MIT Open Access Articles

An interactive expert system for daylighting design exploration

Massachusetts Institute of Technology

An Interactive Performance-Based Expert System for Daylighting Design

Jaime M. L. Gagne[1], Marilyne Andersen[2], Leslie K. Norford[1]

[1]Building Technology Program, Massachusetts Institute of Technology, Cambridge, MA, USA
[2]Interdisciplinary Laboratory of Performance-Integrated Design (LIPID), Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

ABSTRACT: Architects are increasingly using digital tools during the design process, particularly as they approach complex problems such as designing for successful daylighting performance. However, while simulation tools may provide the designer with valuable information, they do not necessarily guide the user towards design changes which will improve performance. This paper proposes an interactive, goal-based expert system for daylighting design, intended for use during the early design phase. The expert system consists of two major components: a daylighting knowledge-base which contains information regarding the effects of a variety of design conditions on resultant daylighting performance, and a fuzzy rule-based decision-making logic which is used to determine those design changes most likely to improve performance for a given design. The system gives the user the ability to input an initial model and a set of daylighting performance goals in the form of illuminance and daylighting-specific glare metrics. The system acts as a "virtual daylighting consultant," guiding the user towards improved performance while maintaining the integrity of the original design and of the design process itself.

Keywords: daylighting, expert system, design process

1. INTRODUCTION

Designers have long considered daylight as an important aid for architectural expression. In recent decades, we have come to understand that daylighting may provide additional benefits, such as reduced energy consumption and improved occupant health and well-being [1,2,3]. Nevertheless, simply providing daylight in a building will not always result in positive results. Daylighting is only as good as its delivery system, so careful design is necessary to ensure that enough light is available and that glare, shadows, and reflections are reduced [4]. Unfortunately, it is often a challenge to create a successfully daylit building.

Digital tools offer new ways of helping architects create or find designs with high levels of daylighting performance using efficient and intelligent guided design exploration methods. Optimization algorithms are a common solution, largely because they have the capabilities necessary to find or generate successful solutions; however, these methods generally do not allow for user-interaction. As it is highly unlikely for a designer to simply accept a design generated by an optimization algorithm, a better approach would be a more interactive search method, which would accept input from a designer and which would grant the designer a larger degree of control.

An example of such an approach is a knowledge-based or expert system. An expert system is one in which human expert knowledge about a specific domain is encoded in an algorithm or computer system [5]. In the daylighting domain, such a system would function as a virtual lighting consultant, guiding the designer towards design modifications which improve overall daylighting performance. Knowledge-based systems have already been successfully implemented for artificial lighting scenarios [6,7]. For daylighting, a few simple expert systems exist. The Leso-DIAL tool provides users with a "qualitative diagnosis" using an expert system based on fuzzy logic rules [8]. The NewFacades approach considers energy and visual comfort based on a prescription energy code for hot climates to suggest a range of facade solutions to the designer [9]. These systems represent first steps in expert systems for daylighting in design, but they do not allow for a comprehensive understanding of daylighting or a large amount of user interactivity.

This paper will describe a user-interactive expert system approach which enables a comprehensive analysis of daylighting. This approach includes two climate-based performance metrics, one for illuminance and one for daylighting-specific glare, in order for the designer to have an understanding of the amount of light and the visual comfort in the space. The method begins with a designer's own initial design and performance goals. It then evaluates the performance of the design and creates a series of suggestions for design changes which are likely to result in improved performance, thus enabling a search process that is highly specific to the user's design problem. Decisions are made using an expert system which is comprised of a pre-calculated database of daylighting-specific information connected to a set of fuzzy daylighting expert rules. Any design decision that the designer chooses to allow will be automatically generated in

the original model and the new performance will be calculated. The designer is allowed to interact with the system during an iterative search process that is both agreeable to the designer and likely to improve the performance of the design.

2. EXPERT SYSTEM FOR DAYLIGHTING

The expert system described in this paper is a fuzzy rule-based system combined with an external database of previously computed daylighting simulation data, called the daylighting knowledge-base. This system has been implemented as a functional tool within the Lightsolve project [10].

2.1. A Daylighting Knowledge-Base

Most expert systems are traditional systems in the sense that they are populated using knowledge from a human expert, and as a result, such systems are restricted in terms of accuracy and complexity. To create an expert system capable of more sophisticated analysis, the expert system described in this paper uses a daylighting-specific database, or "knowledge-base," which has been populated using data from a set of completed daylighting simulations.

These simulations were performed for a set of 512 models with differing facade characteristics, based on the Design of Experiments method [11]. For each model, the illuminance and a model-based approximation of the daylight glare probability (DGPm) [12] were calculated in five different zones within the space (and four different views from within each zone for the glare metric), over the whole year. These climate-based metrics were calculated using the Lightsolve Viewer ("LSV") [13], the simulation engine native to the Lightsolve program. The knowledge-base contains information about the relative effects of ten different facade parameters on each of the two daylighting metrics from the various zones and views within the space. The ten different façade parameters considered are: window area, window height-to-width ratio, vertical and horizontal location of windows on the façade, window distribution (how close or far apart windows are to each other), total number of windows, length of horizontal overhangs and/or vertical fins, glass transmissivity, and glass type (regular or translucent).

By using calculated data rather than heuristics to populate the knowledge base, the expert system can consider highly specific goals and multiple sets of goals for the same design, which can differ based on the daily time period(s), season(s), or zone(s) of interest within a space. It also allows for more logical and accurate comparisons of multiple design options than mere heuristics. A more detailed explanation of the knowledge-base can be found in [14]. The knowledge-base used in this paper used simulations from Boston, MA (USA).

2.2. Expert System User Inputs

The expert system rule base is a decision-making algorithm that assesses specific design situations and creates lists of suggested design changes which should improve the current performance. The rule base uses fuzzy logic [15], which allows it to better emulate the human thought process than classical logic. It has been developed to be a flexible algorithm which can accommodate a wide variety of initial design scenarios. The system was also created in such a way that it requires user interaction and user inputs in order to function.

The major user input is a 3d model of an original design with sensor planes for illuminance and/or glare. Additionally, performance goals for each sensor plane must be specified. For each illuminance sensor plane, the user must specify a desired illuminance goal range in lux, including the actual desired range and a buffer zone of acceptable values. For example, the user may desire the illuminance of a given sensor plane to fall between 400 lux and 1200 lux, but he or she will also accept illuminances as low as 200 lux and as high as 1500 lux. For each glare sensor or glare sensor group, the user must choose a glare tolerance. The glare tolerance options are "zero" (which means that no glare is tolerated), "medium", and "high" (which means that a high amount of glare is allowed). These tolerance values correspond to the three glare ratings of "perceptible", "disturbing", and "intolerable" glare described by Wienold in [16].

In addition to the 3d model and performance goals, the user must also several other inputs. One set of inputs is the set of priority levels for each performance goal. The priority level is a number from 1 to n, where n is the total number of sensors. The highest priority value is 1, and multiple goals may have the same priority. The user must also select a window uniformity scheme from three possible choices: "All windows in the model should look the same", "All windows on a façade should look the same", or "Windows can look different from other windows on the same façade." Finally, the user must indicate times and seasons of interest (the choices are: winter, fall/spring, summer, morning, mid-day, and afternoon) and input the latitude and a weather file for the desired location.

2.3. Fuzzy Sets and Rules

After the user has begun the expert system process, the LSV engine is used to calculate goal-based performance metrics for both illuminance and glare. This information, along with the original user inputs, is used to create sets of fuzzy variables, which help to describe the current scenario. These fuzzy sets are: userPriority (high and low), sensorPerformance (good and bad), illuminanceSensorPerformance (too high and too low), glareSensorPerformance (too high), and distanceFromGoal (close and far). In addition to these fuzzy variables, the system also creates a customized knowledge-base, which is a subset of the knowledge-base described in section 2.1 that contains only the information most relevant to the current design. Based on this customized knowledge-base, each potential design action is given values for the fuzzy set actionResult (Fig 1). These fuzzy variables refer to the likely result of the given design action on a given sensor, for example

"Large Illuminance in Illuminance". Each sensor in the model will have a unique actionResult fuzzy set.

Once the fuzzy variables have been created, they are used to fire a series of fuzzy rules. The result of this process is a set of design actions which has been ordered based on which actions are most likely to improve the performance of the current design based on the user's goals and preferences. The rules are fired in four steps:

1. Determine priority of each sensor. For example, IF SensorPerformance is Bad AND UserPriority is High, THEN SensorPriority is High.

2. Determine which change(s) will improve performance, based on the current scenario. For example, IF SensorPriority is High AND SensorType is Illuminance AND IlluminancePerformance is TooLow: (a) IF distanceFromGoal is Far, THEN DesiredChange is "Increase Illuminance by a Large Amount"; (b) IF distanceFromGoal is Close, THEN DesiredChange is "Increase Illuminance by a Small Amount".

3. Evaluate each possible design action in the customized database using the desired changes determine in Rule Base 2. For example, IF DesiredChange is "Increase Illuminance by a Large Amount" AND ActionResult is LargeIncrease, THEN action is GoodForSensor. These rules are fired once per potential action, and once per sensor.

4. Each potential action is ranked based on how likely it is to improve each sensor and the sensor priorities.

The final step is to sort the set of design actions from highest to lowest rank. The first design actions in the list will be those actions most likely to produce positive performance results in the current design, while those actions at the end of the list are likely to decrease overall performance.
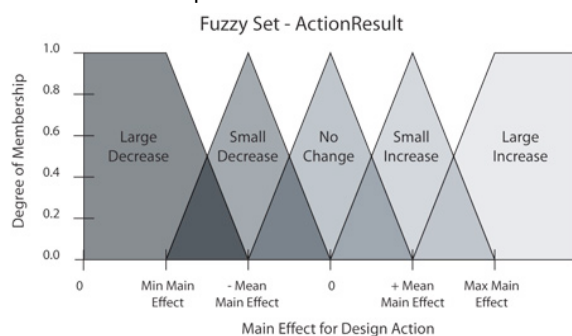

Figure 1: Membership functions for ActionResult fuzzy set.

2.4. System Implementation and Process

The expert system has been implemented within the framework of the Lightsolve project. Google SketchUp [17] is used as the 3d modeller, and the embedded Ruby application programming interface (API) within SketchUp is used to create pop-up interfaces which allow the user to enter the initial inputs and to perform the major processes and calculations. The LSV simulation engine is a stand-alone executable which is called directly from within the SketchUp/Ruby environment.
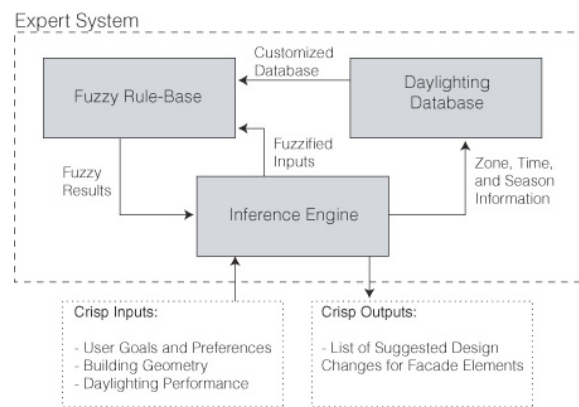

Figure 2: Schematic diagram of expert system process.

The expert system has a functional, stand-alone interface which allows designers to interact with the system (Fig. 3), which has been implemented using Adobe Flash. The interface has been designed to provide an intuitive and clear way of communicating the current performance of a design and the list of changes suggested by the expert system. The interface also allows designers to view the performance of their design over multiple iterations of the exploration process.

The overall expert system process is shown in Figure 2 and consists of the following steps:

1. The user creates an initial 3d model of a design with illuminance and/or glare sensor planes and specifies all necessary initial inputs to the system (using pop-ups in SketchUp).

2. Daylighting performance for the current model is calculated using the LSV engine based on the user's illuminance and glare goals.

3. The knowledge-base described in section 2.1 is used to create a customized database which contains only the information most relevant to the current design.

4. Information about the user's preferences, the original 3d model, the current performance, and the customized knowledge-base is used to create the fuzzy variable sets.

5. Fuzzy rules are fired using the fuzzy variables. The results are a set of suggested design changes that the system will propose to the user in order to improve performance.

6. Results are presented to the user in the user interface (Fig. 3).

7. The user selects a design change to make, and a new 3d model is created automatically. The process begins again starting at step 2.

2.5. An Expert System Design Process

The user's process begins when he or she creates an initial 3d model in SketchUp and initiates the expert system. Once the first set of simulations is complete, the user interface will automatically open. From there, the user's design process is as follows:

1. From the expert system interface, the user can view a list of suggested design changes that can be made to his or her initial model. The user may skip

forward or go backwards between the various options on the list before choosing one.

2. After the user selects one design change to try, the expert system will automatically make the selected change to the 3d model, which should still be open in SketchUp. The expert system will make three different magnitudes of the selected change. For each change, the expert system will create and save a new 3d model, run the LSV engine, and calculate the goal-based performance.

3. After the three different magnitudes of change have been simulated, the expert system will display all three results in the interactive graph within the interface. The user may browse the views of the current design and the temporal maps to see how the performance and design have changed in each of the three options. The user must choose one of the three possibilities before continuing to the next design iteration.

4. After one or more design iterations have been made, the user may then choose either to select a new design change to try from the list presented by the expert system, or the user may return to a previous iteration of the design (including the initial model). If the user elects to make another design change, steps 2 and 3 repeat.

5. After several iterations, the user should be able to view the progressive performance of the design. The user may stop the process at any point.

3.   EXPERT SYSTEM EVALUATION

The main function of the expert system described in this paper is to effectively guide a user towards improved daylighting performance of an original design. It is of critical importance that users have confidence in the advice given to them by the system, so a high level of performance is essential. Although the expert system differs from a traditional optimization algorithm due to its domain-specific and user-interactive nature, it should be capable of performing similarly to an optimization algorithm in a best case scenario.

In order to assess the behaviour of the expert system, a series of case studies were completed which compare the performance of designs found using the expert system to high performing benchmark designs generated using a genetic algorithm (GA). This paper will describe the results of two case studies, which both have two illuminance goals. These case studies were considered for Boston, MA (USA). Although they are not presented here, additional case studies were also completed which consider other situations, such as conflicting illuminance and glare goals. These studies can be found in [18].

The GA used in these case studies was a micro-genetic algorithm [19], which is a GA which uses a very small population size. For comparison purposes, the micro-GA was implemented within the Lightsolve system and uses the same 3d models and performance metrics as the expert system. This system is described in more detail in [20].

3.1.   Case Study Procedure

A set of study procedures was developed to better compare results from the expert system to the GA, given their differences in algorithm type. While a GA is one that generates designs, the expert system always assumes that an initial design is given and suggests design changes based on the current design. The following procedure was used:

• *Micro-GA procedure*: An initial massing model with no windows was used to generate a new model
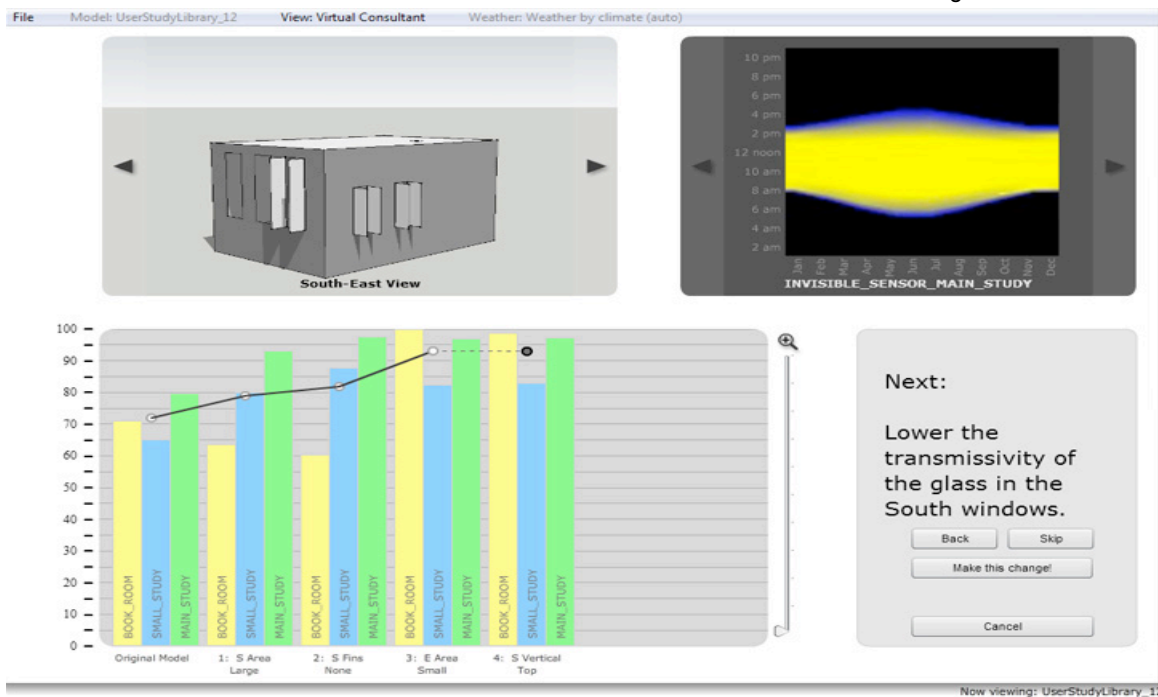

Figure 3: Performance analysis and decision making interface for the expert system. Views of the current design are shown (top left) along with annual performance in temporal map form (top right). Performance over multiple iterations is shown in the interactive graph (lower left). Expert system design suggestions are given in the lower right.

of each generated design. The algorithm was run for ten generations before stopping. If a perfect solution was not found, the best design was considered that with the highest performance found over all generations.

• *Expert system procedure*: An initial model was created with generic rectangular windows. This initial model was designed to be of mediocre performance, so as to avoid starting out with an initial design whose performance was very poor or very good. For these case studies, a "perfect user" was assumed. The "perfect user" was defined as one who would select the first suggested design change at each iteration and the best performing magnitude of each design change. The "perfect user" scenario was also one in which the process continued even if performance decreased after a given design iteration. The algorithm was run for ten design iterations before stopping. As with the GA study, if a perfect solution was not found, the best design was considered that with the highest performance found over all completed iterations.

3.2. Case Studies

Figure 4: Massing model and sensor plane locations for L-shape case study.

This paper will present two case studies, which both have two illuminance goals. The first case study features an L-shaped space (Fig. 4) where the two sensor planes are located roughly parallel to the

facades of interest (west and south). The performance goals for this case study were:

• South zone: 400 lux minimum preferred (200 lux accepted); No maximum.

• West zone: No minimum; 500 lux maximum preferred (800 lux accepted).

Based on these goals, the known design solutions to this problem featured small, shaded windows on the west facade and larger windows on the south façade.


Figure 5: Massing model and sensor plane locations for trapezoidal case study.

The second case study features a trapezoidal space (Fig. 5) where the two facades of interest, north and south, are perpendicular to the two sensor planes. The performance goals for this case study were:

• East zone: 200 lux minimum preferred (100 lux accepted); 800 lux maximum preferred (1200 lux accepted)

• West zone: 400 lux minimum preferred (200 lux accepted); No maximum.

For this case study, it was assumed that good solutions would have windows on both facades shifted towards the west sensor.

For both case studies, the best performing designs found after ten generations or ten design iterations are shown in Figure 6. For the L-shaped space, both the expert system and the micro-GA were able to find designs which were close to
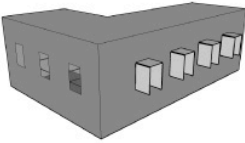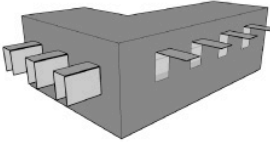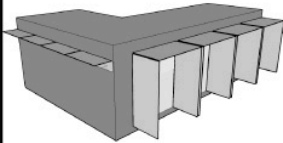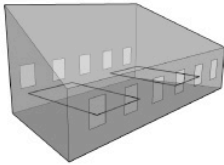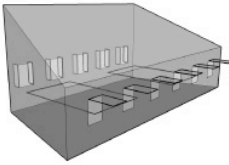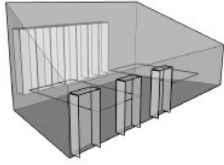
Performance Goals	Expert System Starter Design	Expert System Best Design	Micro-GA Best Design
South Zone: High Illuminance West Zone: Low Illuminance			
	64.5% In Range	96.1% In Range	95.3% In Range
West Zone: High Illuminance East Zone: Constrained Illuminance Range			
	61.3% In Range	82.6% In Range	87.0% In Range

Figure 6. Average performances for the starter expert system design, final expert system design, and final micro-GA design for both case studies.

meeting the performance goals entirely. As expected, both "best" designs have either very small or highly shaded windows on the west facade with larger or less shaded windows on the south facade.

For the trapezoidal case study, both algorithms had more difficulty finding good solutions. In this case study, the micro-GA was able to find a solution which performed about 5% higher than the expert system. This difference is due to the window uniformity scheme selected for the expert system (all windows on the facade must be uniform) and the univariate ("step-by-step") nature of the expert system algorithm. While the micro-GA found a design solution that features windows clustered towards the west end of both facades as expected, the expert system focused on changing the properties of the windows without moving them.

These case studies demonstrate that the expert system is successful at improving the performance of designs for two illuminance goals. The difference in performance between the expert system and the GA was small (4.4% at most) and acceptable given the fact that the expert system was designed with user interactivity in mind, while the GA was not.

4. CONCLUSIONS

This paper presented a new user-interactive expert system approach which enables architects to consider daylighting goals in the early design stages by engaging them in a performance-driven design exploration process. The expert system was shown to be successful at making design decisions which improved the daylighting performance of two case study designs. In both of these case studies, the performances of designs found using the expert system were comparable to those generated by a micro-genetic algorithm (micro-GA).

In addition to the case studies presented in this paper, additional case studies which consider more complex scenarios such as conflicting illuminance and glare goals were also completed. The expert system has also been tested on a group of designers who were asked to complete a design task with the system and to evaluate their experiences using the tool. These additional results will be presented in future papers.

5. REFERENCES

[1] Rashid, M. and Zimring, C., 2008. A review of the empirical literature on the relationships between indoor environment and stress in health care and office settings: Problems and prospects of sharing evidence. *Environment and Behavior*, 40(2), pp.151–190.

[2] Edwards, L. and Torcellini, P., 2002. A literature review of the effects of natural light on building occupants. NREL.

[3] Boyce, P., Hunter, C., and Howlett, O., 2003. The benefits of daylight through windows. U.S. Department of Energy.

[4] Boyce, P., Heerwagen, J., Jones, C., Veitch, J., and Newsham, G., 2003. Lighting quality and office work: A field simulation study. Pacific Northwest National Laboratory.

[5] Luger, G., 2004. Artificial Intelligence: Structures and Strategies for Complex Problem Solving. 5th edition. Boston: Addison-Wesley.

[6] Jung, T., Gross, M, and Do, E., 2003. Light pen: sketching light in 3D. Proceedings of 10th International Conference on Computer Aided Architectural Design Futures, Taiwan.

[7] Guo, B., Belcher, C., and Roddis, W., 1993. RetroLite: an artificial intelligence tool for lighting energy-efficiency upgrade. *Energy and Buildings*, 20(2), pp.115–120.

[8] Paule, B. and Scartezzini, J., 1997. Leso-DIAL, a new computer based daylighting design tool. *Right Light* 4(1), pp.93–97.

[9] Ochoa, C. and Capeluto, I., 2009. Advice tool for early design stages of intelligent facades based on energy and visual comfort approach. *Energy and Buildings*, 41(5), pp.480–488.

[10] Andersen, M., Kleindienst, S., Yi, L., Lee, J., Bodart, M., and Cutler, B., 2008. An intuitive daylighting performance and optimization approach. *Building Research & Information*, 36(6), pp.593–607.

[11] Montgomery, D., 2004. Design and Analysis of Experiments. 6th edition. John Wiley & Sons.

[12] Kleindienst, S. and Andersen, M., 2009. The adaptation of daylight glare probability to dynamic metrics in a computational setting. In Proceedings of LuxEuropa, Istanbul.

[13] Cutler, B., Sheng, Y., Martin, S., Glaser, D., and Andersen, M., 2008. Interactive selection of optimal fenestration materials for schematic architectural daylighting design. Automation in Construction, 17(7), pp.809–823.

[14] Gagne, J., and Andersen, M. A daylighting knowledge-base for performance-driven facade design exploration. *LEUKOS* (submitted).

[15] Siler, W. and Buckley, J., 2005. Fuzzy expert systems and fuzzy reasoning. John Wiley & Sons.

[16] Wienold, J., 2009. Dynamic daylight glare evaluation. In Proceedings of Building Simulation, Glasgow.

[17] Google SketchUp, 2010. [online] Available at: http://sketchup.google.com/ [Accessed November 8, 2010].

[18] Gagne, J., 2011. An interactive performance-based expert system for daylighting in architectural design. PhD thesis, MIT.

[19] Krishnakumar, K., 1989. Micro-genetic algorithms for stationary and non-stationary function optimization/ SPIE Proceedings: Intelligent Control and Adaptive Systems.

[20] Gagne, J., and Andersen, M. A generative façade design method based on daylighting performance goals. Journal of Building Performance Simulation (in press).