Learning about Parametric Model Behavior through Multi-Objective Optimization

Volker Mueller
Bentley Systems, Incorporated, USA
Volker.Mueller@Bentley.com

Abstract

This paper reports about a design process as a case study illustrating different levels of learning that seem required for successful computational design. The learning process occurred during a two-day workshop about parametric design with integrated analysis and multi-objective optimization. First, the design team needs to understand the behavior of the model in order to validate that the model behaves in a way that actually conforms with the project goals; second, the design team needs to learn about potential trade-offs between different project goals, and thus understand the decisions that need to be made, or the additional problems that need to be solved in order to arrive at a better design solution.

Keywords: Simulation and Modeling, Generative Systems, Parametric Design Multi-objective Optimization, Computational Design Learning

Introduction

Computational design approaches employing parametric models are increasingly utilized in design processes because of designers' abilities to embed behaviors in the models and explore variations by changing parameters or behaviors. Incorporation of building performance metrics as design criteria fits into these approaches because performance metrics are mostly computed, which blends seamlessly into an overall computational design framework (Kolarevic and Malkawi, 2005). This enables designers to improve their designs toward better results especially when compared to traditional design processes, given all other constraints like time, budget, or effort remaining equal, because more design variants can be examined with more variations. Support of these explorations of many variants by automating them with optimization processes accelerates this type of investigation so that examination of many variations becomes feasible (Mueller, 2014).

This type of work flow requires translation of the client's and design team's goals into a parametrically controlled model including types of analysis that reflect applicable goals and make them measureable. Because it has been an innovative work flow in architectural design various researchers have investigated several aspects of learning that is involved in acquiring a working skill level of parametric design and its extensions. Topics investigated have been for example how parametric design can be made accessible through different types of views (Aish and Woodbury, 2005), how recasting of parametric models into meaningful pieces, i.e. "design patterns", may help ease their construction (Woodbury et al., 2007; Qian, 2009), how sensitivity analysis unlocks the understanding how specific parameters influence behavior of the parametric model (Hopfe et al., 2007; Hopfe, 2009), or what challenges analysis integration into early design processes may pose (Bleil de Souza, 2009, 2012).

This paper reports about a conceptual parametric design process as a case study to illustrate yet different aspects of learning that are involved or even required for successful computational design incorporating multi-objective optimization. In multi-objective optimization a parametric model is utilized to generate in an automated fashion many candidate solutions. The integrated analyses compute for each solution the corresponding metrics which may be subject to additional computation in order to match more closely the design goals. These values are routed as objective functions to the optimization algorithm which uses them to determine the next set of candidate solutions. For various types of optimization algorithms the details may vary. Crucial is that with the search for the set of best performing solutions having been automated there is an additional level of due diligence necessary to assure that this automated search is leading in the correct direction. The thesis is that this requires understanding of whether and how parametric model behaviors reflect design goals and how selected analyses measure how design goals are being met.

Methodology

The methodology used is an ethnographic approach focusing on participant observation amended by discussions replacing qualified key informant interviews. The material is presented in form of a case study even though similar observations have been made in similar cases; however, there has not been an attempt to formally report on those other cases. It is based on observations of learning processes which occurred during a two-day workshop about parametric design with integrated analysis and multi-objective optimization at a conference for computer aided architectural design. Most participants were designers or researchers with experience in parametric modeling, analysis embedded in design processes, and optimization. As design authoring tool GenerativeComponents (Bentley, 2015a) was used with an optimization add-in based on the DARWIN genetic algorithm (Bentley, 2015b). The workshop participants chose as format a collaborative approach similar to a design team working on a single project.

Terminology

In this paper some terms are used interchangeably: design criterion, performance criterion, objective, fitness function are used to describe computed values that are used to represent design goals as numeric values that can be accepted by the optimization algorithm. Parameters or design variables describe the inputs to the parametric model that are identified to the optimization algorithm as the ones it may vary in order to generate a new design solution. Computational design and parametric design describe an approach to design that uses mathematical methods. The computational design model, parametric design model, or model is the underlying parametric model that expresses its contents as geometry or other outputs. Parametric implies that the model is able to react dynamically to changing inputs. Multi-objective optimization or multi-disciplinary optimization describe a group of optimization algorithms that permit more than one distinct fitness functions to be passed to the algorithm which then computes Pareto optimal solutions. Pareto optimal solutions are those solutions that lie on the Pareto frontier; they are the solutions not dominated by any other solution, in two or more dimensions depending on the presentation as 2d scatter plot or on other presentations; non-dominated solutions in a scatter plot are identified by no other solution being in the space spanned by the rectangle (in the two dimensional case) with the solution and the plot's origin as diagonally opposite corners and edges parallel to the axes.

The Case Study

The starting point for the workshop was a sample script for a very simple massing study with parameters for the length and width of the building footprint, the floor-to-floor height, total height of the building, rotation of the building, a second rotation parameter for twisting the floors incrementally from bottom to top, and additional parameters to control curvature in the vertical edges of the building. Already integrated in the computational model were four performance metrics computed from the design parameters or values computed from the resulting design variants:

1. a proxy analysis for daylight potential based on the ratio of the length of the longer edge of the rectangular floor slab to the length of the shorter edge, with larger ratios better because they indicate a smaller depth of the floor area, which customarily means better penetration of daylight from any window area on the long sides into the depth of the building;
2. a proxy for energy efficiency computed as ratio of the envelope to the volume of the building mass with lower values meaning less envelope enclosing more volume, which with some assumptions about the intended location of the building in this case means better performance;
3. economic viability as available floor area, not considering any subtractions for vertical circulation even though this removes a potential penalty for tall building circulation requirements compared to low-rise buildings; and
4. façade non-planarity as indicator for increased construction cost, with the non-planarity measured as accumulated out-of-plane measurements of one of four points of a quadrangular façade tiling computed against the plane spanned by the tile's other three points.

This starting point for the sample case study had been constructed by the workshop leader and had gone through one workshop iteration already. Starting with this initial, somewhat arbitrary, model state an optimization run permitted a first informative examination of the model's behavior. In this run, the four fitness criteria were the respective proxy fitness criteria for economic viability, energy efficiency, construction cost, and daylighting potential. Width and depth of the building footprint, height, and the rotation of the building were passed as design parameters to the optimization engine which used them to build the genome in the evolutionary optimization process. Any rational parameter theoretically contains infinite variations. Therefore, for effectiveness, all parameter ranges have to be discretized, which means that even though the solution space spans a potentially huge number of solutions, it is nevertheless a finite solution space.

The optimization algorithm that was used in the case study is an evolutionary algorithm which produces iterative generations and evaluates them to develop the genomes for the respective subsequent generations. Without delving into the details of how the specific optimization algorithm works, there are a variety of termination criteria with which users can control when the optimization run terminates. The longer the run, the denser the Pareto frontier will grow. Independent of termination criteria, users can stop an optimization run at any time and explore the solutions that have been generated up to that point.

The result of the multi-objective optimization run is a series of scatter plots mapping the model's performance against two axes showing pairwise fitness criteria, or design goals (figure 1). By selecting a dot in a scatter plot the corresponding geometry is instantiated in a geometric view window (figure 2). Thus it is possible to explore what the

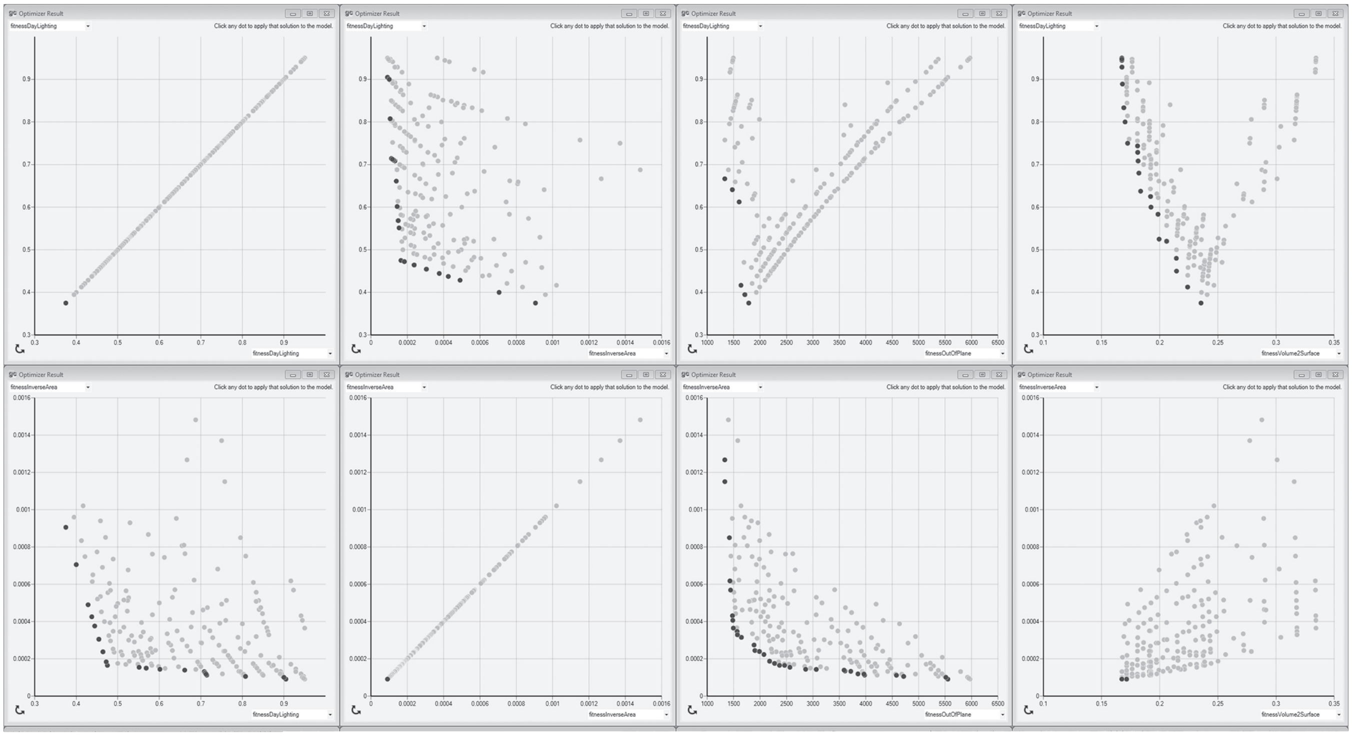

Figure 1: Pareto plots of four fitness criteria. All plotted dots show Pareto optimal solutions in the four-dimensional hyperspace spanned by the four fitness criteria. The darker dots indicate those solutions that are on the two dimensional Pareto frontiers for the respective two fitness criteria on the horizontal and vertical axes of the plots. The plots on the top left to bottom right diagonal are equivalent to single-objective optimization runs and the dark dot in each indicates the best case for that specific criterion, while the other solutions fare worse on that criterion but are still Pareto optimal solutions in terms of other criteria.

various alternatives look like to support visually what in the scatter plot is available only numerically. When plotting the same fitness criterion on both axes, the solution at the bottom left end of the diagonal plot is the optimal solution for that criterion achieved in the specific optimization run, and the top right solution is the worst case generated in the specific optimization run. Note that this is not the worst possible case because the optimization is not interested in generating bad cases. It is just the worst case generated in the specific optimization run. More specifically, in the case of the used software, solutions are only presented if they are on the multi-dimensional Pareto frontier.

Visual inspection of these cases yielded a first recognition: the best solutions for low construction cost are low buildings with small footprints. Examination of how the construction cost proxy value was computed revealed that the outcome would be optimized by minimizing the total non-planar façade area, leading to buildings with the lowest height permitted by the range of the height parameter, as well as, with the smallest footprint permitted by the ranges of the building's horizontal dimensions. This is, of course, logical behavior. The building with the least construction cost is the building that is not being built. In this case, it is the smallest building permitted by the low limits of the parameters involved. The

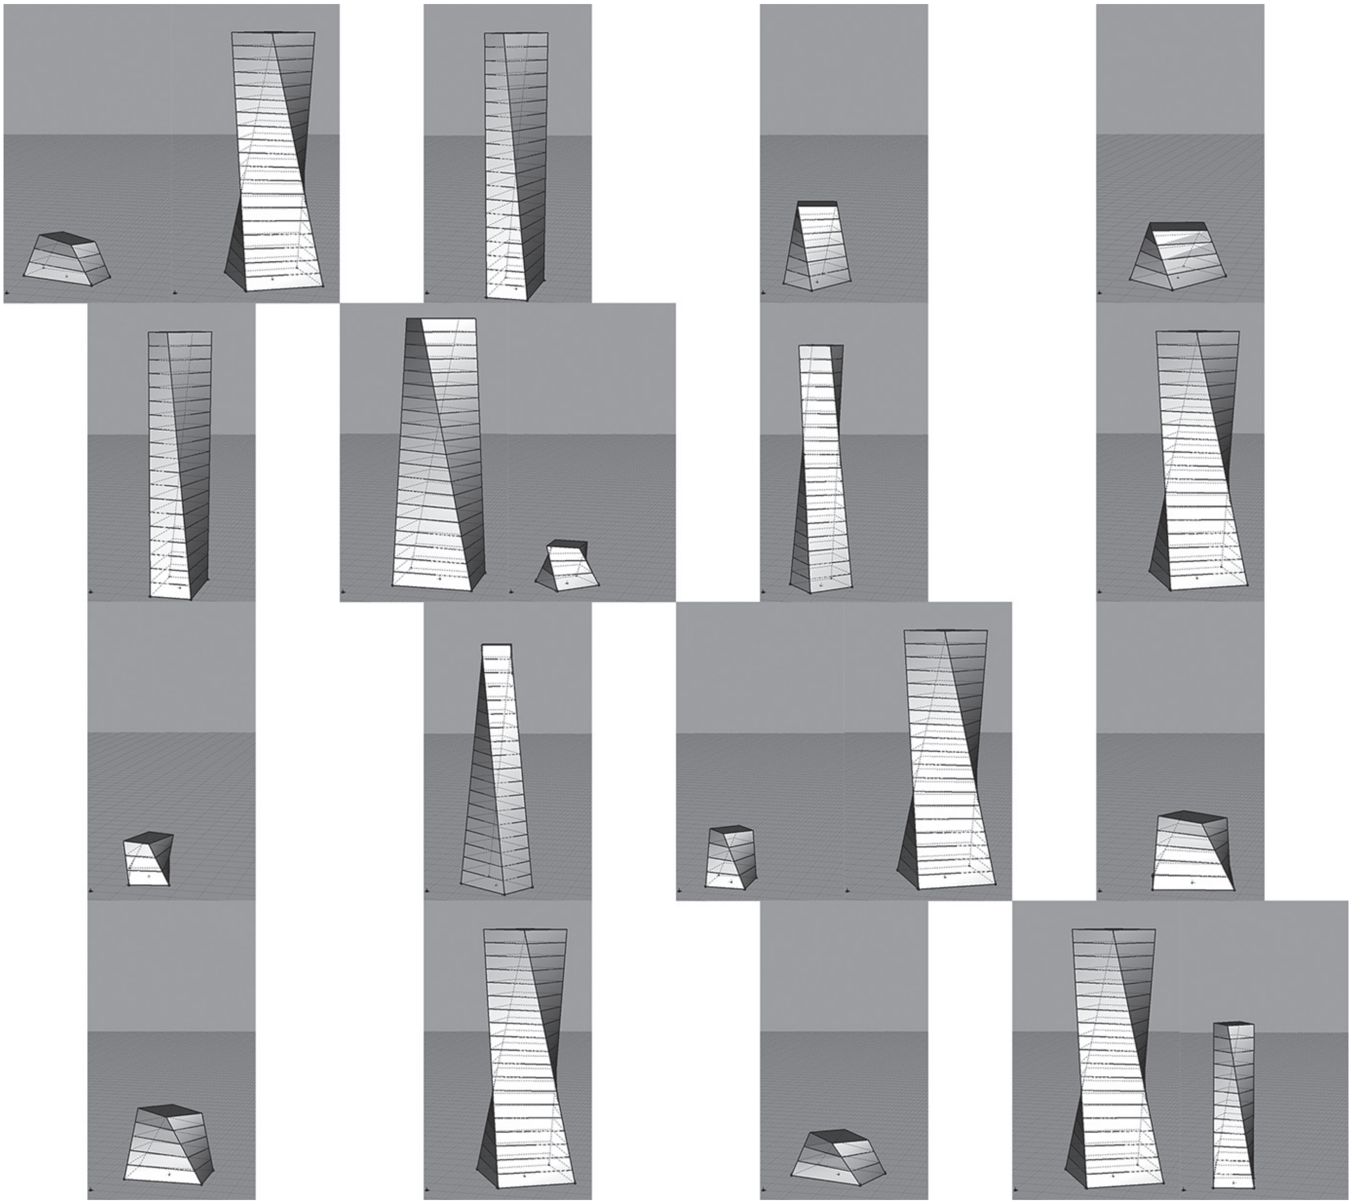

Figure 2: Screen captures of some of the solutions. On the top left to bottom right diagonal are the optimal and worst generated cases side by side for the respective fitness criteria. Optimal cases are at the bottom left end of those scatter plots in figure 1 that map the same fitness criteria on both axes; worst cases are at the top right end of those scatter plots. The other fields show a solution more or less arbitrarily sampled from the area of the Pareto frontier close to its middle in proximity to the plot's origin. The criteria are daylighting, economic viability, construction cost, and energy efficiency.

daylighting fitness exhibited a similar behavior, with the lowest possible building being the optimal solution; however, the solution did include correctly the maximum ratio of the two footprint dimensions. The examination of this behavior revealed that the computation was biased towards the lowest permitted solution.

Further discussion of the results led the workshop attendees to the conclusion that a customary design goal had been included in a somewhat optional fashion instead of as a constraint or a mandatory target. The floor area was included as a fitness criterion to be maximized. With that the shown trade-offs are valid, trading off floor area against construction cost or daylighting potential. However, when revising the design goals to include a specific target floor area, then the parametric model needs to be revised, too.

These two insights illustrate two different types of gained understanding about the parametric model: the first insight was that a correctly working model may not accurately reflect the design goals. The first optimization run revealed that quite quickly without wasting effort proceeding too far along a route based on erroneous assumptions. The second insight was that a working model may not correctly reflect the design goals, too. This was also revealed quickly in the first optimization run. Arguably, a weighted single-objective optimization could have obscured these two shortcomings.

Given the time constraints of the workshop, a simple technique was applied to generate a building mass with the target area, by computing a scale factor from the floor area of the initially generated building mass and the target area, then scaling the initially generated building mass horizontally accordingly so that it meets the target area. In practice a small refactoring of the parametric model would correctly reflect this design change which would be necessary to constrain the building mass to a specific maximum foot print.

Further examination revealed another parameter, rotation, as ineffective regarding the performance criteria considered in the optimization. The initial reasoning behind the rotation parameter, which rotates the entire building, was calculation of a performance parameter that would react to changes in rotation, for example insolation, heat gain, or any other meaningful fitness value that is relevant to early massing studies. Instead of rotation the model included another variable, twist, which rotates each floorplate against each other, this way influencing the non-planarity of the façade, ergo the construction cost proxy. However, this variable was not indicated to the optimization algorithm so that all solutions of the first run were twisted (figure 2). The effects of switching those two parameters showed immediately in a next optimization run (figure 3).

Drawing on the inherent flexibility of a computational system, workshop attendees challenged the proxy analysis for daylighting potential. Instead of considering only the ratio of the floorplates, the suggestion was to compute the areas of the floorplates that are likely to remain without sufficient daylight potential. As parameter for the depth of intrusion of daylight into the floorplate the double value of the floor-to-floor height parameter is used. For each rectangular floorplate first the smallest dimension is calculated and then checked against twice the daylight intrusion depth. If it is larger than twice the daylight intrusion depth the area of the inset rectangle is determined, offset from all exterior edges of the floorplate by the daylight intrusion depth. This inset rectangle is the area that remains without sufficient daylight potential and is used as performance criterion that is to be minimized. Another optimization run confirmed that this improved daylighting potential proxy analysis exposes the behavior of the parametric model as workshop attendees intended. When inspecting the results of this optimization run it also became obvious that in contrast to the first optimization run there was only one criteria pairing that showed a Pareto frontier, the energy efficiency proxy and the daylighting potential proxy (figure 4). Exploration along the Pareto frontier showed the trade-off between energy efficiency and daylighting potential, certainly as implied by the assumptions in the proxy analyses.

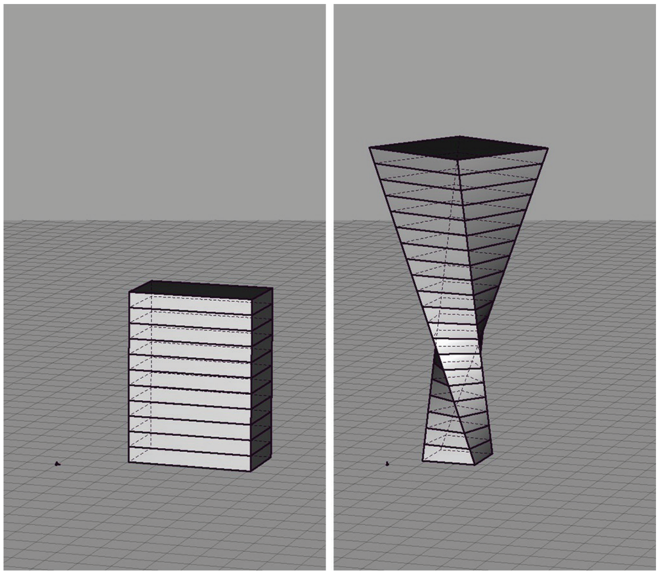

Figure 3: Optimal case and generated worst case for construction cost proxy analysis after using twist parameter.

The solution dots on the vertical axis also indicate that there are several or even many possible solutions with optimal daylighting as defined by the proxy, i.e. with zero dark area. That none of the other pairings show clear Pareto frontiers indicates that these criteria are either aligned which is a banal case for multi-objective optimization, or they are not clearly dependent on each other which is then indifferent to the optimization algorithm, i.e. their result plots just show stochastic scatter. For example, the economic viability proxy had become meaningless because it had been superseded by the target area.

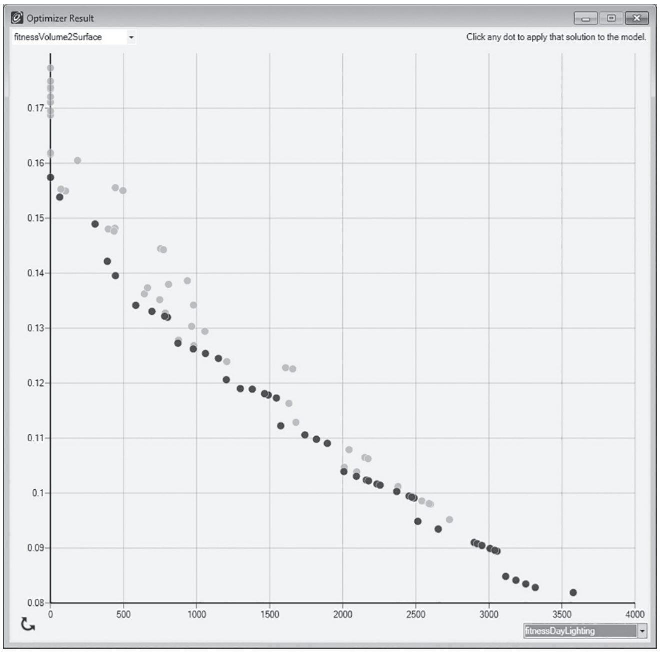

Figure 4: Scatter plot of energy efficiency proxy against daylighting potential proxy.

It became obvious that the solutions were visually less interesting, outright boring. This boredom existed at two levels: in the trade-offs and in the results. In the ensuing discussion one conclusion was that in many higher-end design projects one major trade-off is between aesthetics and costs, with latter driven by increasing complexity of building structure and building envelope. In the case study one of those aspects had already been included, the complexity of the building envelope as measured by an out-of-plane value. It was used as a construction cost proxy and the objective was to minimize it. As shown earlier (see figure 3), this construction cost proxy optimization leads to simple parallelepipeds. There is no conflicting objective that would lead to generation of more complex shapes. The worst construction cost case in the found Pareto optimal solution set arises only from the stochastic probing by the optimization algorithm, an indication that the algorithm searches quite thoroughly the solution space.

In order to generate visually complex solutions more consistently there needs to be an objective that drives solutions toward visual complexity. In the discussion of such a design goal the workshop attendees agreed to avoid a metric that attempts measuring aesthetic quality of a design solution. Rather, prompted by the initial observation of boredom, the objective became quite appropriately minimization of boredom, the inverse of which could be maximization of visual interest which might be a faint proxy for aesthetic quality –or just a response to the drive toward creating unique buildings with the intent to attract tenants. This last observation may be indicative of a potential trade-off between reduced boredom and increased tenant space as measured for the economic viability objective, which in this case does not carry because of the fixed target area; however, given the implementation of the reduction of boredom in the shape-giving part of the algorithm, it is predictable that there will be a trade-off with the construction cost proxy.

Up to the last optimization run of the workshop, the parameter set had been incrementally increased, in the last run including floor-to-floor height, building width, depth, and height, twist, and, for additional control of the edge lines to create visual interest, parameters controlling dampening of the amplitude of the edges' horizontal oscillation along the vertical axis, an amplitude factor and the frequency of the oscillation. This resulted in the expected distribution of solutions from boring to visually interesting (figures 5 and 6).

When changing parameters and objectives in iterative optimizations, new issues will be revealed. In the last optimization run, scatter plots indicated that some objectives may be independent of each other, conflicting with each other –leading to Pareto frontiers, or aligned with each other, which could mean one of them is redundant (figure 7). As anticipated, construction cost and visual interest appeared in conflict (figure 8 and 9). Note that the worst case in the boredom/boredom scatter plot is different from the worst case in the construction cost/boredom scatter plot. As mentioned earlier, this is due to the fact that the context changes for each of the plots due to the solutions' locations on the multi-dimensional Pareto frontier.

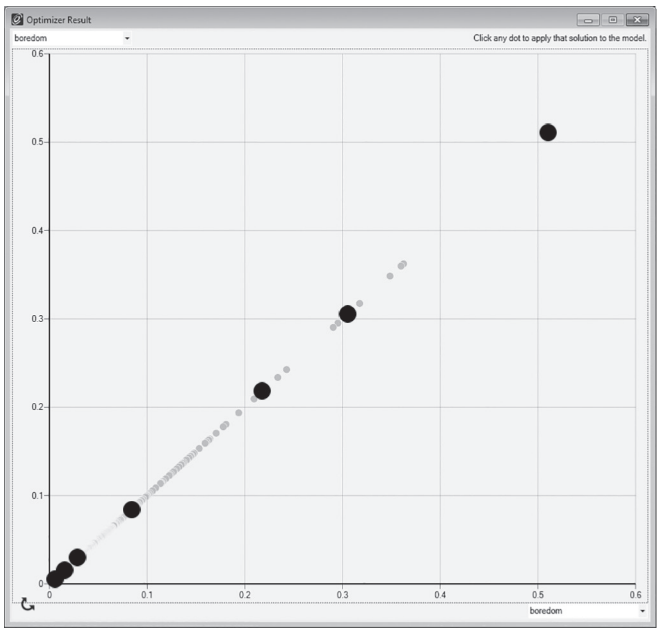

Figure 5: Scatter plot of boredom/boredom. Marked solutions are shown in figure 6.

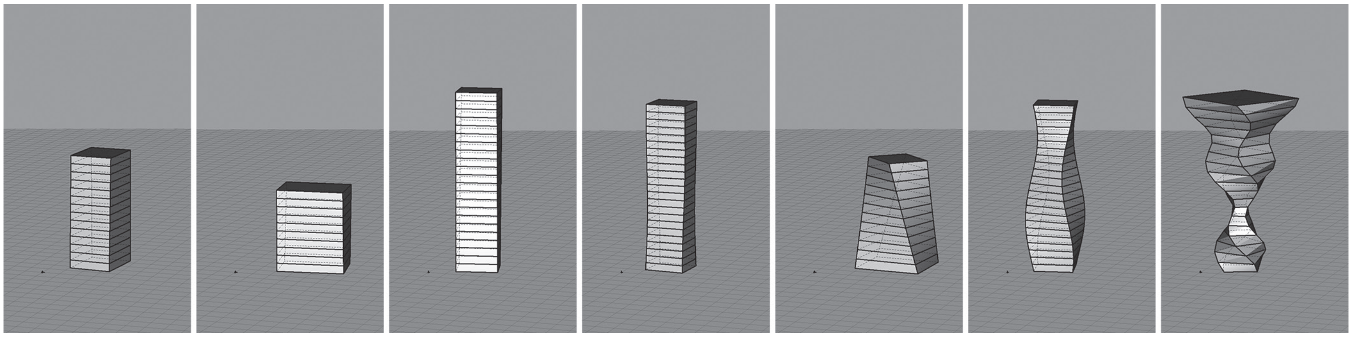

Figure 6: Worst generated solution for boredom objective to best generated solution (optimal/minimal solution) for boredom objective as marked in figure 5.

Discussion

The core issue under investigation in this paper is how multi-objective optimization supports complex design processes. The first observation is that in many complex design processes there are conflicting goals, like maximization of floor area conflicting with minimization of construction cost. This, of course, appears banal and had, therefore, been omitted in favor of the slightly more complex construction cost factor of a non-planar façade.

Inclusion of conflicting goals is per definition not a problem for multi-objective optimization, even if the response in form of Pareto optimal frontiers in scatter plots may seem more like a cop out than a true solution. However, when contrasting multi-objective optimization to single-objective optimization, in order to transform a multi-objective scenario into a single-objective scenario the conflicting goals first need to be made commensurate, and then they need to be weighed against each other a priori, so that a single fitness value can be computed for the optimization algorithm (Flager et al., 2008). This approach conveniently yields a single optimal result, something a multi-objective optimization algorithm can offer when a single fitness functions is provided; however, in order to provide a single fitness function, the comparative values of criteria trade-offs have to be anticipated in order to assign weights.

Assumptions have to be made about trade-offs that cannot be further explored, or only indirectly by changing assigned weights for subsequent optimization runs, each weight reassignment yielding only one additional data point leading to no or only very slow accumulation of knowledge how changing weights may affect outcomes. Such learning, though, is readily supported by the results of multi-objective optimization processes. Each of these processes stochastically samples the space of all possible solutions described by the parametric design. By displaying them in a scatter plot against axes with the fitness functions' own measurement, multi-objective optimization permits exploration of results in terms of those metrics, rather than through some mitigated values that attempt to make those metrics commensurate on a single axis. This means that perhaps multi-objective optimization may lead to the understanding required to define the weights necessary to convert the multi-objective problem into a single-objective issue. At that point this may not be required anymore because design decisions could be based on the understanding gained during the multi-objective optimization exploration.

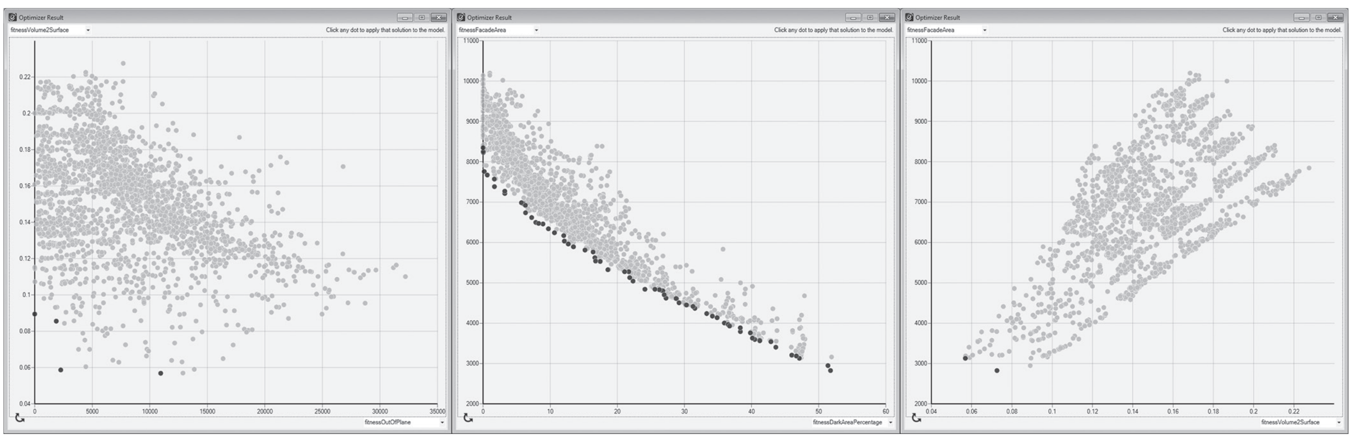

Figure 7: Scatter plots indicating different types of dependencies between plotted fitness functions: on the left the plot of energy efficiency proxy and construction cost proxy indicates that in the current parametric model there is no strong dependency between those two objectives. Center: Scatter plot of façade area fitness and daylighting proxy indicates that there is a conflicting relationship. Right: Scatter plot of façade area and energy efficiency proxy indicates that these two fitness functions are largely aligned and one of them may be redundant.

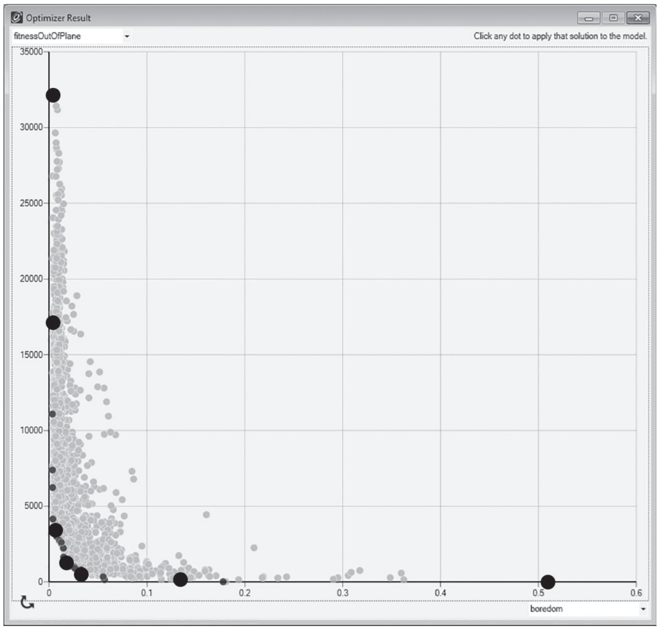

Figure 8: Scatter plot of construction cost proxy and visual interest proxy. The marked solutions are shown in figure 9.

Another observation was that the model may include parameters that do not affect the outcome. The addition of sensitivity analysis could support identification of those parameters (Hopfe et al., 2007). From a different perspective, insights derived from sensitivity analysis may also support investigation of the intended purposes for those parameters and whether they indeed have become obsolete.

Results

The examination of this case study reveals that there are different aspects of learning about a computational design model: first, the design team needs to ascertain that the model behaves in a way that actually conforms with the requirements in order to reach the project goals; second, after this validation of the computational model, the design team can start to understand what potential trade-offs exist between different project goals, and thus understand the decisions that need to be made, or additional problems that need to be solved in order to arrive at a better design solution. This entire process in itself is iterative.

Conclusion

Parametric modeling with integrated analysis and multi-objective optimization is not a trivial matter. A potential pitfall is that design teams are missing to identify gaps in their concepts about computational design and rely on algorithms to bridge those gaps while they may only be hiding them. Therefore, it is necessary to identify where these gaps may occur and raise awareness of these gaps through education, instructional materials, and other means. Nevertheless, computational design offers new opportunities for improving the results of design processes based on measurable criteria. Given the possible breadth of such design investigations the expectation is that given the same amount of effort, this type of computational design process can result in better designs, ultimately in better performing buildings.

References

Aish, R. and Woodbury, R. (2005). Multi-level Interaction in Parametric Design. In A. Butz, B. Fisher, A. Krüger and P. Olivier (Eds.), Smart Graphics: Lecture Notes in Computer Science Volume 3638 (pp. 151-162). Berlin, Germany: Springer.

Bentley (2015a) GenerativeComponents from Bentley Systems: http://www.bentley.com/en-US/promo/Generative+Components.htm, accessed September 21, 2015.

Bentley (2015b) Darwin Optimization (version 0.91) by Dr. Z. Y. Wu, http://communities.bentley.com/communities/other_communities/bentley_applied_research/w/bentley_applied_research__wiki/6584.aspx, accessed September 21, 2015.

Bleil De Souza, C. (2009). A critical and theoretical analysis of current proposals for integrating building thermal simulation tools into the building design process. Journal of Building Performance Simulation, 2:4, 283-297.

Bleil De Souza, C. (2012). Contrasting paradigms of design thinking: The building thermal simulation tool user vs. the building designer. Automation in Construction 22, 112–122

Flager, F., Soremekun, G., Welle, B., Haymaker, J. and Bansal, P. (2008). Multidisciplinary process integration and design optimization of a classroom building, CIFE Technical Report TR175, Stanford University.

Hopfe, C.J., Hensen, J.L.M. & Plokker, W. (2007). Uncertainty and sensitivity analysis for detailed design support. In Jiang, Yi (Ed.), Proceedings of the 10th IBPSA Building Simulation

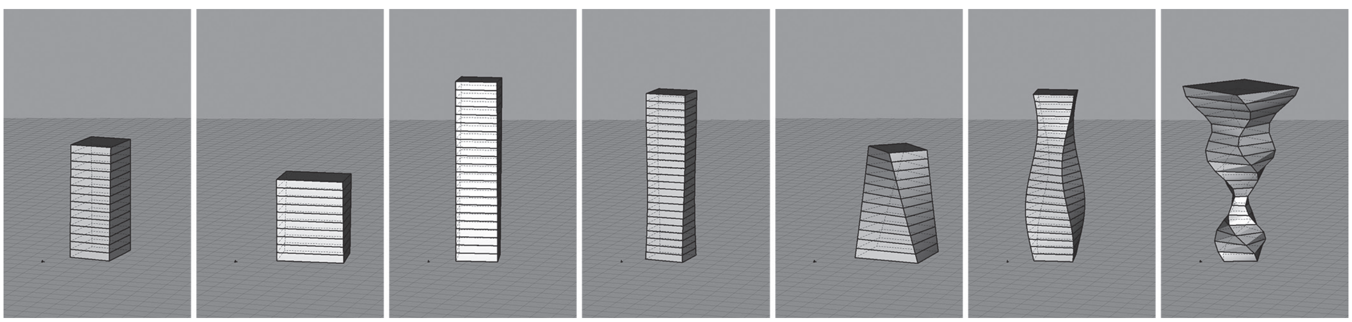

Figure 9: (from left to right) Best generated solution for (lowest amount of) boredom and worst in construction cost to best construction cost with worst (highest amount of) boredom as marked in figure 8.

Conference (pp. 1799-1804), Beijing: Tsinghua University.
Hopfe, C.J. (2009). Uncertainty and sensitivity analysis in building performance simulation for decision support and design optimization, Dissertation at Technische Universiteit Eindhoven, The Netherlands.
Kolarevic, B. and Malkawi, A. (Eds.) (2005). Performative Architecture: Beyond Instrumentality. New York, NY: Spon Press.
Mueller, V. (2014). Second generation prototype of a design performance optimization framework. In Sidawi, Bhzad and Mallasi, Zaki (Eds.), Digital crafting – virtualizing architecture and delivering real built environment, Proceedings of the 7th ASCAAD conference (pp. 199-209). Jeddah, Kingdom of Saudi Arabia: The Arab Society for Computer Aided Architectural Design.
Qian, C. (2009). Design Patterns: Augementing Design Practice in Parametric CAD Systems, PhD Thesis at Simon Fraser University, Surrey, BC, Canada.
Woodbury, R., Aish, R. and Kilian, A. (2007). Some Patterns for Parametric Modeling. In B. Lilley and P. Beesley, (Eds.), Proceedings of 27th ACADIA Conference (pp. 222-229). Halifax, NS, Canada: The Association for Computer Aided Design in Architecture