

Early Building Design Using Multi-Objective Data Approaches

By

Nathan C. Brown

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

S.M. in Building Technology
Massachusetts Institute of Technology, 2016

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

Doctor of Philosophy in Architecture: Building Technology

at the

Massachusetts Institute of Technology

June 2019

© 2019 Nathan C. Brown. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature of Author: _____

Department of Architecture
May 3, 2019

Certified by: _____

Caitlin T. Mueller
Associate Professor of Architecture & Civil and Environmental Engineering

Accepted by: _____

Nasser Rabbat
Aga Khan Professor of Islamic Architecture
Chair, Department Committee on Graduate Studies

Dissertation Committee

Caitlin T. Mueller

Associate 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

Christoph F. Reinhart

Professor of Building Technology
Director, Building Technology Program
Massachusetts Institute of Technology
Thesis Reader

Early Building Design Using Multi-Objective Data Approaches

By

Nathan C. Brown

Submitted to the Department of Architecture on May 3, 2019
in Partial Fulfillment of the Requirements for the
Doctor of Philosophy in Architecture: Building Technology

Abstract

During the design process in architecture, building performance and human experience are increasingly understood through computation. Within this context, this dissertation considers how data science and interactive optimization techniques can be combined to make simulation a more effective component of a natural early design process. It focuses on conceptual design, since technical principles should be considered when global decisions are made concerning the massing, structural system, and other design aspects that affect performance. In this early stage, designers might simulate structure, energy, daylighting, thermal comfort, acoustics, cost, and other quantifiable objectives. While parametric simulations offer the possibility of using a design space exploration framework to make decisions, their resulting feedback must be synthesized together, along with non-quantifiable design goals. Previous research has developed optimization strategies to handle such multi-objective scenarios, but opportunities remain to further adapt optimization for the creative task of early building design, including increasing its interactivity, flexibility, accessibility, and ability to both support divergent brainstorming and enable focused performance improvement.

In response, this dissertation proposes new approaches to parametric design space formulation, interactive optimization, and diversity-based design. These methods span in utility from early ideation, through global design exploration, to local exploration and optimization. The first presented technique uses data science methods to interrogate, transform, and, for specific cases, generate design variables for exploration. The second strategy involves interactive stepping through a design space using estimated gradient information, which offers designers more freedom compared to automated solvers during local exploration. The third method addresses computational measurement of diversity within parametric design and demonstrates how such measurements can be integrated into creative design processes. These contributions are demonstrated on an integrated early design example and preliminarily validated using a design study that provides feedback on the habits and preferences of architects and engineers while engaging with data-driven tools. This study reveals that performance-enabled environments tend to improve simulated design objectives, while designers prefer more flexibility than traditional automated optimization approaches when given the choice. Together, these findings can stimulate further development in the integration of interactive approaches to multi-objective early building design.

Key words: design space exploration, conceptual design, design tradeoffs, interactive design tools, structural design, sustainable design, multi-objective optimization, data science, surrogate modeling

Thesis Supervisor: Caitlin T. Mueller

Title: Associate Professor of Architecture & Civil and Environmental Engineering

Acknowledgements

I am extremely grateful for the many advisors, colleagues, collaborators, funding sources, and friends who helped make this dissertation possible. First, I want to thank my advisor, Professor Caitlin Mueller, who had a profound impact on my development as a researcher and experience as a graduate student. My research is indebted to her own contributions in the area of computational structural design, but more generally to her deep enthusiasm and knowledge of the field. She has been the single greatest professional influence over my time at MIT, both pushing and inspiring me to be creative, thoughtful, and systematic in my research pursuits.

I am also thankful for the other members of my thesis committee: John Ochsendorf and Christoph Reinhart. I had the opportunity to engage with John and Christoph through research and teaching, and I owe much of what I know about both to their dedicated, passionate, and dynamic examples. I especially appreciate their collective vision for the future of design and the built environment, and their corresponding guidance on how to make my research relevant and impactful across architecture and engineering. I also appreciate the thoughtful feedback, technical assistance, and research guidance of the other BT Faculty—John Fernández, Leon Glicksman, Les Norford, and Josephine Carstensen.

Many colleagues and students in the Digital Structures research group also helped sharpen my research, while fostering a friendly, engaging, collaborative environment. Over my PhD years, these students included Renaud Danhaive, Paul Mayencourt, Pierre Cuvilliers, Yijiang Huang, Mohamed Ismail, Demi Fang, Kevin Moreno Gata, Felix Amtsberg, Hanne Cloesen, Jackson Jewett, Judyta Cichocka, Abigail Anderson, and Courtney Stephen. I am also appreciative of the support, friendship, and intellectual engagement from many others in the Building Technology program. Colleagues who helped with this research at one time or another include David Blum, Carlos Cerezo, Irmak Turan, Helen Awino, Andrew Brose, Shreshth Nagpal, Alpha Arsano, Charu Sharma, Johnathan Kongoletos, Justin Lueker, Mariana Liebman-Pelaez, Nicole Liwen, Nikhilesh Ghanta, Norhan Bayomi, and Yu Qian Ang—the longer I spent at MIT, the longer this list grew, and I am immensely appreciative for all of their help.

In addition, this dissertation was strengthened by feedback, comments, and suggestions from experts in related topics involving data science, optimization, and design studies. These include Roy Welsch, Olivier de Weck, Maria Yang, Qifang Bao, Edward Burnell, Zack Xuereb Conti, and Nathaniel Jones. I am also grateful to all of the participants in the user study, who helped illuminate the benefits, drawbacks, and future opportunities for developing performance-driven design tools.

The administrative staff in Building Technology and Architecture was likewise tremendously helpful in keeping me on track. This includes Kathleen Ross, Erin Buckley, and Denita Washington from BT, as well as many others in Architecture HQ. Related to the department, I am also extremely grateful for the MIT

Presidential Fellowship, which provided me with teaching opportunities and the financial support to develop my dissertation.

This research was also enriched by the opportunity to engage with professional practice and test some of these ideas as they developed. First and foremost, I am thankful for the opportunity to work with BuroHappold during my PhD, which provided many of these opportunities. While I had helpful experiences with too many designers and engineers at BuroHappold to mention, I would like to thank Paul Richardson and the Boston Structures team, Emir Pekdemir and the Computational Analytics team, Shrikant Sharma and the Smart Space team, Al Fisher, Rob May, Erik Narhi, Josef Taylor, Isak Näslund, and others with the computational collective, as well as Wolf Mangelsdorf. I am also grateful to Kyle Vansice of SOM and Matt Breidenthal and his team at HOK for fruitful discussions around design tools and processes.

The computational tools created in conjunction with this dissertation were also a collaborative effort, with many developers contributing to the project. These include Renaud Danhaive, Jonathas Felipe, Stavros Tseranidis, Eric Wong, Violetta Jusiega, and Alicia Nimrick, and I am grateful for their involvement.

In closing, the past few years have been greatly enriched by the support of friends and family. In particular, I appreciate the relationships I developed with MIT undergraduates, GRTs, staff, and faculty at Maseeh Hall, as well as many others who listened to practice talks, proofread, or generally helped make MIT a special place for me.

Finally, to my wife, Sarah—I could not have done it without you.

Table of Contents

- 1 Introduction.....23
 - 1.1 Problem statement..... 23
 - 1.2 Research motivation 25
 - 1.3 Research question and purpose 26
 - 1.4 Organization of dissertation 27

- 2 Engaging Performance in Conceptual Design: Background and Framework.....29
 - 2.1 Typical early design ideation: analog and computational..... 32
 - 2.2 The design space 34
 - 2.3 Developments in optimization 36
 - 2.3.1 From structural optimization..... 36
 - 2.3.2 Optimization as a black box for design..... 36
 - 2.3.3 Multi-objective optimization..... 37
 - 2.4 Need and potential for optimization in building design..... 39
 - 2.5 Integration of live feedback and optimization into the building design process..... 40
 - 2.5.1 Interactive design methods based primarily on feedback 40
 - 2.5.2 Design methods based primarily on optimization 44
 - 2.5.3 Interactive design methods using optimization..... 44
 - 2.5.4 Design by shopping and the catalog approach 46
 - 2.5.5 Artificial intelligence and immersive environments..... 47
 - 2.6 Related tools for interaction and optimization in architecture..... 47
 - 2.7 Improving interactive parametric workflows using data science..... 48
 - 2.8 Related data-driven techniques in engineering..... 48
 - 2.9 Related data-driven techniques for building design..... 49
 - 2.10 A discussion of data science applications in this dissertation..... 51
 - 2.11 Vision for performance-based, data-driven design 56
 - 2.11.1 Designer + computer engaging with the design space together 56
 - 2.12 Research goals and responses 58
 - 2.13 Overall map of design process..... 60

3	Design Formulation and Global Exploration.....	63
3.1	Introduction	63
3.2	Approaches to problem formulation and exploration from engineering.....	65
3.3	Data science for tuning the design problem.....	66
3.3.1	Background on variable analysis in architecture	66
3.4	Overview and goals.....	67
3.4.1	Strategies for variable analysis, transformation, and generation.....	67
3.4.2	Case studies and datasets	68
3.5	Methodology and results.....	71
3.5.1	Analyzing existing design variables for importance.....	71
3.5.2	Dimensionality reduction for more meaningful variables	72
3.5.3	Automatic variable generation through reduction.....	79
3.6	Conclusions	82
3.6.1	Contributions towards variable analysis, transformation, and generation.....	82
3.6.2	Future work and concluding remarks.....	83
4	Interactive Optimization and Local Exploration.....	85
4.1	Introduction	85
4.2	Background on gradient-based guidance	88
4.3	Methodology	89
4.4	Gradient stepping for a courtyard building design	95
4.5	Isoperformance stepping for an airport terminal design	98
4.6	Conclusions	100
4.6.1	Contributions towards interactive optimization	100
4.6.2	Future work and concluding remarks.....	101
5	Diversity-driven Ideation and Design.....	103
5.1	Introduction	103
5.1.1	Organization of chapter	105
5.2	Literature review	105
5.2.1	Background on diversity measurement	105
5.2.2	Creativity, novelty, variety, and related concepts in engineering design	108
5.2.3	Diversity in parametric formulations.....	109

5.2.4	Methods for measuring diversity in parametric modeling.....	110
5.2.5	Research contributions.....	114
5.3	Methodology	114
5.3.1	Generation of test data	115
5.3.2	Survey	117
5.3.3	Procedure.....	118
5.4	Results.....	118
5.4.1	Human computer agreement	118
5.4.2	Differences between metrics.....	120
5.4.3	Additional factors in the study: visually dominant variables.....	121
5.4.4	Individual preference strengths for each metric.....	122
5.4.5	Limitations.....	123
5.5	Discussion and conclusions	124
5.5.1	Qualitative comparison of metrics	124
5.5.2	Example application	124
5.5.3	Contributions towards diversity metrics for human-centric design.....	125
5.5.4	Future work and concluding remarks.....	126
6	Implementing Data-Driven Parametric Design with a Flexible Toolbox Approach.....	127
6.1	Introduction	127
6.2	Workflow methodology	128
6.2.1	Design Space Exploration overview	128
6.2.2	Possible workflows with DSE.....	130
6.3	Example of conceptual grid shell roof	136
6.3.1	Design formulation.....	139
6.3.2	Diversity filtering	141
6.3.3	Interactive optimization.....	145
6.3.4	Comparison of designer types and lessons learned.....	148
6.4	Discussion on design example performance compared to benchmarks.....	150
6.5	Qualitative comparison to other tools	157
6.6	Conclusions	159
6.6.1	Contributions towards flexibility & accessibility in data-driven parametric design.....	159
6.6.2	Future work and concluding remarks.....	159

7	User Study of Global and Local Design Exploration.....	161
7.1	Introduction	161
7.2	Literature review	162
7.2.1	Need for interaction during design	162
7.2.2	Background on design studies.....	163
7.2.3	Design studies from product and mechanical design	163
7.2.4	Related design studies in early building design	164
7.3	Study purpose.....	164
7.4	Experiment procedures	165
7.4.1	Overview	165
7.4.2	Design prompt.....	166
7.4.3	Population and recruitment	170
7.4.4	Location and length of subject involvement	170
7.4.5	Study hypotheses and data collection.....	171
7.5	Description of design environments.....	171
7.6	Design outcome results	173
7.6.1	Trends from the design histories	173
7.6.2	Simulated performance results from favorite designs.....	175
7.6.3	Diversity of selected designs	177
7.6.4	Segmented results	179
7.7	Survey results	180
7.7.1	Comparative environment preferences	180
7.7.2	Individual environment reviews.....	183
7.7.3	Tool preferences and general comments	185
7.8	Discussion	187
7.9	Conclusions	189
7.9.1	Contributions towards understanding interactive design processes	189
7.9.2	Future work and concluding remarks.....	189
8	Conclusions.....	191
8.1	Research questions and dissertation goals	191
8.2	Summary of contributions	192
8.2.1	Design space formulation	192

8.2.2	Interactive exploration	192
8.2.3	Diversity-driven design	192
8.2.4	Workflow demonstration and user testing	193
8.3	Potential applications and impact.....	193
8.3.1	In design education.....	194
8.3.2	In practice	194
8.4	Remaining questions and future work.....	194
8.4.1	Next steps of specific methodologies.....	194
8.4.2	The use of data approaches within a traditional project structure.....	196
8.4.3	Natural hierarchies between performance domains in building design.....	197
8.4.4	Further validation of data-driven design methodologies	198
8.4.5	Connections to construction and fabrication	198
8.4.6	Progressive, dynamic, and generalizable surrogate modeling	199
8.4.7	The role of simulation compared to intuition and experience	199
8.5	Final remarks and outlook.....	200
9	References.....	201
10	Appendix A: Design Tool Information.....	221

List of Figures

Figure 1.1: Buildings and structures designed for multiple objectives	26
Figure 2.1: Overall primary energy consumption in the United States from 2017, as well as a simple comparison of embodied versus operational energy usage for a typical building.....	30
Figure 2.2: Design options for River Tower, Ho Chi Minh City, Foster + Partners.....	31
Figure 2.3: Concept sketches by Jörg Schlaich on the left (The Happy Pontist 2009), and concept for a tower by SOM (Chow & Sarkisian 2009)	32
Figure 2.4: An example of a design and objective space, from a two-variable parametric truss, after Mueller (2014)	35
Figure 2.5: Approaches to multi-objective optimization, after Marler & Arora (2004).....	38
Figure 2.6: AcceleradRT by Jones & Reinhart (2018)	41
Figure 2.7: An early real-time structural analysis and design environment by Clune (2010)	42
Figure 2.8: An example of a design generated by RhinoVAULT (Rippmann et al. 2012b).....	43
Figure 2.9: paraGEN by Turrin et al. (2011)	45
Figure 2.10: structureFIT by Mueller & Ochsendorf (2015a)	46
Figure 2.11: Exploring generative 3D shapes using autoencoder networks by Umetani (2017)	49
Figure 2.12: A description of the design space for a trussed roof design example, along with the results of a variable effects analysis using Taguchi methods.....	50
Figure 2.13: A surrogate model of rooftop PV potential, which enables performance-based exploration of form in real time with only small errors	53
Figure 2.14: A clustered design space for an eleven-bar truss	55
Figure 2.15: A view of the typical design process, bogged down by translation issues between software, leading to only a few designs being meaningfully considered	57
Figure 2.16: Humans and computers working together to engage with a dynamic design space in order to find a satisfying design.....	58
Figure 2.17: A map of the overall design process when using a computational parametric approach	61
Figure 3.1: Visual descriptions of the parametric design case studies used for analysis	69
Figure 3.2: Results of the variable analysis procedure for two separate model types	72

Figure 3.3: Variable directions derived from canonical correlation analysis and principal component analysis on the truss case study 74

Figure 3.4: Lesser principal component directions for the truss case study 75

Figure 3.5: Geometric modifications in the directions indicated by canonical correlation analysis..... 76

Figure 3.6: Principal component directions for the SFO case study..... 77

Figure 3.7: A higher resolution visualization of design performance along the CCA structure direction near the starting design 78

Figure 3.8: An example of a live conceptual design environment, which can include performance feedback and new synthetic variables using the transformations described throughout this chapter..... 78

Figure 3.9: 18 representative solutions selected out of the 2000 design data set 79

Figure 3.10: Automatically generated performance-controlling variables for the initial provided geometry 80

Figure 3.11: Potential massing changes that increase the rooftop PV potential for a series of buildings, automatically suggested for a given geometric input 82

Figure 4.1: Two different strategies for manipulating parametric designs 87

Figure 4.2: Visualization of gradient vectors projected into a design environment 92

Figure 4.3: Potential steps in the design space based on gradient information..... 94

Figure 4.4: A visualization of projected outcomes for a single step in the design space 95

Figure 4.5: Two examples of recent or contemporary courtyard buildings..... 96

Figure 4.6: A potential path through the design space based on steps taken for different performance objectives 97

Figure 4.7: Attempted isoperformance exploration using different approaches for step size and choosing directions..... 99

Figure 4.8: Comparison of delayed and live design environments 100

Figure 5.1: A diagram illustrating the difference between measuring distance and measuring the diversity of a set of designs..... 107

Figure 5.2: Diversity of architectural solutions for similar design conditions, in both the physical world and parametric design 110

Figure 5.3: A visualization of the diversity calculation for each design set, as well as the normalized diversity value for each metric..... 116

Figure 5.4: The parameterized shapes of each design set used in the survey..... 117

Figure 5.5: An example survey question as shown to a study participant..... 118

Figure 5.6: The results of the human-computer agreement survey questions, showing the magnitude of the difference in measured diversity versus the percentage of respondents that agreed with all computer metrics. 119

Figure 5.7: An example illustrating the difference between two sets of designs that are easy for humans to select the most diverse set, and two sets that are more difficult to distinguish..... 119

Figure 5.8: A visual description of the different parameterizations for one-variable dominant and regular shapes, along with the shape-separated raw question results, showing no distinct pattern 122

Figure 5.9: The strength and totals for individual preferences of participants in the study, as indicated by their overall best winning percentages. 123

Figure 5.10: The results of an example parametric tower study, showing the novelty of a design set generated using a process based on diversity metrics in this chapter. 125

Figure 6.1: Diagram of the relationships between parametric components and plug-ins..... 129

Figure 6.2: A typical parametric workflow connected to performance simulation 130

Figure 6.3: Workflow for generating a design catalog of options in parametric software 131

Figure 6.4: Workflow for parametric multi-objective optimization 131

Figure 6.5: Workflow that includes a post-processing component for returning a more directed yet diverse sampling of the design space 132

Figure 6.6: Workflow for initial experimentation with synthetic variables in parametric design 133

Figure 6.7: Workflow with data-driven techniques for analyzing the design space 134

Figure 6.8: Workflow for building a predictive surrogate model to provide live objective function estimation 135

Figure 6.9: Similar to the workflow above, but combined with a tool for interactive optimization that allows for direct movement in both the design and objective spaces..... 135

Figure 6.10: Potential options within the design space for this example, along with building precedents they resemble..... 137

Figure 6.11: Geometric design variables for the example..... 138

Figure 6.12: An initial estimation of variable importance to the problem by calculating variable effects. 140

Figure 6.13: CCA variable directions based on SMQ and EUI for the design space 141

Figure 6.14: A comparison of number of qualified designs and diversity of diversity culled sets for brainstorming, at different structural performance targets..... 143

Figure 6.15: Example sets of 12 diverse design possibilities for three different structural performance levels 144

Figure 6.16: Visualization of the accuracy of surrogate models for structure and energy..... 145

Figure 6.17: Live visualization of geometry and performance (according to the surrogate models) together 146

Figure 6.18: The interface for interactive optimization using gradient-based guidance..... 147

Figure 6.19: A combined view of the geometry, feedback, and control window for interactive optimization 148

Figure 6.20: Possible paths through the design space using interactive optimization, for two different designer types..... 149

Figure 6.21: The range of performance improvement generally considered for structural weight per area in this case study, compared to similar buildings from the DeQo database, after De Wolf (2016)..... 152

Figure 6.22: The range of performance improvement generally considered for EUI in this design example, using only geometric variables, compared to similar buildings from Lawrence Berkeley National Lab’s Building Performance Database (Mathew 2015) 154

Figure 6.23: History of energy upgrades for the design example, which are here explored independent of geometry 155

Figure 7.1: A visualization of design objectives and variables for the study. 168

Figure 7.2: Example visualization and design script used for the study 172

Figure 7.3: Number of designs considered and average time spent on designs (in seconds), by phase..... 173

Figure 7.4: The average performance of all designs considered according to the performance surrogate models..... 175

Figure 7.5: Simulated performance of each favorited design from the study. Each data point is a single design. 176

Figure 7.6: Representative designs from the global exploration phases. There is no meaningful difference in the diversity between the two environments. 178

Figure 7.7: Representative designs from the optimization phases of the study..... 179

Figure 7.8: A comparison of design performance between architects and engineers, from their selected favorites. 180

Figure 7.9: Overall composite rankings of design environments using weighted scoring system..... 181

Figure 7.10: The results from a ranked comparison between design environments 182

Figure 7.11: Survey results for each environment individually 184

Figure 7.12: Interface preferences for performance-driven parametric design tools 185

Figure 7.13: A grouping of the environments that allowed for creativity and produced more diversity, versus the environments that tended to improve performance..... 187

Figure 8.1: Examples of when design space exploration techniques may be useful throughout the building delivery process as software becomes better integrated..... 197

Publications Related to this Dissertation

The introduction and examples in the background section were adapted from:

Brown, N., & Mueller, C., 2017. Designing with data: moving beyond the design space catalog. In *Proceedings of ACADIA 2017: Disciplines and Disruption*. Boston, MA.

Chapter 3 was published in part in:

Brown, N., & Mueller, C., 2019. Design variable analysis and generation for performance-based parametric modeling in architecture. *International Journal of Architectural Computing*, 17(1), pp. 36-52.

Chapter 4 was published in part in:

Brown, N., & Mueller, C., 2018. Gradient-based guidance for controlling performance in early design exploration. In *Proceedings of the International Association for Shell and Spatial Structures Symposium 2018*. Boston, MA.

Chapter 5 was published in part in:

Brown, N. & Mueller, C., 2019. Quantifying diversity in parametric design: a comparison of possible metrics. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 33(1), pp. 40–53.

Related work has appeared in:

Brown, N., & Mueller, C., 2017. Automated performance-based design space simplification for parametric structural design. In *Proceedings of the International Association for Shell and Spatial Structures Symposium 2017*. Hamburg, Germany.

1 Introduction

1.1 Problem statement

This dissertation considers how interactive optimization techniques and data science methods can be combined to enable early-stage, multi-objective design exploration for buildings and related structures in the built environment.

It builds on the history of optimization techniques for creative, synthetic contexts, while also engaging with newer techniques for handling data during the design process. The use of large quantities of data for predictive purposes and decision support, whether this data is collected from past actions or generated by high-fidelity simulations, has the potential to revolutionize many fields. Even in professions that are historically dependent on human intuition and reasoning, data can be used to supplement or enhance human activities. Chess and radiation oncology are two examples for which the ability to project future outcomes has fundamentally changed how these activities are pursued (Somers 2017). Based on computers that simulate millions of possibilities in great detail, chess players can see breakdowns of their mistakes, train with customized game scenarios that focus on their weaknesses, and if allowed, view real-time information about how each potential move increases or decreases win probabilities. Similarly, doctors can use physics models of radioactive particles to decide precisely where to focus beams of these particles during treatment. Rather than forcing experts to submit to the will of the computer, data improves their ability to predict future results and make corresponding adjustments.

As computational methods continue to develop, architects and engineers responsible for conceptual design are also taking advantage of the many data streams available to them during design. When compared to other fields, however, some specific difficulties arise while injecting simulation into the design and planning of buildings. Building design is more complex than a game of chess in terms of parameters, stakeholders, or even choosing exactly what “winning” means, and each building design is to some extent a custom process. Nevertheless, researchers have produced a wealth of simulation tools that are available to predict how a building behaves, and formed some broad definitions of design success—energy efficiency, low embodied carbon, high daylight autonomy and thermal comfort, economic and social sustainability—that are both universally recognized and increasingly quantifiable. In order to achieve the goal of data-driven design in the form of human-computer collaboration, these feedback streams and measuring sticks for performance must be intimately linked to the design process itself.

Unfortunately, despite advances in the speed and accuracy of predictive simulations, current technological and organizational barriers often prevent practitioners from fully integrating simulation data, which has traditionally come from engineers or specialists, into their design workflows, especially early in the process. In many cases, building design projects have already moved into design development or are otherwise frozen in terms of massing or geometry before all relevant performance simulations have been conducted. Even for advanced designers comfortable with data-driven design, the state of the art is essentially to generate a variety of potential design options, view visualizations of their performance, and pick the best option. Functionally, the creation of a design catalog resembles older, more analog procedures for design optioning—computation has sped up the process and provided some notion of relative performance between options, but predictive simulations have not yet reached their full potential to change how we design.

When considering the broader convergence of computation and design, optimization can offer a way to move beyond data feedback and push towards guidance, which makes the computer a more active participant in the decision-making process. In mechanical and aerospace engineering, engineering systems, finance, and elsewhere, designers and other professionals have successfully applied optimization to a wide variety of problems. However, most optimization techniques work best when designers know exactly what the objectives are and how much they matter compared to one another, and they are willing to accept whatever answer the computer produces. Optimization techniques can certainly assist in improving upon the catalog, but in order to be most useful in architecture, they must be adjusted to account for non-quantifiable objectives, and be adapted to maintain creative flow and enable satisfying design experiences (Brady 1986; Hoxmeier & DiCesare 2000).

This dissertation adapts traditional optimization techniques with the intention of making them more meaningful and accessible, interactive and flexible, and work in directed or divergent ways, depending on the requirements of a particular design activity. It does this primarily for parametric design environments

through the use of data science techniques and related methods of systematically considering design possibilities by treating each one as a data point in a larger set.

1.2 Research motivation

This research sits at the intersection of three broad trends, which serve to motivate the work:

1. The Performance Imperative | At present, there is a strong demand that future buildings perform well, especially in relation to their human occupants and the environment as a whole. Buildings are constructed to respond to human needs, and as we become better equipped to evaluate, analyze, and plan for those needs, especially through the generation and handling of data, it is important that such considerations drive the design process. The performance imperative is also motivated by the acknowledgement that buildings are responsible for negative environmental impacts, in large part due to their responsibility for a significant percentage of global carbon emissions. Future buildings must work to mitigate this impact through design for low operational energy, efficient use of materials, durability, and resiliency, among other objectives. At the same time, human building designers consider many other design consequences and outcomes outside of what can be quantified or simulated, and performance must be considered in a broader, more holistic design context.

2. Increasing Computational Competency | Designers across architecture and engineering have now been using computers as part of the design process for decades. Engineers are well acquainted with conducting performance simulations, and architects and other design specialists are becoming increasingly comfortable with their usage. Researchers have developed exciting ways of connecting geometry from modeling environments directly to validated simulation engines and managing large datasets related to design and construction. Furthermore, leaders in many firms are seeking designers with skills in parametric design, coding, and simulation. Yet with access to all of this simulation feedback, there is still an ongoing discussion regarding how to best use this information for guidance—in other words, why is one design better than all of the other conceptual possibilities? How can simulation be used to explore these possibilities or even point to new ones that were not originally conceived?

3. Developments in Data Science | Recently, work in the areas of data mining, machine learning, and artificial intelligence have demonstrated extraordinary capacity to uncover helpful patterns and trends for human purposes (Belsley et al. 1980; Gorunescu 2011; Witten et al. 2016). In various professional fields, there have been revolutions prompted by the use of such techniques for predicting outcomes and supporting decision-making. Parametric design and related techniques that automatically generate design possibilities with numerical representations and directly simulatable performance offer the possibility of analyzing these theoretical datasets to provide computational support for performance-based design, especially through interactive optimization.

1.3 Research question and purpose

Flowing from these three developments, a single, broad research question motivates all of the work in this dissertation:

How can creative architects and engineers use computation to manage multiple quantitative objectives as part of a natural design process?

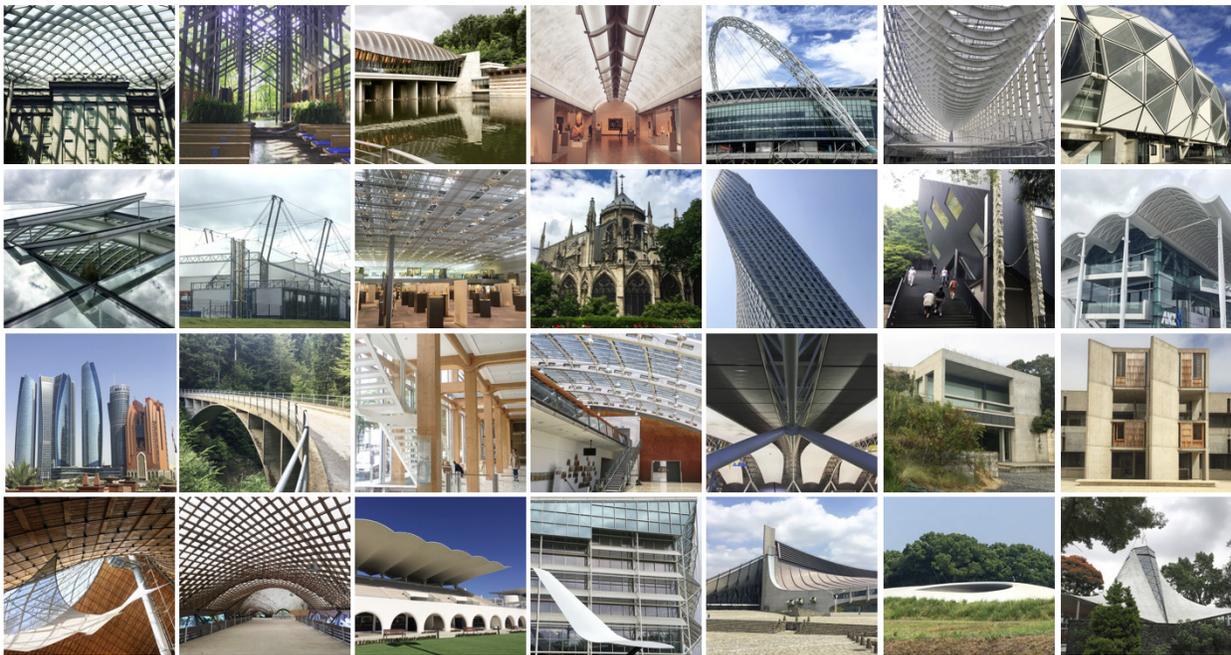


Figure 1.1: Buildings and structures designed for multiple objectives, photographed as part of the SOM Structural Engineering Travel Fellowship (Brown 2016a)

Figure 1.1 shows existing buildings and structures that were designed to respond to multiple conditions, which collectively demonstrate the creative capacity and synthetic abilities of human designers (Brown 2016a). As simulation becomes a more fundamental part of design, it is important to consider how it can be implemented in a way that enhances rather than subtracts from these natural design processes. This research question is broad, and requires careful consideration of the design act itself. To provide context from architecture and engineering, the background chapter of this dissertation contains a review of performance-based and data-driven design approaches along with a conceptual mapping of typical parametric processes. It also contains a discussion of research goals, which stem directly from developments and current practices in computational design. The following chapters then pose and address

more specific research questions relating to design space formulation, interactive optimization, diversity-driven design, and effective digital design environments.

In the end, the goal of this dissertation is to develop design approaches that lead to buildings that are high performing, yet sensitive to designer intent and synthesized through a natural creative process involving human-computer collaboration. It is primarily concerned with relative design decisions that are made in the early stages, between different options, rather than for comparison to absolute metrics or benchmarks. In particular, the dissertation addresses the navigation of design decisions that interact with each other meaningfully. For the examples described in the following chapters, this primarily leads to the exploration of performance metrics that are influenced by global geometry. One hope is that designers will use these approaches to create innovative new forms based on a sound initial concept—in effect, making good designs even better. However, it is also likely that some designers will use such methods to modify willfully form-driven buildings, which can at least lead to relative performance improvements and increase designer awareness about the implications of their decisions. In either case, this dissertation considers how designers compare the conceptual possibilities within all types of parametric datasets, and how they react when quantitative information guides or pushes them towards certain regions of a design space.

These approaches are not (and are not meant to be) a full substitute for experienced, technically proficient designers who can help select appropriate design spaces and interpret their suggested options. Nor are they intended to force or impose a particular outcome on a designer. There are a wide variety of architectural and engineering practitioners across the world, and data-driven design methods could fit into a range of design processes and intentions. At the same time, computational design shows strong potential for the generation of non-standard forms, the ability to give designers a clearer view of the performance impact of their decisions, and even the suggestion of both local and global design directions that can improve performance without sacrificing creativity. As such, the data-driven approaches in this dissertation can enhance conversations between architect, engineer, and computer while helping to mediate perceived tradeoffs in design by quickly providing meaningful, quantitative information and directions for iteration.

1.4 Organization of dissertation

This dissertation is divided into seven chapters. This first chapter serves as the introduction and gives the main problem statement, while also articulating the overall research question that motivates the remaining chapters.

Chapter 2 first provides a critical literature review that contextualizes the new research presented in this dissertation. This review describes typical design processes in early building design, and discusses the trend towards using computation for increasingly consequential design tasks. It then offers background on parametric methods and the concept of the design space, before detailing the use of optimization in both

architecture and engineering. Next, it provides information about how data science techniques can be applied to create interactive, multi-objective environments that are more amenable to the way human designers work creatively and synthetically. Finally, it presents the vision for this dissertation, which includes an overall framework for computational design using data-infused techniques. The chapter concludes by defining research goals and detailing how these goals expand current parametric design methods.

Chapter 3 demonstrates new methods for variable analysis, transformation, and generation during design formulation, which can produce and modify parametric design spaces in ways that stimulate creativity but are still directly linked to performance.

Chapter 4 proposes and interrogates a new method for interactive gradient-based guidance for multi-objective conceptual design problems that allows architects to synthesize both qualitative and quantitative goals simultaneously. When aided by data-driven models, this approach to optimization can inject simulation into live, dynamic design processes while enabling both directed and divergent geometric exploration.

Chapter 5 consolidates existing literature on the measurement of design diversity as it relates to parametric design, and presents the results of a study aiming to understand the relationship between human perceptions of visual diversity and how it is measured computationally. This research provides the foundation for evaluating the effect of performance-driven digital tools on human creativity and ideation. It also demonstrates applications of diversity-based thinking that are relevant to design space formulation and interactive optimization, which are considered in the previous two chapters.

Chapter 6 describes a set of digital tools that have been developed to implement, test, and share the methods described this dissertation. While focused on practical applications, it also contains a discussion that justifies the toolbox approach to data-driven design, and reflects on how this approach builds on existing infrastructure for computational parametric design.

Chapter 7 presents the results of a design study that tests the output and workflow preferences of designers as they engage with live feedback prediction and interactive optimization when compared with performance-free environments and fully automated optimization.

Chapter 8 contains a summary of contributions, a discussion of future work, and closing comments.

2 Engaging Performance in the Conceptual Design Process: Background and Framework

In contemporary building design, considerations of performance have become essential. Performance here is defined broadly, since there are many ways in which a building can be evaluated for effectiveness, efficiency, and experience. Traditional building performance goals may include structural, energy, envelope, cost, and acoustic targets. As additional systems and interactions have become simulatable, designers of high-performance buildings might also consider circulation, views, daylight, embodied energy, lifecycle emissions, and other criteria. Many of these aspects of building performance respond in a direct way to a sustainability imperative, in which designers recognize that buildings account for a large portion of primary energy consumption (~40% in the United States), as well as significant embodied energy (De Wolf et al. 2016). In considering these objectives, designers take responsibility for mitigating negative environmental impacts for both existing and new construction (see Figure 2.1). Other aspects of performance may respond more directly to human experience of buildings, a trend that that will likely grow as increasing availability of occupant-level data results in the development of additional quantitative metrics.

Depending on the complexity of the building, a variety of professionals may be responsible for promoting building performance, including architects, engineers, building scientists, and other specialists. Regardless of how the priorities are distributed, it is clear that technical thinking can have a large influence on the

performance of a design if it is done early, particularly in an integrated fashion. As Mueller (2014) notes, when moving through the design process, knowledge about the eventual building is gained as decisions are made and refined, but flexibility to change the design is gradually lost. When technical thinking only comes in during later phases of the process, its influence is often reduced to small changes or attempts to “optimize” the building performance within an already set system, massing, or even fully documented geometry. Thus, conceptual design becomes a crucial phase for making performance-based decisions.

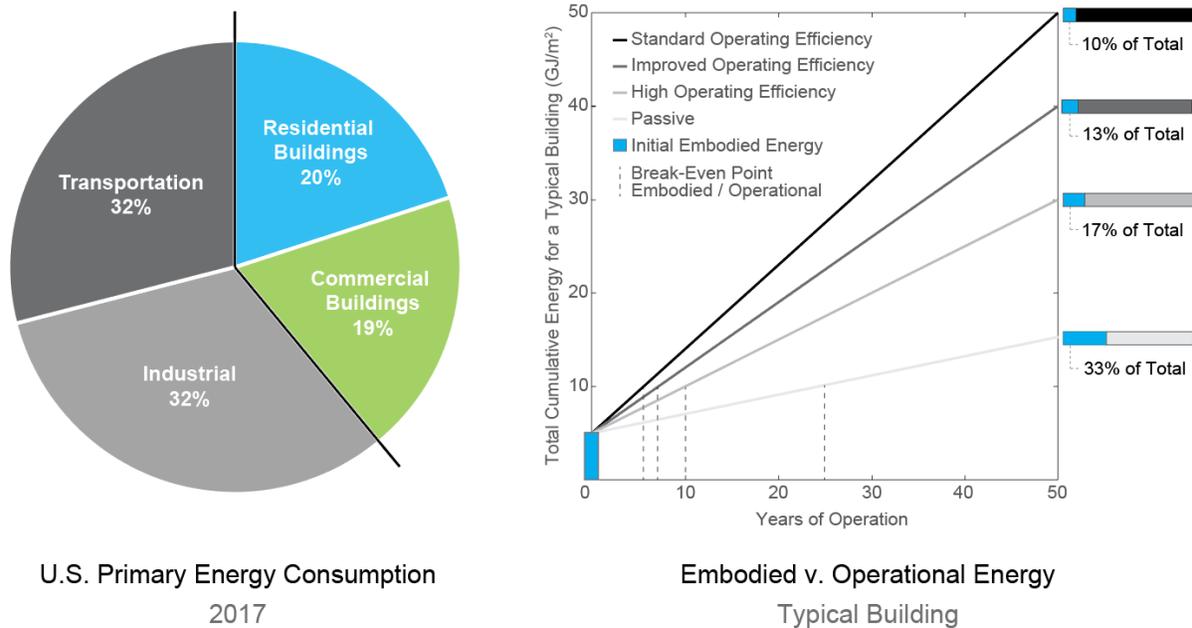


Figure 2.1: Overall primary energy consumption in the United States from 2017 (EIA 2017), as well as a simple comparison of embodied versus operational energy usage for a typical building (Brown & Mueller 2016)

This observation is not meant to diminish research efforts into later design or even the retrofitting of existing buildings, as there are certainly meaningful ways to influence building performance at these later stages, and even for building commissioning and operations. For example, data science and machine learning techniques have proven valuable in using sensor data to tune building energy operations and increase their efficiency (Shaikh et al. 2014). Although it is still a nascent technology, the structural systems of buildings may also eventually be tuned this way, such that they are able to respond to sensor data and self-stress or heal in ways that react to dynamic loading over time, as theorized by Johns et al. (2014) and related research. However, if designers wait until all geometry is already established to make use of computation and simulation data, they are missing clear opportunities to have an enormous influence over the eventual performance of a building, regardless of these other developments.

Moving simulation feedback and guidance earlier into the design process poses significant complications, compared to using these techniques for an already established geometry. As an example, Figure 2.2 shows a series of possibilities for a tower design in Ho Chi Minh City (St John 2019). These designs vary substantially in overall massing, floorplate shape and configuration, structural system, and other characteristics. Using typical workflows for building delivery, a conceptual designer seeking simulation feedback would need to build a separate model for each one, complete the simulations, and compare side by side. In historical design processes this was prohibitively effortful and expensive—even as computer simulations became commonplace in structural and energy analysis, the translation between a massing model and an analytical model for each option could take weeks or months.

As such, typical design concepts are not compared using the quantitative simulation data that is eventually generated for every finished building. For some conceptual designs, expertise and intuition is enough to make good decisions in the early stages. An experienced structural engineer might use an understanding of physical behavior and perhaps some hand calculations to select the best design direction, for example. Yet many buildings have become incredibly complex or at the very least subject to interacting and competing performance objectives. In this context, simulations can certainly help, as long as they are informative while still allowing for the creative freedom and flexibility demonstrated by the design example in Figure 2.2.



Figure 2.2: Design options for River Tower, Ho Chi Minh City, Foster + Partners, 2008 (St John 2019)

2.1 Typical early design ideation: analog and computational

As the addition of digital design tools threatens to overconstrain many design processes, the importance of flexibility in early design has come to require more substantial defense and attention. At its foundation, design is the act of creating a plan for the construction of an objective, system, or action (Design 2017), which can be executed entirely on a conceptual level. The initial stages of any design process often include rapidly generating and considering potential solutions or outcomes through brainstorming. Especially in architecture, or building and structural design more generally, the design process has historically been driven by humans, frequently accompanied by sketching or other forms of representation to communicate and clarify ideas (Goldschmidt 1994). A recent survey of architecture firms found that even though digital representation and even building performance simulation are desirable and accessible tools for architectural practice, most design processes (72%) still start out with hand sketching (Soebarto et al. 2015). Two examples of sketching and physical modeling, which were used during early design processes, are presented below, in Figure 2.3.

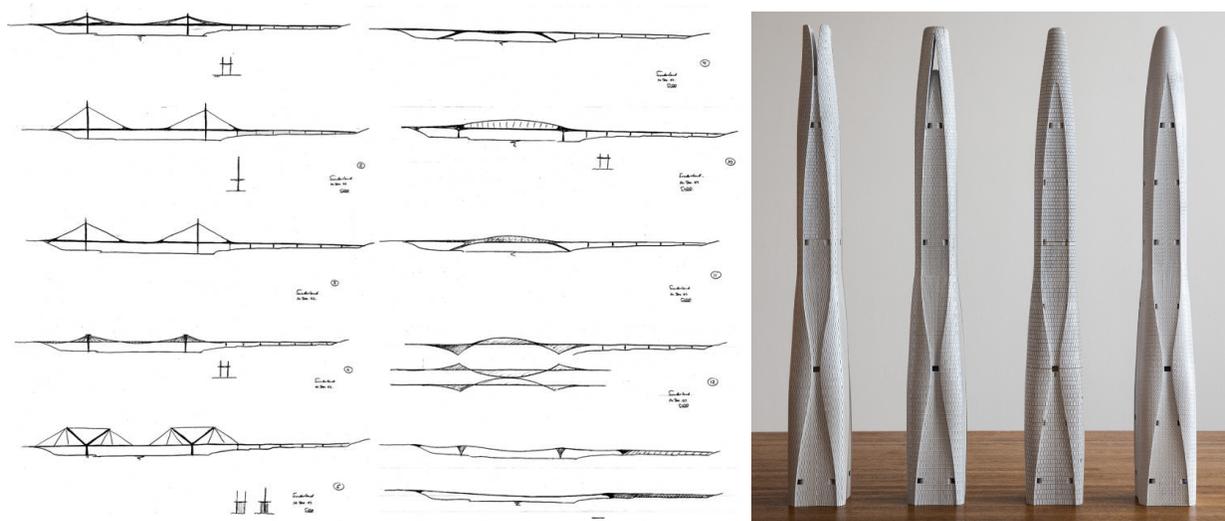


Figure 2.3: Concept sketches by Jörg Schlaich on the left (*The Happy Pontist* 2009), and concept for a tower by SOM (*Chow & Sarkisian* 2009)

These analog methods offer considerably more creative freedom than most computational methods, especially when computational geometry is connected to meaningful performance simulation. While there are obvious limits to what can eventually function as a building design, particularly at larger structural scales, these limits are not imposed in a hard way on an analog process. A designer can draw, sketch, or physically model whatever is imaginatively and geometrically possible, while intuitively making decisions

about which directions to pursue. Yet at the same time, these intuitions may be incorrect, naïve, or at the very least not as refined as later in the design process, when a design has been analyzed more thoroughly. Computational methods seek to augment these intuitions, primarily through performance simulation, at least when technical aspects of building design are considered. In some ways, the bar for adding to or “beating” a piece of sketch paper is extremely low—any simulation feedback, as long as it is accurate enough and relevant to the design problem, is likely an improvement over a performance-free environment when compared solely on the level of available information. Yet it is important to acknowledge and work to mitigate the sacrifices in design freedom that are associated with translating a design concept into one that can be rigorously and usefully simulated.

Due to the clear advantages of computational methods for generating diverse design outcomes and directly connecting those options to performance simulation, many designers have acknowledged their utility during design exploration. Two such computational methods for this purpose are parametric design (Burry 1996; Monedero 2000) and rule-based systems such as shape grammars (Stiny 1980; Knight 1991). These methods save considerable manual effort during design ideation by allowing for the generation of numerous possibilities from a single system, such that a designer need only develop that system rather than individually conceive of and represent every single option. Furthermore, parametric models afford the ability to include required building properties for simulation in the model definition itself, such that each generated option can be simulated and compared directly without additional human effort. While these methods have certainly not replaced the paper and pencil, and may be most useful after an analog method is used to explore vastly more diverse design possibilities (typologically and otherwise), computational design has clear potential for improving the conceptual building design process.

Pioneering contributions towards these computational design methods were made in architecture and related engineering fields as long as 20-30 years ago, meaning that they are no longer novel methods in themselves. While shape grammars have significant generative potential, can overcome certain conceptual limitations of parametric design, and continue to see applications (Wang & Duarte 2002; Muehlbauer et al. 2017), parametric design has become more commonplace as the method of choice in practice. This may be partially due to the prevalence of software such as Grasshopper (Robert McNeel & Associates 2007) and Dynamo (Autodesk 2011), but it may also be caused by a more direct conceptual link to the traditional design process, in which teams are used to making decisions about particular design “parameters” or “drivers”, rather than ceding increased control to a rule-based system. In practice, more architecture and engineering firms have already made the transition to parametric methods or are currently in the process of building capacity for increased participation in computational design. Due to a desire to positively influence design process and practice, this dissertation focuses mostly on parametric modeling methods.

2.2 The design space

Most current practices surrounding performance-driven design are thus based on the concept of the design space, which is described in Figure 2.4. For parametric design, the **design space** (or decision or possibility space) is a theoretical construct in which each instance within a parametric family can be described as occupying a single point, which is located with reference to the characteristics of that design. These design characteristics are most often called design variables or design parameters, which are used to label the axes of the design space. A design vector contains all of the particular settings that describe a single design, such that this numerical vector can itself represent the design. Since the term “parameter” is also commonly used to describe constants or settings within an optimization problem than can be adjusted for separate runs, this dissertation primarily uses the term “design variable” when referring to the elements of a design vector.

In conceptual building design, the design space is typically made of mostly geometric variables, and different options can be visualized accordingly. However, non-geometric properties can also be used as variables, since any independently controlled design decision can be turned into a variable. Such non-geometric variables may include equipment choices, material and envelope properties, or other related design decisions. These variables may be continuous, such as the length or width of a building, and are subjected to specific bounds, which limit the edge settings for a variable. They may also be discontinuous or discrete, which would occur when a design variable represents a decision between materials, for example. Exploring the design space involves changing these different settings to see what design possibilities result.

In addition to the design space as described above, individual designs can also be placed at a point in a theoretical **objective space** (or criterion or cost space), where the location refers to how well the design performs along different axes. In the objective space, the axes refer to numerical objectives, and the location of a design signifies the value of the objective function (also called the fitness, cost, value, or payoff function). Common objectives in conceptual building design include structural criteria such as material weight, stiffness, strain energy or compliance; energy criteria such as cost, operational energy, embodied energy, lifecycle cost or energy, carbon emissions, or resilience metrics (Hasik et al. 2017); daylighting criteria such as spatial daylight autonomy or glare (Reinhart 2014); or anything else that can be measured or simulated for buildings, including acoustics, views, circulation, thermal comfort, mobility, and others. Whereas the location in the design space is defined outright, the location in the objective space must be found, often through simulation at this stage of design. During the design process, objectives may generally be known as performance metrics.

The relationship between the design space and objective space is of great importance to engineers and performance-conscious designers. In reality, they combine to form a high-dimensional topography that can be explored to find designs that perform well and satisfy all other design requirements. Sometimes this

connected space, which includes both the design vector mapped in certain dimensions and the objective results in other dimensions, is together referred to as the design space, and often linked directly to the term design space exploration. Finding a satisfying solution through exploration thus involves moving through the design space with reference to the objective space. Formal optimization uses this framework as well, in which the design space is formed by the set of all possible solutions to a problem, as determined by the design variables and constraints. After formulating the problem in this way, optimization is an automated, algorithmic process for systematically moving through the design space and finding the optimal solution.

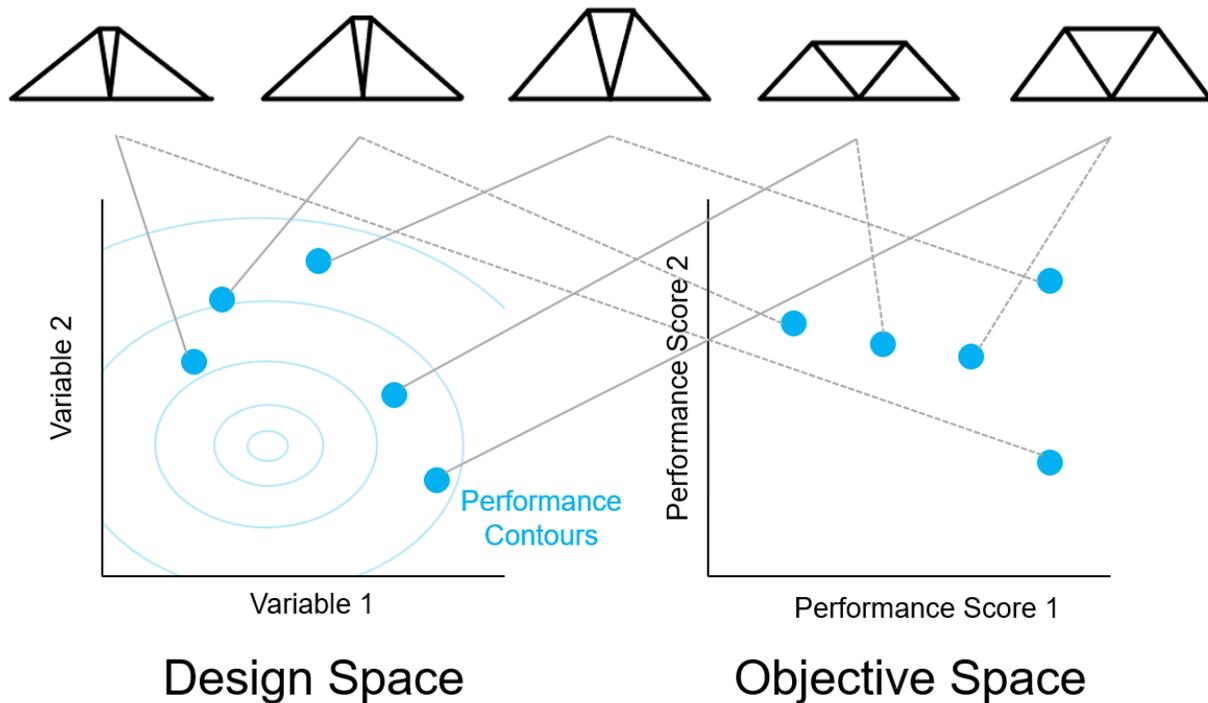


Figure 2.4: An example of a design and objective space, from a two-variable parametric truss, after Mueller (2014)

In design practice, the advent of parametric modeling (Monedero 2000) made it relatively easy to create a parametric logic and evaluate the different possibilities that reside within a design space, especially for early, conceptual design. When compared to traditional modes of design ideation, the presence of many different possibilities creates a more urgent need to systematically explore and consider the design space to find an overall “best” design (Shireen et al. 2017). Under these conditions, a reasonable next step was to combine visualizations of the design space with a presentation of designs in a “catalog” or “design-by-shopping” (Balling 1999) format to designers for selection (Stump et al. 2003). Along with the gradual advancement of performance simulations, the combination of data visualization and catalog presentation has significantly upgraded the traditional decision-making process in architectural design (Tsigkari et al. 2013). Formal automated optimization procedures have also been applied directly to such design situations, with a mixture

of results. The next few sections will describe specific contributions in early architectural and engineering design that range from the catalog method, to interactive methods using performance feedback, to full optimization.

2.3 Developments in optimization

2.3.1 From structural optimization

First, it is relevant to mention historical developments in optimization, particularly related to structural optimization, as many eventual building optimization techniques rose out of this legacy. The roots of mathematical optimization are traced to Newton and Gauss, who proposed methods for iteratively moving towards the optimal value of a function. The rise of computation inspired an explosion of interest in the topic, as computers gained the ability to rapidly execute calculations that allowed for practical applications of optimization. However, the desire to find the minimum amount of material required for a structure to carry a given load is found much earlier than these relatively recent computational developments. Initially, these problems contained analytical solutions, including Galileo's notable optimization of a cantilevered beam (1638), and Michell's (1904) ideal truss configurations for a cantilevered point load. As demonstrated by Beghini et al. (2014), such closed-form solutions are typically not possible for the complicated loading combinations and geometries found in buildings.

When computers began allowing for more numerical, iterative optimizations (Schmit 1981), three major types of structural optimization took form: size, in which discrete structural members are assigned efficient cross sections; shape, in which the connectivity or topology of the structure does not change, but the geometry can otherwise move around (Haftka & Grandhi 1986); and topology, in which continuum or related methods are used to generate optimized geometries without these other restrictions (Bendsøe & Kikuchi 1988). These methods have widespread application and success in diverse engineering fields including mechanical, aerospace, electrical, and others. However, researchers and practitioners alike are still wrestling with the most appropriate uses for optimization in building design, especially as digital fabrication presents opportunities for realizing complex or non-traditional geometries generated through optimization. Despite promising research in the area, structural optimization's application to large, conceptual decisions related to the overall structural system and massing has been largely limited.

2.3.2 Optimization as a black box for design

Despite some challenges associated with applying pure structural optimization to building design, the general approach of using optimization solvers to handle quantitative objectives is attractive to many architects and engineers. Optimization can quickly direct humans to high performance regions of the design space, which is a valuable form of guidance that has gained some popularity.

Part of this interest was catalyzed by the development of heuristic or direct search methods for architectural optimization and their implementation in various design software. As opposed to gradient-based methods, heuristic approaches such as evolutionary strategies (Beyer & Schwefel 2002), simulated annealing (Kirkpatrick et al. 1983), or particle swarm optimization (Eberhart & Kennedy 1994) lead to accessible techniques that can obtain results that are at least “pretty good”. Many practicing designers are not concerned with whether a design problem is convex, or if the presence of a local minima can be guaranteed, or even in some cases, whether or not there is fast convergence or convergence at all. If, as a practical matter, a solver can take in a variety of continuous and discrete variables, run over lunch or overnight, and return a solution that performs considerably better than the starting point or what the designers might have otherwise chosen, it becomes immediately useful. Each building design is custom, at least to some extent, and subject to time and financial pressures, meaning that the details of optimization are often less important than its results. While some care must be taken to ensure designers are not completely misusing optimization algorithms, and have enough understanding to select the most efficient or effective ones, they are often used essentially as black box solvers (Wortmann & Nannicini 2016).

2.3.3 Multi-objective optimization

Yet in conceptual building design, many problems are inherently multi-objective, which requires designers to prioritize their numerical goals when using any type of optimization. Except in rare cases, there are often interrelated and conflicting quantitative objectives that architects and engineers may want to consider. As such, it is worth describing the approaches from different engineering fields when multiple objectives are considered. Marler & Arora (2004) provide a general overview of multi-objective optimization in engineering. It segments these methods into *a priori*, which uses explicit weights to combine all objectives into a single composite function, *a posteriori*, which requires finding Pareto-equivalent solutions in an automated way, and *interactive*, which involves humans in the process (see Figure 2.5). More recently, Coello & Romero (2013) present a review and basic taxonomy of evolutionary multi-objective approaches, which are of particular relevance to building design problems that are not well behaved. Méndez Echenagucia (2013) explores several multi-objective problems in architectural design, and Brown (2016) provides a thorough look at the direct application of multi-objective optimization methods to the conceptual design of buildings and other structures.

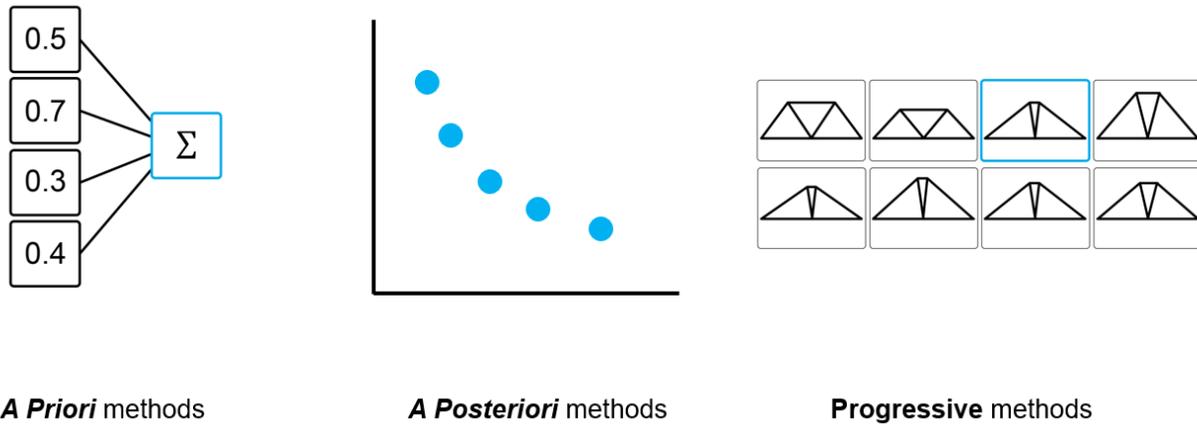


Figure 2.5: Approaches to multi-objective optimization, after Marler & Arora (2004)

When searching through the literature related to building performance, the last twenty years have seen strong growth in the research and application of multi-objective algorithms to building design at various stages (Evins et al. 2012). This development is encouraging, as these techniques seem naturally suited to early design problems. Direct applications of multi-objective optimization have helped designers lower the costs and energy usage of buildings, improve their layout, increase thermal comfort, and otherwise navigate and balance various objectives. However, the complications that arise from using optimization in a field that is not exclusively driven by quantitative metrics cannot be eliminated by switching to multi-objective techniques, or even by finding better or more efficient algorithms.

To make these techniques useful, researchers must ask fundamental questions about the role of computers in early design, and how computation can be used to augment the design process rather than assume control over it. While doing so, it is nevertheless important to maintain the multi-objective lens afforded by these optimization techniques, and to consider salient aspects of their framework. Objective prioritization, managing harmonies or tradeoffs between objectives, satisficing, non-dominated solutions, and other concepts from multi-objective optimization are retained by this dissertation as it explores ways to pursue multi-objective approaches to design without the rigid requirements of traditional optimization.

The following sections detail research contributions that attempt to bring the human and computer together during building design exploration, while drawing from optimization techniques and adapting them for the task at hand. Priority will be given to work that is multi-objective in some way, since it is of particular relevance.

2.4 Need and potential for optimization in building design

Researchers have long acknowledged the potential for optimization in building design and the use of computers for generative and simulation purposes (Radford & Gero 1980). While some of this initial work was speculative, before long researchers were proposing integrated workflows for considering performance and simulation during the generation of building form (Shea et al. 2005). In recent years, a variety of surveys, case studies, and similar methodologies have demonstrated a need for optimization, usually justified by the long list of contemporary design goals and requirements. Related research has also established the significant potential of optimization-based methodologies for improving the built environment. Attia et al. (2013) assess gaps and needs for integrating building performance optimization into the design process, while focusing on examples for net zero energy buildings. Bradner et al. (2014) present a survey of how designers are already using optimization—often as a starting point for future exploration, rather than an end goal, indicating that many human designers are seeking a supportive process rather than a final answer. Attia et al. also note that poor user interfaces, rather than poorly performing algorithms, may present the most significant barriers to their further adoption. Cichocka et al. (2017) provide another survey of practicing architects, a majority of whom (78%) understand multi-objective optimization to be more relevant to early design than single-objective optimization. Designers in this survey also most commonly identified structure, daylight, and geometry as the domain of optimization, which are each considered in this dissertation.

At the same time, there are obvious differences between related engineering fields that have entirely embraced optimization and the more conservative ecosystem surrounding building design. For one, many buildings are essentially custom, as stated before, and this is especially true for buildings designed by architects. Although there are certainly examples of standard solutions involving prefabricated components or systems that are applied repeatedly, new buildings often have particular constraints related to program, site, climate conditions, and other factors that require specific attention rather than a one-size-fits-all approach. When compared to a car, plane, or other product that is manufactured repeatedly, there may be increased effort and diminished returns for formulating and solving an optimization problem that is only relevant to a particular building.

The construction of buildings is also less controlled or centralized, as different contractors bid for a project and in most cases multiple entities are involved in building delivery. The organizational and legal structures of the industry and related factors cause many design decisions to be based on standard elements and geometries, such that commonly available erection or fabrication techniques can be readily employed to realize the result. The scale and complexity of building design, which requires careful coordination between the architect and at least a few separate engineering disciplines, also makes it difficult for a single entity to fully control the process and drive it computationally, without input from a larger design team. As such, considerable research has been conducted into how designers can interact with optimization processes in

ways that are beneficial in the specific context of designing buildings or structures. The next few sections will describe related research aimed at making optimization a productive part of the early design process.

2.5 Integration of live feedback and optimization into the building design process

Earlier, it was mentioned that formal multi-objective optimization procedures have three main ways of achieving prioritization: *a priori*, which involves selecting weights for a composite function and conducting single-objective optimization on this composite function; *a posteriori*, which requires finding Pareto-equivalent solutions and selecting the best design afterward; and *interactive*, in which designers express preferences while an optimization is running. While this characterization is useful for the direct application of multi-objective optimization, it is likewise advantageous to present a related but slightly different taxonomy for a more generalized multi-objective design process in conceptual building design. In this section, recent contributions are divided into workflows that focus primarily on interactive feedback; workflows that are semi- or entirely automated following from an optimization algorithm; and workflows that attempt to combine both feedback and guidance, often facilitated through design catalogues or other means of data visualization.

2.5.1 Interactive design methods based primarily on feedback

In recent years, a number of interactive tools have been developed for specific applications of simulations in building design. Some examples for daylighting include Lightsolve (Andersen, Gagne, et al. 2013; Andersen, Guillemin, et al. 2013), an interactive daylighting tool based on an expert system, and Accelerad (Jones 2017; Jones & Reinhart 2018), which is a GPU-accelerated version of RADIANCE that allows for real-time simulation feedback for interior illumination and visual discomfort (see Figure 2.6). These interactive tools generally rely on a combination of data analysis and rendering of interiors to communicate design information.

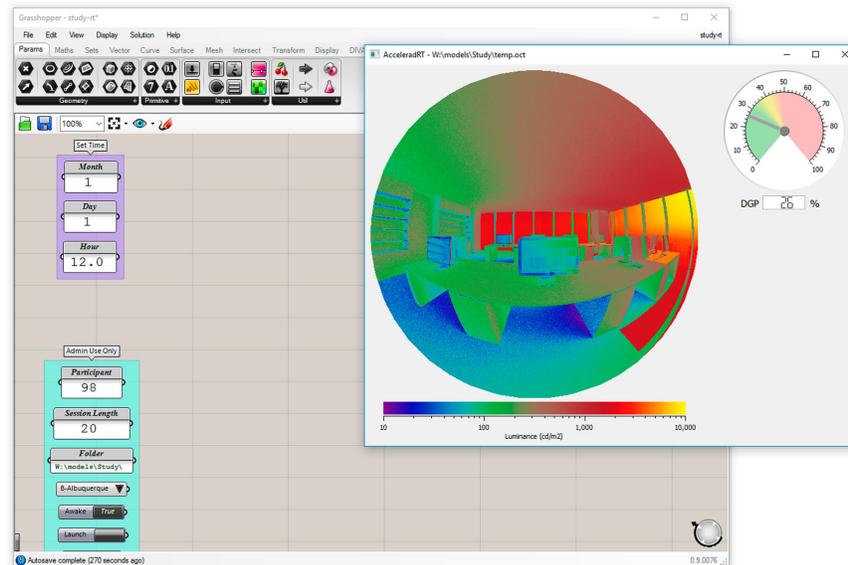


Figure 2.6: *AcceleradRT* by Jones & Reinhart (2018), which provides real-time false color luminance maps and daylight glare probabilities

On the energy modeling side, it seems that many parametric simulations involve discrete settings for the envelope, heating system, and other factors that affect energy performance (Simmons et al. 2015). As such, there is considerable literature on expert systems or other computational advisors for making these decisions (Attia et al. 2012; Abaza 2010), often for pre-calculated designs, in addition to papers involving energy optimization. Sanguinetti et al. (2010) acknowledge the need for rapid feedback and couple analyses for energy, visual comfort, and investment payback to generative processes. There are also a few widely used plugins integrating energy and daylighting analysis into parametric geometry environments, including Archsim (Dogan 2014), DIVA-for-Rhino (Solemna 2012), and Ladybug Tools (Roudsari & Mackey 2013). Rezaee et al. (2015) have developed a plug-in for Revit that helps inform early design decision-making, in reference to uncertainty at this stage and the risks associated with it. While these are not live, they dramatically improve the connection between geometry and simulation for parametric exploration.

In structures, early examples of analysis programs that provide live feedback while exploring geometry or other design properties include Arcade (Martini 2006) and a standalone program developed by Clune (2010) (see Figure 2.7). A plug-in for Grasshopper called Karamba (Preisinger 2014), which layers finite element modeling for linear and shell elements directly on top of Grasshopper geometry, has gained popularity and can generate essentially live feedback depending on the complexity of the model. Other live structural tools exist for specific applications, including formfinding using particle springs (Kilian & Ochsendorf 2006), or RhinoVAULT (Rippmann et al. 2012a) (see Figure 2.8), which is based on Thrust Network Analysis (Block 2009). There is also a long legacy of graphical methods for equilibrium approaches

in the design of structures, which are continuously being updated and employed in innovative ways, often computationally (Fivet & Zastavni 2013; Beghini et al. 2014; D'Acunto et al. 2019).

In practice, it is also common to achieve some measure of feedback by linking parametric design environments directly to commercial structural analysis software. This can be done directly in programs such as recent versions of Robot Structural Analysis (Autodesk 2016), using Dynamo. Many engineers also use third-party plug-ins to connect with their preferred analysis software, including Geometry Gym (Mirtschin 2015), Konstru (Konstru 2017) and Salamander (Jeffries 2019). Firms with in-house software development capabilities can write their own interoperability tools, such as the adapters included in the Buildings and Habitats objects Model (BuroHappold Engineering 2018). While these tools have workflow benefits throughout the design and documentation process, they rarely achieve real-time feedback, and are often not suited to live, interactive optimization techniques.



Figure 2.7: An early real-time structural analysis and design environment by Clune (2010)

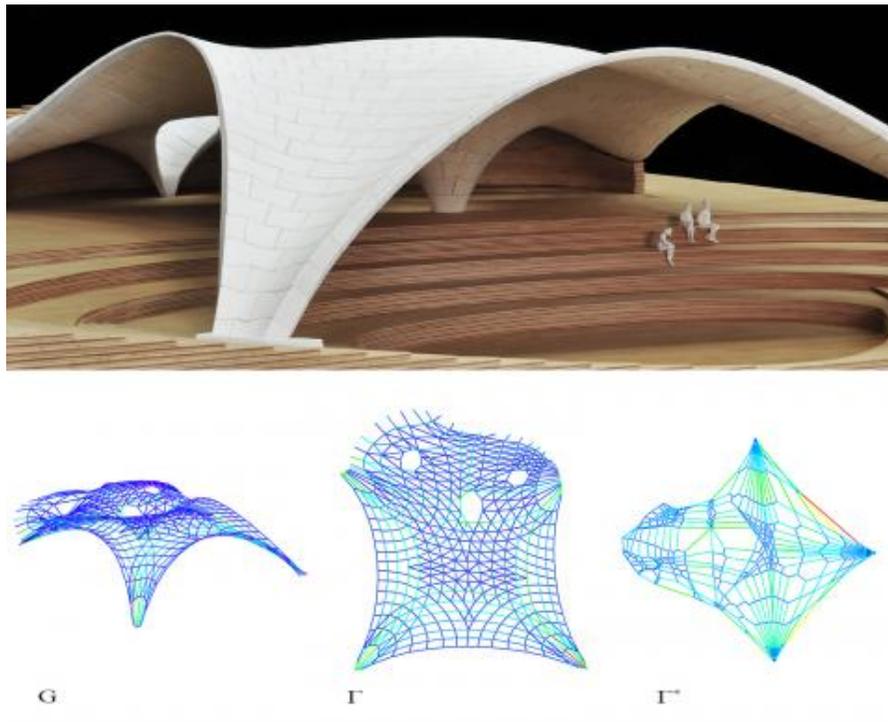


Figure 2.8: An example of a design generated by RhinoVAULT, which enables rapid funicular form-finding (Rippmann et al. 2012b)

Real-time feedback has also been explored in other parts of the design process, such as in Balali et al. (2018), which presents an approach for real-time cost estimation of change orders during construction, which is visualized in VR environments.

More generally, Fisher (2012) describes workflows for fast feedback during design in practice, but mostly focuses on rapid simulations or quick analytical calculations rather than longer simulations or surrogate model predictions. Haymaker et al. (2018) introduce and implement Design Space Construction, which is an extensive conceptual framework that guides teams through problem formulation, the generation of alternatives, impact analysis, and value assessment. Ritter et al. (2013) present a framework called the Design Space Exploration Assistance Method (DSEAM), which intends to gather information on well-performing solutions rather than one single, optimal solution. This system also describes part of the workflow that are automated or semi-automated, which is relevant for the next two categories.

In general, these tools are successful in providing designers with considerable freedom, which is an even greater asset in the cases where the feedback is fast enough to be provided live. Their systematic approach to generating, simulating, and organizing potential design options has many clear advantages, and this dissertation is indebted to their general concept of linking parametric geometry to performance simulations

to make informed design decisions. However, directions to explore are often left totally to the designer, which misses an opportunity to provide some level of guidance at the same time.

2.5.2 Design methods based primarily on optimization

Optimization techniques are one way to provide this guidance. Particularly from the engineering side, there is a large body of existing research involving direct applications of optimization to buildings. This section will mainly give sources for large reviews rather than specific contributions. Evins (2013) provides a review of computational optimization applied to sustainable building design, while mapping trends and the rise of multi-objective optimization as a preferred approach. Asadi & Geem (2015) review the use of meta-heuristic methods for sustainable building design. Waibel (2018) covers many topics related to simulation-based optimization of buildings, while Waibel et al. (2019) provide a benchmarking for different optimization algorithms for building energy optimization. Shi et al. (2016) give a review of optimization applied to energy efficient buildings from the perspective of architects. Shi et al. (2017) explore form generation and optimization at urban scale. While most optimization considered in this section focuses on early design, Evins (2015) demonstrates optimization while considering the building design, system design, and system operations together. In order to analyze which approaches tend to be the most effective, Ren et al. (2011) consider the merits of both MDO and Multi-Agent Systems (MAS) for building design, and Hamdy et al. (2016) compare the performance of MOO algorithms for sustainable buildings. While most of the citations in these reviews are academic, practitioners have also demonstrated the use of optimization in the design of buildings and other structures, such as Luebke & Shea (2005).

These contributions make clear the benefits of optimization, which can lead to high-performing solutions. However, most of the automated approaches do not have the flexibility of allowing for human interaction and preference expression during design.

2.5.3 Interactive design methods using optimization

To mitigate this problem, researchers have developed a broad class of algorithms designed for user interaction while they are running. By having a human in the loop, optimization becomes considerably more intelligent when managing soft design criteria and even some of the constraints, settings, and boundary conditions that might be difficult to dictate or fully understand at the beginning of an optimization run. Scott et al. (2003) provide an early overview of interactive approaches, while testing certain aspects of interactive optimization. In Scott et al., the authors describe finding the best division of labor between humans and computers to be the most important research question, which still motivates portions of this dissertation. While optimization methods can be either gradient-based or heuristic, many of the popular interactive algorithms fall into the latter category.

Takagi (2001) gives an overview of interactive evolutionary optimization methods, which have become particularly popular in the design of buildings and other structures. Parmee et al. (2000) describe ways in which interactive evolutionary optimization can be used to find satisfactory solutions subject to multiple objectives, and Cvetkovic & Parmee (2002) consider the role of designer preferences in interactive optimization. Hofmeyer & Davila Delgado (2015) demonstrate a method for coevolving spatial and structural building designs using topology optimization. Coley & Schukat (2002) offer an example of interactive evolutionary optimization, and also explicitly name “architectural appeal” as an objective that could be managed through the visual presentation of designs during optimization. Machwe et al. (2005) pursue a similar goal, but attempt to learn designers’ subjective aesthetic preferences while an optimization is running.

In architectural and structural design, the idea of interactive evolutionary optimization has been implemented and demonstrated on architectural case studies by paraGEN (von Buelow 2012; Turrin et al. 2010; Turrin et al. 2011), and structureFIT (Mueller & Ochsendorf 2015), which is a tool for 2D truss design that also provides a live interaction mode (see Figure 2.9 and Figure 2.10). The main functionality of structureFIT has been generalized within parametric environments for performance evaluation in the tool Stormcloud (Danhaive & Mueller 2015), which was developed for the Grasshopper environment. Biomorpher (Harding & Olsen 2017) for Grasshopper uses a similar conceptual framework but also adds tools for visualizing and grouping possible design alternatives.

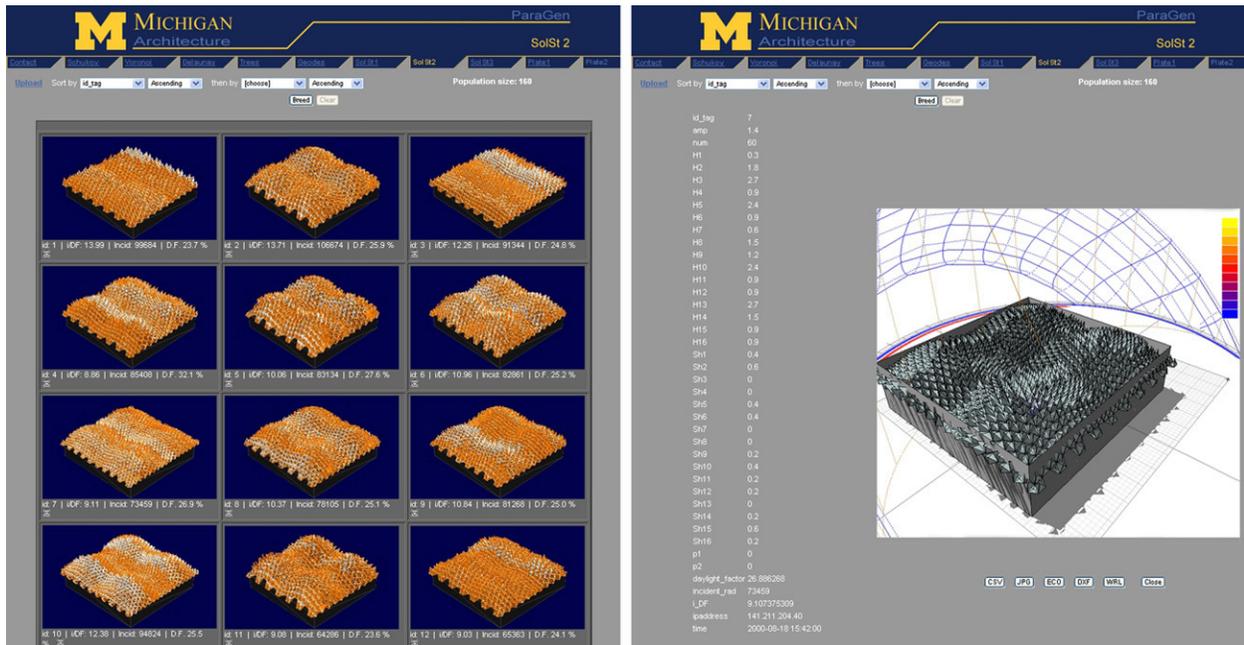


Figure 2.9: paraGEN by Turrin et al. (2011)

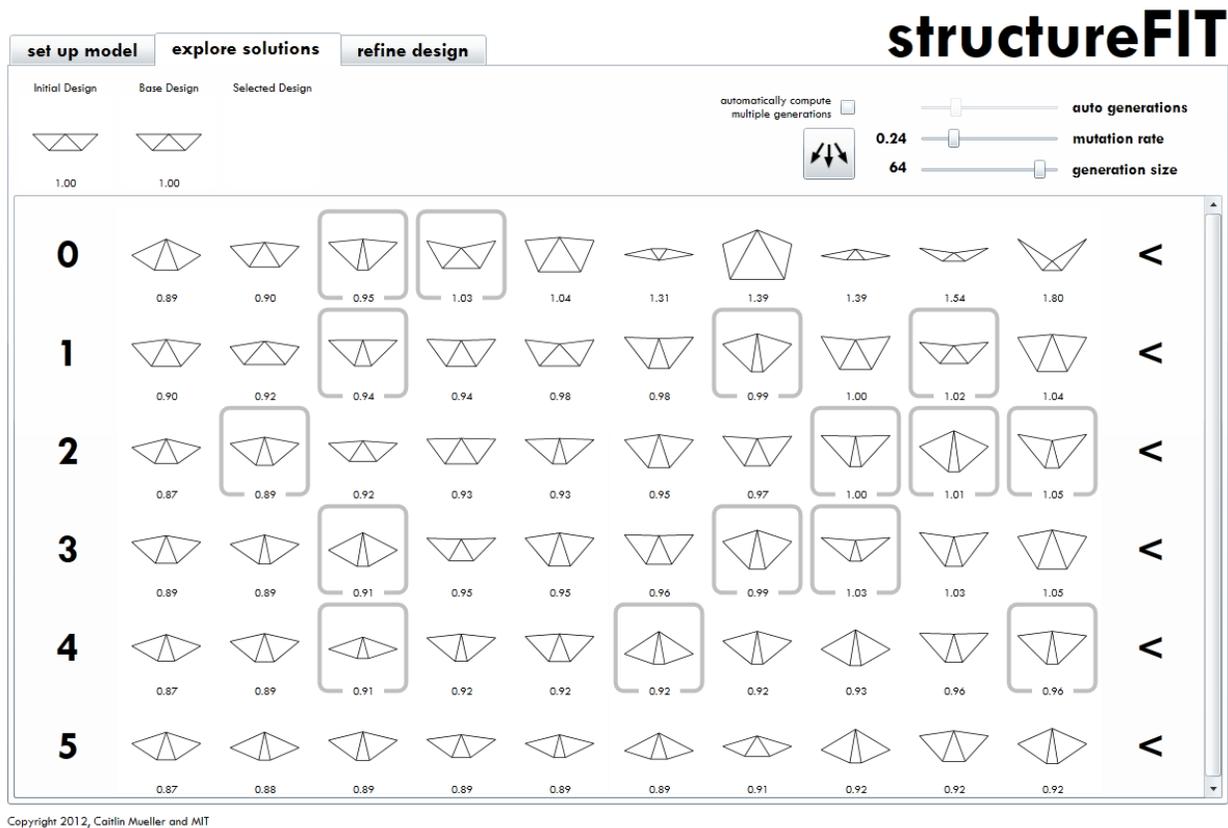


Figure 2.10: structureFIT by Mueller & Ochsendorf (2015a)

Many of these tools are specific to a discipline, and only a few have multi-objective capabilities, either through organizing generated solutions by objective or manually building a composite function. Others could be further extended using model-based methods to replace slow simulations and afford a wider variety of analysis types.

2.5.4 Design by shopping and the catalog approach

In each of these interactive methods, particularly the ones that require direct engagement by human designers, some type of visualization is necessary. This has led to the development of a common methodology called the design-by-shopping approach, as described earlier, in which each design is arrayed in a grid, with a representative glyph of its geometry and information about how well it performs. These catalogs are accompanied by various data visualization techniques. Sometimes design-by-shopping approaches are simply a mechanism within an interactive optimization process, while sometimes they are standalone, and a sampling technique is used to generate the solutions rather than optimization. In either case, it is worth mentioning notable contributions that make use of this technique, in addition to the interactive evolutionary optimization workflows mentioned earlier. The browser-based tool Design

Explorer (Thornton Tomasetti 2016) was created to enable design-by-shopping using previously generated Rhino and Grasshopper models, and later developed into the tool Thread. Ashour (2015) demonstrate a method for visualizing and organizing the designs produced by a multi-objective optimization tool. Gerber & Lin (2014) integrate energy, design, financial considerations by linking simulation tools together and visualizing the results. Freitas et al. (2015) provide an optimization of façade integrated PV panels, while also giving results in this manner.

2.5.5 Artificial intelligence and immersive environments

In addition to the interactive optimization methods described here, there are other ways of automating design tasks that can be integrated into the process. Two prominent examples include Gerber et al. (2017) and Gerber et al. (2015), which pioneer the use of multi-agent systems to solve design problems. While these techniques show exciting promise for the future, this type of artificial intelligence for building design is in its infancy. As such, this dissertation focuses primarily on the integration of human intelligence and computation through interaction.

2.6 Related tools for interaction and optimization in architecture

While many researchers have developed their own workflows and implement them on demonstrative case studies, a few researchers have worked explicitly towards open-source optimization tools with the goal of mass participation. Some of these tools have already been mentioned in the last few sections, since their intellectual contributions fit into specific categories related to interactive optimization (paraGEN, structureFIT, stormcloud, Biomorpher), or design-by-shopping (Design Explorer). This section presents additional tools that help integrate feedback and guidance into the conceptual design process, particularly within parametric environments. Today, many of these tools enjoy wide use by architects and engineers. For optimization in parametric environments, designers can access Galapagos, Goat (Flöry 2013), or if multiple objectives are being considered, Octopus (Vierlinger 2012) or Optimo (Rahmani 2014). (Wortmann 2017a) recently released Opossum, which enables surrogate-model based optimization. Outside of parametric software, Lawrence Berkeley National Laboratory developed GenOpt (Wetter 2001), which is used for minimizing objective functions that are evaluated by an external simulation program. For general engineering design, it is also common to use modeFRONTIER (Etesco 2002), which enables optimization workflows in addition to data analysis and post-processing of results.

It is clear from this overall review that there are viable existing optimization components available for parametric environments. However, opportunities remain to improve on the state of the art in building design. Design Space Exploration (DSE) contains a variety of data-driven tools that allow for catalog generation, multi-objective optimization, objective function approximation, clustering of designs into families, interactive optimization and other approaches, connected directly to parametric design software.

Portions of DSE were developed by the author in coordination with this dissertation, and will be described in Chapter 6 in more detail.

2.7 Improving interactive parametric workflows using data science

This dissertation builds on the many contributions described above regarding interactive optimization and the synthesis of aesthetic, visual, and otherwise qualitative design goals into design processes that are also managing quantitative, technical goals. Primarily, it is concerned with how data science techniques can be combined with interactive optimization to stimulate creativity while enabling live, multi-objective exploration. The conceptual justification for why data science can be a valuable tool is that a parametric design space is essentially a large data set waiting to be analyzed. Unlike in other “Big Data” applications, the design process involves formulating the parametric model itself, generating that dataset iteratively, and exploring it dynamically, rather than trying to gain insights into an already existing phenomena using data mining.

Nevertheless, many typical data science activities are still applicable in a design context. For example, regression can be used to focus designers on important areas of the design space or predict performance more quickly than actual simulations. Classification can help organize designs to make meaningful decisions, and dimensionality reduction can modify or transform the design space for more effective and creative synthesis. Certain statistical approaches can make all of this information concerning design possibilities more robust. Previous contributions that involve these data science applications on existing engineering or architectural design scenarios are described next.

2.8 Related data-driven techniques in engineering

The broadest application of data-driven methods for engineering design is the use of surrogate modeling, sometimes called metamodeling, to improve the performance of optimization processes that rely on heavy external simulations. As demonstrated by review papers spanning back to nearly two decades ago (Simpson, Peplinski, et al. 2001; Sóbester et al. 2014), surrogate modeling has been used for many years. While a more substantial discussion of the merits and deficiencies of surrogate modeling techniques for architecture will be provided over the next few sections, the general approach shows considerable promise for both automated and interactive optimization.

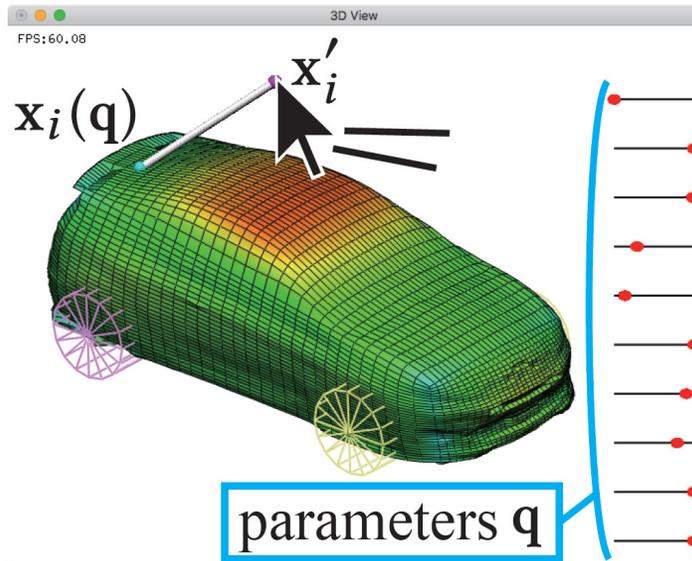


Figure 2.11: Exploring generative 3D shapes using autoencoder networks by Umetani (2017)

More recently, researchers have proposed other applications of data science and machine learning to parametric engineering design. For example, Umetani (2017) introduces an automated parameterization approach that robustly constructs a consistent parameterization for machine learning of 3D shapes, and then demonstrates the construction of a shape manifold using an autoencoder. The end result is a design interface that allows for essentially any major car shape to be found by adjusting ten sliders, as shown in Figure 2.11.

2.9 Related data-driven techniques for building design

Certain surrogate modeling techniques have also been applied to building design. Machairas et al. (2014) gives an overall review for building optimization research, and contains a full section on papers that use surrogate-based methods. Ramallo-González & Coley (2014) created an adaptive optimization process that uses surrogate modeling in the beginning and then transitions to more accurate simulations when getting closer to a final solution. Brownlee & Wright (2015) describe an implementation of a surrogate modeling enabled NSGA-II multi-objective algorithm for constrained, mixed integer optimization problems in building design.

Other data-driven techniques in parametric design attempt to understand which variables tend to control performance, and thus require special attention, while often also predicting performance objectives during exploration. Asadi et al. (2014) use linear regression to predict energy usage for early building design. Kristensen & Petersen (2016) provide guidance on using sensitivity analysis to understand which input parameters affect models in general for building energy, an approach that can also be used in parametric

design situations. Brown & Mueller (2017a) adapt methods from quality control and manufacturing to gain a rapid sense of parametric variable importance with only a few simulations, shown in Figure 2.12. In this image, a quick analysis of variable importance indicates to the designer of a space truss roof which variables tend to most affect the efficiency and steel required to build the roof, and where depth and curvature can be manipulated visually with less effect on performance. In general, it appears that more depth is required near the central supports, which resists the large global moments, especially compared to the backspan of the cantilever. More involved methods for interrogating the importance of design variables are described later in this dissertation. Samuelson et al. (2016) also use a sensitivity analysis method to find important building parameters.

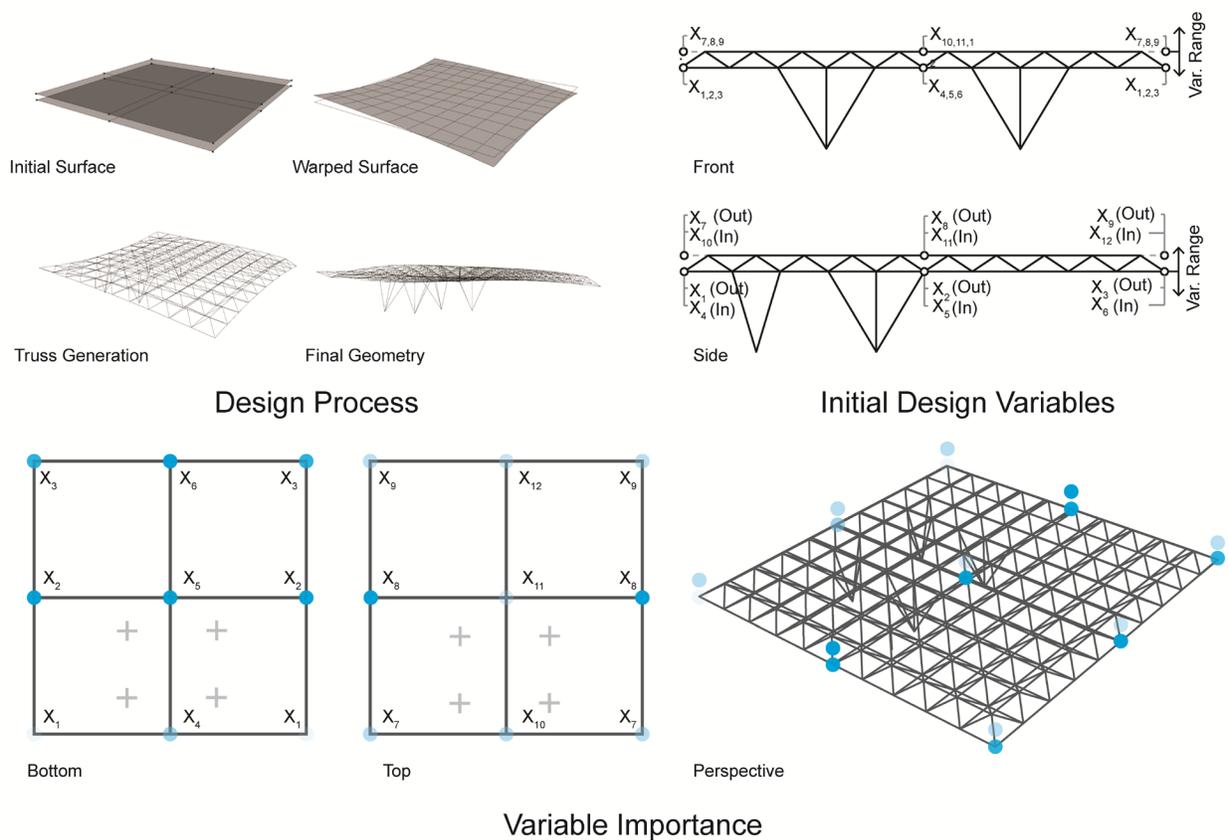


Figure 2.12: A description of the design space for a trussed roof design example, along with the results of a variable effects analysis using Taguchi methods. The strength of the color represents the importance of each variable based on its ability to influence performance, which here is the weight of the structure (Brown & Mueller 2017b).

Various predictive models have been proposed for fast guidance on building performance, across building engineering disciplines. Nault et al. (2017) predict solar daylight potential at the neighborhood scale. Østergård et al. (2017) settle on a two-step solution for making practical use model-based methods in design: a significant amount of model preparation and simulation by engineers and data scientists, and then

a multi-agent meeting with architects to make design decisions. Reynolds et al. (2015) introduce a surrogate modeling method for providing information about a predicted structural system from early massing models of large buildings. Geyer & Schlüter (2014) automatically build metamodels for energy prediction in new construction and retrofits, such that designers can quickly change aspects of the design and receive real-time performance feedback.

It is also worth mentioning a few innovative examples of data methods in early design, even if some of them do not fit neatly into variable importance calculation or objective value prediction. Danhaive & Mueller (2018) demonstrate a technique for modeling deformations of shells during live design. Conti & Kaijima (2017) apply Bayesian inference to help designers understand the probable outcomes of various design variable settings. Harding (2016) demonstrates the use of dimensionality reduction for parametric design visualization, such that a high dimensional design space can be organized on a 2D access such that similar designs are nearby, but possibilities can still be presented on a page or screen. Reynolds et al. (2018) have experimented with machine learning for aesthetic preferences in design, a task that could point towards smart suggestion engines for early design. Fuhrmann et al. (2018) demonstrate a method for generating, representing, and evaluating a large solution space of networks in equilibrium using self-organizing maps and a Uniform Manifold Approximation and Projection (UMAP) algorithm.

While many of these examples provide a compelling case for the applicability of data science in design, and also proposing innovative new approaches, they have not collectively integrated data methods into early stage design with enough functionality and accessibility to stimulate widespread usage. Significant opportunities remain to make additional contributions towards this ultimate goal.

2.10 A discussion of data science applications in this dissertation

While a variety of possible data science applications exist within contemporary parametric workflows, this dissertation embraces two main concepts that can enhance the typical computational design process. These concepts are:

- 1. Surrogate modeling for interactivity and exploration**
- 2. Data analysis for interpretability and accessibility**

The foundations of these concepts have been developed for numerous engineering applications, as demonstrated in this literature review. The first idea is that predictive surrogate models can rapidly approximate heavy simulations, which sacrifices some accuracy while speeding up performance feedback to essentially real-time rates. Although these techniques are commonly employed in optimization settings, where surrogate models replace the hundreds or thousands of required simulations for an optimization run, this dissertation explores ways in which the use of real-time feedback can lead to more interactive, collaborative, performance-driven workflows. Thus, the idea of surrogate modeling for model interaction

itself is not a contribution, but it is central to many of the proposed and tested live design strategies. The process of using parametric modeling techniques to generate a dataset and then using that data to modify, interpret, and ultimately explore the design space, often with the assistance of surrogate models, is consequently a first step to many of the proposed workflows in this dissertation.

Surrogate modeling has many advantages and disadvantages, and although these are mentioned specifically in other sections, it is worth a brief discussion of possible obstacles to its widespread usage in design. First, since interactive performance modeling is being proposed as an alternative or improvement to the design space catalog, in which a designer sifts through a set of previously generated designs, one might raise the objection that the pre-computation required to generate a surrogate model takes too long, or severely restricts its utility. In other words, it may be more useful to simply choose from previously simulated options and rerun new simulations when necessary rather than use surrogate modeling techniques and run the risk of being guided in the wrong direction by an inaccurate interpolation between simulations.

While there are conceptual similarities between an interactive visualization tool that chooses from pre-computed options and using those options to generate a predictive model, this dissertation demonstrates clear advantages of having a live geometry that can be manipulated based on guidance from a computer. Furthermore, in design practice, a few timescales have been observed for the use of these techniques: colloquially, these might be called the weekly, overnight, lunch break, and coffee break timescales. For real designers in either architecture or engineering, there are times during the daily rhythm in which humans need to do something away from their desks, at which point computers can run background simulations and build surrogate models which will augment exploration when they reengage. In this way, running simulations overnight and returning in the morning to a live, dynamic, performance-enhanced design environment is both feasible and beneficial for many design processes. Furthermore, ongoing and future research is currently directed at building progressively more accurate or localized surrogate models during exploration or generalizing data models that can be shared by certain common design problems. These techniques will be discussed further in future work, but they have the potential to drastically reduce the need to rerun many simulations every time an aspect of the parametric model changes, which would greatly improve the usability of surrogate models for design in general.

To demonstrate how a surrogate modeling technique could function in a live design setting, consider the example of a parametric exploration of rooftop photovoltaic (PV) energy harvesting potential for a courtyard building in Figure 2.13 (Brown & Mueller 2017b). For this geometry, the designers seek to explore the massing of a typical courtyard building while understanding both visual and geometric implications. First, a parametric model was formulated which allows for moving the corner points of the roof in the z -direction. Different variable combinations allow for tilt and double curvature, while also creating self-shading and otherwise changing the amount of solar radiation incident on different roof surfaces. For visualization purposes, the design space was dimensionally reduced such that different design

options, which contain particular relationships between the variables, can be represented along a single axis that sweeps through various geometries.

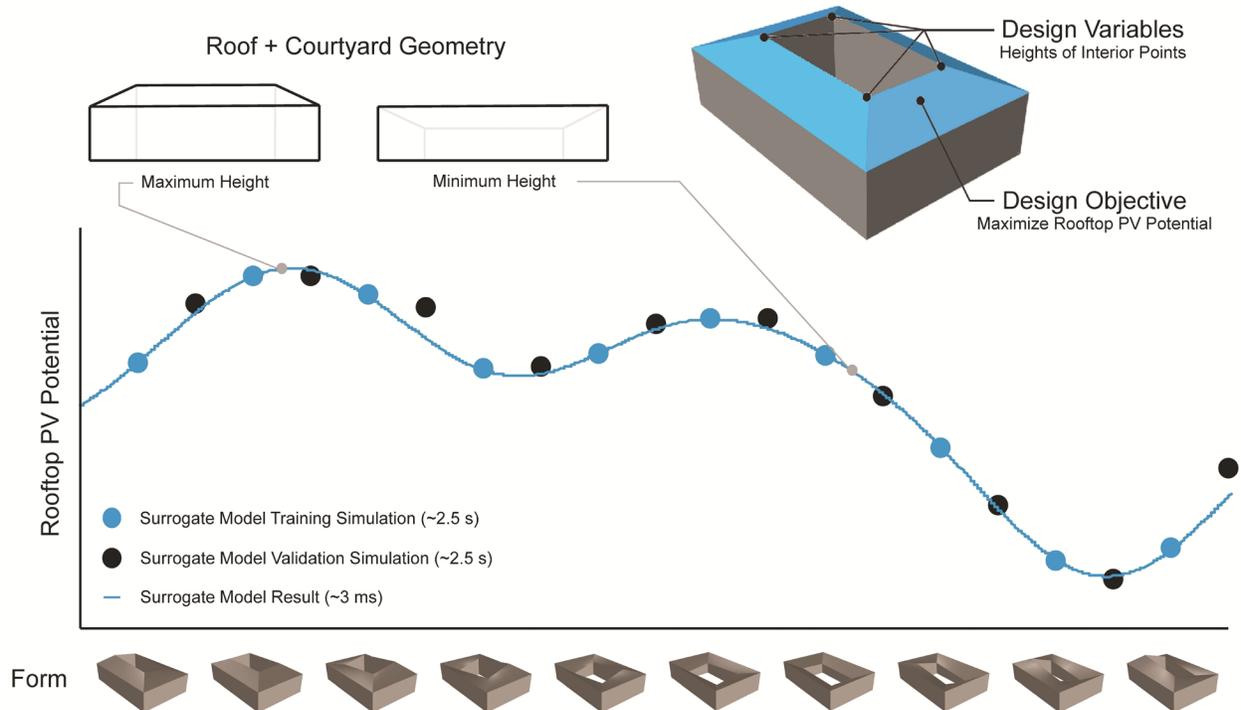


Figure 2.13: A surrogate model of rooftop PV potential, which enables performance-based exploration of form in real time with only small errors. Although the degree of error depends on specific problems and methods, this approximation of rooftop PV potential closely matches the simulated results (Brown & Mueller 2017b).

As the pitches of the roof sections are manipulated, the amount of energy that can be produced by a photovoltaic installation on their surfaces changes, as demonstrated by the associated graphic. This effect on annual PV potential was modeled using the parametric plug-in Archsim, assuming a location in Boston and a north-south orientation of the long axis. Each simulation takes approximately 3 seconds on a standard desktop computer. The author conducted 20 simulations and used the data points to train and validate a surrogate model of the objective function using ensemble neural networks, after which the approximation provided feedback effectively in real time (within a few milliseconds). While significantly more data is required for a more complicated problem, this brief study demonstrates a typical relationship between the computation required and the predicative abilities of surrogate modeling methods. The approximate results for this problem are fairly accurate, and the faster response time makes the difference between a set of static options versus a truly interactive tool, an effect that is magnified for more complicated designs. As such approximation techniques become more viable for early-stage exploration, the set of side-by-side,

static options of the design catalog can give way to a more dynamic, interactive model with live performance feedback in multiple dimensions.

A second set of limitations for surrogate modeling relates to their often-opaque structure and potential for misuse. While advanced data techniques are remarkably effective at describing a variety of relationships, their mechanisms can be hard to understand. In addition, certain problems in architectural and engineering design are very difficult to model accurately. Especially in cases where both are true, it may be counterproductive or even dangerous to the design process to hand an essentially meaningless black box model to a designer for exploration. As such, it is important to acknowledge that surrogate models are not a perfect solution for every design problem, and researchers and tool developers must work to effectively communicate modeling errors and their implications to designers.

Nevertheless, surrogate modeling has been demonstrated to be effective in representing many types of simulation results associated with building design, to a degree of accuracy that is useful for making early design decisions in the face of large uncertainties. These uncertainties refer to both the accuracy of the simulation models themselves but also the specificity of design assumptions that may change later on, such that designers often need only a relative understanding of which design options are likely to be “better” than others, and by roughly how much. In these situations, early stage performance feedback from surrogate models can be extremely useful.

The second general concept explored is the notion of interpretability, which becomes significant when computational methods can generate many options and pair each considered design option with multiple simulation feedback streams. In this context, the historical requirements of early design ideation are in some ways inverted. When the main goal of generating creative possibilities is achieved, at some point the pertinent question becomes: how does one sift through all the possibilities and find the “best” design? This need for interpretability is also magnified by the comparatively passive ways in which parametric methods are used to generate design options. In more traditional design processes, although designers would likely have fewer options to consider, they might sketch out each option out directly and deliberately, while thinking through the implications of each one.

In a mixed human-computation process the designer would certainly retain authorship and deep knowledge of the entire parametric setup and workflow, but would not necessarily have considered each option during its creation. In order to be better acquainted with these parametric possibilities, data science techniques can assist in the interpretation of broad trends within the design space, which contextualize individual options and lead to a richer experience than simple, repeated, side-by-side comparison of generated designs. Improved interpretation capabilities can also make these methods more accessible by reducing the potential for designers to be overwhelmed when first encountering parametric methods.

As a brief example of the benefits of data science methods for design space interpretability, consider the following 11-bar truss loaded at its three interior nodes (see Figure 2.14). The 11-bar truss is a simplified design example, but trusses are found in many real-life structures, including roofs and bridges. The geometric layout of a truss has a strong relationship with its performance, which makes visual exploration of form valuable in this case. The goal of this design process is to find a minimum (or at least a low) weight design that also responds to secondary requirements that are difficult to communicate to the computer. There are three design variables, which correspond to the width of the corner node, the inverse height of this node, and the height of the central node, with symmetry enforced. The height of the central node is defined with reference to the starting point of the corner node, meaning that high values always generate deeper trusses. This eliminates many unrealistic truss configurations in which the corner and central nodes are on different sides of the straight chord of the truss. The 11-bar truss problem was selected because the outcomes should be intuitive to structural designers—the depth of the truss affects its performance, and there are local minimums both above and below the constant chord.

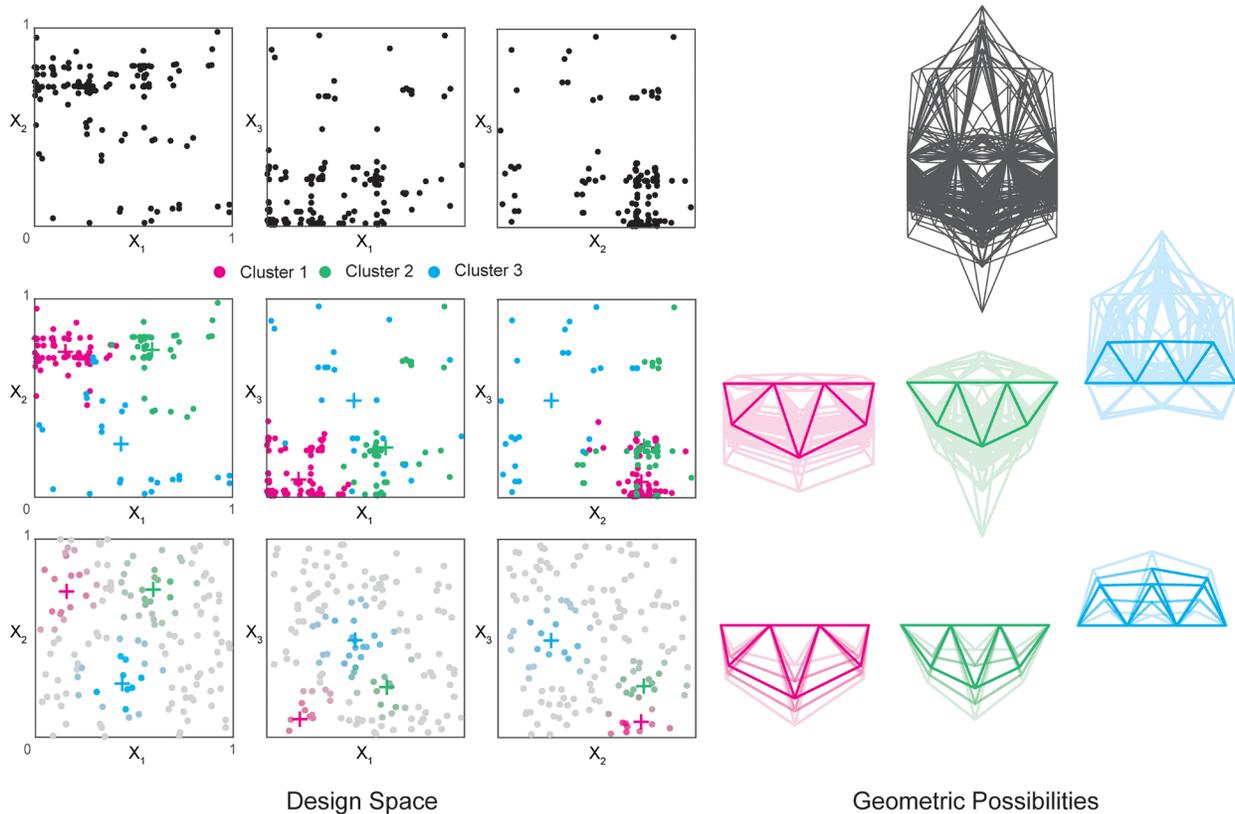


Figure 2.14: A clustered design space for an eleven-bar truss. The first row shows an optimization run history, which exhibits some patterns but is still visually complex. The second row breaks the data into clusters, which are then sorted by performance. These groups can be used to set new bounds for cluster-based exploration, as shown in the third row. From Brown & Mueller (2017a) and Brown & Mueller (2017b).

Given the presence of a numerical objective, one initial design strategy might be to run an optimization. The figure below shows the history of designs considered during an optimization run using an evolutionary solver for 20 generations. While selecting the optimal truss might be desirable for some applications, other design process might call for the consideration of “near” optimal designs that satisfy some other requirement. The history of the run itself leaves many similar designs, which might be cognitively challenging to meaningfully consider and compare side by side. In response, a data science technique called clustering is used to create groups (or categories, families with representative archetypes) based on the design vectors of each solution, such that designs which look similar should belong to the same group. These groups are then organized based on median performance, such that designers can cycle through groups and meaningfully compare them based on their merits rather than get lost in a hundred nearly similar designs. The clusters are then combined with a user-defined flexibility rating to create a more focused design space boundary for each cluster, in which a designer can more quickly explore a region that likely contains higher performing designs. While this example presents a simplified design space, such interpretation techniques have significant potential to improve parametric design processes with substantially more complexity, which can place a higher cognitive load on designers.

Both of these concepts can be layered onto interactive optimization procedures in different ways to enable more effective design space exploration for early building design.

2.11 Vision for performance-based, data-driven design

2.11.1 Designer + computer engaging with the design space together

Based on these technical developments and conceptual possibilities, the vision of this dissertation is a process in which a human designer and a computer work together while engaging with the design space during conceptual design, using a combination of optimization concepts and data science methods to enhance interactivity and promote interpretability of the design space. The more natural combination of designer and computer alleviates the tendency to consider only a few options in conceptual design, a problem which became rampant through design processes in which architects hand off a model to engineers, who then build their own simulation models (Flager et al. 2009). The tedious translation between modeling types and associated logistical complications have prevented not only more flexible exploration (see Figure 2.15), but also limited discussions of the interplay between competing performance objectives when early, consequential decisions are being made. Now that researchers across disciplines have expended considerable effort to bring design and analysis software together, and also begun to explore the applications of interactive optimization and data science to design, it is vital to develop, study, and test ways in which their effective integration into design processes can continue.

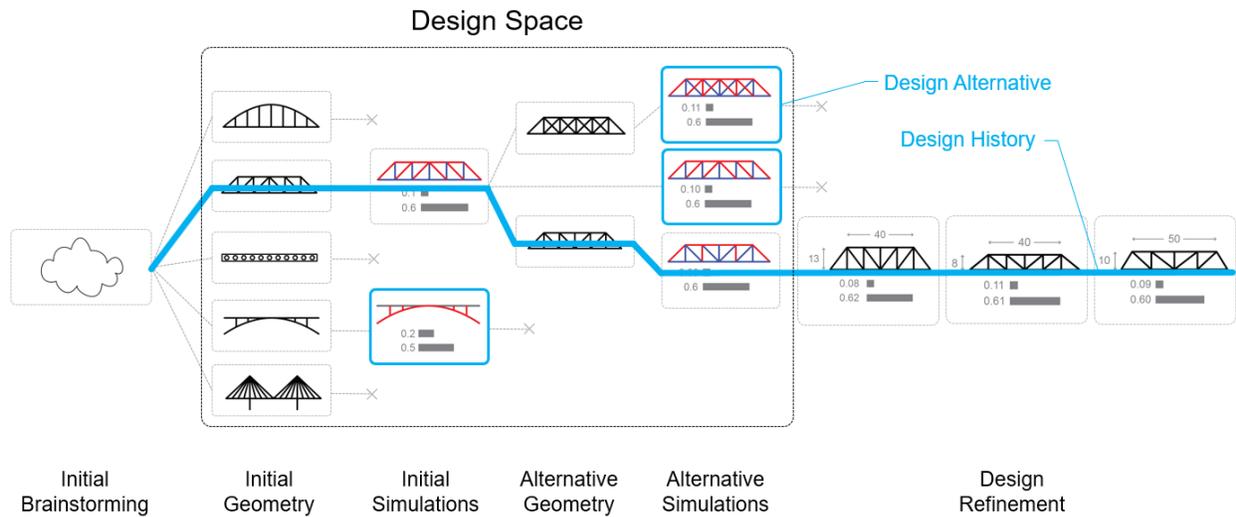


Figure 2.15: A view of the typical design process, bogged down by translation issues between software, leading to only a few designs being meaningfully considered

Although this relationship between human, computer, and the design space has previously been enabled in various forms, this dissertation seeks to extend earlier contributions concerning interactive optimization by proposing a generalized framework for multi-objective early building design that can enable the navigation the complex interactions and tradeoffs between quantitative and qualitative design goals (see Figure 2.16). In addition, it seeks to move beyond the typical interface of the design space catalogue towards more flexible means of manipulating geometry and other design characteristics, closer to the freedom allowed by a paper and pencil or other analog design methods, but with the significant added benefit of computational simulation feedback and guidance. It also expands the scope of typical research into design space exploration by proposing ways in which concepts from data science can be used earlier in the process, even when the design space is being initially formulated. Although this dissertation does not solve every limitation of parametric design or barrier towards adopting optimization in a flexible manner, its contributions have the potential to stimulate more widespread adoption of computational methods in research and practice, while overcoming fixations on standard solutions and enabling creative, performance-conscious design outcomes for the built environment. The next section describes specific research goals that address early building design using multi-objective data approaches.

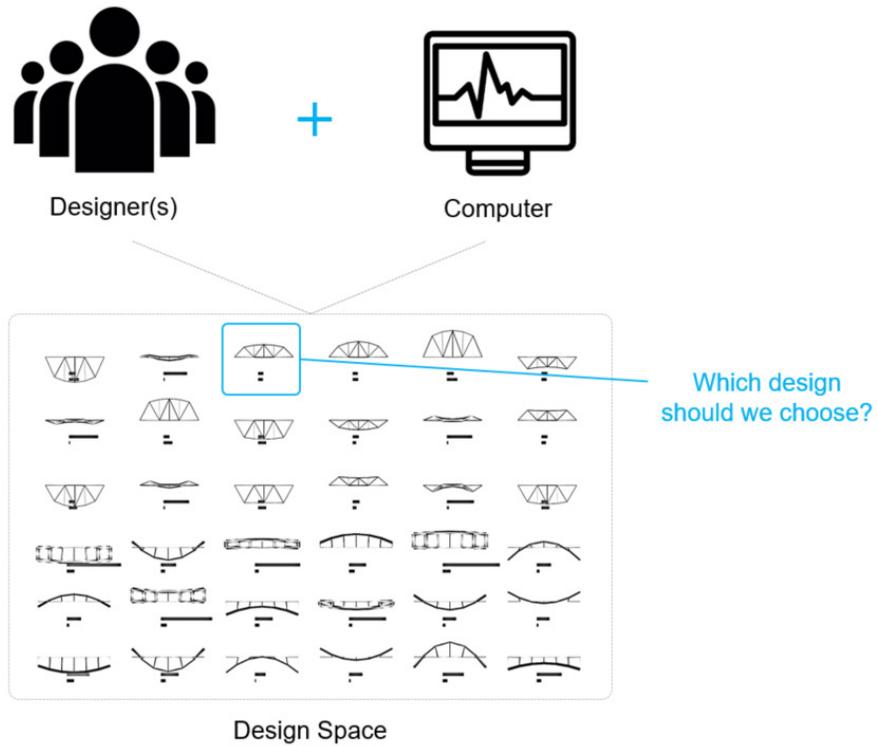


Figure 2.16: Humans and computers working together to engage with a dynamic design space in order to find a satisfying design

2.12 Research goals and responses

Recalling the introduction, all areas of inquiry in this dissertation are meant to address one broad question: how can creative architects and engineers use computation to navigate various performance objectives as part of the natural human design process? However, as this background chapter has detailed, other researchers have made considerable progress towards answering this overarching question. Upon reviewing the literature, this dissertation presents a series of additional goals for how current parametric methods can be improved or extended to make them viable as a common early design tool. These goals are to make computational design processes more:

1. **Meaningful and Attainable**
2. **Flexible and Interactive**
3. **Directed and Divergent**

The concept of a more meaningful design space aims to prevent the tendency of human designers to get overwhelmed by the large volume of options and simulations that can be generated. The whole goal of such

parametric methods is to provide feedback to make the most important design decisions, not drown in small details. To avoid cognitive overloading and to ensure that these parametric models are valuable, can we propose ways to direct designers to only important areas of the design space, without sacrificing too much creative freedom? Can we analyze design spaces to uncover patterns that suggest a successful outcome?

Attainability is the practical goal of enabling the effective use of design space exploration methods for those who are not specialists and experts in optimization. Can difficult parts of the interactive optimization process be automated, such that the results are still trustworthy but accessible to a wide variety of architects and engineers? What methods, interfaces, or tools can be created and validated that promote functional understanding of interactive optimization methods for design?

Flexibility means the ability to move beyond standardized solutions, and consider high performing designs that are not simply the optimal or obvious solution. It also refers to specific tools and workflows—are there multiple ways to employ data-driven design techniques? Does an interface dictate how the design process must go, or are designers able to move freely from concept to concept?

Interactive workflows involve providing input into a design process that still generates feedback and guidance based on valuable simulations. It is contrasted with fully automated procedures in which designers cede total control to the computer while only pursuing quantifiable goals.

Directed and divergent are opposites, but both may be beneficial during various parts of the design process. In many phases, a designer might want to achieve a specific quantitative target, while still being able to consider secondary implications but staying true to an original intent. At other times, the designer might want the computer to be a brainstorming partner, and consider quantitative goals while using the computer to generate diverse solutions. What kind of multi-objective data approaches allow for both directed and divergent functionality, and how can they be implemented in digital tools?

While there is some correspondence of these goals to the specific order of contributions of the next three sections, the goals are not meant to be isolated. Certain chapters address multiple goals in different ways.

Chapter 3 presents methods for analyzing, transforming, modifying, and eventually automatically generating design variables while formulating a parametric design space. Analysis and transformation aims to make the process more meaningful and lead to more satisfying exploration later on, while automation seeks to stimulate creativity in new directions, and empower non-experts or time-pressed designers to use parametric definitions with less effort.

Chapter 4 demonstrates gradient-based guidance for optimization, which allows for increased control over design space directions during exploration. It affords more flexibility than many optimization methods, and can be conducted interactively with the assistance of data models in place of performance simulations.

The presented guidance methods include functionality for both directed design by continuously improving objectives, and divergent design by suggesting designs that perform similarly but may look different than a current solution.

Chapter 5 synthesizes, proposes, and tests measurements for diversity in design, which can ultimately help enable more divergent and creative brainstorming, even within a particular parametric definition.

When used together, in succession or iteratively, the various methods explored in this dissertation make contributions towards achieving its stated goals.

2.13 Overall map of design process

In order to understand the relationship between each of these next sections, an overall map of the conceptual design process is presented in Figure 2.17. This figure describes a typical approach when computational modeling, and in particular parametric modeling, are used along with performance considerations to arrive at a design concept. While this design concept is a final product for this design stage, it would almost certainly be further explored and refined during later stages, as decisions continue to be made and more information is known about the design. The design process starts on the left of the diagram, when the concept is geometrically nebulous in that it can take many different directions. As decisions are made, the design is gradually refined and comes into greater clarity, while certain possibilities are discarded. This process is not always linear, and often includes considerable iteration, but on the whole must progress towards an eventual, viable design concept.

Along the way, designers must complete discrete tasks when using a particular computational approach, especially for parametric design. For example, unless a parameterization is cleverly developed to be trans-typological, this general typology will need to be decided prior to other explorations. In other words, like deciding on lunch, a pizza can be manipulated in terms of size, toppings, and other variables, and eventually be optimized within this decision space (Williams 2016). However, it is very difficult to use continuous or ordered variables to turn a good pizza into an effective taco, despite the culinary promises of contemporary fast-casual dining establishments. Thus, although a pizza and a taco could each easily be explored parametrically beyond this point, it would likely be most effective to do so in some parallel form, and probably at two different restaurants. The same is true for many structural concepts or massing schemes for early building design.

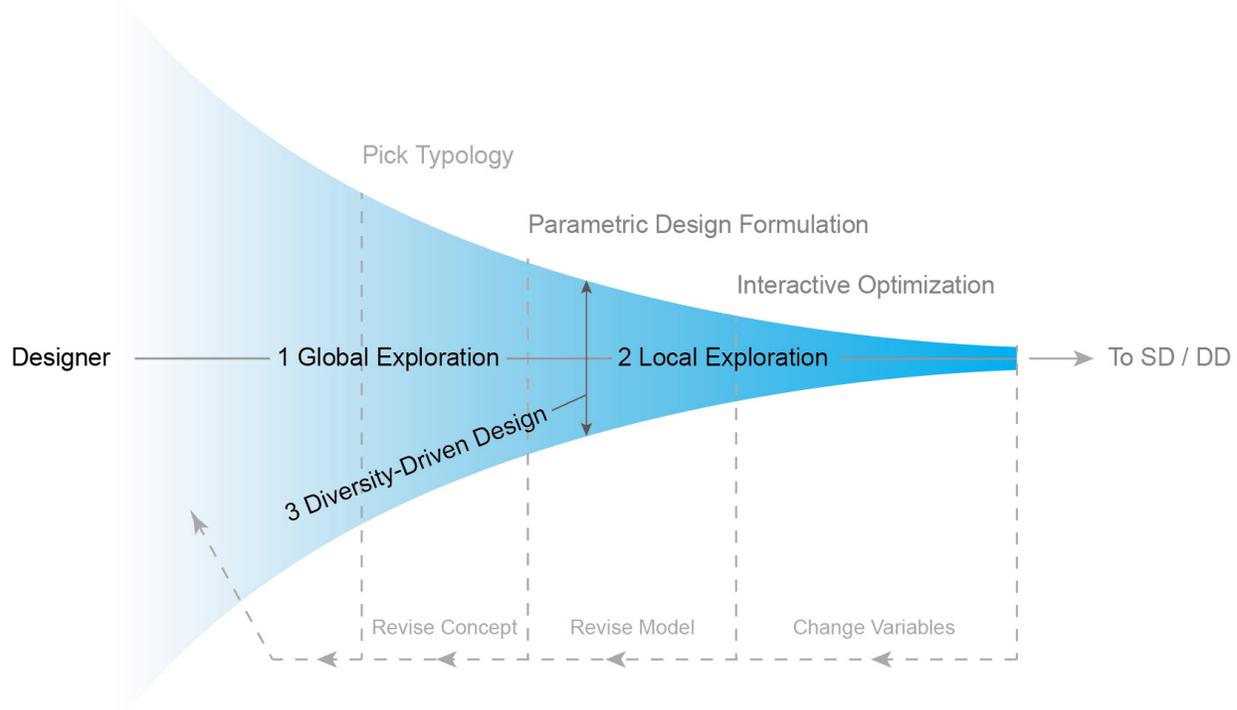


Figure 2.17: A map of the overall design process when using a computational parametric approach

After picking a typology, the designer encounters the task of formulating the design space. While doing so, the designer makes numerous decisions while picking which aspects of the design should be flexible, what should stay constant, how variables related to one another, what simulations are valuable, and what objective functions accurately describe the design goals for the project. This task often takes a mixture of experience and creativity, and it is immensely consequential to the rest of the exploration because it in some sense sets all design possibilities from the beginning. Although it is certainly possible to encounter a surprising solution within a human-formulated design space, especially with complex parameterizations, it is vital that the designer chooses variables and their relationships in a way that will lead to meaningful exploration.

More will be written later on the challenges and limitations of parametric design, particularly in relation to previous research that seeks to break the typical structure of models and make them more flexible. However, parametric frameworks and the design and objective spaces they describe have advantages for optimization, data mining, diversity measurement, and other concepts in this dissertation. Furthermore, these approaches are gradually being adopted as a valuable tool for early design exploration in practice, even if first by specialists highly skilled at this specific task. Thus, parametric design space formulation remains as a crucial step for many computational workflows, and remains a critical topic of this dissertation.

After setting up a typical design space, the next task is to explore this space to arrive at a final concept. Significant research exists in this area of design space exploration across many domains, due to the widespread advantages of parametric design approaches to engineering and other fields. However, this dissertation will focus on methods that are interactive, which allow for the ongoing expression of designer preferences throughout engagement with the design space, which is a requirement of truly natural design processes. This interactive optimization, and the final concept it produces, is the last task in the scope of this dissertation. It is again worth mentioning that although these tasks were described linearly, they can of course be iterative—for example, it is very common for lessons learned from interactive optimization that cause the designer to reformulate, or even go back to an earlier “drawing board”. Nevertheless, these tasks are in many ways discrete, and inform the organization of this dissertation.

The final area of research examined here, called diversity-driven design, does not match a particular task but instead strengthens the ability to modulate the directness of a design approach at different points in the process. While optimization and refinement to a single outcome is for the most part the inevitable goal of building design, at various points it is useful to consider more options. In many brainstorming situations, the objective is actually to generate many, meaningfully different possible solutions. A major section of this dissertation is concerned with this goal of diversity-driven design within a parametric framework, which fits into the overall vision of the design process as a way of controlling modulating the specificity of the concept during design.

The following chapters describe specific contributions in each of these areas related to the overall design process: formulating the design space (also called global design exploration), interactive optimization (local design exploration), and diversity-driven design. The dissertation then presents a set of computational tools implemented for a parametric design platform that enables such approaches, before describing a design study meant to understand how such ideas and tools affect the overall building design process.

3 Design Formulation and Global Exploration

3.1 Introduction

In early stage architectural design, it is becoming increasingly common to employ parametric models to consider different design options. Although parametric design is a general strategy with wide ranging formal and geometric applications, many designers recognize its potential as a viable approach to performance-driven design (Oxman 2008c; Shea et al. 2005; Turrin et al. 2011; Holzer et al. 2008). Given the contemporary computational tools mentioned in the previous section, it is now possible to use simulation and approximation techniques to rapidly estimate performance goals such as building energy usage, material quantity, structural deflections, daylight availability, and others. These simulations often correspond directly to design goals that are of global concern to architects—including cost, carbon footprint, resource efficiency, thermal comfort, and resiliency. This chapter addresses the enhancement of interactive approaches to parametric design formulation using data science techniques, with the goal of assisting designers in an early, and often ignored pre-exploration phase of parametric design¹.

While engineers have long used parametric simulation techniques to evaluate and ultimately optimize designs (Sobieszcanski-Sobieski 1995; Anderl & Mendgen 1995), recent advances that connect geometric design environments with analysis programs have made simulation more accessible to architects, even early

¹ A version of this chapter has been published in:

Brown, N., & Mueller, C., 2019. Design variable analysis and generation for performance-based parametric modeling in architecture. *International Journal of Architectural Computing*, 17(1), pp. 36-52.

in the design process. The strength of parametric design is that when considering even just a few variables, designers can automatically generate hundreds or thousands of possibilities. In this context of ubiquitous simulation, it becomes both possible and necessary for designers to analyze the vast quantity of potential solutions and distill this information into actionable, specific feedback and guidance that can simplify and ultimately enhance the design process. Drawing on recent advances, it has also become clear that data science techniques can help designers identify patterns, discover useful information, suggest conclusions, and generally support decision making during design space exploration.

However, when using such strategies, it can take considerable time and effort to set up a quality parametric design space in the first place. Deciding which variables to use and how their relationships dictate geometry or other model properties is an important initial step, since it sets the edges of the possibility space. In engineering, this process is difficult because it can require considerable domain knowledge, as well as an understanding of scaling, constraints, and algorithms if optimization related procedures are being used for exploration. In architecture and building design, the task of establishing the design and objective spaces can be even more significant, since designers might seek more geometric flexibility while considering a multitude of non-quantifiable design goals. The setup process is likely to be iterative itself, and an integral part of the design inquiry—in some cases, formulating variables might be so frustrating or unsatisfying that it leads designers to abandon this technique for more design freedom, which limits the benefits of simulation feedback in general.

It is thus crucial to consider how machine learning and data analysis techniques can be used to improve or assist during design space formulation, and how this might affect eventual exploration. In architecture, numerous design considerations can be translated into data, which designers must process and synthesize while making decisions. This chapter focuses on parametric design while considering building performance, which can greatly benefit from directed data analysis as early as possible. In this setting, such techniques can give designers more direct control of performance, provide ways of manipulating geometry that are more meaningful than variable-by-variable exploration, and uncover performance-conscious geometric transformations that are surprising to the designer and would not have been explored otherwise. In addition, data analysis could help automate parts of the design space formulation process, especially for key building typologies, which could drastically increase the ease of use of parametric modeling for both practicing architects and engineers. Any of these potential techniques are then applicable to live design workflows, directed optimization, or any interactive, performance-driven design approaches in between.

Given the needs and opportunities related to effective design space formulation, this chapter first provides a brief overview of methods from engineering that have been adapted for architectural design, while also describing relevant and applicable concepts from data science. It then identifies three situations in which data analysis can be beneficial while designers are establishing parametric variables and setting up the design and objective spaces. Next, the analysis techniques are demonstrated on two design case studies

involving trussed spanning structures. Finally, a discussion is provided concerning when each technique may be applied, how it affects the design process, and how it might be further generalized to other types of buildings and structures.

3.2 Approaches to problem formulation and exploration from engineering

Most engineering design processes involve running computer analyses in order to find a relationship between design inputs and responses (Simpson, Peplinski, et al. 2001). In engineering fields related to buildings, including structural and mechanical design, the process often includes considerable guess and check by an engineer, who manipulates each iteration directly. However, in manufacturing situations where a single design is mass-produced, engineers have long been concerned with finding an optimal design within a pre-established set of possibilities. In this case, the inputs for a simulation are specific design variables, which act to bound the problem and allow for a more robust design approach. Such engineering problems can be further divided into categories based on whether they are low or high dimensional, computationally cheap or expensive, and contain explicit analytical or black-box simulation functions (Shan & Wang 2010). A variety of approaches exist for solving each type of problem (Zandieh et al. 2009).

Especially for design problems that are high dimensional or require expensive simulations, the first explorative step is often choosing an experimental process for systematically sampling the design space at an appropriate resolution. This involves creating a sequence of *factors* (design variables) and *levels* (values for these variables) that can be simulated first, after which designers can analyze the data and decide what to do next (Box et al. 2005). Typical methods for design of experiments include Factorial Design, Central Composite or Box-Behnken Design, or Taguchi Orthogonal Arrays (Tsui 1992). Many of these were developed for quality control purposes in manufacturing (Kackar 1985), but can be applied to early stage design as well. In some cases, it is possible with only a few simulations to gain a basic understanding of which variables affect performance most, and which settings tend to lead to good performance. These methods can be effective as a first pass, and have been proposed directly on architectural applications before (Brown & Mueller 2017a). However, they have significant weaknesses when used as a standalone method in conceptual design, since they cannot capture complex, nonlinear relationships or discontinuities while mapping between inputs and outputs.

As computational techniques continued to advance, designers began creating surrogate models, sometimes called metamodels, that attempt to approximate the general performance of an entire design space from these limited data points (Queipo et al. 2005). If additional computation is allowable, a denser sampling technique is used to produce a more accurate surrogate model. Methods for surrogate modeling include response surface methodology (Myers et al. 2009), neural networks (Hagan et al. 1996), Kriging (Simpson, Mauery, et al. 2001), and others. If this technique is applied directly to architectural conceptual design, architects may choose to explore the design space immediately after the surrogate model is built, either

through optimization or using sliders to select different variable combinations and gain real-time feedback about the performance response of each possibility. However, the initial variable selection and problem setup may be inadequate for a variety of reasons. While considering the setup process itself to be iterative, there are additional analyses that can provide further information and value to the designer at this stage.

3.3 Data science for tuning the design problem

During problem parameterization, a basic analysis can first inform designers which variables are most important to the performance outcome. This task is not necessarily straightforward, as it can be difficult to define variable importance, although many surrogate modeling techniques include internal procedures for its calculation. These calculations, which more generally identify the controlling features of a dataset, are immensely valuable for exploratory data mining and analysis applications. However, if the goal is to keep controlling variables and throw out (or leave flexible) unimportant variables, it may also be useful to reduce their dimension. In data exploration, visualization, and optimization, dimensionality reduction can make computational processes more efficient, and sometimes more meaningful and accurate.

A common approach for dimensionality reduction is principal component analysis (PCA), which performs a linear mapping of the data to a lower dimensional space in such a way that the variance of the data is maximized (Pearson 1901). This technique has been used in conjunction with engineering processes (Su & Tong 1997; Tong et al. 2005), and even directly for variable reduction or creation during optimization (Yonekura & Watanabe 2014; Khaled & Smaili 2005). Specific to the case of determining relationships between design variables and performance outcomes, a similar technique called Canonical Correlation Analysis (CCA) can find a linear set of coefficients for input variables that maximize correlation with another dataset (Hotelling 1936). It is also possible to conduct PCA in a way that captures non-linear relationships (Karhunen & Joutsensalo 1995; Schölkopf et al. 1998; Scholz et al. 2008), or conduct CCA when mapping input variables to an objective function which does not have linear behavior (Hardoon et al. 2004). Other dimensionality reduction techniques, which can be especially useful when associations between data are nonlinear, include self-organizing maps (Kohonen 1990) and Sammon mapping (Sammon 1969). While different techniques range in levels of complexity or specialization, this chapter focuses on simple, fundamental approaches that are broadly applicable and could be improved or modified for specific design applications.

3.3.1 Background on variable analysis in architecture

Increasingly, researchers are exploring ways to implement data analysis and modeling techniques during the architectural design process (Tamke et al. 2017; Derix & Jagannath 2014; Chen et al. 2015). One prominent application is data visualization for high-dimensional design spaces (Harding 2016; Wortmann 2017b). Others have proposed using statistical methods to consider relationships between variables and

objectives, including Bayesian inference (Conti & Kaijima 2017). Chaszar & Joyce (2016) note the limitations of a design process in which the original variables are frozen, and Harding & Shepherd (2016) propose evolving the directed acyclic graphs for parametric models themselves. In computer graphics, there is considerable existing research into breaking the structure of traditional parametric modeling and manipulating or retrieving geometry more directly (Sederberg & Parry 1986; Schulz et al. 2017), but these changes do not necessarily maintain the required model inputs for performance simulations.

With a few exceptions, most of this research in the field of architecture emphasizes exploration or selecting between design options after variables have been established, rather than investigating the variables themselves or considering how they might be modified. During parametric design processes, designers often add or remove variables, conduct sensitivity studies, and adjust parameter ranges (Bradner et al. 2014). Furthermore, there is a rich history in both architectural theory and practice that considers computational design (Oxman 2006; Kolarevic 2004; Terzidis 2006), design space exploration (Woodbury & Burrow 2006), parametric design (Woodbury 2010), and their relationships to the general drivers, goals, constraints, and other forces acting on the architectural design process (Lee et al. 2014; Rowe 1986; Oxman 2008b). Even the terms “design parameter” and “parametric design” have been understood in different ways (Schumacher 2009; Kotnik 2010), and only sometimes correspond directly to the formal design variables used in engineering (Haymaker 2012) and optimization.

While some assumptions about the role of designers can be shared across disciplines (Cross et al. 1981; Cross et al. 1992; Dorst 2011), architectural design often presents challenges particular to its discipline. Especially when adapting methods primarily from engineering, it is important to acknowledge this context, and consider parts of the architectural design process that may be outside the scope of a traditional parametric engineering problem. Within architecture, difficulties arise due to the reality that architectural design problems are often considerably fuzzier and ill-formed than their engineering counterparts. Nevertheless, by considering portions of the process that could be considered pre-parametric design, or focusing on interactive rather than automated workflows, it is possible to extend the relevance of data analysis techniques to a much wider portion of the design process. Using data analysis to question, modify, relate, transform, or generate design variables is one such endeavor.

3.4 Overview and goals

3.4.1 Strategies for variable analysis, transformation, and generation

This chapter proposes and tests three strategies for applying data analysis techniques to performance-driven parametric design problems. The strategies are:

1. Analyzing existing design variables for performance
2. Analyzing existing design variables and reducing their dimension to more meaningful variables

3. Starting with a fixed geometry and generating initial variables automatically

While the case studies shown later emphasize structural and daylighting simulations, these methods can be implemented on any quantitative objective involved in building design.

The first strategy addresses the situation where an architectural designer has already built an initial parametric model, but cares to identify variables that have a large influence on performance, as well as variables that have very little effect on performance. Unlike other applications of data analysis or modeling, both variable types are significant in conceptual design, since performance-important variables must be carefully considered, while non-important variables may offer designers freedom in terms of visual impact or secondary design goals. Information generated during a variable analysis could focus design exploration on particular areas of the design space or help to adjust the active variables, bounds, and connectivity of the parametric model itself.

The second strategy is applicable when a designer builds a parametric model for exploration, but manipulating variables directly is essentially meaningless with respect to performance. In these cases, transforming initial variables into meaningful, synthetic ones enables more effective interaction with the design model. The initial condition of meaningless variables could occur because there are too many to consider individually, or it is difficult to intuitively improve performance through manual manipulation. For example, a parametric design for a truss might have numerous design variables related to the coordinates of each node on a truss, but if the truss contains many nodes, moving them one at a time is not useful for exploration.

The third strategy is meant to save time and manual effort in setting up the design space, and possibly lead to the creation of design alternatives that were not initially imaged as possibilities. Within such a workflow, designers could essentially sketch a possible geometry using a computer, provide a desired number of design variables (or threshold for importance), and after a heavy initial simulation, have the ability to control meaningful geometric variables that morph the design in ways that are compelling both visually and in terms of performance. While the implementation of automatic variable generation in this chapter is limited to a single, intuitive case, this methodology could be widely extended in the future.

3.4.2 Case studies and datasets

This chapter primarily makes use of two design case studies and corresponding datasets to demonstrate the potential utility of variable analysis in early stage design (see Figure 3.1). The first case study is a simply supported truss, which is loaded at each of its bottom nodes. The truss is made of triangulated linear steel elements that can only transfer axial forces, in either tension or compression, which are here visualized in blue and pink respectively. For the sections involving an existing dataset, the variables are the vertical locations of the top and bottom nodes, with symmetry enforced. The resulting six variables are numbered 1-3 on the top and bottom, with 3 referring to the central node. Two performance objectives are considered:

the weight (or mass) of structural material required to resist the loads, and the maximum deflection experienced by the truss. This case study was selected because its results are intuitive to structural designers—general behavior and expected optimal geometries are understood visually.

The second case study was inspired by a built project and thus connects this chapter to practical architectural applications. The starting geometry for this case study is modeled after the SFO International Terminal, designed by SOM and completed in 2000 (SOM 2000). The model contains two external trusses resting on two columns each, with the outside trusses supported a third in the middle. This model explores the interplay between two design objectives for the case study, which are the weight of the structure versus the spatial daylight autonomy (shortened to daylight autonomy in the graphic keys) of the interior space. Six different structural load cases are considered, in addition to the weight of the structure itself: a dead load, vertical live load distributed along the top nodes of the trusses, asymmetric live loading on either half of the roof, and lateral loads in each direction with the forces acting on the outermost point of the truss. To assess the amount of daylight in the space, all surfaces of the geometry were modeled as either opaque surfaces or translucent panels, which occupy skylights located above the central truss and clerestory windows surrounding the perimeter of the building. The design variables allow for the basic geometry to be modified by adjusting truss depths or heights relative to one another, which creates considerable possibilities involving increased truss depth, as well as tilts and general massing adjustments.

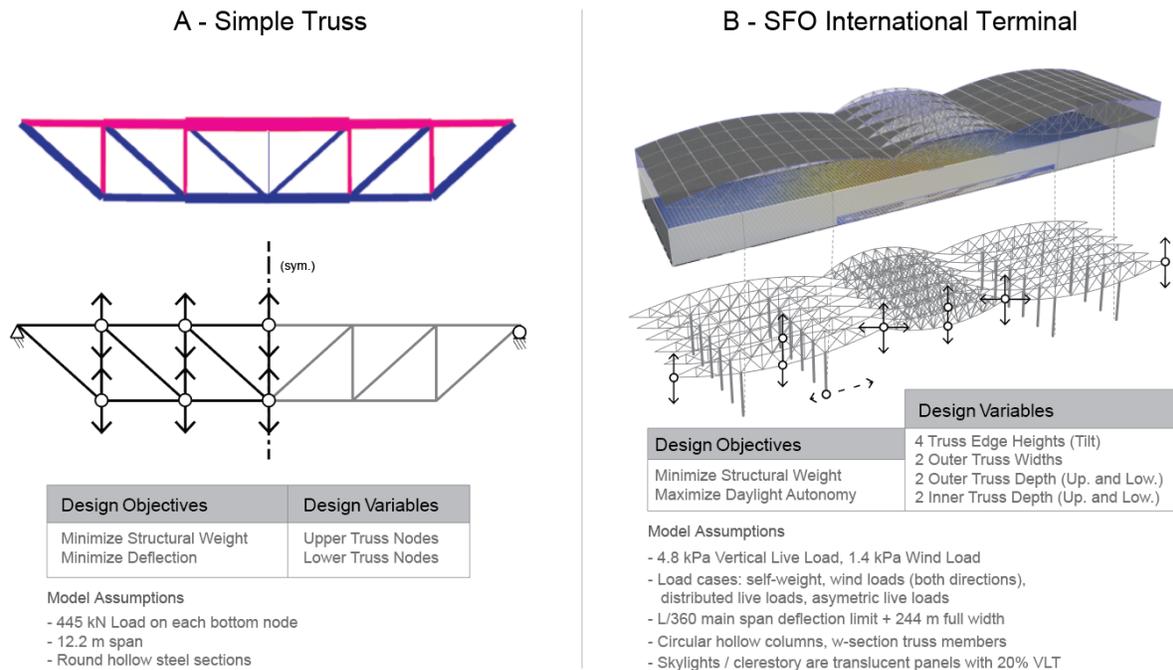


Figure 3.1: Visual descriptions of the parametric design case studies used for analysis

The design case studies were modeled using Rhino and Grasshopper, while some of the data analysis was conducted in MATLAB and R (The R Foundation 1997). In Grasshopper, each structural design is evaluated by simulating the loads, calculating internal forces, and generating required sizes for each member. This process is automated using Karamba, which relies on a cross section sizing algorithm based on European structural design codes. The final outcome of the structural simulation uses these sizes to calculate the required steel weight and maximum deflection of each design iteration. The plug-in DIVA, which is based on Radiance (Ward 1994) and Daysim (Reinhart & Herkel 2000; Reinhart 2011) and enables raytracing simulations directly in Grasshopper, was used to calculate spatial daylight autonomy. Since this case study demonstrates early design decisions and experiments with overall performance metrics rather than high-precision rendering, the “low” simulation quality was used, corresponding with the Radiance settings in Table 3.1. With these settings, each sDA simulation could be completed in around three minutes on a desktop computer.

The initial datasets for variable analysis and dimensionality reduction consisted of design vectors and corresponding performance results, which were created using Latin Hypercube sampling on each parametric model. Although datasets of various sizes were utilized, it was generally possible to find patterns and achieve reasonably accurate approximations for daylighting with fewer than 500 results, while 2-10x this amount was used for structure. In practice, this initial creation of a dataset could often be completed at the timescale of a break during the workday or overnight, after which data analysis or exploration can be conducted. Applications requiring longer, better quality simulations or higher resolution sampling of the design space might require initial simulations to run for several days or longer. These decisions are at the discretion of the designer based on time costs, risks, and potential benefits. For this chapter, surrogate models were generated as part of or alongside the analyses in the following sections, which enabled essentially real-time manipulation of geometry and feedback during exploration for even the more extensive SFO case study.

Table 3.1: Additional simulation settings for the SFO case study

SFO Structural Model Assumptions	SFO Daylight Model Assumptions
4.80 kPa vertical live load, distributed to nodes of truss	DIVA sDA “Low” Quality Radiance Parameters
1.44 kPa horizontal wind load, applied in both x-directions	Ambient accuracy (-aa) : 0.15
L/360 main span deflection limit	Ray reflection limit (-ab) : 2
244 m wide x 61 m deep footprint	Ambient divisions (-ad) : 512
Circular hollow columns, w-section truss members	Ambient resolution (-ar) : 256
Karamba optimize cross section feature	Ambient super-samples (-as) : 128
	Direct relays (-dr) : 2

3.5 Methodology and results

3.5.1 Analyzing existing design variables for importance

First, the existing parametric models were analyzed to determine which variables had a large effect on performance, and which variables have less of an effect, thus offering more design freedom. While initial calculations of levels and effects provided some insight into the behavior of the design problem, more rigorous analysis involved creating full surrogate models of the design space samples. Two of the most successful models were regression trees and linear regression, which both offer internal metrics for judging variable importance. Results of these analyses are provided in Figure 3.2.

In general, the methods agree on ranks between variable importance measurements for both case studies, at least for the most important variables. The simple truss shows that nodes two spots in from the support matter most, while the locations of the bottom interior nodes tend to control deflection. For the SFO case study, the top truss heights are most important for structure, whereas the width of the central truss matters most for daylight. These results are intuitive, since truss depths generally control structural material requirements, and the size of the skylights follow the interior SFO truss. A designer might use this information to concentrate on these variables when considering performance, but understand there is geometric flexibility in tilting the roof of each bay, which may be advantageous for a secondary reason such as allowing additional daylight into the space. Between these model types and methods, the deciding factor for whether one should be used depends on how well they match actual simulations and explain overall relationships.

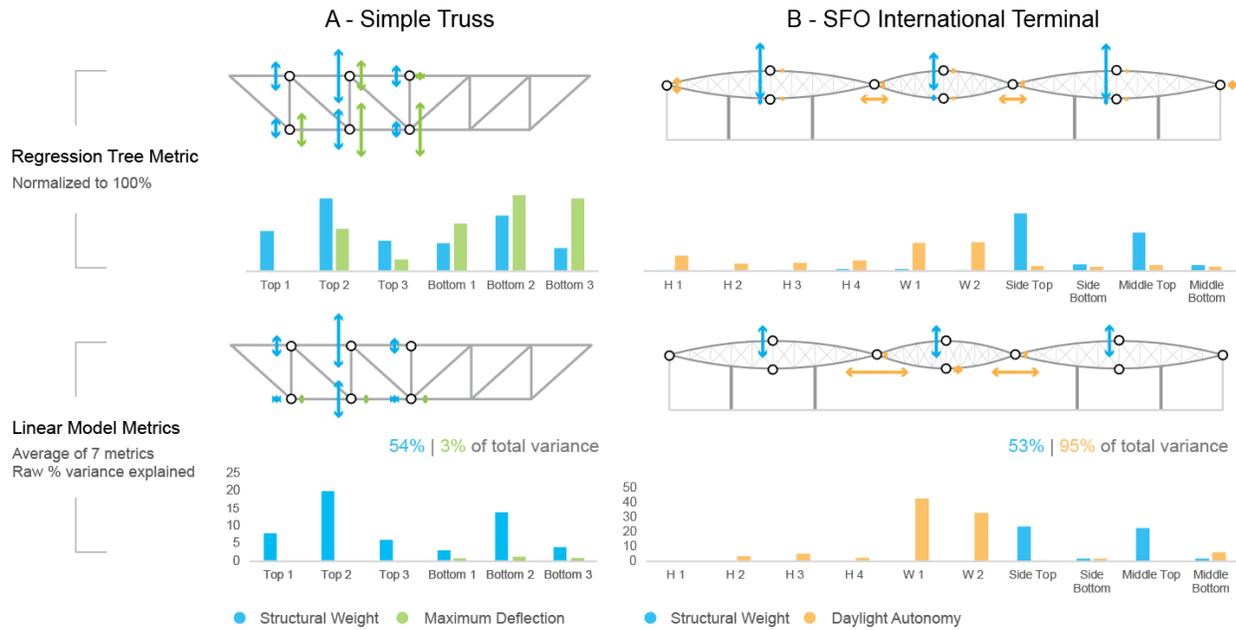


Figure 3.2: Results of the variable analysis procedure for two separate model types

For example, the linear models accounted for only around half of the variance in the structural weight models, and very little of the variance for deflection. Architects may encounter a wide variety of simulation types, some of which may be difficult to model accurately. In these cases, designers should be careful about the use of data models in general, perhaps restricting them to a first pass for obtaining an initial sense of variable importance and design behavior. Nevertheless, these variable strategies could be useful for designers looking to learn more about relationships between variables and performance objectives during model setup. After running a variable analysis, designers might respond by eliminating certain variables by locking their values, focusing on performance-important variables first before moving on to less important variables, adjusting the bounds and parametric logic before repeating the exercise, or using the data model for predictive purposes during live exploration.

3.5.2 Dimensionality reduction for more meaningful variables

In this chapter, two main strategies are tested for dimensionality reduction in the design variable space. In each case, the variables are transformed by calculating coefficients for existing variables that correspond to the performance-controlling variable. The intention is to provide designers with “sliders” or “knobs” that allow for more direct control of performance during exploration. The employed strategies are: (1) canonical correlation analysis and (2) principal component analysis. Canonical correlation analysis finds weights for each initial variable that maximize the correlation with another dataset, in this case each objective (considered separately). PCA similarly finds weights, but only operates on the independent variables,

attempting to maintain the largest percentage of total variability of the system while finding orthogonal components in descending order of variance explained. As such, it must operate on a dataset that is biased in some way by performance, rather than uniformly sampled. Although a variety of strategies were tested for generating this biased dataset, including slicing regions of the sampled objective space, this chapter shows results for the history of a formal optimization runs. These datasets in theory include designs that have been pushed towards specific, high-performance regions of the design space. While variable analysis could be combined with design optimization procedures, the goal here was to generate a dataset rather than to most efficiently or effectively find the optimal solution. As such, evolutionary solvers (genetic algorithms, NSGA-II) and a gradient-free algorithm (COBYLA) were all tested, with the former used on most examples and the later utilized on the SFO structure for this chapter.

In the case of CCA, only one master variable is created, which would essentially act as a design slider or knob specifically related to an objective. There could be multiple sliders in a multi-objective problem, one for each objective, and they could be combined with initial variables in a hierarchical design approach. Such an approach would allow designers to primarily consider the performance knob to move towards a desirable region of the design space, and then conduct local exploration within this area. Importantly, the CCA slider would not be the only direction to improve performance, but should be a reliable one. For the PCA analysis, ideally designers would be presented with a few sliders ranked in terms of their ability to influence overall system performance, which would allow for a similar hierarchical, interactive workflow.

Figure 3.3-Figure 3.6 illustrate CCA and PCA variable transformations on both case studies. Each row or column of designs represent moving through the design space in the direction of the transformed variable, starting from a central base design. The scale in each image refers to how far has been traveled in the design space in that direction. Each variable coefficient returned in the analysis is between -1 and 1, which is then multiplied by the variable range and added to the initial value, such that the maximum possible movement is equal to the edge of the initial design space. However, especially when more variables are added, the coefficients are in practice closer to zero—as such, the transformations have been scaled in some cases that would be reasonable for further exploration.

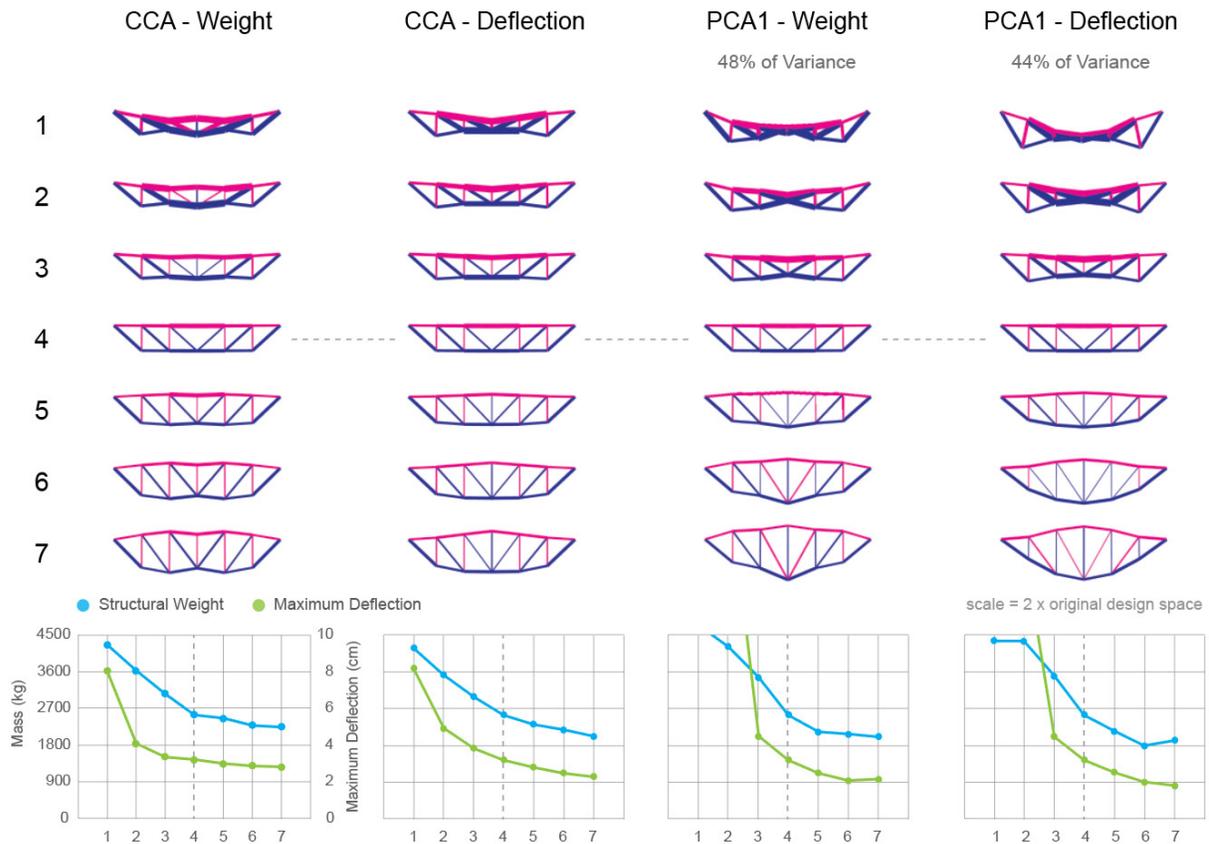


Figure 3.3: Variable directions derived from canonical correlation analysis and principal component analysis on the truss case study. In each example, the analysis finds a way of manipulating geometry that allows the designer to improve performance. Directions of improvement generally correspond to deeper trusses, although the exact shape changes depending on the considered objective.

In the truss models, it appears that CCA generates synthetic design variables that roughly correspond to truss depth, which controls both required weight and deflection for a truss. Although this relationship is not linear, a slider based on this modified variable would at least allow designers to control whether performance is improving or getting worse. The first PCA directions, which use the history of an optimization run as the dataset, show similar behavior. In some cases, moving in the direction of “improving performance” trends directly towards an optimal shape, which provides immediately useful information. When this is true, the PCA slider could be used to navigate towards the optimal design, after which designers could explore how geometry is changing away from that optima. The less important PCA directions do not seem as tangibly useful in a parametric design setting, since many do not correspond to meaningful geometric transformations. However, they could be useful in a hierarchical exploration, in which designers consider the directions in order, first navigating to a general desired performance level with the primary PCA slider.

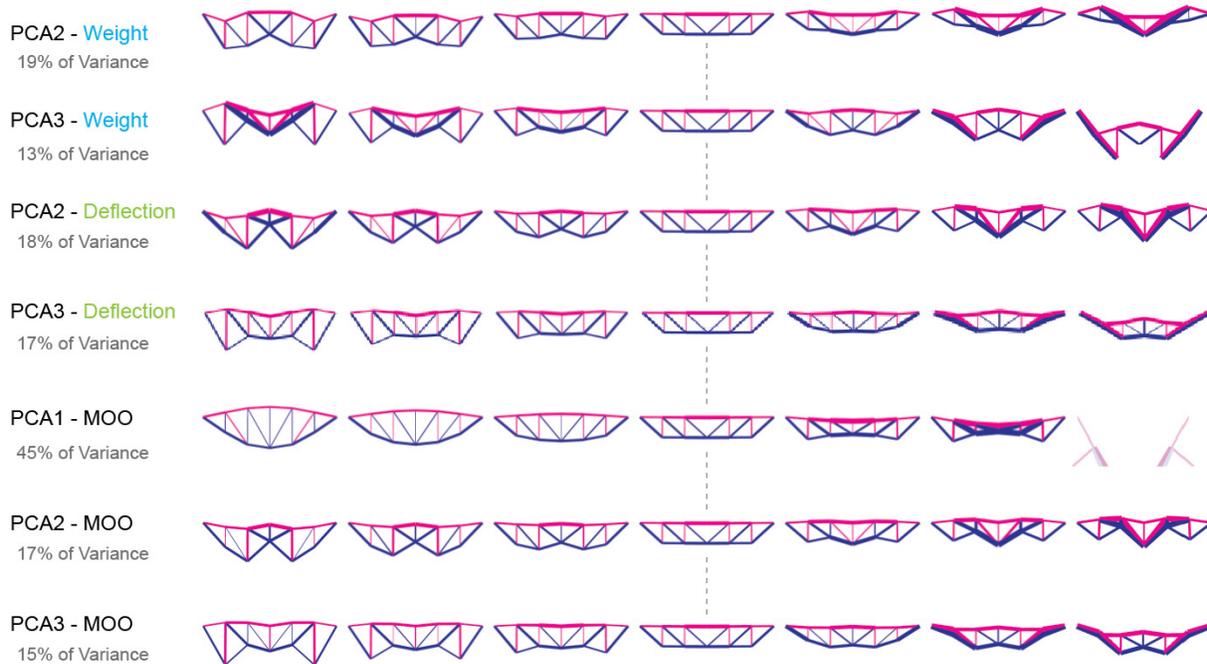


Figure 3.4: Lesser principal component directions for the truss case study. While the first component leads directly towards a visually rational, high-performance design, the less important ones are not as meaningful on their own. However, these secondary components could still be used for pure exploration, or hierarchically in combination with the primary components.

For the SFO case study, CCA chose directions that primarily corresponded to truss depth for structure and skylight width for daylight, while not meaningfully affecting roof tilt. This result could indicate tilt to be an area of creative freedom, while truss depths are more dictated by performance concerns. Curiously, the depths for the interior and exterior trusses moved opposite one another along the primary variable direction—it is possible given the initial dimensions that one truss starts out too deep, and the other too shallow for optimal performance. It is also worth noting some complications that can arise with scaling and choosing how far to move in the design space during exploration. While the direction of improving daylight remains true across the scale shown in Figure 3.5, the structural weight is a more complicated objective function, and the starting design is near a local minimum. Thus, the direction of improving structural performance calculated from the correlation analysis only holds in a limited region of the design space (see Figure 3.7), and slightly better structural solutions can be obtained by first moving in the direction of better daylight. The sensitivity of a particular objective space to scaling, sampling resolution, and step size during exploration is specific to that design—while these features may require careful consideration for problems like the SFO case study, such data exploration techniques can help designers better understand their models as they decide which design directions to pursue.

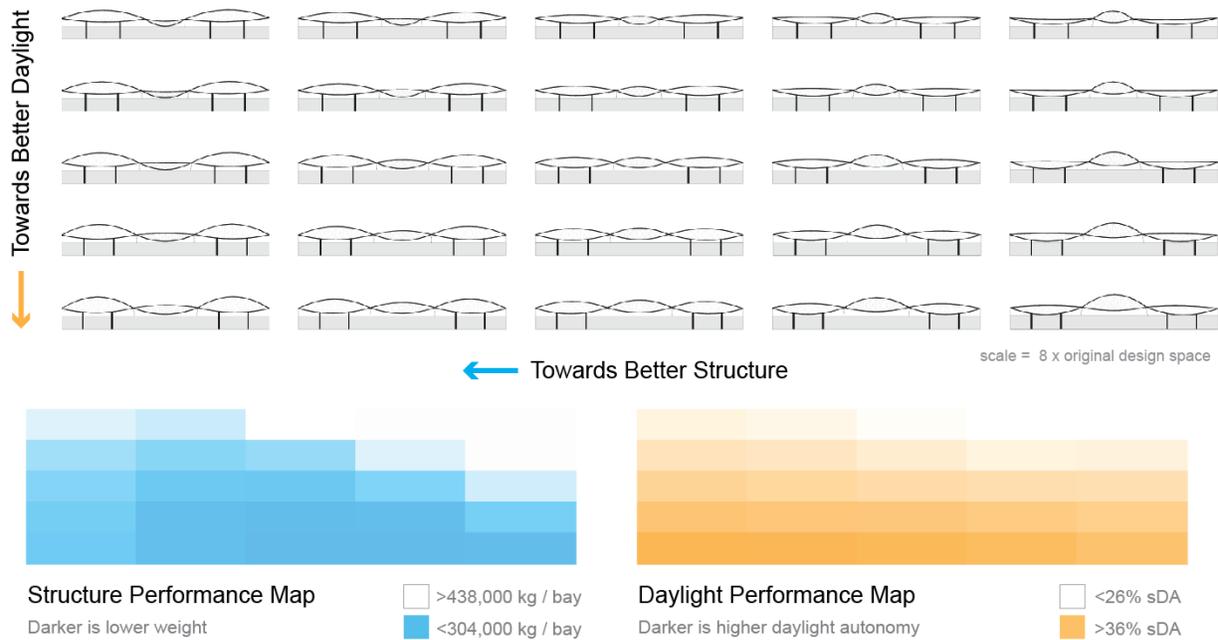


Figure 3.5: Geometric modifications in the directions indicated by canonical correlation analysis. Moving horizontally should control structural weight, while spatial daylight autonomy is controlled by moving vertically. The structural direction primarily changes truss depth, while the daylighting objective tends to make the central skylights larger. Such explorations often require scaling the design space appropriately—in this visualization, the correlation directions maintain expected relationships for daylight, but move beyond the best performing structures.

Compared to the simpler truss, this complex case study indicates that PCA-generated variable directions might be more interesting for actual design, as they can correspond to surprising geometric modifications that are worth considering, but not part of the initial parameterization. Staggered tilting, arching, or opening up the walls may all be design directions that are more tedious to arrive at while adjusting the initial variables one by one.

It is worth noting that each of these visualizations are centered on an initial mean value for consistency and clarity. This is directly appropriate as a strategy for design exploration with CCA, since the entire design space is considered and sampled. However, actual PCA-based explorations should likely be mean-centered depending on their own dataset, which is often skewed or biased. One could also imagine a design strategy in which the design space is clustered in some way, and then the designer cycles through clusters and explores each one by moving in PCA directions away from the center of that cluster. While the current visualizations show the variable relationships for each PCA direction, it is likely that their implementation would be most useful when using dynamic center points for geometric manipulation.

In general, the case studies nevertheless suggest broad applicability of these data analysis techniques to parametric problems. While not every attempted transformation is immediately meaningful, the idea of

having sliders in the performance space in addition to the variable space is potentially powerful. CCA and PCA appear to have some success in generating these performance sliders, although further research should investigate whether more specialized techniques can vastly improve their accuracy or functionality. Any such information could be mapped back into a live, interactive design workflow, such as the one shown in Figure 3.8.

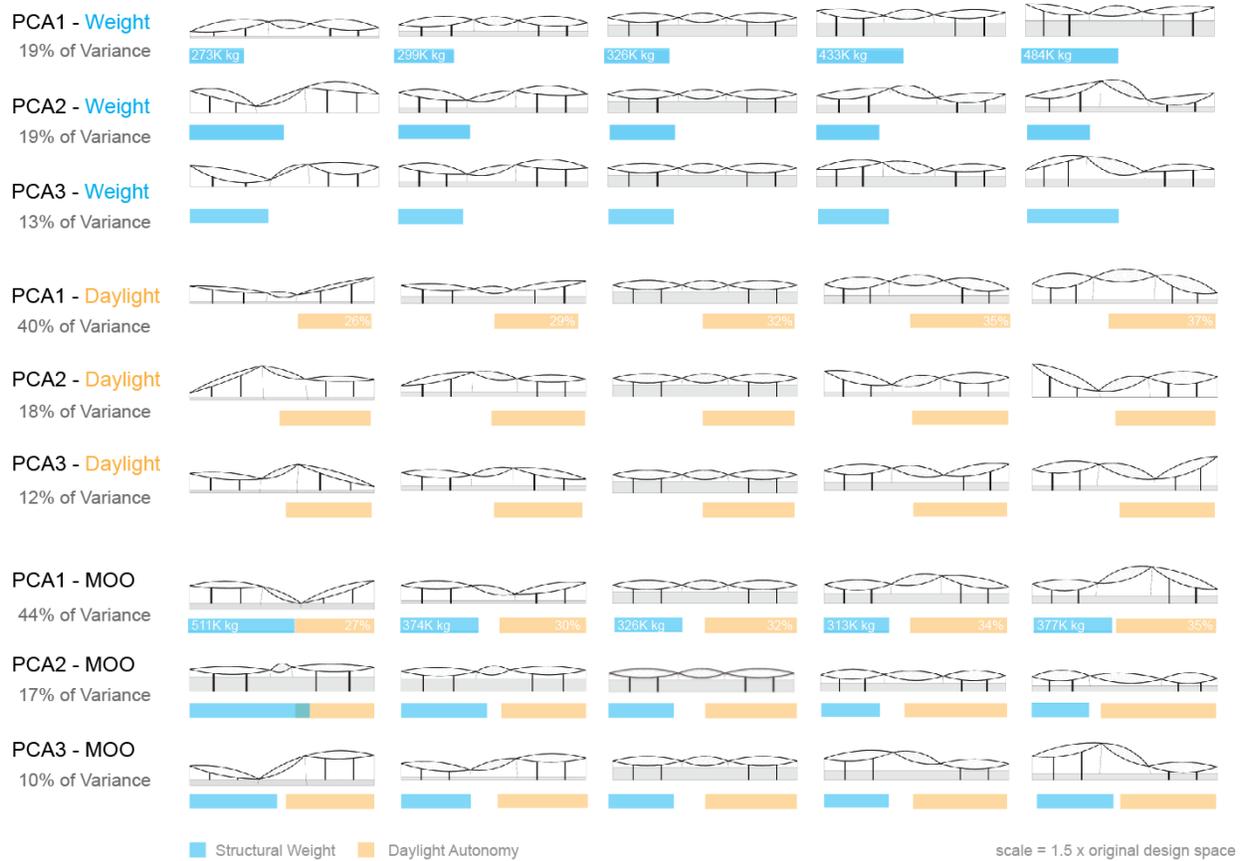


Figure 3.6: Principal component directions for the SFO case study, involving both structural weight and spatial daylight autonomy. As with the truss, the first components lead towards high performing designs, while the later components are difficult to interpret but could be useful for exploration.

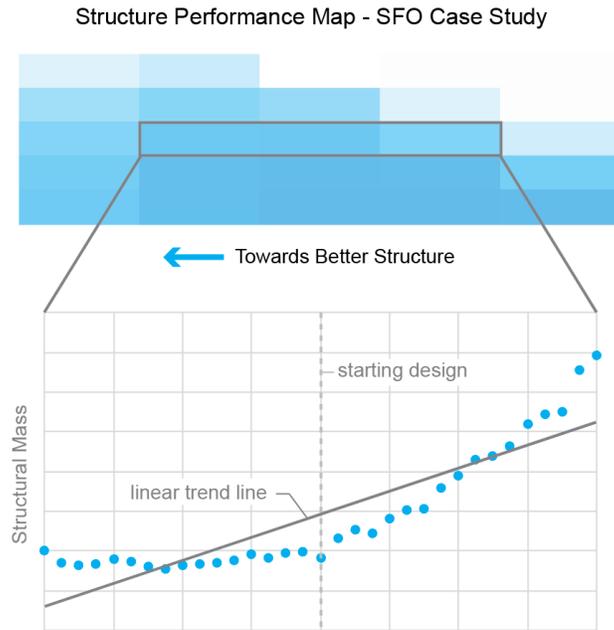


Figure 3.7: A higher resolution visualization of design performance along the CCA structure direction near the starting design. Since the starting design is near a local minimum in this direction, it is difficult to directly find better solutions, even as moving opposite projected improvement definitively decreases performance.

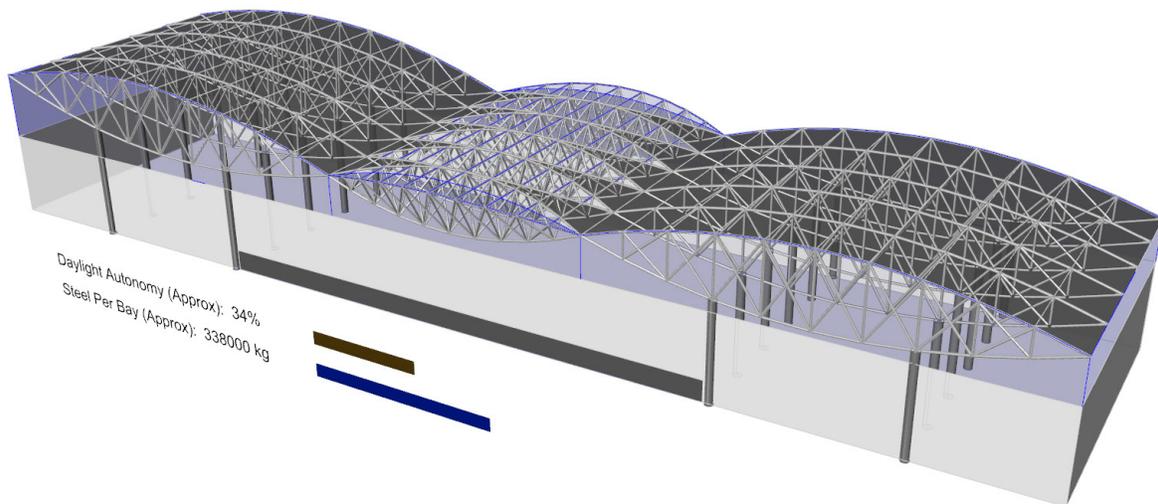


Figure 3.8: An example of a live conceptual design environment, which can include performance feedback and new synthetic variables using the transformations described throughout this chapter

3.5.3 Automatic variable generation through reduction

Finally, the possibility of using data analysis techniques from the previous section to automatically generate parametric relationships is considered. The strategy presented in this chapter is not true auto-generation since it is limited to restricted typologies, in which initial “trial” variables are first attempted and sampled before being transformed into more meaningful, controlling parameters. However, it is increasingly possible to use raw geometric properties common to many typologies—node locations for trusses or other connected graphs, control points for curves and surfaces—for this initial sampling, which makes such a procedure of potential use for numerous architectural applications. To test its feasibility, a workflow was first created for automatically parameterizing the simple truss case study. This workflow incorporates reading in a static truss geometry, identifying nodes, creating dummy design variables corresponding to the location of each node, asking for basic design information (supports, loads, boundary conditions), sampling the initial design space, and then analyzing the resulting dataset for patterns that could then become variables. For the truss case study, example designs in the automatically generated dataset are given in Figure 3.9. Very few of the sample designs are viable solutions in themselves, due to the complete freedom of each node to move independently of one another.

Example Designs in Data

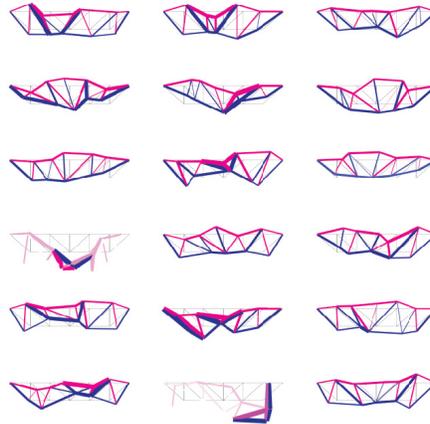


Figure 3.9: 18 representative solutions selected out of the 2000 design data set. While these designs are mostly meaningless on their own, considering them together uncovers patterns that can be useful for performance-based design.

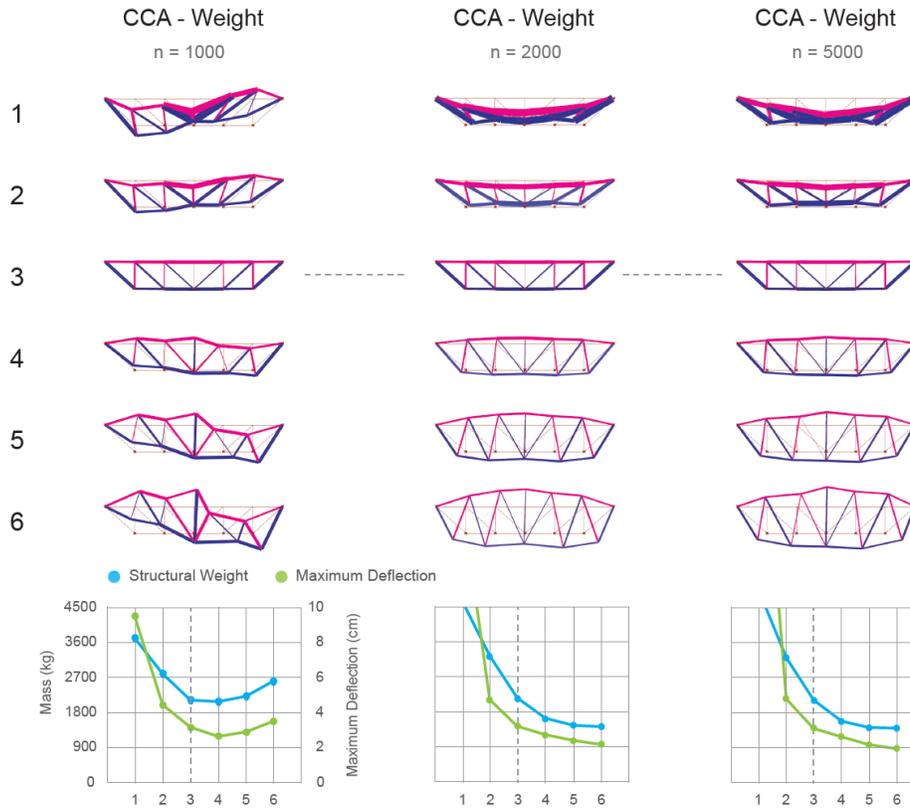


Figure 3.10: Automatically generated performance-controlling variables for the initial provided geometry. Despite the freedom on the initial sampling, the two datasets containing at least 2000 points found a variable direction corresponding to overall truss depth.

However, it is possible to uncover meaningful relationships even from this data source, as demonstrated in Figure 3.10, which shows the CCA directions for datasets of varying sizes. Although other transformations were attempted, the CAA results were the most promising in terms of their ability to generate meaningful results. For the weight objective, CCA was able to find variable directions that generally corresponded to truss depth, which tends to control required weight. A reasonably smooth transformation occurred for both the 2000 and 5000 sample datasets. This suggests that a fairly large amount of data is needed to extract a discernable pattern, but there are diminishing returns in the smoothness that can be eventually created. In order to arrive at even better results, one possible strategy is to use the initial data points to first generate a prediction model, and then resample the prediction model (which may be more continuous overall) to calculate variable directions. However, the truss directions in Figure 3.10 are encouraging in themselves—no symmetry, constraints, or reasonable bounds were imposed on the design initially, since the raw x/y node locations were used directly as trial variables.

A separate case study was also developed to demonstrate the potential of this automated variable generation procedure at the urban scale. In this case, the designer provides the computer with an initial plan for an urban building complex with varying floorplates and heights. The goal is to use computational methods to understand how the geometry might be compelled to change if additional rooftop PV potential is desired. A similar procedure involving dummy parameterization, sampling, and analysis is used to determine geometric transformations that correlate with increased PV potential. The simulations for PV potential were conducted for the Boston climate using Archsim. The resulting directions are shown in Figure 3.11, which clearly include changes to the design that affect the overall appearance and spatial sequence of the buildings.

Some aspects of this geometric transformation make it clear that the workflow is still in its infancy—a more refined artificial intelligence process will eventually include additional assumptions that set constraints on what is feasible in a real urban setting, and build smart rules and capabilities into the design process, such as the ability to pick which walls should stay frozen, and recognizing that buildings should be combined when they overlap. Nevertheless, the results stimulate potentially compelling directions for design exploration. For example, the direction for increasing PV potential seems to indicate that the tall, southernmost building, which blocks many of the others, should be peeled away, while all other buildings for the most part should spread out to maximize surface area. At the urban scale, such suggestions become immediately more useful due to the geometric complexity of separate buildings, and can provide a starting point for urban designers for where to move next, beyond simulation feedback that does not itself provide guidance. It is assumed that such guidance would likely be generated for multiple competing interests in the design, providing a rich exploration of both the design and objective spaces.

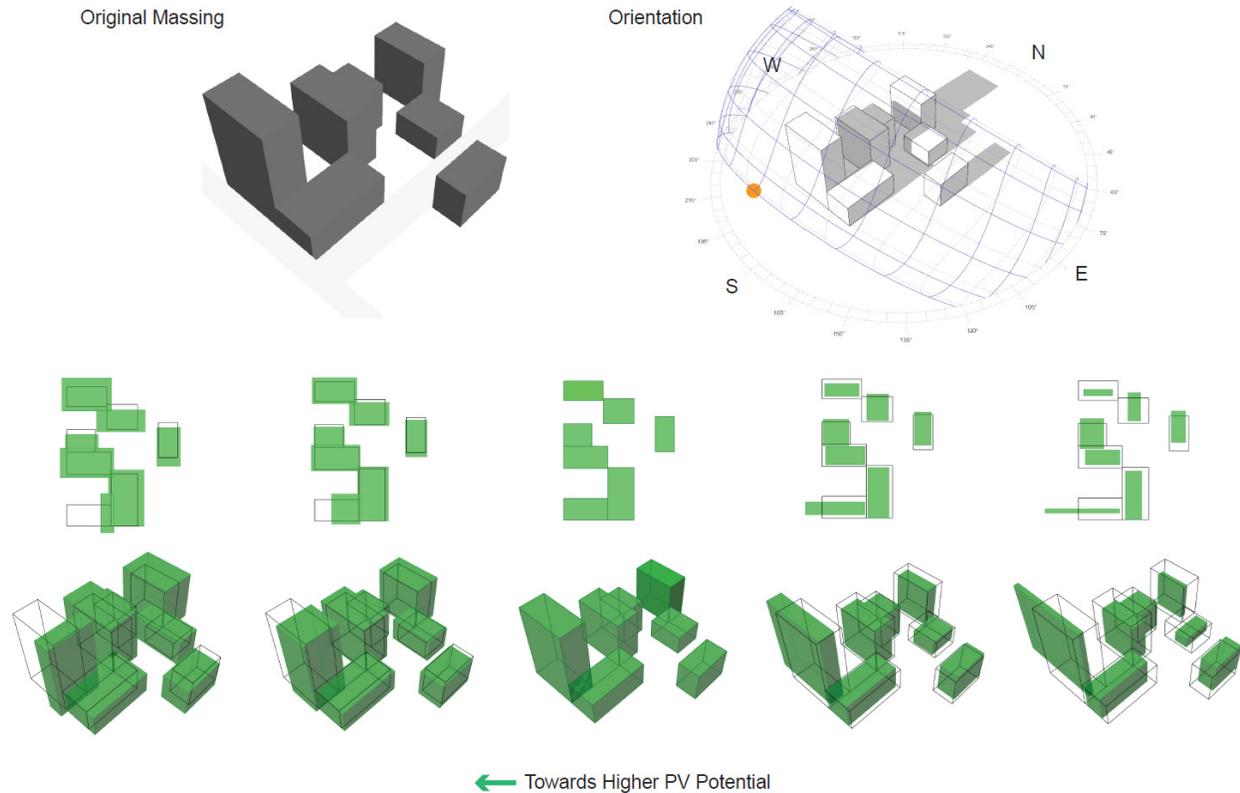


Figure 3.11: Potential massing changes that increase the rooftop PV potential for a series of buildings, automatically suggested for a given geometric input

At the same time, a human is generally required to consider the process that generated these directions, meaning that although tedious parts of the workflow are automated, design itself can still be a collaboration between human and computer. Not only must a person choose a final design, but the designer must also answer questions such as: do these results make intuitive sense, based on what is known about solar radiation in this location? Is more data required for smoother or at least intelligible results? Is there an assumption that may be missing, or a design intervention that is worth pursuing that was not found in the dataset? Consequently, this automated process is still subject to approval and modification by a human designer, even as more autonomy is given to the machine to dictate where the design should go.

3.6 Conclusions

3.6.1 Contributions towards variable analysis, transformation, and generation

This chapter makes several contributions towards more effective use of performance-based parametric modeling during architectural design. First, it demonstrates initial workflows for determining variable importance during model setup, which can assist designers during the often iterative, interactive process of

setting up a design space. Next, it provides and tests methods for generating synthetic, dimensionally reduced design variables that give designers more control over performance and global geometric form compared to typical variable-by-variable exploration. Finally, it proposes a workflow for automatically generating performance-based design variables for commonly used architectural typologies. These contributions rely on new applications of data science techniques for dimensionality reduction and pattern recognition in early building design. Such applications allow for transformation of the design space in a way that is intimately linked to performance, but can stimulate new creative avenues for parametric exploration outside of the variables originally established by the designer.

3.6.2 Future work and concluding remarks

While the methods described in this chapter show promise, there are some limitations regarding their ability to be immediately and extensively applied in architectural practice. First and foremost, some objective functions may simply be too complicated to control with a single slider. Although nonlinear functions (such as the one demonstrated in Figure 3.7) could be more accurately mapped using extensions to the techniques presented here, if the gradient of an objective function switches directions repeatedly throughout the design space, then a dimensionally reduced performance slider is only useful for small segments of that design space. However, as these tools are intended for exploration, the only risk in conducting data analysis is potentially drawing an incorrect assumption and missing a compelling area of the design space. In most cases, the discovery of complex performance response is likely to encourage further consideration and a modified design approach, which is more probably beneficial than detrimental.

In addition, simulations used to generate performance-based datasets often require expert input to ensure data quality, and the proper workflow for collaboration between designers and consulting specialists is still an open question in the field, especially when parametric modeling is involved. In general when conducting data analysis, there are initial steps—including exploration and cleaning, scaling, calculation of model error and cross validation—that require both manual effort as well as a basic background in data science. Furthermore, many powerful programs for data analysis are not fully integrated into architectural software at present, which necessitates pipelining between different software. However, many designers are becoming increasingly proficient in both data and simulation, and further efforts could be made to automate or simplify the process of data generation and cleaning.

This research also focuses on foundational data analysis techniques, with the goal of broad applicability to various simulation types and applications. Many complex or specialized algorithms for variable analysis and dimensionality reduction exist in addition to those tested here, and they might be more effective in finding useful patterns in design data. In the future, these advanced techniques should be explored and tested on case studies involving other types of performance simulations. Yet in early stage design, there is considerable uncertainty in both simulation accuracy as well as the likelihood of a given assumption to

remain true throughout the design and construction process. The goal of the variable analyses and transformations in this chapter are not necessarily to resolve those uncertainties, which are common to all early stage design, but to enhance the design space formulation process by providing decision-makers with valuable information and increased control over performance. This chapter has thus demonstrated viable techniques for achieving more effective use of data analysis during the iterative formulation process, which can lead to more effective exploration in the future.

4 Interactive Optimization and Local Exploration

4.1 Introduction

Since the previous chapter focused on design space formulation, it did not yet offer a method for higher resolution decisions while continuing to refine a parametric design. In consideration of the overall design map at the end of Chapter 2, the topic of this chapter is local optimization, which is a logical next step for parametric computational design procedures². The term “local optimization” is used here broadly with reference to the entire design process, rather than explicitly to describe exploration or optimization towards a local minimum or maximum within a multimodal objective function, as it might in mathematical literature. A local exploration in this context occurs after certain decisions regarding design variables and their connectivity have already been fixed. This situation lends itself to systematic consideration of the possible design vectors and their corresponding geometries. Such optimization processes within an established design space can be iterative, and cause the designer to reformulate and resolve in ways that are advantageous to the design problem at hand. Regardless of how a design space was established, the methods in this chapter are suitable for learning about its properties and moving towards specific high-performing areas. Out of the topics in this dissertation, the local optimization methods here have perhaps the widest

² *A version of this chapter has been published in:*

Brown, N., & Mueller, C., 2018. Gradient-based guidance for controlling performance in early design exploration. In *Proceedings of the International Association for Shell and Spatial Structures Symposium 2018*. Boston, MA.

applicability. A systematic approach to finding an effective geometry out of a range of possibilities can be immediately useful for design tasks involving non-traditional geometry, but also to answer everyday questions such as bay spacing or lateral system placement.

To implement a formal optimization technique for early building design space exploration, many designers have begun using heuristic methods such as genetic algorithms that are natively available within parametric modeling software. However, the most common workflow of setting up a parametric model, clicking a button to run an optimization, and then returning later to find a single solution has considerable limitations in early design settings. Users miss many potential solutions, have no control over intermediate steps in the process, and gain little context or intuition concerning the final solution that is determined by the computer. Moreover, secondary objectives that are not quantifiable must largely be ignored. This approach to optimization is essentially passive, unless it is repeated with different starting points, conditions, objective priorities, or algorithm settings.

Recently, evolutionary methods have been extended to allow for increased user interaction, which results in more effective preference expression during parametric design. An extensive review of such methods is available in Chapter 2. Interactive optimization techniques become particularly relevant when rapid evaluations can provide enough information to make sound decisions in conceptual design. Many early proposals of interactive optimization were thus demonstrated on simple geometry and correspondingly fast simulations, which cannot easily be extended or generalized to accept any type of evaluation required by a designer. However, the use of surrogate modeling to rapidly approximate simulation results offers a partial solution to slow simulations, and helps extend these discipline- or problem-specific contributions towards a generalized multi-objective approach. Ongoing developments in GPU processing and other techniques for improving computational speed will continue to make interactive optimization more viable moving forward.

Despite this promising outlook, the development of interactive, performance-based design tools has largely ignored the class of optimization algorithms that involves calculating or estimating the gradient for performance objectives. In many formal optimization procedures, the objective space gradient is calculated repeatedly, and this information is used to dictate the next step of an automated algorithm. The gradient shows how objective functions are changing with respect to design variables at a given location in the design space, and thus is used to move towards an optimal solution, which is the general goal of an automated optimization process.

In contrast to full optimization, gradient-based guidance can also be utilized interactively by a designer, rather than as part of an automated procedure. This chapter explores two main methods for interacting with gradient information during design: live vector visualization and objective space stepping. In the first method, a vector indicating the magnitude and direction of the gradient with respect to each variable and objective is visually projected into the modeling environment. Such visualizations have previously been

implemented in computer graphics and specific design tools. However, they could also serve as a generalized strategy for multi-objective architectural design problems, especially for models with mostly geometric variables and objective functions that can be calculated or estimated in real-time, since they can be explored interactively through continuous visualization.

The second method uses finite differences to take individual steps in the objective space that attempt to improve performance or hold it constant, based on the direction of the gradient. Objective space stepping thus enables simultaneous control in both the design space and objective space, as shown in Figure 4.1. While direct control of an objective “slider” is not necessarily possible, this stepping strategy enables exploration in which designers can improve performance while constantly adjusting step size, switching objectives, interrogating tradeoffs, and shifting regions of the design space by finding isoperforming solutions. In traditional parametric design, a human would only have control over the design space sliders on the left. A designer could move these sliders, generate geometry, and then run a simulation to see how well that combination of variables performs. Objective space stepping instead uses gradient information to map the objective space back into the design space, such that a human can make a move to improve performance, and then see how the geometry responds, inverting the traditional parametric process.

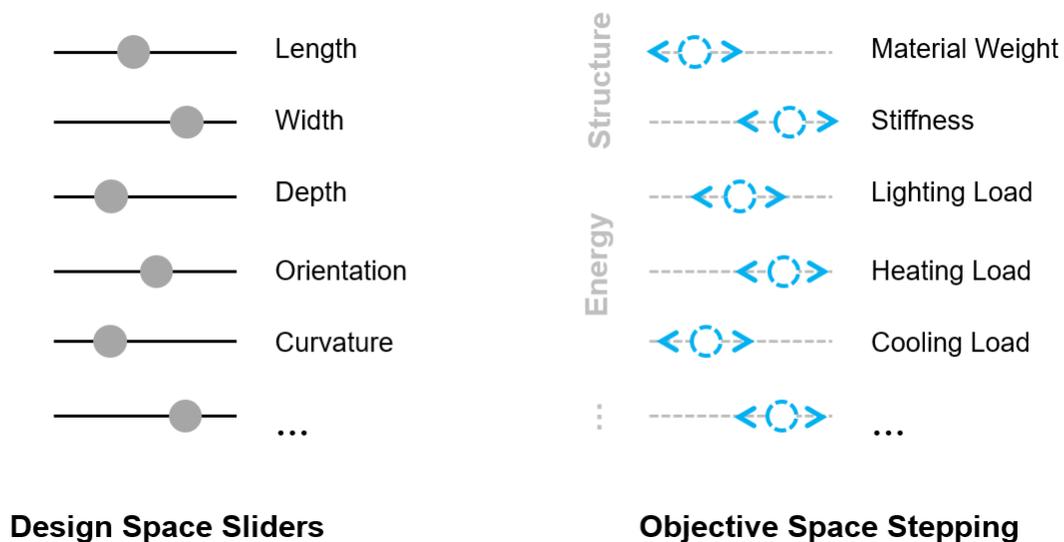


Figure 4.1: Two different strategies for manipulating parametric designs. Designers typically only control their models with independent sliders, as on the left. Using gradient information, they could also have direct control over movement in the objective space, as on the right.

Although live vector visualization and design space stepping have similar requirements to other gradient-based or interactive methods, including continuous or ordered design variables and fast or approximated performance simulations, these methods offer an alternative to passive optimization workflows by giving

designers more control. As such, gradient-based guidance is well suited for early geometric design problems with multiple performance objectives and potential design tradeoffs.

4.2 Background on gradient-based guidance

Many optimization algorithms exist for addressing design problems that have multiple or competing objectives, which is often the case even in the early stages of architectural or structural design. As described in Chapter 2, common heuristic approaches include simulated annealing, particle swarm, and evolutionary algorithms (Deb 2001). Other methods include Normal Boundary Intersection (Das & Dennis 1998) and weighted sum methods (Marler & Arora 2010), which can incorporate gradient-based search. Numerous optimization algorithms involving gradient information also exist, including line search methods, conjugate gradient methods, and quasi-Newton methods (Nocedal & Wright 2006). Heuristic and gradient methods have been hybridized in different ways (Bosman & Jong 2005; Goh et al. 2008). A review of applications in building design indicates the popularity of heuristic methods over gradient-based approaches (Evins 2013), especially when user interaction is involved, leaving the opportunity to apply gradient approaches in new ways.

The need for interactive visualization and similar techniques for supporting decision-making during multi-objective conceptual design has been addressed in other fields, including groundwater monitoring (Kollat & Reed 2007), aircraft design (Sun 2014), and product design (Baril et al. 2013). Various tools have been created for interactive performance visualization in engineering, such as RAVE (Daskilewicz & German 2013) and CityPlot (Knerr & Selva 2016). Similarly, the concept of isoperformance has been used extensively by others, including de Weck & Jones (2006) in aerospace engineering. In architecture, multi-criteria visualization tools such as Design Explorer by Thornton Tomasetti, which relies on precomputation and uploading solutions to an interactive database, and Octopus, which implements multi-objective optimization inside Grasshopper, have recently gained popularity. More information about specific tools related to these design tasks can be found in Chapters 2 and 6.

Other research specifically involving gradient visualization and stepping through the design space is also relevant. The direct visualization of gradients on top of geometric models has been employed by Whiting (2012) to direct architects towards feasible structural solutions, and by Tacit.Blue (Burnell 2014), an app for 2D engineering design problems. McHugh (2017) experimented with visualizing changes in objective performance using sensitivity analysis for multi-objective design problems. Similarly, interactive and iterative exploration of specific design space regions has been demonstrated in adjacent engineering disciplines (Tappeta et al. 2000).

In architecture, stepping through the design space based on separate performance indicators has been proposed by Kesik & Stern (2008) for the design of passive solar houses. Michalek & Papalambros (2002)

offer a system for optimizing architectural layouts interactively, while adding, deleting, or modifying variables, objectives, and constraints. Clune et al. (2012) describe a method for essentially conducting interactive gradient-based optimization, but the interaction involves designers selecting a starting point in the design space, rather than step-by-step control. The tool Performance Explorer (Wortmann 2018) allows for interactive movement in both the design and objective spaces using surrogate modeling and a dimensionally reduced visualization, which involves many concepts similarly used in this dissertation.

However, the literature does not contain implementations of gradient-based guidance for interactive conceptual design that are generalized to any simulation or geometry. This chapter proposes a new application for such general cases and demonstrates an implementation within Grasshopper for interactive gradient-based design exploration. Performance-conscious architects are thus able to use these methods to pursue quantitative design targets, while simultaneously managing non-quantifiable preferences. Due to their generalizability, these techniques are also available to structural engineers, mechanical engineers, daylighting experts, energy modelers, or any other designers familiar with parametric workflows. Since they are implemented within a parametric environment that provides access to various simulation types, they can help facilitate productive conversations between specialists when used in multidisciplinary settings.

4.3 Methodology

This section describes the procedure for implementing live gradient visualization and gradient-based design space stepping within parametric design. The mathematics behind gradients are first described, and then a simple 2D truss example demonstrates how these techniques are applied during early design exploration involving performance objectives as design goals.

The ability to provide gradient-based guidance requires the calculation or estimation of the gradients for each objective under consideration. Directions of objective improvement or isoperformance follow from this initial gradient calculation. For the applications in architectural or structural design envisioned by this chapter, which include interactive multi-objective exploration, analytical solutions are rarely possible. In order to realize this methodology as a general approach within parametric design for any type of simulation, a central finite differences approximation for the gradient is used:

$$\frac{\partial J_{z,1}}{\partial x_n} \approx \frac{J(x+\frac{1}{2}h) - J(x-\frac{1}{2}h)}{h} \quad (4.1)$$

In the above equation, $J_{z,1}(x) \dots J_{z,m}(x)$ represent all objective functions for a design, while $x_1 \dots x_n$ represent the design space variables, collectively called the design vector. The objective functions can involve any

calculation or simulation that generates a quantity that should be minimized, maximized, or targeted during design exploration. While the central finite differences approximation is utilized in the examples in this chapter, forward finite differences are also used for the tools in Chapter 6, since they require fewer function evaluations. Regardless of which finite difference approximation is used, the gradient of the objective function is given as:

$$\nabla J_{z,1} = \begin{bmatrix} \frac{\partial J_{z,1}}{\partial x_1} \\ \vdots \\ \frac{\partial J_{z,1}}{\partial x_n} \end{bmatrix} \quad (4.2)$$

The gradient at a given point is the direction of maximum rate of change of a function at that point. Thus, moving through the design space either opposite the direction of the gradient (when minimizing) or in the direction of the gradient (when maximizing) will in most cases improve the performance of the design. For multi-objective problems, the system Jacobian must be calculated, which is formed by the gradients of each separate objective:

$$\nabla J_{z,m} = \begin{bmatrix} \frac{\partial J_{z,1}}{\partial x_1} & \cdots & \frac{\partial J_{z,m}}{\partial x_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial J_{z,1}}{\partial x_n} & \cdots & \frac{\partial J_{z,m}}{\partial x_n} \end{bmatrix} \quad (4.3)$$

From the Jacobian, it is also possible to complete a singular decomposition and calculate the null space of the Jacobian:

$$U\Sigma V^T = \nabla J_z^T \quad (4.4)$$

$$V = [v_1 \cdots v_m \quad v_{m+1} \cdots v_n] \quad (4.5)$$

In these equations, $v_1 \dots v_m$ form the column space, and $v_{m+1} \dots v_n$ form the null space. From de Weck & Jones (2006), the null space of the Jacobian for a multi-objective problem contains $n - m$ vectors, and any linear combination of these vectors points in a performance invariant direction. For the specific case with only

two design variables, it is also possible to move in a direction orthogonal to the gradient in search of designs that do not change performance, since only two of these directions exist.

As described in the introduction, designers can use the information provided by the gradient in two ways. First, the gradient can be calculated and then projected into the design environment, especially with respect to geometric variables. An example of this technique is illustrated in Figure 4.2. In this figure, the design task involves a cable-supported roof, which can change shape based on support locations, cable geometry, and surface curvature. A structure with this geometry could serve a variety of purposes, including a porch or vestibule for a building entrance, or an overhang for a sidewalk or bus stop. A parametric model of the design space is linked to simulations and estimates for the deflection, carbon emissions due to structural material, shaded area, and number of connections. As a designer moves around certain nodes or other variables in the design, gradient information can be calculated for each of the objectives. This example was employed for a user study asking participants to design an overhang for a restaurant's outdoor seating area. As such, more information about the objective models and assumptions can be found in Brown & Mueller (2016).

More generally, when a geometric design variable describes the location of a node for a structure or building, a visual vector can be projected originating from that node, pointing in the direction that the node should move in order to improve performance. The lengths of the vectors represent the relative magnitudes of how quickly the objective is changing at that point, with a longer vector showing that movement in that direction can make performance much better than for a shorter vector. For variables that are not geometric but still have a meaningful gradient, additional visualizations could be beneficial, such as projecting a scaled vector over the control slider of non-geometric variables, to indicate which direction it should move. Although the provided example is static, this gradient information is ideally dynamic, so that designers can adjust variables in real-time and gain feedback on how the objective functions change throughout the design space.

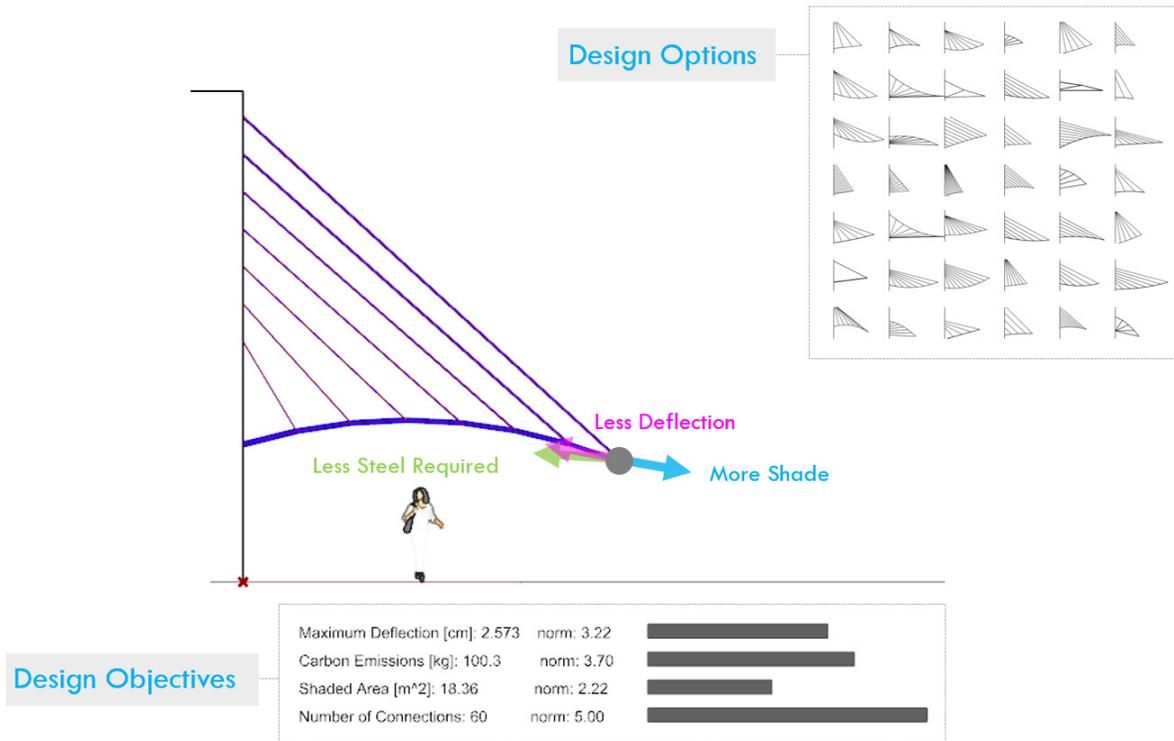


Figure 4.2: Visualization of gradient vectors projected into a design environment

The second design approach involves using gradient information interactively to take discrete steps through the design space based on the objective space, while exploring both quantitative and qualitative design implications. This method is similar to optimization algorithms, but allows designers the control and flexibility to switch or add objectives, change step sizes, or otherwise modify the design at any point, all while visualizing how the design is morphing and continually assessing its quality in ways that are not reducible to objective functions. The author has implemented this approach as a compiled C# component in Grasshopper, which was used to generate results for the design case studies in this chapter. The component, called Stepper, allows a designer to find directions of improving performance or isoperformance for any parametric model and any objective functions modeled within the Rhino/Grasshopper environment. More information about specific tools related to interactive optimization can be found in Chapter 6.

Stepper asks the user to provide variable sliders, which dictate the values and bounds of the design vector, as well as a list of objectives and an objective index, step size, and direction for moving through the design space. The direction can be either up, down, or isoperformance. When activated, Stepper will first calculate the gradient through finite differences. The component then takes one step in the direction of the gradient (or its opposite, if direction is negative) by changing each variable slider value from:

$$(x_n) \text{ to } (x_n + \frac{\partial J_{z,m}}{\partial x_n} * h)$$

This is done using the step size normalized by each variable range, denoted as h . If an isoperformance direction is desired, Stepper calculates the null space of the system Jacobian and selects a path to move based on these vectors. In the two-variable case, Stepper moves either left or right with respect to the gradient. For higher dimensional problems, selecting appropriate null space vectors is more complicated, and will be discussed later.

To better visualize how these gradient-based design strategies are applied in conceptual design, consider the simple seven bar truss illustrated in Figure 4.3 and Figure 4.4. The variables for the problem dictate the height and width of the truss, while the objective is to minimize its weight while carrying a single point load at the center. The simulations for the truss were completed using the Optimize Cross Section feature within Karamba. For this seven bar truss, the right side of Figure 4.3 shows steps taken away from an initial design based on the estimated gradient, either in the direction of improving performance or along what are intended to be isoperformance contours. The paths through the design space are visualized on top of the contours of the objective space and its optimum, which were approximated prior to initiating exploration. Since the sizing algorithm leads to a discontinuous design space, approximation techniques do not always lead to the desired solution, especially if inappropriate step sizes are chosen. However, by allowing users to interact with the design process, they are able to continuously test, adjust, consider, and move, while gaining a better understanding of the design space and ultimately arriving at an acceptable solution.

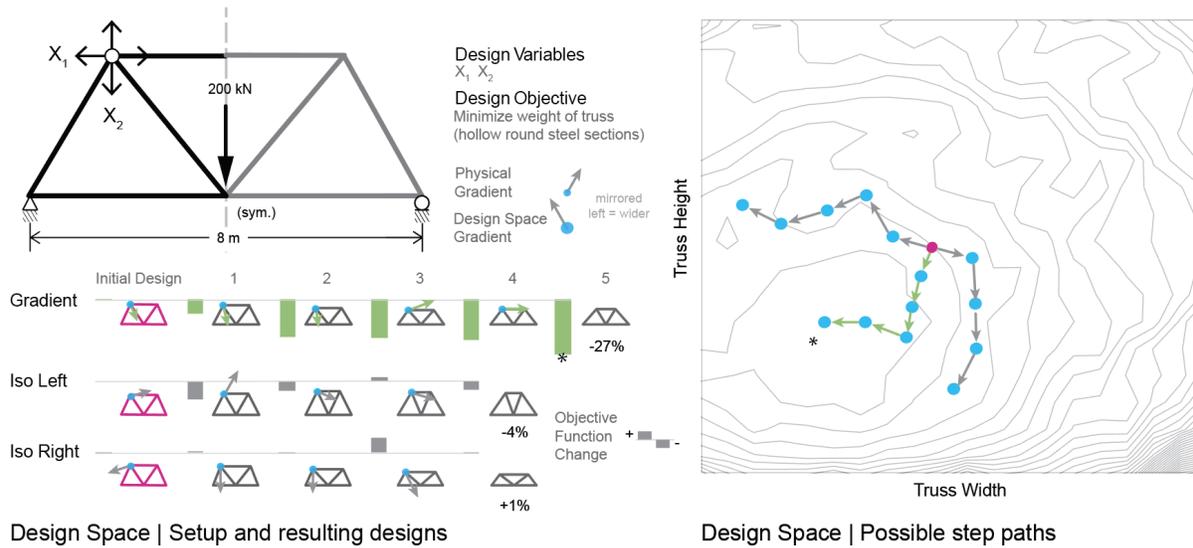


Figure 4.3: Potential steps in the design space based on gradient information about how the design is performing

The left side of Figure 4.4 zooms in on a single step of the truss problem, depicting the level of information that could be provided to a designer at this point. Once the gradient is estimated, this figure shows a visualization of potential steps in directions of improving performance or static performance. While moving farther away from the initial point reduces the likelihood that additional steps in that direction will continue to influence objective functions in the same way, this information can all be helpful to understanding the design space. While stepping around to explore different designs, the local shape of the objective space and its gradient changes, as demonstrated by the actual simulations depicted as blue dots on the right side of Figure 4.4.

In this graphic, steps in any direction closely match the expected behavior, either improving or worsening along with the gradient, or remaining flat when moving in an estimated performance-invariant direction. This relationship fades when moving away from the initial point—however, an interactive visualization can be updated with each step, gradually informing the designer of where best to move and how performance is likely to change while exploring the design space. These concepts are tested further in the following case studies, which consider performance-drive designs related to energy, structures, and daylighting. In these examples, objective function results are shown as a percent change from a baseline performance.

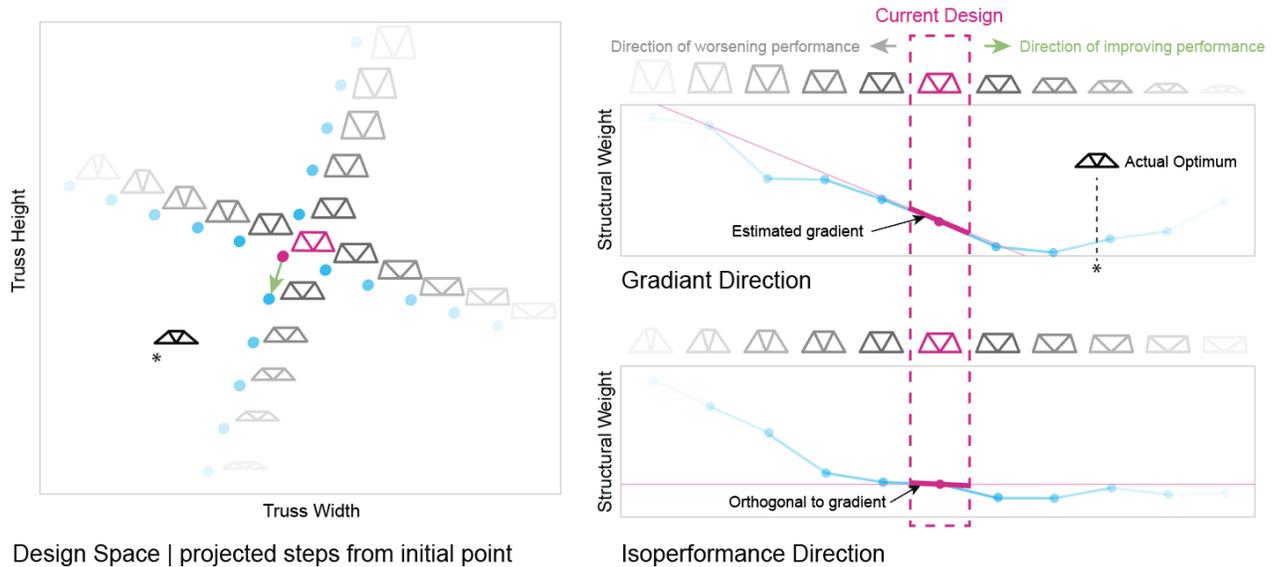


Figure 4.4: A visualization of projected outcomes for a single step in the design space

4.4 Gradient stepping for a courtyard building design

The next case study involves a conceptual massing design for a rectangular courtyard building. This problem relates to early massing studies regularly completed by architects when developing building forms across a variety of scales, locations, and uses. It was primarily inspired by two existing buildings: “The Whale”, designed by de Architekten Cie. and constructed in Amsterdam in 2000, and VIA 57 West, designed by BIG and built in New York City in 2016 (see Figure 4.5). Both buildings are essentially a courtyard typology, but the designers pushed and pulled on different portions of the massing in ways that generated considerable visual expressiveness and variation in floor plan. Although the designers may have used computational methods during the design process, it is unlikely that the designers had live simulation and corresponding performance feedback when experimenting flexibly with this geometry. This case study reimagines an early design process in which interactive, gradient-based optimization is used to move in the objective space while simultaneously understanding the geometric implications.

The simplified building is three stories tall, with a footprint of 30 x 40 meters. The design variables are the corner heights of both the interior and exterior facades, which can move both up and down, causing double curvature, self-shading, and other conditions that affect energy performance. The objectives are the PV potential of the roof geometry, which should be maximized, as well as the lighting, heating, and cooling loads calculated for the geometry, which should be minimized. Together, the combined loads minus the PV potential represent the total energy requirements that depend on the building massing, which neglects hot water, equipment, and other loads. Although this example is for demonstration purposes and the model was not as refined as later design stages would require, basic model settings are provided in Table 4.1. A

more comprehensive case study involving building energy modeling is given in Chapter 6. Despite the courtyard example containing one overall objective for the entire building, it is still useful to give designers control of these separate loads in the conceptual stage. For example, designers might have specific non-geometric mitigation strategies for certain loads that will be added in a later design phases, and want to first focus on lowering certain impacts through geometry. The unmodified form closely resembles a standard building used for building physics simulations, but geometric complexity is added when manipulating the design variables, since this leads to different pitches and even curvature on the roof.



“The Whale”, Cie., Amsterdam (2000)



VIA 57 West, BIG, New York (2016)

Figure 4.5: Two examples of recent or contemporary courtyard buildings that inspired the massing case study

Table 4.1: Basic model settings for the PV and energy simulations

Model Type	Model Parameter	Model Setting
PV Simulation	Panel efficiency	15%
	Effective area	0.80
Energy Simulation	Floor to floor height	2.7 m
	Window to wall ratio	30%
	Location	Boston
	Heating / cooling set points	20° C / 26° C
	Occupancy, equipment, lighting schedule	Always on
	Occupant density	0.2 p / m ²
	Equipment power	12 w / m ²
	Lighting power	12 w / m ²
Dimming	Continuous, at 500 lux	

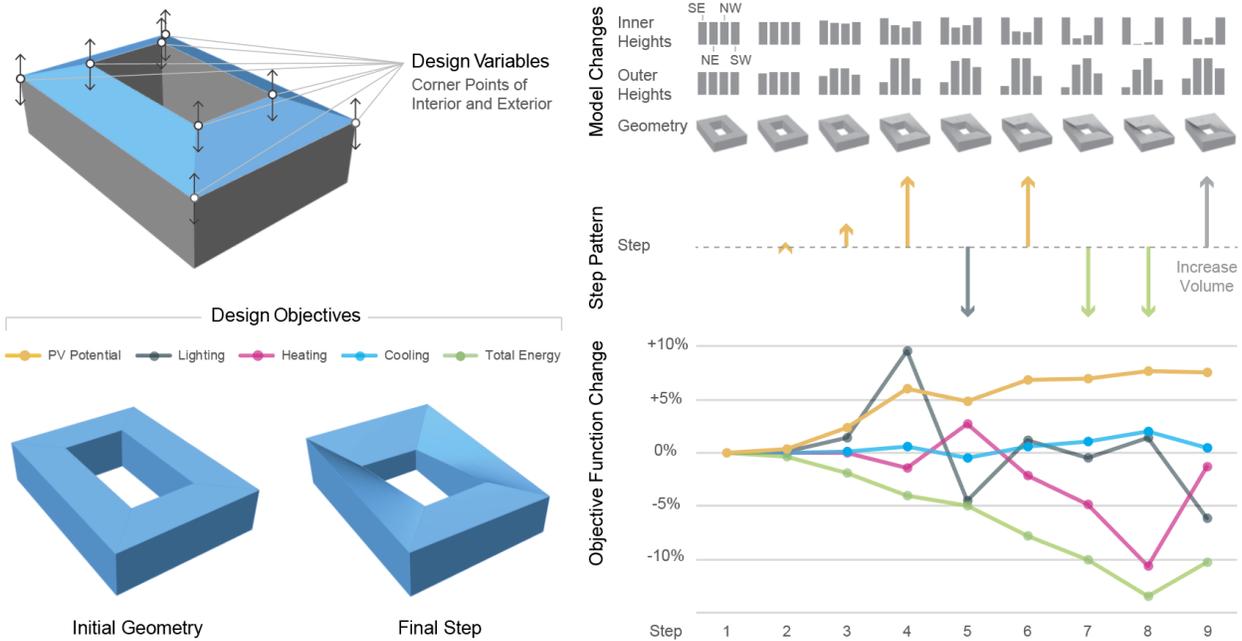


Figure 4.6: A potential path through the design space based on steps taken for different performance objectives

The geometry of the building was modeled in Grasshopper, and the PV and energy simulations were completed using the analysis plug-in DIVA. The massing is broken into three floors and four zones for each floor for the energy models. A Boston weather file was used to complete the annual energy simulations. In order to achieve full interactivity for exploration, a surrogate model of each load was constructed using the Tilde component within the Design Space Exploration suite of tools. This component took in 200 simulations that were precomputed and generated an Ensemble Neural Network model for each load. Upon completion of this case study, every design considered was re-simulated to ensure that the surrogate models provided acceptable accuracy for design exploration. The vast majority of surrogate values used were less than 3% from a simulated value, which is essentially within the error of the energy simulations themselves and considered reasonable for an early design study comparing the relative performance of different options.

Figure 4.6 demonstrates one possible design space stepping history for the problem created by the author, out of essentially infinite possibilities. This example begins with a basic flat roof geometry, establishing a baseline energy usage. At first, steps are attempted to increase the PV potential of the building, in different sizes to gain an understanding of what step size leads to a worthwhile design response. After taking three steps to improve PV, the designer experiments with different objectives to understand how the geometry changes in response. Realizing that most steps are reducing the overall building volume, which may be

cutting into usable floor area in some third-floor locations, the designers decided to add volume as an objective and step in a larger direction to test sensitivity.

Such considerations must be built into an automated optimization process as a constraint, which can be complicated to manage or dictate prior to exploration. However, interaction allows for all secondary design requirements and goals to be considered by the designer during the process and respond accordingly. In addition, any design in this history may be worth considering, and at any point, a different direction could be taken to explore another interesting area of the design space. In each step in this example, the objective functions provided the desired response, essentially giving the user an ability to control the design by adjusting objectives as well as variables, which is a potentially powerful approach to performance-driven design.

4.5 Isoperformance stepping for an airport terminal design

The second case study explores tradeoffs between daylight and structural efficiency for a long span roof design inspired by the SFO International Terminal in San Francisco, CA. While the last case study focused on stepping with the gradient, this section considers interactive stepping for isoperformance. The mathematical framework for isoperformance in this chapter comes from de Weck & Jones (2006), but is here applied interactively for design rather than in an automated fashion. De Weck et al. (2002) develop three algorithms for identifying isoperformance designs: branch-and-bound design space exploration, gradient-based contour following, and progressive vector spline approximation. This chapter explores the interactive application of gradient-based contour following, since it is conceptually linked to strategies that rely on gradient information for direct objective space movement.

The design problem has 11 variables related to the global geometry that affect both structural form and availability of interior daylight. The structural objective is to minimize steel weight subject to a 4.80 kN/m² vertical load and 1.44 kN/m² lateral load, again calculated using the sizing algorithm of Karamba to find members that are large enough to adequately handle all load combinations. The daylight objective is to maximize Spatial Daylight Autonomy (sDA), simulated by DIVA assuming a generic white interior and translucent panels for all skylights and clerestory glazing. This case study was also explored in Chapter 3, and thus more information about the modeling assumptions can be found in Section 3.4.2.

As with the courtyard case study, a surrogate model of each performance objective was created and used for analysis. The 11 variables allow for testing two effects related to the curse of dimensionality that arise with this methodology. The first potential issue is that steps that are technically similar in size to those in lower dimensional problems may lead to differences that are less visually and geometrically meaningful, since the distance travelled in each variable direction contributes to the total overall step length. The second

issue is that many null space vectors calculated from the Jacobian can point almost entirely in the direction of a single variable, which is not necessarily a compelling direction for exploration.

Both issues relate to the probable purpose of isoperformance design exploration: taking large enough steps that the resulting designs are meaningfully different, while ideally maintaining similar performance. Consequently, this chapter tests the design response of various step sizes, as well as different strategies for mitigating the null space vector issue. The exploration step sizes represent faster and slower paths through the design space. For the selection of isoperformance directions, three possibilities are tested: randomly selected vectors, cycling through the null space vectors and choosing the first vector for which the difference between the first and second variable changes is above a provided threshold, and a hybrid of the two methods.

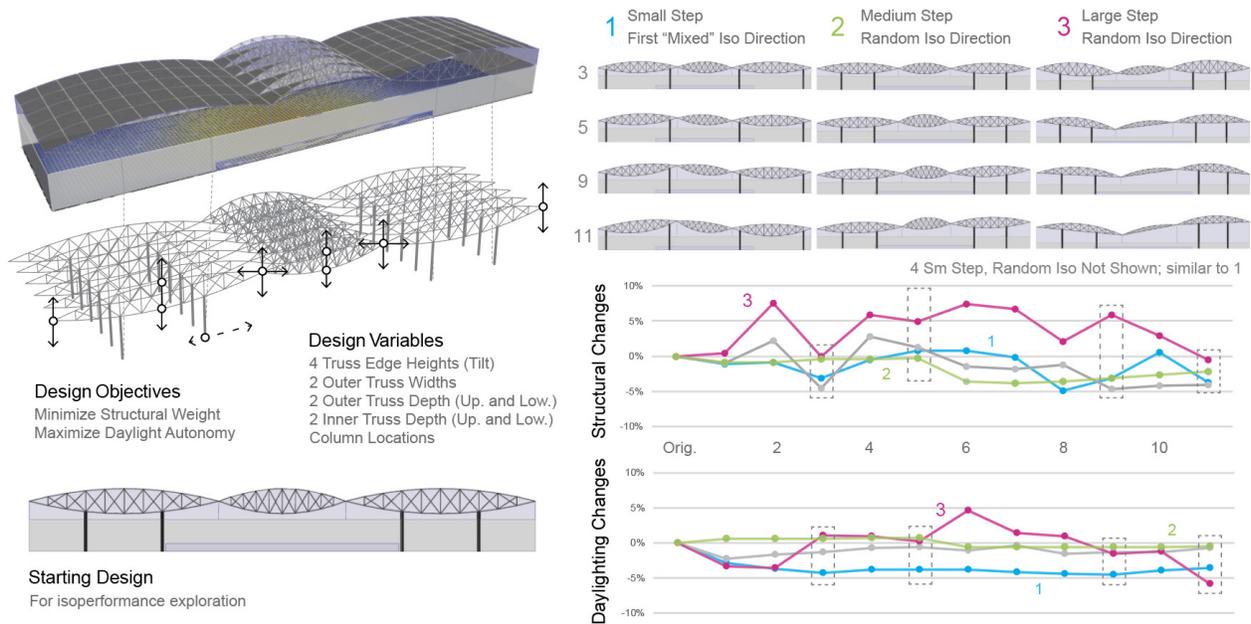


Figure 4.7: Attempted isoperformance exploration using different approaches for step size and choosing directions

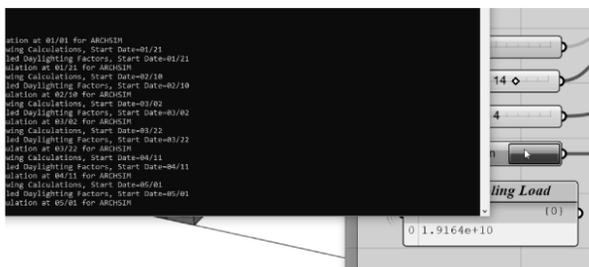
The histories of example isoperformance stepping paths are provided in Figure 4.7. Even on initial consideration, the calculated design space directions do provide visually diverse possibilities for the roof geometry. However, the ability of this method to find designs with constant performance is mixed, as some steps result in larger performance changes than is likely desirable. In the specific case of the largest step size, the surrogate model likely does not perfectly predict the actual performance in this region, as the shallow central truss should have a much larger structural performance penalty. Although designers could pick different step sizes to suit their own thresholds for visual and objective difference, it is clear that further experimentation is required to pick best default strategies for an effective isoperformance design tool.

While this isoperformance methodology is included in the interactive design tools described throughout this dissertation, and was available to participants of the design study described later, there are areas for future work that can potentially improve the approaches described here.

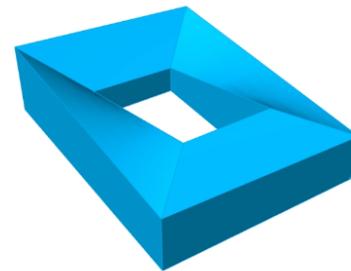
4.6 Conclusions

4.6.1 Contributions towards interactive optimization

The main contribution of this chapter is the proposal of a gradient-based interactive optimization method for generalized, multi-objective, conceptual exploration. This method has been demonstrated on three case studies of different scopes and dimensions that are relevant to early building design. When using surrogate models for objective functions, designers can conduct design space exploration by moving directly in both the design space and objective space, while synthesizing all non-quantifiable objectives at the same time. In terms of maintaining focus and creative flow, this truly live interactive experience enables a more dynamic design process than the state-of-the-art in rapid simulation, even when such simulations are connected directly to parametric design environments. Rather than tools or methods that have been developed to aid decisions made by designers of one particular discipline, such an approach is agnostic to simulation types, and can thus be used to study design objectives including structure, energy, daylighting, and others. Further validation of the effect of this interactive optimization framework on the design process is presented in Chapter 7.



Repeated Simulation



Live Exploration

Figure 4.8: Comparison of delayed and live design environments. Although difficult to demonstrate on a static page, the environment on the left includes repeated simulations, which run every time the geometry is adjusted. In the environment on the right, the geometry can be manipulated live, providing direct feedback on geometric changes.

4.6.2 Future work and concluding remarks

This study has a number of areas for future work. The author and collaborators have created an initial Grasshopper tool for design space stepping, which essentially provides the interactive environment demonstrated in Figure 4.6. The tool contains controls for stepping up, down, and in an isoperformance direction at any point, modifying the step sizes, switching objectives, leaving certain variables out of the step calculation, and the simultaneous ability to control design space variables and objective space directions. Along with these controls, the tool interface includes a visualization of changing performance, as well as features for saving and returning to previous designs. Further details of this tool will be described in Chapter 6, and it was used for the design study in Chapter 7.

However, there is more than can be done to improve the interface, provide additional visualizations, and ensure greater accessibility to such a tool within parametric environments. Future tools could also integrate additional functionality, including stepping in directions based on linear combinations of objectives, which can lead to more direct multi-objective exploration. The author has also implemented a preliminary automated procedure for the special case of bi-objective problems, in which the two gradients can be bisected repeatedly until the Pareto front is found, at which point the individual gradients can be followed to locally explore this front. Early results indicate that this method can be quicker and require fewer evaluations than evolutionary procedures for the same task, but this concept requires further exploration and analysis.

In addition, the isoperformance stepping strategies for higher dimensional problems require further development and testing before this functionality can be a robust design strategy. It is also worth noting the magnitudes of possible errors introduced into the simulations by attempts to make this workflow interactive. For the case studies, changes in performance were often within 5-10%, which may be in the range of error for the initial simulations, especially when using surrogate modeling. However, such modeling techniques can be useful for early stage design and creative exploration of form, especially when they use relative simulations between alternatives within a single design space to rank these alternatives rather than calculate their absolute performance. In these cases, and especially for multi-objective geometric problems in which the designer seeks to balance both qualitative and quantitative design goals, gradient-based interactive strategies can provide a useful alternative to passive optimization routines.

5 Diversity-driven Ideation and Design

5.1 Introduction

This chapter supplements an overall multi-objective data approach to computational design by providing a framework for how diversity, and consequently divergent portions of a design process, can be measured, implemented, and understood³. As mentioned in Chapter 2, the early stages of any design process often include rapidly generating a variety of potential solutions to a problem through brainstorming. Historically, this part of design has been driven exclusively by humans, frequently accompanied by sketching or other forms of representation to help communicate and clarify ideas (Goldschmidt 1994). Especially in the field of architecture, a premium has often been placed on creativity and innovation during this process, valuing the ability to come up with many different ideas for possible buildings and structures, as well as ideas that have not been proposed before. Design in architecture requires balancing aesthetic, technical, functional, economical, and other concerns, with different priorities taking precedence depending on the building and its context. Increasingly, computers are able to help analyze how a design is performing, navigate the relationships between different design goals, or even overcome human cognitive obstacles (McCaffrey &

³ A version of this chapter has been published in:

Brown, N. & Mueller, C., 2019. Quantifying diversity in parametric design: a comparison of possible metrics. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 33(1), pp. 40–53.

Spector 2017). Despite these advances in how design priorities are pursued, fundamental aspects of design remain synthetic, creative, and human.

As such, many computational design methods require interaction with humans primarily in order to be effective, but also to appeal to human designers as a potential creative aid in the first place. In facilitating effective interaction, computers can be much faster than humans at certain tasks such as rapidly generating design variations within a provided scheme and evaluating their performance (Cvetkovic & Parmee 2002), but they must do so in a way that enhances the creative process rather than restricting it. In many cases, this means that computers are used to create sets of potential design options, and these sets are presented to a human for comparison. Unless explicitly implementing optimization in the rare case that a single numerical objective overshadows all other considerations, the possibilities generated by computers must be varied and different, rather than all essentially representing the same or similar options. Otherwise, they cannot provide human designers with the flexibility to exert preferences or simultaneously consider design goals that are more difficult to quantify.

The concept of diversity, which here describes a group that is composed of different elements or qualities, is thus helpful in computational research and practice. It is theorized by many that more diverse design spaces lead to better and more compelling design outcomes (Rusch 1966; Kolarevic 2004; Gane 2004). Although much of the recent focus on design diversity has been catalyzed by developments in digital technology, this assumption undergirds traditional, non-numeric approaches to brainstorming as well. In these analog situations, humans have often intuitively judged the diversity of potential groups of designs. Yet since computers work more effectively with numbers, there is a need to translate the concept of design diversity into the language of algorithms and code by formulating a way it can be measured numerically. This is especially true as architecture and engineering move towards performance-based design, in which rapid generation and analysis are increasingly integrated into the early-stage design process.

One potential application for diversity measurement in computational design is the ability to build mechanisms into algorithms that ensure the results they produce are diverse enough to be interesting to designers. For example, in genetic algorithms, diversity is persevered in the process itself by using crossover and mutation (Srinivas & Patnaik 1994) or some other means. A popular algorithm in this category, NSGA-II, also operates a crowding distance check on results in performance space to ensure a representative spread along a Pareto front (Deb et al. 2002). However, this distance is only considered in objective space, and does not directly translate to diverse solutions in the design space, meaning that designs with different performances could still look the same and not be diverse enough for explorative design. Many optimization techniques are similarly limited by focusing only on differences in objective function values. Direct measurement of diversity based on a design set's characteristics rather than its performance could help future algorithms generate solution sets that are more valuable for creative designers.

One could also envision a process in which designers are not only given the best performing solution for a given objective, but also able to tune the algorithm to produce a more diverse set while understanding how much of a performance penalty is taken for this higher diversity. In addition, diversity metrics could be used in research, such as comparing which algorithms are more effective as creative design tools by determining which ones yield more diversity of solutions. In each of these cases, the existence of a numerical diversity metric that is trusted to accurately represent a human understanding of the concept could greatly enhance interactive computational design processes.

5.1.1 Organization of chapter

Based on the need for a more comprehensive understanding of diversity measurement in computational design, this chapter consolidates existing information on the topic and offers new data concerning the link between quantitative diversity metrics and human perception of design diversity. Section 5.2 provides a literature review of diversity measurement and potential metrics for use in parametric design, before detailing the contributions of the chapter. Section 5.3 describes the methodology of a survey-driven research study that tests how well these measurements match human judgments of design diversity and the relative performance of each measurement technique. Section 5.4 presents the results of the research study, while Section 5.5 offers a discussion of the results and contributions, as well as an example application. Finally, the last section indicates future work and gives concluding remarks.

5.2 Literature review

5.2.1 Background on diversity measurement

Diversity, especially when it is being measured, is also called variation, dissimilarity, or novelty in academic literature, although there can be slight differences between definitions. In comparison to other disciplines, architecture and related design fields have only recently become concerned with quantifying diversity, as parametric and other forms of computational design have begun using algorithms and processes to generate high-performing yet expressive geometries (Brown & Mueller 2016). As such, many mathematical methods for measuring diversity are more widely used in other fields, such as biology, where measurement of population genetic diversity has been studied for much longer (Patil & Taillie 1982). Nevertheless, many of the concepts created to compare species or phenetic differences within a species are mathematical or geometric in nature, and can be applied to conceptual design, especially if design options are defined in a way that can be represented numerically.

However, a key difference between the measurements of biological diversity and parametric design diversity makes it difficult to use biological methods directly in design settings. When used in the context of biology or ecology, a diversity index is a quantitative measure of how many different types exist in a dataset, and how evenly the entities in the dataset are distributed across these types (Hill 1973). Sociologists can similarly

use these existing metrics to compare types of people (Agresti & Agresti 1978), and information scientists can use them to distinguish between characters in a password, for example. Common diversity metrics in these situations include the Gini-Simpson Index (Rao 1981), which indicates the probability that two entities randomly chosen from a set represent different types, the related Inverse Simpson Index, and the Shannon Entropy Index (Shannon & Weaver 1949), which quantifies the uncertainty of being able to correctly predict the type of the next entity in a given set (Jost 2006). Rényi's entropy (Rényi 1961) is also used for similar purposes. Within the field of engineering design, Kan and Gero (2017) utilize the concept of entropy to quantify the creativity and productivity of a design session by scrutinizing linkographs from protocol analysis.

In contrast, the different designs being compared within a parametric definition often cannot clearly fit into distinct "types", especially if there are many continuous variables. It is possible to manually group designs within a set, or use clustering algorithms to achieve categorization automatically, but parametric design spaces frequently transform too gradually and subtly for an approach relying on a first-step classification to yield meaningful results. Such sorting-based metrics could be used when comparing typologically different conceptual architectural sketches generated during a brainstorming process, but this comparison is less directly relevant for implementation in computational settings, and is thus outside the scope of this chapter. Although clustering is a promising idea and will be discussed later, this chapter focuses on faster, single-step, geometric measurements that can be applied for early-stage conceptual design models. Nevertheless, despite the difficulty in using probabilistic biological measurements directly, many mathematical techniques related to their calculations can be applied to quantify diversity in parametric design.

It is possible to borrow concepts from other fields because most diversity metrics rely on distance measurement. Distance measurement is itself a distinct topic, and although it is used to create the diversity indexes mentioned above, it exists within the more fundamental fields of geometry and statistics. Since the question of diversity in parametric design is not necessarily how *many* different types are present, or how well distributed these distinct types are across a set, but rather involves characterizing how *far apart* different designs are in relation to all other designs, distance measurements can be especially useful for design settings. However, in most cases a distance measurement must be combined with another geometric or statistical concept to create a meaningful metric relating to an entire set of designs. The difference between measuring raw distance, versus attempting to quantify the diversity of an entire design set based on combinations of distances or other geometric properties, is illustrated in Figure 5.1.

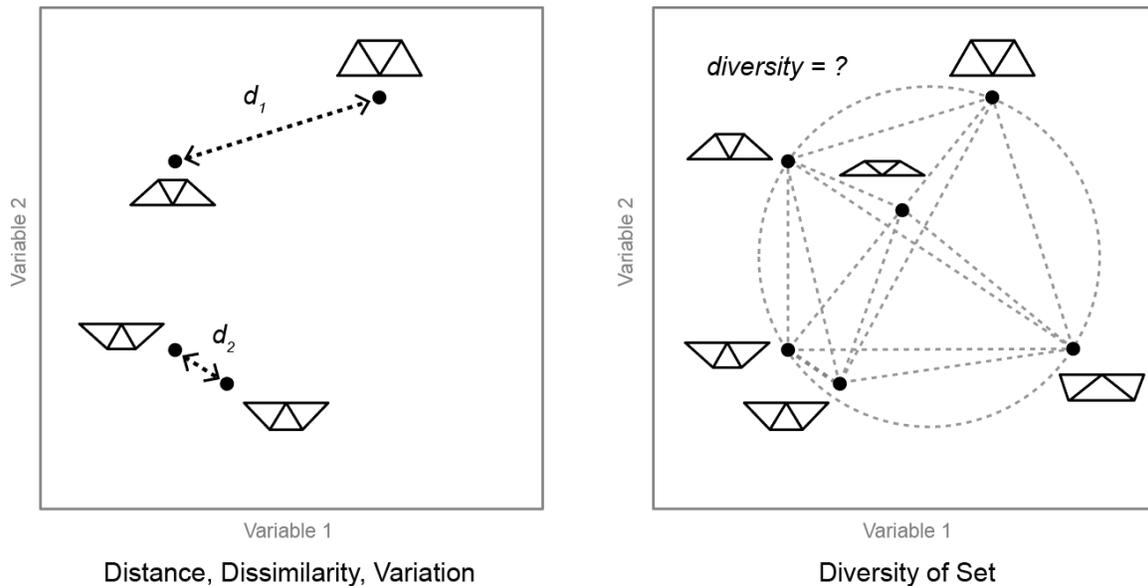


Figure 5.1: A diagram illustrating the difference between measuring distance and measuring the diversity of a set of designs. In the image on the left, the distance between two sets of parametric designs can represent their dissimilarity or variation, as closer designs look more similar than those far away— d_1 is clearly greater than d_2 . However, if trying to quantify the diversity for the set on the right, judgements must be made about which geometric properties to use.

Relevant distance measures include Euclidean, City-block, Correlation, Geodesic, and when distance is between groups rather than points, Chi-squared and Mahalanobis distance (McCune & Grace 2002). Distance metrics are widely studied and characterized in computer science, as the choice of method can have a significant effect on algorithmic performance in applications such as pattern recognition, data mining, segmentation, and signal analysis (Perlibakas 2004; Kokare et al. 2003; Wang et al. 2007). Although there are important differences to each, many of these metrics attempt to normalize for the size of the entire possibility space or adjust based on the relative importance of the axes being considered (in the case of ecology, these axes generally represent species). As such, they often collapse to Euclidean distance in simple cases when axes are equal and already unitized. Yet each of these measurement techniques, although not necessarily directly applicable to measuring set diversity rather than pair similarity in parametric design, are worth considering because they have the potential to inspire modified versions that can be applied in architecture or structures. More sophisticated distance metrics, such as the Manhattan distance, may also be required to increase the accuracy or effectiveness of diversity measurement for some applications, including higher dimensional problems (Aggarwal et al. 2001). This chapter makes use of distance measurement in each diversity metric considered.

5.2.2 Creativity, novelty, variety, and related concepts in engineering design

In engineering design, there has been considerable research into idea generation during early-stage brainstorming, often called product ideation (Shah et al. 2000). While attempting to improve ideation methods, researchers have created empirical laboratory or simulated environment studies that require characterizing and measuring ideation effectiveness (Shah et al. 2003). Common metrics for evaluating the outcomes of brainstorming sessions include creativity, novelty, and variety. *Novelty* is defined as being original, rare, unusual, or unexpected when compared to other ideas, while *Variety* is a characterization of the explored solution space.

Measuring novelty often involves defining what is not novel, decomposing a problem into key functions and characteristics, assessing how an individual design fulfills these functions, and then grading its novelty with reference to all other possible designs. Although variety applies to a group of ideas rather than an individual idea, it is often measured by examining how different functions are satisfied and calculating differences in physical/working principles, followed by less relevant details such as dimensions. Creativity is defined in a variety of ways, including pure notions of novelty or originality (Runco & Charles 1993), or adding aspects such as quality, workability, and relevance (Dean et al. 2006), or importance and affect (Horn & Salvendy 2009). Others have proposed defining creativity in an entirely subjective manner based on observer response (Amabile 1982).

Diversity in the context of architectural computational design aims to capture some combination of these interrelated concepts of novelty, creativity, and variety. However, within parametric design there is a direct need for numerical measurements that do not depend on expert judges or qualitative assessments, and can thus be applied immediately in current design environments. Many architects and engineers working in building design are already relying on strategies such as creating catalogs of the design space or using multi-criteria optimization to obtain sets of designs for final selection (Balling 1999; Stump et al. 2003). These designs are frequently generated rapidly using parametric scripts, as described in the next section. Although performance simulation is regularly attached to decisions between options (Tsigkari et al. 2013), experts do not necessarily sort designs into hard categories of feasibility, functionality, or other traditional engineering metrics. This is because it is common for parametric options in later conceptual design to be typologically or functionally consistent, yet contain subtle visual or geometric differences that are still relevant to the designer as he or she weighs the quantitative performance and qualitative experience of a future building. Thus, metrics for diversity measurement in this chapter are valuable because they can be used to tune algorithms and processes in human-computer collaborative design, leading to creative, innovative, yet high-performance design outcomes.

5.2.3 Diversity in parametric formulations

It has become increasingly common in architectural and engineering design to define designs parametrically, using digital software to generate a geometry based on a set of variables (Burry 1996; Monedero 2000; Holzer et al. 2008). This approach lends itself to both rapid design space exploration, since many potential designs exist for a given parametric formulation, and optimization, since different design possibilities can be evaluated for performance and these variables can be adjusted accordingly using a guiding optimization algorithm (Oxman 2008a). Parametric software, such as Grasshopper for Rhinoceros or Dynamo for Autodesk enable parametric formulations as part of an integrated conceptual design process. Since parametric formulations at their root contain a numerical design vector, which denotes a point in space that is mathematically distinct from other vectors, the diversity of a parametric set of designs can be measured using mathematical or geometric techniques.

Parametric design also has limitations, which can relate to how diversity is quantified computationally. Although parametric design is more flexible than earlier digital design tools, which focused on documenting a single idea, it requires considerable input and knowledge by the designer to properly formulate the parametric logic and set appropriate variables and constraints. Parametric definitions are also commonly frozen to a specific typology, which can be limiting in architectural or structural design, and their results can be dependent on specific decisions made during the formulation of the design problem itself. In this way, a human could easily draw a suspension, arch, cable-stay, and truss bridge on a sheet of paper, and declare that this collection of designs is more diverse than any parametric set offered in this chapter. Computers would not only have a difficult time generating each of these structures from a single definition, but also measuring diversity without a hierarchical rubric of function, aesthetics, or other properties devised by a human designer.

Nevertheless, despite still being reliant on human creativity and labor, parametric design is a useful tool for all types of designers looking to rapidly iterate, prototype, and explore. Consider Figure 5.2, which illustrates diverse design spaces in the physical world and a digital design context, both of which are relevant to architecture. Diversity during the design process can stimulate creativity and imagination, leading to a more diverse built environment. At the same time, it also allows designers to consider more ideas during brainstorming, which can improve design outcomes. The left image shows bridges over the Sumida River in Tokyo, where there are 26 bridges in 27 km. Despite each site having a similar span and loading, the bridge designers responded with a diverse range of creative structural solutions rather than one standard form.

The image on the right shows a sampling of options generated by a parametric script for a hanging canopy structure. Despite limitations to parametric design, the visual diversity of the canopy solutions on the right would give designers substantial creative freedom to select an effective and expressive design. When such diversity is a desired outcome, the underlying mathematics behind design vectors in parametric modeling

becomes an advantage, since it can be measured computationally using geometric concepts. The next section describes existing methods for this purpose. First, metrics that are evaluated comparatively in the survey portion of the chapter will be presented, after which other related or composite methods will be described.

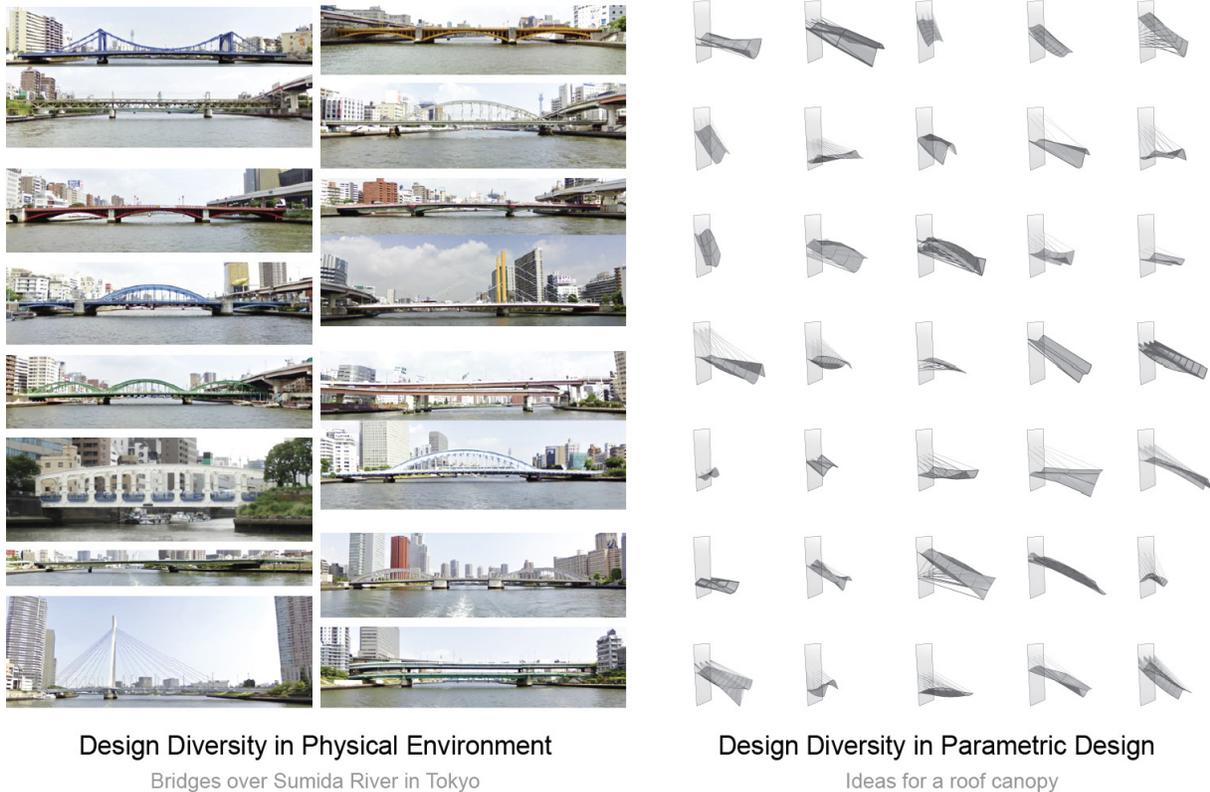


Figure 5.2: Diversity of architectural solutions for similar design conditions, in both the physical world and parametric design. The bridges in Japan demonstrate the creative instincts of architects and engineers during design, which leads them to consider larger design spaces. Although the designs on the right were generated with a parametric script, there are many visually different options to consider for the final solution. Images of the bridges are from Google Maps (2017).

5.2.4 Methods for measuring diversity in parametric modeling

This chapter consolidates and synthesizes existing literature on diversity measurement to highlight the following metrics for use in parametric design. Some of these metrics have been directly proposed and used by others before, and are cited accordingly. Each metric measures the diversity of a set of designs, which is sometimes required to contain a minimum number of designs in order for the measurement to be possible. The study does not assign priority to design variables or classify designs into species or types, and the vectors are normalized to the unit square. The problems in this chapter are two-dimensional—it is assumed that

the proposed metrics would apply to early-stage parametric designs, which generally have low dimensionality so that designers can meaningfully interact with all variables that influence visual appearance and performance. As such, standard Euclidean distance was chosen as an intuitive, commonly accepted distance measurement for these applications. Euclidean distance between two points in space is defined as:

$$dist(p, q) = \sqrt{(p_1 - q_1)^2 + (p_2 - q_2)^2 + \dots + (p_n - q_n)^2} \quad (5.1)$$

Distance from centroid to outlier

The first diversity metric under consideration is the maximum distance between the centroid of the group and each of the individual designs, effectively measuring the furthest outlier in the set:

$$D = \max[dist(x_{centroid}, x_i)] \quad (5.2)$$

This metric was proposed by Mueller and Ochsendorf (2015) and is used in Brown et al. (2015). Although a single outlier can have a large effect on the value of the metric, the distance to this outlier is only considered once, rather than multiple times, as is the case with other methods.

Sparseness at center (centroid and median)

The second metric involves a similar distance measurement, but considers each of the points in the set simultaneously. It does so by measuring the sparseness at the center of the design set, which is the average distance between a given point and each of the surrounding points. Thus, the diversity is equal to the measured sparseness:

$$D = \frac{1}{k} \sum_{i=0}^k dist(x_{centroid}, x_i) \quad (5.3)$$

This concept is taken from the novelty search algorithm developed in Risi et al. (2009) and Lehman & Stanley (2011). This algorithm, designed to overcome difficulties in finding intermediate, stepping stone solutions in fitness-driven evolutionary optimization, repeatedly measures sparseness and rewards points in sparser areas of the design space. This chapter considers two different metrics involving sparseness to evaluate the diversity of a set of parametric designs. The first takes the average distance from the centroid, as described in Equation 3 above. The second takes an average distance from the median value for each dimension of the design set:

$$D = \frac{1}{k} \sum_{i=0}^k \text{dist}(x_{median}, x_i) \quad (5.4)$$

Although these two measurements are often similar, they have different sensitivities to outliers, and would yield drastically different results in the case where the centroid is not near the median.

Enclosing Hypersphere

The next two diversity metrics involve geometric properties of the entire set rather than direct distance measurements. The first metric is the radius of the smallest hypersphere in n -dimensional space that contains each of the designs in the set:

$$D = \min(r_{hypersphere}) \quad (5.5)$$

Finding the enclosing hypersphere, sometimes called the bounding or enclosing ball, is a common problem in computer graphics, computational geometry, and statistics. Consequently, various algorithms exist for this purpose, including Ritter (1990) and Welzl (1991). These algorithms have different runtimes and degrees of accuracy, but for diversity measurement, faster approximations are likely acceptable. The idea for employing the enclosing ball to quantify diversity comes from Pavoine et al. (2005), which uses enclosing hypersphere radius while discussing quadratic entropy as a measurement for diversity when the entities being considered are divided among a fixed set of categories.

Convex Hull

The final diversity metric studied in this chapter is the volume of the convex hull formed by the designs in n -dimensional space:

$$D = \text{volume}_{\text{convex hull}} \quad (5.6)$$

The convex hull is often described in two dimensions as the shape a rubber band would take if stretched around all points in the set. Convex hulls are used by Cornwell et al. (2006) to look at habitat filtering, by Layman et al. (2007) to measure trophic diversity in food webs, and by Villéger et al. (2008) to quantify functional richness of communities. A description of the suitability of convex hulls for these tasks, as well as their pros and cons in diversity measurement, is given in Podani (2009). As with the enclosing hypersphere, various algorithms exist for calculating its volume for a given set of points, including the “convhull” and “convhulln” functions in MATLAB.

Related metrics

In addition to these metrics, other measurements exist that are either similar to those already presented, or less relevant to interactive, conceptual design. Nevertheless, they are worth mentioning, as they are potentially useful for researchers and designers with slightly different applications. One such metric is measuring the average distance between each point and every other point in the design set, while ignoring the centroid or median (Smaling 2005). However, this metric gives similar results to the sparseness measurement at the centroid, and has more potential to be thrown off by outliers, since the distance to a faraway point would be counted multiple times. Another potential metric is counting the number of points on a boundary, or taking the average distance between points only on that boundary. These concepts are related to crowding distance, which is employed by various optimization algorithms to ensure an even distribution of points along a significant edge of the design space, such as the Pareto front in multi-objective problems. Such metrics are more useful when diversity is desired along a specific boundary, or when larger number of designs are present.

As mentioned before, many other diversity metrics could be employed if the user is willing to categorize the designs in the set either manually, through clustering (Berkhin, 2006; Jain, 2010), or by creating a rubric. However, a manual or rubric approach is more labor-intensive and susceptible to the inherent bias of the categorization process. If the design space is of a relatively high dimension and there are many designs, an automated clustering algorithm could be used along with available statistical methods for diversity or

entropy mentioned in the literature review. Yet this method has issues as well, since many common clustering algorithms are not deterministic and could give results that are not replicable, especially if design space clusters are not particularly distinct. Furthermore, it is not immediately clear that an interactive design processes where a human must simultaneously consider many tens or even hundreds of design variables and iterations is effective. Although the concept of choice fatigue has been extensively studied and hotly debated (Scheibehenne et al., 2010), smaller, more manageable sets are likely the preferred approach of designers who want creative choice, but do not want to be overwhelmed by options. The metrics selected for study in this chapter focus on these types of conceptual design sets.

5.2.5 Research contributions

After reviewing existing methods for measuring diversity within parametric formulations, it is clear that although a number of possible metrics exist, there is no consensus on the single best metric. Furthermore, when considering the human-centric context of architectural or structural design, there has been no rigorous investigation as to whether or not the given metrics accurately translate to human perception of design diversity, and how the use of different diversity metrics might influence the conceptual design process. When design computation applications require such a diversity metric, many designers and researchers have simply borrowed an existing methodology or proposed their own, as the metric itself was not the focus of their research.

In response, this dissertation contributes to existing scholarship on the topic of diversity measurement and connects it to computational design in a number of ways. First, it consolidates methodologies proposed in different fields and provides an overview of available measurements for future use by both designers and researchers. Second, it demonstrates the extent to which existing diversity measurement metrics match up with human perception, which is essential to understanding how humans and computers can best collaborate while creatively brainstorming during the architectural or engineering conceptual design process. Finally, it provides a quantitative comparison between each of the methods as well as a qualitative discussion of their merits, which can assist tool developers and even designers in selecting the most appropriate diversity metric for a given application. The following sections describe the methodology and results leading to these last two contributions.

5.3 Methodology

Relying on the methods described in Equations 5.2-6, a visual survey was developed and conducted to determine the agreement of human and computer perceptions of design diversity and to compare metrics directly. This research utilized Mechanical Turk (Buhrmester et al., 2011), which is a service provided by Amazon that allows users to distribute tasks requiring human intelligence to a diverse, on-demand, temporary workforce. The survey contained a series of questions asking participants to pick the most

diverse group between two sets of designs presented visually on a computer. An analysis of these head-to-head results over a large number of participants yielded significant insight into human perceptions of diversity in relation to geometric representations of design concepts.

Mechanical Turk is commonly used for social science research, since it allows for a quick, high volume of responses that tend to be representative of the general population (Paolacci et al., 2010). Due to its advantages such as supportive infrastructure, accessibility, and anonymity, Mechanical Turk is increasingly employed in engineering and product design as well (Häggman et al., 2015; Macomber & Yang, 2011). Although Mechanical Turk is most often used in product design to obtain crowdsourced feedback from potential end-users or related stakeholders external to the design process, it has been demonstrated that unlike design quality metrics, there is a strong correlation between novelty as judged by online respondents and design experts (Kudrowitz & Wallace 2017). In that study, the researchers propose a three-attribute metric of novelty, usefulness, and feasibility for narrowing a large pool of early-stage product ideas down to manageable number for design exploration. Although experts are required to adequately judge the second two metrics, the researchers tested crowdsourcing strategies like Mechanical Turk and found them to be effective for assessing novelty. Since the evaluation of parametric design diversity is strongly related to novelty, and it involves helping computers to be more effective brainstorming partners for creative architectural design, feedback from Mechanical Turk should provide information about general human perceptions of diversity that are relevant even for trained designers.

5.3.1 Generation of test data

To create the survey content, ten sets of six design vectors were randomly generated, with each design vector containing two variable settings between 0 and 1. The diversity of each set was then calculated using the methods described in the previous sections. A visualization of these diversity measurements is provided in Figure 5.3. In these figures, each dot represents a design with variable settings corresponding to its location on the two variable axes. The variety of lines, circles, and hatches describe the geometric measurements used in the diversity calculations. A “smaller” grouping of points, such as in Set 8, has a lower measured diversity than a “wider” grouping, such as Set 10, although the metrics can disagree. For example, a long, narrow grouping could have opposite points that on their own dictate the size of an enclosing ball, which would cause this metric to measure the same diversity as in a spread grouping where each point lies on the boundary of the same ball. However, the other metrics would judge the spread grouping to be considerably more diverse than the narrow grouping. In addition, although each metric is based on some form of geometric measurement, their units are mixed, and their original values are of vastly different scales. As such, each diversity value is normalized based on the minimum and maximum of the set, which can also be seen in the bar graphs in Figure 5.3. For simplicity, this normalization is consistent throughout this chapter, and diversity measurements are presented in terms of relative percentages between sets rather than raw scores.

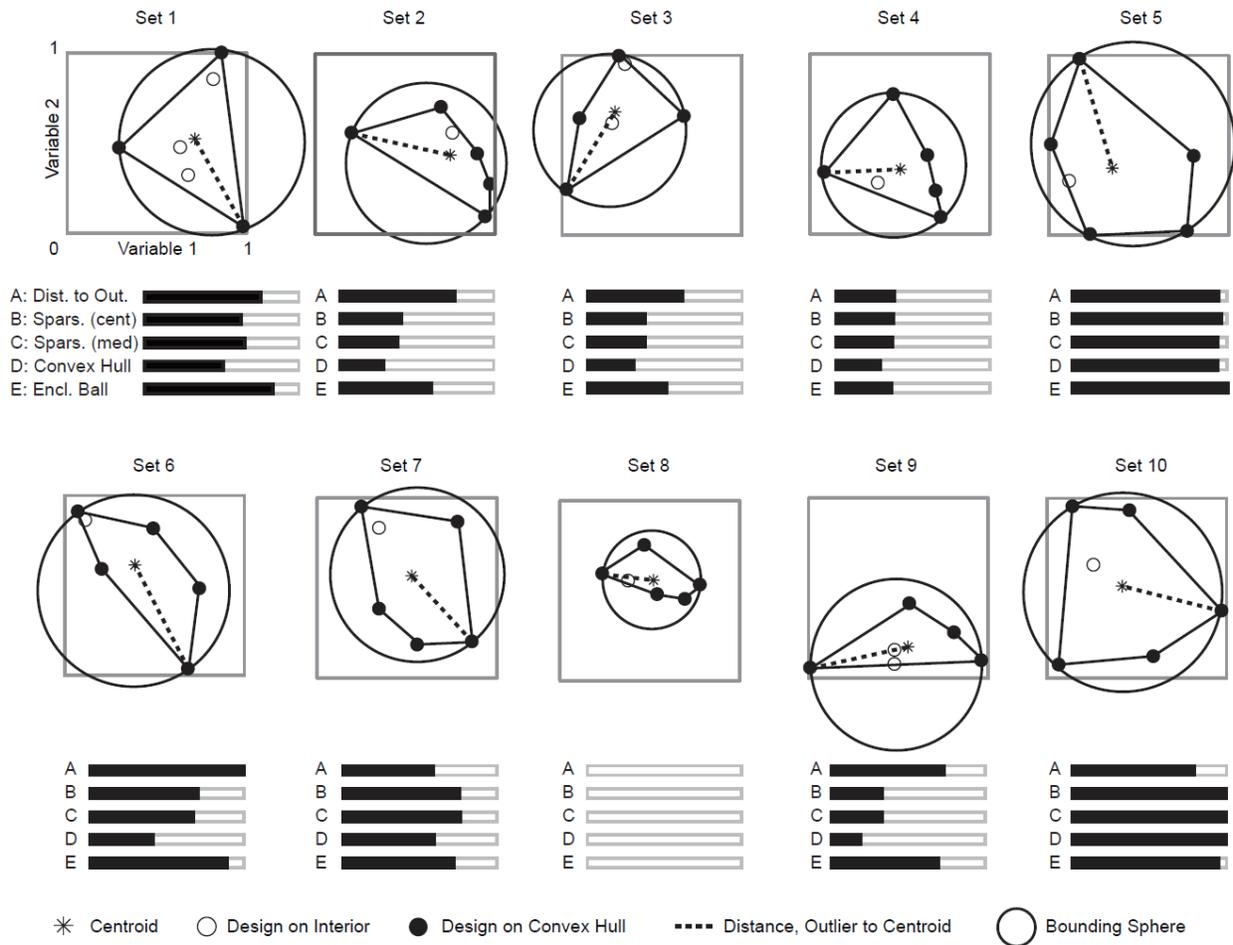


Figure 5.3: A visualization of the diversity calculation for each design set, as well as the normalized diversity value for each metric. Each dot represents a design at the corresponding variable settings, while the lines and hatches show the different calculations, described in the key at the bottom. The normalized diversity is shown by the horizontal bar graphs where the dark bar represents the diversity of the set, and the white bar shows the maximum normalized diversity calculated for any of the sets.

Next, the design vectors were fed into parametric scripts in Grasshopper that generated simple shapes, and the resulting images were recorded. Four different geometric shapes were used in the study to reduce the possibility of one noticeable, visual effect skewing the survey data, and to test for effects such as variable dominance. The chosen shapes were arches, boxes, trapezoids, and trusses, which simulate simple architectural possibilities that might be sketched as part of a conceptual design exercise. The variables in each design vector set global properties such as width and height for each shape. Although the sets of design vectors were generated randomly, the same random sets were used throughout the survey, ensuring that each participant answered similar questions. Figure 5.4 presents the set of shapes used for the survey.

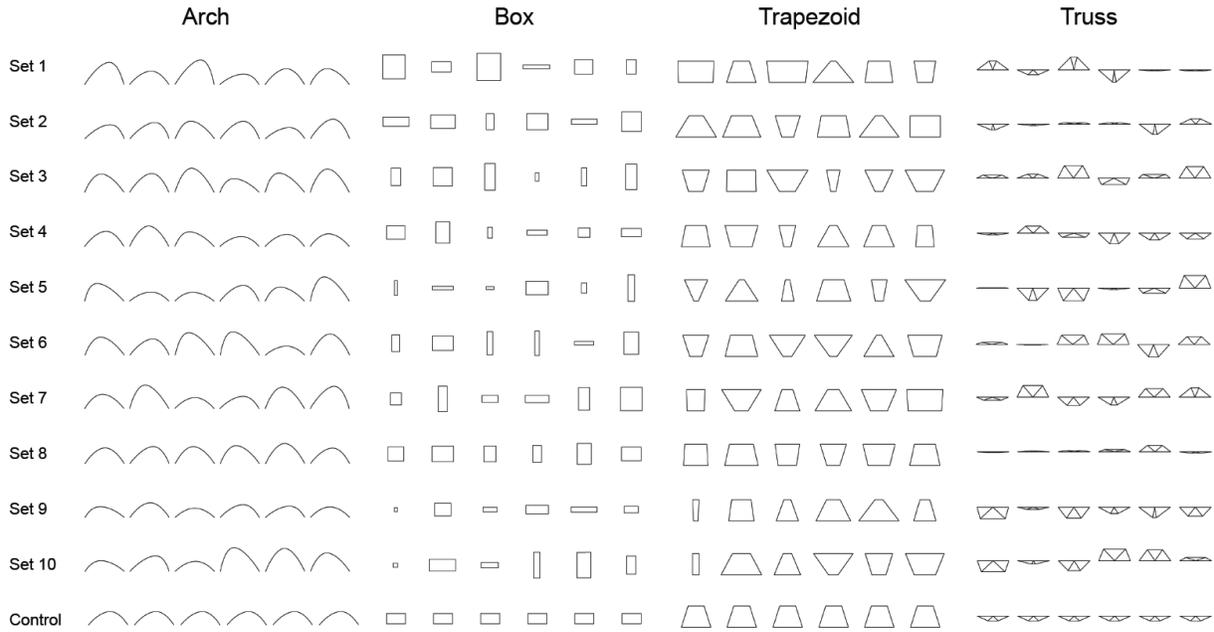


Figure 5.4: The parameterized shapes of each design set used in the survey

5.3.2 Survey

The survey itself consisted of 18 separate questions, each of which asked participants to examine two sets of geometric shapes and determine which is more diverse (see Figure 5.5). The 18 questions were broken down into three types: human-computer agreement, distinction between metrics, and controls. For the six human-computer agreement questions, each diversity metric agreed on which set should be considered more diverse. However, these questions ranged in terms of magnitude differences between diversity scores, from questions in which the more diverse set should be easily determined by visual inspection to more ambiguous comparisons. In each of the eight distinction questions, one or more of the metrics disagreed with the others on which set should be considered more diverse. Every metric disagreed with each other one at least once, allowing for head-to-head evaluations for all of the metrics. As such, an overall analysis of how often each metric was chosen can determine if there is a significant difference between metrics, and if so, which metric best approximates human perception of diversity. The remaining four control questions displayed a set of identical shapes opposite a randomly selected set, which is trivial for a human to answer correctly. The control questions ensured data quality by confirming that participants understood the instructions properly and were actively engaged in the task. The order of questions, whether each set appeared on the left or right, and which shapes were displayed were all randomized.

Diversity Survey

Question 4 out of 18

Please choose the set whose population you find MORE diverse.

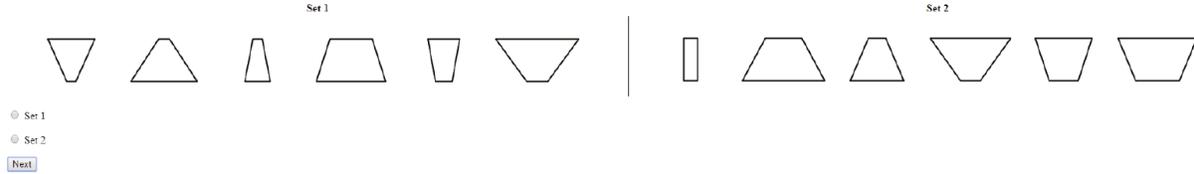


Figure 5.5: An example survey question as shown to a study participant

5.3.3 Procedure

The survey was distributed to workers through Mechanical Turk, and 400 quality survey responses were recorded. No additional qualifications, other than being a Mechanical Turk worker, were required to participate in the survey. In the Mechanical Turk interface, workers were given a brief description of the task and a link to the survey, which was hosted on a private server. If the worker passed each of the four control questions, he or she was given a unique code to input back into Mechanical Turk for payment. Upon successful completion of the survey, workers were compensated \$0.50. The survey recorded 65 survey attempts that were not accepted along with the 400 accepted responses, for a success rate of 86%. The surveys were taken anonymously, and there is no way for the authors to directly identify the participants. Each participant used his or her own computer, and no additional resources were provided other than the survey itself. The preferred set for each question was collected, along with the time spent, as well as the shape and position that the participants viewed. The average time for survey completion was just under 3 minutes, with participants spending an average of 22 seconds on the first question presented to them and around 7 seconds on the last several, indicating a slight learning curve and suggesting that a longer survey likely would have led to diminishing concentration.

5.4 Results

5.4.1 Human computer agreement

The results for the human-computer agreement portion of the survey are presented in Figure 5.6. The cumulative percentage of how many respondents agreed with all metrics for each question are plotted against the average higher diversity of the correct set as measured by the metrics. The x-axis refers to the relative diversity of the higher set, or the diversity of higher set divided by the diversity of the lower set, as a percentage. As expected, a greater difference in diversity between the two sets led to more participants identifying the correct set as more diverse. For the question testing the largest difference in diversity, 87% of respondents agreed with the computer, with a confidence interval of 3.3. On the other end, when the computational metrics barely agree on which set is more diverse, many humans were not able to correctly identify the more diverse set—survey responses for these questions were close to 50%. These results indicate

that in general, computer metrics are able to match human perceptions of diversity, a finding that encourages their usage in interactive, creative design. Yet at a certain point, the distinctions made by diversity metrics become almost meaningless to humans (see Figure 5.7). While these metrics still may be useful at small differences in measured diversity for certain applications, the magnitude of relative diversity measurements must be considered in interactive computational design, and perhaps even controlled by the user, in order to yield meaningful results.

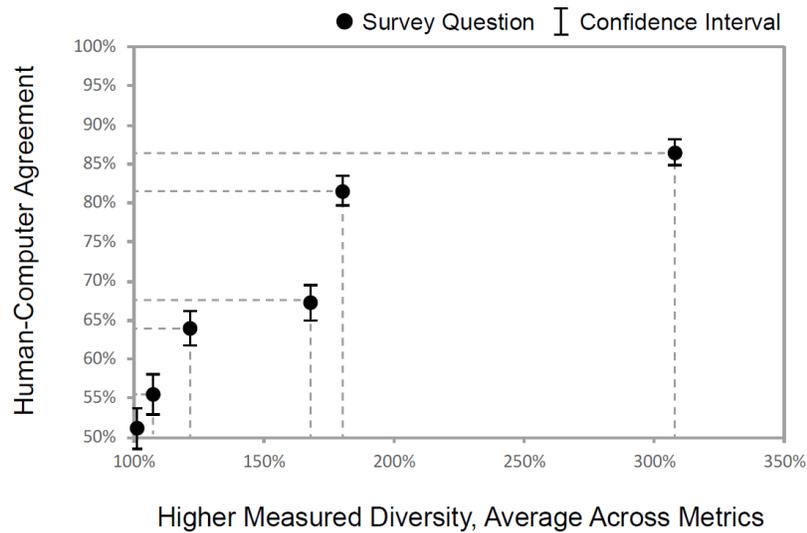


Figure 5.6: The results of the human-computer agreement survey questions, showing the magnitude of the difference in measured diversity versus the percentage of respondents that agreed with all computer metrics

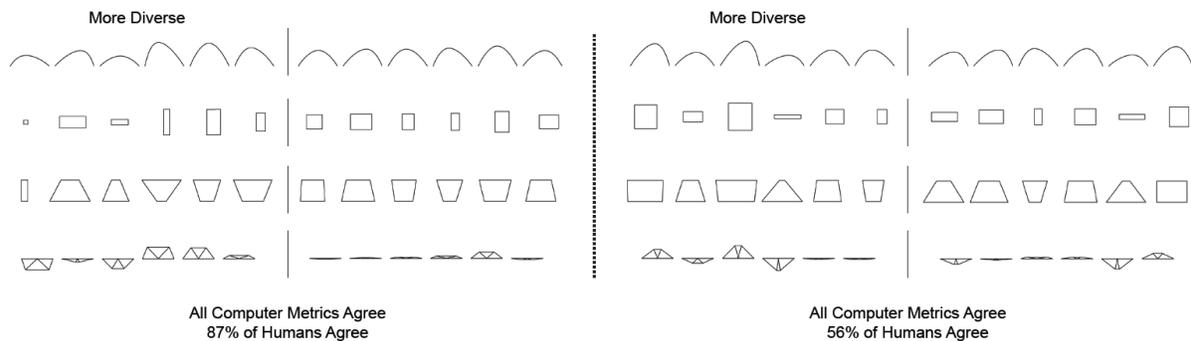


Figure 5.7: An example illustrating the difference between two sets of designs that are easy for humans to select the most diverse set, and two sets that are more difficult to distinguish

5.4.2 Differences between metrics

To determine the overall best metric as judged by human participants, the results of head-to-head matchups between the metrics were recorded, and the overall winning percentage for each metric was calculated. These results are provided in Table 5.1, where each entry in the matrix represents the winning percentage of the row's metric over the column's metric. The overall winning percentage is given as both a simple average and a weighted average based on the number of matchups conducted for each entry. However, due to the nature of the survey questions and their ability to make specific distinctions, metrics could face off against a relatively strong or weak slate. To compensate for these differences, an adjusted score was also calculated using the Colley matrix ranking system (Colley 2002). This system, which was developed for ranking college football teams but is increasingly used for academic research such as Häggman et al. (2015), adjusts the winning percentage for each metric based on its strength of schedule, or how many times it faced off against strong metrics versus weak metrics. The adjusted Colley rankings are also given in Table 5.1, and produced similar results to the weighted average.

Table 5.1: The final results of the faceoff question analysis, showing the winning percentage of each diversity metric, the raw and weighted overall average winning percentage, and the Colley Matrix winning percentage

Winning Metric	Distance to Outlier (A)	Sparseness Centroid (B)	Sparseness Median (C)	Convex Hull (D)	Enclosing Ball (E)	Winning Percentage (raw)	Winning Percentage (weighted)	Colley Matrix Rating
A	-	0.505	0.457	0.437	0.430	0.457	0.460	0.480
B	0.496	-	0.220	0.325	0.477	0.379	0.429	0.436
C	0.543	0.780	-	0.378	0.538	0.559	0.534	0.526
D	0.563	0.675	0.623	-	0.582	0.610	0.593	0.572
E	0.570	0.523	0.463	0.418	-	0.493	0.468	0.486

As demonstrated by this survey, the convex hull metric performed the best, followed by the sparseness at the median, enclosing hypersphere, distance to outlier, and sparseness at the centroid. Since the two sparseness metrics are so closely related, they had the fewest numbers of face-to-face matchups. The median metric performed surprisingly well in these matchups, which significantly impacted the overall average winning percentages. Yet despite the weighted average and Colley matrix method adjusting for this situation, the ranking among all of the metrics remains the same. However, it is important to interpret these results in light of the human-computer agreement questions. No metric proved to be by far the best choice, as the highest adjusted winning percentage was only 57%. Furthermore, the distinctions between metrics in many of these questions seem to be below the threshold of what humans can accurately and

repeatedly perceive. As such, any of these metrics could likely be used in applications where the magnitude of diversity difference is high enough to be meaningful for a designer.

5.4.3 Additional factors in the study: visually dominant variables

Various factors that are not necessarily related to numerical design vectors can affect human perception of visual diversity. In the case of parametric design, the specific parameterization can determine to a substantial degree which variables are most noticeable when adjusted. Some global variables can have a large visual impact, while some less consequential variables could be barely perceptible. A designer could intelligently correct for these differences by scaling the different variables in the diversity calculation. However, for this study, the variables are weighted equality. The four parameterizations developed were designed such that two had one visually dominant variable (the height of the arch versus its skew, and the depth of the truss versus its outer chord width), while the other two had equally notable variable changes due to symmetry (the dimensions of the boxes and trapezoids). It was theorized that the survey results might be different for these two categories, as participants being shown arches or trusses might neglect the second variable when considering diversity. To test this conjecture, the raw question responses for the survey are presented in Figure 5.8, separated by shape. As can be observed by comparing each shape's response to the overall average, the study did not uncover any systematic effect between the two different types of parameterizations.

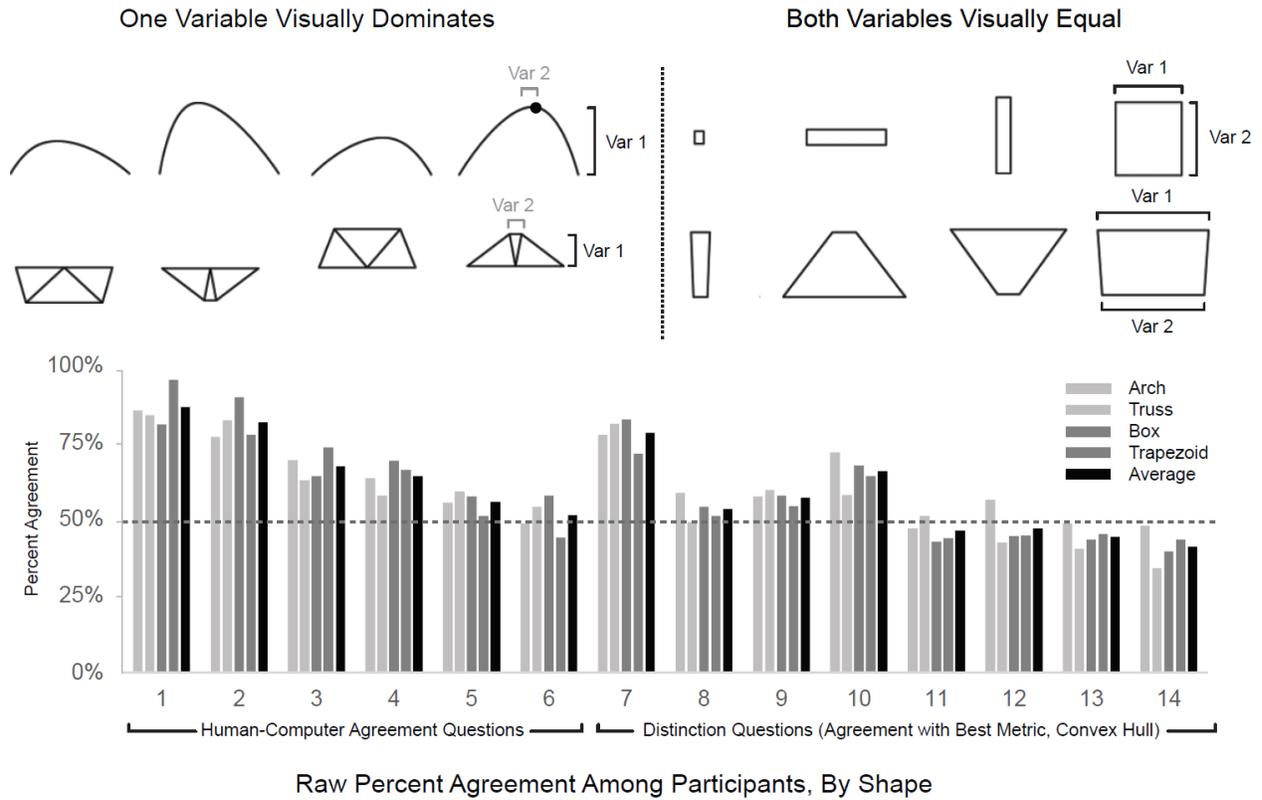


Figure 5.8: A visual description of the different parameterizations for one-variable dominant and regular shapes, along with the shape-separated raw question results, showing no distinct pattern. A value above 50% in the distinction questions indicates agreement with the convex hull metric.

5.4.4 Individual preference strengths for each metric

Although the average and Colley matrix winning percentages yield the most complete picture of the relative performance between different metrics, these averages could hide strong individual preferences for different types of people. During the study, a participant might have developed a rule, such as finding the biggest visual outlier or punishing designs that appear to be duplicates, which may have led to a noticeable preference for a particular metric. Both the strength of these preferences, as well as the total number of people who preferred a particular metric, are potentially useful results. Figure 5.9 presents histograms of the top overall raw winning percentage for each individual participant, as well as a breakdown of proposed metrics. A high overall winning percentage shows that a participant agreed with a specific metric on most distinction questions, while a lower score indicates no strong preference. Thus, a histogram with many participants showing strong preferences, indicated by high bars closer to 100%, would show that a metric is both popular and strongly preferred. In addition to having the best overall winning percentage, the convex

hull metric was also the most popular preference, at just over a third of the participants, and had higher preference strengths in general.

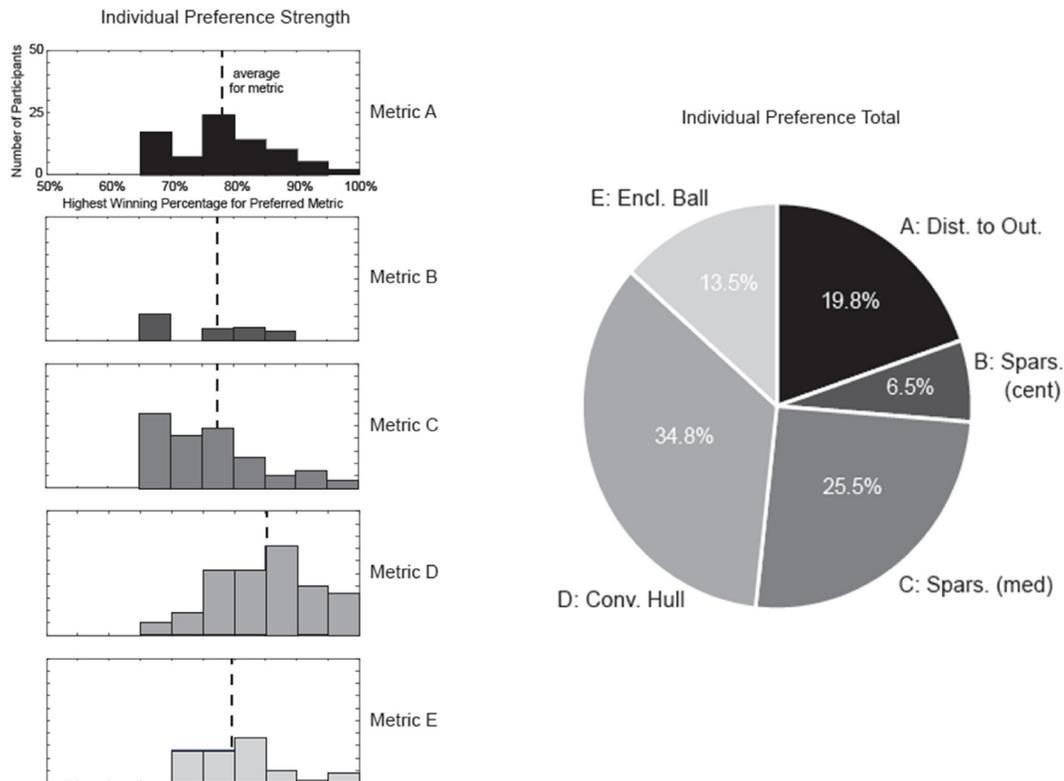


Figure 5.9: The strength and totals for individual preferences of participants in the study, as indicated by their overall best winning percentages

5.4.5 Limitations

It is worth noting potential limitations to the study. First, the participants were not expert designers, but were rather part of the general population. It is possible that experts would be more perceptive of visual diversity. These design spaces are also of a low complexity, both in terms of a limited number of variables and examples, as well as their simple representation. This simplicity has advantages, since all of the metrics can be easily visualized and understood, and the designs can be evaluated quickly. As mentioned earlier, there is some research to suggest that with higher numbers of choices or more subtle differences, the effect of higher diversity on design decision-making might even be negative. Nevertheless, a disadvantage of the study is some doubt about how these concepts can translate to more complicated designs, even if they are demonstrably useful for very early-stage design exploration. The order within the sets was also ignored during the study, although the order of the sets themselves, the questions, and the shapes were all

randomized. Finally, in order for the metrics to disagree in the distinction section, the diversity of the sets had to be fairly close, in the range that the human-computer agreement questions generally showed to be difficult for humans to accurately and repeatedly distinguish. Yet when taken as a whole, these results still provide insight into different methods for quantifying diversity in a computational design setting.

5.5 Discussion and conclusions

5.5.1 Qualitative comparison of metrics

In addition to the results of the study, there are clear advantages and disadvantages to some of the diversity metrics. For example, duplicates do not affect the geometric bounding techniques in the same way that they influence the metrics that measure distance for each point individually. Each metric is sensitive to extreme outliers in the set, although some may be affected more than others. Geometric measurements like the convex hull and enclosing ball are dependent on the specific algorithm used to approach the problem, and may be less accurate or more computationally intensive, especially compared to the sparseness measurements. Of the tested metrics, the convex hull calculation also requires a minimum number of unique points, which can be problematic for design spaces of higher dimensions, a problem that is highlighted in the example shown in Section 5.2. Especially in these situations, the convex hull volume can also yield results on different orders of magnitude, which may require its own special normalization. The methods tested in this chapter are deterministic, but for other methods that include randomness, such as categorization by clustering, accuracy and reproducibility may be an issue. Each of these properties may be worth considering when selecting a diversity metric for measurement depending on its specific use.

5.5.2 Example application

To demonstrate potential applications of this research, diversity metrics were applied to a tower study, illustrated in Figure 5.10, that is more complex than the 2D shapes but is still of relatively low dimension. The parametric definition of the high-rise tower contains ten variables related to the floorplate shape, size, and twist of the building. For this example, it is assumed that the designers and developers have a target building volume, roughly corresponding to available square footage, but desire a signature, atypical geometric architectural solution for their site. First, an evolutionary solver was applied to the parametric design space to find a set of ten potential designs that are near to the desired volume. The entire design space was then randomly sampled at a resolution of 1,000 designs, and the ten designs closest to the target volume were chosen as possibilities. Finally, the 50 designs closest to the target volume were attempted in 10 different combinations.

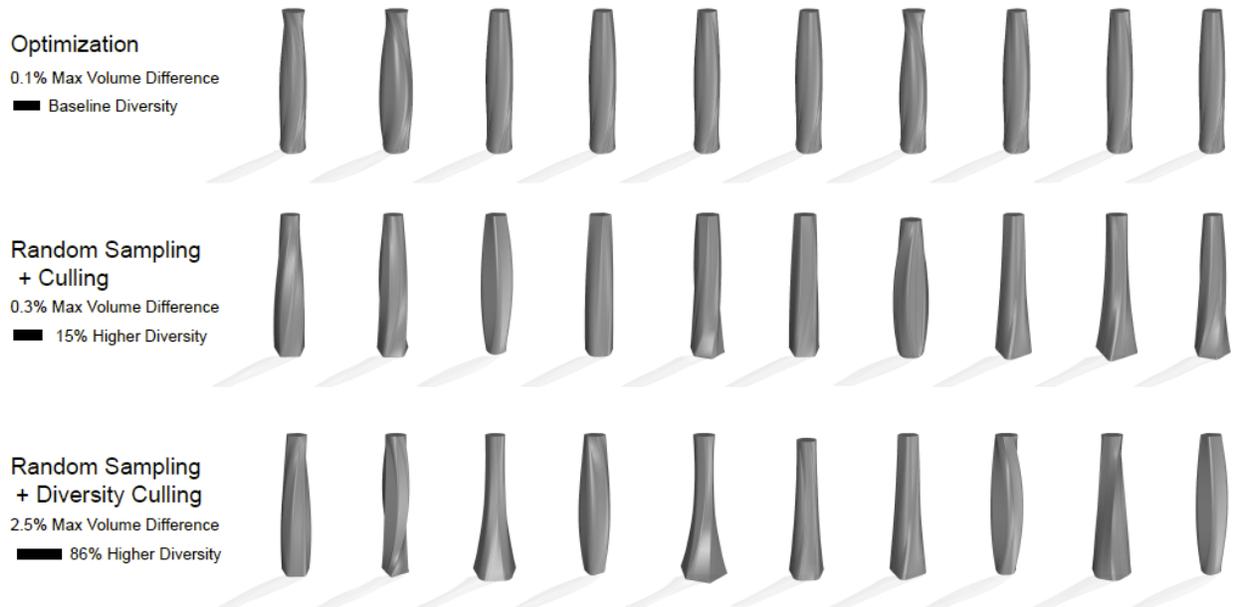


Figure 5.10: The results of an example parametric tower study, showing the novelty of a design set generated using a process based on diversity metrics in this chapter

The diversity of these sets was measured using each metric in this chapter except the convex hull, which requires more points to run for a problem of this dimension. The process then returned the most diverse set returned during the exercise. This crude diversity forcing technique does not require any additional performance evaluations from the sampling approach, but only some extra computation for the diversity calculations. As can be observed visually, the process driven by diversity metrics sacrifices some in terms of performance, but exhibits considerably more variety and expressiveness than optimization or random sampling alone. Although this example is a rough approximation of more sophisticated algorithms, it nevertheless demonstrates the creative potential of diversity-based computational design properties.

5.5.3 Contributions towards diversity metrics for human-centric design

In summary, this chapter has consolidated literature from different areas and extracted techniques that could be useful for designers in architecture, engineering, and related fields looking to consider performance and visual appearance simultaneously. It has demonstrated agreement between visual perception and digital calculation of diversity for sets of simple geometries that are often considered during conceptual design, although this agreement becomes unreliable when differences in diversity are not large enough. It also offers both a quantitative and qualitative comparison of five different metrics, which have various strengths and weaknesses, but all appear to be acceptable at the magnitudes in relative diversity that concern human designers. Finally, this chapter has provided ideas and examples for researchers looking to integrate considerations of visual diversity and novelty into interactive, computational design processes.

5.5.4 Future work and concluding remarks

Further research into the topic of quantifying diversity could take a number of paths. First, the concepts in this chapter could be tested on more complicated designs, which could also make some metrics mentioned but not explored in this chapter more relevant. Yet it is again worth noting the general difficulty in using distance-based metrics in higher dimensions (Aggarwal et al. 2001). As such, these metrics are mostly applicable to problems with a small number of variables (<15-20, often even smaller). However, problems of this size are very common in the conceptual design of buildings and other structures, where variables control global parameters such as building massing (Marks, 1997; Caldas & Norford, 2003), envelope properties (Coley & Schukat 2002), and structural geometry (Mueller & Ochsendorf, 2015). Thus, architects and engineers working on such problems can use the diversity metrics presented here immediately.

The diversity metrics in this chapter are also restricted to parametric design, which may be a common computational approach, but it is not the only one. Although it is difficult to translate pure geometrical measurement metrics to rule-based or other methods of defining and generating design options, future research could be made into diversity measurements that are more suitable to design definitions that do not contain numerical design vectors. In general, many digital design processes focused on diversity and performance are also limited by problems associated with data visualization for higher dimension design spaces, and/or complex designs. Additional research into new, interactive modes of data visualization could influence how designers understand and interact with diversity, visual or otherwise.

Beyond quantification of parametric design diversity, though, lies the question of how designers actually interact with design processes that are based on computational measurements of diversity. Future studies should consider how designers engage with diversity-based conceptual design, and how this affects design outcomes. These design studies would involve specific architectural problems, of relatively low dimension, that are common during early exploration. Such studies may reveal nuances in how expert designers most effectively use computational brainstorming assistants. Nevertheless, design studies of this nature require a starting point for understanding and measuring computational diversity, so that it can be implemented in various algorithms or workflows, which can subsequently be tested. A combination of the metrics proposed in this chapter are adequate for such design studies, as demonstrated by the ability of computers to increasingly match human perception of visual diversity at the magnitude of design differences that are noticeable for humans. Current designers can also integrate diversity-based techniques into their processes right away, and explore what value these approaches bring to their practice. As such, this chapter provides researchers and designers interested in creative, novel design with an initial basis for measuring diversity in a computational setting.

6 Implementing Data-Driven Parametric Design with a Flexible Toolbox Approach

6.1 Introduction

A central motivation of this dissertation is increasing accessibility of design tools, and efforts in this area can multiply the usage of data-driven approaches and their subsequent impact in practice. Ideally, such tools should be flexible in terms of how easily they can be integrated into existing design approaches. Certain technologically savvy designers already generate design space catalogs, conduct architectural optimization, and even implement surrogate modeling and other techniques in their workflows. Those designers who are comfortable with coding have found considerable support through open-source libraries and the ability to access them through integrated development environments or other scripting methods. The particular language or libraries may change as computational environments are further developed, but there will always be a segment of the design community that prefers to work on the cutting edge, often manipulating code in a raw form. There have also been substantial efforts in both academia and practice to educate architects and engineers and increase their capacity for computational design through visual programming and coding, which in the long term is likely to have a substantial effect on how buildings are designed.

This democratization and customization of computational design approaches has many benefits, and has created a constantly evolving and improving shared framework for performance-based parametric design. On the other hand, it is clear that many practicing architects and engineers do not have the background,

interest, or time to become full-time software developers themselves, especially in the near future. Many prefer to spend their time concentrating on creating new designs, rather than improving computational workflows. These designers, both professional and academic, will always depend on existing software to suit their needs. Yet the risk for them is that existing software may not be exactly what is needed, and it may exert more of an influence on their process than they desire, down to the particulars of a digital interface. For these designers, it is worthwhile to find a happy medium that balances flexibility and accessibility.

This chapter proposes a set of computational design tools that seek to achieve such a balance by bringing data-driven approaches into common parametric environments. Like other recent research contributions, it acknowledges that visual programming is becoming a viable medium for widespread parametric design exploration, and it has a greater level accessibility for many practicing designers than raw code. Such data-driven approaches can also take advantage of the considerable effort within the design community to directly connect geometry and simulation through this shared parametric environment. This chapter reveals how individual components with specific functions are combined to pursue data-driven design strategies and generate the results of the case studies in this dissertation. It then presents an integrated design example tracing possible progressions through design space formulation, diversity-based brainstorming, and interactive optimization.

The culminating design example described here shows, in practical terms, the advantages and possibilities of using these tools, as well as the limitations and complications that can come about when applying them on particular design problems. While these contributions are primarily about implementation, they required intentional consideration of a democratized design process, in which designers are provided supportive functionality yet still have freedom in connecting particular data techniques within a parametric logic.

6.2 Workflow methodology

6.2.1 Design Space Exploration overview

This section first describes the individual components developed as part of the Design Space Exploration (DSE) toolkit, which is the practical backbone of this dissertation. It then outlines selected workflows and corresponding modes of interacting with parametric design that are enabled by the tools. This list of workflows is not exhaustive, and others have used components in the DSE toolbox on a variety of applications, for both design projects and published research (Nagpal et al. 2019). Since the intention is to create a flexible mixture of data science and optimization functions to allow designers to create their own approaches, it is anticipated that the tools will not always be used in the ways described in this chapter. However, the listed workflows include the most natural and obvious configurations for using the toolkit

components. These particular arrangements are demonstrated on the case studies throughout this dissertation.

Design Space Exploration was developed in collaboration with students and researchers affiliated with the Digital Structures Research Group in the MIT Department of Architecture. It is an open-source plug-in for Grasshopper and consists of components for creating design space catalogs and conducting machine learning, optimization, and design space organization. Some of the code, including significant portions related to surrogate modeling, was developed as part of Mueller (2014). However, the other workflows listed here were proposed and designed by the author, with some development support from collaborators and research assistants in architecture, computer science, and civil engineering⁴.

DSE is not a simulation engine itself for predicting and understanding building performance, such as EnergyPlus or structural finite element modelers. Rather, it is designed to connect to any numerical design variables, geometry, and corresponding simulations, as described in Figure 6.1. Various DSE components rely on external libraries, including Accord.NET (Souza 2014), Math.NET (Ruegg et al. 2009), and JMetal (Durillo & Nebro 2011). Design Space Exploration is freely available online for users of Grasshopper and could be developed for other parametric software in the future. While the author of this dissertation is not the sole software contributor to the entire toolkit, many of the components were designed explicitly to enable and test the workflows and case studies described in the previous three chapters. The toolkit is in constant development, but versions of it have been downloaded over 2,000 times at the time this dissertation was completed.

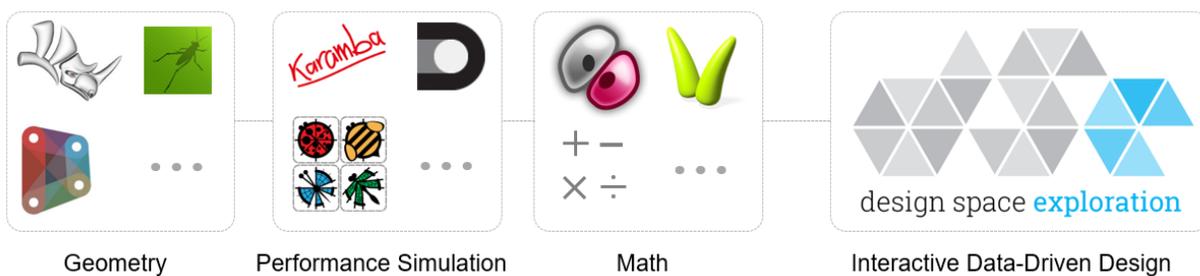


Figure 6.1: Diagram of the relationships between parametric components and plug-ins. Design space exploration is a plug-in for data-driven design in Grasshopper, which works with other simulation engines and parametric components native to the Grasshopper software.

The specific components and functionality that are available as part of DSE are described in Appendix A.

⁴ Contributors to DSE include Caitlin Mueller, Renaud Danhaive, Jonathas Felipe, Stavros Tseranidis, Anthony McHugh, Violetta Jusiega, and Alicia Nimrick

6.2.2 Possible workflows with DSE

These separate DSE components developed for visual programming can operate on a parametric design space in a variety of ways. As described throughout this dissertation, the standard method for interacting with a performance-based parametric model is to modify the sliders, view the geometry, run a simulation, and then see the results, as in Figure 6.2. Depending on which aspect of building performance is being evaluated, this combination may include a manual step to actually run the simulation. For very rapid evaluations, the simulation and corresponding visualization may be completed automatically and update every time the slider is adjusted. This base relationship between variables, geometry, simulation, and output remain the fundamental building block of all other workflows.

1 | Standard Parametric Design

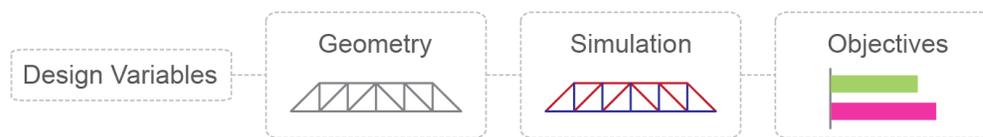


Figure 6.2: A typical parametric workflow connected to performance simulation

A catalog approach to parametric design repeats the basic structure by automatically generating multiple designs and simulating their geometry in sequence, as in Figure 6.3. While conducting this automatic generation, decisions must be made about the method and resolution of the sampling in the design space. The most straightforward sampling technique is grid sampling, in which designs are selected at even increments across the range of each design variable. While this functionality is included in a few iterative sampling tools for Grasshopper, a designer often desires randomized possibilities or more specific control over how many designs are being considered for practical, creative, or data quality reasons. DSE separates the selection of sampling type from its more general iterator, which provides increased control over the sampling method and resolution. At the same time, this separation allows the iterator to operate on any list of designs recorded in the same format—for example, to take screenshots of every design along a Pareto front as found by a multi-objective optimization algorithm, as described in Figure 6.4.

2 | Design Catalog

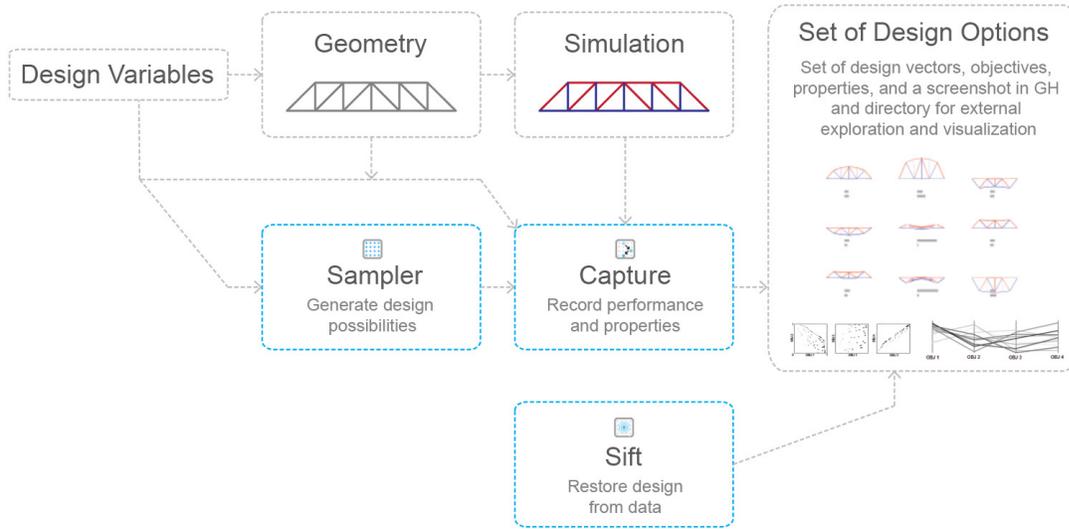


Figure 6.3: Workflow for generating a design catalog of options in parametric software

3 | Multi-Objective Optimization

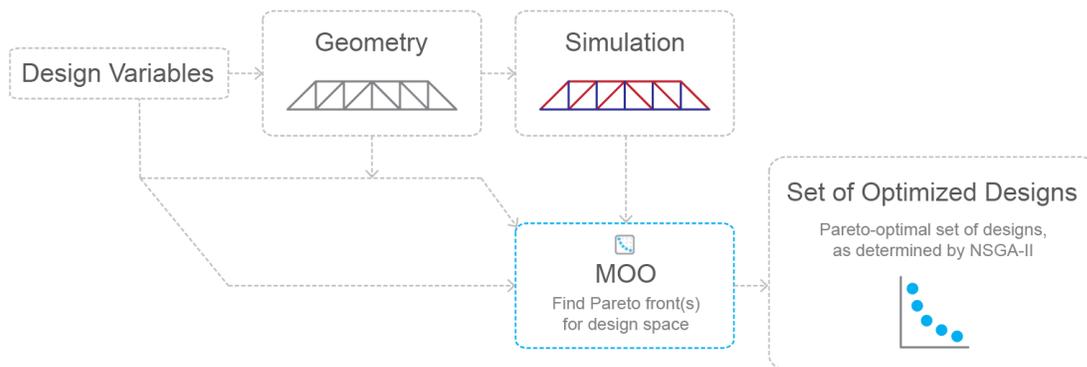


Figure 6.4: Workflow for parametric multi-objective optimization

Other workflows build on the basics of design space exploration and catalog creation by enabling the supplemental methods described in this dissertation. As an example, Figure 6.5 shows how one might use a diversity filter to assist with meaningful brainstorming within parametric software. In this example workflow, the designer can use a dataset generated either through sampling, the history of an optimization run, or any other means. This dataset may contain hundreds or thousands of designs, which are not all worth considering individually, especially in early design. The diversity component first allows the users to

select target performance objectives and filter out all designs that are not within a specified range of the target. Then, it asks the designer for a number of representative designs he or she would like to consider, before using the diversity measurements in Chapter 5 to generate a highly diverse, curated set of samples for consideration.

4 | Design with Diversity Filter

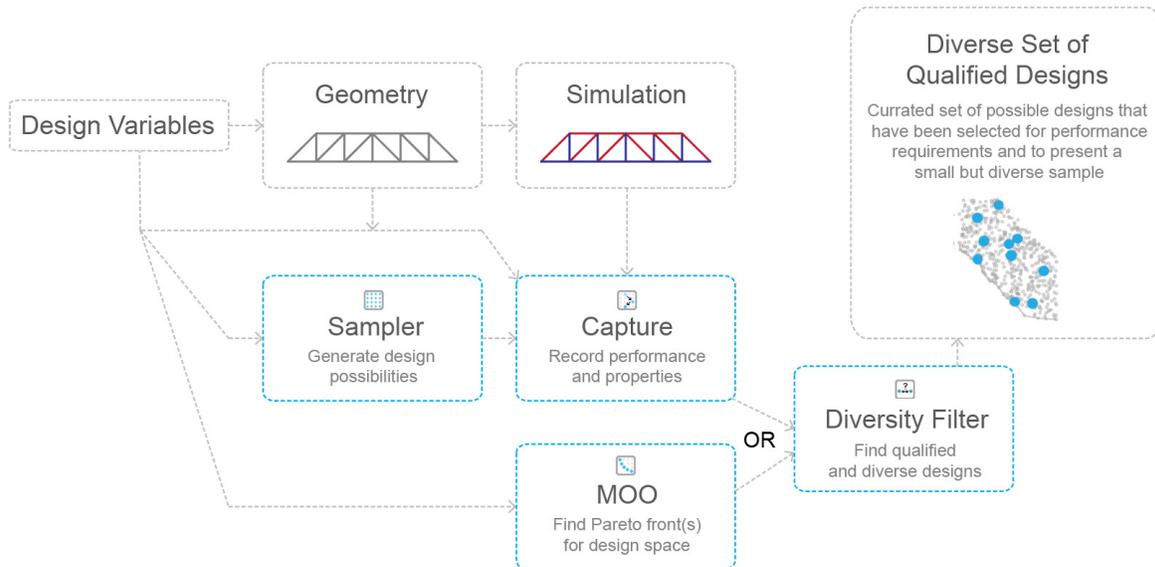


Figure 6.5: Workflow that includes a post-processing component for returning a more directed yet diverse sampling of the design space

The next possible workflow saves time, preserves original parametric relationships, and prevents clutter on a visual scripting canvas while transforming design variables as described in Chapter 3. For this workflow, a DSE component reads in sliders, a design space scale, and a set of coefficients, which may be calculated within or outside Grasshopper (see Figure 6.6). These coefficients can be numbers, which allows for linear mapping of the original variables, but can also be dynamic and depend on the script itself. Once these coefficients are established, a user can create separate synthetic variable sliders to control the design, which override Grasshopper’s main solution structure and adjust the original sliders that are still connected to the geometry and simulation. Since new sliders do not have to be reconnected each time, designers can use this workflow to rapidly cycle through synthetic variables that are discovered using different data science techniques.

5 | Variable Transformation

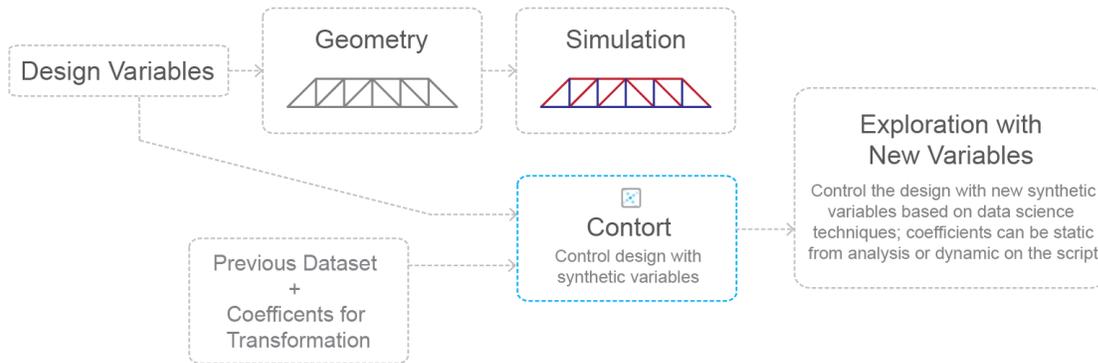


Figure 6.6: Workflow for initial experimentation with synthetic variables in parametric design

The remaining workflows described here feature direct applications of classification and supervised learning for design space exploration. In the first one, shown in Figure 6.7, the designer can use a DSE component to get an initial sense of the design space and how it behaves, and also find broader trends that might inform overall decisions. First, the calculation of effects uses only a few simulations to discover which variables influence the performance functions most severely, and by roughly how much. Although the calculation of effects is dependent on two or three discrete settings in the design space, which will not capture complex objective function behavior, it can be an effective first pass method. Designers can then use the clustering component on a previous dataset to find groups or families in the design space, which can then be viewed in order of their average performance. Although a dataset for clustering should be intentionally biased in some way, since uniform sampling will not lead to meaningful groups, a biased design set can be created through optimization, taking a “slice” of the design space at an isoperformance contour, or through other means. The clustering component can also automatically reset the original sliders to the bounds of each cluster, with an affinity rating provided by the user, such that designers can explore in a restricted area of the design space carved out by a single cluster.

6 | Early Design Space Analysis

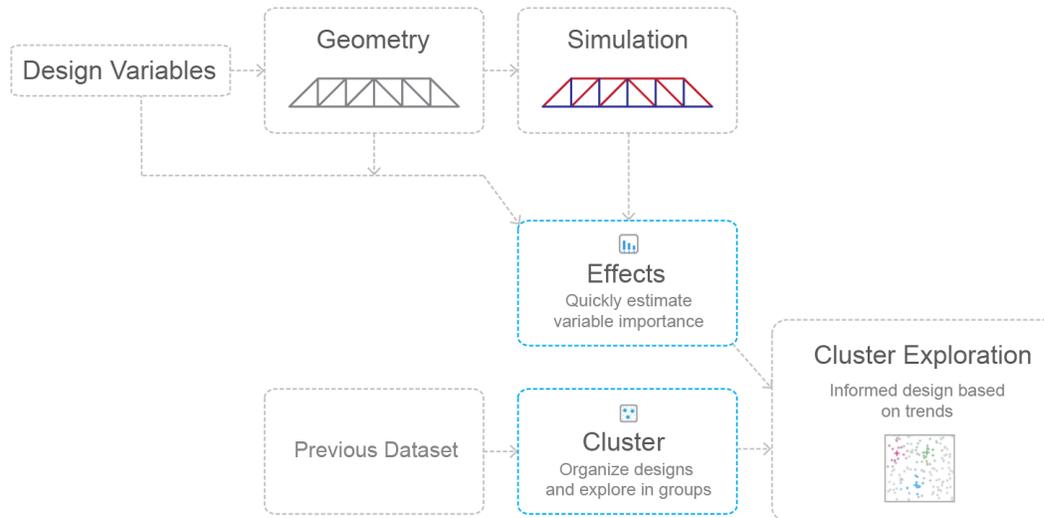


Figure 6.7: Workflow with data-driven techniques for analyzing the design space

Figure 6.8 describes how to use a surrogate modeling component to train a predictive model of objective performance based on an existing dataset. Once the model has been appropriately trained, the simulation can be turned off, and the geometry can be manipulated with only the surrogate model information showing. While the model takes some time to train, the predictions are essentially real-time, which allows for a flexible model with direct performance feedback. In most cases, the geometry generation and rendering are considerably slower than the prediction components, which means that the DSE components do not meaningfully slow down most parametric models. In Figure 6.9, these live predictions have been plugged into a general interface for gradient-based interactive optimization, as described in Chapter 3. The interactive optimization component of DSE offers the most engaging separate interface for moving in both the design and objective spaces. However, it will most likely be used after or in conjunction with other components in the tool suite in order to be most effective. Examples of many of these workflows are demonstrated next.

The primary intellectual contribution of this section is the mapping of component relationships and how they should be structured to enable effective integration of interactive optimization and data science in parametric design. These components and collective workflows extend the individual theoretical proposals of the previous three chapters by bringing them together in a coherent yet flexible framework. In combination with the conceptual contributions, the implementation and testing of these tools on comprehensive examples and human designers yields significantly more insight into data-driven early design processes than theory alone.

7 | Design with Live Feedback

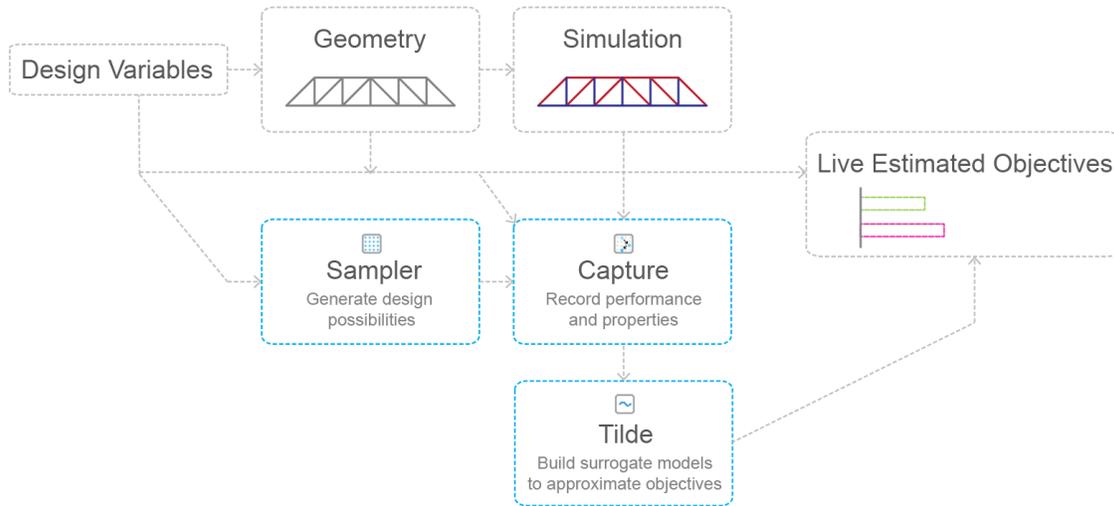


Figure 6.8: Workflow for building a predictive surrogate model to provide live objective function estimation

8 | Interactive Optimization

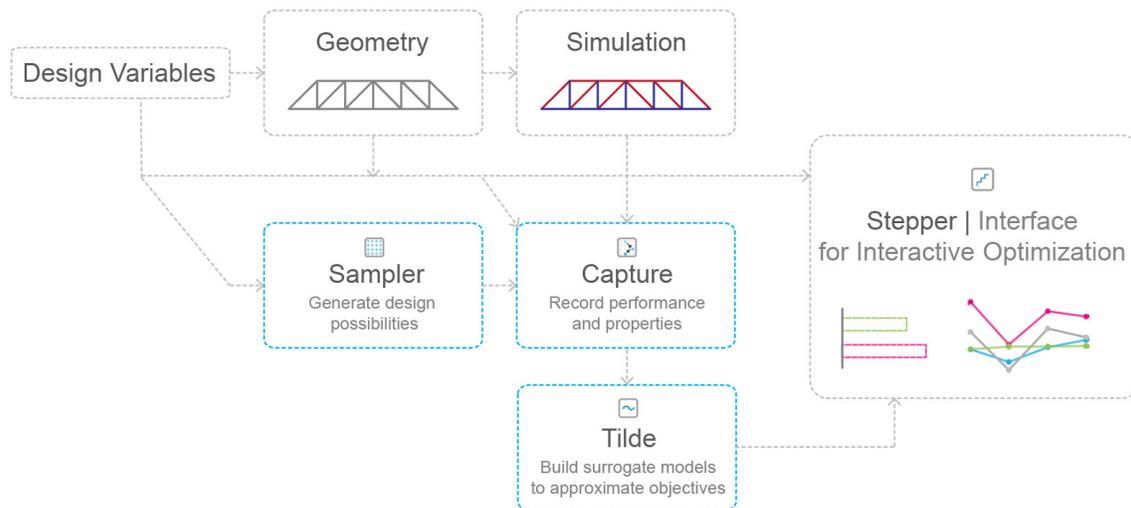


Figure 6.9: Similar to the workflow above, but combined with a tool for interactive optimization that allows for direct movement in both the design and objective spaces

6.3 Example of conceptual grid shell roof

This section gives an example of how the DSE components could be used in sequence to pursue parametric design in accordance with the goals of this dissertation. It begins by analyzing and modifying a parametric design space during formulation, then uses a diversity filter to select specific directions for further exploration, and finally demonstrates how interactive optimization can be used locally on a selected working design concept. There are some parallel aspects or repeated steps in the process for demonstrative purposes, as a designer would not necessarily need to consider every individual approach simultaneously. However, this example design shows how the flexible approach to tool and interface development offers designers a buffet of data-driven methods that can stimulate creativity and support design ideation throughout a computational process.

The selected case study is the design of an athletic center for a campus environment in Boston, MA. For the case study, it is assumed that the design team decided on a hybrid structural system involving a curved grid shell roof, which can be supported on large external columns, as well as directly on the ground. The grid shell has four main sides, although its main axis of curvature can be either N-S or E-W within the design space. When the edge of the grid shell is lifted off of the ground, the resulting gap can be filled with a mixture of opaque wall and transparent glazing. Thus, in many configurations, the primary structural action shifts from compressive arching to spanning in bending between columns. As described in more detail later, structural models for this example account for bending by allowing members to grow deeper, which approximates the effect of adding global depth to the roof surface when required.

For many of the possible designs, the columns form clusters or tripods, which assist with lateral loads as well as gravity loads. The column clusters are especially important for variants with flatter roofs that are entirely column-supported, but less so for arching structures that transfer loads directly to the ground. During exploration, it is assumed that superfluous columns could be removed, since they occur often but are not turned off in the parametric script. Designers can also slide all columns to the corners, which negates their influence if the corner is supported. The overall massing can be explored by considering different boundary conditions, curvatures, and orientations within the original design concept. Due to its adjustable variables, the design space contains considerable freedom involving arches, cantilevers, double curvature, and column geometry. These massing and structural system decisions have both performance and visual implications.

In addition to the desire for a visually expressive structure, the case study has four numerical objectives. Two are related to structural performance (overall structural weight and weight normalized by covered area), and two are related to energy performance (annual energy consumption and energy use intensity for the enclosed portion). This example will primarily focus on the normalized objectives, which are weight of steel per area, referred to as structural material quantity (SMQ), and energy use intensity (EUI). The goal

of the case study is to achieve a high-performance design according to these four objectives, but in the context of a natural design process in which many other aspects of the design may influence decision-making and should be considered simultaneously. Since this grid shell example is also used in Chapter 7, more information about the modeling of these performance objectives can be found there.

A few distinct building types within this design space are shown in Figure 6.10, along with a comparable precedent for each geometry. These possibilities include cantilevered spanning structures, arches, vaults, and other variations of a typical long span design. Although the column configurations change the force flow for some of these examples, many behave in a similar manner to their precedents. As is true in the built environment, some of these designs have clear structural logic and will likely perform well in that domain, while others will not. By considering such a wide design space, it is possible to see how a designer might use the approaches in this dissertation at different scales for global brainstorming, local optimization within an already sound concept, or a combination of the two.

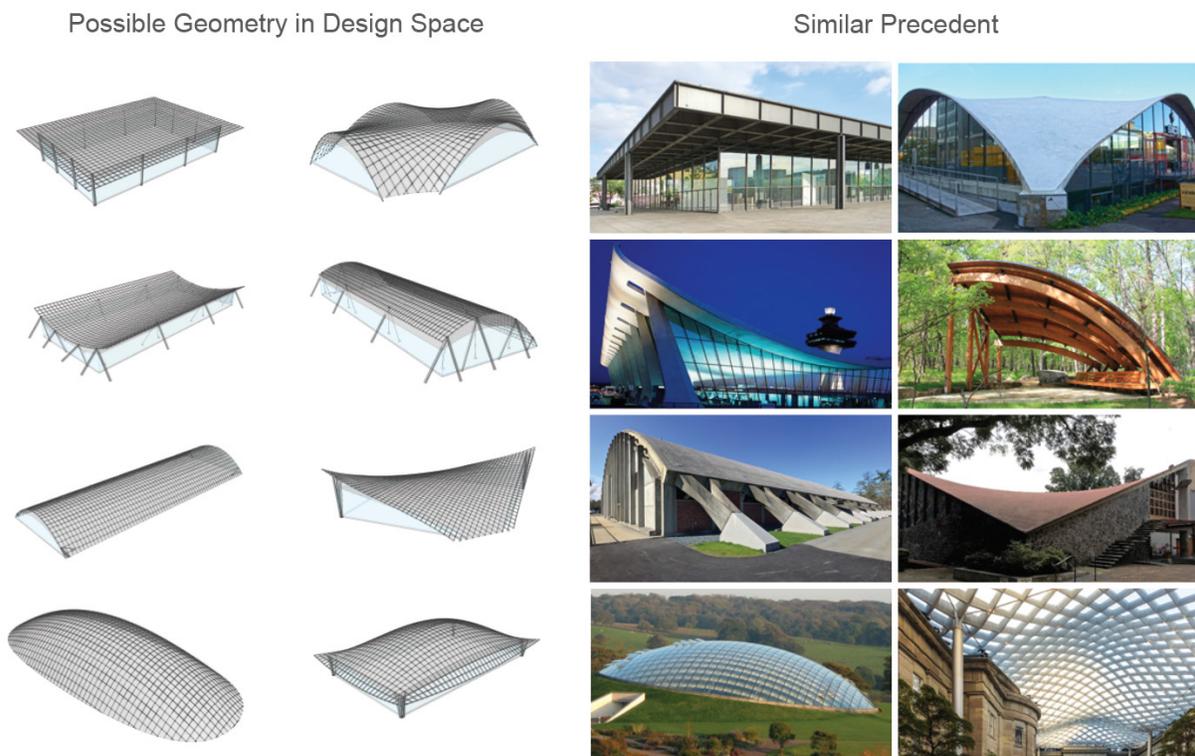


Figure 6.10: Potential options within the design space for this example, along with building precedents they resemble. Clockwise from top left are the Neue Nationalgalerie in Berlin (Dalbera 2013); SICLI Company Building in Switzerland (Janberg 2015); Tulip Tree Shelter in Bentonville, AR (Crystal Bridges 2013); El Altillo in Mexico City; Kogod Courtyard in D.C. (Acroterion 2012); Great Glasshouse in the National Botanic Garden of Wales (Ford & de Vere 2010); Thompson Area in Hanover, NH (Klack 2018); and Dulles International Airport in Virginia (Falkenpost 2016). Photos not credited are by the author.

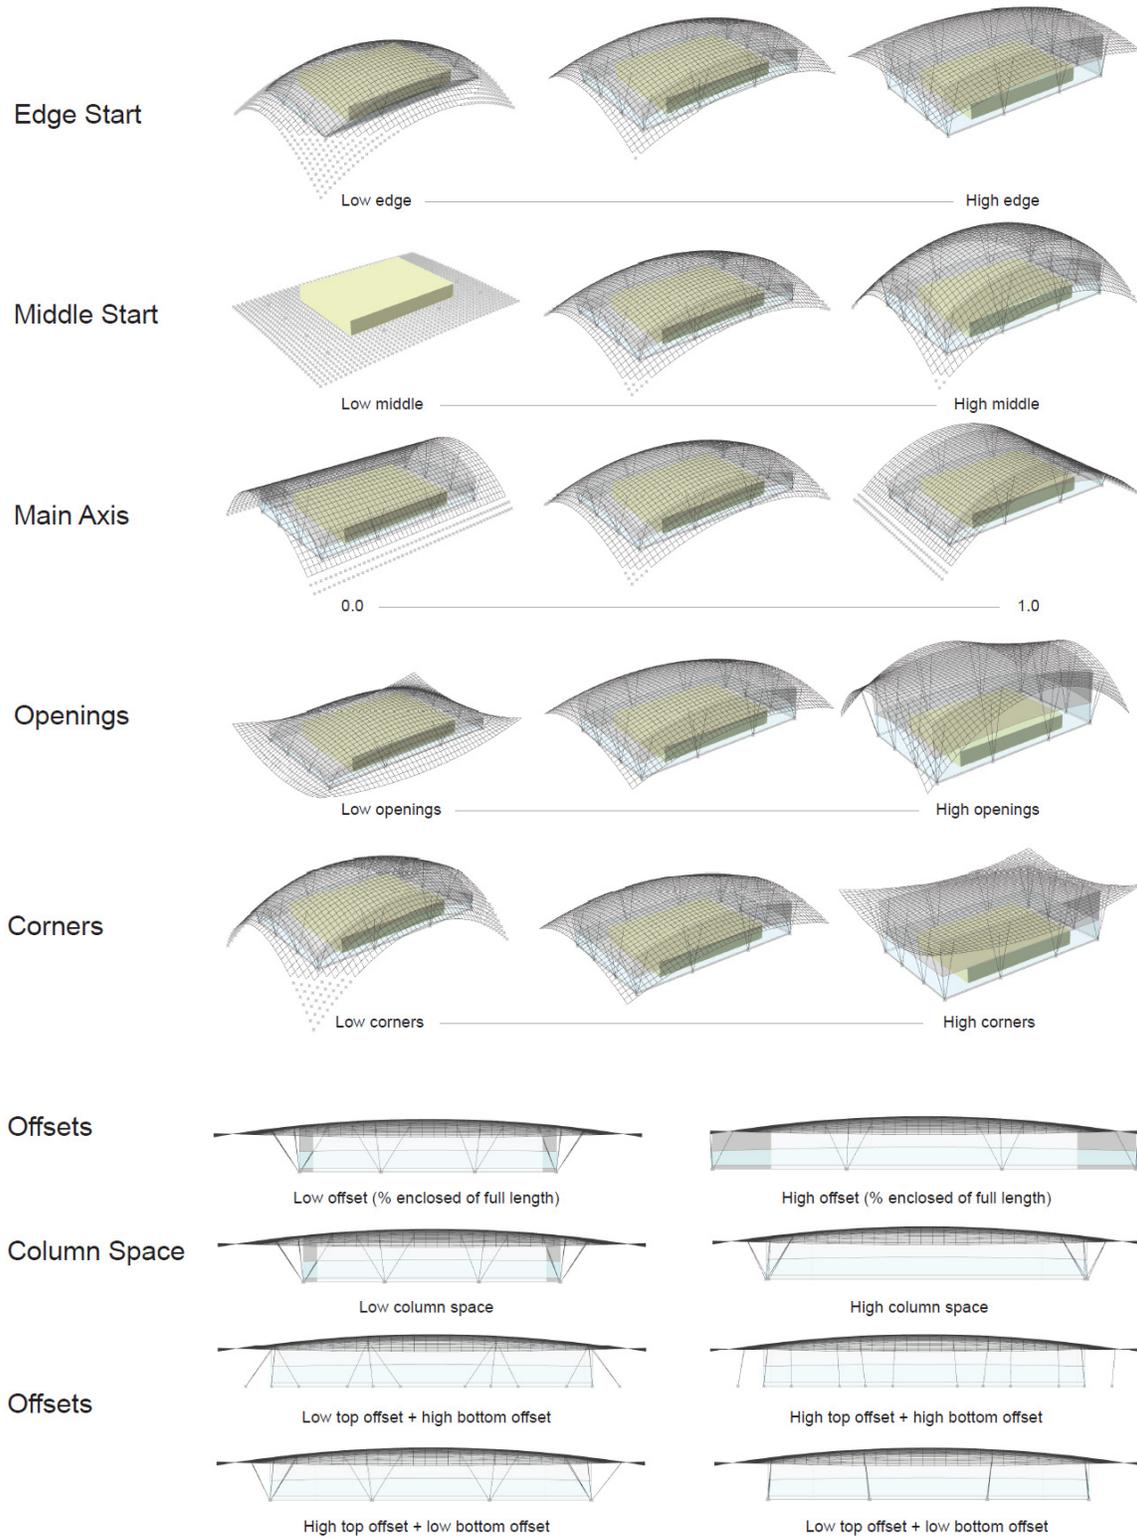


Figure 6.11: Geometric design variables for the example

Visualizations of each degree of geometric freedom are provided in Figure 6.11, and a full list of variables is given in Figure 7.1. Although the design variable settings show sweeps through the design space from an average starting point, this central design should not be thought of as a single, default design with potential variations. Rather, this example should be thought of as an entire design space, with a wide range of potential outcomes. The geometries within this design space are also not completed designs—especially in early brainstorming situations, it is assumed that designers would continue to further refine the resulting forms, both through small changes within the parametric model and larger, intuitive manipulations such as enforcing or relaxing symmetry, removing or shifting columns to more logical locations, or adding global depth to the single layer that was modelled to allow for crossing typologies. The structural and energy models are both approximate and contain considerable assumptions that permit parametric flexibility, but would be updated during further design refinement. Even though these flexible assumptions do not hold perfectly for every option, they enable quantitative and data-driven comparison and manipulation, which are meaningful in early design.

6.3.1 Design formulation

Initially, a designer might dive into the design space through direct slider manipulation, sampling, or optimization, in accordance with common methodologies described throughout this dissertation. While direct slider manipulation is often a natural first step, it does not systematically explore the design space. When deciding on the whether or not the parameterization is effective, prior to a more exhaustive sampling or optimization procedure, one useful approach is to gain a quick understanding of how the variables affect the problem in terms of performance. Intuitively, the designer might decide to begin this exploration by focusing on one objective first, due to experience, interest, or prioritization. In this case, structure is considered first, since there is a strong relationship between geometry and structural weight that will become clearer as the example progresses. A first pass method for calculating variable importance for structure, which has a faster simulation than energy, is demonstrated in Figure 6.12. This effects calculation was conducted using three levels at 0.25, 0.5, and 0.75 of each variable range.

Since there were more than 13 variables, which is a limit of the orthogonal array implemented in DSE, two separate calculations were conducted and then normalized using an effects calculation that included variables from both original groups. The results of this analysis indicate that column spacing, edge start, and overall size have a considerable influence over the structural performance. Column spacing mostly dictates the largest span for a given footprint, and it can remove intermediate columns for arching geometries, while the edge start generally prescribes the boundary condition along the outside of the roof. Conceptually, the overall size variables should control structural performance, as larger structures and corresponding spans require more area. However, designers must be careful to not assume too much from this analysis, as there is clearly noise in the data, even as rough knowledge of the main relationships can be useful.

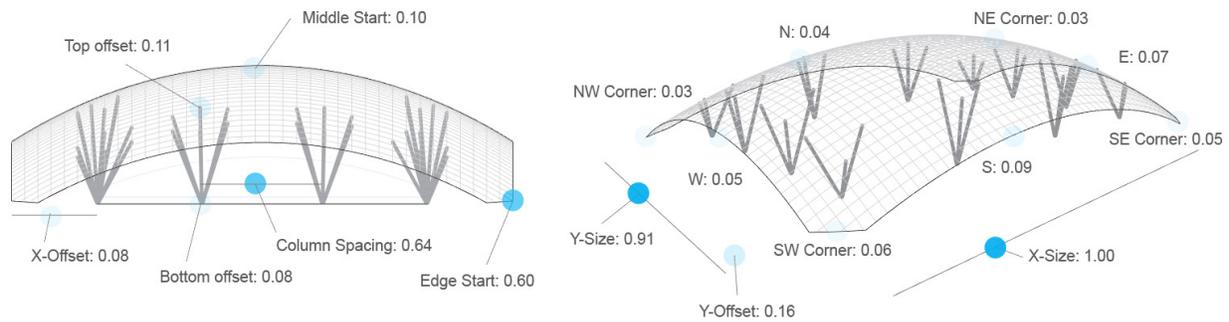


Figure 6.12: An initial estimation of variable importance to the problem by calculating variable effects

After understanding variable importance, the designer might seek creative solutions within the design space, but be willing to expend more computational energy to discover them. At this point, they could sample the design space at a resolution that fits the practical pressures on the design process. As described in the literature review, the designer could simply compare each of these options side by side in catalog form. However, this dissertation demonstrates new ways in which an initial dataset can be used to pursue more natural, interactive, and flexible approaches to design exploration. One particular method is to mine the existing dataset for patterns and find new ways of manipulating geometry that connect more directly with performance. Similar to Chapter 3, Figure 6.13 provides two example directions for morphing geometry that are meant to correlate with performance for both structure and energy. In the structural direction, very large spans and a relatively flat roof give way to a much smaller, curved roof that is supported at the corners. Along this continuum, the structure transitions from acting primarily in bending to behaving primarily in compression, which is more efficient. For energy, a tall, high-surface area design transitions towards larger, lower surface area structures before finding a design that is too low to be feasible for the programmatic requirements.

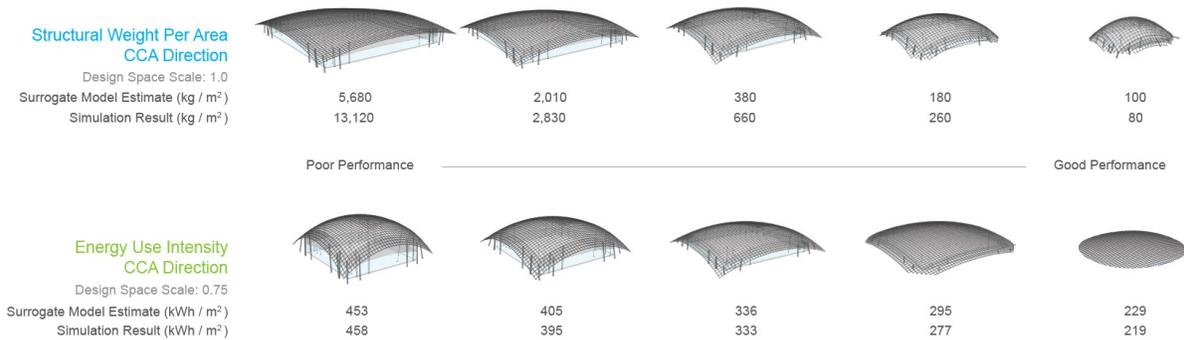


Figure 6.13: CCA variable directions based on SMQ and EUI for the design space

This generation of synthetic variables has a few potential benefits. First, the designer might want to use these directions during live exploration, mostly likely in conjunction with the original sliders to provide more flexibility. In this way, a designer can move to a particular area of the objective space and then explore locally. In addition, this slider essentially provides a composite visualization of which design variables matter, and how they should trend in order to improve performance. For the CCA structural variable, it is clear that better performing structures tend to have supported edges and more curvature. They also have more evenly spaced columns, and are smaller. While again this does not tell the whole story for complex objective functions, it can certainly provide additional information, and in some cases give greater control over the design. While this example is demonstrated for a parametric design in which there are already existing variables, this methodology becomes more advantageous when parts of the variable generation process are automated, as in the end of Chapter 3. In these future cases, computers will be considerably more active participants in the computational design process.

6.3.2 Diversity filtering

Next, or at another point in a parametric design process independent of these other forms of exploration, a designer might want to brainstorm and find a diverse range of possibilities, before further drilling down on one potential option. Working again from a generated dataset that contains either original or synthetic variables, a designer could achieve greater diversity while considering fewer results by using a diversity filter. In this example, the designer first starts with the large dataset that was employed for variable analysis and will eventually be used to build performance surrogate models. To get an initial starting point for further design, the designer first puts in a set of target numerical objectives. Rather than look at all designs in the dataset that meet this qualification simultaneously, which are often too numerous for a human to fully consider unless they are organized in a useful way, the designer can use a diversity filter to return 12 sufficiently different designs from the set. The goal is to produce a wide range of potential options for

further refinement, within a small enough group of designs to meaningfully consider each one. The outcomes of this section are thus not final solutions, but intermediate geometries meant to inspire creativity.

As an exploratory procedure, this dissertation considers progressively lower isoperformance levels for the design space, to understand how much diversity is sacrificed by moving towards better structural performance within the dataset. Figure 6.14 shows the number of qualified designs, diversity of a random sample within each performance level, and the diversity of a specifically filtered set for structural performance targets ranging from 80-600 kg / m². The number of qualified designs shows a downward trend with lower structural material quantities—for example, within this dataset, fewer designs can be found at 100 kg / m² than at 200 or 300 kg / m².

However, a similar trend is not clearly found for design diversity. To arrive at the diversity ranges in Figure 6.14, a random sample of 12 designs was first taken from within each qualified set 5 times, and the diversity of this culled set was measured using an average of the outlier and sparseness diversity metrics from Chapter 5. Next, the DSE diversity filter was used to find 5 sets of 12 qualified designs that have measurably higher diversity, using the same metrics. There is randomness in each procedure, hence the ranges and repeated sampling for research interest. Yet in each case, the diversity filter leads to more diverse designs than a random selection, which might be the default way of initiating a similar interactive workflow. Comparing across performance targets indicates that little diversity is sacrificed when moving towards 100-150 kg / m² from the poorer performing levels. Armed with this knowledge, a performance-conscious designer still in the creative brainstorming phase might begin by considering potential options at these lower ranges.

A visual comparison of example isoperformance sets from this exploratory analysis is provided in Figure 6.15. This image shows 12-design sets at three separate performance levels: 400, 200, and 150 kg / m², along with corresponding overall steel quantities, which helps to ensure structures of similar size. Following from the measured diversity, each of these sets provides noticeably dissimilar directions for further exploration and refinement. Level 3 contains considerable double curvature, but also has many designs supported primarily by branching columns, as well as long spans in both directions. The level 2 designs trend towards increasing curvature and often form compressive load paths the whole way to the ground, as many edges are entirely supported. In level 1, there are similar arching structures, but infeasible geometric solutions become more of a problem, which must be handled properly while sampling such a wide design space. While each of these levels satisfy the need for geometric diversity, it may be wisest to initiate further exploration at one of the lower performance levels, depending on the needs of a project.

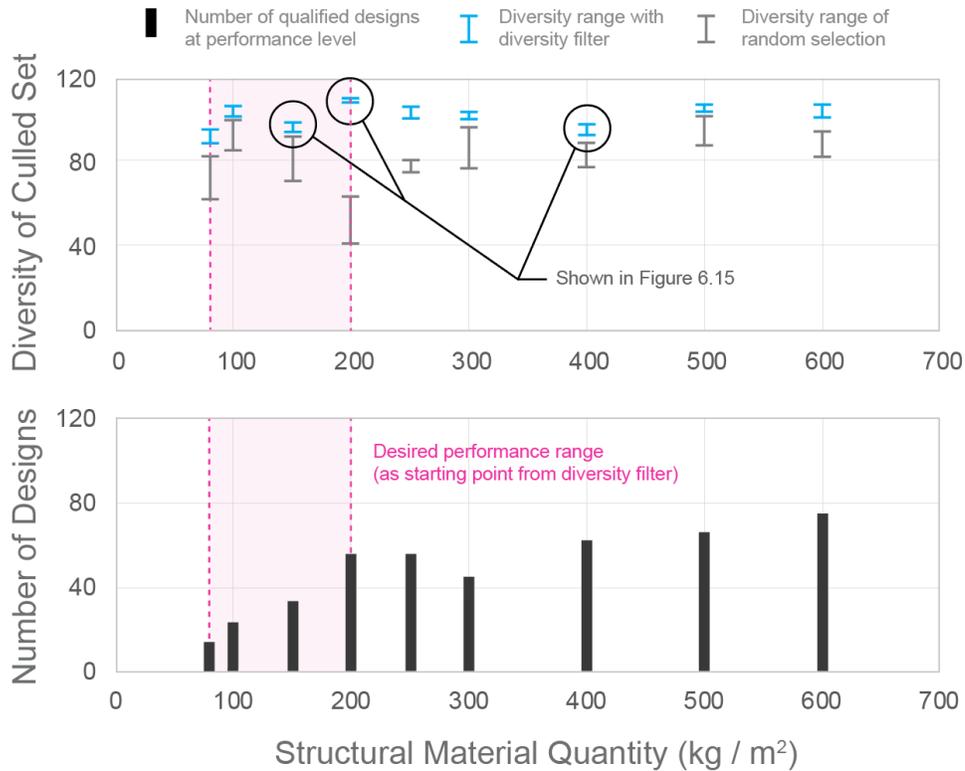


Figure 6.14: A comparison of number of qualified designs and diversity of diversity culled sets for brainstorming, at different structural performance targets

Regardless of which level or individual designs would be selected, it is clear from the large differences and unorthodox geometries that these designs represent non-standard solutions, which expands the possible options during a brainstorming phase. As mentioned previously, this diversity depends on the resolution of the sampled design space, and in some ways by the creativity of the original parameterization. Nevertheless, it is clear that designers would have had trouble rapidly generating each of these options with only a chosen target threshold and an automated optimization process. A practical comparison of user study results in Chapter 7 demonstrates such differences between solutions generated by humans using interactive methods and automated optimization.

While a static catalog approach would still contain all of these designs, the diversity filter makes sure to eliminate designs that are too similar from consideration and allow designers to meaningfully engage with their preferred number of options. Future versions of this workflow could provide more control over the process, such as setting inequality constraints for targets or dynamically sampling based on ongoing user exploration, such that diversity is not so dependent on original sampling. However, the next section offers another path beyond these original starting points, since interactive exploration allows for live exploration between initial samples with surrogate modeling performance feedback.

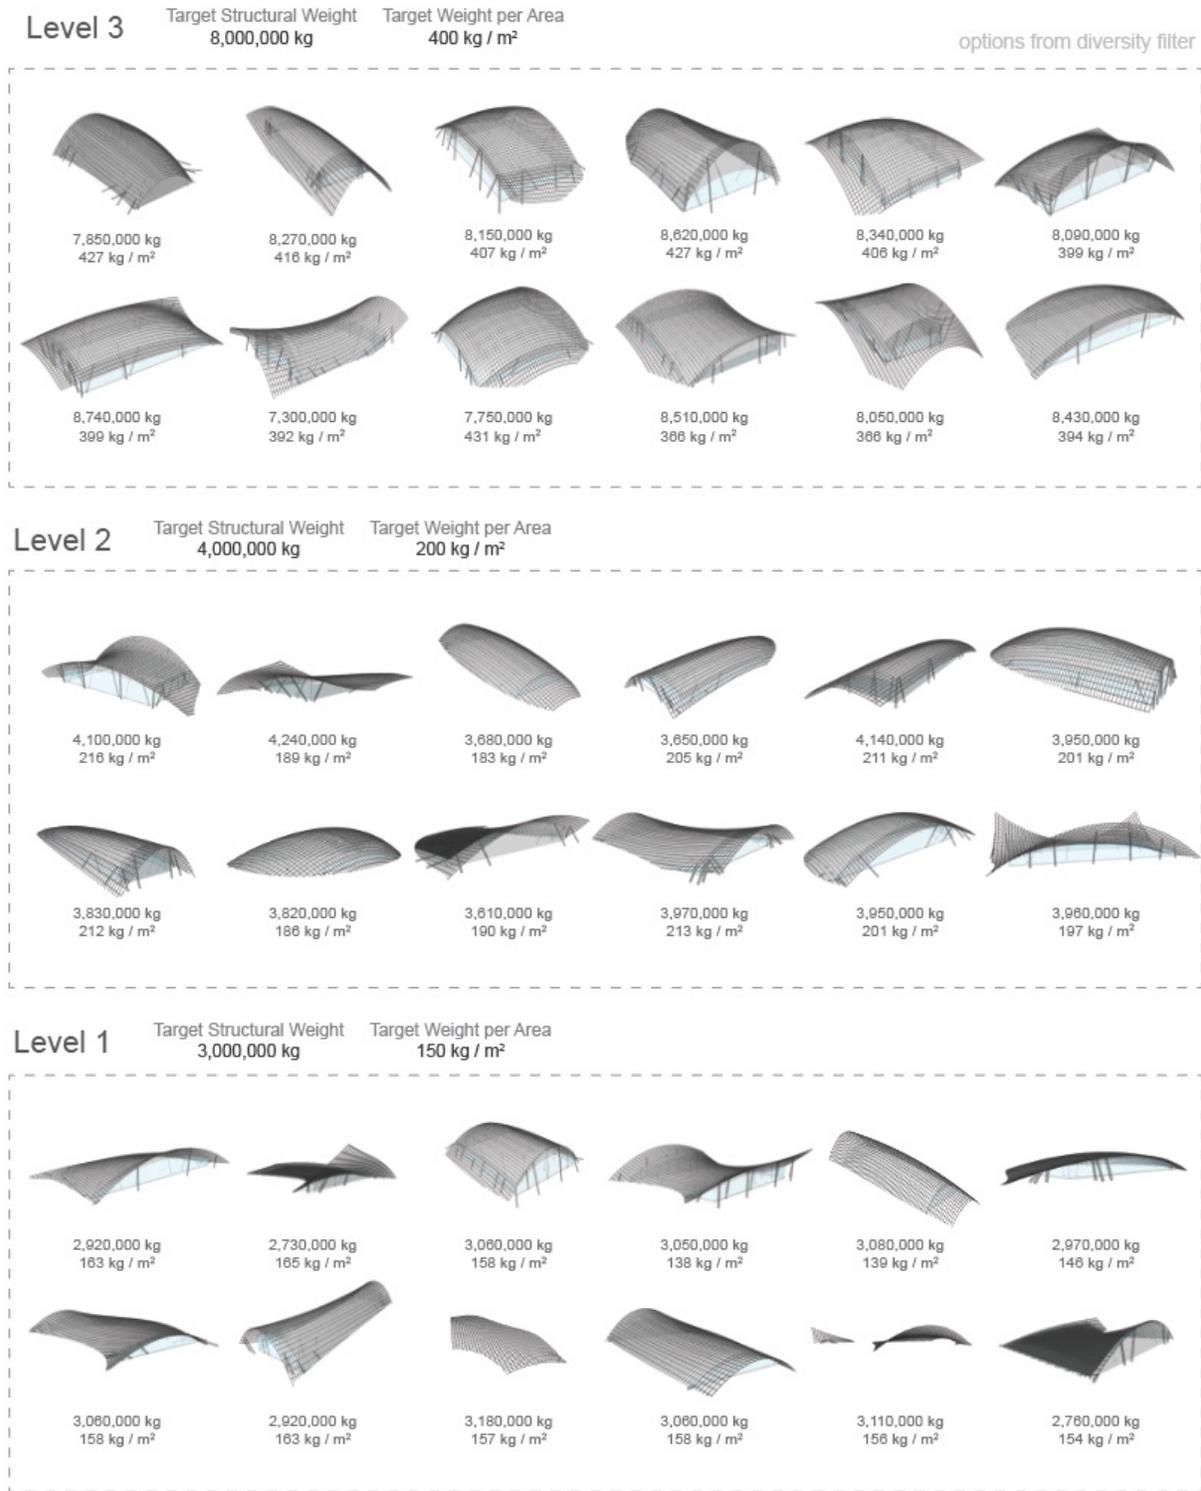


Figure 6.15: Example sets of 12 diverse design possibilities for three different structural performance levels. Moving towards Level 1 improves performance without meaningfully sacrificing too much diversity, but also leads to some geometrically unqualified designs.

6.3.3 Interactive optimization

After adjusting variables and solidifying a useful design space, or selecting a particular design from a brainstorming activity, a user would likely want to continue refining the design. With traditional parametric methods, it is only possible to adjust the original sliders. Furthermore, simulations must be completed at each iteration, or not at all. Instead, using the interactive optimization framework described in Chapter 4, designers can adjust the design while moving in the design space and objective space simultaneously. First, with the previously generated dataset, the designer could train surrogate models that predict in real-time the estimated performance of the building.

For this example, different surrogate models were attempted by splitting the dataset into training and validation data, before the most accurate models were selected for each of the objectives. Random Forest models were found to be the most accurate in each case, and were thus selected and further analyzed using clean test data. The structural model required 12,000 initial simulations for training and validation, while the energy model required 1,000. Example plots of actual versus predicted performance for structural weight per area (500 test points) and Energy Use Intensity (288 test points) are provided in Figure 6.16. In this figure, the dotted lines represent 10% difference from the actual simulation values. It is clear that the energy model is more accurate than the structural model overall. However, in the feasible structural range for the design problem, there is a strong relationship between predicted and actual data, which can still help support designers making live geometric decisions.

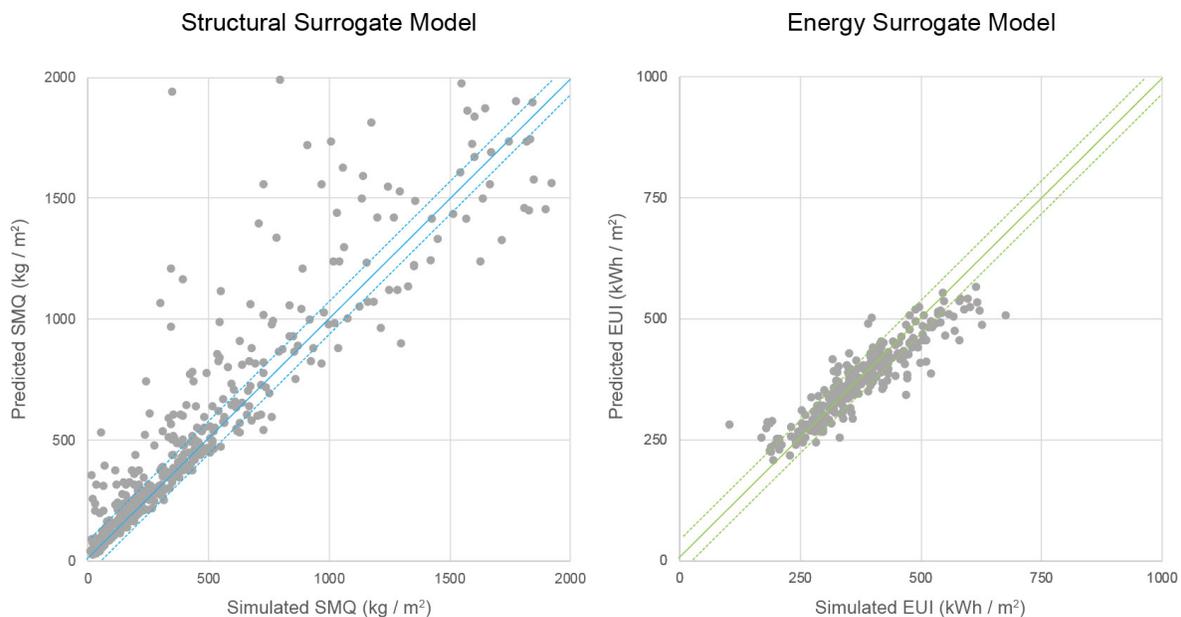


Figure 6.16: Visualization of the accuracy of surrogate models for structure and energy

The output of these surrogate models can then be projected onto the screen along with other geometric design information, as shown in Figure 6.17. As sliders are moved, the results of the surrogate models update in real-time, which the designer can use to further build intuition about the design space and understand if local movements improve or worsen performance. Various visualization techniques can be used to provide the performance feedback—these graphics were developed quickly using native Grasshopper components. As demonstrated in the next chapter, the presence of performance feedback itself can have a significant, positive effect on design experience and outcomes, even without optimization tools and the inevitable error from surrogate models.

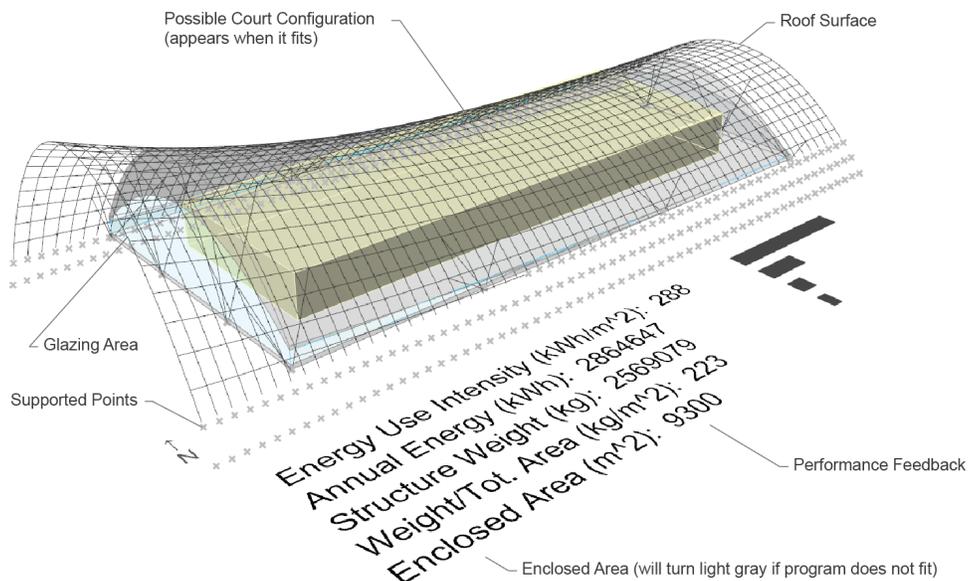


Figure 6.17: Live visualization of geometry and performance (according to the surrogate models) together

The design can also be manipulated interactively using gradient information, which provides the user with specific directions to improve performance rather than a guidance-free exploration. The starting point for this local exploration can be a design already optimized for one objective, a general concept with room for improvement selected using the diversity filter, or any other preferred point in the design space. Using the Stepper component, which has a separate interface (see Figure 6.18), designers are able to select an objective, pick a step size, and click to move in the direction of the gradient or its opposite. Figure 6.19 shows a history of how the objective functions have changed throughout the exploration. In this stage, the designer can begin more meaningfully engaging with all of the design objectives simultaneously, to understand their relationships as they make subtle adjustments and ultimately refine the design.

In this design example, the changing objective functions tend to trend together, but this is often not true as certain objectives trade off for a particular design. Yet the designers might not want to keep improving a given objective until it flattens out, since following a single path through the design space might lead to ostensibly better performing designs that deviate too far from the original design intent, fail to balance competing objectives, or begin to violate obvious spatial constraints. In an automated optimization, such geometric constraints must be manually coded into the problem, which can take time and expertise to do properly. The threshold at which a design might deviate from the original intent is also totally ignored by a computer. By gradually optimizing step-by-step, these issues can be directly managed by the designer without any additional coding.

Through this interface, it is also possible to attempt to move in isoperformance directions, which can be used to traverse between or depart certain areas of the design space or more generally for brainstorming support. In addition, users can select only certain variables to include in the gradient calculation and subsequent movement through the design space. These tools, when used sequentially or partially in parallel, enable a rich, data-driven, multi-objective approach to early computational design.

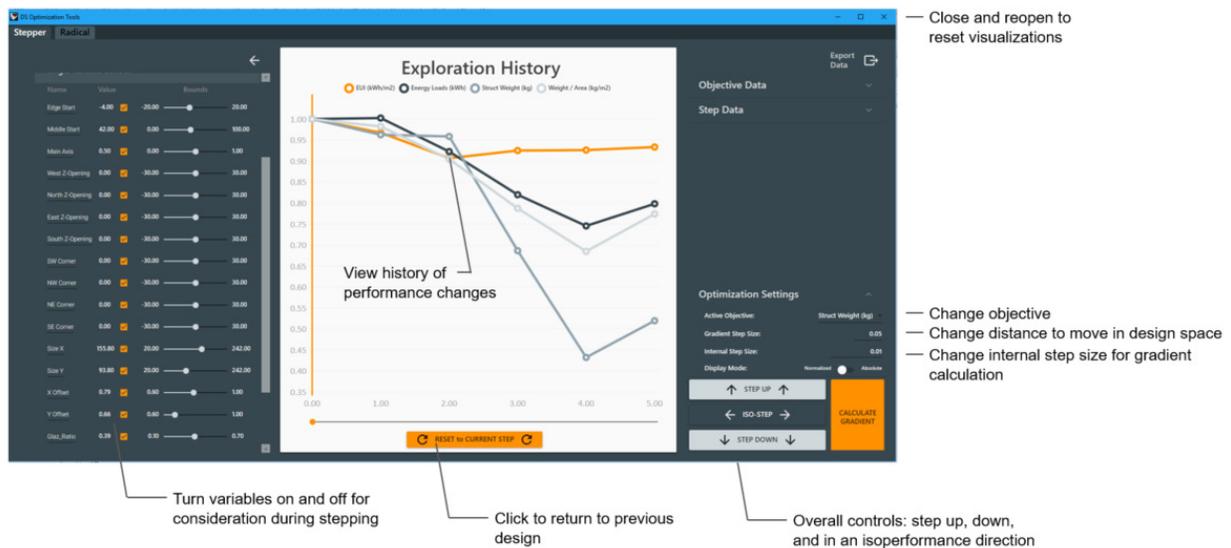


Figure 6.18: The interface for interactive optimization using gradient-based guidance

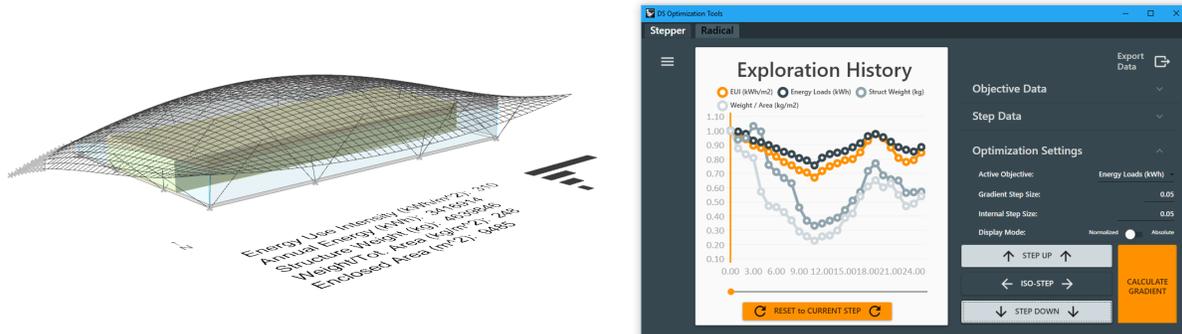


Figure 6.19: A combined view of the geometry, feedback, and control window for interactive optimization

6.3.4 Comparison of designer types and lessons learned

Since these data-driven methods are applicable to many numerical design spaces and quantitative objectives, they might be employed in a variety of design situations. In some cases, they could be used to explore large geometric design differences for expressive, non-traditional applications, with an eye towards specific quantitative performance goals along with other design criteria that must be considered and balanced. In these settings, the designers might still be open to a wide range of possibilities, and not yet have a sense for which aspects of their design should remain constant as they interact with the design and objective spaces. For other situations, designers might already have a sound concept, and simply seek to refine it while considering multiple objectives simultaneously, perhaps because secondary objectives were ignored during the initial ideation. In either case, there are potential benefits and drawbacks to using these interactive methods compared to other parametric workflows that are available.

Figure 6.20 presents two paths through the design and objective spaces for this case study, based on the hypothetical needs of two different design teams. The dotted lines represent the changes in objective functions through discrete steps taken in the objective space, and the start and end points represent actual simulations of these two designs, which give a more accurate evaluation of overall performance changes. These paths were generated using a combination of design space sliding and objective space stepping, which is enabled through DSE, allowing for both direct and indirect geometric manipulation.

In the first example, the starting point is a form that seemed compelling from a visual perspective, but does not perform well. By moving through the design space, the designer notices ways in which the SMQ and EUI can both be reduced, while still maintaining visual aspects of the original concept. While conducting this method, the designers can also consider qualitative feedback and even hard constraints, such as required programmatic area, without needing to code them directly for an automated solver. In this region of the design space, the surrogate model tends to be fairly accurate, and the final result is a substantially better performing concept as measured by simulation.

Large Geometric Changes



Subtle Geometric Changes

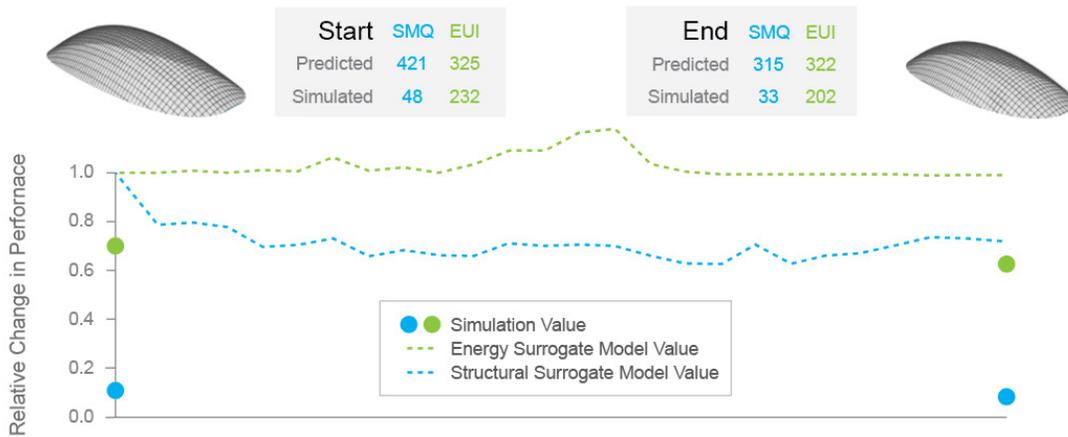


Figure 6.20: Possible paths through the design space using interactive optimization, for two different designer types

In the second design space path, the designers already have a structurally efficient concept. The initial oval grid shell, which is symmetrical and supported along all edges, is a higher performing design than most other options in the space, weighing in at under 50 kg / m². This particular structural surrogate model, which maps relationships between variables and objectives across a much broader design space, severely overestimates the amount of steel required to build it compared to the simulation. The energy surrogate model provides a similar overestimate. These inaccuracies make the stepping itself more difficult—attempting to improve performance in what is clearly near a local minimum leads to more back and forth. In the end, the simulated performance turns out to be better for both objectives after moving through the design space, but not with the change in magnitude for the freer exercise above.

In addition to the paths themselves, which are in both cases ultimately beneficial, this exercise raises other issues that require further reflection. It is obvious from the second example that surrogate model accuracy matters. The goals of this case study necessitated an extremely broad design space, which made it difficult for a surrogate model to properly capture structural performance at the extremes. Such methods must be properly tuned to the resolution of the question at hand, which was not necessarily the case for the second design path. Even for problems in which surrogate models are generally more accurate, the risks propagated by their uncertainty never go away. In other words, this approach does not always lead to an efficient form, due to either user preference, user error, or relationships that are simply too difficult to model.

This is also true of using complicated performance models at all in the early design stages. Especially for structural and energy modeling, which many practitioners and researchers are now doing parametrically, complex design spaces require that it be done carefully to ensure meaningful results. In a context where very subtle design decisions are being made and one or two performance metrics dominate the conversation, it may make more sense to use a catalog or run an optimization overnight, and skip the surrogate model and interactivity altogether. The choice of proper tool or method can also be discipline specific, and a general multi-objective data framework is not always the answer.

Yet the utility of these methods for supplementing creative brainstorming processes, understanding and interpreting the design space, and making important design decisions about overall geometry is clearly demonstrated through this design example.

6.4 Discussion on design example performance compared to benchmarks

While the main goal of this chapter remains the exploration of relative options within a (possibly dynamic) design space, it is worthwhile to compare the performance of the options mentioned here to databases of actual buildings. This comparison illuminates the magnitudes at which it is possible and practical to improve performance in the domains of structure and energy, which has implications for which performance objectives are useful to consider in such a design environment, and how they might be prioritized. It also helps clarify ways in which already performance-conscious designers might apply these tools to supplement their workflows, as well as when these approaches may be used in buildings that are less ambitious in terms of embodied carbon, operational energy, or are constrained from hitting generally accepted performance targets for one reason or another.

First, the range of structural options in this chapter is compared to buildings in the DeQo database (De Wolf 2016), which contains an extensive record of structural material quantities for built projects around the world (De Wolf et al. 2016). Figure 6.21 compares examples in this case study to other buildings based on construction material, number of stories, and longest clear span. The meaningful range of the case study contains an upper bound of around 600 kg/m^2 , which may be an outer limit of what a designer might want

to allow for initial brainstorming, with the knowledge that general concepts at this level of performance could be substantially improved through various interactive optimization techniques. The lower bound considered is around 100 kg/m^2 , and a few carefully chosen structures in the design example are even lower.

These higher performance designs are definitely possible to achieve within this design space, as evidenced by the diversity-filtering options, the results of the stepping exercise, and even the favorite designs found in Chapter 7. It must be stressed that the designs generated in this grid shell example are only estimates, and have considerable uncertainty in their values. The options are more valuable for relative comparison, since the model does not include every loadcase that would be considered in later design, and the geometries generated from the parametric model would almost certainly be refined considerably. Nevertheless, the SMQ of these designs generally places them in a better performance range than most of what is available in the DeQo database.

This result must be understood cautiously, since a single-story roof that can behave primarily in tension or compression should be much lighter per usable square meter than a tall tower (Khan & Rankine 1981), and the DeQo database only allows for the basic segmentations offered in Figure 6.21. Similar efforts to quantify the environmental impact of structures often only report carbon equivalents of structural material (Simonen et al. 2017), which makes direct comparison difficult for this example. In the design context of shells or other lightweight geometries, 100 kg/m^2 is a more reasonable target for a long span structure, even if this is much lower than many designs in the database. Capable structural designers around the world have pushed even lower on SMQ for high-performance roofs. For comparison, the designs for the SFO case study in Chapter 3 are around $90\text{-}100 \text{ kg/m}^2$ as initially modeled, and further modifications through design space movement could obtain values in the range of $70\text{-}80 \text{ kg/m}^2$. These targets should be kept in mind when considering the results of the user study in Chapter 7, which shows where participants generally landed in terms of their favorite designs.

Yet these graphs point to another major conclusion, especially compared to the potential for energy improvement that will be discussed next. For structural designs, even with a given footprint and parametric model, it is easy to find options that are 10x, 20x, or even 30x worse in terms of material quantities. Since stakes are quite high for designers in the structural domain, it is important that these data-driven tools be used by or in conjunction with experienced designers, who can identify potentially good and bad designs based on an understanding of structural behavior. It also follows that users must be careful to not become fixated on one of the first designs considered and work to refine initial concepts, particularly if using divergence-inducing workflows such as the diversity filter described previously. However, even for these experienced designers, data-driven techniques can help discover potentially new forms that perform similarly to well-known shapes, locally optimize particular geometry, or consider structural decisions with the benefit of multidisciplinary performance feedback.

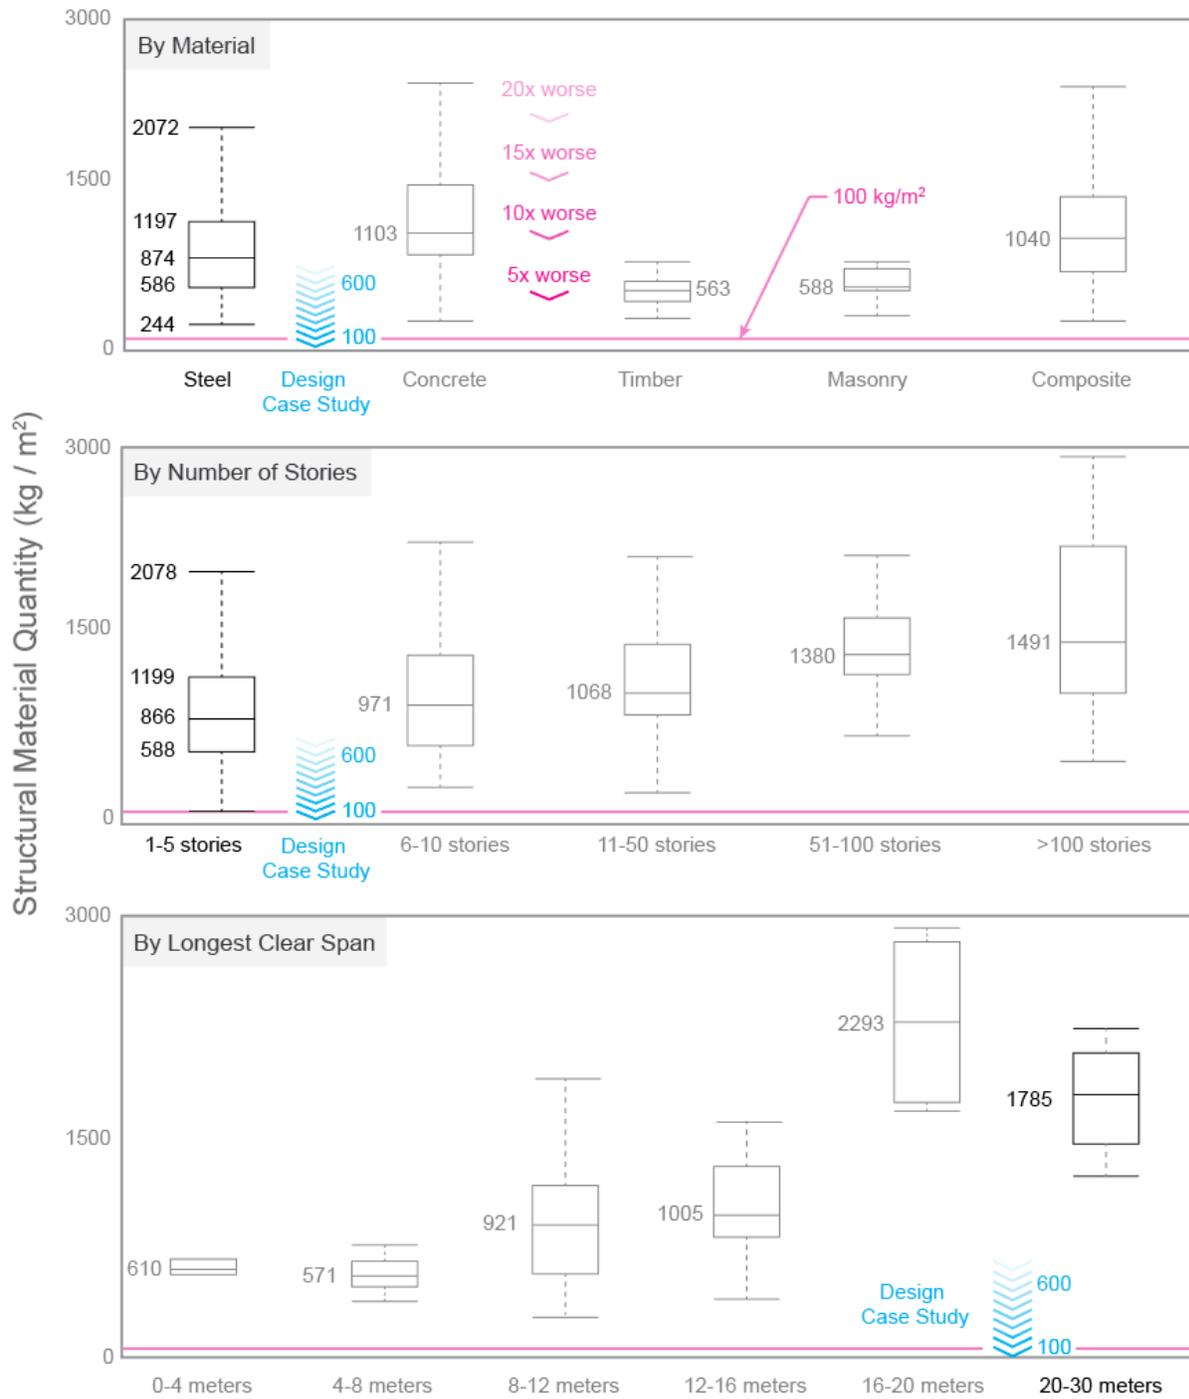


Figure 6.21: The range of performance improvement generally considered for structural weight per area in this case study, compared to similar buildings from the DeQo database, after De Wolf (2016)

For the energy model, two types of comparisons reveal that geometric design decisions have a much smaller performance multiplier than does structure at this stage of design. The first comparison is to similar buildings found in the Building Performance Database (BPD) (Mathew 2015), which was created by Lawrence Berkeley National Laboratory and is the largest dataset of information about the energy-related characteristics of commercial and residential buildings in the United States (Mathew et al. 2015). Figure 6.22 includes all buildings of comparable use in the BPD constructed after 2000, in the same climate zone (5) as Boston. Although there is no specific category for an athletic center, the selection included buildings categorized as education, healthcare, lodging, nursing home, office, public assembly, retail, service, and transformation, while avoiding energy outliers such as convenient stores, data centers, grocery stores, laboratories, and parking garages.

In the case study described here, it is generally possible to move from slightly over the median building to substantially better than the 25th percentile, by making geometric changes using data-driven tools in DSE. This range is a rough estimate based on the interactive design space paths and the general bounds of the sample simulations, and should perhaps be reduced since that design space contains many poorly performing options. Similar to the discussion on structural metrics, being considerably better than the median does not in itself indicate a high-performance building. Depending on use, climate, and other factors, contemporary practicing designers routinely set more ambitious source EUI targets to reduce the environmental impacts of their buildings. Often, these targets are a direct response to the Architecture 2030 Challenge (Architecture - 2030 2006), or a similar push for better performance in buildings. For example, the Zero Tool from Architecture 2030 indicates a baseline source EUI of 305 kWh / m² for a new construction fitness center in Boston with a similar square footage. A 70% reduction would set a target of 92 kWh / m², and an 80% reduction would leave a target of 61.

The geometric changes explored in this example alone cannot reach these targets. Nevertheless, it does appear that early stage energy modeling during geometric exploration has some noticeable benefit, in that it allows designers to move towards geometries that are relatively more efficient, and would likely lead to lower energy uses as the design is further refined. These changes are not nearly as substantial by percentage as the steel that can be saved by designing a more efficient structure. At the same time, annual energy over the lifetime of a building has a much larger relative effect than the one-time initial decision of structural embodied energy for typical buildings, although this relationship is changing with the movement in architecture towards net-zero operation. In either case, it is clear that some early design processes may allow structure to drive early discussions, and consider the energy implications to be a secondary objective.

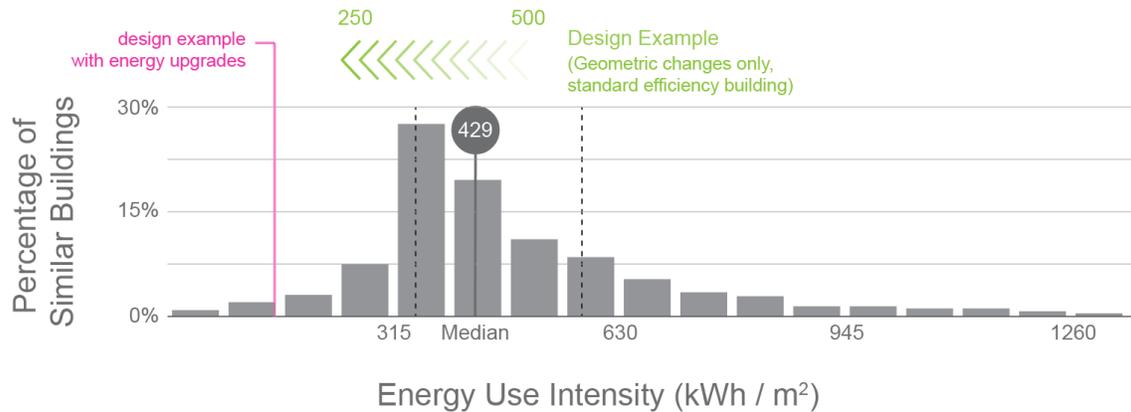


Figure 6.22: The range of performance improvement generally considered for EUI in this design example, using only geometric variables, compared to similar buildings from Lawrence Berkeley National Lab’s Building Performance Database (Mathew 2015)

The second comparison considers the final design obtained from interactive stepping as it was modeled, versus how much the EUI could be reduced by upgrading non-geometric settings for energy performance. For reasons that are described throughout Chapters 6 and 7, which primarily stem from a desire to allow users to perceive differences in energy performance during geometric exploration, the case study was conducted with a code-compliant building rather than one that contains substantial energy efficiency features. This leads to average performance within the design space that is better than the median, but not as highly efficient as it could be. To understand the consequences of this decision, a test model was conducted with better insulation and glazing, more efficient lighting, and less substantial requirements for showers and equipment power density. Figure 6.23 shows consecutive upgrades in these domains, according to the values in Table 6.1. Overall, the energy efficiency features drop the baseline EUI of the selected design to ~150 kWh/m², or around 55% of what was originally shown in the model.

Table 6.1: Energy upgrades and corresponding model settings for the example

Setting	Original	Upgrade
Envelope Upgrades	R-2.75 Walls + R-3.67 Roof Double Pane Clear Glass	R-5.5 Walls + R-9.17 Roof Double Pane Low-E Glazing
Lighting Upgrades	12 w/m ²	5 w/m ²
Equipment & Shower Reduction	12 w/m ²	5 w/m ² , 50% reduction in showers

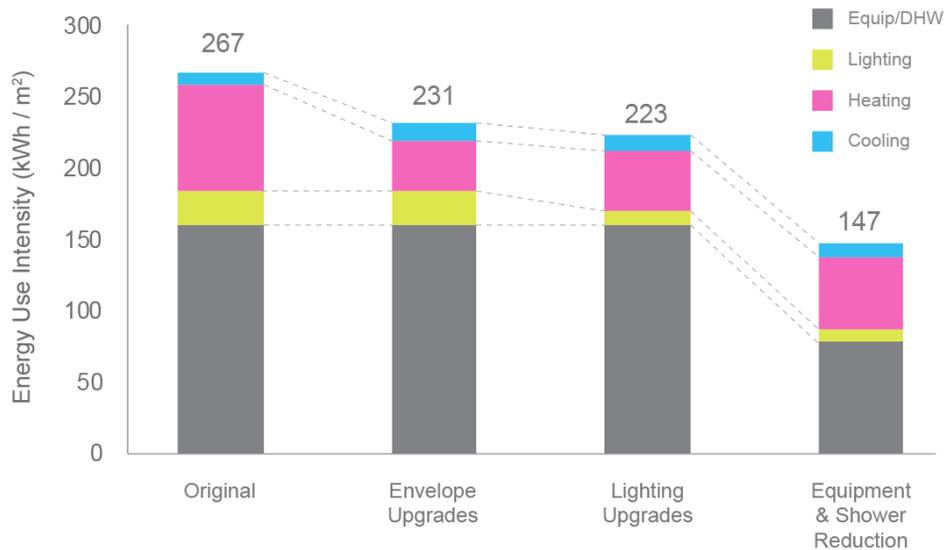


Figure 6.23: History of energy upgrades for the design example, which are here explored independent of geometry

Clearly, these design decisions have a substantial influence over the performance of the building, and might be worthwhile to pursue using data-driven approaches in some fashion. The energy assumptions and requirements of building energy modeling are rapidly changing as designers respond to the contemporary sustainability imperative. As lower EUIs are expected, designers must grapple with whether or not certain aggressive assumptions are justifiable with available technology and information about occupant behavior. However, such decisions can be made independently of the geometric exploration in this case study, and they often involve discrete options with associated costs, which is a different conversation than a geometric exploration of form.

The results from this exploration are further evidence that early energy models are not always as effective “form-generators” as models for daylight, structure, and other objectives, depending on the design variables in question and the resolution of the model. In general, building energy use responds to envelope conditions, ventilation conditions, and internal conditions (Brown & DeKay 2001). Broad decisions such as orientation, ratios between the length, width and height, and shading strategies do affect this first category, and an early energy model with geometric variables captures these changes. The model for this grid shell example would not have found major changes in ventilation or internal loads, as they are outside the scope of the formal exploration. While the example in this chapter does explore significant adjustments in surface area, glazing, and overhangs, which certainly influence energy performance, these geometric changes lead to relatively smaller changes in simulated operational energy compared to structural efficiency. This is why this demonstration started by considering structural form first, before leaving changes in energy usage to be managed during the last step of interactive optimization.

In many similar cases, a design team would be better served by considering a simplified shoebox model to set all energy parameters, and then explore form with those settings already established, along with the knowledge that subtle geometric differences will not have a huge effect on overall EUI. This is an opposite procedure than what was used in this dissertation for demonstrative and research purposes. The actual case study data in this section should be interpreted with this context in mind, and the strategy of modeling a standard efficiency building to better capture geometric differences before moving on to energy upgrades should not be taken as prescriptive for future workflows. In the user study in Chapter 7, the same design example is presented to participants without clear direction on which objectives are more important in early design, leading users to make their own decisions about performance feedback. The proper approach to prioritization when dealing with multiple objectives is worth further consideration, and will be discussed more extensively in Chapter 8.

In closing this section, the relative importance of form on structural and energy models is further demonstrated by the simulations used to generate the surrogate models in Figure 6.16, in which EUIs span from roughly 250-500 kWh/m² for most of the simulations, while many structures use 5-10x as much steel as the majority of the designs that could be selected. These results follow the rough magnitudes of multipliers for structure that exist in real examples, which can be read in the aggregate using DeQo or understood through key examples such as the Beijing Olympic Stadium, which results in over 10x as much embodied carbon per seat than the London Olympic Stadium (De Wolf 2017). Thus, as was the case for the example in this chapter, it might make sense to concentrate heavily on one objective, and acknowledging that another is secondary. The approaches in this dissertation generally entrust these prioritization decisions to the designer, since relationships between objectives can change subject to the needs of a site or program. Even while allowing designers the flexibility to prioritize, data-driven methods provide considerably more multi-objective feedback and guidance than is typical in conceptual design.

This section highlights the fact that the DSE tools, and data-driven approaches more generally, are meant to enhance or supplement designers' abilities rather than replace them. These tools can be used to improve the performance of standard designs, depending on the variables that are selected, or push high-performance designs to use even less steel and energy, while also understanding how these objectives might interact with a long list of related goals. Within a given design space, there are likely to be both good and bad options. Depending on how it is formulated, that design space might need to be continually adjusted and refined to lead to any good solutions at all. This is true of most early stage modeling platforms or comparable design tools, even as they begin to employ techniques like data science and machine learning—an interactive system cannot do all of the work alone, from problem formulation through final selection. In the search for non-standard forms and creative freedom especially, which is an important component of this dissertation, designers still bring their own experiences and sensibilities to the process, often in a more direct, tangible way than they would to automated optimization.

6.5 Qualitative comparison to other tools

While the previous section delved into details concerning performance objectives and their context, this chapter is still primarily about design processes and methods. In response to the widespread need for new approaches to exploring and visualizing the design space using quantitative objectives, researchers have developed a number of existing tools that offer capabilities related to those in DSE. To distinguish the specific contributions of DSE towards data-driven parametric design within the shared environment of Grasshopper, a brief summary of overlapping functionality is provided in Table 6.2. The tools described in this table primarily fit into three categories, with some of them spanning between: (1) parametric toolboxes, (2) optimization solvers, and (3) separate visualization interfaces. DSE contains elements of each category, leading to some shared functionality across the board. Although there are specific advantages and corresponding applications for all of the tools, DSE combines many data-driven components together in a shared format, which allows them to be easily linked together during the explorative design processes described in this dissertation. The breadth of available methods and easy transfer of information between components is not always possible in parametric design, especially for the tools that rely on a specific interface, which do not necessarily function properly with other plug-ins.

Although these examples are the most prominent among Grasshopper users, this list is not exhaustive. Since platforms like Grasshopper encourage continuous coding and modification, others may have developed similar functionality on their own. In particular, many architecture or engineering firms have a group of developers creating digital tools for design exploration among their own design studios or teams. These groups may be formal or informal, and concentrated or distributed across a large firm. However, with a few notable exceptions, many privately developed tools are not publicized by the firm, or open source and available for broad use within parametric design.

Despite the specific purpose of each tool described in Table 6.2, another advantage of the toolbox approach within a shared parametric environment is that designers do not necessarily have to choose between DSE and other plug-ins. While the shared format within DSE has advantages, this format can be connected to other tools using native data and list management components. If a different tool works better for a given application, designers can mix and match functionality as needed, with little additional effort. For example, a designer might use early analysis tools from DSE to initially test the problem, before adjusting the problem formulation and deciding that an automated optimization is more appropriate. At this point, the designer might use another tool such as Goat, Galapagos, or Opossum to find an optimized geometry for further consideration, especially if the design space is better suited for a particular algorithm. Similarly, a designer might want a specific regression type included in Lunchbox to predict performance, and use this model in conjunction with the interactive optimization interface in DSE. Overall, the toolbox approach towards design achieves greater flexibility and accessibility than other tools that require specialized knowledge to operate in a precise way.

Table 6.2: A comparison of functionality within data-driven parametric tools

Tool	Sampling and Iteration	Machine Learning or Surrogate Modeling	Classification	Automated Single Objective Optimization	Multi-Objective Optimization	Integrated Visualization	Interactive Optimization	Inspection and reconstruction of designs
Design Space Exploration	Grid, random, LHS	Ensemble neural network, random forest	Clustering and cluster-based design exploration	Constrained optimization through Radical, a tool that is included with DSE	NSGA-II		Based on gradient estimation; within a new interface	On canvas, through sift
Stormcloud/ StructureFIT				Evolutionary		For interactive optimization	Interactive Evolutionary	
Biomorpher			Designs grouped by similarity	Evolutionary		For interactive optimization	Interactive Evolutionary	
Octopus					SPEA-2 and HypeE	Of design space		
Goat				NLOpt; gradient-based and direct search methods				
Galapagos				Evolutionary and Annealing		Within interface		Within interface
Silvereye (Cichocka et al. 2017)				Particle Swarm				
Opossum		For optimization		Using RBFOpt		With Performance Explorer	With Performance Explorer	Within interface
Lunchbox (Proving Ground 2011)		Regression, Hidden Markov Model, Neural Network	Clustering and other classification					
TT Toolbox (Thornton Tomasetti 2017)	Brute force, and grid through Colibri					Easy export to Design Explorer, Thread		
Generation (Turiello 2013)	Random/ Mixed							
Genoform (Genometri Ltd. 2013)	Random/ Mixed					Through external interface		Through external interface
Dodo (Greco 2015)		Neural Network		NLOpt + Swarm				

6.6 Conclusions

6.6.1 Contributions towards flexibility & accessibility in data-driven parametric design

This chapter describes several practical contributions towards data-driven, multi-objective, parametric design. Primarily, it proposes and justifies a toolbox approach, called Design Space Exploration, by explaining natural relationships between data science and optimization components that enable new data-driven strategies beyond the typical creation of design catalogs. This chapter demonstrates these approaches on a case study involving complex geometry, considerable design freedom, and a mixture of qualitative and quantitative design goals. It also offers a comparison of the functionality in DSE to related optimization tools within a common parametric design environment, which can help designers understand which tools or component combinations may be advantageous for their specific applications. By enabling the workflows described in this dissertation, these freely available tools provide the basis for further implementation of data-driven methods within both research and building design practice, which increases their potential for broader impact.

6.6.2 Future work and concluding remarks

Since the toolbox is in constant development, there are many areas for future work. For one, the tools can gradually be improved in terms of user interfaces, robustness, and additional functionality. Another large topic of research and software development would be to connect these methods directly to a generalized data visualization platform. Although used continuously to clarify ideas, data visualization techniques themselves are largely out of the scope of this chapter and wider dissertation. There are many existing tools for design space visualization as mentioned in the literature review, and some have been connected directly to Grasshopper in compelling ways. While there are certainly opportunities to develop new methods for effective data visualization related to computational design, it was a conscious decision to separate this part of the process from the components that control and operate on the actual information found throughout the design space.

This research decision was partially motivated by already existing visualization tools, but it also stems from an acknowledgement that customized scripts or groups can be written parametrically to convey data within Rhino and Grasshopper. As such, decisions about data visualization should sometimes be left to the designer, and they offer a similar opportunity for creativity during design. Nevertheless, a similar toolbox approach to design space visualization that connects directly to the DSE format would make these tools more immediately useful, especially for quick design tasks. At the same time, the general ideas around data-driven design proposed here are platform agnostic and not married to certain software. As such, future work could extend their functionality to whatever programs designers adopt over the next ten or twenty years and beyond. Each of these areas for future work would build on the conceptual framework for data-driven, multi-objective design as described in this chapter.

7 User Study of Global and Local Design Exploration

7.1 Introduction

As demonstrated throughout this dissertation, the topic of optimization for architectural and structural design has seen significant interest and advancement in recent years, particularly as parametric design methods are becoming more accessible and widespread. Researchers have developed tools for implementing both heuristic and gradient-based optimization techniques in parametric environments, which can be employed as an automated process, as traditionally used in adjacent engineering fields. Depending on how performance is simulated, optimization methods can also be executed interactively, such that designers are able to express preferences while still accessing performance feedback and guidance. To facilitate improved computational speed in automated optimization as well as live approximations for interactive processes, researchers have turned to surrogate modeling. Surrogate models can use prior computation to replace long simulations with a level of accuracy that is often reasonable for the relative sense of performance required when making decisions during conceptual design. Interest in surrogate modeling and related data science methods has led to proposals for a variety of live design exploration techniques. Given that much of this research relies on speculative case studies rather than genuine design processes, there is significant need for further investigation of how designers typically interact with these data-enriched environments.

This chapter presents the initial results of a design study testing the output and workflow preferences of designers as they engage with a parametric model that relies on live prediction of performance. In the study,

participants with backgrounds in architecture, engineering, and building science were provided with a prompt to generate a conceptual design of a long-span athletic center, which they explored using a previously constructed model. The design prompt, which is based on the example from the previous chapter, included information that could realistically be considered in early stage design, such as the site, program, and other building requirements, as well as guidance on constraints and desired performance outcomes. In addition, the computational environment contained surrogate models trained on prior simulations that predicted the structural material quantity and energy usage of the proposed building with effectively real-time response.

Participants were given access to different modeling environments in a semi-randomized order: (1) a basic parametric model and (2) a parametric model with live feedback for global exploration, and then (3) a tool for interactive gradient-based optimization and (4) a tool for composite, multi-objective, automated evolutionary optimization. The design histories for each exploration were tracked along with participants' "favorite" designs, after which the participants were given a survey to describe their experiences. The results of this study yield significant insights into the effects of surrogate-based exploration and optimization techniques on the performance and diversity of generated designs, as well as the preferences of digital designers, which can aid in future tool development.

7.2 Literature review

7.2.1 Need for interaction during design

This user study is meant to quantify and better understand the benefits of interactive design tools for the specific task of performance-based conceptual building design. As mentioned in previous chapters, there are a variety of explanations for why and how live, fast, interactive feedback influences human-computer collaborative tasks. The concept of System Response Time (SRT) is helpful here, which is the time a user waits after entering input until the system provides results back to the user (Doherty & Kelisky 1979). Research by Brady (1986) establishes a relationship between SRT and productivity, while Hoxmeier & DiCesare (2000) show that faster response yields higher levels of satisfaction. These findings are especially true for tasks requiring considerable creative flow (Csikszentmihalyi 1996), such as the engaging task of conceptual building design.

In engineering, researchers have long recognized the utility of fast simulation response. In 1998, the National Research Council declared that effective computer design interfaces should be integrative, visual, and fast, which ideally enables real-time response (National Research Council 1998). Simpson et al. (2007) conduct a user study showing that even delays as small as 1.5 seconds significantly increase error, completion time, and perceived workload for engineering design tasks. Similarly, Faas et al. (2014) find that a level of presence for a designer can be an indicator of how well engineering designs perform. While

long, computationally expensive simulations play an important role in many engineering analysis and design applications, it is clear that for conceptual design, faster feedback is often worth the lower simulation resolution.

7.2.2 Background on design studies

This chapter builds on an extensive body of knowledge concerning the design process in general (Dorst 2011), while at the same time providing new insights into the specifics of performance-based conceptual building design. While contributing towards the establishment of the field of design studies, Cross et al. (1981) argue that design should be regarded as a technology, since it includes the application of various forms of knowledge. Finger & Dixon (1989) provide an early review of design theory and methodology, and Shah et al. (2000) propose experimental guidelines for evaluating conceptual design strategies, primarily in mechanical design. This chapter uses similar methodology while also relying on the new concepts proposed throughout this dissertation.

7.2.3 Design studies from product and mechanical design

Due to the historical development of design thinking, there is a wealth of design studies from fields adjacent to architecture including product, industrial, and mechanical design. Much of this research uses protocol analysis (Ericsson & Simon 1984) and a procedure for measuring the creativity and quality of designs developed by Cross et al. (1996). In this general research area, Dorst & Cross (2001) conduct a design study and describe a model for design as the coevolution of the problem and solution spaces, which is conceptually similar to the formal design and objective space of parametric design. Kruger & Cross (2006) look at solution, information, or knowledge versus problem-driven design processes. More recently, Bao et al. (2018) have compared sketching to prototyping during early design.

Of particular relevance to this dissertation are studies focused on digital tools and CAD environments. Jonson (2005) provides foundational work on the emergent use of digital tools across design fields, finding that for many people, verbalization in combination with other design aides was a common starting point for generating design ideas. Related studies compare the effectiveness of various computational design environments, typically against more traditional tools for ideation. Robertson & Radcliffe (2006) find that CAD can enhance visualization and communication, but also lead to premature fixation and bounded ideation. Robertson & Radcliffe (2009) follow their initial user study with an extensive survey documenting how widespread these occurrences are in design practice. Dorta et al. (2009) compare analog tools (sketching and physical models) to CAD and hybrid environments for industrial design. Moving beyond feedback, Burnell et al. (2017) consider the integration of design, analysis, and optimization in a single geometric CAD environment.

7.2.4 Related design studies in early building design

There is also research specifically pertaining to the conceptual building design process, which is distinct from other disciplines due to its scale, complexity, and breadth of specialists involved. Salman & Laing (2014) present an extensive study of the effect of CAAD on the conceptual design process by recording the verbalizations and design outcomes generated by students in response to a prompt. Similar user studies compare how groups collaborate using technology during conceptual design (Leon, Doolan, et al. 2014; Leon, Laing, et al. 2014).

However, these studies do not consider building performance. The concept of performance-based design (Oxman 2008a) is central to many further studies that seek to understand not only individual or group dynamics during design, but also how well the selected design outcomes perform based on simulation. In this vein, Arnaud (2013) considers the user satisfaction of architects and engineers with current software compared to interactive evolutionary optimization, as well as how both tools are used to generate ideas. Jones & Reinhart (2018) study live design environments for visual comfort. Wortmann (2018) includes a similar user study comparing surrogate-modeling enabled manual, automated, and interactive manipulation of options while designing a pavilion. Brown & Mueller (2016) examine early multi-objective design processes, but use analytical solutions or simple analysis models that are fast enough predict performance without relying on heavy simulation and surrogate modeling, which may be required for more comprehensive conceptual design tasks.

While many of these studies are similar in methodology or even scope, they mostly do not consider early multi-objective design processes at the entire building scale, in which participants must balance a variety of quantitative and qualitative outcomes. In response, this chapter presents an extensive user study involving architects and engineers as they engage with the conceptual design task of generating overall massing and the structural system of a building, while receiving live feedback and optimization opportunities for both domains.

7.3 Study purpose

The purpose of this study is to test the effects of specific performance-driven computational techniques described in this dissertation on the conceptual building design process. The three separate approaches compared in the study are:

1. **Feedback Only** | Free exploration with live performance feedback from surrogate models
2. **Feedback and Guidance** | Interactive, multi-objective, gradient-based design space stepping
3. **Guidance Only** | Semi-constrained automated optimization

All of these techniques are being tested within the context of early architectural parametric modeling, and are being compared to a baseline parametric model without access to performance simulations. As such,

the study primarily tests the effect of surrogate-modeling based feedback and methods for interactive optimization compared to both free exploration and automated optimization, which have not fully been studied together in the context of early building design.

In previous chapters, this dissertation contains additional contributions in the areas of design space formulation and diversity-driven design. However, while Chapter 6 describes how each major data-driven approach can be used together, in sequence, or iteratively for an integrated approach to multi-objective design, this user study scope was reduced to consider interactive optimization only. Attempting to test the effect of every dissertation contribution simultaneously would potentially muddle or mask important results. Thus, although this chapter employs concepts from design formulation and diversity-driven design while creating and evaluating the design study, it only considers the results of interactive optimization as a general workflow, which was described in Chapter 4. Future studies are planned to further interrogate the effects of the other major contributions.

Although there is some overlap between the techniques being tested in that they are to some extent progressive (for example, interactive optimization can include live on-screen performance feedback), each study participant tested the techniques in a randomized sequence. The testing involved two successive sections: global exploration first, and then local optimization. The primary aspects of each environment being tested were:

1. Their effect on performance outcomes, as measured by simulation-based design histories
2. Their effect on design creativity, as measured by the diversity of produced designs
3. Their perceived utility and usability by current designers

In addition, secondary aspects and features of performance-driven design tools were also considered as part of the investigation into usability, which can help future researchers develop effective interfaces for design.

7.4 Experiment procedures

7.4.1 Overview

There were two main components of the study: an interactive design exercise, and a short survey, developed in consultation with and approved by an Institutional Review Board. Participants were first given a design prompt along with a previously constructed parametric model to facilitate their exploration of potential designs. The design prompt included information that could realistically be considered in very early stage design, as well as guidance on constraints and desired performance outcomes.

Prior to taking part in the design study, participants were given a brief (15-30 minute) introduction to the concepts and tool interfaces being tested. After this introduction, they were provided with a computer and

an open Grasshopper file for design exploration. Once they were familiar with the file and the study instructions, they began a self-timed session involving four 12-minute intervals that provided them access to design environments described above, in a partially randomized order. The first environment was a basic parametric model that could be adjusted through pre-arranged design sliders, which controlled geometry that updated on screen. The second environment provided performance feedback on screen in addition to the geometry. The third environment opened a separate interface, which allowed designers to directly improve design “objectives” through rapid gradient estimation and subsequent modifications to the design geometry. Finally, designers were directed to another tool that enabled constrained, automated optimization, using a composite function of surrogate model predictions of building performance with the constraints related to geometric architectural requirements coded as a penalty function.

The study was designed so that global exploration of the design space occurred first, before optimization techniques were applied during local exploration. Global exploration was measured with and without performance feedback (environments 1 and 2), and local exploration was measured for interactive and automated optimization (environments 3 and 4). However, to avoid bias and reduce learning effects, the order of environments 1 and 2 was randomized, and then the order of environments 3 and 4 was again randomized. This allowed the participant to discuss and compare all four environments while benefitting from what was learned during global exploration later in the design study, but offered fairer comparisons to the controls being considered. Free parametric exploration was considered a baseline for comparison to all sections, while the performance-driven environments were also considered directly.

In each design environment, participants were free to adjust the geometry with the provided tools to experiment with the design, and were asked to record up to three “favorite” designs during the session. Following their participation in the design activity, they were automatically directed to an online survey that asked questions about their experiences, tool preferences, and backgrounds.

7.4.2 Design prompt

The design prompt is the same as the example used in the previous chapter. Participants were asked to design a long span structure to serve as an athletic center. The hypothetical building needed to allow room for at least 6 athletic courts, as well as an additional 1,000 square meters for other programmatic needs. The site for this building was Boston, a coastal city with humid, hot summers and cold, wet winters. Participants were asked to assume that the construction site is large enough that context does not meaningfully impact the building in terms of shading, obstructing views, or other aspects of performance and experience. The conceptual athletic center was sited in a prominent location in an educational setting, and participants were told that the client would give considerable latitude to create a signature design. However, the theoretical client was also interested in environmentally conscious design, and desired a high-performance building as well.

The study assumed that a design team had decided on a hybrid structural system involving a curved grid shell roof, which can be supported on large external columns, as well as directly on the ground. When the edge of the grid shell is lifted off of the ground, the resulting wall is filled with a mixture of opaque wall and transparent glazing. Each participant was tasked with exploring early massing possibilities by considering different boundary conditions, curvatures, and orientations within the original design concept. These decisions about massing had both significant performance and visual implications, and was evaluated in a variety of ways.

To assist in this design task, participants were provided with a flexible parametric model that allowed for geometric manipulation. The model was visualized on screen in an abstracted way that delineated the outline and potential glazing area of the envelope, the location of columns, and the curvature of the roof. Variables defining the height of the roof at different locations generated arches, cantilevered edges, and anything in between. It was also possible to define an interior offset for the main envelope, using the roof as a shading structure with free edges. Various global parameters including the width, length, connection to the ground, and glazing ratio were also adjustable.

The parametric model also provided feedback about expected building performance based on machine learning algorithms trained by extensive simulations previously computed for this geometry. These trained surrogate models allowed for the rapid prediction of required structural material, overall energy use, and normalized per area metrics for both structure and energy. Although there is some error in approximating the simulations (average error of ~8% for energy models and ~20% for structural models, when normalized to the mean design), and certain regions of the design space had inaccurate predictions, this information supplemented designers' existing intuition about building performance. Since the same datasets and surrogate models were used for the previous chapter, more information about their speed and accuracy is provided in Section 6.3.

Considering the feedback from these quantitative models, the study asked participants to explore the design space in order to move towards a satisfying design, based on their own values and judgment as a designer. Participants were explicitly told that the exercise was not an optimization game with a clear numerical goal, but rather a digital environment in which they could engage with performance tools based on their own synthetic approach to design, considering all implications of the given prompt. Although their activity was tracked, at various points participants were asked to record up to three "favorite" designs they would want to refine further as the design schedule progressed. Their justification for what makes these designs "good" was left to their own design sensibilities and reactions to the theoretical design task at hand.

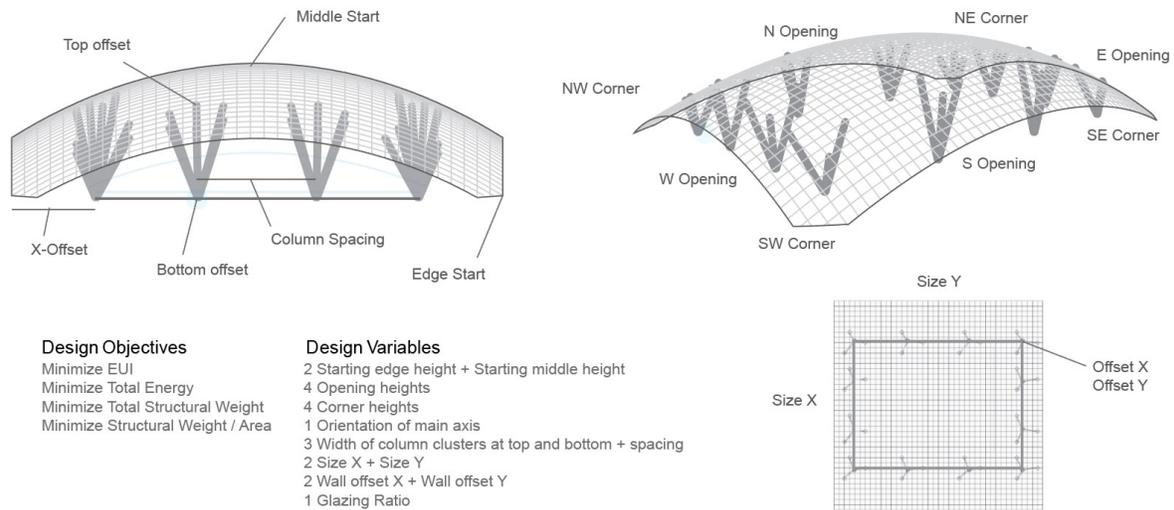


Figure 7.1: A visualization of design objectives and variables for the study. Participants were able to manipulate the variables to produce a wide variety of geometries, and were instructed in the prompt to consider the performance objectives in combination with their own values as designers.

All geometry was modeled using the Grasshopper platform. To arrive at objective function values used to train the surrogate models, separate structural analyses and energy analyses were conducted using Karamba and Diva, respectively. The structural model took advantage of Karamba’s optimize cross section feature such that for each new geometry, the software would run a finite element analysis and perform an iterative sizing procedure that ensured each member was sufficient for axial, bending, shear, and local buckling, according to a procedure for steel beams from Eurocode 1993-1-1. This allowed for an overall estimate for steel tonnage required, which was then normalized by covered area for the second objective function.

The model assumed standard HSS sections as well as larger, deeper custom hollow boxes when standard sections were inadequate, which occasionally occurred for columns or members near interior supports. While these sections are not necessarily what would be selected for the eventual design, they allow for an early estimation of structural performance. Ten load cases were applied, assuming self-weight, a live load of 2.39 kPa (50 psf), and wind loads of 1.44 kPa (30 psf). The load cases involved a combination of dead and live loads, asymmetrical live loads on every possible half of the roof, and lateral wind loads coming from all four directions. The goal was not to have a fully designed structure by the end, but to provide an indication of performance rapidly enough to enable surrogate modeling in a reasonable timeframe and accurately enough to support early design decision-making, especially in a relative manner between design options.

A similar approach was taken for the energy objective functions. First, the geometry was prepared for the energy model by discretizing the curved boundaries into a series of planar surfaces and boxes. A total of nine adjacent zones were generated in a 3 x 3 grid: a central zone offset 10 meters from the exterior, and a corresponding 8 zones surrounding this central zone, such that the corner zones were 10 m x 10 m and the side zones were 10 m deep along the remaining side widths. Heights of each zone were scaled to the average height of the roof within that zone. Rectangular windows were placed along each exterior zone scaled to the area dictated by the glazing ratio.

The roof overhangs were discretized into a triangulated mesh and included as shading elements in the energy model. Walls were assigned an R-Value of 2.75 K*m²/W, and the roof was assigned to be 3.67 K*m²/W. This decision offered a slight magnification of geometric effects on energy performance than if extremely good insulation was used, to aid in making relative design decisions between geometries. Considerable discussion of the impact of using high efficiency versus standard models during geometric exploration is provided in section 6.4. Minimum ventilation rates followed the prescription for health club/aerobics room in ASHRAE 62.1. For lighting, a target of 500 lux was assumed (for athletic facility usage) along with continuous dimming for natural daylight to have an influence on lighting loads. Again, these assumptions were selected to enable early stage relative comparison. Additional information about the model assumptions are provided in Table 7.1.

Table 7.1: Additional model settings for the case study

Model Type	Model Parameter	Model Setting
Structural Simulation	Vertical live load, wind load	2.39 kPa, 1.44 kPa
	Cross sections	HSS and custom
	Loads considered	Dead, live, asymmetrical on every half of roof, wind-all directions
Energy Simulation	Perimeter zone depth	10 m
	Window to wall ratio	variable
	Wall construction	2.75 K*m ² / W
	Roof construction	3.67 K*m ² / W
	Heating / cooling set points	20° C / 26° C
	Occupancy, equipment, lighting schedule	40, 55, 70 h / wk
	Service hot water peak flow	0.0015 m ³ / h / m ²
	Occupant density	0.2 p / m ²
	Fresh Air (per person)	10 L / s / person
	Fresh Air (per area)	0.3 L / s / m ²
	Equipment power	12 w / m ²
	Lighting power	12 w / m ²
	Dimming	Continuous, at 500 lux
Sensible heat recovery ratio	0.5	

As mentioned in Section 6.3, a sample of these structural and energy models within the design space provided the data set for training predictive surrogate models that informed designers of estimated performance changes in real-time. For this study, Random Forest models were found to be the most accurate in each case through testing different prediction model types by splitting the data into training and validation sets. In the end, the chosen structural model required 12,000 initial simulations for training and validation, while the energy model required 1,000. Example plots of actual versus predicted performance for structural weight per area (500 test points) and Energy Use Intensity (288 test points) are provided back in Figure 6.16. Since the structural behavior is much more difficult to model accurately, despite the larger sample set, the energy model is almost always more accurate, which might have had an effect on the design study itself as participants were aware of this discrepancy. However, in the feasible structural range for the design problem, there is a strong relationship between predicted and actual data, which can still help support designers making live geometric decisions.

7.4.3 Population and recruitment

This study included 34 graduate students as participants. These students were recruited through masters and PhD-level Architectural Design, Building Technology, and Masters of Engineering programs. An effort was made to balance students with backgrounds in architecture, engineering, and building science, although in the end 20 students reported experience in architecture only. All but five students reported professional experience in either architectural or engineering design, with an average of 2.6 years of experience. Students were offered \$30 as compensation for their participation. Participants were screened for a minimum skill level with Grasshopper, such that frustrations with the basic interface did not affect the overall results. On a 5-point scale with range “None / Beginner / Moderate / Frequent User / Expert”, participants reported an average familiarity of 4.0 for Rhino, 3.3 for Grasshopper, 2.0 for structural modeling, 2.5 for energy modeling, and 2.1 for optimization.

7.4.4 Location and length of subject involvement

The study took place in a large university computer lab over a series of evenings in February and March 2019, offering standard computing environments (27-inch iMacs) for each participant. The design study was conducted using Rhinoceros 5.0 and Grasshopper, along with custom Grasshopper components developed to enable data-driven design exploration and manage the study itself. The design script was personally installed on each computer and checked by the study administrator. No more than five participants were present at any of the sessions, allowing for all questions to be answered before and during the study. Prior to the exercise, participants were given an opportunity to read through the study and ask questions, in addition to the 15-30 minute introduction to the concepts and custom tool interfaces.

Students were provided with a randomly generated participant ID that linked their survey responses with their design histories. During the study, data about the exploration was collected, including all designs

considered, the time they were considered, and which designs were designated as favorites. The data was recorded in the format of design variable settings, performance simulation results, and timestamps through a custom-written data collection component in Grasshopper. No identifying information was captured by the study. Participants were provided with a consent form at the beginning to certify that they understood the potential risks and benefits of being involved.

7.4.5 Study hypotheses and data collection

Prior to conducting the study, a series of specific hypotheses were developed, with corresponding methods for testing them through the data collection. A summary of these hypotheses is given here:

Primary

1. **Data-driven tools improve design performance** | tested by comparing simulated favorite designs; supplemented by trends in the design history
2. **Data-driven tools are preferred to performance-less environments when both are available** | tested through survey answers
3. **Data-driven design tools do not have a significant negative impact on design diversity** | measured through comparison of diversity of favorites

Secondary

1. Current iteration of interactive design tools is faster than long simulations, but their speed could still be improved
2. Performance-driven design environments influence designers to feel more confident in their final design outcomes, while interactive environments cause them to feel more engaged and creative
3. Specific aspects of tool interfaces (i.e. a pop-up window from grasshopper and/or the ability to exclude variables during optimization) are beneficial

As mentioned earlier, these hypotheses are tested through analysis of the design exploration histories, simulated performance of “favorite designs”, and the results of a post-study survey. By combining both design histories and final outcomes, it is possible to evaluate the different environments by average and most meaningful performance, as in Girotra et al. (2010). The data for these analyses were collected through the Grasshopper interface using a custom component that recorded the design vectors and surrogate model results of every design considered at the resolution of ~1 second. This component also recorded the phase, favorite designs, whether or not a geometric constraint was violated, and other study information.

7.5 Description of design environments

Each of the design environments used either standard Grasshopper functionality or DSE tools as interfaces for manipulating geometry. For phases 1 and 2, participants interacted with the design using 19 sliders to

control the variables. Basic custom Grasshopper visualizations projected the surrogate model predictions directly into the modeling environment, along with visualizations that informed the designer of where courts might fit, if there was enough space in plan, and where the roof was supported on the ground. An image of the Grasshopper script and corresponding visualization used for Phases 1-2 is provided in Figure 7.2. Phase 3 allowed for direct manipulation of sliders, but added the Stepper tool from Chapter 6. Through this environment, it was possible to directly increase or decrease the objective functions, or attempt to move in an isoperformance direction. The tool also allowed designers to return to earlier design options directly, visualize changing performance, or include and exclude certain variables from the performance-based manipulation of the design.

In Phase 4, participants were given a scaled composite objective function in Grasshopper, which could be adjusted based on numerical objective weights. This function included a stepwise penalty for designs that did not satisfy the geometric constraints, either by being too small in plan or not offering enough vertical space for the athletic courts. The composite function was connected to Galapagos, which designers were instructed to use repeatedly to run optimizations for different prioritizations between the objectives. Galapagos was chosen over other solvers in order to compare the proposed methodologies to a native and widely-used tool in Grasshopper. Since Galapagos was being used interactively, the visualization and design reinstatement afforded by its interface created a fairer comparison between interactive and automated optimization than if a more stripped-down solver was used. The next several sections describe the effects of each environment on design quality, diversity, and user satisfaction.

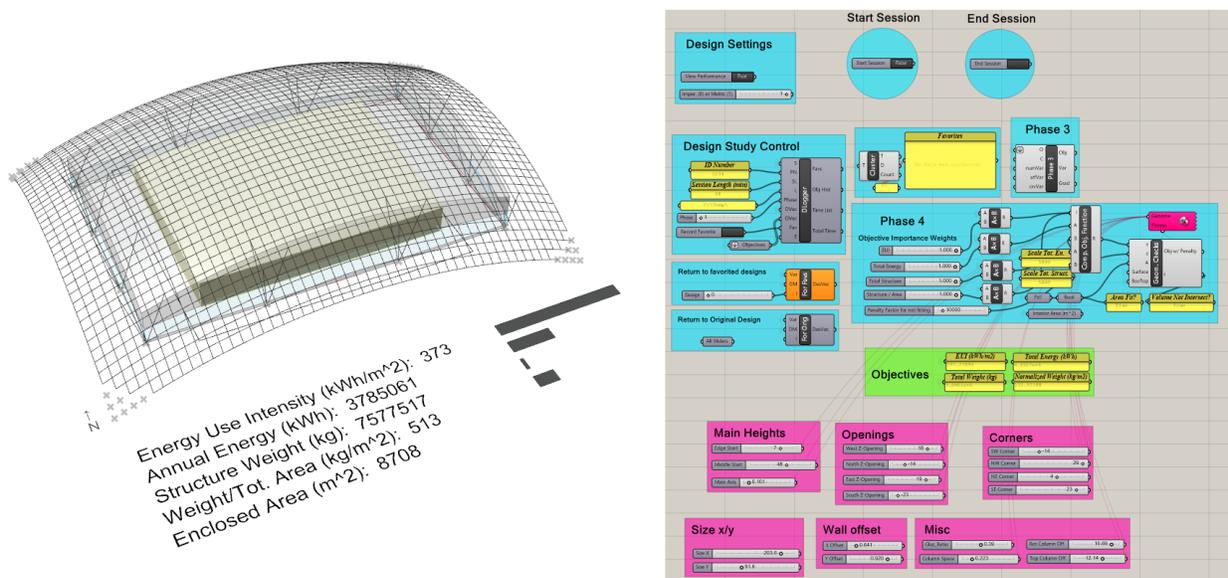


Figure 7.2: Example visualization and design script used for the study

7.6 Design outcome results

7.6.1 Trends from the design histories

First, the general histories of the participants are considered to gain insight into their overall processes rather than just their favorite designs. For the following visualizations, a single participant is considered as one data point, and all activity is condensed into averages for that participant. Figure 7.3 shows the volume of designs considered, as well as the average time per design. For both metrics, the most significant difference occurs between the interactive sessions and the automated session. Optimization allows the computer to consider the most options by far—in reality, the actual number evaluated is much larger than reported here, because the recorder only triggered at the resolution of one second. However, these designs were not meaningfully considered by the designer in terms of non-numerical objectives, which negates any preference expression by the designer, except through the means of explicit objective prioritization while running optimizations.

Among interactive methods, participants took longer and considered fewer designs while using the Stepper tool in Phase 3. Given responses in the survey, much of this additional time was likely spent figuring out a new interface, and thus it cannot be attributed purely to additional design contemplation. At the same time, the comparison between free exploration and exploration with feedback, which were randomized such that any learning effect would be negated, offers a potentially surprising result. Despite providing participants with additional information to synthesize, performance feedback allowed designers to explore slightly more quickly. This might suggest a minor dip in creative engagement with all solutions, as designers may have rapidly moved on from poor performing designs rather than fully considering them.

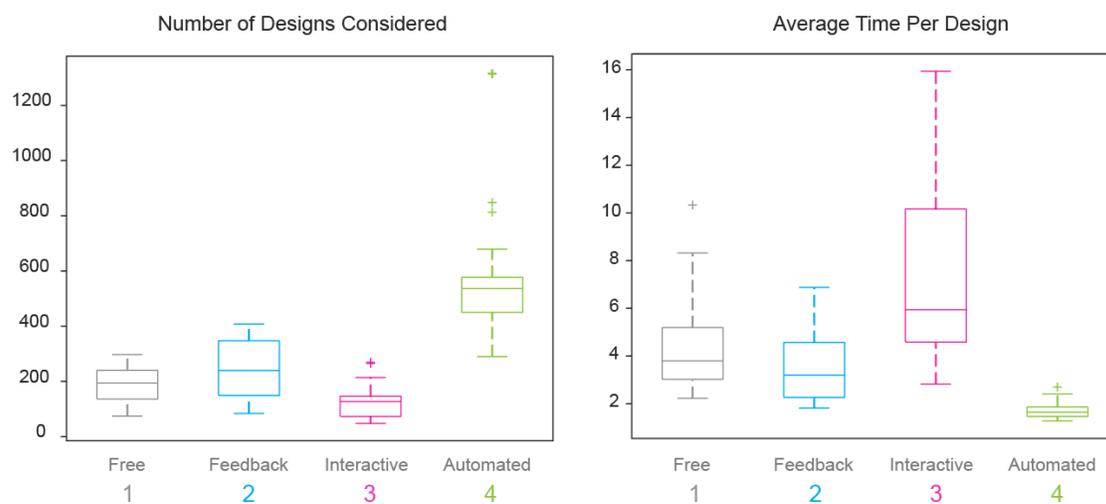


Figure 7.3: Number of designs considered and average time spent on designs in seconds, by phase

Figure 7.4 provides the average performance of all designs considered in a given phase as predicted by the surrogate model. While the actual simulated performance is collected for all favorited designs and described in the next section, this analysis of the surrogate model results reflects the feedback that participants received during the study. For the two energy objectives (given in the top row), the interactive optimization environment seemed to push designers into considerably better performing regions of the design space. Due to its relationship with surface area and volume, the energy model was likely easier for controlling performance than the structural model, and perhaps more intuitive to designers. For the most part, considering designs that were better performing on average according to the surrogate models led to better simulated energy performance when designers were asked to pick a final form.

Differences between phases seem smaller for the structural objectives, although that is partially a visual consequence of the fact that outliers could have much worse performance, which scale the plots differently. In general, each of the three performance-driven phases led to consideration of design spaces that were better performing overall than the Phase 1 performance-free environment, according to surrogate model data. This difference can be partially attributed to learning effects for Phases 3-4, since they occurred after global design exploration from the earlier phases, but this effect is not present for the randomized comparison between Phases 1-2. For the structural metrics, which had more complex objective spaces leading to considerably more error in their surrogate models, it is more meaningful to compare simulated results of favorited designs, even if they do not give a full picture of exploration history.

Average Surrogate Model Values of Every Design Considered, By Phase

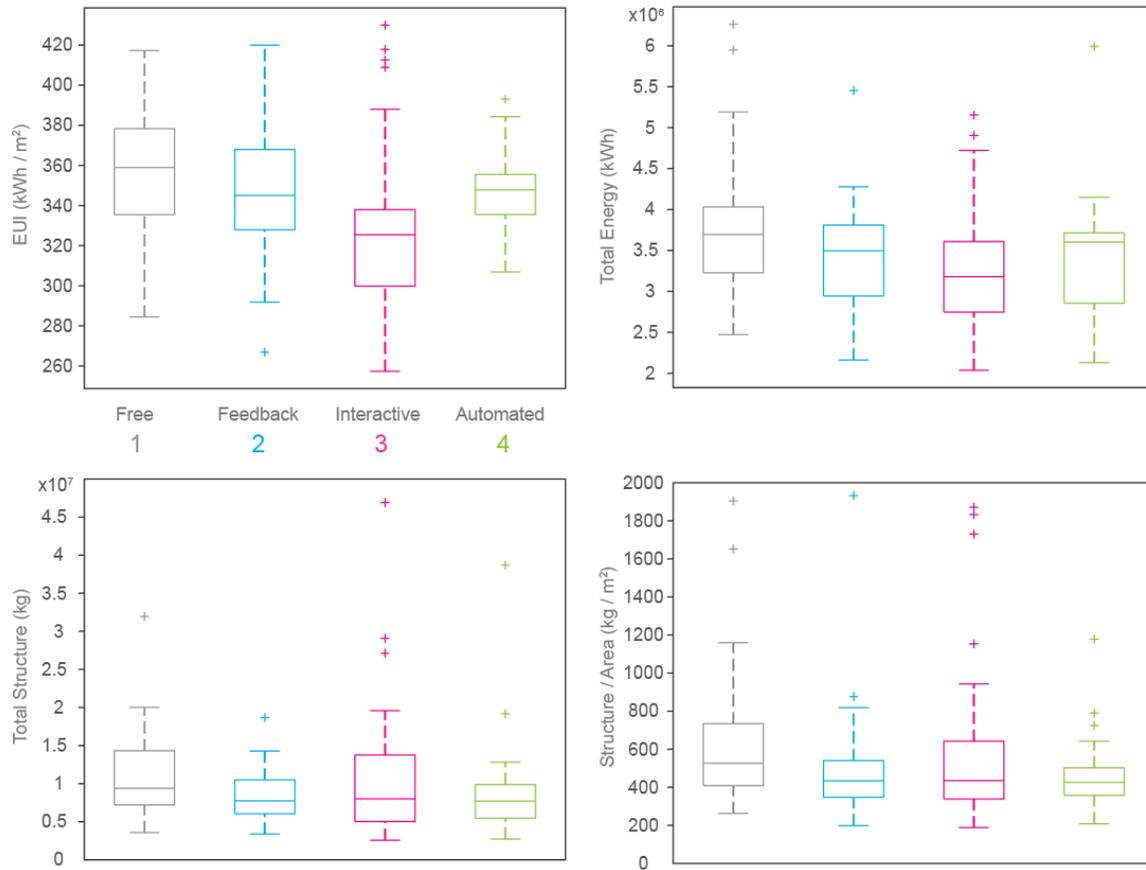


Figure 7.4: The average performance of all designs considered according to the performance surrogate models. Each data point in this set is an average for the history from a single designer, which gives an indication of the areas of the objective space being explored.

7.6.2 Simulated performance results from favorite designs

Next, the actual simulated performance of participants’ selected designs is compared across phases. Up to three favorite designs were allowed per participant, per phase. The analysis in this section provides an indication of whether or not surrogate-model-based methods lead to better design performance based on actual simulations, despite their error present during exploration. Conceptually, this is similar to measuring the performance of an optimization algorithm that uses surrogate models. While error is acknowledged during the process, the final result is more important in the end. Figure 7.5 provides boxplots of the objective performance of favorited designs broken down by phase. To compare the results of each phase, a single factor ANOVA test was first conducted for each objective, testing the null hypothesis that the performance results are statistically equal. Each of these ANOVA tests indicated that at least one of the

phases was significantly different than the others. Consequently, a series of two-tail t-tests were completed comparing the populations of each phase, to find out which performance categories tended to improve in a given design environment. For these statistical tests, each individual design was considered as a single data point. However, to be conservative, an n of 34 (the number of independently observed participants) is used when reporting the standard error. Table 7.2 provides a summary of these average values and statistical tests.

Simulated Performance of all Favorite Designs, By Phase

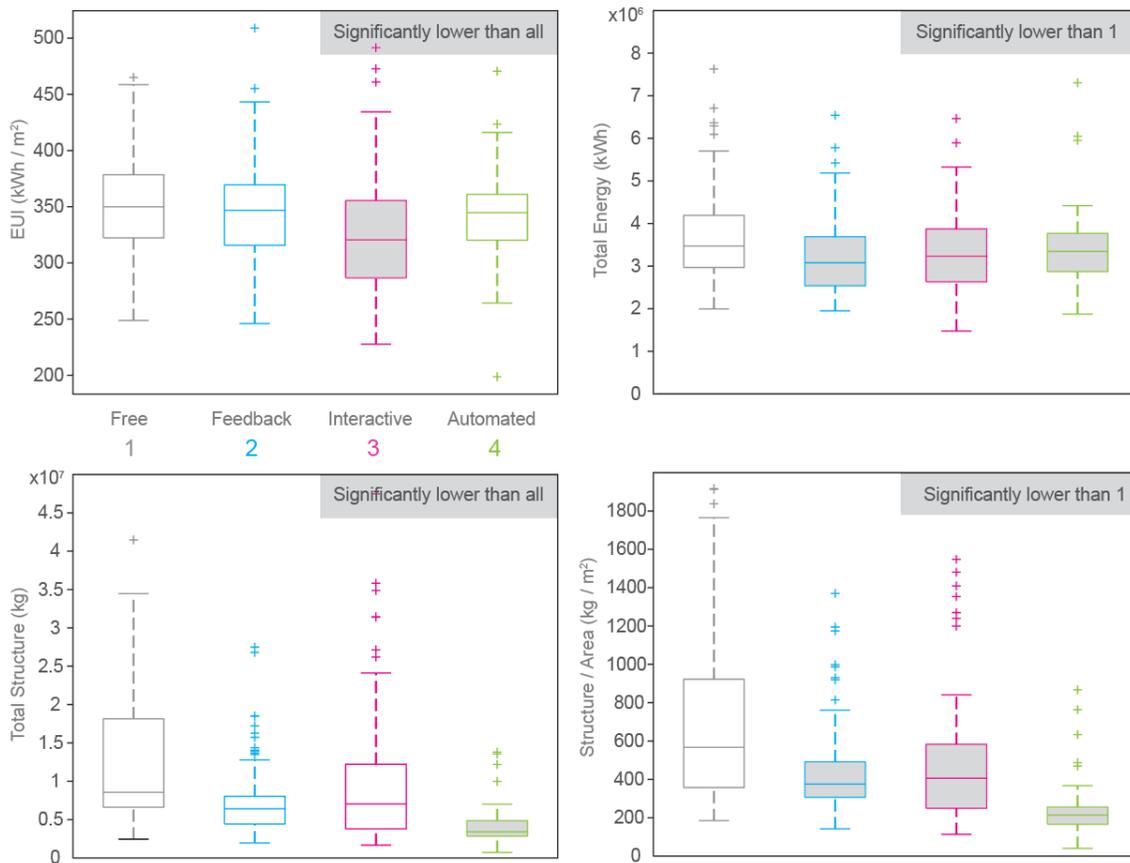


Figure 7.5: Simulated performance of each favorited design from the study. Each data point is a single design.

Table 7.2: Comparison of average performance of selected designs across phases

Group Mean + SEM	EUI (kWh / m ²)	Energy (kWh)	Structure (kg)	Structural Material Quantity (kg / m ²)
1 – Free Exploration	355.3 ± 8.6	3.73E+06 ± 2.04E+05	2.45E+07 ± 1.10E+07	802.9 ± 149.3
2 – Performance Feedback	346.6 ± 8.0	3.20E+06 ± 1.48E+05	9.08E+06 ± 2.85E+06	457.7 ± 56.4
3 – Interactive Optimization	325.9 ± 9.0	3.30E+06 ± 1.47E+05	1.20E+07 ± 2.56E+06	580.9 ± 104.6
4 – Automated Optimization	341.8 ± 7.5	3.39E+06 ± 1.53E+05	4.11E+06 ± 4.33E+05	239.7 ± 24.1
ANOVA: Single Factor (F > F _{cr})	6.44 > 2.63	5.59 > 2.63	5.68 > 2.63	14.85 > 2.63
ANOVA: P-value	0.000294	0.000926	0.00082	3.81E-09
t-Test Summary – Two-Sample Assuming Unequal Variances	Phase 3 significantly lower than all others	Phases 2, 3, 4 significantly lower than Phase 1	Phase 4 significantly lower than all; Phase 2 significantly lower than Phase 1	Phase 4 significantly lower than all; Phases 2, 3, 4 significantly lower than phase 1

For EUI, participants in Phase 3 selected favorite designs that were significantly better performing than in all other phases. For overall energy use of the building (which is not scaled by area), all favorite designs selected with any performance-enabled phase were better than those selected during free exploration. For both structural objectives, automated optimization lead to significantly better performing designs. The other performance-enabled environments also significantly improved the objectives in a few cases. It should be kept in mind that these results offer a fair comparison between Phase 1 and 2, since their order was randomized, but that Phase 3 and 4 occurred after this global exploration in their own randomized sequence. It is notable that the structural performance, which may have been less intuitive and difficult to control for many of the participants, tended to be much better for the automated optimization phase than any of the others. On the other hand, EUI, which is likely the most common metric for building performance and the easiest to control by adjusting the surface area and volume of the building, was considerably lower when using the interactive optimization tool than with any of the other environments.

7.6.3 Diversity of selected designs

Next, the geometric diversity of the favorited designs is considered, which often corresponds to visual diversity. This section addresses the effect of performance-driven design environments on creativity. While creativity can be defined in different ways, as described in Chapter 5, it is generally assumed that measurements of diversity correspond with increased creative freedom, while the tendency towards standard solutions indicates less creative freedom. Figure 7.6 and Figure 7.7 show selected samples from the different design study phases, along with a measurement of their overall diversity. Since the design

space is of a relatively high dimension, an average of the sparseness and outlier methods from Chapter 5 was used to provide this diversity measurement. Although the absolute units of the diversity measurement are meaningless, their relative values provide a worthwhile comparison. The designs shown in these figures are representative across the histories for each phase, showing roughly a quarter of all designs selected. Although designers had complete freedom to submit three entirely different designs, it was common for the favorites from a single designer to be similar, likely stemming from a gradual refinement of an initial concept throughout the phase. As such, an effort was made to select designs from separate participants for visual inspection. These results indicate a similar visual diversity among designs in the first three phases. However, the automated optimization tended to select slightly more standard solutions.

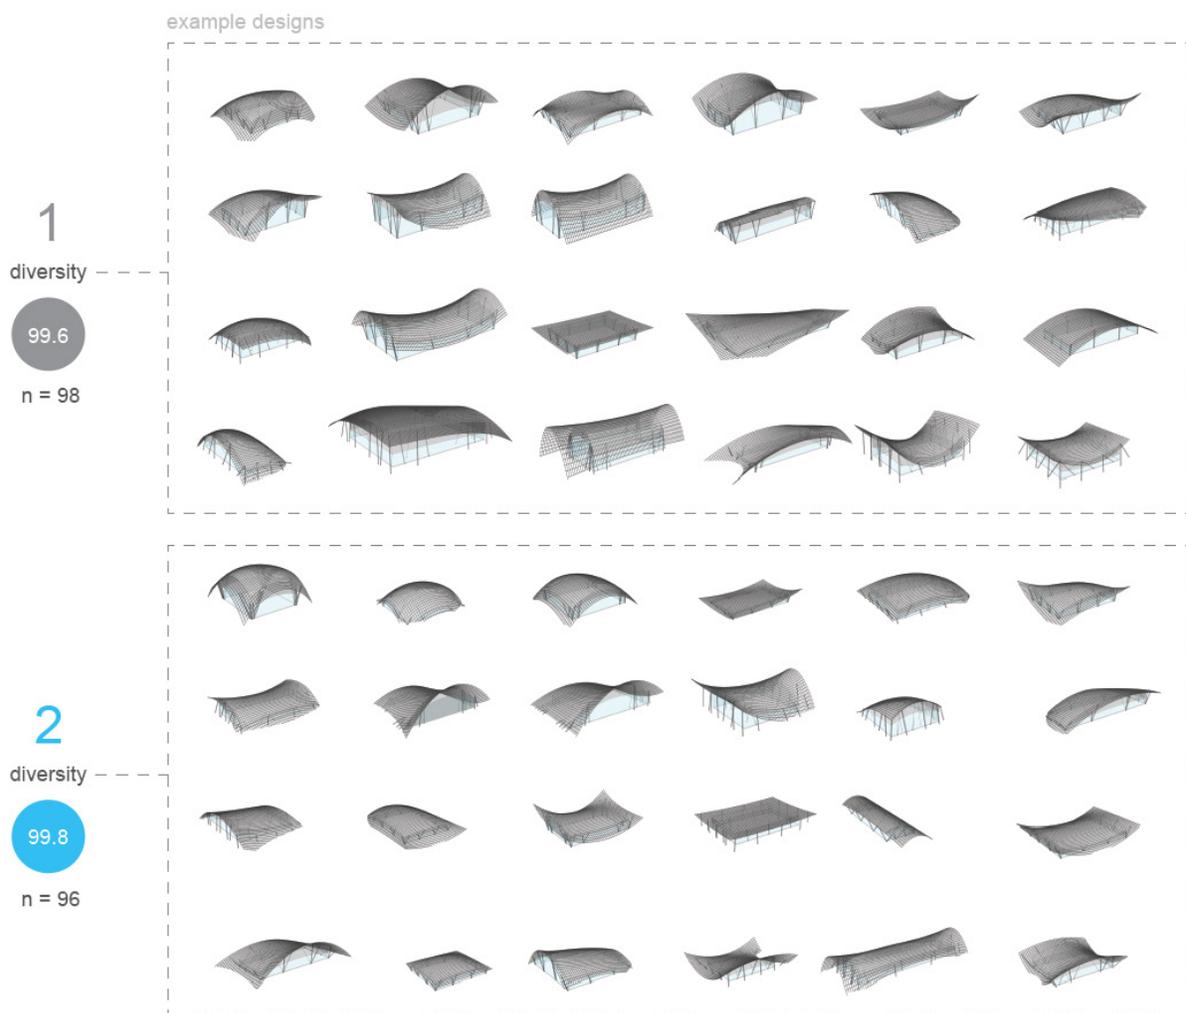


Figure 7.6: Representative designs from the global exploration phases. There is no meaningful difference in the diversity between the two environments.

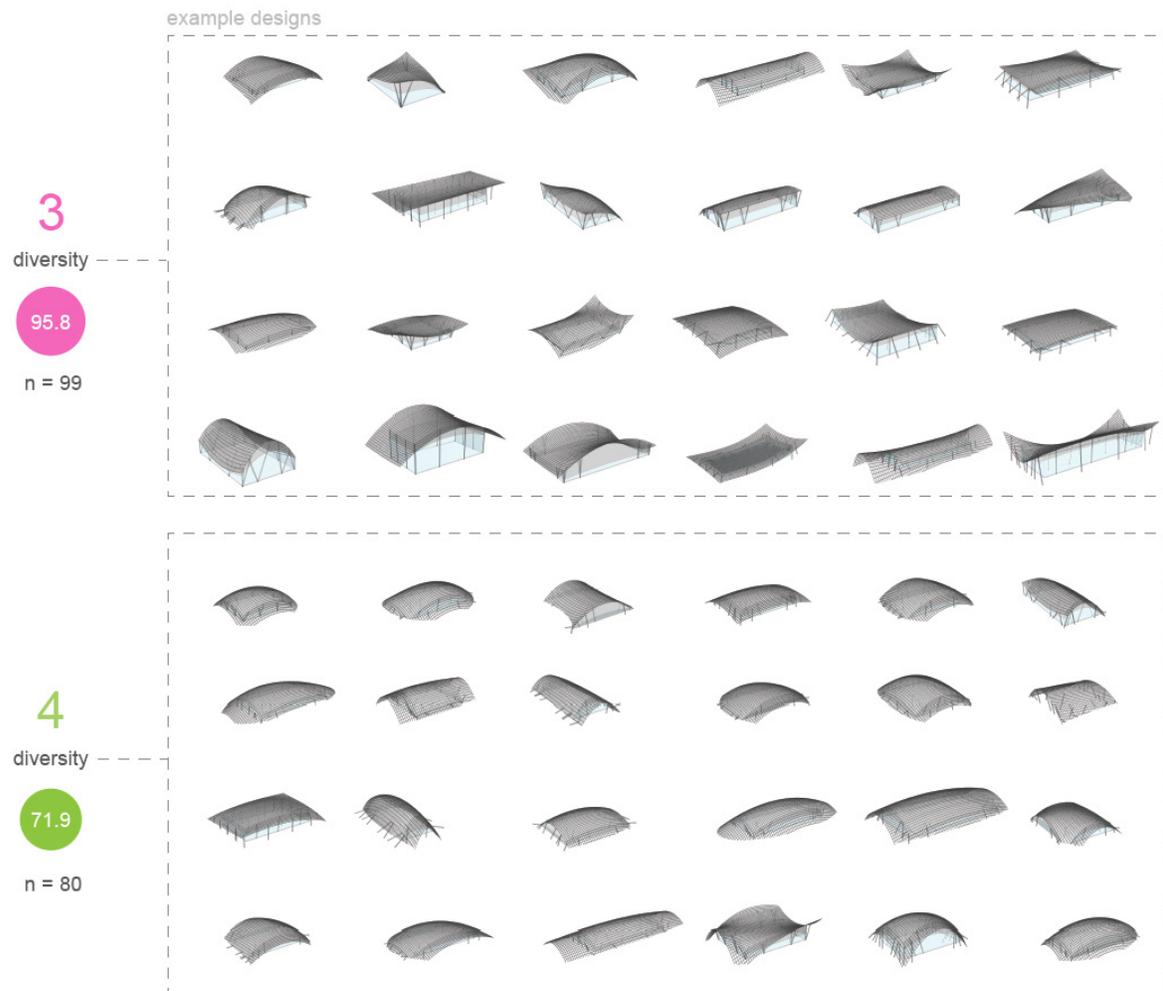


Figure 7.7: Representative designs from the optimization phases of the study. While the selected designs for interactive optimization maintain a similar diversity to the earlier exploration, automated optimization led to slightly more standard solutions.

7.6.4 Segmented results

In analyzing these results, it is plausible that designers with different backgrounds might systematically prefer certain types of designs. Since technical design requires the input of both architects and related engineers and specialists, this study sought participants across these disciplines. One might hypothesize, for example, that architects would generate considerably more expressive or atypical designs, whereas engineers would concentrate more closely on performance. The recruited participants all had experience with computational and parametric tools, which likely self-selected for the types of architects and engineers who share an interest in early stage exploration, and they may have overlapping design experiences and preferences. Nevertheless, it is worth considering the differences between these populations.

Figure 7.8 offers a comparison between the performance of favorited designs by pure architects versus those with some engineering or building science experience. There were 20 architects and 14 engineers or building scientists in the study. Of these 14 technical designers, at least 8 also had some training in architecture. In this graphic, bars above the line represent better (lower) average performance scores from engineers and scientists, while bars below the line represent better average performance from the architects. Across all phases, the architects produced a better average EUI. The specialists tended to produce on average lower total structural material and total energy usage. However, there is no obvious trend that would indicate either architects or engineers were more successful in improving building performance during early design exploration.

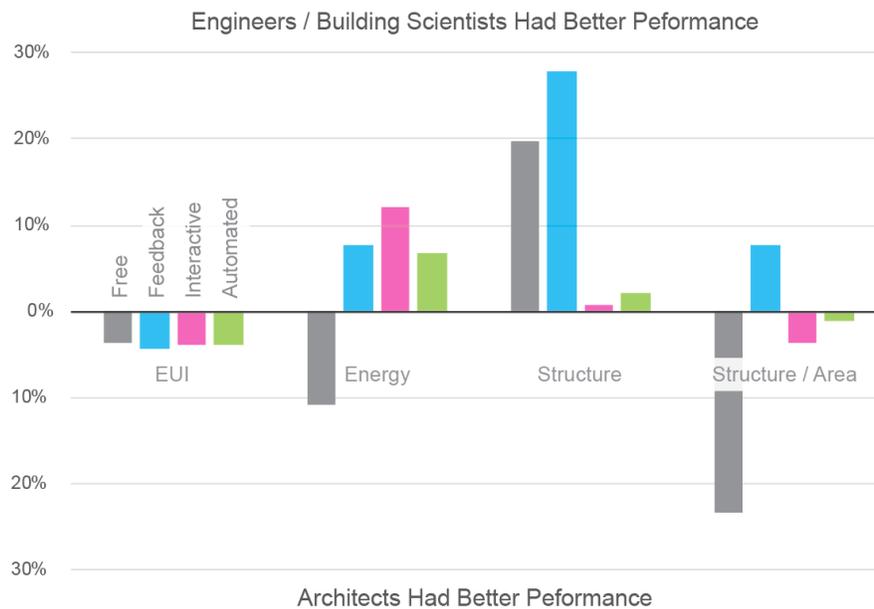


Figure 7.8: A comparison of design performance between architects and engineers, from their selected favorites

7.7 Survey results

7.7.1 Comparative environment preferences

Following completion of the design study, participants were asked to complete a survey regarding their experiences. The first set of questions asked them to directly compare the four phases across seven statements or metrics. A composite summary of this ranking is provided in Figure 7.9, and the full results are given in Figure 7.10. For the composite summary, a weighted sum was assigned such that an environment received more points if a participant ranked it in greater agreement with the statement that was presented to them. First place received four points, second place received three points, and so on. The

overall score has been normalized in the visualization to the maximum and minimum overall score. In Figure 7.10, each colored bar represents the number of participants who selected a particular phase in that rank, such that the first column shows the proportion of phases that were ranked the best. For example, in the upper left graph, the tall blue (second) bar indicates that the most participants felt most confident about their design performance during exploration with feedback, and the short (first) grey bar shows that the least number of participants felt the same about free exploration. The bottom left graph shows that the most participants preferred exploration with performance feedback overall, followed by interactive optimization and then free exploration and automated optimization.

A few notable observations can be made from these results. First, most participants would prefer to conduct freer global exploration of the design space in general before moving to optimization techniques. Automated optimization scored poorly in terms of creativity and focus, which is expected given that it is often used passively during design to generate a single or few options for selection, rather than as an interactive design tool. In general, having at least performance feedback (Phases 2-4) made designers more confident about the final performance of their designs, while the two interactive performance-based phases seemed to teach participants the most about the design space.

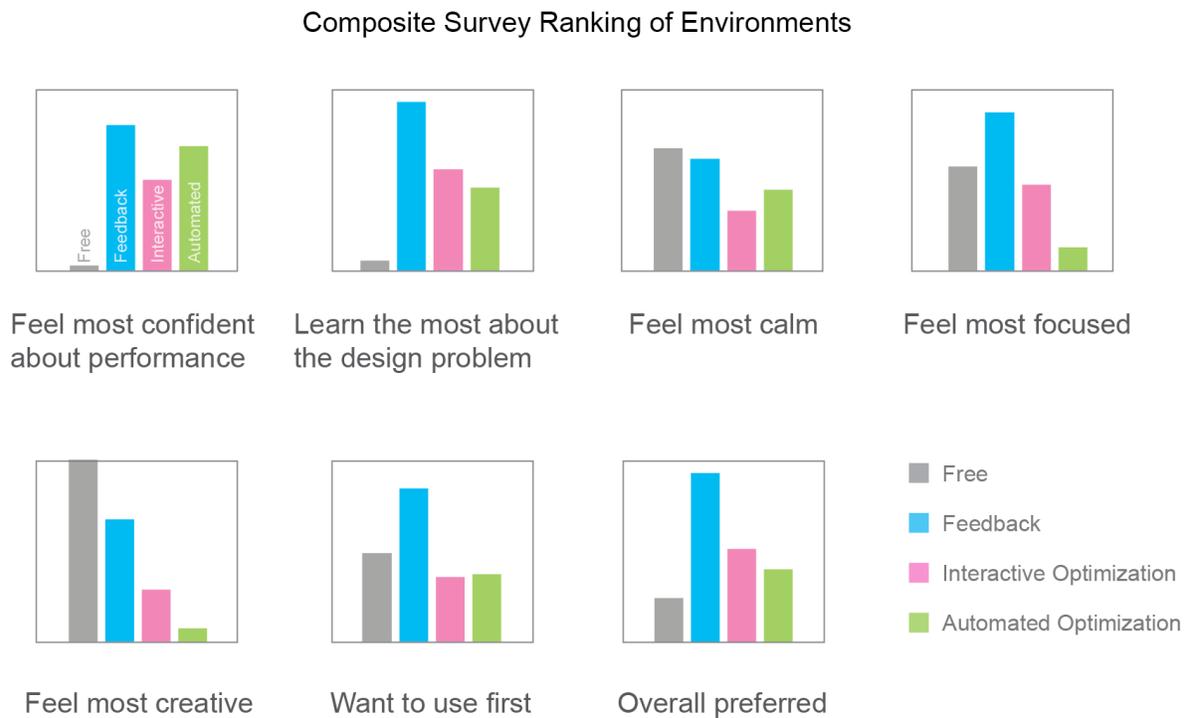


Figure 7.9: Overall composite rankings of design environments using weighted scoring system

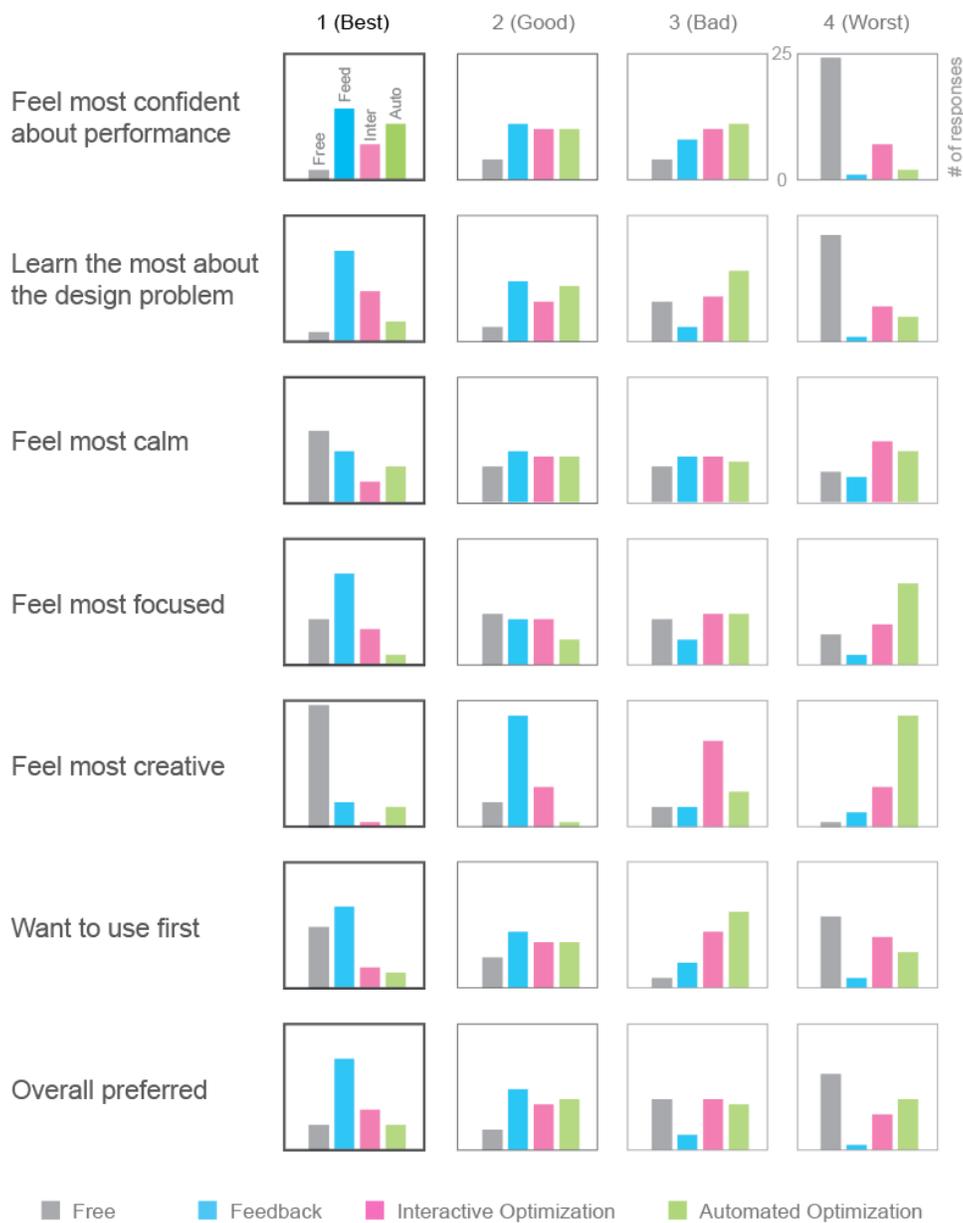


Figure 7.10: The results from a ranked comparison between design environments. The first column indicates the number of participants who chose various environments as best representing the statements on the left.

Designers were the least calm when using the interactive optimization approach. As can be seen from the answers in the next section, this may be partially due to the interface for this phase being newer and more complicated than the others. All three other phases relied primarily on standard components or interfaces designers might have seen before, which did not require as much effort and concentration to understand for the first time. Overall, the most notable result is the clear ranking between which phases made designers

feel most creative. In order, the most creative phases were overwhelmingly free exploration, then exploration with feedback, followed by interactive optimization and automated optimization. This result underscores the necessity of developing interactive performance-based tools for situations in which designers desire creative freedom during computational exploration. Although their exact implementation requires refinement, interactive optimization techniques offer potential for injecting simulations into flexible design environments without sacrificing this creative flexibility.

7.7.2 Individual environment reviews

The next set of questions asked participants to evaluate each design environment individually. Figure 7.11 provides a summary of these results, in the form of average answers and interquartile ranges across a five-point scale. Again, a number of trends become apparent when looking at the cumulative results. First, participants tended to strongly agree that they needed more training to fully utilize the interactive optimization tool, but they were also most interested in receiving additional training for this tool. In general, the exploration with feedback environment scored the best for leading to a quality outcome, most likely to use in practice or studio, most desirable to mention to a client or critic, and computational response being fast enough to maintain interest. These results bolster the overall preference for this environment from the ranking question. Free exploration scored worst in most of these categories, again confirming that designers felt more confident when knowing at least some information about the estimated performance of each design option. All phases except the automated optimization scored well in terms of perceived speed of computational response, which likely enabled participants to remain engaged throughout most of the design task.

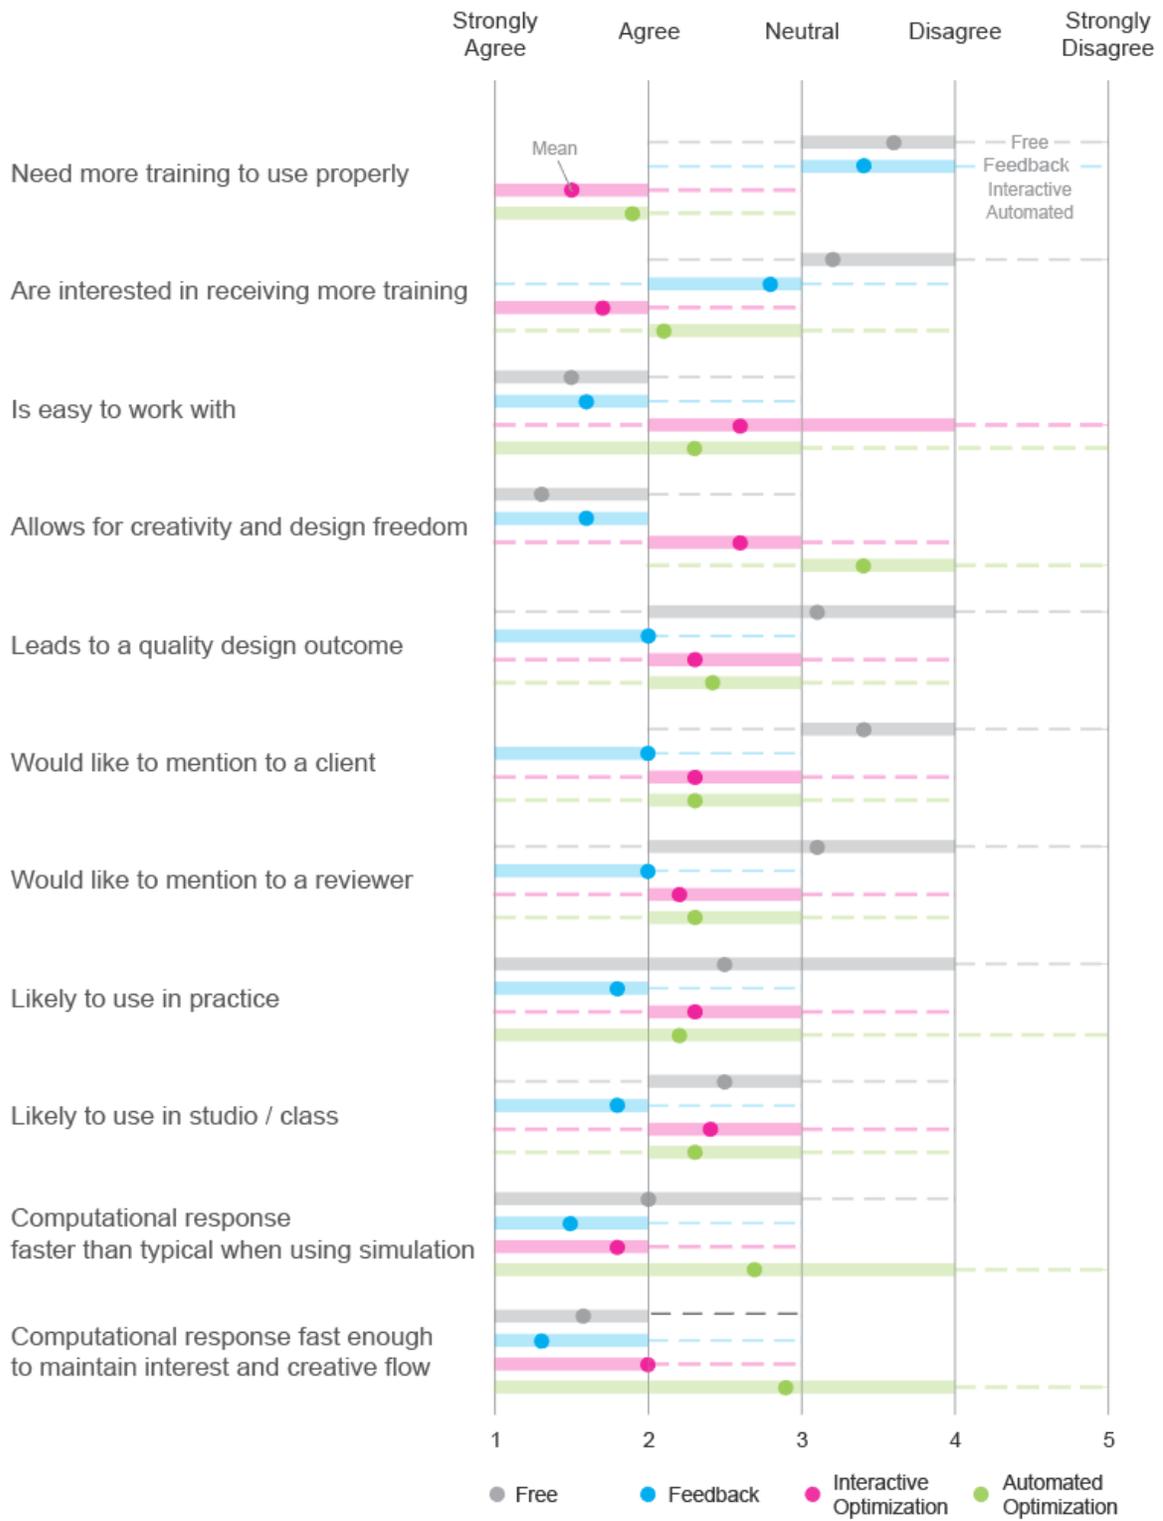


Figure 7.11: Survey results for each environment individually. The dark circle indicates the mean across participants, and the shaded area is the interquartile range (25th percentile to 75th percentile) of the responses.

7.7.3 Tool preferences and general comments

The survey also contained more specific questions about user interfaces for interactive design tools. Especially when working on plug-ins for existing software, developers make choices about what designers see on screen and how they interact with these visual representations. While not comprehensive, Figure 7.12 provides results on a few key questions related to the interactive optimization tools used in the study. Participants overwhelmingly responded that seeing both live geometry changes and direct performance feedback on screen significantly improved their design experience. The ability to easily adjust specific aspects of the optimization was also generally found to be helpful. More rigorous data on this question should be sought during future studies, but tool designers working on interactive optimization techniques must always find a balance between ease of use and customization. Whereas some designers want easy plug and play, others want total control over advanced options. There are thus competing risks that the former may misuse techniques if too many assumptions are made, and the latter may abandon the tool for one with more flexible functionality.

Having a separate pop-up interface gained less support, although over half of participants felt this element was at least somewhat helpful compared to conducting all design manipulation within the Grasshopper interface. While the study was conducted on iMacs with large 27-inch displays, the additional pop up window can crowd a screen that already must show the geometry through Rhino and the Grasshopper script at the same time. While this is not an issue for optimization tools that contain geometric visualizations within their own interfaces, or ones that automatically search through options with no intention of showing each geometry, it is a consideration for interactive tools that rely on the main Rhino screen to communicate information.

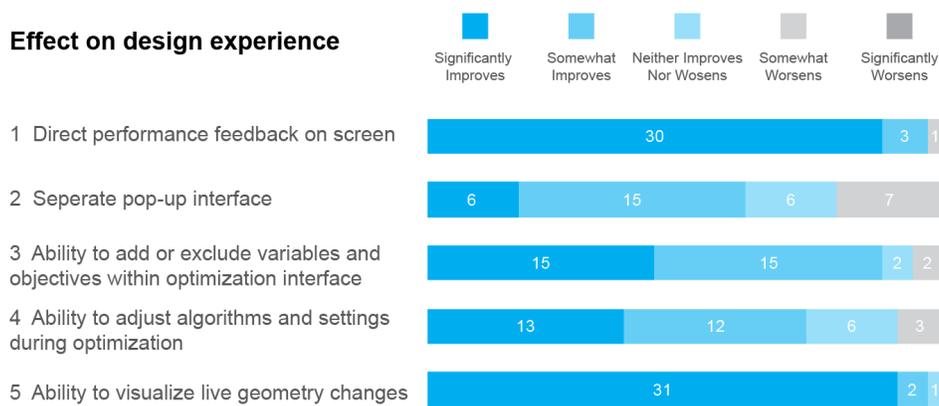


Figure 7.12: Interface preferences for performance-driven parametric design tools

Much was also learned through direct remarks from participants, who were asked to make both positive and negative comments about any of the design environments. These comments should be considered in light of the overall rankings and evaluations provided in the previous sections. On the negative side, some participants felt that the interactive optimization interface was complicated, unintuitive, or confusing. However, a number of these same participants acknowledged that it was the newest interface to them, that they were excited by the concept, and were interested in learning more about how to use it. Participants also complained that the automated optimization process was too slow, leading them to feel bored, disengaged, and not at all creative.

It is worth considering that some of these complaints may be responding to the artificial environment of the study—most optimization users would not be given a short, discrete timeframe and asked to repeatedly run optimizations, but would rather set up the problem, run the optimization, and return to the result later. Nevertheless, when using the surrogate models, an evolutionary solver would often hit a minimum within a minute or two (and sometimes faster), with limited improvement after that. Thus, while using this technique, designers would frequently need to iterate every few minutes anyway, occupying their time in this semi-engaged state, unless they somehow stage repeated optimizations with different conditions. This reality makes for a reasonable comparison of interactive and automated optimization on a fixed timeframe.

Other participants also noted that the exercise felt too complicated in general upon entering the optimization phases, since there were so many objectives and variables to consider. The design problem was intentionally chosen to test this limit of perceived complexity, as well as other limits inherent in the tools, such as the ability of surrogate models to provide functionally accurate feedback with a reasonable amount of data. In fact, the final design space formulation included manually reduced variables compared to other iterations of the design study script. As such, this case study may represent an upper bound of complexity in terms of variables and objectives that can be meaningfully and interactively considered during design using the techniques explored in this dissertation.

On the positive side, some participants reaffirmed their support for either interactive optimization or automated optimization—the interactive technique because it gave them creative freedom and clear control over specific objectives, or the automated technique because it explored the most possible designs, and they trusted it to produce an effective, valid outcome. Others mentioned being pleasantly surprised by designs within the script, ones they would not have thought of independently even when looking at the variable ranges. While most comments simply expressed a desire to use one of the environments further, some specifically mentioned that the interactive ones helped them better understand variables, learn from “mistakes”, and make faster decisions.

7.8 Discussion

While there are many small insights into the early design process found throughout the datasets in this study, the primary objective was to test the hypotheses related to improved performance, designer preferences, and design diversity. There is evidence from the study that in many cases, performance information led to better simulated outcomes, while interactive methods were preferred by designers and lead to the most diverse solutions proposed across the study. To some extent, the interactive optimization technique from this dissertation sits at the intersection of these two trends (see Figure 7.13). Environment two was preferred overall, followed by environment three, according to survey preferences. Both environments enabled freedom of movement through the design space, which is associated with high amounts of creativity. However, environment three is an extension of the surrogate modeling functionality from environment two, as it enables intentional adjustment of the performance metrics. This interactive optimization environment significantly improved performance for more objectives, and it resulted in comparable overall diversity as found for environments one and two.

The study also made it clear that aspects of interactive, performance-based design are still new to many architects and engineers, even ones who are computationally savvy and technically oriented. Researchers must continue to carefully consider interfaces and desired workstyles while developing new tools and commissioning them into the design field. This reality is underscored by the combined study results, in which participants preferred less functionality (environment two) to more functionality (environment three), as least in part due to unfamiliar concepts and seemingly complicated interfaces, even as environment three often lead to comparable or better results.

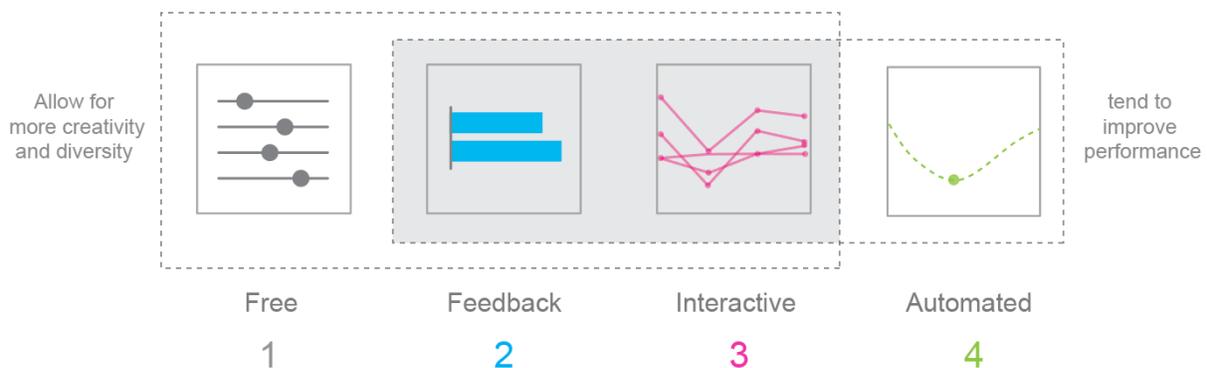


Figure 7.13: A grouping of the environments that allowed for creativity and produced more diversity, versus the environments that tended to improve performance

There are also some potential limitations of the study. The population of designers, though all trained in design and most having professional experience, was recruited from a pool of graduate students, which may

skew the data in various ways. As with any simulated design study, there were little actual consequences for the design choices being made. Although designers were given a realistic prompt and told to approach the design using their own sensibilities and intuition, they were aware of the performance-driven context of the study. In the future, the results of these simulated design studies should be considered in conjunction with investigations of real design processes and their outcomes. In theory, this task should become easier as firms add to their growing digital repositories of information about design options, histories, and built projects, as long as they make that information available to researchers.

Some aspects of structural and energy modeling may also slightly restrict the generalizability of these results to any type of performance objective. These simulation models were discussed at length in Section 6.3 and Section 6.4, but in short, there are natural hierarchies that arise between the importance of performance objectives, when they can be most influenced, and what variables are prudent to consider in early stage design. This study included an abstracted structural form that can transition between efficient arching behavior and less efficient bending behavior. Since many participants found hybrid solutions between clear building typologies, it might have been difficult to find high-performing designs in an absolute sense. At the same time, it was likely clear to participants that they had a substantial range of structural performance within the design space. In contrast, the energy performance did not change with such great magnitudes across the potential geometries of the study, despite using standard efficiency model settings rather than high-performance settings to magnify geometric effects. While geometry did noticeably influence all objectives, participants were given no specific guidance about these relationships, and were left to prioritize and set performance targets on their own. In an actual design setting, those involved would have prior knowledge of these relationships, and make more active decisions about when and where certain objectives should be considered.

In addition, the study scope was mostly restricted to answering the question of whether or not these computational tools positively affect design outcomes. The results do indicate some clear benefits of performance-driven workflows. However, it is more difficult to answer whether the time and effort required to set up such design environments are worth these benefits, especially in the short term when they require specialized knowledge. While the tools described in the previous chapter aim to make the process of creating a simulation-driven environment more accessible, it is still time-consuming at present, especially when constraints need to be formulated manually and data needs to be carefully scaled and normalized, as they were for the automated optimization phase. Yet the time and level of engagement required to build surrogate models and plug into interactive optimization environments is only slightly beyond what is currently required to generate the design options themselves. As designers continue to learn about new tools and machine learning methods, this gap between creating catalogs and creating interactive models will continue to get smaller, increasing the ratio of benefit to time cost inherent in creating interactive environments.

7.9 Conclusions

7.9.1 Contributions towards understanding interactive design processes

This chapter describes a user study involving 34 participants engaging with an early stage design task, before answering questions about their design workflow preferences. The study demonstrates that the population of architecture and engineering design students tended to prefer flexible design environment that allowed for more creative control and freedom over design options, when compared to an automated environment. At the same time, the results indicate that performance feedback and optimization tools significantly improved the simulated performance of selected designs when compared to the standard parametric environment in most cases. Together, this new data encourages the further development of design approaches that seek to integrate human intuition along with computational feedback and guidance. Finally, the study offers information concerning the current capabilities and desired design methods of contemporary designers, which can help researchers and developers create more effective tools.

7.9.2 Future work and concluding remarks

The results of this study are primarily a first step towards understanding the effects of gradient-based interactive optimization on early design processes. There is much still to be learned about how designers use computational tools in general. Future work will study the direct impact of other design approaches proposed in this dissertation to give a fuller picture of an integrated approach to data-driven design throughout initial ideation, global exploration, and local optimization. As mentioned in the discussion, design studies in this area should also be extended to professional populations, on both design prompts and real projects. Another area for further investigation is how such computational tools can be used collaboratively by multiple users, which will require a significantly different design methodology.

In a few cases, participants themselves commented on other directions for tool development and exploration of design approaches. For example, one participant suggested using linear combinations of objectives for stepping, which could provide users with even more control and robust navigation of tradeoffs, which corresponds to proposed future work in Chapter 4. It may also be worthwhile to test interactive methods along with formal constraints, which could lead to a satisfying design experience on certain problems. While these and other suggestions could be continually added to the tools and methods, this chapter provides strong evidence that interactive optimization techniques are worth pursuing further, as they offer the potential for designers to more effectively integrate quantitative and qualitative design objectives through a natural, creative, synthetic design process.

8 Conclusions

This dissertation has presented new approaches that combine data science and interactive optimization to enable multi-objective exploration during early building design. It first motivates the need for integrating simulation more effectively into natural design processes through a literature review, before presenting a vision of designer-computer collaboration during parametric design. It then describes specific goals for improving on the current state of the art in performance-driven conceptual design and proposes strategies for design space formulation, interactive optimization, and diversity-based ideation. Finally, it presents an integrated design example and a design study that measures the effect of interactive, data-enhanced computational strategies on the conceptual design process. In conclusion, this chapter summarizes contributions, discusses remaining questions, and offers final comments.

8.1 Research questions and dissertation goals

The broad research question motivating this dissertation is: how can creative architects and engineers use computation to navigate multiple performance objectives as part of a natural human design process? This question stems from the intersection of three trends: a need for buildings that perform well in relation to human occupants and the environment; increasing computational competency among design architects and engineers; and developments in data science and optimization. In this context, a series of goals were proposed to improve on existing computational methodologies: making parametric exploration methods

more meaningful and attainable, flexible and interactive, and directed and divergent. The following contributions together address these interconnected research goals.

8.2 Summary of contributions

8.2.1 Design space formulation

In the area of formulation, this dissertation demonstrates initial workflows for determining variable importance during model setup, which is a new application of a traditional data science technique that can assist designers during the difficult, iterative process of setting up an appropriate parametric design space. It also provides methods for generating synthetic, dimensionally reduced design variables, and tests these methods on a series of case studies ranging from individual building components to overall massing and structural system development. These synthetic variables can stimulate creativity by proposing new ways of interacting with a geometry that are not entirely dependent on the initial manual parameterization. In many cases, these variables provide valuable insight into trends in the design space while allowing some direct control over performance. In addition, this dissertation proposes a workflow for automatically generating performance-based design variables for commonly used architectural typologies, which an initial step towards automated design parameterization. With further work, such automation can free designers to spend more time interacting with possibilities rather than manually coding parametric relationships.

8.2.2 Interactive exploration

Building on research into human-in-the-loop optimization processes, this dissertation proposes a gradient-based interactive method for generalized, multi-objective, conceptual design exploration. While exploring locally in a design space, this method includes functionality for moving in directions that improve various performance objectives while also offering the possibility of moving in isoperformance directions for increased design flexibility. When used in conjunction with surrogate models or other fast evaluations of performance, this methodology allows a designer to simultaneously move directly in both the design and objective spaces, while evaluating the geometric and performance implications of every design decision. Rather than approaches that are discipline specific, interactive design space stepping takes advantage of the connection between parametric environments and numerous performance simulations in the early design stages, enabling a multi-objective and often multi-disciplinary lens at this point in the process.

8.2.3 Diversity-driven design

The chapter on diversity in design first consolidates and synthesizes techniques for the numerical measurement of diversity from various fields, and extracts relevant metrics for architects and engineers who want to consider performance and visual appearance together through computation. It then establishes agreement between visual perception and digital calculation of diversity at the thresholds of difference that

become useful for early design ideation. In addition, it offers a comparison of five different diversity metrics based on geometric relationships between parametric design vectors, and clarifies situations in which the metrics may be more or less useful for computational design. Finally, it provides examples for researchers looking to integrate diversity metrics into generative design processes.

8.2.4 Workflow demonstration and user testing

The demonstration and discussion of the toolbox approach reveals how data science and optimization components can be connected in new ways to enable data-driven parametric design methods. It explains these approaches using an example multi-objective conceptual design problem, showing the integration of multiple techniques proposed throughout Chapters 3-5. Through investigation of a realistic early design prompt, it reveals strengths, limitations, important issues, potential pitfalls, and future opportunities for data-driven parametric design. It also compares the novel and overlapping functionality of the tools developed in accordance with this dissertation to existing plug-ins for parametric exploration and optimization.

The user study builds on this demonstration by testing and validating the effects of various environments on the conceptual design processes of both architects and engineers. The extensive design study involving 34 participants yields significant new data on the design histories, performance outcomes, and tool preferences of computational designers. In general, the results reveal that surrogate-model-enabled feedback, interactive optimization, or both often improve the simulated performance of selected designs when compared to a traditional parametric environment. However, designers generally preferred the creative freedom of dynamic design methods compared to using automated optimization at this stage. Although tools and interfaces in this area must be refined, the interactive optimization methods explored in this dissertation can potentially bridge these two realities.

8.3 Potential applications and impact

Since they are generalized and agnostic to specific objective functions, these methods are viable for any designers who interact with geometry, simulation, multiple performance objectives, and have a desire to integrate qualitative and quantitative goals. In most cases the methods were conceived with reference to the exploration of early massing and structural systems for buildings, which generally share a scale, a software environment, and rough simulation times in early stage design. However, these methods could also be used for product design, mechanical or aerospace design, urban design, and other related disciplines in engineering and the built environment. The methods can transition between scales as well, depending on the resolution of simulations required to describe a system, which means they can help create a structural connection, lay out a neighborhood, or everything in between. They are particularly well suited to problems that are parameterized geometrically and visualized accordingly.

8.3.1 In design education

In some cases, the approaches in this dissertation are likely to be first adopted by students for both research and design projects, due to their general affinity for adopting new computation methods. However, these approaches can also be directly used as teaching aides to communicate the structural and energy behavior of buildings; common tradeoffs in design, such as embodied and operational energy; and the principles of optimization. When demonstrated on relevant examples, they can also facilitate discussion between students of architecture and engineering, who do not always overlap during their educations but need to work together effectively while designing buildings.

8.3.2 In practice

Aside from a few of the speculative case studies and demonstrations, most of the proposed workflows in this dissertation have been developed into free, open-source design tools. As such, designers in architecture or engineering offices can begin using these methods right away to integrate simulation more effectively into their early design approaches. For some, designing with data might represent a fundamentally new way of thinking about buildings and early ideation, which could lead to more solutions being meaningfully considered in conceptual design. For others already engaged with parametric design space exploration, the advanced methods here might allow increased flexibility and the ability to move to non-traditional areas of the design space with confidence that performance is still being factored into the process. At minimum, such tools can supplement useful conversations between various stakeholders during design, whether these conversations include live manipulation of design possibilities or discussion over options generated with data-driven methods. Ultimately, the mass adoption of tools that support design decision-making with data concerning multi-disciplinary outcomes can facilitate more effective integrated design and lead to higher performance buildings.

8.4 Remaining questions and future work

Moving beyond the contributions of the dissertation, this section succinctly consolidates the specific discussions regarding immediate next steps found at the end of each chapter. It then moves to broader questions concerning the future use of data-driven approaches in the design and delivery of high-performance buildings.

8.4.1 Next steps of specific methodologies

The following next steps for future work are provided for each task or concept related to early design exploration:

Variable Transformation and Generation (Chapter 3)

- Discovering more robust and systematic ways for segmenting and sampling data to create meaningful coefficients using data analysis techniques
- Implementing and testing automatic variable generation on additional problems involving nodal locations of trussed members, control points of surface structures, and urban massing
- Exploring non-linear or dynamic mapping of coefficients to test their effect on the design space

Interactive Optimization (Chapter 4)

- Improving the algorithm for selecting isoperformance directions during local movement, while also considering how to adequately handle the curse of dimensionality as design variables are added
- Providing more freedom in the controlling step direction, including linear combinations of objectives or continuous vector bisection for bi-objective problems
- Integrating additional visualizations of gradient information into a fully developed parametric software tool, including live gradient visualization in the geometric design space and nearby potential outcomes
- Continuing to benchmark and compare the design outcomes in terms of performance and creative freedom when compared to fully automated methods (a continuation of Chapter 7)

Diversity-driven Design (Chapter 5)

- Quantifying diversity within computationally generated design options that are not parametric and do not contain an explicit design vector
- Combining the concepts of diversity measurement and interactive isoperformance search, such that a diversity-driven design technique is not as dependent on culling from an original sampling resolution
- Testing metrics on more complex geometries and for expert designers rather than the general human population
- Performing design studies that validate the impact of diversity-based early design approaches

Workflow demonstration and validation (Chapters 6-7)

- Adding ongoing functionality updates for data-driven approaches to parametric design
- Creating components for interactive visualization that connect directly to the previously developed tools, and generalizing to other platforms
- Extending interfaces to facilitate multi-user interaction, conversation, and collaboration
- Further validating the effect of all methods proposed in this dissertation on early stage design processes

In general, the largest additional area for future work involves implementing these techniques on pressing design questions in the built environment where both quantitative and qualitative goals are significant, and simulations can supplement creative design processes. In addition to the case studies demonstrated here, which focused on discretized structures and building massing, these tools can be further applied at different scales, from building components, to a whole building system, to the urban scale. They can also be used in adjacent fields such as product or mechanical design, in both research and practice.

8.4.2 The use of data approaches within a traditional project structure

Given the presence of silos that exist between architecture, engineering, and related disciplines, who precisely uses these tools is still an open question. A number of clear possibilities exist, almost always involving an architect, since they tend to control many of the decisions these conceptual design methods were developed to address. One option is that architects might sketch initial ideas, and then turn them over to a computational specialist in their own office to parameterize and generate options. The computational specialist might interface directly with engineers, who can assist in creating useful simulations and mapping the design space possibilities to an objective space. The process can be iterative and conversational, such that architect and engineer are passing back and forth flexible models, which each entity can explore, rather than rigid geometries or full simulation results. Some design firms around the world are already implementing this general process, but it could always be improved.

Another approach might involve a similar transfer of ideas from architect to engineer at the beginning, such that an architect hands off a concept to an engineer, who then fleshes out the model to include performance considerations. However, the next step would be to manipulate the geometry live in meetings, while trying to collectively capture many considerations during conceptual design. This interaction of full human teams with a conceptual model could even be done remotely, though web-enabled means, and involve voting, live manipulation, or other techniques for arriving at a satisfying solution.

In some cases, it may be possible for architects themselves to learn enough about the simulations to control the entire process, from sketching through parametric model formulation, simulation generation, and interactive optimization. A single designer (or smaller group) might prefer this level of autonomy during very early design. In the future, various parts of this process might be automated, such that designers spend most of their precious time modifying geometry and considering possibilities rather than coding.

Regardless of the preferences of different architects, it is generally observed that these methods are valuable when decisions are made. Speaking practically, data-driven approaches are of particular use when one team member responsible for the laborious parts of the design process must present options to a person who controls the design decisions, either by virtue of their status within a firm or their legal responsibility to the process. Data-driven models can justify choices, suggest new solutions, and generally provide more confidence that higher-performance outcomes are being achieved. Although this dissertation focuses on

conceptual design, Figure 8.1 below describes examples of when such an exchange occurs throughout the building delivery process, between architects, engineers, contractors, clients, and others. In many of these cases, having a flexible, data-driven model that considers both quantitative and qualitative objectives can enrich early design.

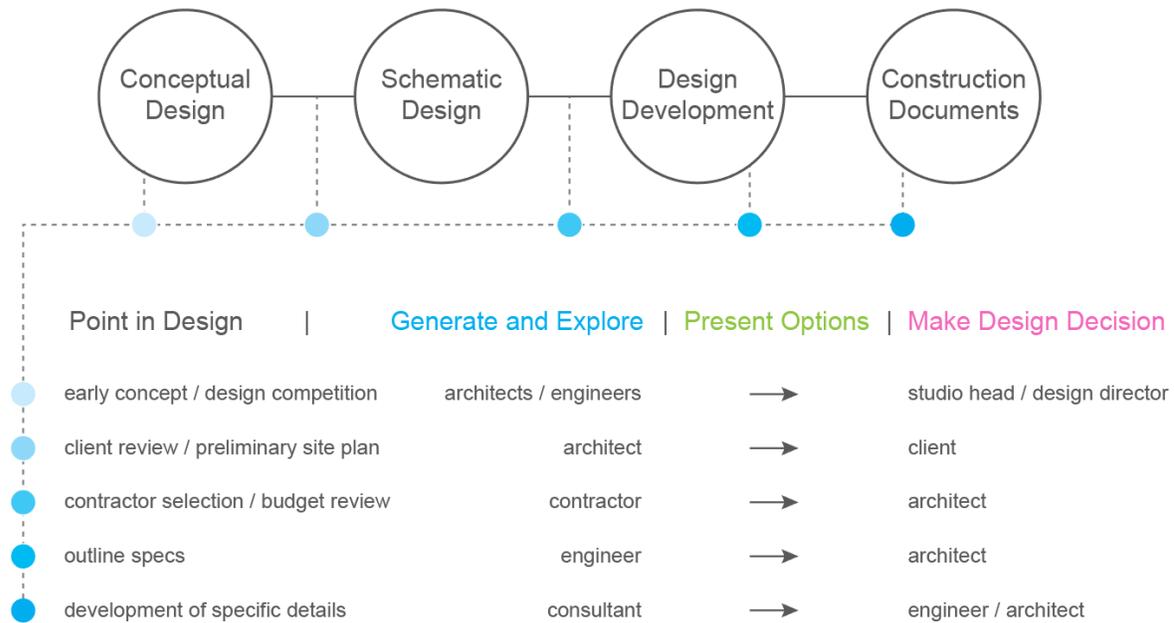


Figure 8.1: Examples of when design space exploration techniques may be useful throughout the building delivery process as software becomes better integrated

8.4.3 Natural hierarchies between performance domains in building design

While this dissertation does not systematically identify the relative importance of objectives, and instead presents workflows and tools that are agnostic to the type of quantitative metric, experienced designers already understand that not all objectives are equal. This reality has implications for how multi-objective data approaches fit into the broader design process. In some situations, it may be prudent to run simultaneous parallel simulations in structure, daylight, energy, acoustics, views, or other objectives, and use them all to enable one large interactive design. This approach is certainly now possible with current parametric software. However, it may not always be the most logical. In plenty of real-world applications, a designer might initiate interactive design methods for one or more driving objectives, before gradually adding in others, or only run certain simulations after the form is largely set. To some extent, these decisions about workflow sequence are particular to the building, and do not always translate from one project to the next.

Yet at the same time, there are common trends in building design within relationships between the main performance objectives that are universally considered. Certain objectives often trade off, harmonize, or require decisions in one domain that require adjustments to the model of the other. Engineers and other specialists have discovered increasingly efficient methods for simultaneously navigating the needs of structural, mechanical, lighting, plumbing, and other systems, particularly during design development and construction. Now that designers are becoming more comfortable with using common software and languages for parametric design exploration, they should develop similar intuitions for how the natural hierarchies between these disciplines play out in early design exploration as well.

8.4.4 Further validation of data-driven design methodologies

In various ways, this dissertation sought to validate the proposed concepts and design methodologies on actual humans, to provide insights into how they interact with data-driven computational methodologies. The design study in Chapter 7 was the largest effort in this regard, and demonstrated the effects of interactive, multi-objective, gradient-based optimization within parametric design. For this study, an intentional decision was made to carefully analyze the interactive framework to other traditional parametric methods while capturing its specific effects, rather than attempt to combine all proposed approaches simultaneously and risk designer fatigue, difficulty separating influencing factors, or other problems that could have arisen with such an involved study. Chapter 5 also internally contained a study involving visual human perception of diversity, which is a form of validation of the computational metrics that chapter proposes.

However, many opportunities remain for further design studies in this area, even directly related to the proposals in this dissertation. Further study and testing are required for the design workflow possibilities for transforming the design space with synthetic variables, especially if these variables are automatically generated. Studies should also be conducted in a more practice-oriented setting, which could lead to different results than testing these methodologies on graduate students, even if many of them have professional experience. Future studies might also capture the design processes of entire teams who are engaging with such data-driven methods, perhaps through innovative technological solutions like a digital model shared between multiple computers or interfaces, rather than individuals who are in full control over decisions. Such explorations could involve analysis of team dynamics and other factors at play that might be specific to collaborative design processes involving both quantitative and qualitative drivers.

8.4.5 Connections to construction and fabrication

The architectural and building engineering communities are currently engaged in exciting research at the intersection of optimization and fabrication. Their basic argument is that historically, many optimized solutions may have been discarded because they were too difficult to build. Constant advancements in additive manufacturing (Jared et al. 2017) and digital fabrication, however, offer the potential to more easily

build new, innovative, non-traditional, yet high-performance components (Chang & Tang 2001; Mayencourt et al. 2017; Ismail & Mueller 2019; Jewett & Carstensen 2019) and forms (Beghini et al. 2015; Tam et al. 2018; Huang et al. 2018).

The data-driven, multi-objective approaches in this dissertation intend to enable such design possibilities, and also catalyze their more widespread usage in architectural and engineering practice. Given the prevailing interest in this topic, the development and realization of creative new geometries should be expected in the future. Especially when such data-driven methods lead to non-traditional solutions, future efforts should be made to physically test design outcomes at various scales to ensure the value of simulation methods in the early stages of design.

8.4.6 Progressive, dynamic, and generalizable surrogate modeling

In this dissertation, the live design approaches generally relied on static surrogate models. Although they could be constructed iteratively, and almost always used an ensemble method, the predictive models for each case study were built at the beginning of the exploration and then not modified. Furthermore, anytime the structure of a parametric model changes, the design space must be resampled and a new surrogate model created. If the changes to the parametric model involved transformations of variables that could be mapped back to the original variables used to build the data model, re-simulation can be avoided, but this is not always the case.

However, there are significant improvements on the horizon for surrogate modeling techniques that are either dynamic or generalizable. Dynamic surrogate models would likely run parallel simulations during exploration in the background and continue updating the accuracy of predictions. These models could focus on areas of the design space that are selected by the designer, effectively learning where better accuracy is needed. In a similar concept to the idea proposed here for automated parameterization through reduction of initial dummy variables, there are also efforts to generate models with fundamental parameters, such that designers can reuse data rather than building a new dataset every time the problem changes slightly. These would likely be specific to certain disciplines, such as finding and modeling the fundamental relationships that inform computational structural analysis. Such developments could drastically improve the use of surrogate-model enhanced interactive design, since designers would no longer have to wait for training data to be generated if the problem changes.

8.4.7 The role of simulation compared to intuition and experience

As mentioned previously, these tools are not a substitute for experienced designers who understand the technical principles on which a simulation is based. Instead, they are meant to supplement human design processes and enhance design conversations. By focusing on relative design decisions between options, especially within a parametric framework, there is a risk that users might become satisfied that a solution is

good, without a broader contextual understanding of different design options outside the set of parametric choices. Any similar design tool might be misused in the wrong hands, if simulations are not properly understood and applied. Finding the best methods for training designers on how to apply technical principles and engage with computers during the design process requires a broader educational, social, and cultural discussion. Nevertheless, when done appropriately, the infusion of data-rich feedback and guidance into multi-objective early design processes can uncover new design directions and performance trends in the design space, which significantly improves on the traditional approach of sketching and then individually modeling a few discrete options for comparison.

8.5 Final remarks and outlook

Ongoing advances in computation are driving massive changes across industries, as well as human endeavors more generally. At the same time, design remains a central human activity—not just because some of us do it better than computers at this point, but because it captures our imaginations and expresses our values. While it is imperative to respond to legitimate social, environmental, and economic pressure to design buildings that perform well, the design community should at the same time assert the importance of other non-quantifiable qualities during design.

This dissertation considers ways to take advantage of improvements in computational speed and power along with developments in data science and artificial intelligence, without losing the human aspects of the design process. Interactive optimization methods, or collaborative designer-computer methods more generally, have the potential to improve building performance while at the same time inspiring creative and novel design solutions that push architecture and engineering forward.

9 References

- Abaza, H., 2010. An interactive design advisor for energy efficient buildings. *Journal of Green Building*, 3(1), pp.112–125.
- Acroterion, 2012. Photo of the National Portrait Gallery Courtyard. Available at: commons.wikimedia.org/wiki/File:National_Portrait_Gallery_courtyard_DC1.jpg.
- Aggarwal, C.C., Hinneburg, A. & Keim, D.A., 2001. On the surprising behavior of distance metrics in high dimensional space. *ICDT*, 1, pp.420–434.
- Agresti, A. & Agresti, B.F., 1978. Statistical analysis of qualitative variation. *Sociological methodology*, 9, pp.204–237.
- Amabile, T.M., 1982. Social psychology of creativity: a consensual assessment technique. *Journal of Personality and Social Psychology*, 43(5), pp.997–1013.
- Anderl, R. & Mendgen, R., 1995. Parametric design and its impact on solid modeling applications. In *Proceedings of the third ACM symposium on solid modeling and applications*. pp. 1–12.
- Andersen, M., Guillemin, A., et al., 2013. Beyond illumination: an interactive simulation framework for non-visual and perceptual aspects of daylighting performance. In *BS2013 - 13th International Conference of the International Building Performance Simulation Association*. Chambéry, France.
- Andersen, M., Gagne, J.M.L. & Kleindienst, S., 2013. Interactive expert support for early stage full-year daylighting design: a user's perspective on Lightsolve. *Automation in Construction*, 35, pp.338–352.
- Architecture - 2030, 2006. Architecture 2030 Challenge. Available at: architecture2030.org/2030_challenges/2030-challenge/.
- Arnaud, V.B., 2013. *Quantifying architects' and engineers' use of structural design software*. M.Eng. Thesis, Massachusetts Institute of Technology.
- Asadi, S., Amiri, S.S. & Mottahedi, M., 2014. On the development of multi-linear regression analysis to assess energy consumption in the early stages of building design. *Energy and Buildings*, 85, pp.246–255.
- Asadi, S. & Geem, Z.W., 2015. Sustainable building design: a review on recent meta-heuristic methods. In X.-S. Yang, ed. *Recent Advances in Swarm Intelligence and Evolutionary Computation. Studies in Computational Intelligence*. Springer, pp. 203–223.
- Ashour, Y., 2015. Optimizing creatively in multi-objective optimization. In *Proceedings of the Symposium on Simulation for Architecture & Urban Design. Society for Computer Simulation International*. Alexandria, VA, pp. 128–135.
- Attia, S. et al., 2013. Assessing gaps and needs for integrating building performance optimization tools in

- net zero energy buildings design. *Energy and Buildings*, 60, pp.110–124.
- Attia, S. et al., 2012. Simulation-based decision support tool for early stages of zero-energy building design. *Energy and Buildings*, 49, pp.2–15.
- Autodesk, 2011. Dynamo. Available at: dynamobim.org.
- Autodesk, 2016. Robot Structural Analysis Professional. Available at: www.autodesk.com/products/robot-structural-analysis/overview.
- Balali, V. et al., 2018. Improved stakeholder communication and visualizations: real-time interaction and cost estimation within immersive virtual environments. *Construction Research Congress 2018*, pp.522–529.
- Balling, R., 1999. Design by shopping: a new paradigm? In *Proceedings of the Third World Congress of structural and multidisciplinary optimization (WCSMO-3)*. Berlin: International Soc. for Structural and Multidisciplinary Optimization.
- Bao, Q., Faas, D. & Yang, M., 2018. Interplay of sketching & prototyping in early stage product design. *International Journal of Design Creativity and Innovation*, 6(3–4), pp.146–168.
- Baril, C., Yacout, S. & Clément, B., 2013. Collaborative decision-making in product design: an interactive multiobjective approach. *Journal of Industrial Engineering*, pp.1–12.
- Beghini, L.L. et al., 2015. Integrated discrete/continuum topology optimization framework for stiffness or global stability of high-rise buildings. *Journal of Structural Engineering*, 141(8), pp.1–10.
- Beghini, L.L. et al., 2014. Structural optimization using graphic statics. *Structural and Multidisciplinary Optimization*, 49(3), pp.351–366.
- Belsley, D.A., Kuh, E. & Welsch, R.E., 1980. *Regression diagnostics: identifying influential data and sources of collinearity*, John Wiley & Sons.
- Bendsøe, M.P. & Kikuchi, N., 1988. Generating optimal topologies in structural design using a homogenization method. *Computer Methods in Applied Mechanics and Engineering*, 71(2), pp.197–224.
- Berkhin, P., 2006. Survey of clustering data mining techniques. In J. Kogan, C. Nicholas, & M. Teboulle, eds. *Grouping Multidimensional Data*. Berlin: Springer, pp. 25–71.
- Beyer, H. & Schwefel, H., 2002. Evolution strategies: a comprehensive introduction. *Natural Computing*, 1, pp.3–52.
- Block, P., 2009. *Thrust network analysis: exploring three-dimensional equilibrium*. Ph.D. Thesis, Massachusetts Institute of Technology.
- Bosman, P.A.N. & Jong, E.D. De, 2005. Exploiting gradient information in numerical multi-objective evolutionary optimization. In *Proceedings of the 7th annual conference on genetic and evolutionary computation*. ACM, pp. 755–762.
- Box, G.E.P., Hunter, J.S. & Hunter, W.G., 2005. *Statistics for experimenters: design, innovation, and*

- discovery, 2nd Edition*, John Wiley & Sons.
- Bradner, E., Iorio, F. & Davis, M., 2014. Parameters tell the design story: ideation and abstraction in design optimization. *Simulation Series*, 46(7), pp.172–197.
- Brady, J.T., 1986. A theory of productivity in the creative process. *IEEE Computer Graphics and Applications Magazine*, pp.25–34.
- Brown, G.Z. & DeKay, M., 2001. *Sun, wind & light*, New York: John Wiley & Sons.
- Brown, N.C., 2016a. *Integrating secondary goals into structural design*, SOM Structural Engineering Travel Fellowship Report. Available at: somfoundation.som.com/fellow/nathan-brown.
- Brown, N.C., 2016b. *Multi-objective optimization for the conceptual design of structures*. SMTB Thesis, Massachusetts Institute of Technology.
- Brown, N.C. & Mueller, C.T., 2017a. Automated performance-based design space simplification for parametric structural design. In *Proceedings of the IASS Annual Symposium 2017*. Hamburg, Germany.
- Brown, N.C. & Mueller, C.T., 2016. Design for structural and energy performance of long span buildings using geometric multi-objective optimization. *Energy and Buildings*, 127, pp.748–761.
- Brown, N.C. & Mueller, C.T., 2017b. Designing with data: moving beyond the design space catalog. In *Proceedings of ACADIA 2017: Disciplines and Disruption*. Boston, MA.
- Brown, N.C. & 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 (IASS) Symposium 2016*. Tokyo.
- Brown, N.C., Tseranidis, S. & Mueller, C.T., 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.
- Brownlee, A.E.I. & Wright, J.A., 2015. Constrained, mixed-integer and multi-objective optimisation of building designs by NSGA-II with fitness approximation. *Applied Soft Computing Journal*, 33, pp.114–126.
- 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.
- Buhrmester, M., Kwang, T. & Gosling, S.D., 2011. Amazon's Mechanical Turk: a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1), pp.3–5.
- Burnell, E. et al., 2017. Integrating design and optimization tools: a designer centered study. In *Proceedings of the ASME 2017 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. Cleveland, OH.
- Burnell, E., 2014. Tacit. Available at: github.com/bqpd/tacit.
- BuroHappold Engineering, 2018. BHoM. Available at: github.com/BHoM/BHoM.

- Burry, M., 1996. Parametric design and the Sagrada Familia. *arq: Architectural Research Quarterly*, 1(4), pp.70–81.
- 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.
- Chang, K. & Tang, P., 2001. Integration of design and manufacturing for structural shape optimization. *Advances in Engineering Software*, 32, pp.555–567.
- Chaszar, A. & Joyce, S.C., 2016. Generating freedom: questions of flexibility in digital design and architectural computation. *International Journal of Architectural Computing*, 14(2), pp.167–181.
- Chen, K.W., Janssen, P. & Schlueter, A., 2015. Analysing populations of design variants using clustering and archetypal analysis. *Proceedings of the 33rd eCAADe Conference*, 1, pp.251–260.
- Chow, L. & Sarkisian, M., 2009. SOM | Super Tall Buildings | Venues, Structures and Construction. Available at: wemp.app/posts/c37ca648-c0e6-41df-a49f-0545f50a798c.
- Cichocka, J. et al., 2017. SILVEREYE - the implementation of Particle Swarm Optimization algorithm in a design optimization tool. In G. Çağdaş et al., eds. *Computer-Aided Architectural Design. Future Trajectories. CAADFutures 2017. Communications in Computer and Information Science*. Singapore: Springer Science & Business Media.
- Cichocka, J.M., Browne, W.N. & Rodriguez, E., 2017. Optimization in the architectural practice. In *CAADRIA 2017 - 22nd International Conference on Computer-Aided Architectural Design; Research in Asia: Protocols, Flows and Glitches*. pp. 387–396.
- Clune, R. et al., 2012. An object-oriented architecture for extensible structural design software. *Computers and Structures*, 100–101, pp.1–17.
- Clune, R.P., 2010. *Explorative structural design*. S.M. Thesis, Massachusetts Institute of Technology.
- Coello Coello, C.A. & Mariano Romero, C.E., 2013. Evolutionary algorithms and multiple objective optimization. In M. Ehrgott & X. Gandibleux, eds. *Multiple Criteria Optimization: State of the Art, Annotated Bibliographic Surveys*. Boston, MA: Springer.
- Coley, D.A. & Schukat, S., 2002. Low-energy design: combining computer-based optimisation and human judgement. *Building and Environment*, 37(12), pp.1241–1247.
- Colley, W.N., 2002. *Colley's bias free college football ranking method: The Colley matrix explained*, Princeton, NJ.
- Conti, Z.X. & Kajijima, S., 2017. Enabling inference in performance-driven design exploration. In K. De Rycke et al., eds. *Humanizing Digital Reality: Design Modelling Symposium*. Paris, pp. 177–188.
- Cornwell, W.K., Schwilk, D.W. & Ackerly, D.D., 2006. A trait-based test for habitat filtering: convex hull volume. *Ecology*, 87(6), pp.1465–1471.
- Cross, N., Dorst, K. & Christiaans, H., 1996. *Analysing design activity*, Wiley.
- Cross, N., Dorst, K. & Roozenburg, N. eds., 1992. *Research in design thinking*, Delft, The Netherlands: Delft

University Press.

- Cross, N., Naughton, J. & Walker, D., 1981. Design method and scientific method. *Design Studies*, 2(4), pp.195–201.
- Crystal Bridges, 2013. Photo of the Tulip Tree Shelter. Available at: crystalbridges.org/blog/art-and-architecture-the-tulip-tree-shelter-2/.
- Csikszentmihalyi, M., 1996. *Flow and the psychology of discovery and invention*, New York: Harper Collins.
- Cvetkovic, D. & Parmee, I.C., 2002. Preferences and their application in evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 6(1), pp.42–57.
- D’Acunto, P. et al., 2019. Vector-based 3D graphic statics: a framework for the design of spatial structures based on the relation between form and forces. *International Journal of Solids and Structures*.
- Dalbera, J.-P., 2013. Photo of the Neue Nationalgalerie. Available at: [commons.wikimedia.org/wiki/File:La_nouvelle_galerie_nationale_\(Berlin\)_\(11478130635\).jpg](https://commons.wikimedia.org/wiki/File:La_nouvelle_galerie_nationale_(Berlin)_(11478130635).jpg).
- Danhaive, R.A. & Mueller, C.T., 2015. Combining parametric modeling and interactive optimization for high-performance and creative structural design. In *Proceedings of the IASS Annual Symposium 2015*. Amsterdam.
- Danhaive, R. & Mueller, C., 2018. Structural metamodelling of shells. In *Proceedings of the IASS Symposium 2018*. Boston, MA.
- Das, I. & Dennis, J.E., 1998. Normal-Boundary Intersection: a new method for generating the Pareto surface in nonlinear multicriteria optimization problems. *SIAM Journal on Optimization*, 8(3), pp.631–657.
- Daskilewicz, M.J. & German, B.J., 2013. Nondomination-level coordinate system for parameterizing and exploring Pareto frontiers. *AIAA Journal*, 51(8), pp.1946–1959.
- 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.
- de Weck, O.L., Miller, D.W. & Mosier, G.E., 2002. Multivariate isoperformance methodology for precision opto-mechanical systems. In *43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. Denver, CO.
- De Wolf, C. et al., 2016. Material quantities and embodied carbon dioxide in structures. *Proceedings of the Institution of Civil Engineers: Engineering Sustainability*, 169(4), pp.150–161.
- De Wolf, C., 2016. Database of embodied Quantity outputs (DeQo). Available at: www.carbondeqo.com [Accessed April 10, 2019].
- De Wolf, C., 2017. *Low carbon pathways for structural design: embodied life cycle impacts of building structures*. Ph.D. Thesis, Massachusetts Institute of Technology.
- Dean, D.L. et al., 2006. Identifying good ideas: constructs and scales for idea evaluation. *Journal of Association for Information Systems*, 7(10), pp.646–699.

- 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.
- Derix, C. & Jagannath, P., 2014. Near futures: associative archetypes. *Architectural Design*, 84(5), pp.130–135.
- Design, 2017. Design. *Cambridge Dictionary of American English*.
- Dogan, T., 2014. ARCHSIM. Available at: www.solemma.com.
- Doherty, W.J. & Kelisky, R.P., 1979. Managing VM/CMS systems for user effectiveness. *IBM Systems Journal*, 18(1), pp.143–163.
- Dorst, K., 2011. The core of “design thinking” and its application. *Design Studies*, 32(6), pp.521–532.
- Dorst, K. & Cross, N., 2001. Creativity in the design process: co-evolution of problem–solution. *Design Studies*, 22, pp.341–361.
- Dorta, T., Lesage, A. & Perez, E., 2009. Design tools and collaborative ideation. *Joining Languages, Cultures and Visions: CAADFutures*, pp.65–79.
- Durillo, J.J. & Nebro, A.J., 2011. jMetal: a Java framework for multi-objective optimization. *Advances in Engineering Software*, 42, pp.760–771.
- Eberhart, R. & Kennedy, J., 1994. Particle swarm optimization. *Proceedings of the IEEE international conference on neural networks*, 4, pp.1942–1948.
- EIA, U.S., 2017. How the United States uses energy. Available at: https://www.eia.gov/energyexplained/index.php?page=us_energy_use.
- Ericsson, K.A. & Simon, H.A., 1984. *Protocol analysis: verbal reports as data*, Cambridge, MA: The MIT Press.
- Etesco, 2002. modeFRONTIER. Available at: www.esteco.com/modelfrontier.
- Evins, R., 2013. A review of computational optimisation methods applied to sustainable building design. *Renewable and Sustainable Energy Reviews*, 22, pp.230–245.
- Evins, R., 2015. Multi-level optimization of building design, energy system sizing and operation. *Energy*, 90, pp.1775–1789.
- Evins, R. et al., 2012. Multi-objective design optimisation: getting more for less. *ICE Proceedings Civil Engineering*, 165(CE5), pp.5–10.
- Faas, D. et al., 2014. The influence of immersion and presence in early stage engineering designing and building. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, (Special Issue 02), pp.139–151.
- Falkenpost, 2016. Photos of Dulles International Airport. Available at: pixabay.com/photos/dulles-airport-building-1651554/.

- 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.
- Fisher, A., 2012. Engineering integration: real-time approaches to performative computational design. *Architectural Design*, 82(2), pp.112–117.
- Fivet, C. & Zastavni, D., 2013. Constraint-based graphic statics: new paradigms of computer-aided structural equilibrium design. In *Proceedings of the International Association for Shell and Spatial Structures (IASS) Symposium 2013*. Wroclaw University of Technology, Poland.
- Flager, F., Adya, A. & Haymaker, J., 2009. AEC multidisciplinary design optimization: impact of high performance computing. *CIFE Technical Report #TR175*.
- Flöry, S., 2013. Goat. Available at: www.rechenraum.com/goat/download.html.
- Ford, C. & de Vere, N., 2010. Photo of the Great Glasshouse. Available at: commons.wikimedia.org/wiki/File:Great_Glass_House.jpg.
- Freitas, S., Serra, F. & Brito, M.C., 2015. Pv layout optimization: string tiling using a multi-objective genetic algorithm. *Solar Energy*, 118, pp.562–574.
- Fuhrmann, L. et al., 2018. Data-driven design: exploring new structural forms using machine learning and graphic statics. In *Proceedings of the IASS Symposium 2018*. Boston, MA.
- Gane, V., 2004. *Parametric Design - a Paradigm Shift?* S.M.Arch.S. Thesis, Massachusetts Institute of Technology.
- Genometri Ltd., 2013. Genoform. Available at: www.food4rhino.com/app/genoform.
- Gerber, D.J. & Lin, S.H.E., 2014. Designing in complexity: simulation, integration, and multidisciplinary design optimization for architecture. *Simulation*, 90(8), pp.936–959.
- Gerber, D.J., Pantazis, E. & Marcolino, L.S., 2015. Design agency prototyping multi-agent systems in architecture. In *CAAD Futures 2015, CCIS 527*. Sao Paulo, Brazil, pp. 213–235.
- Gerber, D.J., Pantazis, E. & Wang, A., 2017. A multi-agent approach for performance based architecture: design exploring geometry, user, and environmental agencies in façades. *Automation in Construction*, 76, pp.45–58.
- Geyer, P. & Schlüter, A., 2014. Automated metamodel generation for Design Space Exploration and decision-making - A novel method supporting performance-oriented building design and retrofitting. *Applied Energy*, 119, pp.537–556.
- Girotra, K., Terwiesch, C. & Ulrich, K.T., 2010. Idea generation and the quality of the best idea. *Management Science*, 56(4), pp.591–605.
- Goh, C.K. et al., 2008. An investigation on evolutionary gradient search for multi-objective optimization. In *Proceedings of the IEEE World Congress on Computational Intelligence*. Hong Kong, pp. 3741–3746.

- Goldschmidt, G., 1994. On visual design thinking: the vis kids of architecture. *Design Studies*, 15(2), pp.158–174.
- Gorunescu, F., 2011. *Data mining: concepts, models and techniques*, Berlin Heidelberg: Springer Science & Business Media.
- Greco, L., 2015. Dodo. Available at: www.food4rhino.com/app/dodo.
- Haftka, R.T. & Grandhi, R. V., 1986. Structural shape optimization-a survey. *Computer Methods in Applied Mechanics and Engineering*, 57, pp.91–106.
- Hagan, M.T., Demuth, H.B. & Beale, M.H., 1996. *Neural network design*, Boston, MA: PWS Pub.
- Häggman, A. et al., 2015. Connections between the design tool, design attributes, and user preferences in early stage design. *Journal of Mechanical Design*, 137(7).
- Hamdy, M., Nguyen, A.T. & Hensen, J.L.M., 2016. A performance comparison of multi-objective optimization algorithms for solving nearly-zero-energy-building design problems. *Energy and Buildings*, 121, pp.57–71.
- Harding, J., 2016. Dimensionality reduction for parametric design exploration. In S. Adriaenssens et al., eds. *Advances in Architectural Geometry 2016*. Zurich, Switzerland, pp. 204–221.
- Harding, J. & Olsen, C.B., 2017. Biomorpher. Available at: www.food4rhino.com/app/biomorpher.
- Harding, J. & Shepherd, P., 2016. Meta-parametric design. *Design Studies*, 52, pp.73–95.
- Hardoon, D.R., Szedmak, S. & Shawe-Taylor, J., 2004. Canonical Correlation Analysis: an overview with application to learning methods. *Neural Computation*, 16(12), pp.2639–2664.
- Hasik, V. et al., 2017. Investigation of the sustainability and resilience characteristics of buildings including existing and potential assessment metrics. In *AEI 2017: Resilience of the Integrated Building - Proceedings of the Architectural Engineering National Conference 2017*. Oklahoma City, OK, pp. 1019–1033.
- Haymaker, J. et al., 2018. Design space construction: A framework to support collaborative, parametric decision making. *Journal of Information Technology in Construction*, 23, pp.157–178.
- Haymaker, J., 2012. Expanding design spaces. In *Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2011 Symposium*. Washington, D.C.: The National Academies Press, pp. 89–96.
- Hill, M.O., 1973. Diversity and evenness: a unifying notation and its consequences. *Ecology*, 54(2), pp.427–432.
- Hofmeyer, H. & Davila Delgado, J.M., 2015. Coevolutionary and genetic algorithm based building spatial and structural design. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 29(4), pp.351–370.
- Holzer, D., Hough, R. & Burry, M., 2008. Parametric design and structural optimisation for early design exploration. *International Journal of Architectural Computing*, 5(4), pp.625–644.

- Horn, D. & Salvendy, G., 2009. Measuring consumer perception of product creativity: impact on satisfaction and purchasability. *Human Factors and Ergonomics in Manufacturing*, 19(3), pp.223–240.
- Hotelling, H., 1936. Relations between two sets of variates. *Biometrika*, 28(3–4), pp.321–377.
- Hoxmeier, J.A. & DiCesare, C., 2000. System Response Time and user satisfaction: an experimental study of browser-based applications. In *AMCIS 2000 Proceedings*. pp. 140–145.
- Huang, Y., Garrett, C.R. & Mueller, C.T., 2018. Automated sequence and motion planning for robotic spatial extrusion of 3D trusses. *Construction Robotics*, 2(1–4), pp.15–39.
- Ismail, M.A. & Mueller, C.T., 2019. A platform of design strategies for the optimization of concrete floor systems in India. In *International Conference on Structures and Architecture*. Lisbon, Portugal.
- Jain, A.K., 2010. Data clustering: 50 years beyond K-means. *Pattern Recognition Letters*, 31(8), pp.651–666.
- Janberg, N., 2015. Photo of the SICLI Company Building. Available at: structurae.net/photos/242468-sicli-company-building.
- Jared, B. et al., 2017. Additive manufacturing: toward holistic design. *Scripta Materialia*, 135, pp.141–147.
- Jeffries, P., 2019. Salamander 3. Available at: www.food4rhino.com/app/salamander-3.
- Jewett, J.L. & Carstensen, J. V, 2019. Topology-optimized design, construction and experimental evaluation of concrete beams. *Automation in Construction*, 102, pp.59–67.
- Johns, R.L., Kilian, A. & Foley, N., 2014. Design approaches through augmented materiality and embodied computation. In W. McGee & M. Ponce de Leon, eds. *Robotic Fabrication in Architecture, Art and Design*. Springer, pp. 305–317.
- Jones, N.L., 2017. *Validated interactive daylighting analysis for architectural design*. Ph.D. Thesis, Massachusetts Institute of Technology.
- Jones, N.L. & Reinhart, C.F., 2018. Effects of real-time simulation feedback on design for visual comfort. *Journal of Building Performance Simulation*, 1493, pp.1–19.
- Jonson, B., 2005. Design ideation: the conceptual sketch in the digital age. *Design Studies*, 26(6), pp.613–624.
- Jost, L., 2006. Entropy and diversity. *Oikos*, 113(2), pp.363–375.
- Kackar, R.N., 1985. Off-line quality control, parameter design, and the Taguchi Method. *Journal of Quality Technology*, 17(4), pp.176–188.
- Kan, J.W.T. & Gero, J.S., 2018. Characterizing innovative processes in design spaces through measuring the information entropy of empirical data from protocol studies. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 32(1), pp.32–43.
- Karhunen, J. & Joutsensalo, J., 1995. Generalizations of Principal Component Analysis, optimization problems, and neural networks. *Neural Networks*, 8(4), pp.549–562.

- Kesik, T. & Stern, L., 2008. Representation of performance indicators for the conceptual design of passive solar houses. In *The 5th IBPSA - Canada Conference*. pp. 49–56.
- Khaled, N. & Smaili, A., 2005. Curve representation using principal component analysis for shape optimization of path generating mechanisms. In *Proceedings of IDETC/CIE 2005 ASME 2005 International Design Engineering Technical Conferences & Computers*. pp. 1–8.
- Khan, F.R. & Rankine, J., 1981. *Tall building systems and concepts*, New York: American Society of Civil Engineering.
- Kilian, A. & Ochsendorf, J., 2006. Particle-spring systems for structural form finding. *Journal of the International Association for Shell and Spatial Structures*, 148(2), pp.77–84.
- Kirkpatrick, S., Gelatt, C.D. & Vecchi, M.P., 1983. Optimization by simulated annealing. *Science*, 220, pp.671–680.
- Klack, G., 2018. Photo of Thompson Ice Arena. Available at: commons.wikimedia.org/wiki/File:Rupert-Thompson-Ice-Arena-Dartmouth-College-Hanover-New-Hampshire-05-2018a.jpg.
- Knerr, N. & Selva, D., 2016. Cityplot: visualization of high-dimensional design spaces with multiple criteria. *Journal of Mechanical Design*, 138(9), pp.1–9.
- Knight, T.W., 1991. Designing with grammars. *CAAD futures Digital Proceedings, 1991*, pp.33–48.
- Kohonen, T., 1990. The self-organizing map. *Proceedings of the IEEE*, 78(9), pp.1464–1480.
- Kokare, M., Chatterji, B.N. & Biswas, P.K., 2003. Comparison of similarity metrics for texture image retrieval. In *TENCON 2003: Conference on Convergent Technologies for the Asia-Pacific Region*. Bangalore, India.
- Kolarevic, B. ed., 2004. *Architecture in the digital age: design and manufacturing*, Taylor & Francis.
- Kollat, J.B. & Reed, P., 2007. A framework for visually interactive decision-making and design using evolutionary multi-objective optimization (VIDEO). *Environmental Modelling & Software*, 22, pp.1691–1704.
- Konstru, 2017. Konstru. Available at: www.konstru.com.
- Kotnik, T., 2010. Digital architectural design as exploration of computable functions. *International Journal of Architectural Computing*, 8(1), pp.1–16.
- Kristensen, M.H. & Petersen, S., 2016. Choosing the appropriate sensitivity analysis method for building energy model-based investigations. *Energy and Buildings*, 130, pp.166–176.
- Kruger, C. & Cross, N., 2006. Solution driven versus problem driven design: strategies and outcomes. *Design Studies*, 27(5), pp.527–548.
- Kudrowitz, B.M. & Wallace, D., 2017. Assessing the quality of ideas from prolific, early-stage product ideation. *Journal of Engineering Design*, 24(2), pp.120–139.
- Layman, C.A. et al., 2007. Can stable isotope ratios provide for community-wide measures of trophic

- structure. *Ecology*, 88(1), pp.42–48.
- Lee, J., Gu, N. & Williams, A.P., 2014. Parametric design strategies for the generation of creative designs. *International Journal of Architectural Computing*, 12(3), pp.263–282.
- Lehman, J. & Stanley, K.O., 2011. Abandoning objectives: evolution through the search for novelty alone. *Evolutionary Computation*, 19(2), pp.189–223.
- Leon, M., Doolan, D.C., et al., 2014. Application of interactive surfaces to support computer mediated collaborative design environment. *Proceedings of the International Conference on Information Visualisation*, pp.281–286.
- Leon, M., Laing, R., et al., 2014. Development and testing of a design protocol for computer mediated multidisciplinary collaboration during the concept stages with application to the built environment. *Procedia Environmental Sciences*, 22, pp.108–119.
- Luebke, C. & Shea, K., 2005. CDO: Computational Design + Optimisation in building practice. *The Arup Journal*, 40(3), pp.17–21.
- Machairas, V., Tsangrassoulis, A. & Axarli, K., 2014. Algorithms for optimization of building design: a review. *Renewable and Sustainable Energy Reviews*, 31(1364), pp.101–112.
- Machwe, A.T., Parmee, I.C. & Miles, J.C., 2005. Integrating aesthetic criteria with a user-centric representation. *International Conference on Engineering Design*, pp.1–15.
- Macomber, B. & Yang, M., 2011. The role of sketch finish and style in user responses to early stage design concepts. In *Proceedings of the ASME 2011 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference*. Washington, D.C.
- 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.
- Marler, R.T. & Arora, J.S., 2010. The weighted sum method for multi-objective optimization: new insights. *Structural and Multidisciplinary Optimization*, 41, pp.853–862.
- Martini, K., 2006. A new kind of software for teaching structural behavior and design. In *Building Technology Educators' Symposium*. pp. 1–10.
- Mathew, P., 2015. Building Performance Database. Available at: bpd.lbl.gov/ [Accessed April 10, 2019].
- Mathew, P.A. et al., 2015. Big-data for building energy performance: lessons from assembling a very large national database of building energy use. *Applied Energy*, 140, pp.85–93.
- Mayencourt, P. et al., 2017. Structural optimization and digital fabrication of timber beams. In *Proceedings of the IASS Annual Symposium 2017*. Hamburg, Germany.
- McCaffrey, T. & Spector, L., 2017. An approach to human–machine collaboration in innovation. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, pp.1–15.

- McCune, B. & Grace, J.B., 2002. *Analysis of Ecological Communities*, Glenden Beach, Oregon, USA: MjM Software.
- McHugh, A., 2017. *Visualizing variable sensitivity in structural design*. S.M. Thesis, Massachusetts Institute of Technology.
- Méndez Echenagucia, T., 2013. *Computational search in architectural design*. Ph.D. Thesis, Politecnico di Torino.
- Michalek, J.J. & Papalambros, P.Y., 2002. Interactive design optimization of architectural layouts. *Engineering Optimization*, 34(5), pp.485–501.
- 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., 2015. Geometry gym. Available at: geometrygym.wordpress.com.
- Monedero, J., 2000. Parametric design: a review and some experiences. *Automation in Construction*, 9, pp.369–377.
- Muehlbauer, M., Burry, J. & Song, A., 2017. Automated shape design by grammatical evolution. In J. Correia, V. Ciesielski, & A. Liapis, eds. *EvoMUSART 2017: Computational Intelligence in Music, Sound, Art and Design*. pp. 217–219.
- 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.
- Myers, R.H., Montgomery, D.C. & Anderson-Cook, C.M., 2009. *Response surface methodology: process and product optimization using designed experiments*, Hoboken, NJ: John Wiley & Sons.
- Nagpal, S. et al., 2019. A methodology for auto-calibrating urban building energy models using surrogate modeling techniques. *Journal of Building Performance Simulation*, 12(1), pp.1–16.
- National Research Council, 1998. *Visionary Manufacturing Challenges for 2020*, Washington, D.C.
- Nault, E. et al., 2017. Predictive models for assessing the passive solar and daylight potential of neighborhood designs: A comparative proof-of-concept study. *Building and Environment*, 116, pp.1–16.
- Nocedal, J. & Wright, S.J., 2006. *Numerical optimization* 2nd Ed., New York: Springer.
- Østergård, T., Jensen, R.L. & Maagaard, S.E., 2017. Early building design: informed decision-making by exploring multidimensional design space using sensitivity analysis. *Energy and Buildings*, 142, pp.8–22.
- Oxman, R., 2008a. Design: current practices and research issues. *International Journal of Architectural Computing*, 6(1), pp.1–17.

- Oxman, R., 2008b. Digital architecture as a challenge for design pedagogy: theory, knowledge, models and medium. *Design Studies*, 29, pp.99–120.
- Oxman, R., 2008c. Performance-based design: current practices and research issues. *International Journal of Architectural Computing*, 6(1), pp.1–17.
- Oxman, R., 2006. Theory and design in the first digital age. *Design Studies*, 27(3), pp.229–265.
- Paolacci, G., Chandler, J. & Ipeirotis, P., 2010. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5(5), pp.411–419.
- Parmee, I.C. et al., 2000. Multiobjective satisfaction within an interactive evolutionary design environment. *Evolutionary Computation*, 8(2), pp.197–222.
- Patil, A.G.P. & Taillie, C., 1982. Diversity as a concept and its measurement. *Journal of the American Statistical Association*, 77(379), pp.548–561.
- Pavoine, S., Ollier, S. & Pontier, D., 2005. Measuring diversity from dissimilarities with Rao's quadratic entropy: are any dissimilarities suitable? *Theoretical Population Biology*, 67, pp.231–239.
- Pearson, K., 1901. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11), pp.559–572.
- Perlibakas, V., 2004. Distance measures for PCA-based face recognition. *Pattern Recognition Letters*, 25, pp.711–724.
- Podani, J., 2009. Convex hulls, habitat filtering, and functional diversity: mathematical elegance versus ecological interpretability. *Community Ecology*, 10(2), pp.244–250.
- Preisinger, C., 2014. Karamba. Available at: www.karamba3d.com.
- Proving Ground, 2011. Lunchbox. Available at: provingground.io/tools/lunchbox/.
- Queipo, N. V. et al., 2005. Surrogate-based analysis and optimization. *Progress in Aerospace Sciences*, 41(1), pp.1–28.
- Radford, A.D. & Gero, J.S., 1980. On optimization in computer aided architectural design. *Building and Environment*, 15, pp.73–80.
- Rahmani, M., 2014. Optimo. Available at: dynamobim.com/optimo/.
- Ramallo-González, A.P. & Coley, D.A., 2014. Using self-adaptive optimisation methods to perform sequential optimisation for low-energy building design. *Energy and Buildings*, 81, pp.18–29.
- Rao, C.R., 1981. *Gini-Simpson index of diversity: A characterization, generalization and applications*, University of Pittsburgh.
- Reinhart, C., 2014. *Daylighting handbook I*, Building Technology Press.
- Reinhart, C.F., 2011. Simulation-based daylight performance predictions. In J. Hensen & R. Lamberts, eds. *Building performance simulation for design and operation*. Taylor & Francis.
- Reinhart, C.F. & Herkel, S., 2000. The simulation of annual daylight illuminance distributions — a state-of-

- the-art comparison of six RADIANCE-based methods. *Energy and Buildings*, 32, pp.167–187.
- Ren, Z. et al., 2011. Multi-disciplinary collaborative building design-A comparative study between multi-agent systems and multi-disciplinary optimisation approaches. *Automation in Construction*, 20(5), pp.537–549.
- Rényi, A., 1961. On measures of entropy and information. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*. pp. 547–561.
- Reynolds, D., Ghantous, K. & Otani, R.K., 2015. Performance measures from architectural massing using machine learning. In *Proceedings of the International Association for Shell and Spatial Structures (IASS) Symposium 2015, Future Visions*. Amsterdam.
- Reynolds, D., Wang, M. & Ghantous, K., 2018. Structural design taste test. In *Proceedings of the IASS Symposium 2018*. Boston, MA.
- Rezaee, R. et al., 2015. Assessment of uncertainty and confidence in building design exploration. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 29(4), pp.429–441.
- Rippmann, M., Lachauer, L. & Block, P., 2012a. Interactive vault design. *International Journal of Space Structures*, 27(4), pp.219–230.
- Rippmann, M., Lachauer, L. & Block, P., 2012b. RhinoVAULT. Available at: block.arch.ethz.ch/brg/content/tool/rhinovault/.
- Risi, S. et al., 2009. How novelty search escapes the deceptive trap of learning to learn. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2009)*. New York: ACM.
- Ritter, F., Geyer, P. & Borrmann, A., 2013. The Design Space Exploration Assistance Method: constraints and objectives. In *Proc. of the 13th Conference on Construction and Virtual Reality (CONVR)*. pp. 72–81.
- Ritter, J., 1990. An efficient bounding sphere. In A. S. Glassner, ed. *Graphics Gems 1*. pp. 301–303.
- Robert McNeel & Associates, 2007. Grasshopper. Available at: www.grasshopper3d.com.
- Robertson, B.F. & Radcliffe, D.F., 2009. Impact of CAD tools on creative problem solving in engineering design. *CAD Computer Aided Design*, 41(3), pp.136–146.
- Robertson, B.F. & Radcliffe, D.F., 2006. The role of software tools in influencing creative problem solving in engineering design and education. In *International design engineering technical conference, ASME 2006*. Philadelphia, PA, pp. 999–1007.
- Roudsari, M. & Mackey, C., 2013. Ladybug Tools. Available at: www.ladybug.tools.
- Rowe, P.G., 1986. *Design thinking*, Cambridge, MA: MIT Press.
- Ruegg, C., Larsson, J. & Van Gael, J., 2009. Math.NET. Available at: www.mathdotnet.com.
- Runco, M.A. & Charles, R.E., 1993. Judgments of originality and appropriateness as predictors of creativity. *Personality and Individual Differences*, 15(5), pp.537–546.

- Rusch, C., 1966. *The psychological basis for an incremental approach to architecture*. Ph.D. Thesis, University of California at Berkeley.
- Salman, H.S. & Laing, R., 2014. The impact of computer aided architectural design programs on conceptual design in an educational context. *Design Studies*, 35, pp.412–439.
- Sammon, J.W., 1969. A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers*, C-18(5), pp.401–409.
- Samuelson, H. et al., 2016. Parametric energy simulation in early design: high-rise residential buildings in urban contexts. *Building and Environment*, 101, pp.19–31.
- Sanguinetti, P. et al., 2010. Real-time design feedback: coupling performance-knowledge with design iteration for decision-making. *Proceedings of the 2010 Spring Simulation Multiconference*.
- Scheibehenne, B., Greifeneder, R. & Todd, P.M., 2010. Can there ever be too many options? A meta-analytic review of choice overload. *Journal of Consumer Research*, 37(3), pp.409–425.
- Schmit, L.A., 1981. Structural synthesis - its genesis and development. *AIAA Journal*, 19(10), pp.1249–1263.
- Schölkopf, B., Smola, A. & Müller, K.-R., 1998. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5), pp.1299–1319.
- Scholz, M., Fraunholz, M. & Selbig, J., 2008. Nonlinear Principal Component Analysis: neural network models and applications. In A. Gorban et al., eds. *Principal Manifolds for Data Visualization and Dimension Reduction*. Springer.
- Schulz, A. et al., 2017. Retrieval on parametric shape collections. *ACM Transactions on Graphics*, 36(1), pp.1–14.
- Schumacher, P., 2009. Parametricism: a new global style for architecture and urban design. *Architectural Design*, 79(4), pp.14–23.
- Scott, S.D., Lesh, N. & Klau, G.W., 2003. Investigating human-computer optimization. *Proceedings of the SIGCHI conference on Human factors in computing systems*, 4(1), pp.155–162.
- Sederberg, T.W. & Parry, S.R., 1986. Free-form deformation of solid geometric models. *Proceedings of the 13th annual conference on computer graphics and interactive techniques - SIGGRAPH '86*, 20(4), pp.151–160.
- Shah, J.J., Kulkarni, S.V. & 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., Smith, S.M. & Vargas-Hernandez, N., 2003. Metrics for measuring ideation effectiveness. *Design Studies*, 24(2), pp.111–134.
- Shaikh, P.H. et al., 2014. A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renewable and Sustainable Energy Reviews*, 34, pp.409–429.
- Shan, S. & Wang, G.G., 2010. Survey of modeling and optimization strategies to solve high-dimensional

- design problems with computationally-expensive black-box functions. *Structural and Multidisciplinary Optimization*, 41(2), pp.219–241.
- Shannon, C.E. & Weaver, W., 1949. *The mathematical theory of information*, Urbana: University of Illinois Press.
- Shea, K., Aish, R. & Gourtovaia, M., 2005. Towards integrated performance-driven generative design tools. *Automation in Construction*, 14(2), pp.253–264.
- Shi, X. et al., 2016. A review on building energy efficient design optimization from the perspective of architects. *Renewable and Sustainable Energy Reviews*, 65, pp.872–884.
- Shi, Z., Fonseca, J.A. & Schlueter, A., 2017. A review for simulation-based urban form generation and optimization modeling for energy-driven urban design at the district scale. *Building and Environment*, 121, pp.119–129.
- Shireen, N. et al., 2017. Making sense of design space. In *International Conference on Computer-Aided Architectural Design Futures*. Singapore, pp. 191–211.
- Simmons, B. et al., 2015. Determining the cost optimum among a discrete set of building technologies to satisfy stringent energy targets. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 29(4), pp.417–427.
- Simonen, K. et al., 2017. Embodied carbon benchmark study: LCA for low carbon construction. Available at: hdl.handle.net/1773/38017.
- Simpson, T.W. et al., 2007. Impact of response delay and training on user performance with text-based and graphical user interfaces for engineering design. *Research in Engineering Design*, 18(2), pp.49–65.
- Simpson, T.W., Mauery, T.M., et al., 2001. Kriging models for global approximation in simulation-based multidisciplinary design optimization. *AIAA Journal*, 39(12), pp.2233–2241.
- Simpson, T.W., Peplinski, J.D., et al., 2001. Metamodels for computer-based engineering design: Survey and recommendations. *Engineering with Computers*, 17(2), pp.129–150.
- Smaling, R., 2005. *System architecture analysis and selection under uncertainty*. Ph.D. Thesis, Massachusetts Institute of Technology.
- Sóbestor, A. et al., 2014. Engineering design applications of surrogate-assisted optimization techniques. *Optimization and Engineering*, 15(1), pp.243–265.
- Sobieszczanski-Sobieski, J., 1995. Multidisciplinary design optimization: an emerging new engineering discipline. In J. Herskovits, ed. *Advances in Structural Optimization, Solid Mechanics and Its Applications*. Dordrecht: Springer.
- Soebarto, V. et al., 2015. Capturing the views of architects about building performance simulation to be used during design processes. In *Proceedings of BS2015: 14th Conference of International Building Performance Simulation Association*. Hyderabad, India, pp. 1480–1487.
- Solemna, 2012. Diva-for-Rhino. Available at: solemma.net/Diva.html.

- SOM, 2000. San Francisco International Airport - International Terminal. Available at: www.som.com/projects/san_francisco_international_airport__international_terminal.
- Somers, J., 2017. Don't use the force, Luke—use the targeting computer. *The Atlantic*.
- Souza, C.R., 2014. The Accord.NET Framework. Available at: accord-framework.net/index.html.
- Srinivas, M. & Patnaik, L.M., 1994. Adaptive probabilities of crossover and mutation in genetic algorithms. *IEEE Transactions on Systems, Man and Cybernetics*, 24(4), pp.656–667.
- St John, N., 2019. River Tower, Ho Chi Minh City, Vietnam 2008. Available at: www.nathanstjohn.com/Mixed-Use-Buildings.
- Stiny, G., 1980. Introduction to shape and shape grammars. *Environment and planning B: planning and design* 7.3, 7(3), pp.343–351.
- Stump, G. et al., 2003. Design space visualization and its application to a design by shopping paradigm. In *Proceedings of DETC'03 ASME 2003 Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. Chicago, IL.
- Su, C.-T. & Tong, L.-I., 1997. Multi-response robust design by principal component analysis. *Total Quality Management*, 8(6), pp.409–416.
- Sun, X., 2014. Incorporating multicriteria decision analysis techniques in aircraft conceptual design process. *Journal of Aircraft*, 51(3), pp.861–869.
- Takagi, H., 2001. Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. *Proceedings of the IEEE*, 89(9), pp.1275–1296.
- Tam, K.-M.M. et al., 2018. Fabrication-aware structural optimisation of lattice additive-manufactured with robot-arm. *International Journal of Rapid Manufacturing*, 7(2–3), pp.120–168.
- Tamke, M. et al., 2017. Effect and affect: designing and making with data-integrated techniques. *International Journal of Architectural Computing*, 15(1), pp.3–5.
- Tappeta, R. V et al., 2000. Interactive physical programming: tradeoff analysis and decision making in multicriteria optimization. *AIAA Journal*, 38(5), pp.917–926.
- Terzidis, K., 2006. *Algorithmic Architecture*, Oxford: Architectural Press.
- The Happy Pontist, 2009. Bridge competition debris part 12: River Wear. Available at: happyontist.blogspot.com/2009/03/bridge-competition-debris-part-12-river.html.
- The R Foundation, 1997. R. Available at: www.r-project.org.
- Thornton Tomasetti, 2016. Design Explorer. Available at: core.thorntontomasetti.com/design-explorer/.
- Thornton Tomasetti, 2017. TT Toolbox. Available at: www.food4rhino.com/app/tt-toolbox.
- Tong, L.I., Wang, C.H. & Chen, H.C., 2005. Optimization of multiple responses using principal component analysis and technique for order preference by similarity to ideal solution. *International Journal of Advanced Manufacturing Technology*, 27(3–4), pp.407–414.

- Tsigkari, M. et al., 2013. Integrated design in the simulation process. In *Proceedings of the Symposium on Simulation for Architecture & Urban Design*.
- Tsui, K.L., 1992. An Overview of Taguchi Method and newly developed statistical methods for robust design. *IIE Transactions*, 24(5), pp.44–57.
- Turiello, A., 2013. Generation. Available at: www.food4rhino.com/app/generation.
- Turrin, M. et al., 2010. Performance-oriented design of large passive solar roofs: a method for the integration of parametric modelling and genetic algorithms. *28th eCAADe Conference Proceedings*, pp.321–330.
- Turrin, M., Von Buelow, P. & Stouffs, R., 2011. Design explorations of performance driven geometry in architectural design using parametric modeling and genetic algorithms. *Advanced Engineering Informatics*, 25(4), pp.656–675.
- Umetani, N., 2017. Exploring generative 3D shapes using autoencoder networks. In *SA '17 SIGGRAPH Asia 2017 Technical Briefs*. Bangkok, Thailand.
- Vierlinger, R., 2012. Octopus. Available at: www.food4rhino.com/app/octopus.
- Villéger, S., Mason, N.W.H. & Mouillot, D., 2008. New multidimensional functional diversity indices for a multifaceted framework in functional ecology. *Ecology*, 89(8), pp.2290–2301.
- Waibel, C. et al., 2019. Building energy optimization: an extensive benchmark of global search algorithms. *Energy & Buildings*, 187, pp.218–240.
- Waibel, C., 2018. *Simulation-based optimization of buildings and multi-energy systems*. Dr. Sc. Thesis, ETH Zurich.
- Wang, J., Neskovic, P. & Cooper, L.N., 2007. Improving nearest neighbor rule with a simple adaptive distance measure. In *Pattern Recognition Letters*. pp. 207–213.
- Wang, Y. & Duarte, J.P., 2002. Automatic generation and fabrication of designs. *Automation in Construction*, 11, pp.291–302.
- Ward, G.J., 1994. The RADIANCE lighting simulation and rendering system. In *Proceedings of the 21st annual conference on computer graphics and interactive techniques*.
- Welzl, E., 1991. Smallest enclosing disks (balls and ellipsoids) H. Maurer, ed. *New Results and New Trends in Computer Science*, 555, pp.359–370.
- Wetter, M., 2001. GenOpt - a generic optimization program. In *Proc. IBPSA's Building Simulation 2001 Conference*. Rio de Janeiro, pp. 601–608.
- Whiting, E., 2012. *Design of structurally-sound masonry buildings using 3D static analysis*. Ph.D. Thesis, Massachusetts Institute of Technology.
- Williams, C.J., 2016. Conversation with Chris Williams. Bath, UK.
- Witten, I.H. et al., 2016. *Data mining: practical machine learning tools and techniques*, Burlington, MA:

Morgan Kaufmann.

Woodbury, R., 2010. *Elements of parametric design*, London: Routledge.

Woodbury, R.F. & Burrow, A.L., 2006. Whither design space? *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 20, pp.63–82.

Wortmann, T., 2018. *Efficient, visual, and interactive architectural design optimization with model-based methods*. Ph.D. Thesis, Singapore University of Technology and Design.

Wortmann, T., 2017a. Opossum: introducing and evaluating a model-based optimization tool for grasshopper. In P. Janssen et al., eds. *Protocols, Flows and Glitches, Proceedings of the 22nd International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA)*. Hong Kong, pp. 283–292.

Wortmann, T., 2017b. Surveying design spaces with performance maps: a multivariate visualization method for parametric design and architectural design optimization. *International Journal of Architectural Computing*, 15(1), pp.38–53.

Wortmann, T. & Nannicini, G., 2016. Black-box optimisation methods for architectural design. In S. Chien et al., eds. *Living Systems and Micro-Utopias: Towards Continuous Designing, Proceedings of the 21st International Conference of the Association for Computer-Aided Architectural Design Research in Asia*. Hong Kong, pp. 177–186.

Yonekura, K. & Watanabe, O., 2014. A shape parameterization method using principal component analysis in applications to parametric shape optimization. *Journal of Mechanical Design*, 136, pp.1–7.

Zandieh, M. et al., 2009. A robust parameter design for multi-response problems. *Journal of Computational and Applied Mathematics*, 230(2), pp.463–476.

10 Appendix A: Design Tool Information

This section describes specific components available as part of the DSE toolkit introduced in Chapter 6. At submission, the tools can be downloaded at:

www.food4rhino.com/app/design-space-exploration

The tools are open-source, and the code for the most recent released version can be found at:

github.com/Digital-Structures/gh-design-space-exploration

Sampler | This component automatically generates a list of parametric design vectors, called a “design map”, based on user-defined design variable properties. Different modes enable Grid, Random, and Latin Hypercube Sampling. It outputs this design map as a nested list in Grasshopper, which can be used directly or plugged into other DSE tools, and saves it as a .csv file for documentation and interfacing with outside software.

Sift | This component allows the user to select a specific design from the map and access the geometry or other properties generated by the script for this design vector. Although the next component (Capture) allows for cycling through all the designs, Sift changes variable settings based on the given index input, which saves the user from setting these variables manually while considering a single design.

Capture | This component is a general iterator, which allows the user to automatically generate many different design options and record an image, the performance, and other properties of each design. It contains options for exporting this data to external software, accessing it through Grasshopper, or completing separate parametric tasks at each iteration.

Cluster | This component executes two separate actions: the first is clustering design samples based on their design variables using k-means, and the second is using this information to support cluster-based live exploration. Using this component, a designer can cycle through different clusters organized by performance and explore within the bounds of each cluster, while gaining knowledge about the relationship between design decisions and performance.

Effects | This component gives a quick measurement of variable importance to overall performance by sampling an orthogonal set of design points and attempting to isolate the contribution of each setting. It reads in variables, objectives, and experiment settings and evaluates a small number of designs before outputting the average and raw effect of each variable. Each “effect” is a number that signifies the effect

on the objective value that results from the “cause” of that variable setting. By testing a few different settings and averaging the magnitudes of the effects for a variable, designers can get a sense of what variables have the most influence on the problem.

Tilde | This component builds a surrogate model to approximate an objective function using data generated by sampling a design space or provided by the user. Tilde takes in variables and a design map + objectives previously generated for a problem, trains the model when told to “build”, and thereafter outputs an estimated value for the objective function based on the current variable settings. Thus, after spending time to produce the original dataset and train the model, users can gain essentially instant feedback about the performance of any design vector within a parametric model. Tilde works by attempting to fit different Ensemble Neural Network or Random Forest models to the data, validating each model by calculating its errors, and finally choosing the best model for approximation.

Diversity | This component helps users filter sampled design space results based on objective values, while also generating a reduced, diverse, representative sample of qualified designs. This process has two steps. First, the user sets target objectives and acceptability limits for removing unwanted solutions. Next, he or she provides a number of desired representative samples, at which point the component iteratively searches for a set of solutions with a high measured diversity. Together, these functions help simplify post-simulation analysis, while producing a manageable design set with meaningfully different solutions for visual inspection.

MOO | This component implements the NSGA-II multi-objective optimization algorithm (Deb et al. 2002), which is a non-dominated sorting genetic algorithm. NSGA-II approximates the Pareto front in a given design problem by using crossover and mutation to iteratively breed successive, higher-performing generations of designs.

Stepper | This component is an interactive gradient-based optimization tool that allows users to move in the objective space in addition to the design space. It includes a separate interface from Grasshopper that connects to the original design variables and a rapid objective function simulation or approximation. When in operation, designers can select any objective function and attempt to make it increase or decrease from the current point in the design space. They can also select an isoperformance direction and attempt to find a similarly performing design that is nearby in the objective space. The interface includes a history of the changing objective function and options for adjusting the step size and included variables of the problem.

Contort | This component connects to the original design variables and drives the design based on new synthetic sliders mapped to the original design space using coefficients. While this functionality can also be achieved manually by switching sliders in Grasshopper, this component skips a few steps and enables faster exploration during design space formulation and transformation.