. The second location, Sydney, is in a temperate climate zone where the required cooling load is slightly higher than in Boston, but the heating load is very small. The arid location is Abu Dhabi, which has no heating load but is dominated by cooling loads that are roughly five times higher than those required by Boston and Sydney. The selected tropical location is Singapore, which is also dominated by high temperatures year round. However, the average high temperatures are less extreme than in Abu Dhabi and do not swing as much either monthly or daily. A summary of these climate characteristics is given in Table 7. Each model was simulated with a west-east orientation of the longer axis, allowing for maximum daylighting effects on the north and south sides of the building where the majority of windows are located. When placed in different climates, the structural performance does not change—however, the way in which structure trades off with operational energy is different, as is the best energy form for each location.

Climate Characteristics				
	Abu Dhabi	**Boston**	**Singapore**	**Sydney**
Köppen-Geiger Classification	BWh	DFa	Af	CFa
IECC Climate Zone	1	5	1	3
Latitude	24.43 N	42.37 N	1.37 N	33.95 S
Longitude	54.65 E	71.03 W	103.98 E	151.18 E
Elevation	7 m	6 m	15 m	6 m
HDD18	54	3,134	0	751
CDD10	6,417	1,609	6,664	2,922
Heating Design Temp (ºC)	13	-14	23	6
Cooling Design Temp (ºC) (Dry Bulb)	43	31	32	29

Table 7: Climate characteristics for each case study location. Values from (ASHRAE 2004)

5.4.3 Performance evaluation

The performance of each design was measured using Grasshopper plug-ins that enable simulations using Rhinoceros geometry as an input. The structural optimization objective was to minimize the amount of steel required, and the energy optimization objective was to minimize the annual operational energy of the building, which includes requirements for lighting, heating, and cooling. To convert the amount of steel into embodied energy, the mass of each design was multiplied by an embodied energy coefficient. Although the model generates idealized operational energy loads already in the correct energy units, these loads were converted to primary energy requirements using assumed equipment and transmission efficiencies. The conversion processes for both embodied and operational energy are explained in greater detail in a later section.

Structural performance was quantified using Karamba, a finite element analysis tool embedded in Grasshopper. Dead load, symmetrical live load, and asymmetrical live loads on each side of the structure were all considered as part of combination structural load cases. To calculate the overall weight of structural material, Karamba's sizing optimization feature was used. This feature checks the allowable axial, bending, and buckling loads for each member and then searches through a structural section library to determine the smallest member that can adequately handle each load before outputting the total weight of the structural system. Lateral wind loads were neglected, since they did not substantially affect the sizing already governed in most cases by the asymmetrical live loads.

Structural Model Assumptions	
Steel Beam Elements	
X-Brace and PI: Wide Flange Beam Sections	
Arch: Round Hollow Steel Sections	
Dead Load	7.18 kPa
Live Load	3.48 kPa
Load Cases	Dead Load Dead Load + Full Live Load Dead Load + Live Load applied asymmetrically to each side

Table 8: **Modeling assumptions for the structural models**

The energy evaluation was performed using Archsim (Dogan), a Grasshopper plug-in that connects Rhinoceros geometry with EnergyPlus, a widely used energy analysis and thermal load simulation program made available by the U.S. Department of Energy. Archsim assigns material properties to each surface of the model, creates thermal zones, and runs a whole-year simulation, returning the total amount of energy required to keep the building appropriately lit and temperature-controlled for an entire year. This model accounts for both solar and temperature effects. The envelope assumptions were created based on an International Energy Conservation Code (IECC) compliant building in each of the separate regions (International Code Council 2012). Other model settings were based on ASHRAE standards, including required heat recovery effectiveness, mechanical ventilation rate, and lighting power density. Although ASHRAE 62.1 specifies a minimum combined outdoor air rate of 4.1 L/s/person for "transportation waiting", airports and related buildings could contain spaces with other uses and occupant densities, such as "retail", which requires a combined 7.8 L/s/person (ASHRAE 2013). As such, a conservative assumption of 7 L/s/person was used for the case studies. To allow for direct comparison of architectural form across climates, assumptions about structural loads were held constant.

Energy Model Assumptions				
	Abu Dhabi	**Boston**	**Singapore**	**Sydney**
Roof R-Value	3.67 K*m^2/W	4.52 K*m^2/W	3.67 K*m^2/W	3.67 K*m^2/W
Wall R-Value	2.29 K*m^2/W	2.75 K*m^2/W	2.29 K*m^2/W	2.75 K*m^2/W
Window U-Value	2.84 W/K*m^2	2.16 W/K*m^2	2.84 W/K*m^2	2.61 W/K*m^2
Window SHGC	0.25	0.40	0.25	0.25
Window-to-Wall Ratio	0% for Arch / 30% for PI and X-brace			
Schedule	Equipment, lighting, ventilation all on (Airport Usage)			
Heating Set Point	20^o C			
Cooling Set Point	26^o C			
Mechanical Ventilation	7 L/s/person			
Heat Recovery	Sensible, 50% recovery effectiveness			
Infiltration	0.1 ACH			
Occupancy	0.2 p/m^2			
Equipment	12 W/m^2			
Lighting Power Density	12 W/m^2			
Daylighting	Continuous dimming, 500 Lux Target			

Table 9: Modeling assumptions for the energy models. Values from (ASHRAE 2004), (International Code Council 2012), (ASHRAE 2013)

5.4.4 Optimization method

In multi-objective optimization, there are a number of potential methods for selecting the best design based on designer preference. These include using *a priori* articulation of preferences, where a composite weighted objective function is created at the outset, *a posteriori* optimization articulation of preferences, where the designer chooses from a number of equivalently optimal designs after an optimization has been run, and interactive articulation of preferences, where a designer gives input while the algorithm is running (Marler & Arora 2004). This chapter contains sets of solutions that would be presented to a designer as part of an *a posteriori* process.

Specifically, the multi-objective genetic algorithm NSGA-II (Non-Dominated Sorting Genetic Algorithm II) (Deb et al. 2002) was used to iteratively approach the Pareto front over multiple generations of design alternatives. This algorithm starts with a population of design samples, evaluates their performance, and breeds the next generation through crossover and mutation between the highest performing designs. The NSGA-II also employs a diversity preservation mechanism to ensure a representative spread along the entire Pareto front, and uses the concept of elitism to speed up computation. Because the algorithm is stochastic, there is no guarantee that the exact Pareto solution will be found, but this method has been demonstrated to perform well on a wide range of engineering problems. The algorithm was implemented within the Grasshopper platform using the open source JMetal library (Durillo & Nebro 2011). This implementation was then applied to each of the parametric models to approximate the Pareto front of the objective space. For each case study, the algorithm used a generation size of 50 and was stopped at a maximum of 20 generations due to computational constraints. This generation size led to non-dominated sets of approximately 10-40 designs across the different case studies, which is a resolution that clearly shows the transition of building geometries along the Pareto front between the best structural and energy performers.

Figure 16 shows an example of this implementation of the NSGA-II by indicating successive generations of solutions with increasingly darker points. As the algorithm runs, it gradually moves in objective space from a random cloud of performances towards a set of Pareto optimal points. In addition to the NSGA-II optimization, the best structural solution for each case study was computed using the gradient-based algorithms available in Goat. The ability of the NSGA-II to locate a near structurally optimal design at one extreme of the Pareto front varied across the case studies, and thus the independently determined structural optimum is also presented in the results.

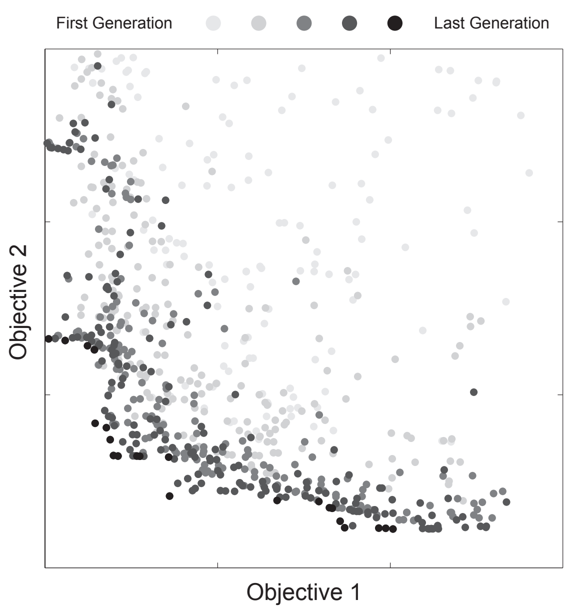

Figure 16: Demonstration of the NSGA-II multi-objective optimization algorithm progressively finding solutions closer to the actual Pareto front using crossover and mutation of high-performing designs in the previous generation

Although a range of theories exists concerning how to select the best design out of a given Pareto front, this research presents the entire fronts without a single chosen solution. It is assumed that in an architectural application, designers would be tasked with selecting a final design while relying on additional inputs, priorities, and preferences that exist outside of the optimization problem itself. Thus, the entire Pareto set represents a final result of the MOO process.

5.4.5 Data processing

For the purpose of direct comparison, the results of each performance simulation are normalized by the constant enclosed floor area, and thus presented in GJ per m^2 for both the embodied energy of the structure and the annual operational energy. To calculate the embodied energy of the structural system, the weight of the sized steel structure was multiplied by an embodied energy coefficient given by Hammond & Jones (2010). The operational energy simulation results are given in Joules, but require a conversion from idealized lighting, cooling, and heating loads to primary energy requirements. The efficiencies and coefficients of performance used in this chapter match those detailed in Bourgeois et al. (2006).

Conversion Assumptions	
Embodied Energy Coefficient for Steel	24.4 MJ / kg
Heating Assumptions	100% Site Efficient 33% Conversion Efficiency
Cooling Assumptions	Coefficient of Performance - 3 33% Conversion Efficiency
Lighting Assumptions	85% Site Efficient 10% Transmission Losses

Table 10: Assumptions for conversion to embodied and primary energy. Values from (Hammond & Jones 2010) and (Bourgeois et al. 2006)

5.5 Case study results

5.5.1 Shape and location of Pareto fronts

The performance and optimization results for all three case studies are presented in this section. Figure 17 gives the Pareto front found by the algorithm between the two objectives for each structural configuration and location. In each of these plots, a single design generated at any point during the algorithm is represented as a grey circle, while the Pareto optimal points are colored black. In a few cases, these plots display one or two additional Pareto points than were not returned in the last generation of the NSGA-II, but were included in the overall dataset and determined to be Pareto optimal during post-simulation analysis. Since results are normalized and given in GJ per m^2, all plots have the same aspect ratio and can be compared directly, although the locations of the axes change based on the specific climate and structural system being modeled. For conciseness, this results section will often refer to embodied energy simply as structure, and operational energy as energy. Thus, structural efficiency refers to performance on the horizontal axis, and energy efficiency refers to performance on the vertical axis. The general directions of left, right, top, and bottom within this established coordinate system will be used to explain performance throughout this section.

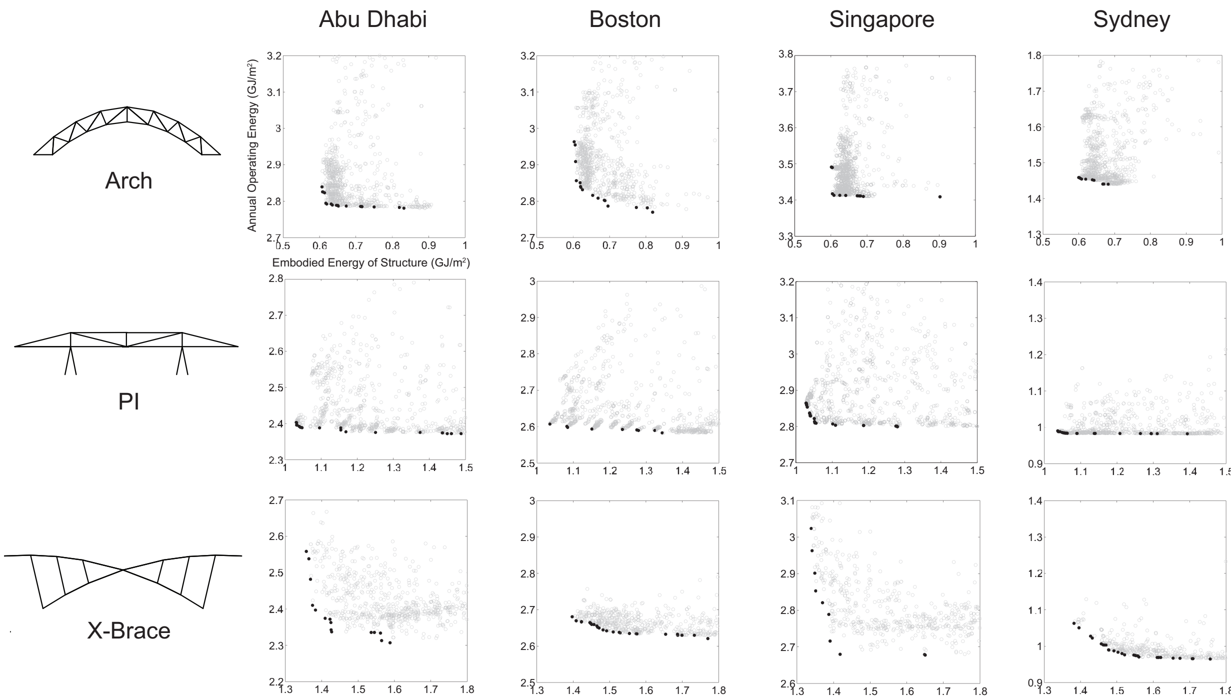

Figure 17: Approximations of the Pareto front between the embodied energy of the structure and the annual operational energy of each building design

In order to provide context, Figure 18 shows each of these Pareto fronts on the same plot. In this figure, the performance of the structurally optimal solution for each set is represented by an asterisk (*). Since structure is independent of location, the structurally optimal solutions are always stacked vertically. In some cases, the NSGA-II algorithm was able to generate a solution on the approximate Pareto front close to the optimal; in other cases it was not. The distance between the highest performing structural solution in each generated set and the overall structural optimum is signified by a dotted line. For each structural system, Singapore had the highest energy loads, followed by Boston, Abu Dhabi, and Sydney. The lowest required structural material was found in the arch, and the highest amount of structure was required by the x-brace. This trend generally indicates that the arch, PI, and x-brace are ranked in order from most to least structurally efficient.

Although each case study demonstrates a tradeoff between structure and energy when a minimum is desired for both, the shape of this tradeoff changes substantially in different contexts. In a few of the plots, there is a gradually sloping curve signifying that the building geometry is transitioning between equally affected performance objectives. In other cases, the Pareto front appears to be nearly horizontal, indicating that the embodied energy of the structure changes more significantly than the operational energy of the building within the optimal set. This is especially true for the PI structure, although a horizontal front occurs in the other systems as well.

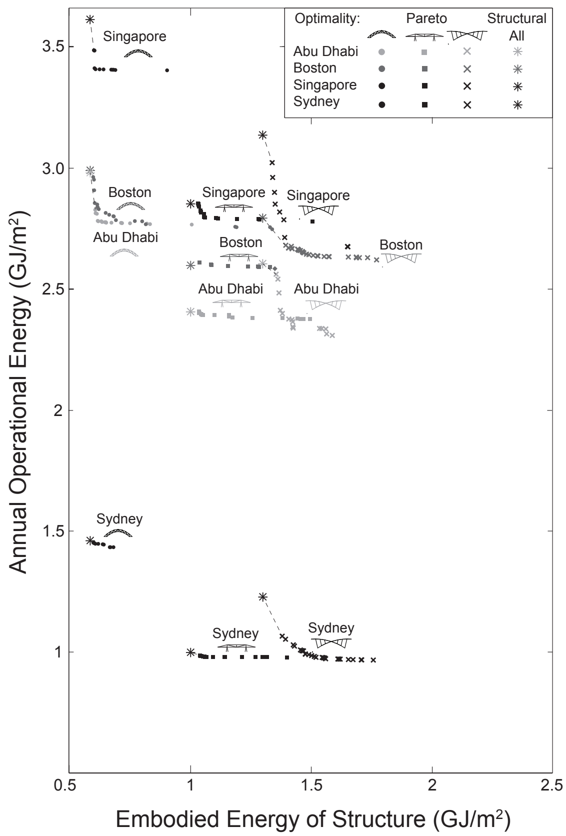

Figure 18: A plot showing the absolute location of each Pareto front along with the structurally optimal configuration for direct comparison

In horizontal Pareto front cases, a designer aiming for the best performing geometry would likely select from the solutions farthest to the left of the front. In the different PI locations, this far left point trends near the structural optimum. However a near horizontally Pareto front could also simply indicate a weakness of the evolutionary approach and its ability to find optima at an extreme edge of performance space. In the cases where there is a flat Pareto front approximation and a structural optimum that performs much worse in terms of energy than the entire approximation set, the actual Pareto front is not flat. Rather, the structural optimum is far away enough in the design space from the rest of the high-performing designs that an evolutionary process is unlikely to find it.

On the other extreme, the Singapore x-brace Pareto front is sharply vertical, with only a small 'hook' towards a few solutions that are slightly better in terms of energy but much worse in terms of structure. This shape indicates that geometric differences have a larger effect on the energy performance than on the structure. Intuitively, this Pareto front shape is more likely to occur in a climate dominated by large energy loads rather than in more temperate locations.

When the Pareto front is sloped sharply in both axes, there is an easily identified 'knee point'. A wide variety of definitions exist for this term, but broadly speaking, a knee point requires an unfavorably large sacrifice in one objective to gain a slightly better performance in the other objective (Deb & Gupta 2011). As such, it is almost always a preferred solution if it exists in a bi-criteria optimization problem, although this is not necessarily the case when taking aesthetic and other architectural conditions into account. Due to their prominence as preferred solutions in pure optimization problems, the geometries of various knee points will be discussed in the next section. Often in this chapter, the clear knee point performs almost as well structurally as the optimum, but the energy performance of the actual structural optimum is sharply worse than the designs nearest to it on the generated Pareto front.

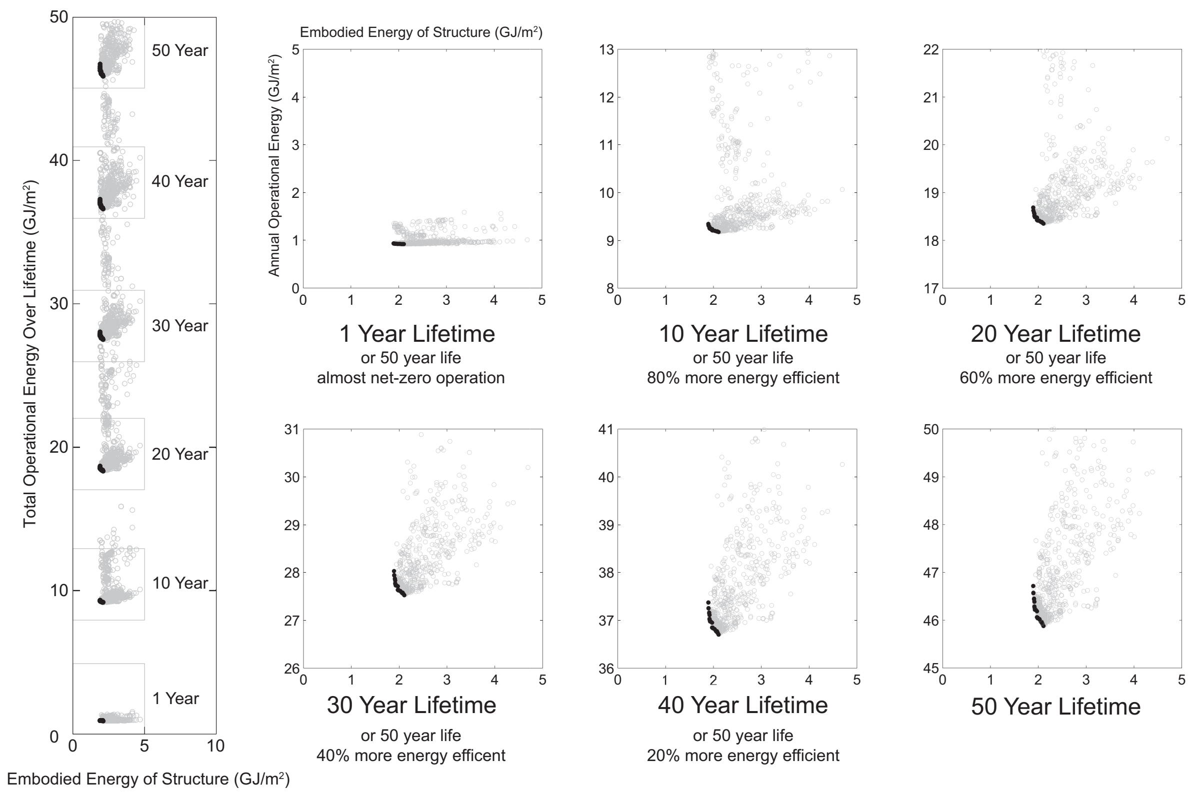

Figure 19: A comparison of the shape and location of the structure-energy Pareto front for different design assumptions concerning the lifetime and operating efficiency of a typical building

It is important to consider that the shape of the Pareto front of these conceptual designs can be stretched or compressed depending on design assumptions. For example, Figure 19 shows structure plotted against simple, whole lifetime energy requirements for different building lifetime assumptions in a typical structure-energy optimization. These Pareto front shapes, which appear nearly flat when only one year is considered but spread out for longer lifespans, stand in contrast with those presented for annual energy requirements in the rest of the chapter. For buildings with longer lifetimes, the Pareto front stretches vertically, illustrating the increased significance of operational energy in total lifecycle usage. Similarly, these same plots could each show a 50-year lifetime building, but with different assumptions about the efficiency characteristics of the envelope and mechanical systems. As building specifications and codes move towards net-zero operating design goals, future conceptual design tradeoffs may become flatter in the energy axis even when the entire lifecycle is considered, increasing the importance of embodied energy in the overall sustainability of a structure.

5.5.2 Geometric implications

Although objective space plots provide valuable insight into the behavior of parametric models, it may be more useful to a visually conscious designer to view the geometric configurations of high-performing designs. Figure 20 provides samples of changing geometry along the Pareto front for each case study and location. Although the number of designs in each Pareto optimal set varies, five representative samples from each set are displayed using a 2D visualization of the structural geometry. The visualizations are arranged from the most structurally efficient design at the top to the most energy efficient design at the bottom. For reference, each design's performance is plotted in black and illustrated by the gray bar graphs located to the right of each glyph.

In the arch case study, there are clear tradeoffs in both the depth of the truss and the overall height of the structure. Each of the locations produces a solution on the far left of the Pareto front that looks visually comparable to the structurally optimal solution. However, the most energy efficient solutions vary widely. In Abu Dhabi, Boston, and Singapore, the arch truss depth and height decrease when moving towards more energy efficient solutions. This minimizes interior conditioned volume and envelope surface area at the expense of a deeper, higher sloped, more structurally efficient design. On the other hand, the Sydney arch becomes larger towards the energy efficient side of its Pareto front, which maximizes its surface area. Sydney's climate is mild enough that although the cooling loads are more substantial than the heating loads, the cooling set point in the model (26 degrees Celsius) is above the outside air temperature most of the time during the year. However, a building's cooling system in Sydney must still offset the heat loads due to people, equipment, and other interior elements that produce heat. In this context, a greater surface area allows for more exchange with the cooler outside air, which ultimately lowers annual cooling loads.

It is possible that certain climates have a tall arch as the most energy efficient Pareto solution and others have a shallow arch because the structurally optimal configuration is generated by truss depths and heights located in the middle of their variable ranges. Thus, it is possible to get structurally worse performance on both sides of the optimum, either by making the arch too tall or too shallow. In Abu Dhabi, Boston, and Singapore, the structural performance gets significantly worse as the truss depth and arch rise go to minimum values. However, there is a sharp knee point in the Pareto front for the hot climates, which Figure 20 shows to occur at a design that is visually near the structural optimal, but also slightly smaller in terms of surface area and volume.

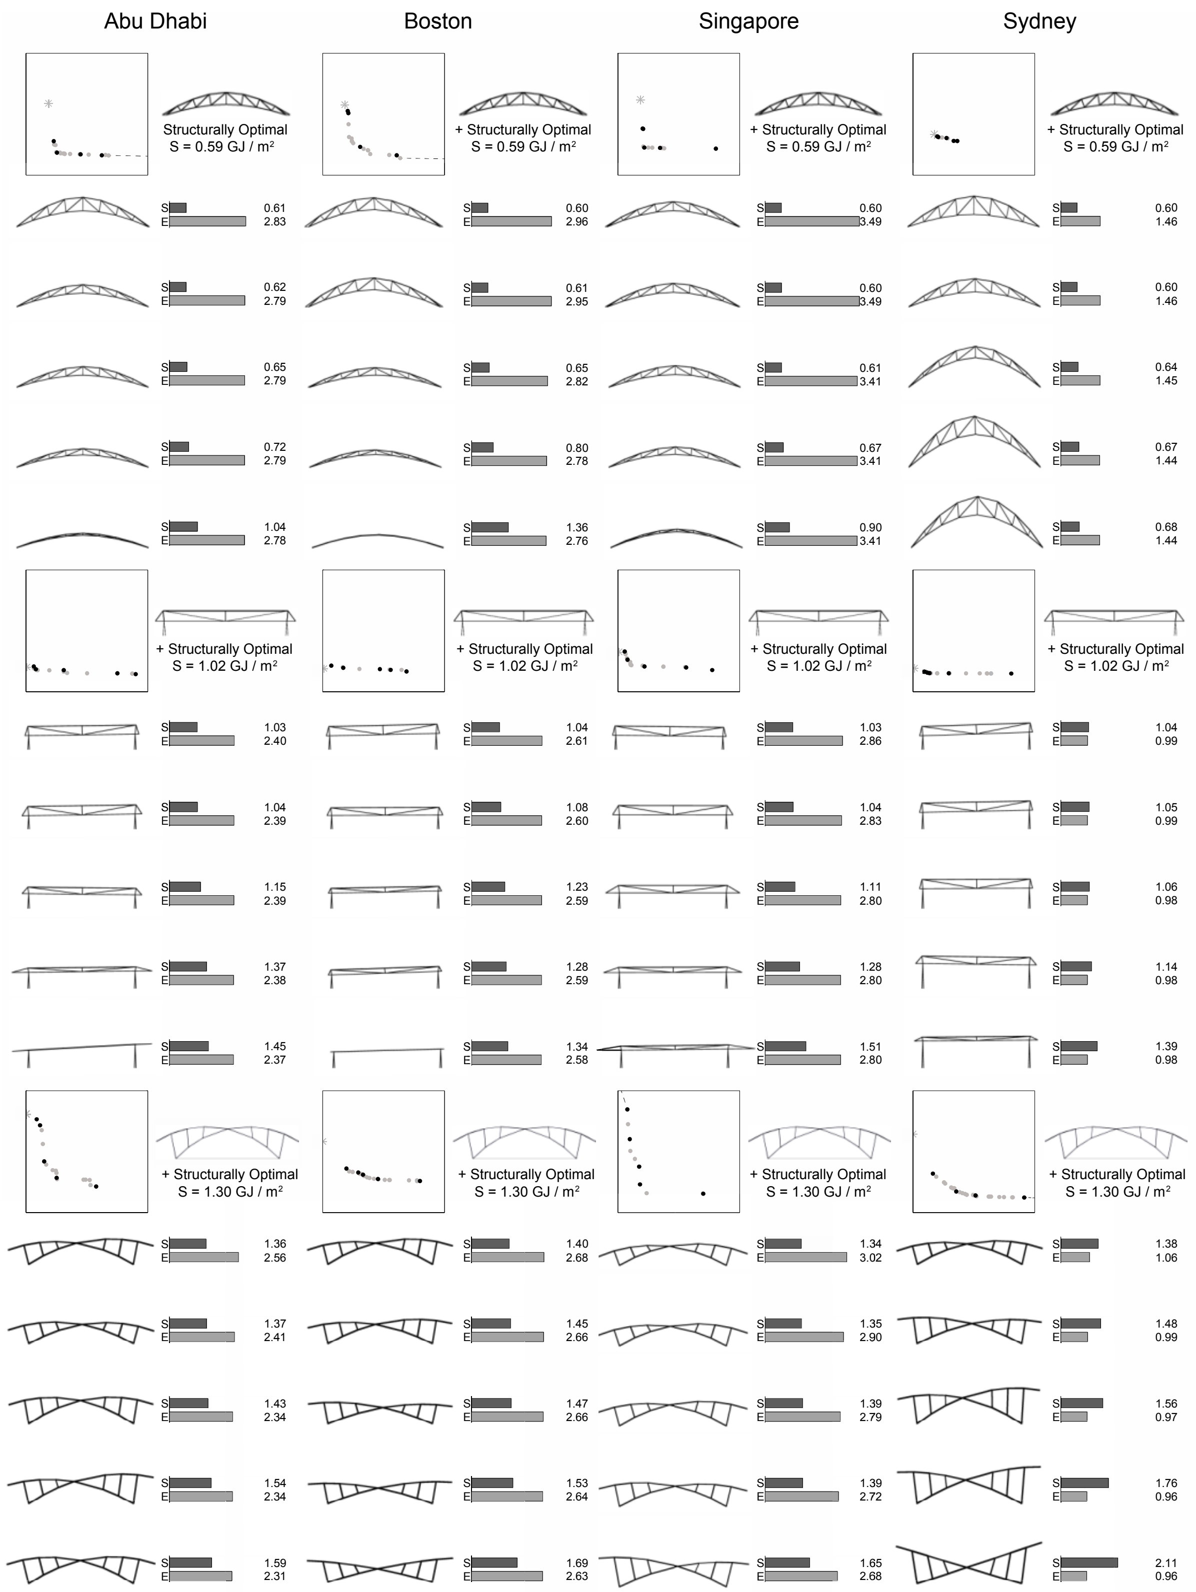

Figure 20: Overall building geometry representations of each Pareto front

In the PI configuration case study, the NSGA-II is also able to find nearly structurally optimal solutions, but the overall nature of the transition to most energy efficient design changes for different climates. The parameterization of the PI model allows for similar tradeoffs in truss depth compared to the arch, but because the surface area of the envelope changes only slightly with a deeper truss, this single variable has much less ability to influence the overall magnitude of operational energy requirements. Especially in Abu Dhabi, Boston, and Sydney, the Pareto fronts are nearly horizontal, suggesting to a designer that there is little energy penalty for finding a more structurally efficient solution. Although Abu Dhabi, Boston, and Singapore all trend towards a shallower truss while moving right along the Pareto front, there is a notable difference in the overhang dimensions between the hot and cold climates.

In Boston, the overhangs remain fairly close to the building envelope, indicating that structural efficiency, daylighting potential, and helpful solar gains outweigh the benefits of static shading. In terms of architectural program, a large overhang over an entrance or airport drop-off lane may actually decrease the energy efficiency of the design in these climates. The situation in Abu Dhabi and Singapore is the opposite—a large overhang cuts out solar gains and leads to lower cooling loads and greater overall efficiency. Sydney's more energy efficient solutions share this characteristic of larger overhangs, but the effect is visually insignificant compared to the increased height and surface area of the best energy performers. Overall, Sydney's best PI energy solutions rely on taller columns to increase surface area rather than greater truss depth.

The PI case study also shows the positive performance effects of asymmetry and lean. Due to the combination of dead and live structural loading being modeled, a structurally optimal solution is always symmetrical, and there is a potentially significant structural penalty to be paid for introducing asymmetry into the system. However, the energy implications of shading, solar gain, and daylighting may push the Pareto optimal solutions towards asymmetry in a bi-objective optimization. In each modeled case study, the left side of the structure (when looking at the 2D visualization) faces south, and the right side faces north. In Abu Dhabi and to a lesser extent Singapore, the most energy efficient solutions lean noticeably to the south, which leads to smaller wall and glazing areas and more effective shading on the south side of the building that sees the most sun. This result stands in contrast to the arch case study, where there are no windows and consequently no substantial solar gains or daylighting opportunities, leading to symmetrical solutions.

The x-brace model allows for a combination of arch and overhang effects to influence its form. In each climate, the combination of these effects leads to a gradual transition between best structure and energy performers. The structurally optimal solution is a symmetrical, tall x-brace with highly curved spanning structural members. Although the best x-brace structural Pareto solutions are a significant distance away from optimal in a few of the climates, each generated front curved towards the structural optimum, rather than having a sharp hook like in other examples. In Abu Dhabi, the geometric transition is subtle and difficult to visualize, as all of the solutions are relatively shallow, will small surface areas and shading edges that curve down towards the windows they protect.

Singapore, which is also in a hot climate, shows a similar response. However, the transition in Boston is much more noticeable, as the main arch members become less curved while moving down the front. The flatter shading elements allow more sunlight to enter, which can help offset the dominant heating loads in Boston. In Sydney, the main structural members also become flat but at a higher angle, which generates much taller walls and windows. Again, this leads the Sydney best energy performers to have a higher overall surface area and volume than in other climates.

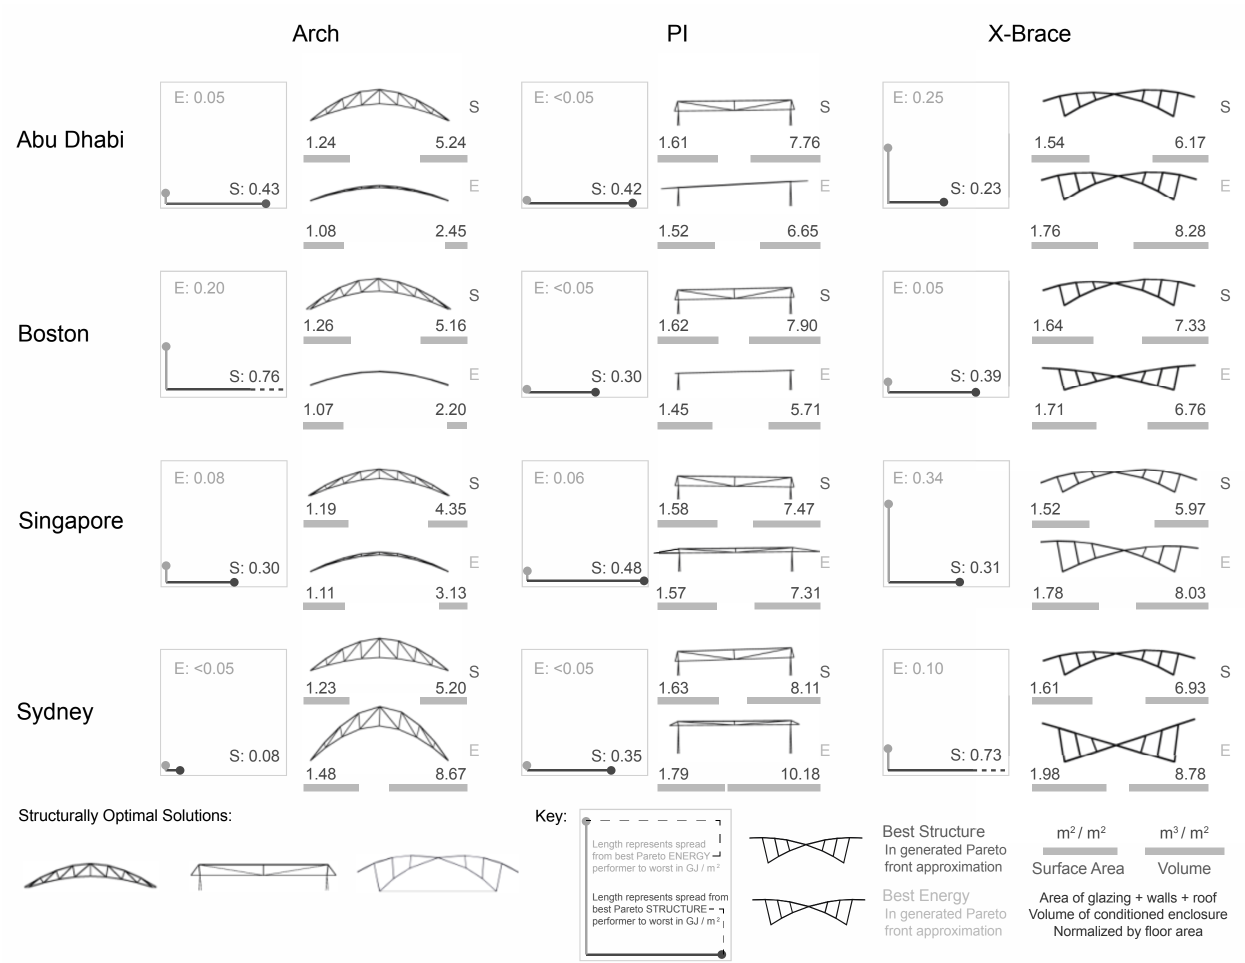

Figure 21: The best structure and energy performers from each Pareto optimal set, as well as the range of performances generated by the NSGA-II algorithm in each location

Although these geometric transitions happen gradually along the Pareto front, even at a resolution approximated by a population size of 50 designs, the contrast between structural and energy high performers is more noticeable when considering only the edges of the Pareto front. Figure 21 shows the best structure and energy performers generated for each case study, while also giving an indication of the performance range of the front in absolute units. From these images, it is clear that the first three climates prefer energy solutions with smaller envelope surface areas, while Sydney prefers the opposite. Large overhangs have the most substantial effect on Abu Dhabi and

Singapore, but are not found in the Pareto optimal set for Boston. Asymmetry is also most prominent in Abu Dhabi, which is consistent with its increased sensitivity to solar effects. The Pareto sets for the PI structure have large ranges of structural performance with only minimal changes in energy. There is no similar trend for the other two systems, as the relationship between the structure range and energy range changes for different climates. When considering the entire modeling results, there is a substantial amount of geometric diversity among the optimal solution sets, as well as a significant difference in the design responses to varying climates.

5.5.3 Frequency of solutions

The methodological choices of the research presented in this chapter have some influence on the design results, and a brief discussion of the optimization method is given here. In the performance evaluations for each model, the structural calculation was a finite element analysis of linear elements and occurred almost instantaneously. The annual energy simulation required substantially more time—although actual computational time depends on hardware, on a standard desktop computer each design simulation took around 15-20 seconds. Including some computational time for the implementation of the NSGA-II, this resulted in overall runtimes of around 12 hours for each case study in each context. As such, the case studies were generated for a finite number of 20 generations rather than until a certain threshold for change in performance was reached.

Figure 22 shows a histogram of the performance results for every solution generated through the course of the case studies. In some of the structural configurations, the NSGA-II was able to find the zone of high performance solutions within a few generations. The ease with which certain problems could be mapped is indicated by the large spike in frequency of solutions near the edge of optimality; after an initial random seed of the entire design space, sometimes as many as 800 of the 1,000 solutions considered throughout the entire process fell within the best interval of the histogram. In other cases, the more gradual trend towards optimality indicates a higher degree of computational difficulty in surveying the objective space for a given architectural problem. This observation is important to keep in mind when considering which tradeoffs can be most effectively navigated through the use of multi-objective optimization techniques.

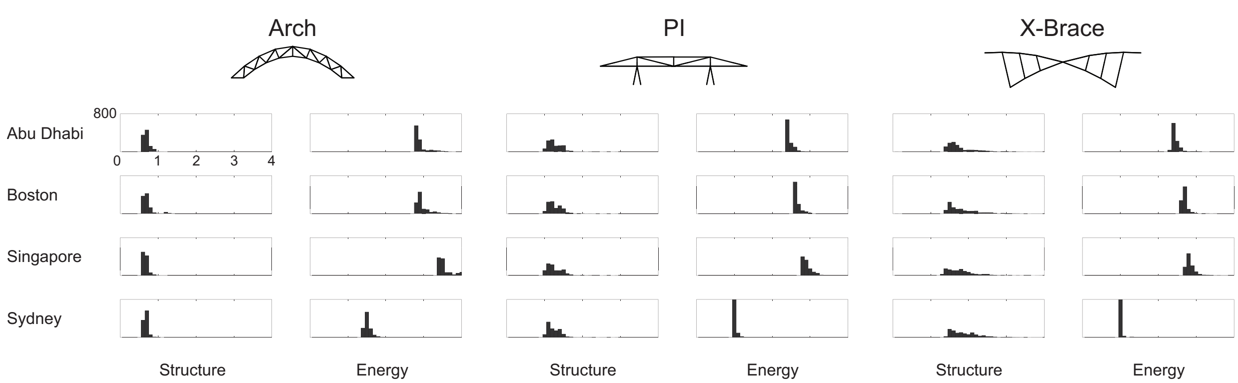

Figure 22: Histograms showing the frequency of performance scores within the generated solutions

5.6 Discussion and summary of contributions

This chapter makes three key contributions to the broad goal of encouraging optimization in integrated conceptual building design.

First is the demonstration of how a multi-objective optimization methodology for structural efficiency and operating energy efficiency can be used to generate a geometrically diverse range of high-performing designs. The results of the various case studies presented in this chapter show that although finding expressive architectural forms and setting up an intelligent design space requires designer intuition on the front end, a MOO approach can push high-performing solutions in interesting, unexpected ways. With contemporary computational tools that generate rapid performance feedback in conceptual design, this process can already be integrated into a typical design workflow. As future theoretical advances are made and software develops, the ease of using MOO interactively will only increase. At the same time, the complexity of these results also highlights the role a designer plays throughout the process, as a human must prioritize design goals and consider visual impact. In this way, optimization in design does not have to be reductive or overly constraining to a creative architect—it can be used as a sophisticated tool to support decisions rather than make them.

Second, there is value in the broad lessons of the case studies concerning the relationship between geometry and performance. As can be seen by the wide variety of Pareto optimal results, the precise nature of tradeoffs between structure and energy are not always intuitive, and they are extremely sensitive to context. However, by restricting this research to the architectural typology of a long span building and selecting problems that are common to this typology, many of the geometric lessons learned through this study could be applied to future designs.

Third, this chapter presents new data showing how considerations of structural material efficiency and embodied energy compare in magnitude to total building energy usage, and how this knowledge affects sustainable design strategies. The case studies in this chapter show that structural embodied energy and annual operating energy requirements are on the same order of magnitude, with slight variations due to structural system and location. This is a simplified comparison that does not include all other embodied energy in the building, allowances for renovations, changing equipment efficiencies, or many other details that would be included in a full lifecycle energy analysis. Nevertheless, the absolute energy comparison shows that at current building code levels and reasonably long building lifetimes, structure is entirely overshadowed by energy. Yet as future buildings reach increasingly stringent energy efficiency standards, structure may become the dominant objective in MOO methodologies for particular building types.

5.7 Future work and concluding remarks

This research could be extended in a number of ways. One next clear step is to look at tradeoffs in terms of monetary cost in addition to energy usage. Although cost and energy are largely correlated, the nature of specific tradeoffs could fundamentally change with this different perspective. Analyzing cost in addition to energy could also allow for more specific study of the effects of upfront costs and price changes over time on optimized architectural form. In addition, there are many aspects of the energy model that could be added into the parameterization. Although this study focused on changing geometry, the effect of envelope attributes, mechanical system efficiency, and even energy modeling assumptions could be explored in an MOO workflow. Similarly, efforts could be made to use more area-specific coefficients for material and energy production, as opposed to the broad, standardized assumptions used in this thesis. Additional design objectives such as constructability could be added to the MOO process, either quantitatively if possible or at least articulated in the discussion of user-evaluated design criteria. Any of this future work could use similar case studies to build on the results of this thesis, or focus on new problems that are relevant to architectural design.

In conclusion, this chapter is another step in demonstrating how MOO can be applied to geometrically, architecturally interesting conceptual design problems, leading to diverse, high-performing results that are difficult to generate through other means. These case studies can combine with other chapters to increase the prevalence and success of optimization in sustainable conceptual design, which can yield significant improvements in the overall performance of future buildings.

6 Structure + Light in Gothic Cathedral Design

The designers of the great cathedrals of Europe have long been praised for their practical understanding of stone structures, which allowed them to build soaring, durable monuments full of characteristic pointed arches, flying buttresses, cross vaults, and stained glass windows. Yet despite the physical existence of many examples of this architectural style, the original builders' reliance on rules of thumb and inherited design knowledge has left the door open for a scholarly debate concerning the structural behavior of Gothic churches (Huerta 2009). Methods of calculation throughout the years have included the use of static equilibrium to find thrust lines, photoelasticity, finite element models, and limit analysis, but many of these techniques fall outside of the typical training and practice of modern engineers. Due to the current age of these structures and the difficulty of analysis, many are at risk of failure due to dynamic loading or being retrofitted in a way that is architecturally destructive. Thus, there is a clear need for the assessment of the safety and stability of many existing cathedrals (Block et al. 2006).

In addition to height and the structural achievements that allow for it, builders of Gothic cathedrals had another significant design objective: natural daylighting. Since events in these churches were a prominent feature in the daily schedules of many medieval Europeans, and interior images and iconography played an important theological role, proper use of daylighting was key for designers. Many earlier Romanesque churches had difficulty with interior lighting due to their relatively small windows and thick walls, but the invention of the flying buttress represented a structural innovation that radically improved the daylighting potential of religious architecture. Although Gothic cathedrals have been celebrated for their control of light (Baker et al. 1993), there is a lack of research into the interplay between light and structure. In the specific case of exterior buttresses, the structural form has a direct impact on how much light is allowed into adjacent windows, but this impact has not been quantified in a meaningful way.

In response to the need for the study of Gothic structural systems and their implications for daylighting behavior, this chapter presents an analysis of the main buttressing system of Amiens Cathedral. First, a safety factor against failure is calculated for a number of different load cases and sets of structural assumptions. Then, by reducing the geometry of the main pier to a parametric model, this research determines how the structural stability and daylight autonomy (DA) change in response to different pier shapes, representing the performance effects of the design decisions that were made by medieval cathedral builders. Each analysis has value in itself, since structural calculations can address how the cathedral stands up and how likely it is to keep doing so, while the daylighting analysis attempts to quantify a performance aspect of a historical structure that is widely acknowledged but not well understood. Together, these simultaneous analyses lead to a greater understanding of how Gothic builders navigated trade-offs between their two most significant design objectives.

6.1 Background: analysis of Gothic structures

In the introduction to *Experiments in Gothic Structure (1984)*, Robert Mark calls the structural design and performance of Gothic cathedrals "the most enigmatic aspect of the technology of historic buildings". While the spread of Gothic architecture through Europe can largely be attributed to its wide acceptance as a building style, it was made possible through the development of radical new technologies that allowed builders to achieve height, transparency, and interior volume at scales not previously attained. Many art historians have explored the religious, aesthetic, symbolic, and even social developments that lead to the building of large Gothic cathedrals, but few of these researchers have the technical expertise to understand the engineering aspects of these buildings. Knowledge of how these structures perform is further obscured by the fact that medieval builders likely relied on models, rules of thumb, and design knowledge passed down through generations of skilled masons and carpenters, leaving only a few records that could give insight into their methods. As this way of building is fundamentally different from the modern approach, even many engineers with today's sophisticated analysis techniques would struggle to explain how Gothic cathedrals were built, and how they stand up (Mark 1984).

The designers' inability to clearly express analytical behavior has led to different historical explanations as researchers have attempted to discover both what Gothic builders intended, and what they achieved. The most prominent, early academic argument was structural rationalism, in which the development of the Gothic form represents a progressive problem-solving exercise. In this view, every part of a cathedral must have a structural function, because otherwise it would not have developed as it did. Robert Willis and Viollet-le-Duc most famously argued this view, and the contemporary understanding of cathedral performance owes much to their explanations, writing, and drawing (Mark 1977). Yet as some researchers have pointed out, there are many elements of

Gothic cathedrals, such as openwork flying buttresses, that do not in fact successfully achieve their purported function, undermining the rationalist argument (Bork et al. 1997). This realization means researchers must go beyond the assumptions of pure structural function to determine how cathedrals actually behave and think more deeply about other factors that may have influenced designers and their shaping of various stone elements.

In the last century, there have been a few key contributions made to the general understanding of structural performance of Gothic cathedrals. Most recent researchers have sought to confirm analytically what Gothic builders likely understood intuitively and empirically, but they are not constrained by a dogmatic desire to confirm pure structural rationalism and perfection, which has lead to more accurate descriptions of behavior (Courtenay 1997). Although these researchers have analyzed arches, vaults, towers, spires, and other Gothic elements, this review will focus on buttressing systems and the forces that act on them, including vault thrusts.

One of the most significant researchers of Gothic structure from an engineering perspective is Robert Mark, who dedicated a large portion of his academic career to understanding how cathedrals behave. Mark published *Experiments in Gothic Structure* in 1984, and this book included the major innovation of using photoelastic models to analyze the magnitude of stresses found across the section of a cathedral. These experiments were conducted by loading a physical model made of epoxy-plastic, cooling it to lock in the deformations, and viewing the resulting stress lines using a polarizing-light instrument. Mark generally created models of cathedral sections taken at one of the bays of the main nave, since this portion of the cathedral is most critical for overall stability as there are no buttresses in other directions for added stiffness, and because the nave and adjacent aisles are repeated longitudinally along the main axis of the church. Later on, Mark also used computational methods to arrive at similar results.

Jacques Heyman analyzed the behavior of cathedrals contemporaneously to Mark's work on historical structures, but with different methods. Heyman stressed plasticity theory, which assumes for masonry structures that stone has no tensile strength, compressive stresses are so low that stone has effectively infinite compressive strength, and sliding failure does not occur (Heyman 1966). These assumptions lead to basic stability calculations and limit analysis since masonry is highly indeterminate. Heyman published the bulk of his ideas in *The Stone Skeleton (1997),* and many of his concepts will be used in this chapter. Heyman also references Ungewitter (1901), who developed a method for estimating the vault thrusts based on a geometric aspect ratio and the specific material of the cathedral. This work has given researchers the ability to calculate reasonable values for the main forces that a buttressing system must resist. An example of more contemporary research into the structural aspects of Gothic cathedrals is given by Coccia et al. (2015), which presents a detailed analysis of their wind strength. While all of these contributions have increased our knowledge of the structural behavior of these cathedrals, none have addressed how this behavior interacts with lighting considerations. To understand how Gothic builders pursued and achieved both height and light, it is necessary to studying natural daylighting of cathedrals as well.

Other researchers have addressed the daylighting performance of Gothic cathedrals, although this has been done in a quantitative manner only recently. *Daylighting In Architecture: A European Reference Book* (Baker et al. 1993) notes that considerations of daylighting have been around since the writings of Vitruvius. However, design for daylighting became more technologically possible in Northern Europe when glass production became more affordable, and it grew more necessary with the noticeable lighting shortage brought about by an early medieval climate shift. This design condition contributed to the high ceilings and open plans favored in many examples of European architecture. Ramzy (2013) speaks to this result in an analysis of the lighting quality of five examples of sacred architecture. Although Ramzy's research focuses on the religious and experiential qualities of the light, it determined the Gothic form to have the highest empirical levels of natural daylighting out of the examples considered.

Extensive analysis of the daylighting quality of Gothic cathedrals has also been conducted by Simmons & Mysak (2010), who took a survey of the particular arrangements of stained glass and their effects on relative transmissivity and interior illuminance. This research addresses the design shift in France from deep red and blue pieces of glass to more transmissive whites, yellows, and near-monochrome grisailles over the Gothic period. The grisailles in later Gothic French cathedrals resemble what had previously been built in England, further north. The researchers have theorized that a combination of geographic latitude, a climatic shift towards cooler temperatures and cloudier days, and a variety of social and technical factors may have influenced this change. Simmons' research was helpful in the development of the daylighting model presented in this chapter, but it does not address the effects of structural form on daylighting.

6.1.1 Amiens Cathedral

This chapter uses Amiens Cathedral as a case study, and thus it is important to explain the reasons behind its selection and acknowledge previous research into its history and behavior. Amiens Cathedral is located about 75 miles north of Paris, and its initial building phase was from 1220-1270, making it an example of the early Gothic style. As with many Gothic cathedrals, portions of the church were added on or modified over the following centuries. In terms of interior height and volume, it is the largest church in France. It was virtually untouched during the French Revolution, and it received only minimal restorations during the period it was under the care of Viollet-le-Duc, which began in 1849 (UNESCO World Heritage Centre).

Amiens was selected as a case study for a variety of reasons. First, there is a large amount of documentation concerning the cathedral's dimensions, including drawings by Viollet-le-Duc and a diagram of its section and structural loading from Dehio & Bezold (1892). Many photographs of the cathedral also exist in Tallon and Murray's *Mapping Gothic France,* and relatively much is known about the master builders who oversaw its construction. All of these documents were relied upon heavily in the analysis of the cathedral. Second, there are a number of peculiar aspects of the church that suggest that its structural condition should be further analyzed. Bork et al. (2011) have shown that the initial rear openwork buttresses do not adequately take the thrusts of the vaults, and thus an array of secondary flying buttresses had to be added during construction. Furthermore, the stability of some of the vaults were questioned by the medieval builders, who decided to install a post-tensioned iron chain around the entire clerestory level to help carry these forces. Tallon and Murray have also used laser-scanned drawings to observe a number of interior columns that are displacing at alarming rates. In light of these discoveries, it seems that the stability of the main buttressing system should be analyzed to give a more complete picture of the current safety of the church.

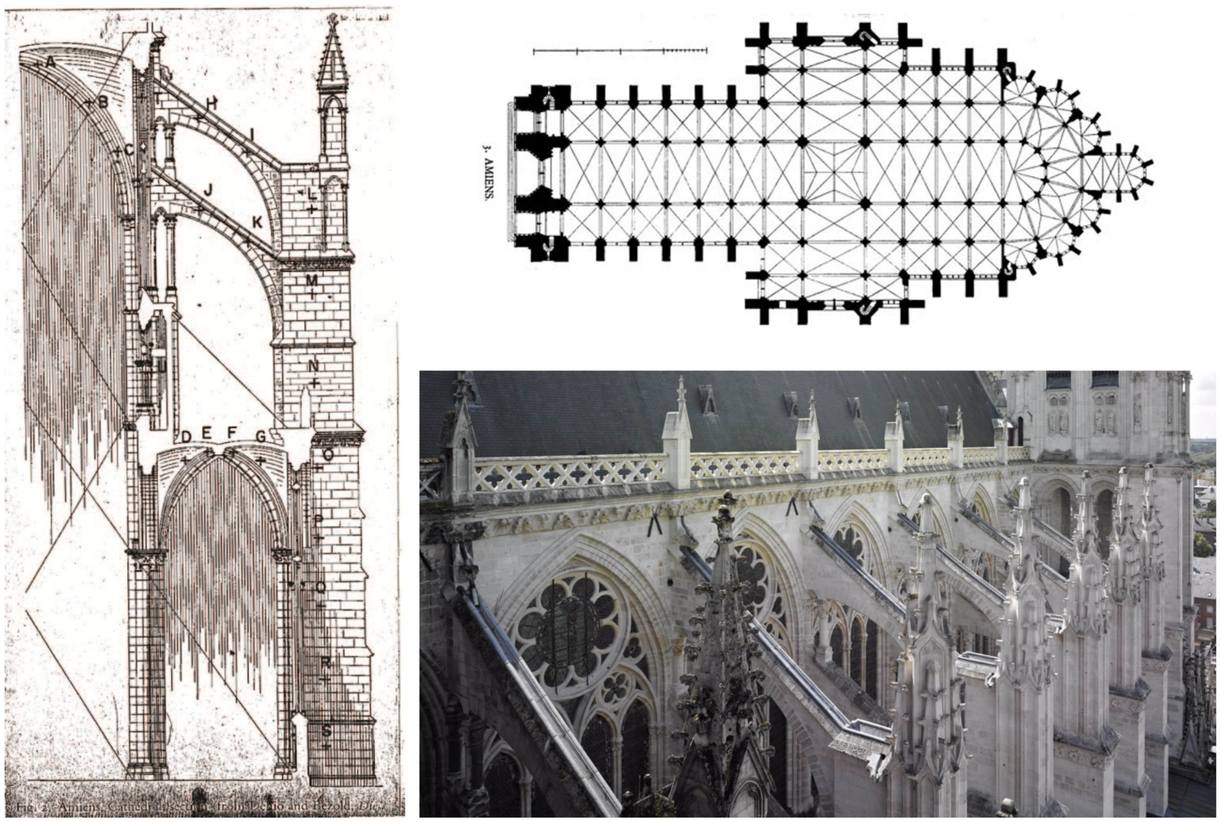

Figure 23: Section, plan, and photograph of the buttressing system in Amiens Cathedral. Section from (Nelson 1971), plan and photo from (Tallon & Murray)

While conducting a structural study of Amiens, much can also be learned through a simultaneous and comparative daylighting analysis. Thus, there are two main goals of this research. The first is to analyze the structural behavior of the Amiens buttressing system and determine relevant safety factors against failure. The second is to use multi-objective optimization techniques to explore the interplay between the structural and lighting performance of the buttresses in response to retroactively applied design modifications, with the intent of revealing what Gothic designers may have known about structures and daylighting, as well as what they actually achieved.

6.2 Methodology

The methodology for this research is broken into two main sections: the structural calculation procedure and the parametric daylighting analysis procedure. The structural calculations determine a factor of safety against failure of the main pier buttresses based on different cases and assumptions. The parametric daylighting analysis compares the impact of different pier shapes on both the structural stability and daylighting performance of the cathedral, and thus relies on the structural calculation methodology for the performance evaluation. Both calculation methods are described in this section.

6.2.1 Structural calculations

The structural evaluation focuses on the stability of the pier buttresses supporting the main three-aisled nave, shown in the section drawing in Figure 23. This is the most significant part of the building for an overall stability analysis, since it represents the repeated module of the narrowest portion of the church, where no other out-of-plane buttresses exist. The main pier is situated against the lower vault of the outside aisle, where it rises above the lower vault to meet a pair of flying buttresses. These flyers in turn are attached to the clerestory wall, taking the thrust of the upper vault as well as contributing to the cathedral's resistance of lateral wind loads. The buttress is stepped in from the outer edge by about 2.4 meters and is topped with an ornate pinnacle.

Since Gothic cathedrals are primarily unreinforced masonry built of a material with high compressive strength, a plastic analysis is more appropriate than one based on elasticity methods. As such, the main failure mode considered is overturning, which can be calculated with a simple stability analysis. The safety factor against overturning will be defined as:

$$\mathrm{SF} = \frac{\text{Restoring Moment}}{\text{Overturning Moment}} = \frac{\mathrm{M_R}}{\mathrm{M_O}} \tag{1}$$

Where M_R is the restoring moment, consisting of all of the forces preventing the pier from rotating outward, and M_O is the overturning moment, consisting of all the forces that could push the pier over. Broadly speaking, arch thrusts and wind loads are the overturning forces the buttresses were built to resist, which they do primarily through their own weight. Since the cathedral is symmetric in section at this location, only half of the system was considered for analysis. Three different geometric assumptions about how the pier would fail were tested and will be explained later in detail.

6.2.2 Load estimates

The first step in this safety factor calculation was to estimate the loads being placed on the pier. These loads consist mainly of vault thrust, wind loads, and the self-weight of the flyers, pinnacle, and pier itself (Figure 24).

Most of these calculations rely on weight estimates of different portions of the cathedral catalogued by Nelson (1971). These are the same weight estimates used by Mark in his photoelastic models. In order to first determine the thrust of the cathedral, the overall mass of the vault was calculated, along of the centroid of this mass. There must be a constant horizontal force in the vault, and it was assumed that the magnitude of this force could be determined by drawing a line from the centroid of the vault mass to the point of force transfer to the lower flyer.

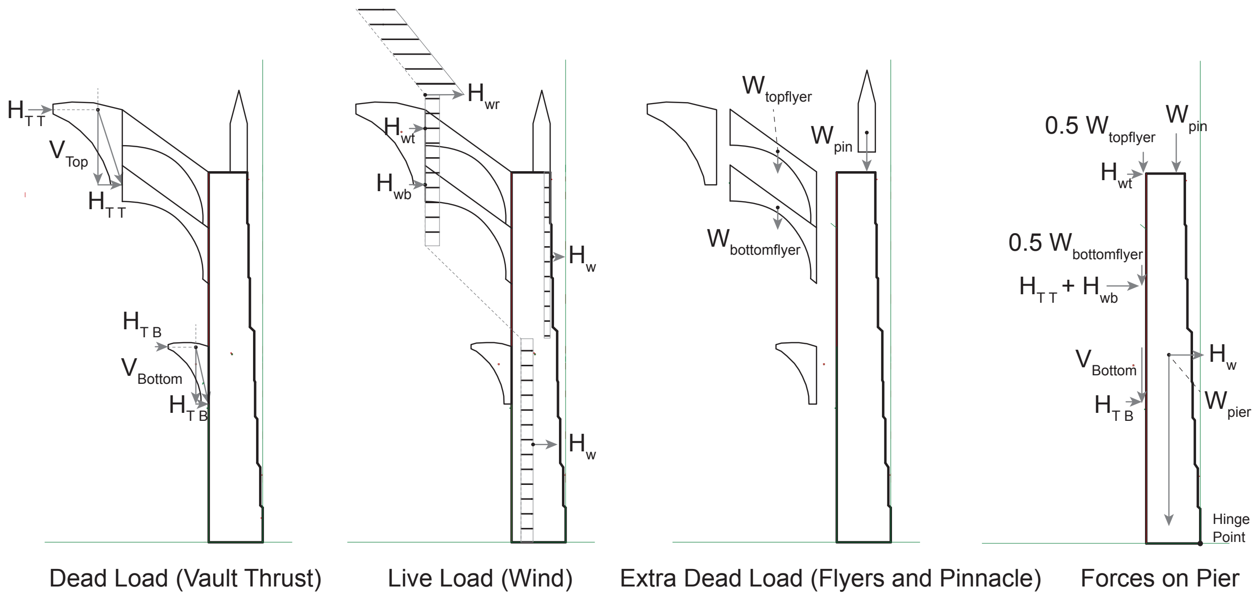

Figure 24: Diagrams of loading on the main pier buttresses

The vertical component of this line represents the weight of the vault, which was conservatively assumed to be taken entirely by the interior column for the upper vault. This leaves only the horizontal component to be transferred through the flyer to the buttress. The values for horizontal thrust obtained by this method were 48.1 kips for the upper vault and 30 kips for the lower vault.

Since this method relies heavily on assumptions, the final values for vault thrust were compared to values given by other researchers. In Mark, a variety of ratios of vertical weight to horizontal thrust were obtained using physical models. By comparing the geometry of Amiens Cathedral to a similar cathedral, Mark determined a weight/thrust ratio of 0.39, leading to a thrust estimate of 48.4 kips. Another researcher, Ungewitter (1901), developed ranges for this horizontal/vertical force ratio based on differences in material and arch aspect ratio. Using Ungewitter's methods, the horizontal thrust of the main vault would fall between 50-54 kips. Table 11 shows that each of these methods leads to similar results, confirming that the initial assumptions are reasonable.

Comparison of Historical Load Estimates			
	Value Used	**Mark**	**Ungewitter**
Calculation Method	Graphical	Thrust/Weight Ratio Physical Modeling	Height/Span Ratio Derived Values
Top Vault Thrust	214.0 kN	215.3 kN	222.4-240.2 kN
Bottom Vault Thrust	133.4 kN	133.4 kN	102.3-137.9 kN
Wind Loads	Used Mark's Values		

Table 11: Comparison of historical load estimation methods for main pier buttresses

The next group of forces to consider was wind loads. Although the wind loads themselves are reasonably easy to calculate, their exact effect on the isolated pier is not. Wind loads on the cathedral include positive pressure on the windward side and negative pressure on the leeward side, acting on the walls, roof, and piers. However, the response of the structure to these loads is more complicated, with Mark theorizing that wind loads are taken by the entire mass of the section as a braced frame. For this analysis, which is focused on the buttresses only, some assumptions are required concerning how the wind forces flow through the structure. This thesis considers two conditions in addition to the dead load: the case where only wind forces on the clerestory walls are being transferred through the flyers to the pier, and the case where these flyer forces are combined with direct suction forces acting on the lower walls and the pier itself. This choice of load cases was influenced by various researchers' conclusions on the impact of wind loads and the behavior of upper flyers, including a paper by Fitchen (1955). The worst-case scenario occurs when the pier is on the leeward side and wind forces are contributing to the overturning moment. Diagrams of the different loading cases are given in Figure 25.

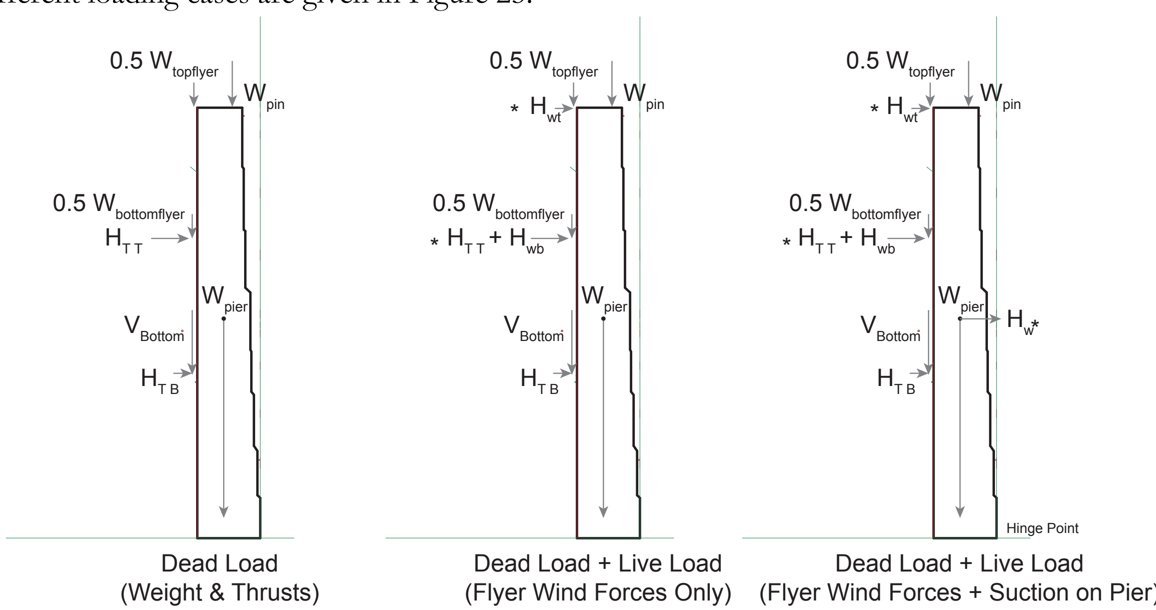

Figure 25: Diagram of considered load cases

In addition, half of the vertical weight of each flyer and the full weight of the pinnacle were calculated and added to the safety factor equation. Table 12 gives a full list of forces used in the safety factor analysis.

Forces Used in Analysis				
Vault Forces				
Force	V_{top}	V_{bottom}	H_{TT}	H_{TB}
Description	Weight of top vault	Weight of bottom vault	Thrust of top vault	Thrust of bottom vault
Value Used	556.0 kN	413.7 kN	214.0 kN	133.4 kN
Moment Arm (Full Pier)	5.33 m	5.33 m	29.30 m	14.30 m
Wind Forces				
Force	H_{wt}	H_{wb}	H_w	
Description	Horizontal force from wind carried through the flyers		Suction force on pier and wall	
Value Used	80.1 kN	80.1 kN	213.5 kN (wall) + 53.4 kN (upper pier)	
Moment Arm (Full Pier)	35.05 m	14.30 m	10.97 m (wall) & 29.87 m (upper pier)	
Element Weight Forces				
Force	$W_{topflyer}$	$W_{bottomflyer}$	W_{pin}	W_{pier}
Description	Weight of top flyer	Weight of bottom flyer	Weight of pinnacle	Weight of pier
Value Used	209.1 kN	200.2 kN	321.6 kN	7,295.1 kN
Moment Arm (Full Pier)	5.33 m	5.33 m	2.47 m	3.41 m

Table 12: Description of force estimates used in analysis

6.2.3 Failure condition assumptions

Once the loads were calculated, they were applied at the point of transfer to the pier to determine a safety factor against failure. To calculate the safety factor, three different conditions were considered (Figure 26). First, the entire pier was modeled as one rigid body (1) that rotates about its outside corner. Next, the possibility of pier breakage (2) was considered. In this situation, a 45-degree slice was taken through the pier from the overturning point, and only the pier mass above this slice contributed to the restoring moment. Finally, the safety factor against the development of tensile forces (3) in the pier was considered, with the overturning point moved in by one third of the width of the pier. The assumed failure cases get less conservative from (1) to (3), and each will be evaluated in the results section.

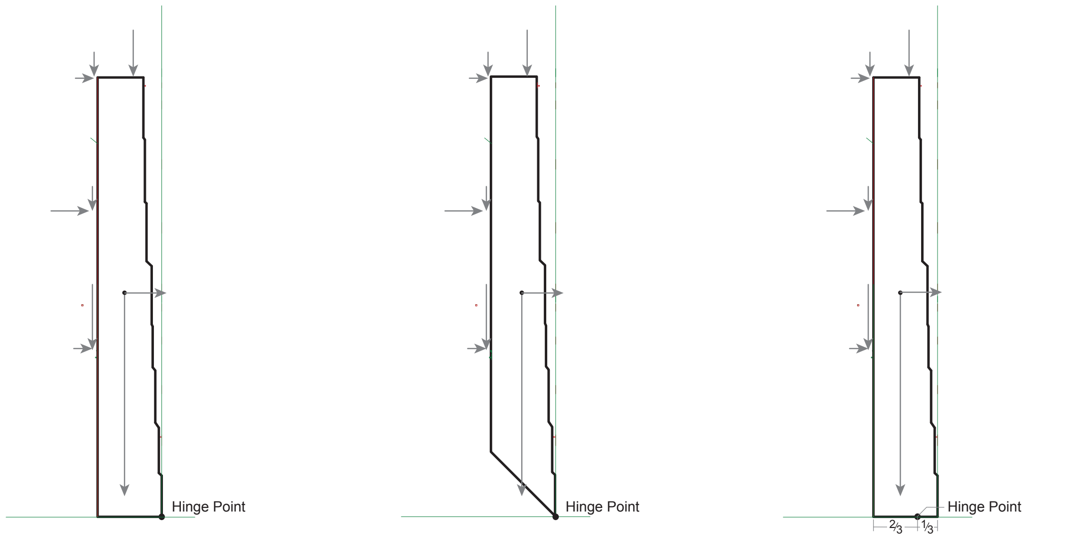

Figure 26: Different failure assumptions considered in the structural analysis

6.2.4 Daylighting calculations

The next portion of the research studies why Gothic designers developed and shaped the pier buttressing systems in a particular way. There are many different examples of cathedral buttress shapes, ranging in terms of size, height, flyer dimensions, and relationship to the wall. Although there is little documentation on the actual design thinking of the time, many observers have attempted to justify these differences retroactively. Unfortunately, there is considerable disagreement on the design decisions that shaped these buttresses as well as the effects these decisions had on building performance. This section uses simulations to test how the cathedral would have responded had the designers made the piers taller, deeper, wider, or more inwardly sloped.

6.2.5 Parametric model setup

To address the question of buttress design decisions and their implications, a computational parametric model was created. This model relies on four variables to generate the geometry of the pier: height, depth (out from the wall of the cathedral), width (in the direction of the wall), and step-in, which is the difference in depth between the bottom of the pier and the top and captures the inward slope of the outer edge of the pier (Figure 27). In most cathedrals this line is stepped rather than sloped, but the approximation was made in order to simplify the performance evaluation. The parametric model was developed using Rhinoceros and Grasshopper so that the geometry could be evaluated using existing plug-ins in addition to the structural calculations mentioned earlier in this chapter.

A script was written in Grasshopper to loop through a range of design samples within this parametric design space, evaluate the structural and daylighting performance of the design, and record the results. Since this research examines an existing case study with actual buttress dimensions, the samples were not chosen randomly or even to represent the entire design space. Rather, while changing one design variable, the other variables were held constant at the dimensions found on the actual cathedral. This single variable sampling technique tests the performance sensitivity of each variable separately. Although the resulting geometry may become less reasonable as samples move towards the edges of the design space, the slopes and ranges of the performance results provide insight into how cathedral performance responds to geometric changes applied to the existing pier. For each variable, 11 design samples were taken at even intervals spanning the variable ranges given in Table 13.

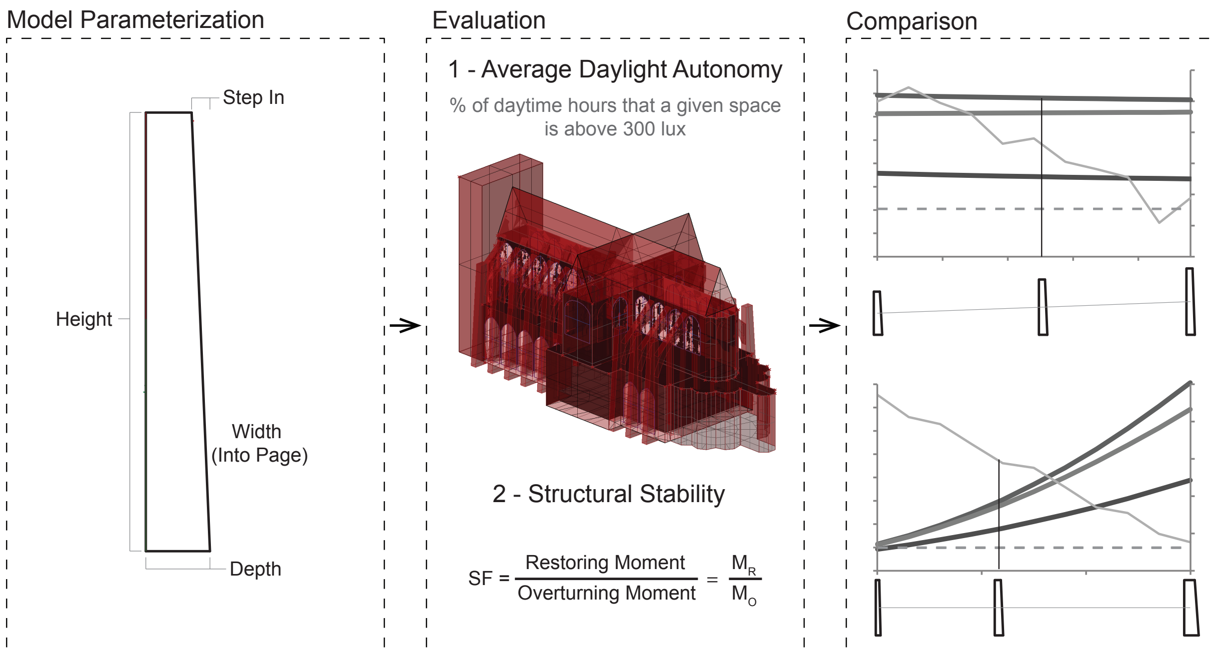

Figure 27: Methodology overview for structure + daylighting analysis

Parameter Ranges			
Pier Parameter	**Low (m)**	**Actual (m)**	**High (m)**
Depth	3.05	5.33	9.14
Height	27.34	35.05	42.06
Step-In	0	2.47	4.57
Width	0.61	2.10	4.88

Table 13: Ranges and default values for the parametric pier model

6.2.6 Model development

The next step was to develop a 3D daylighting analysis model of the Amiens Cathedral, which was done in Rhinoceros. This model was based on dimensions of the cathedral as well as historical drawings of the plan and section. Since daylighting models are computationally intensive and model complexity increases the time required for each analysis, the geometry was simplified to represent only the main surfaces of the cathedral. The model geometry is formed by a collection of surfaces and volumes representing the walls, windows, floor, and ceiling of the cathedral. The tower, front façade, first level, and clerestory level are modeled as a group of rectangular boxes. Windows with locations corresponding to those found on the cathedral are carved out of the walls. The floor of the cathedral was assumed to be the location where daylighting matters most for occupants, and so a test mesh consisting of 116 analysis nodes was created in a regular grid.

Along with the base model, a parametric pier volume generated by the four main design variables was set up using Grasshopper. Inside the Grasshopper script, the pier volumes were arranged to correspond to the actual pier locations of the Amiens cathedral. Thus, as different parametric samples are generated, the pier size changes for every instance in the daylighting model. In addition to the shape of the pier, the flyer geometries also change with the 'height' variable, since higher and lower piers change where these flyers are attached. Initial flyer and inner column dimensions were based on the shape of the existing cathedral, meaning that the flyers on the front and rear of the church are different. This daylighting analysis does not include the ambulatory and seven radiating chapels found behind the altar, instead focusing only on the main space of the church. This means that the piers and windows are only present if they have a direct effect on the analysis mesh located in the main interior region of the cathedral.

Once the geometry of the cathedral was established, a daylighting analysis of the interior space was conducted using DIVA-for-Rhino (Solemma). DIVA is a daylighting and building energy modeling plug-in developed for the Rhino NURBS modeler. In this particular model, DIVA first reads in all of the geometry and assigns a radiance material to each surface. It then establishes a test node at each point of a user-defined grid. Finally, it uses a climate file, assumed schedule, and lighting requirements to calculate the average daylight autonomy of each test node throughout the course of a year. Daylight autonomy (DA) is one of the annual dynamic daylight metrics proposed by Reinhart et al. (2006), and it is represented as the percentage of annual daytime hours that a given point in a space is above a specified illumination level. An average of this metric taken across all analysis nodes will be used as the daylighting performance measurement for the parametric model.

The daylighting model required a number of important assumptions. The piers and walls, made of stone but in some places ornately carved in a way that would diffuse light, were assumed to have radiance properties similar to contemporary exterior masonry facades. For the windows, it is difficult to determine the historical transmissivity despite considerable research into the topic, since few examples of original glass currently exist and windows contained stained glass images full of different colors. After consultation of existing research into the glass properties of cathedrals, a conservative transmissivity of 0.47 (the lowest default value offered by DIVA) was used. This transmissivity likely represents an upper bound for how much light was historically allowed into the cathedral since it is based on a modern window.

A range of required thresholds for interior illuminance levels could also be applied to a historical cathedral. In contemporary architecture, it is recommended to keep a home, warehouse, or theater at 150 lux, while offices, libraries, and classrooms require in between 250-500 lux. Religious services may include some amount of reading, by either the congregants (in historical cases where they were literate) or the clergy, but these reading tasks are likely not as visually demanding as a typical office job. However, due to light's religious significance, the designers likely wanted a more visually illuminated space than a modern theater. Taking all of this into account, the model considers daily use from 8am to 6pm with no artificial lighting and a required daylighting threshold of 300 lux. The climate file for Paris was used because it was the closest available file to Amiens.

When modeling a historical structure, there are difficulties in finding accurate and precise thresholds and measurements for performance. There is no building code to rely on, and much of the original quality of the building has long since been modified. Nevertheless, there is still value in this level of analysis, as the broad relationships and trade-offs between performance objectives remain, as do their relationships to the geometric properties of the building.

6.3 Results and discussion

The results of both the structural and daylighting analyses are provided in this section. Figure 28 shows the pier safety factor (SF) against failure for three load cases and three failure mode assumptions. The safety factor under dead load for the initial whole-pier assumption case is 3.5, which is within 5% of Mark's calculation of 3.7. When considering less favorable failure cases, the safety factor drops to 2.5 for failure with breakage and to 1.5 for the development of tension in the pier. The safety factor also lowers as horizontal wind loads are added to the pier, both through the buttresses and also directly acting on the pier itself. Thus, although the structure is fairly safe against collapse under dead load thrust, the live loads merit further investigation.

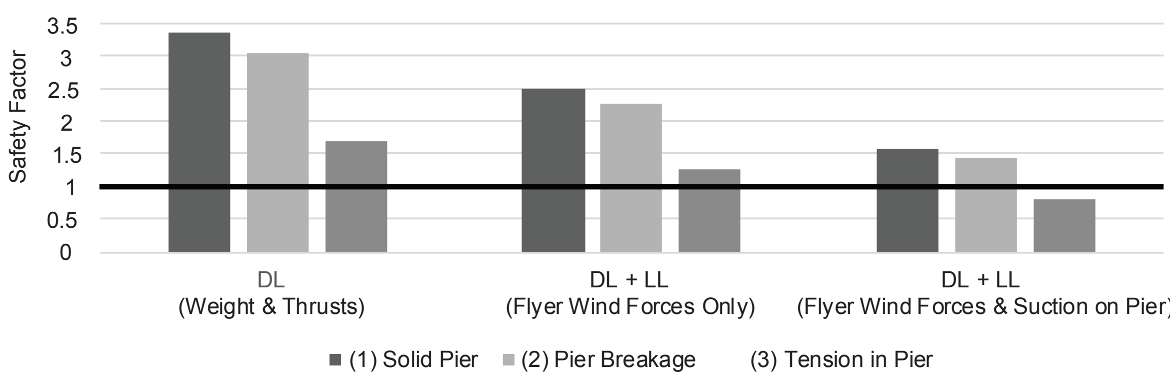

Figure 28: Safety factor calculation results for each loading and failure assumption

It is clear from the graph that assumptions made about the structure's response to wind loading have a dramatic effect on the safety factor. If just the wind loads on the clerestory level transferred through the flying buttresses are considered, the safety factor only slightly changes from the dead load case. However, when leeward suction on the both the pier itself and the lower wall is considered, the safety factors drop dramatically, even below 1.0 for the development of tensile forces in the pier. This result is somewhat surprising given that the magnitude of gravity loads acting on the structural system is over 18 times higher than the magnitude of total wind loads (roughly 2.8 million kilograms versus 150,000 kilograms per bay). A more sophisticated analysis is necessary to accurately account for any rigid frame behavior across the entire cathedral section. Still, when isolating the pier, these results indicate that during very large gusts, there is the possibility of tension developing on the inner pier edge on the leeward side of the cathedral.

The initial safety factor results indicate the overall stability for the existing cathedral dimensions, but the second part of the analysis compares this safety factor with daylight autonomy as the dimensions of the pier change. Figure 29 shows each of the sampled pier design variables plotted against safety factor on the left vertical axis and daylight autonomy on the right vertical axis. All three sets of safety factor assumptions are shown for the dead load case in solid lines, while the results of the individual daylighting simulations are indicated by dots, in most cases forming a fairly smooth curve across the 11 samples.

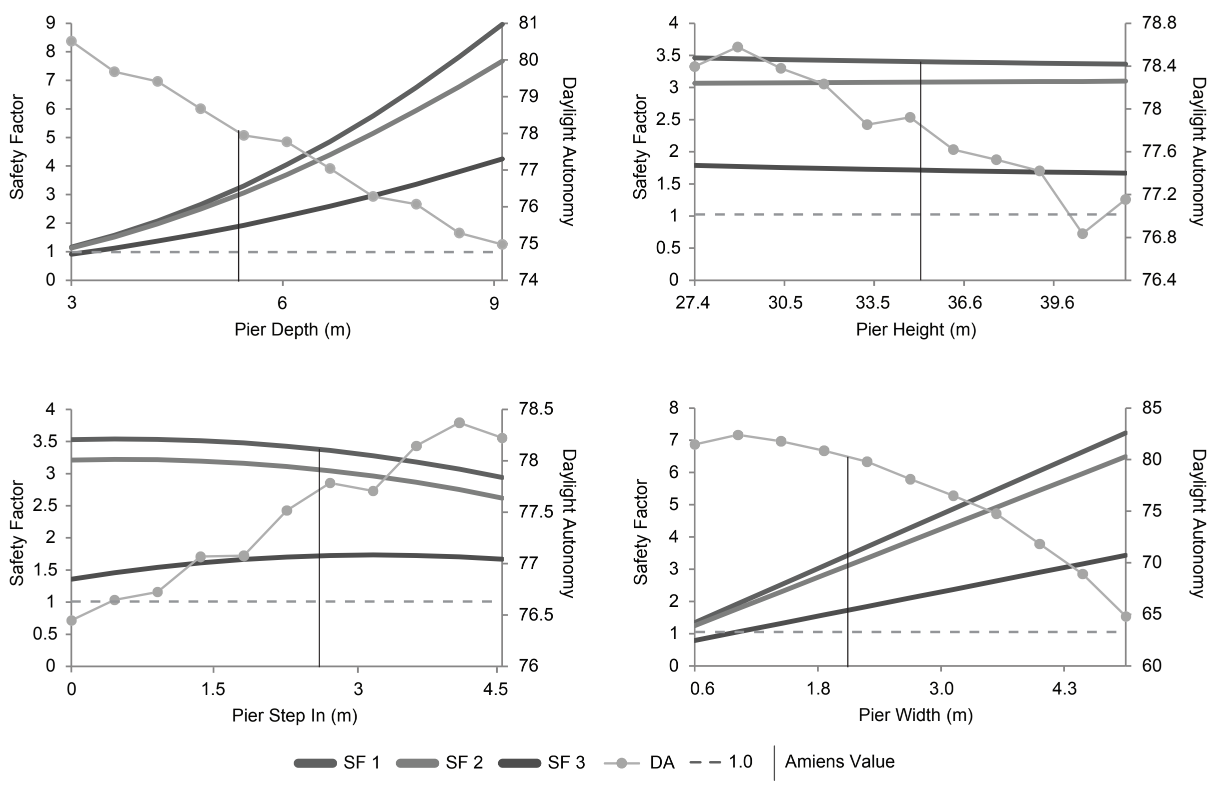

Figure 29: Change in safety factor and daylight autonomy in response to different variable changes

For each of the four graphs, there is a trade-off when the safety factor and daylight autonomy lines have opposite slopes. These results align with the broad assumption that more structural mass leads to better stability but less useful interior daylight. However, the nature of this trade-off is different for each of the variables. As the piers are made deeper, their stability grows exponentially, but the daylight autonomy drops linearly. Pier height has only a slight effect on stability and daylight, with both dropping as height is increased. This means that the additional weight of the taller pier essentially compensates for the additional height at which thrusts and wind loads are applied to it. A taller pier would likely result in a dramatically different flyer design, since the thrusts would still have to be transferred from the base of the upper vault to a new point on the pier, but the differences in weight of this new flyer would not substantially affect the properties of the pier itself.

The slopes of the safety factor lines come to a local maximum as the step-in is changed, suggesting that there may be an optimal pier shape. The location of this maximum changes for the different assumptions, since the step-in determines the load line of the pinnacle, which can contribute to either the restoring or failing moment in the case of the development of tensile forces. Daylight autonomy increases slightly with higher step-in as some of the solid material of the pier is removed, allowing more light to access the interior. Safety factor increases linearly as the pier becomes wider, and daylight decreases more substantially across the range of sampled widths than for any other variable.

In addition to these trends, it is also important to compare the overall ranges and slope magnitudes for the graphs, which indicate the extent to which each design variable is able to influence both performance factors. Figure 30 gives a better indication of these ranges by comparing the safety factor versus daylight autonomy results for each of the four variables in the dead load case. This plot shows only performance rather than variables, and each dot represents one design sample in which the indicated variable has been changed. When read together as a chain, the broad trend of these data points shows the sensitivity of performance response to the different variable changes. For example, by changing only the depth or the width, it is possible to quickly increase the safety factor, up to a value of over 7.0 within the variable range selected for this problem. In contrast, the safety factor changes very slowly with height and step-in, ranging from only 3-3.75 even at the most extreme samples. On the daylighting side, increasing the pier width had the largest effect, with the widest piers and consequently smallest windows only allowing for 65% daylight autonomy compared to over 80% for the narrowest piers. Each of the other three variables yielded smaller ranges, from about 77-80% daylighting autonomy.

Although it is difficult to know precisely the design influences on Gothic cathedrals, it is likely that their builders understood these trade-offs qualitatively. One of the most significant formal aspects of Gothic piers is their depth/width aspect ratio, in which piers are allowed a larger depth so that width can be decreased for the sake of larger windows. The heights of Gothic piers probably had more to do with the configuration of the flyers and the direction of their forces than overall pier stability, but the pinnacle represents an opportunity to pile more weight (and height) on the pier in the interest of a higher safety factor. The presence of step-ins on cathedral designs show that builders understood more pier mass needed to be concentrated closer to the walls and further away from the exterior overturning point. Thus, if pier material were to be removed to allow additional light into the interior, it could be removed in steps on the outside edge. All of these measured trade-offs correspond with a way in which Gothic designers may have been influenced in their dimensional decisions.

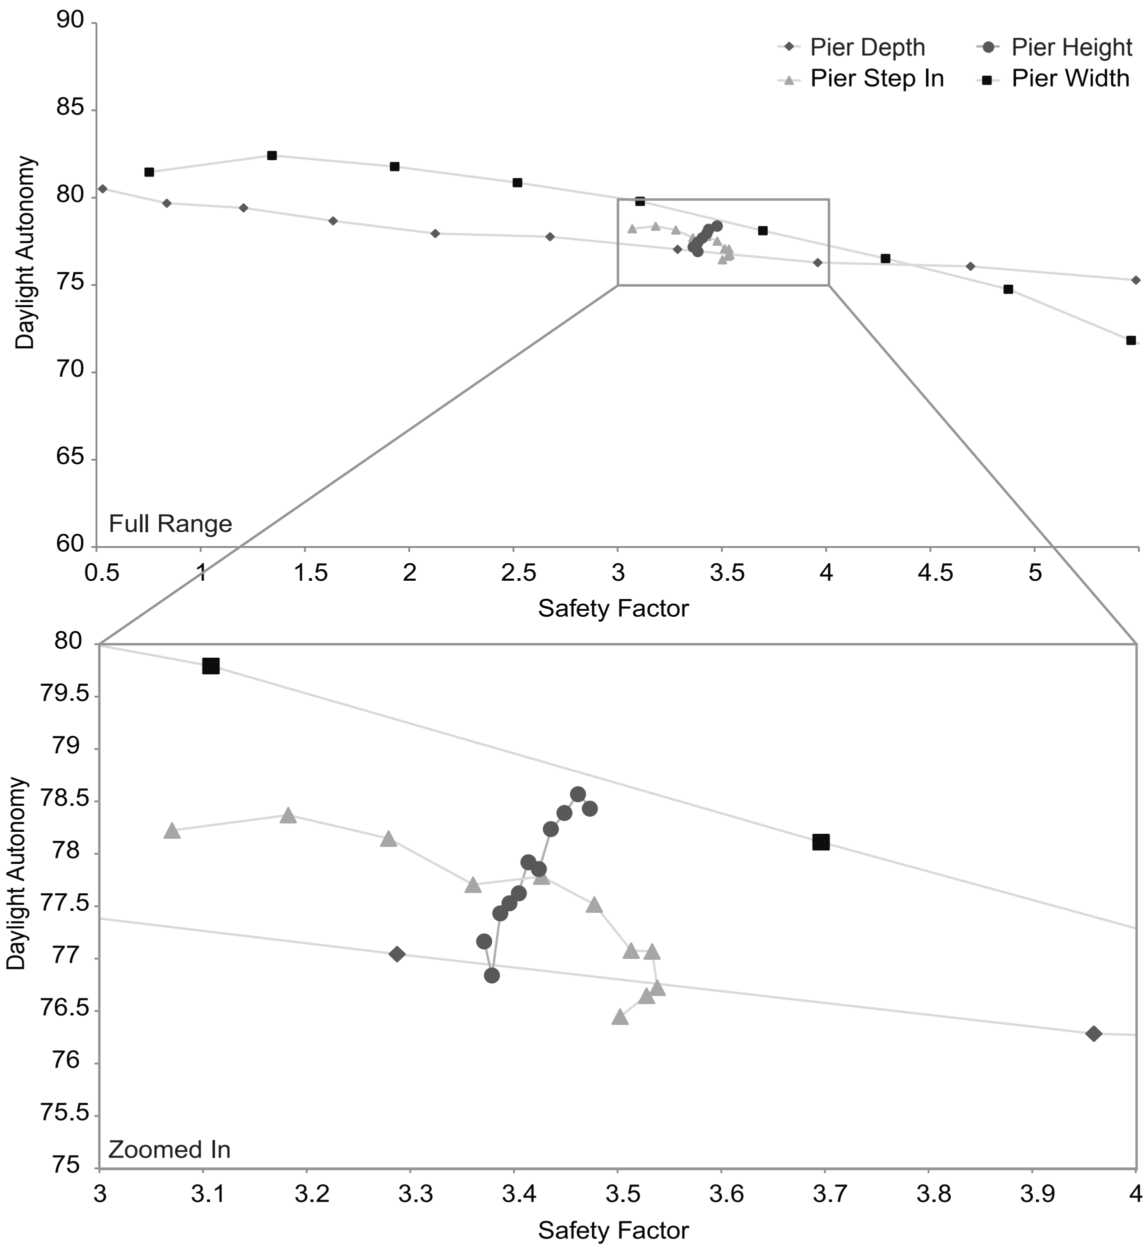

Figure 30: Ranges and slopes for safety factor versus daylight autonomy in response to different variable changes (dead load)

As a standalone comparison, it is also useful to contrast the current flying buttress design with one that contains full pier walls to numerically test the conclusion that the less massive form of flying buttresses allowed for more interior daylight than older full wall piers. Figure 31 shows the results of this comparison, indicating that the flying buttress design allows in substantially more daylight. Although this comparison is not perfect in that the precursors to Gothic designs also had smaller windows and other significant differences, it bolsters the common claim of the design implications of the flying buttress, and it has direct implications for other cathedrals such as the Cathedral at Mallorca that have much fuller pier buttresses.

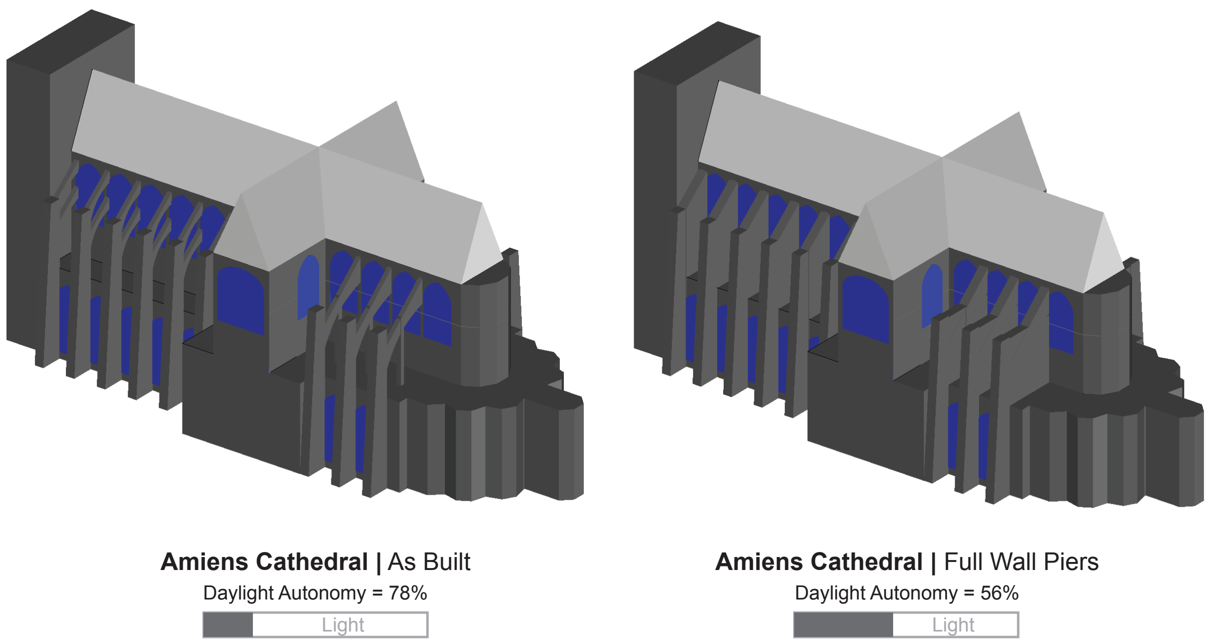

Figure 31: A comparison of daylight autonomy for flying buttresses and a full pier wall

6.4 Conclusions and further work

This chapter has presented both a structural and comparative daylighting analysis of the main buttressing system of the Amiens Cathedral. The structural study calculated safety factors against a number of possible failure modes while exploring a variety of load cases and their influences on the safety factor result. The daylighting analysis developed a parametric model of the pier buttress and determined the impact of changing each design parameter on the structural and daylighting performance of the cathedral. These simultaneous analyses each serve to provide a better understanding of the design goals and achievements of Gothic builders, and together give insight into how the trade-offs between structural stability and interior daylight were navigated. However, since only one case study was considered, this research scratches the surface of what could be learned through more extensive quantitative analysis of the general relationship between Gothic structural form and usable daylight.

There are a number of possibilities for future work, both within this specific case study and by extending the methodology to other cathedrals. First, the evaluation methods of the comparative analysis could be made more accurate. In the structural study, the precise behavior of the pier under live wind loading has a significant effect on the overall safety factor. Much could be learned through a more extensive analysis of the cathedral's overall response to wind, which would include additional engineering calculations and more elaborate modeling. Similarly, the daylighting analysis relied on key assumptions concerning the materials present in the cathedral and their behavior concerning light transmittance, reflectivity, and color. A more comprehensive research effort would include physical testing of the cathedral's actual material properties as well as comparisons to other historical materials tests, since many aspects of the cathedral have been modified. Each value obtained through further testing would add a degree of accuracy to the model.

In addition to improving the evaluation methods, more expansive study and sampling of the design space for a Gothic buttressing system could yield interesting results. Only a single variable at a time was modified in this model, partly because a structure with existing dimensions was being analyzed, and partly due to the computational demands of the daylighting model. However, an actual Gothic designer had a much less restricted design space for the buttress, as he could have made the pier taller, wider, deeper, or more stepped at the same time. Furthermore, the final dimensions of a medieval design may have been decided upon based on historical precedents, geometry, aesthetics, or even ancient and biblical proportions rather than empirical studies of stability. Additional sampling of the potential design space, as well as research into influential historical texts, could be completed in the future. Each of these methodological additions could be tested on the Amiens Cathedral or extended to other case studies to tell a richer story of interplay between structure and light, and possibly show a historical progression of design ideas or clear formal differences based on geography or other significant factors.

Nevertheless, this research presents a new method of examining the performance of Gothic cathedrals based on their two most significant design objectives. It does this by implementing a multi-objective methodology for quantitatively addressing a common hypothesis of architectural historians, which is that the design of the Gothic flying buttress system developed in response to a desire for more height and more light. This methodology and its results, combined with the standalone analysis of the structural stability of the cathedral piers, add to existing knowledge concerning the design and behavior of one of history's most enigmatic architectural typologies.

7 Conclusions

In conclusion, this thesis has demonstrated that multi-objective optimization techniques can be flexible and interactive for creative designers to use during brainstorming in the conceptual design of structures. At the same time, MOO can guide design space exploration in useful and sometimes unexpected ways and generate feedback with enough accuracy to push towards high-performance design outcomes.

7.1 Summary of contributions

There are a number of specific research contributions presented in this thesis. The first is guidance on how to effectively formulate and navigate an architectural design space using multi-objective methods and interactivity. In many published MOO case studies, the design space is presented but rarely explained thoroughly. No current roadmap exists to guide designers towards selecting parameterizations, variables, and even objectives in a way that leads to desired results. Similarly, certain methods within the larger umbrella of MOO may be more or less suited to specific design problems. A general discussion of the advantages and usability of various methods, as well as demonstrative case studies, are presented to assist architectural designers in the application of MOO.

This thesis also introduces a conceptual map and the development of design tools built to support MOO functionality within parametric design software. These tools allow for the direct integration of MOO into conceptual design workflows to guide users towards high-performing, preference-infused solutions. They also facilitate a variety of organization and visualization techniques that support decision-making for managing performance tradeoffs in conceptual design.

The third contribution is new data concerning the relationship between multi-objective performance feedback, optimization tools, and conceptual design outcomes. This data, collected through a study of how designers interact with a series of performance-based conceptual design tools, shows that feedback and guidance can help designers improve on performance metrics that are not intuitive to non-experts, without sacrificing the ability to prioritize different objectives or explore diverse design spaces. In design problems involving both significant geometric freedom and specific structural requirements, the presence of performance feedback can dramatically alter the process and outcome of form generation. This drastic effect on form and performance was demonstrated by the over 50% reduction in carbon emissions when study participants were provided with either real-time feedback or optimization tools compared to free exploration.

The fourth outcome is the demonstration of how an optimization methodology for multiple objectives can be used to generate a geometrically diverse range of high-performing designs. Many existing MOO case studies are constrained to regular, rectangular geometries. With parametric modeling, contemporary computational tools that generate rapid performance feedback, and software that seamlessly links the two, a MOO process can be integrated into a typical conceptual design workflow for even the most expressive designers. As future theoretical advances are made and software develops, the ease of using MOO interactively and flexibly will only increase. At the same time, this thesis demonstrates that a human must currently prioritize design goals and consider visual impact to arrive at a satisfying design, which could make MOO more attractive for creative designers than deterministic optimization methods.

Fifth, through case studies that explore generalizable architectural problems such as arch shapes and static overhangs, this thesis presents a number of broad lessons about the relationship between geometry and performance. Especially in structurally demanding building programs, there are key tradeoffs between structural efficiency, energy efficiency, and other aspects of performance that can change with building geometry. In some cases, these results are intuitive, but when climate, orientation, or other building characteristics are modified, the simultaneous consideration of multiple design objectives can push "optimal" building forms in unexpected ways. Knowledge of these tradeoffs can have direct impact on future designs.

Finally, this thesis demonstrates how multi-objective thinking can be applied in historical analysis. While studying how different pier shapes affect both structural stability and daylight availability in a Gothic cathedral, this research assesses the quality of Gothic designs while also providing insight into how designers of old structures may have synthesized multiple inputs and requirements into successful designs.

7.2 Potential impact

The contributions made in this thesis respond to the limitations of current conceptual design software and the potential of multi-objective optimization techniques to overcome some of these limitations. Thus, much of the possible impact is assumed to be on architectural and structural practice. However, some contributions could also be useful for academic purposes in both teaching and research.

In practice, architects, structural engineers, and other building specialists involved in conceptual design could make use of these contributions by setting up their own parametric models and implementing the multi-objective optimization and visualization strategies offered in this thesis. In the case of signature, formally expressive structures in which the designers have freedom and geometric flexibility, they could use MOO techniques for performance-based design space exploration at the building scale. In more traditional buildings, these techniques could still be used immediately for smaller architectural elements, such as the canopy described in Chapter 4. When many people are involved in the conceptual design phase, these methods could be used to facilitate collaborative discussion and support team-based decision-making. Finally, the tradeoffs revealed in this thesis could be of direct use to designers working on architectural typologies that contain geometries similar to the case studies.

In an academic setting, the multi-objective optimization techniques demonstrated in this thesis could be used to help students learn how to synthesizing multiple feedback streams and develop intuition about relationships between form and performance. On the research side, this thesis provides supporting evidence for the further development of performance-based conceptual design tools, adds information to the ongoing conversation concerning embodied and operational energy, and offers a new multi-criteria method for analyzing historical structures, which could all be utilized and applied in research projects moving forward.

7.3 Future work

Although multi-objective optimization techniques has been moving from pure engineering applications into architectural design problems, there are still significant opportunities to expand MOO methods, demonstrate their effectiveness in performance-based conceptual design, and improve their accessibility and usability among more creative architects. As such, future work can include contributions in three main areas related to the integration of performance feedback into the conceptual design process: theory, implementation, and development. The future work described here would still be focused on creative, expressive architectural design and non-traditional building geometries while concurrently using optimization to capture interactions between big picture building performance objectives. Especially when compared to traditional structural optimization, this holistic view of design that allows both performance considerations from different disciplines and non-quantifiable preferences to affect best solutions would build on the research in this thesis and continue to strengthen the link between architectural design and building technology. Although some of the previous chapters have mentioned more directed future work, the next sections take a broad view and consolidate this suggested further research.

7.3.1 Theory

The first area of potential future work, theory, addresses the identification of fundamental characteristics and features of computational, multi-objective design methods and workflows for early stage design. There is a significant need at the front end of any MOO workflow to characterize variables, objectives, and constraints in a way that assists designers in formulating the design problem itself. When emphasizing the capabilities of MOO techniques to generate a wide range of high-performance designs, another need is to more fully understand the concept of diversity. There is currently no single established method for computationally measuring diversity, and no explicit understanding of how greater optimization output diversity leads to better architectural outcomes. In addition, theoretical contributions could be made by modifying post-Pareto analysis and other techniques for decision-making in fuzzy multi-objective problems to more adequately match the needs of creative, exploratory, interactive architectural design.

7.3.2 Implementation

The second area of future work, implementation, could involve employing MOO on additional architectural case studies to further demonstrate its potential and generate new, generalizable design knowledge about structural tradeoffs in architecture. Although some research in this area already exists, future work could extend MOO techniques to non-traditional building geometries and emphasize diversity of solutions, which makes them more suitable to creative, interactive design. Due to their significance to contemporary sustainable design, additional case studies of the tradeoffs between structure and operational energy could be implemented for different building typologies, structural systems, and non-geometric energy variables. There are many other design tradeoffs that could be explored in this context, including between structural efficiency and thermal mass, daylighting performance, or aspects of constructability. Other multi-objective performance problems at smaller than building scale could also be investigated, including thermal bridging in structural connections and façade details. Finally, the addition of cost as an objective in future MOO studies could prove immensely valuable to the field, although the difficulty of obtaining accurate cost information in an academic setting should be managed by adding uncertainty to the analysis.

7.3.3 Development

Additional future work could be concentrated on developing computational tools that enable MOO for architectural design. As mentioned in Chapter 3, some such programs for early design already exist. This thesis also develops a conceptual workflow for how a future set of tools (some of which have already been created by the author and collaborators) could be configured to give maximum flexibility and usability. However, there is still no best, established method for implementing a MOO workflow that is seamlessly integrated within typical design software, gives the user total flexibility for different modes of interaction, and does not require complicated translation of model information for data visualization and choosing between design options. Future work could address this need by developing the additional components needed to complete the proposed workflow, or by offering new strategies altogether. These tools could be bolstered by additional user experience testing to determine which approaches work best for creative design and measure the magnitude of performance improvements made possible through MOO when compared to typical, performance-free design exploration.

7.4 Concluding remarks

Together, the contributions in this thesis have the potential to encourage more widespread and effective use of optimization in sustainable conceptual design. Although the architecture, engineering and construction industries often require considerable time to make the transition to new technologies, many researchers and practitioners have already worked towards integrated design and analysis environments, while adopting such practices as parametric modeling. Increased accessibility to multi-objective optimization for designers, as well as the demonstration of its usefulness and applicability to the conceptual design of structures, can accelerate the adoption of these advanced design technologies. Acceptance and implementation of such methods can lead to future buildings that are high performing and well constructed, without sacrificing an appreciation for creative architectural design.

8 Publications and Presentations Related to Thesis

Brown, N. & Mueller, C., 2015. Optimization for structural performance and energy usage in conceptual building design. Presented at *EMI 2015: Engineering Mechanics Institute Conference*, Stanford University: ASCE.

Brown, N., Tseranidis, S. & Mueller, C., 2015. Multi-objective optimization for diversity and performance in conceptual structural design. In *Proceedings of the International Association for Shell and Spatial Structures (IASS) Symposium 2015: Future Visions*. Amsterdam: IASS.

Brown N. & Mueller C., 2016. Early-stage integration of architectural and structural performance in a parametric multi-objective design tool. In *Proceedings of the International Conference on Structures and Architecture,* Guimarães, Portugal.

Brown, N. & Mueller, C., 2016. The effect of performance feedback and optimization on the conceptual design process. In *Proceedings of the International Association for Shell and Spatial Structures (IASS) Symposium 2016*. Tokyo: IASS (abstract accepted).

Brown N. & Mueller C., 2016. Design for structural and energy performance of long span buildings using geometric multi-objective optimization (in preparation).

Tseranidis S., Brown N., & Mueller C., 2016. Data-driven approximation algorithms for rapid performance evaluation and optimization of civil structures. *Automation in Construction*. Available at: http://dx.doi.org/10.1016/j.autcon.2016.02.002.

9 References

ACADIA, 2015. Call for Submissions. In *Computational Ecologies: Design in the Anthropocene*. Cincinnati, OH.

Allen, E. & Zalewski, W., 2010. *Form and Forces*, Hoboken, NJ: John Wiley & Sons.

Arnaud, V.B., 2013. *Quantifying Architects ' and Engineers ' Use of Structural Design Software*. M.Eng. Thesis, Massachusetts Institute of Technology.

Arora, J.S., 2004. *Introduction to Optimum Design* 2nd ed., San Diego, CA: Elsevier Academic Press.

ASHRAE, 2013. *ANSI/ASHRAE STANDARD 62.1-2013. Ventilation for acceptable indoor air quality*, Atlanta, GA.

ASHRAE, 2004. *ANSI/ASHRAE/IESNA Standard 90.1-2004. Energy Standard for Buildings Except Low-Rise Residential Buildings*, Atlanta, GA.

Asl, M.R. et al., 2014. BIM-based parametric building energy performance multi-objective optimization. *Education and Research in Computer Aided Architectural Design in Europe 32*, 224, pp.1–10.

Autodesk, Dynamo. Available at: http://dynamobim.org/download/.

Autodesk, 2016a. React Structures. Available at: http://react.autodesk.com/.

Autodesk, 2016b. Revit. Available at: http://www.autodesk.com/products/revit-family/overview.

Autodesk, 2016c. Robot Structural Analysis Professional. Available at: http://www.autodesk.com/products/robot-structural-analysis/overview.

Baker, N., Fanchiotti, A., & Steemers, K., 1993. *Daylighting in Architecture: A European Reference Book*, Commission of the European Communities Directorate-General XII for Science, Research, and Development.

Bayer, C. et al., 2010. *AIA Guide to Building Life Cycle Assessment in Practice*, Washington, D.C.: The American Institute of Architects.

Beghini, L.L. et al., 2013. Structural optimization using graphic statics. *Structural and Multidisciplinary Optimization*, 49(3), pp.351–366. Available at: http://link.springer.com/10.1007/s00158-013-1002-x.

Bentley Systems Inc., 2016. Generative Components.

Block, P., Ciblac, T. & Ochsendorf, J., 2006. Real-time limit analysis of vaulted masonry buildings. *Computers and Structures*, 84, pp.1841–1852.

Boeing, A. & Bräunl, T., 2007. Evaluation of real-time physics simulation systems. In *Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and Southeast Asia*. pp. 281–288.

Bork, R., Mark, R. & Murray, S., 1997. The openwork flying buttresses of Amiens Cathedral: "Postmodern Gothic" and the limits of structural rationalism. *Journal of the Society of Architectural Historians*, 56(4), pp.478– 493.

Bourgeois, D., Reinhart, C. & Macdonald, I., 2006. Adding advanced behavioural models in whole building energy simulation: A study on the total energy impact of manual and automated lighting control. *Energy and Buildings*, 38(7), pp.814–823.

Brown, N. et al., 2016. Early-stage integration of architectural and structural performance in a parametric multi-objective design tool. In *International Conference on Structures and Architecture*. Guimarães, Portugal.

Brown, N. & Mueller, C.T., 2016. The effect of performance feedback and optimization on the conceptual design process. In *Proceedings of the International Association for Shell and Spatial Structures Symposium 2016*. Tokyo: IASS (abstract accepted).

Brown, N., Tseranidis, S. & Mueller, C., 2015. Multi-objective optimization for diversity and performance in conceptual structural design. In *Proceedings of the International Association for Shell and Spatial Structures Symposium 2015: Future Visions*. Amsterdam: IASS.

von Buelow, P., 2012. Paragen: performative exploration of generative systems. *Journal of the International Association For Shell and Spatial Structures*, 53(4), pp.271–284.

Caldas, L.G. & Norford, L.K., 2003. Genetic algorithms for optimization of building envelopes and the design and control of HVAC systems. *Journal of Solar Energy Engineering*, 125(3), pp.343–351. Available at: http://solarenergyengineering.asmedigitalcollection.asme.org/article.aspx?articleid=1456824.

Céa, J. et al., 2000. The shape and topological optimizations connection. *Computer Methods in Applied Mechanics and Engineering*, 188(4), pp.713–726. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0045782599003576.

Clune, R.P., 2010. *Explorative Structural Design*. S.M. Thesis, Massachusetts Institute of Technology.

Coccia, S., Como, M. & Di Carlo, F., 2015. Wind strength of Gothic cathedrals. *Engineering Failure Analysis*, 55, pp.1–25. Available at: http://dx.doi.org/10.1016/j.engfailanal.2015.04.019.

Coley, D.A. & Schukat, S., 2002. Low-energy design: combining computer-based optimisation and human judgement. *Building and Environment*, 37(12), pp.1241–1247. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0360132301001068.

Courtenay, L.T. ed., 1997. The Engineering of Medieval Cathedrals. *Studies in the History of Civil Engineering*, pp.xii–xxxiv.

Crawley, D.B. et al., 2001. EnergyPlus: Creating a new-generation building energy simulation program. *Energy and Buildings*, 33, pp.319–331.

Cross, N., 2004. Expertise in design : an overview. *Design Studies*, 25(5), pp.427–441.

Cross, N., Naughton, J. & Walker, D., 1981. Design method and scientific method. *Design Studies*, 2(4), pp.195–201.

Cvetkovic, D. & Parmee, I.C., 2002. Preferences and their application in evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 6(1), pp.42–57. Available at: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=985691.

Danhaive, R.A., 2015. *Integrating Interactive Evolutionary Exploration and Parametric Structural Design*. M.Eng. Thesis, Massachusetts Institute of Technology.

De Wolf, C., 2014. *Material quantities in building structures and their environmental impact*. S.M. Thesis, Massachusetts Institute of Technology.

Deb, K. et al., 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), pp.182–197.

Deb, K., 2001. *Multi-Objective Optimization Using Evolutionary Algorithms*, New York: John Wiley & Sons.

Deb, K. & Gupta, S., 2011. Understanding knee points in bicriteria problems and their implications as preferred solution principles. *Engineering Optimization*, 43(11).

Dehio, G. & Bezold, G., 1892. *Die kirchliche Baukunst des Abendlandes*, Stuttgart: J.G. Cotta.

Diakaki, C., Grigoroudis, E. & Kolokotsa, D., 2008. Towards a multi-objective optimization approach for improving energy efficiency in buildings. *Energy and Buildings*, 40, pp.1747–1754.

Dogan, T., ArchSim. Available at: www.archsim.com.

Dogan, T. & Reinhart, C., 2013. Automated conversion of architectural massing models into thermal "shoebox" models. In *Proceedings of BS2013: 13th Conference of International Building Performance Simulation Association*. Chambery, France: IBPSA, pp. 3745–3752.

Dorsey, J. et al., 2007. The mental canvas: A tool for conceptual architectural design and analysis. In *15th Pacific Conference on Computer Graphics and Applications (PG'07)*. Maui, HI: IEEE, pp. 201–210. Available at: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4392730.

Dr. Software, 2009. Dr. Frame3D. Available at: http://www.drsoftware-home.com/df3d.html.

Durillo, J.J. & Nebro, A.J., 2011. jMetal: a Java framework for multi-objective optimization. *Advances in Engineering Software*, 42, pp.760–771.

Faas, D. et al., 2014. The influence of immersion and presence in early stage engineering designing and building. *Artifcial Intelligence for Engineering Design, Analysis and Manufacturing*, (Special Issue 02), pp.139–151.

Fialho, A., Hamadi, Y. & Schoenauer, M., 2011. Optimizing architectural and structural aspects of buildings towards higher energy efficiency. In *Proceedings of the 13th annual conference companion on Genetic and evolutionary computation*. pp. 727–732.

Finger, S. & Dixon, J.R., 1989. A review of research in mechanical engineering design. part I: descriptive, prescriptive, and computer-based models of design processes. *Research in Engineering Design*, 1, pp.51–67.

Fitchen, J.F., 1955. A comment on the function of the upper flying buttress in French Gothic architecture. *Gazette des Beaux-Arts*, 6th(45), pp.69–90.

Flager, F., Welle, B., et al., 2009. Multidisciplinary process integration and design optimization of a classroom building. *Electronic Journal of Information Technology in Construction*, 14, pp.595–612.

Flager, F., Adya, A. & Haymaker, J., 2009. AEC Multidisciplinary Design Optimization: impact of high performance computing. *CIFE Technical Report #TR175*, (December).

Flöry, S., Goat. Available at: http://www.rechenraum.com/goat/download.html.

Goldschmidt, G., 1994. On visual design thinking: the vis kids of architecture. *Design Studies*, 15(2), pp.158–174.

GRAITEC, 2016. Advance Design: Professional Finite Element Structural Analysis Software.

Hammond, G. & Jones, C., 2010. *Inventory of Carbon & Energy (ICE), Version 1.6a*. University of Bath.

Heyman, J., 1966. The stone skeleton. *International Journal of Solids and Structures*, 2, pp.249–279.

Heyman, J., 1997. *The Stone Skeleton: Structural Engineering of Masonry Architecture*, Cambridge, England: Cambridge University Press.

Holzer, D., Hough, R. & Burry, M., 2008. Parametric design and structural optimisation for early design exploration. *International Journal of Architectural Computing*, 05(04), pp.625–644.

Huerta, S., 2009. The debate about the structural behaviour of gothic vaults: from Viollet-le-Duc to Heyman. *Construction History*, (May), pp.837–844.

International Code Council, 2012. *International Energy Conservation Code*, Section C402 Building Envelope Requirements. Available at: http://publicecodes.cyberregs.com/icod/iecc/.

International Energy Agency, 2013. *Transition to Sustainable Buildings*, Paris, France: Energy Technology Perspectives Series.

James J. Hirsch & Associates & Lawrence Berkeley National Laboratory, eQuest. Available at: http://www.doe2.com/equest/.

Kaethner, S. & Burridge, J., 2012. Embodied CO2 of structural frames. *The Structural Engineer*, pp.33–40.

Khajehpour, G. & Grierson, D.E., 2002. Method for conceptual design applied to office buildings. *Journal of Computing in Civil Engineering*, 6(2), pp.83–103.

Khajehpour, S. & Grierson, D.E., 2003. Profitability versus safety of high-rise office buildings. *Structural and Multidisciplinary Optimization*, 25(4), pp.279–293.

Kottek, M. et al., 2006. World Map of the Köppen-Geiger climate classification updated. *Precipitation Climatology Wetterdienst, Deutscher*, 15(3), pp.259–263.

Krem, M. et al., 2013. Structural configuration and building energy performance. *Journal of Architectural Engineering*, 19(1), pp.29–40.

Machwe, A.T., Parmee, I.C. & Miles, J.C., 2005. Integrating aesthetic criteria with a user-centric evolutionary system via a component-based design representation. In *International Conference on Engineering Design*. Melbourne, pp. 1–15.

Magnier, L. & Haghighat, F., 2010. Multiobjective optimization of building design using TRNSYS simulations, genetic algorithm, and Artificial Neural Network. *Building and Environment*, 45(3), pp.739–746. Available at: http://dx.doi.org/10.1016/j.buildenv.2009.08.016.

MaP3, 2014. In Progress: The Qingdaobei Station / AREP Architect + MaP3 Structural Engineering. *arch daily*. Available at: http://www.archdaily.com/439741/work-in-progress-the-qingdaobei-station-arep-architect-map3-structural-engineering.

Mark, R., 1984. *Experiments in Gothic Structure*, Cambridge, MA: The MIT Press.

Mark, R., 1977. Robert Willis, Viollet-le-Duc and the structural approach to Gothic architecture. *Architectura*, 7(2), pp.52–64.

Marks, W., 1997. Multicriteria Optimisation Energy-Saving Buildings. *Building and Environment*, 32(4), pp.331–339.

Marler, R.T. & Arora, J.S., 2004. Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26(6), pp.369–395. Available at: http://link.springer.com/10.1007/s00158-003-0368-6.

Martini, K., 2006. A new kind of software for teaching structural behavior and design. In *Building Technology Educators' Symposium*. University of Maryland: BTES, pp. 1–10.

McKay, M., Beckman, R.J. & Conover, W.J., 1979. Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2), pp.239–245.

Mckoy, F.L. et al., 2001. Influence of design representation on effectiveness of idea generation. In *Proceedings of DETC'01: ASME 2001 Design Engineering Technical Conferences*. Pittsburgh, PA: ASME, pp. 1–10.

Méndez Echenagucia, T., 2013. *Computational Search in Architectural Design*. Ph.D. Thesis, Politecnico di Torino.

Méndez Echenagucia, T. et al., 2014. Multi-objective acoustic and structural design of shell structures for concert halls. In *Symposia of the International Association for Shells and Spatial Structures*. Brasilia: IASS.

Méndez Echenagucia, T. et al., 2015. The early design stage of a building envelope : multi-objective search through heating , cooling and lighting energy performance analysis. *Applied Energy*, 154, pp.577–591.

Michell, A.G.M., 1904. The limits of economy of material in frame- structures. *Philosophical Magazine Series 6*, 8(47), pp.589–597.

Mirtschin, J., Geometry Gym. Available at: https://geometrygym.wordpress.com/.

Morbitzer, C. et al., 2001. Integration of building simulation into the design process of an architecture practice. In *Seventh International IBPSA Conference*. Rio de Janeiro, Brazil: IBPSA, pp. 697–704.

Moreno, D.P. et al., 2014. Fundamental Studies in Design by Analogy: a Focus on Domain-Knowledge Experts and Applications to Transactional Design. *Design Studies*, 35, pp.232–272.

Mueller, C.T., 2014. *Computational Exploration of the Structural Design Space*. Ph.D. Thesis, Massachusetts Institute of Technology.

Mueller, C.T. & Ochsendorf, J.A., 2015. Combining structural performance and designer preferences in evolutionary design space exploration. *Automation in Construction*, 52, pp.70–82.

Nelson, B., 1971. *Structural Principles Behind the Failure of Gothic Cathedrals*, Cambridge, MA: Class notes for MIT Framed Structures II, Prof. Traum.

Oxman, R., 2008. Digital architecture as a challenge for design pedagogy: theory, knowledge, models and medium. *Design Studies*, 29, pp.99–120.

Parker, E.G. & O'Brien, J.F., 2009. Real-time deformation and fracture in a game environment. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. New York, pp. 165–175.

Preisinger, C. & Bollinger-Grohmann-Schneider ZT GmbH, Karamba. Available at: www.karamba3d.com.

Quaglia, C.P. et al., 2014. Balancing energy efficiency and structural performance through multi-objective shape optimization : Case study of a rapidly deployable origami-inspired shelter. *Energy and Buildings*, 82, pp.733–745.

Raeissi, S. & Taheri, M., 1998. Optimum Overhang Dimensions for Energy Saving. *Building and Environment*, 33(5), pp.293–302.

Rahmani, M., Optimo. Available at: dynamobim.com/optimo/.

Ramzy, N.S., 2013. Perceptual dimension of interior daylight in sacred architecture: analytical study of the lighting programs in five sacred buildings of different styles. *International Journal of Architecture, Engineering and Construction*, 2(4), pp.219–233. Available at: http://www.iasdm.org/journals/index.php/ijaec/article/view/200.

Reinhart, C.F., Mardaljevic, J. & Rogers, Z., 2006. Dynamic Daylight Performance Metrics for Sustainable Building Design. *LEUKOS*, 3(1), pp.7–31.

Robert McNeel & Associates, Grasshopper. Available at: www.grasshopper3d.com.

Sartori, I. & Hestnes, A.G., 2007. Energy use in the life cycle of conventional and low-energy buildings: A review article. *Energy and Buildings*, 39(3), pp.249–257.

Schlecht, L.A., 2013. *Impact of prototyping resource environments on idea generation in product design*. S.M. Thesis, Massachusetts Institute of Technology.

Schmit, L. a., 1981. Structural synthesis - Its genesis and development. *AIAA Journal*, 19(10), pp.1249–1263. Available at: http://arc.aiaa.org/doi/abs/10.2514/3.7859.

Shah, J., Kulkarni, S. & Vargas-Hernandez, N., 2000. Guidelines for experimental evaluation of idea generation methods in conceptual design. *Journal of Mechanical Design*, 122(4), pp.337–384.

Shah, J.J., Vargas-hernandez, N. & Smith, S.M., 2003. Metrics for measuring ideation effectiveness. *Design Studies*, 24, pp.111–134.

Shea, K., Aish, R. & Gourtovaia, M., 2005. Towards integrated performance-driven generative design tools. *Automation in Construction*, 14(2), pp.253–264. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0926580504000809.

Simmons, C.T. & Mysak, L. a., 2010. Transmissive properties of Medieval and Renaissance stained glass in European churches. *Architectural Science Review*, 53(2), pp.251–274.

Solemma, Diva-for-Rhino. Available at: www.diva4rhino.com.

Sommerschield, H.F. et al., 1977. Letters: Montreal olympic structures. *Civil Engineering - ASCE*, 47(4), pp.5–36.

Suga, K., Kato, S. & Hiyama, K., 2010. Structural analysis of Pareto-optimal solution sets for multi-objective optimization: An application to outer window design problems using Multiple Objective Genetic Algorithms. *Building and Environment*, 45(5), pp.1144–1152. Available at: http://dx.doi.org/10.1016/j.buildenv.2009.10.021.

Suwa, M. & Tversky, B., 1997. What do architects and students perceive in their design sketches? A protocol analysis. *Design Studies*, 18, pp.385–403.

Tallon, A. & Murray, S., Mapping Gothic France. Available at: mappinggothic.org.

Tedesko, A., 1976. Experts Slam the Olympic Structures of Montreal. *Civil Engineering - ASCE*, 46(12), pp.50–54.

The Core Studio at Thorton Tomasetti, TT Toolbox. Available at: www.food4rhino.com/project/tttoolbox?ufh.

Tseranidis, S., Brown, N.C. & Mueller, C.T., 2016. Data-driven approximation algorithms for rapid performance evaluation and optimization of civil structures. *Automation in Construction*. Available at: http://dx.doi.org/10.1016/j.autcon.2016.02.002.

Tuhus-Dubrow, D. & Krarti, M., 2010. Genetic-algorithm based approach to optimize building envelope design for residential buildings. *Building and Environment*, 45(7), pp.1574–1581. Available at: http://dx.doi.org/10.1016/j.buildenv.2010.01.005.

U.S. Department of Energy, 2011. Buildings Energy Data Book. Available at: http://buildingsdatabook.eren.doe.gov/.

UNESCO World Heritage Centre, Amiens Cathedral. Available at: http://whc.unesco.org/en/list/162/.

Ungewitter, G., 1901. *Lehrbuch der Gotischen Konstruktionen* 4th ed. K. Mohrmann, ed., Leipzig.

Vanderplaats, G.N., 1999. *Numerical Optimization Techniques for Engineering Design* 3rd ed., Colorado Springs, CO: Vanderplaats Research and Development.

Vierlinger, R., Octopus. Available at: www.food4rhino.com/project/octopus?ufh.

Wang, W., Zmeureanu, R. & Rivard, H., 2005. Applying multi-objective genetic algorithms in green building design optimization. *Building and Environment*, 40(11), pp.1512–1525.

de Weck, O.L. & Jones, M.B., 2006. Isoperformance: Analysis and Design of Complex Systems with Desired Outcomes. *Systems Engineering*, 9(1), pp.45–61.

Wikipedia User: Wikigod, 2009. Suvarnabhumi Airport: Inner courtyard and terminal structure by night. *Wikipedia*. Available at: https://en.wikipedia.org/wiki/Suvarnabhumi_Airport#/media/File:Suvarnabhumi_Airport_Bangkok_by_Night.jpg.

Wikipedia_User:Tolivero, 2006. Olympic Stadium (Montreal). *Wikipedia*. Available at: https://en.wikipedia.org/wiki/Olympic_Stadium_(Montreal)#/media/File:Le_Stade_Olympique_3.jpg.

Zhai, Z. & Previtali, J.M., 2010. Ancient vernacular architecture: characteristics categorization and energy performance evaluation. *Energy and Buildings*, 42, pp.357–365.