Improvements to Building Energy Usage Modeling During Early Design Stages and Retrofits

Andrew Mandelbaum

B.S.E., Mechanical and Aerospace Engineering
Princeton University, 2012

Submitted to the Department of Mechanical Engineering
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Mechanical Engineering
at the
Massachusetts Institute of Technology

June 2014

Signature of Author: Signature redacted

Andrew Mandelbaum

Department of Mechanical Engineering

May 19, 2014

Certified by: Signature redacted

Leon R. Glicksman

Professor of Architecture and Mechanical Engineering

Thesis Supervisor

Accepted by: Signature redacted

David E. Hardt

Professor of Mechanical Engineering

Chairman, Department Committee on Graduate Students

Improvements to Building Energy Usage Modeling During Early Design Stages and Retrofits

by

Andrew Mandelbaum

Submitted to the Department of Mechanical Engineering

on May 19, 2014 in Partial Fulfillment of the Requirements for the

Degree of Master of Science in Mechanical Engineering

Abstract

A variety of improvements to the MIT Design Advisor, a whole-building energy usage modeling tool intended for use during early design stages, are investigated. These include changes to the thermal mass temperature distribution and lighting requirement calculations, the formulation of a representative day scheme to reduce the number of days that need to be simulated to find energy usage for a full year, the creation of an optimization algorithm to allow users to improve on their designs, and the addition of an algorithm to predict potential savings from retrocommissioning (RCx) using a limited set of simple inputs. Design Advisor itself is also tested for usability, speed, and accuracy using three existing buildings.

The frequency of thermal mass-related calculations is reduced by finding the limits of the semi-implicit Crank-Nicolson method before it begins to return physically implausible oscillatory temperature profiles. An effort is made to speed up lighting calculations using a multivariate regression analysis in place of a multiple-reflection-based illuminance model. Representative days are formed by creating an average and two extreme weather days per month using existing climatological data, reducing the number of simulated days per year from 365 to 72 (three per month, repeated once for training). Combined, these changes lead to reductions in run time of up to 50% with roughly 10% loss of accuracy.

The optimizer leverages these run time improvements to rapidly find optimal building designs given a set of input constraints. Initially, a multistep multivariate regression is used to reduce the given search space and tighten the constraints. Then, a genetic algorithm is used to find the target solution. Initial tests of this combination have led to average reductions in energy usage of 25% given 6 minutes of calculation.

To extend Design Advisor's applicability to existing buildings, an algorithm for predicting potential energy savings from RCx is implemented and tested. A database of 90 buildings that have undergone an RCx process and had their resulting energy savings documented has been collected. A k-

nearest neighbors algorithm is used to evaluate the potential savings of test buildings based on this data set, operating on the assumption that similar buildings (in terms of location, size, and energy usage intensity) will present similar faults or opportunities for savings. While the average savings percentage prediction error is 0.02, the root-mean-square error is 12.4, which is greater than the actual savings potential of many buildings.

Model validation is performed using three existing buildings; two in the Philadelphia area and one on MIT's campus. For energy types for which no building faults or other issues were later found (as in the MIT building), final usage predictions are found to be accurate to within a mean bias error of -11.2% to 2.6%. To improve upon these accuracies, further details about key building parameters and modes of operation would be required. These studies also inform further usability improvements, including reporting site (rather than primary) energy usage and expending reported electricity usage to include loads other than lighting.

Thesis Supervisor: Leon Glicksman

Title: Professor of Architecture and Mechanical Engineering

Acknowledgments

Professor Glicksman, for introducing me to the world of building technology and thoughtfully guiding me through my research while giving me the freedom to make discoveries on my own.

My labmates and coworkers, for their myriad insights and endless practical knowledge.

Previous students who have worked on or used Design Advisor, for laying the groundwork for the research presented here.

MIT Sport Taekwondo, for being a much-needed source of stability and fun.

My parents, for instilling in me a love of learning that, with their support, has carried me this far.

Susan, for her love and help in times of need.

This work was sponsored by the Department of Energy through the Energy Efficient Building Hub at The Pennsylvania State University, under the Greater Philadelphia Innovation Cluster for Energy Efficient Building (GPIC) (Sponsor Reference #4762-MIT-DOE-4261).

Table of Contents

Table of Figures

Table of Tables

Introduction

Motivation

The buildings sector is one of the most significant drivers of energy usage in the developed and developing world today. In 2010, energy consumption by buildings in the United States alone (for heating, cooling, lighting, and other activities) accounted for 7% of global primary energy consumption (and over 40% of primary energy consumption nationwide) [1]. With the value of new construction in the United States increasing year-on-year [2] and the famed economic surges of the BRIC nations (among others) yielding entire new cities, moderating this energy usage will be vital if we as a global community are to set and meet effective energy efficiency goals.

The first step in controlling this energy usage is to understand it – what factors lead to energy consumption in buildings and how these factors interact with one another. The best way to accomplish this is through building modeling and simulation. While individual building subsystems or construction materials can be tested empirically using an environmental test chamber, this is not always feasible. These experiments can be costly and time-consuming, and it can be challenging if not impossible to scale results for an entire building. Furthermore, pressure from government regulations as well as changes to industry code have pushed designers and property managers more towards modeling and simulation in recent years [3].

There are numerous software platforms available that simulate the behavior of whole buildings and building subsystems. The US Department of Energy lists over 140 under the category of "Whole Building Analysis: Energy Simulation" alone [4]. Many are paid tools and some are now defunct or have not been updated recently. Much of the industry and development attention is focused on a select few programs, such as DOE-2, eQUEST, and EnergyPlus. We have found, both in our own experience and through conversations with building designers, that non-expert users are often put off by the initial complexity of these programs. Many can be simplified using additional frontend software, but the default skin can still be overwhelming. One example is the default launch window (and associated building definition file) for EnergyPlus, shown below.

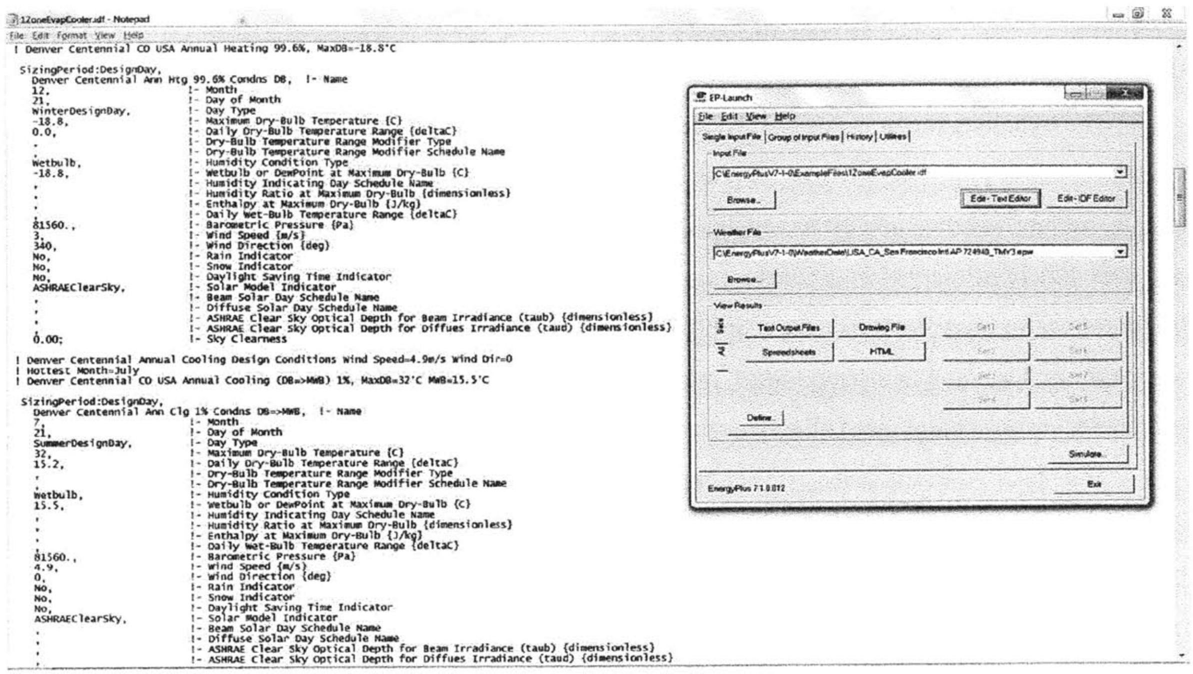

Figure 1: Default EnergyPlus launch window

While an experienced user might, just as with any other complex piece of software, know beforehand which options are germane to their work or how to install an additional frontend, not every user has the time or expertise to do that. Even if they do, the level of detail required for many of these models and the amount of time it can take to collect and input this information can lead to simulation being put off until late in the design process. While these late-phase simulations may be needed for regulatory compliance, "when modeling is only used late in design—after the massing, orientation, envelope and glazing design, and mechanical systems in a building are already specified, and hundreds of hours of work have already been put into those designs—the modeling might have little value beyond keeping score" [3]. Ideally, modeling and simulation is a part of design process from the very beginning, where it can provide valuable input that can be acted upon despite the number of still-unknown variables [5], [6]. Therefore, we believe that there is a need for simplified building energy modeling software that limits the number and specificity of the variables required from the user while retaining the most important first-order effects.

Building energy usage concerns do not disappear once a design is finalized and construction is complete. They must still be considered throughout the lifespan of a building, as systems degrade and occupant practices change. One way for these concerns to be addressed is through retrocommissioning (RCx). Building commissioning is analogous to ship commissioning, wherein a commissioned ship is "one deemed ready for service" through a variety of tests and checks. Retrocommissioning applies this notion to already-existing buildings, and is aimed at finding and fixing faults in components and procedures that arise naturally over time. Another method of reducing energy usage in existing buildings is through retrofitting, in which machinery and infrastructure is upgraded using more efficient or more reliable technology. While an RCx investigation can lead to a retrofit, the initial intention is often to keep the building as is but fix what is necessary. As a result of this purposeful limitation of scope, RCx initiatives generally have very short payback periods and rarely lose money for the building owner in the long term.

As with building energy usage simulation, there are certain barriers to be overcome in the world of retrocommissioning. One is knowledge of the process itself. While building commissioning is not an especially new concept (for example, ASHRAE first formed a committee on the topic in 1982 [7] and the National Conference on Building Commissioning is currently in its 22nd year [8]), its application to existing buildings is a relatively recent development (ASHRAE's first guideline committee for existing building commissioning was formed in 2007 [7]). As a result, RCx is very much in a growth phase and it is easy to find a variety of publications from trade magazines [9], non-profits [10], and government labs and agencies [11], [12] explaining RCx and promoting its use. Property owners and managers are unlikely to make use of RCx without understanding the process and its benefits, so this multi-pronged approach is necessary to reach the widest swath of the target audience.

That said, knowledge of RCx alone is not enough. The second obstacle is motivation. A successful RCx project that yields long-lasting results requires multiple steps: a well-researched planning phase, a thorough investigation phase that involves testing, monitoring, and documentation, implementation of the identified faults, and a hand-off phase that helps property managers and their workers to properly maintain the building for years to come [13]. Given the increasing number of RCx providers going into business, simply starting this process can be daunting for the inexperienced, especially if they are unsure of their potential return on investment (of both time and money). To solve this, both governments and energy providers have been working on requirements and incentives to encourage RCx.

Of particular interest is a recent legislative push that has led to the integration of RCx into building performance standards in New York City and San Francisco. The New York law, Local Law 87, was enacted in 2009 and "mandates that buildings over 50,000 gross square feet undergo periodic energy audit and retro-commissioning measures" [14]. An energy report must be submitted to the city every 10 years. A previous law, Local Law 84, mandated annual energy benchmarking [15], but LL87 explicitly discusses RCx and provides resources for building owners and managers to learn about and get in touch with approved energy auditors and retrocommissioning providers.

San Francisco's 2011 law, the Energy Commercial Buildings Energy Performance Ordinance, requires an "energy efficiency audit once every 5 years identifying specific cost-effective measures that would save energy" for "existing nonresidential buildings 10,000 square feet and larger" [16]. While the law, unlike New York's, does not call for an actual RCx process, the city does accept a variety of retrocommissioning certifications as proof of qualification and competence for the auditors in question [17]. This fact, combined with the savings identification requirement from the law, should encourage property managers to seek out RCx providers on their own.

The U.S. Department of Energy has seized upon these two laws as well-implemented examples of state and local laws intended to promote retrocommissioning [18]. Local governments can enact policies that best fit the buildings and services in their own jurisdictions and can also test out these policies on their own buildings (which are, on average, "nearly 25% more energy-intensive than non-government-owned buildings" [18]). On the other hand, each jurisdiction must write and pass these laws through their own legislative systems, meaning that successful trailblazing attempts (especially by such large cities) are important for the instruction they provide.

As previously mentioned, utilities have also been working on incentive programs, even in areas where there has yet to be any legislative pressure. Three examples from across the country are programs by Connecticut Light and Power [19], ComEd in Illinois [20], and Pacific Gas and Electric in California [21]. Plans like these offer property owners payments based on energy savings if they undergo and implement the fixes recommended by a well-documented RCx project. For example, through PG&E, participants can be paid roughly "$0.09/kWh, $1.00/therm, and $100/on-peak kW, capped at 50% of the total project cost" (based on achieved savings) as well as receive some engineering support during the process. To qualify, the building must be of a certain size or consume a certain amount of energy per year. The participants must also meet certain standards, such as a willingness to "spend typically between 8-16 hours documenting their facility's energy usage involving multiple site visits" and "spend a minimum of $25,000 on all reasonable and eligible RCx measures identified as having a simple payback less than or equal to one year" [21]. There are additional guidelines for the RCx process itself. The other utility companies have similar guidelines (with varying rebate rates and eligibility requirements). Given that, in the case of PG&E, current commercial electricity rates range from $0.15 to $0.23/kWh and commercial gas costs roughly $0.90/therm [22], these refunds can be powerful economic incentives. Additionally, given the clause concerning simple payback time, property owners are encouraged to act on issues for which they will see results the fastest, mitigating some of the uncertainty they may feel.

To further disseminate knowledge and understanding of RCx and to promote its use, we would like to apply the same philosophy of simplicity and usability that drives Design Advisor to the world of commissioning. If, using a minimal set of inputs, building owners could be given a rough sense of potential energy or monetary savings through RCx and then pointed to further resources, the previously discussed barriers to implementation would be lowered even further. The specifics of both this work as well as the intended improvements to Design Advisor are detailed in the "Goals" section below.

Existing Work

The MIT Building Technology Lab has done work in the past on promoting and facilitating the consideration of energy usage during the design and planning stages of building construction. This work culminated in the development of the MIT Design Advisor [23], a simplified and streamlined building modeling tool. Available free of charge online, this web-based tool is intended for use by architects and other non-expert users who may not have the time or resources to learn and use the more complex modeling suites previously mentioned. As has been noted, there are many other modeling programs available that simulate the whole or parts of a building. As the goal of the building modeling portion of this work is to further develop and improve Design Advisor, it will be the focus of this section.

Design Advisor is a white box simulation platform – its predictions are based primarily on the fundamental energy balances and heat transfer processes governing the behavior of the building, and the outside assumptions and approximations that are made are reported. This, combined with the fact that only those inputs that affect these calculations to first order are considered, allows for much of Design Advisor's simplicity. As all of the major design factors of the building are retained and the basic equations are sound, Design Advisor is capable of producing energy usage predictions comparable to those from Energy Plus – for an idealized building, the two have been shown to agree to within 10% [23]. Avoiding the use of highly specific and complex correlations and reducing the number of inputs required from the user does limit Design Advisor's accuracy, especially when modeling real-world buildings. Factors such as building geometry, air distribution, and human behavior can only be simplified so far, especially considering the fact that uncertainties in these inputs can lead to prediction errors much larger than 10% even when using more complex pieces of software. However, if these limitations are understood and accepted, Design Advisor can still be a very powerful and instructive tool.

Design Advisor's setup screen and a sample results page are shown in Figure 2.

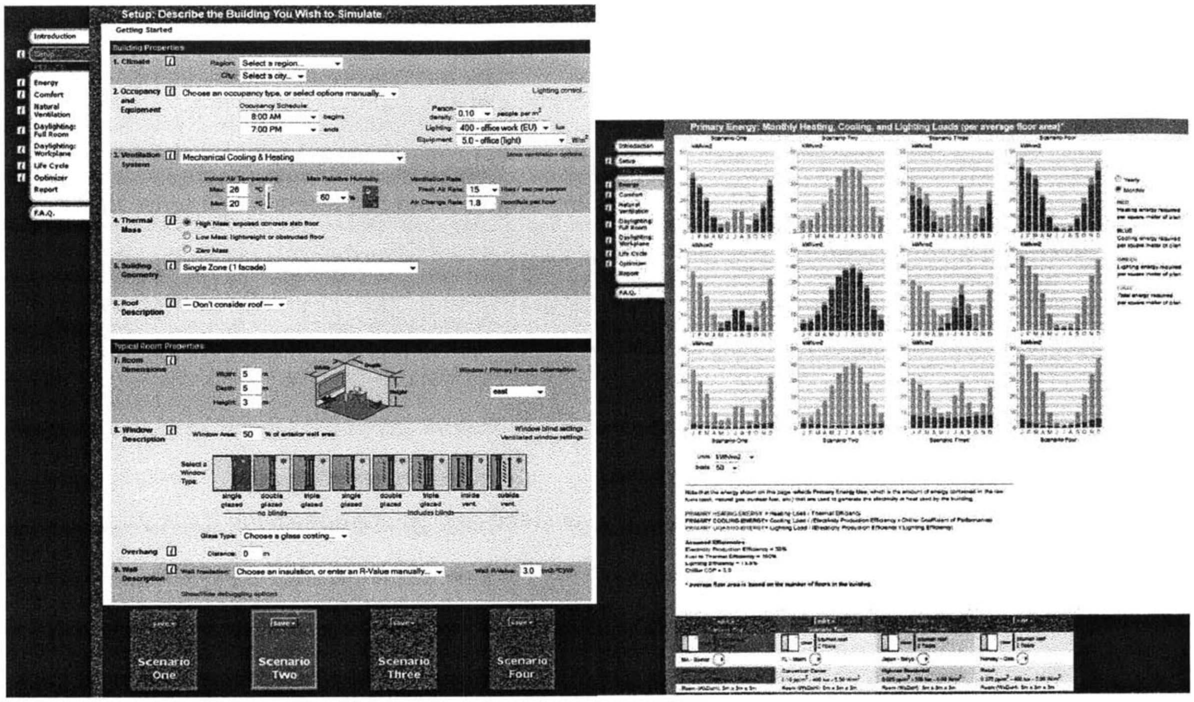

Figure 2: Setup and energy usage results screen from Design Advisor

The various inputs are broken down into groups by type, with presets for many of them based on building type (for occupancy schedule, plug loads, person-density, and lighting and air change requirements) or window type (for blinds and ventilated window unit settings). Generalized drop-down menus are used where applicable to guide the user and avoid specific details that may be unknown at the time of simulation. Users select a region and input information about the building usage, ventilation, thermal mass, building geometry, roof, and a typical room (that is tiled around the perimeter of the building for four-sided simulations). Up to four simulations can be saved concurrently, allowing for direct comparison between different designs, as seen in Figure 2. Viewable results include yearly and monthly energy usage (broken down by heating, cooling, and lighting), occupant comfort metrics, natural ventilation temperature histograms, daylighting simulations, CO_2 production, and life cycle costs. Users can also download text-based reports offering a greater level of detail. Relative to popular simulation suites, Design Advisor is fast, easy to use, and reasonably accurate.

Retrocommissioning, owing to its relative youth as a concept, has not been explored as fully as building simulation. Much of the work on savings prediction has been done by RCx providers themselves. Their systems, once installed, monitor building performance and aid in the detection of potential faults and notification of building managers if the building is not behaving as expected. This is known by a multitude of names, including ongoing/continuous commissioning [24], [25] and monitoring-based commissioning (MBCx) [20], [26]. These systems can be extremely powerful, customizable, and informative, allowing building managers to find and fix problems before they worsen. For example, KGS Building's system, when installed in a wet laboratory, identified problems whose repairs "led to $286,000 per year in operational energy cost savings, with additional benefits to the customer derived from utility incentive payments from documented savings." This involved "monitoring and analysis on more than 13,000 data points, including 24 air-handlers, over 700 terminal boxes and fume hoods, and many other pieces of HVAC equipment" [24]. The benefits, especially for large, energy-intensive buildings, are clear. However, implementation of such a system can be very complex (depending on existing infrastructure) and time-consuming for managers of small buildings or limited property portfolios. It also does not solve the problem of building managers not being able to quantify possible savings before contacting and working with a commissioning provider.

One interesting piece of work is a spreadsheet that was put together by the California Commissioning Collaborative. It "targets commercial buildings and allows providers to calculate energy and peak demand savings for thirteen common controls- and schedule-based commercial building optimization measures" [27]. Based on a variety of inputs, including basic characteristics of the building and its systems as well as baseline performance measures and metric concerning proposed changes, the spreadsheet calculates potential savings from correcting various system sensors, schedules, and settings. While this can be a useful tool, it also fails to address the problem at hand. Not only is it intended for use by commissioning providers (and asks for proposed system changes in order to make its estimates), but it also produces its estimate using a large lookup table containing the results from an exhaustive parametric set of building simulations. This limits the accuracy of the estimates to the accuracy of the simulations themselves. Such a program that makes use of historical RCx data rather than simulations has yet to be created. This could be due to the lack of a comprehensive data set of verified real-world RCx savings, as addressed in the "Data Collection" section.

Goals

Given all of this, there are three primary goals of this work. The first is to reduce the run time required for a single simulation. This would have multiple benefits, chief among them being improved usability. Multiple studies have shown that web users are put off by seemingly minor delays in information retrieval and feedback. [28], [29] The run times for a variety of simulations of a classroom building located in Boston in the current version of Design Advisor are shown below in Table 1. For all non-specified variables, Design Advisor's defaults were used.

Cooling	Geometry	Floors	Roof	Window	Run 1 (s)	Run 2 (s)	Run 3 (s)	Average Run Time
Mech	Single zone	One	None	Single glazed	9.9	9.3	9.6	**9.6**
Mech	Single zone	One	None	Inside vent.	69.5	72.4	79.9	**73.9**
Mech	Single zone	Five	Bitumen	Single glazed	23.7	24.2	24.2	**24.0**
Mech	Single zone	Five	Bitumen	Inside vent.	149.4	145.9	165.2	**153.5**
Mech	4 zones+core, well-mixed	One	None	Single glazed	48.9	45.4	46.7	**47.0**
Mech	4 zones+core, well-mixed	One	None	Inside vent.	300.4	293.1	343.7	**312.4**
Mech	4 zones+core, well-mixed	Five	Bitumen	Single glazed	119.8	122.2	121.3	**121.1**
Mech	4 zones+core, well-mixed	Five	Bitumen	Inside vent.	709.4	708.5	610.5	**676.1**
Natural	Single zone	One	None	Single glazed	10.5	9.9	10.1	**10.2**
Natural	Single zone	One	None	Inside vent.	71.9	80.2	70.3	**74.1**
Natural	Single zone	Five	Bitumen	Single glazed	25.3	25.5	25.4	**25.4**
Natural	Single zone	Five	Bitumen	Inside vent.	148.1	150.1	153.7	**150.6**
Natural	4 zones+core, well-mixed	One	None	Single glazed	49.4	49.7	50.7	**49.9**
Natural	4 zones+core, well-mixed	One	None	Inside vent.	292.9	347.3	306.5	**315.6**
Natural	4 zones+core, well-mixed	Five	Bitumen	Single glazed	129.5	126.9	134.1	**130.2**
Natural	4 zones+core, well-mixed	Five	Bitumen	Inside vent.	720.2	623.5	644.2	**662.6**

Table 1: Run times for selected simulations in current version of Design Advisor (seconds)

The simplest simulations (of a single shoe box zone with no roof) can take 10 seconds or less. More complex simulations, involving heat transfer through a roof as well as air mixing between four zones and a central core, can take several minutes. Ventilated window units, included for completeness despite being a niche application, increase simulation times by a factor of five or more. While it is understandable that a full-year energy usage simulation would take longer than loading a web page, that is no reason to not strive to improve the user experience. A more convenient tool is more likely to be used, especially when other alternatives are available. This is important considering Design Advisor's original purpose of facilitating and popularizing early-stage building modeling.

Improving Design Advisor's performance feeds into the second goal: implementing a form of model optimization. Optimization can allow the user to reach the most energy-efficient design, both for new buildings and retrofits. It can provide suggestions for the most efficient new design or for retrofit actions that the user can explore in further depth. Finding even a local energy usage minimum, however, can require thousands of iterations for a system as complex as a building. While users may be willing to wait several minutes for a single simulation, it is unreasonable to expect them to leave Design Advisor open for a day or longer without any interruption between their computer and a remote server. Therefore, a "fuzzy" optimization method must be developed, one which sacrifices some accuracy (within the spirit under which Design Advisor was developed) and may not reach the absolute most efficient design by reducing the number of simulations that must be performed. If the remaining simulations can be made faster, the total computation time may be reduced to a more acceptable level. The keys are then to define meaningful target criteria and select an acceptable final error.

The final goal is the development and implementation of an algorithm to predict potential energy and monetary savings from retrocommissioning based on real building data. While the development of Design Advisor has focused in the past on improving inchoate building designs, expanding its functionality and enhancing its appeal to other elements of the buildings industry necessitates examining existing buildings as well. As previously discussed, while retrocommissioning has been rapidly expanding for several years, city and state governments around the country are beginning to recognize its value and mandate regularly scheduled RCx projects for certain types of buildings. This has prompted a growth in the number of commissioning firms as well as renewed attention from certification and standards groups. The RCx savings prediction algorithm is therefore intended to aid property owners and managers in navigating these new regulations and interacting with commissioning professionals. It is also, much like Design Advisor itself, intended to generate interest (especially among non-expert users) in energy savings measures. Lastly, it is our goal that the prediction algorithm derive its results using data from completed and documented RCx projects. Simulations, as already discussed, are limited by a wide variety of factors. Historical data implicitly include these factors and reflect actual savings.

Combined, these measures represent a variety of improvements and additions to the Design Advisor simulation engine. Their implementation and the results from the changes made are detailed in the following sections.

Reducing Run Time

Thermal Mass Calculations

The most computationally-intensive portion of a given day's simulation involves calculating the heat flux in and out of the thermal mass in the floor and roof (if applicable) of the building. For a typical building, it is assumed that each floor is a concrete slab and comprises the bulk of the thermal mass in any given room. The slab itself is assumed to be semi-infinite along the expanse of the room (that is, edge effects near the walls are neglected) and of finite depth (specified by the user). This depth is discretized such that each slice can be assumed to have a uniform temperature at any given point in time as per the Biot constant:

$$Bi = \frac{hd}{k} \ll 1 \tag{1}$$

$$Fo = \frac{k\Delta t}{\rho c d^2} < \frac{1}{2} \tag{2}$$

$$d_{max} = Bi * \frac{k}{h} \tag{3}$$

$$\Delta t_{max} = Fo * \frac{\rho c d_{max}^2}{k} \tag{4}$$

Using these slices, the temperature distribution in the thermal mass is then calculated using the Crank-Nicolson method. The Crank-Nicolson method is one of several schemes for numerically evaluating the one-dimensional diffusion equation:

$$\frac{\partial T}{\partial t} = D\frac{\partial^2 T}{\partial x^2} = \frac{k}{c_p \rho}\left(\frac{\partial^2 T}{\partial x^2}\right) \tag{5}$$

The one-dimensional form of the equation is used because the concrete slab is assumed to be significantly more expansive in the horizontal plane than it is deep. The thermal diffusivity of the slab is assumed to be constant in both time and space.

A simple and common numerical method that can be used to find an approximate solution to this equation is Euler's method. Euler's method comes in two main forms: forward time centered space (FTCS) and backward time centered space (BTCS) [30].

FTCS:

$$\frac{T_j^{n+1} - T_j^n}{\Delta t} = D\left[\frac{T_{j+1}^n - 2T_j^n + T_{j-1}^n}{(\Delta x)^2}\right] \tag{6}$$

BTCS:

$$\frac{T_j^{n+1} - T_j^n}{\Delta t} = D\left[\frac{T_{j+1}^{n+1} - 2T_j^{n+1} + T_{j-1}^{n+1}}{(\Delta x)^2}\right] \tag{7}$$

Where n is the time step index and j is the slice index.

As the names imply, the difference between the two lies in the dependence on the change in temperature across a time step (on the left-hand side of the equation) on the temperature difference within the slab during the current (FTCS) or next (BTCS) time step. As a result, the FTCS Euler method is fully explicit, while the BTCS form is fully implicit. Both have their own advantageous and disadvantages. The FTCS scheme is simple to understand, implement, and compute, even by hand. With that simplicity comes major shortcomings – the error in its approximation is proportional to the time step used and it can become unstable when the time step or slice thickness becomes too large [31]. The BTCS method, on the other hand, is unconditionally stable but requires solving a system of equations at each time step [32]. This requires more computation time and is generally more difficult to implement. It is also still subject to the same accuracy concerns as the FTCS version.

To improve on the accuracy of these methods, we turn to the Crank-Nicolson method. The Crank-Nicolson method is, in essence, an average of the FTCS and BTCS methods described above:

$$\frac{T_j^{n+1}-T_j^n}{\Delta t} = \frac{D}{2}\left[\frac{\left(T_{j+1}^{n+1}-2T_j^{n+1}+T_{j-1}^{n+1}\right)+\left(T_{j+1}^{n}-2T_j^{n}+T_{j-1}^{n}\right)}{(\Delta x)^2}\right] \tag{8}$$

Like the BTCS scheme, the semi-implicit Crank-Nicolson method is unconditionally stable and requires solving a system of equations at each time step [33]. Unlike either Euler method, it is susceptible to predicting non-physical (but not destabilizing) temperature oscillations at larger time steps or slice thicknesses [23]. The Biot criteria discussed above will prevent these oscillations by limiting the size of both of these factors.

In order to form the actual approximation of the temperature distribution at each time step, the n equations (one for each thermal mass slice) are combined into a matrix. The matrix equation used in Design Advisor is as follows [23]:

$$T^{t+\Delta t} = B^{-1}[S * T^t + Q\Delta t] \tag{9}$$

$$B = \begin{bmatrix} C + \frac{hA\Delta t}{2} + \frac{kA\Delta t}{2d} & -\frac{kA\Delta t}{2d} & 0 & 0 & 0 \\ -\frac{kA\Delta t}{2d} & C + \frac{kA\Delta t}{d} & -\frac{kA\Delta t}{2d} & 0 & 0 \\ 0 & -\frac{kA\Delta t}{2d} & C + \frac{kA\Delta t}{d} & -\frac{kA\Delta t}{2d} & 0 \\ 0 & 0 & -\frac{kA\Delta t}{2d} & C + \frac{kA\Delta t}{d} & -\frac{kA\Delta t}{2d} \\ 0 & 0 & 0 & -\frac{kA\Delta t}{2d} & C + \frac{kA\Delta t}{d} \end{bmatrix} \tag{10}$$

$$S = \begin{bmatrix} C - \frac{hA\Delta t}{2} - \frac{kA\Delta t}{2d} & \frac{kA\Delta t}{2d} & 0 & 0 & 0 \\ \frac{kA\Delta t}{2d} & C - \frac{kA\Delta t}{d} & \frac{kA\Delta t}{2d} & 0 & 0 \\ 0 & \frac{kA\Delta t}{2d} & C - \frac{kA\Delta t}{d} & \frac{kA\Delta t}{2d} & 0 \\ 0 & 0 & \frac{kA\Delta t}{2d} & C - \frac{kA\Delta t}{d} & \frac{kA\Delta t}{2d} \\ 0 & 0 & 0 & \frac{kA\Delta t}{2d} & C - \frac{kA\Delta t}{d} \end{bmatrix} \tag{11}$$

$$Q = \begin{bmatrix} \alpha Q_{solar\ incident} \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \tag{12}$$

Where C is the specific heat capacity of the slab, h is the convection heat transfer coefficient between the slab and the air in the room, k is the thermal conductivity of the concrete, d is the slice thickness, A is the area of the room, Δt is the time step size, α is the thermal absorptivity of the slab, and Q is the heat flux.

In the original Design Advisor code, conservative values are used for the Fourier modulus and the Biot constant (0.5 and 0.05, respectively) to ensure stable, physically plausible results. The table of values that determined this is shown below.

α	0.80
c	880 kJ/kg-K
ρ	2300 kg/m^3
k	1.4 W/m-k
h	10 W/m^2-K

Table 2: Physical properties of concrete slab [23]

Using these values in the relevant equations for the Biot criterion, this yields a maximum slice thickness of 0.007m.

	Thermal Mass Setting	
	High	Low
Slab thickness (m)	0.2	0.02
Max slice thickness (m)	0.007	0.007
Number of slices	28.57	2.86
Whole number of slices	29	3
Slice thickness (m)	0.006897	0.006667
Max time step (s)	34.38	32.13
Time step (s)	30	30

Table 3: Discretization of concrete slab [23]

As shown in Table 3, the values used in Design Advisor yield a time step of slightly more than 30 seconds. Both the number of slices and the time step size were rounded to allow for simpler whole-number calculations and to allow the thermal mass time steps to fit into the 60-second time step used in the rest of the program. However, given the long timescale over which the relevant heat transfer processes occur, these constants (and thus the computation time step) can be increased without a significant loss of accuracy. Through testing, it was found that, rather than twice every minute, heat transfer to and from the concrete slab could safely be calculated as infrequently as once every 15 minutes without instability. The effects from making this change are summarized in the figures and tables below.

A single-sided one-floor shoebox model room was simulated for a full year using the original Design Advisor code. The ambient air temperature and the midpoint temperature of the top thermal mass slice were extracted at the start of each hour on January 15 and July 15. From that point, to allow for a fair comparison, the only change that was made to the code was to reduce the number of times per hour that the heat flux from the thermal mass was recalculated (to 10 times per hour and 4 times per hour).

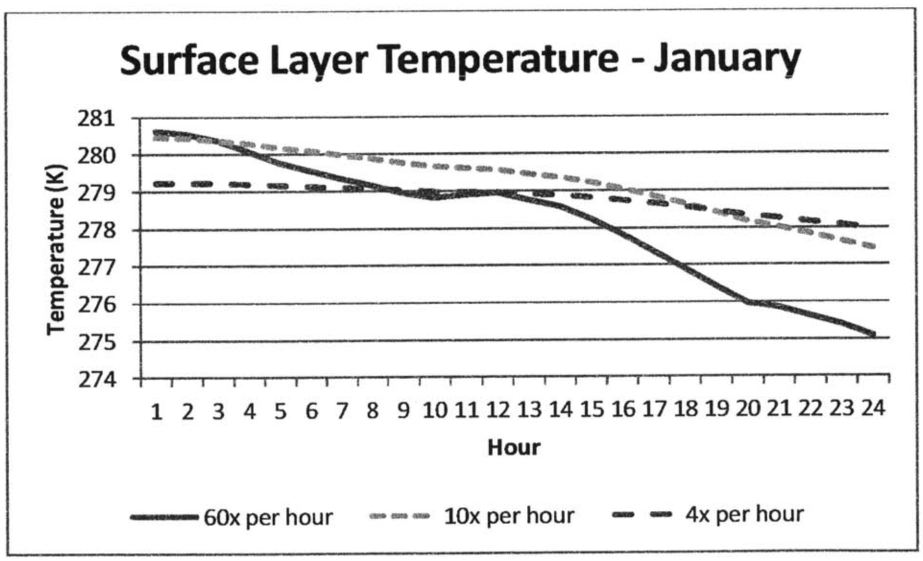

Figure 3: Thermal mass surface layer temperature in January

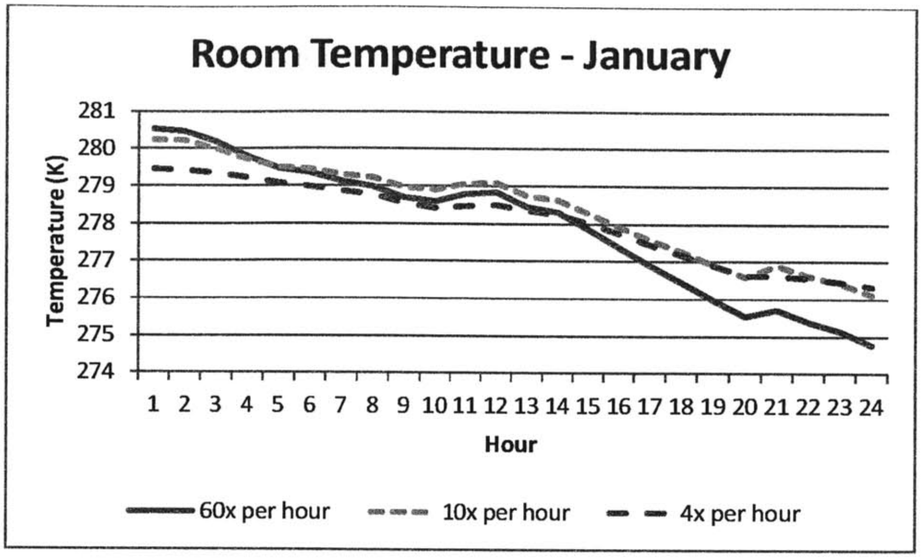

Figure 4: Ambient room temperature in January

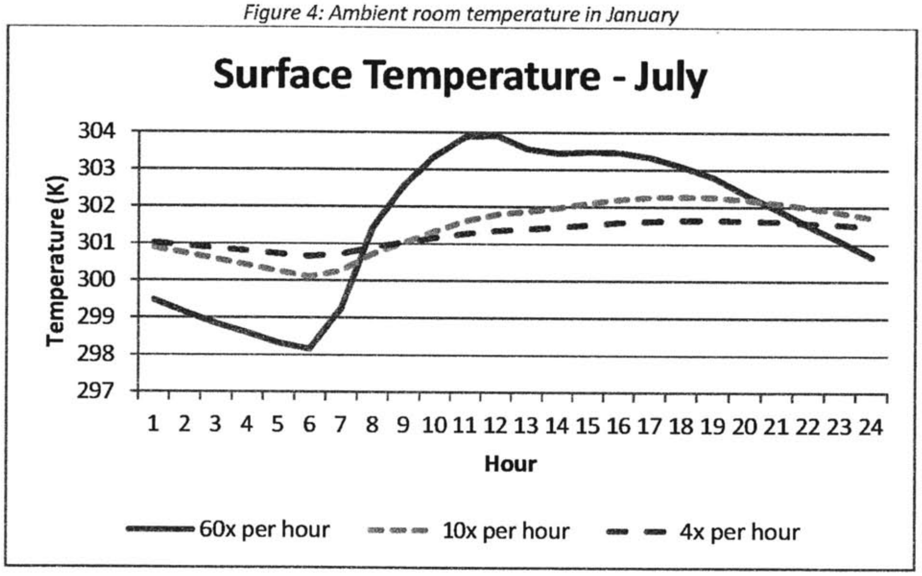

Figure 5: Thermal mass surface layer temperature in July

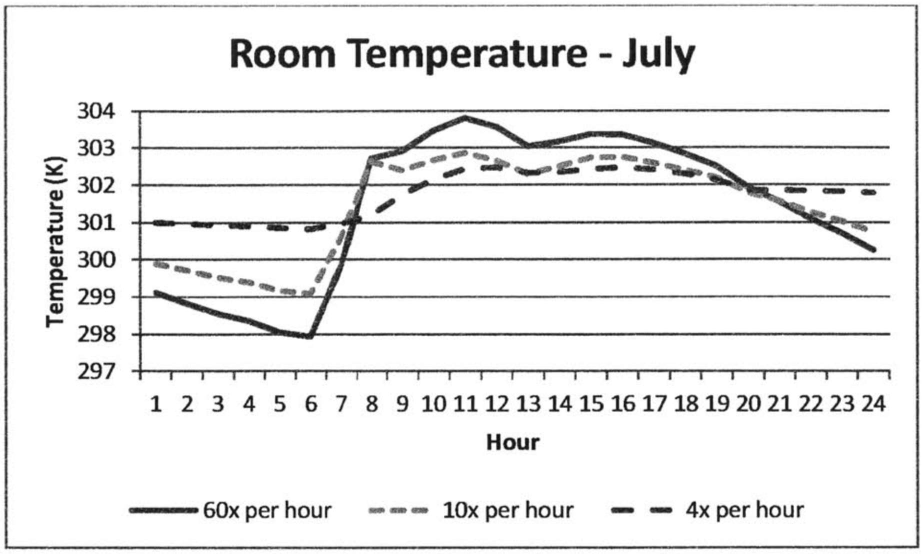

Figure 6: Ambient room temperature in July

January

	Max Error - Surface	Avg Error - Surface	Max Error - Room	Avg Error - Room
10x per hour	2.35	1.05	1.35	0.46
4x per hour	2.88	0.59	1.56	0.12

July

	Max Error - Surface	Avg Error - Surface	Max Error - Room	Avg Error - Room
10x per hour	1.95	-0.13	1.15	0.01
4x per hour	2.50	-0.32	2.88	0.33

Table 4: Temperature difference when compared to original method when using reduced calculations

Runtime

	Single Sided	Four Sided
60x per hour	9.02	38.22
10x per hour	8.14	33.25
4x per hour	6.78	27.14

Table 5: Runtime comparison between original method and reduced calculations

Variable	Value
City	Boston
Room Width	5
Room Depth	5
Room Height	3
Thermal Mass Thickness	Low
Window Area %	50

Window Type	Single glazed, no blinds
Glass Coating	Clear
Wall R-Value	1
Roof Type	Bitumen
Roof R-Value	2
Ventilation Type	Mechanical

Table 6: Important design parameters for thermal mass example

As demonstrated by Table 4 and Table 5, the potential reduction in runtime is effectively proportional to the level of error one is willing to accept. While on average the temperature predictions that arise from calculating the heat flux from the floor every 15 minutes are not significantly different from those when the flux is calculated every 6, the maximum error can be more than a degree larger. However, using 4 calculations per hour reduced the simulation time of a sample multi-floor building by 25% (single-sided) to 30% (four-sided), compared to 10% to 13%.

It must be noted that, despite these temperature differences, the resulting heating and cooling energy usage predictions (the only information reported to the user) for each of the modes were nearly identical:

Heating Energy (kWh/m^2)

	January	February	March	April	May	June
60x per hour	55.188	47.450	35.708	18.000	6.202	0.000
10x per hour	55.135	47.444	35.617	17.669	5.875	0.005
% Error	-0.096	-0.012	-0.255	-1.837	-5.277	1400.377
4x per hour	55.078	47.444	35.590	17.536	5.847	0.000
% Error	-0.200	-0.013	-0.329	-2.576	-5.731	-100.000

	July	August	September	October	November	December
60x per hour	0.000	0.000	0.223	13.101	26.909	49.137
10x per hour	0.004	0.004	0.104	12.951	26.813	49.132
% Error	Inf	Inf	-53.467	-1.145	-0.358	-0.011
4x per hour	0.000	0.000	0.000	12.542	26.763	49.129
% Error	Inf	Inf	-99.986	-4.264	-0.542	-0.016

Cooling Energy (kWh/m^2)

	January	February	March	April	May	June
60x per hour	0.011	0.000	0.000	0.000	0.106	4.335
10x per hour	0.011	0.000	0.000	0.000	0.048	3.238
% Error	0.000	Inf	Inf	Inf	-54.263	-25.301
4x per hour	0.011	0.000	0.000	0.000	0.042	3.303
% Error	0.000	Inf	Inf	Inf	-60.345	-23.804

	July	August	September	October	November	December
60x per hour	13.815	13.689	1.843	1.244	0.495	0.000
10x per hour	12.515	12.905	1.302	1.215	0.545	0.000
% Error	-9.407	-5.723	-29.348	-2.285	10.167	Inf
4x per hour	12.703	13.076	1.236	1.227	0.567	0.000
% Error	-8.046	-4.474	-32.914	-1.329	14.503	Inf

Table 7: Heating and cooling energy usage predictions using reduced thermal mass calculations

During the heating and cooling seasons (i.e., months during which heating or cooling energy usage is significant, respectively), the prediction error for either calculation reduction scheme (to 10 times per hour or 4 times per hour) does not exceed 8%. Given the simulation time comparison in Table 5, it appears to be a safe and logical choice to reduce thermal mass-related calculations to four per hour from 60, especially if, as in the case of the optimizer, the primary concern is speed.

An effort was also made to fully replace the matrix-based Crank-Nicolson calculations with a fully explicit forward time method adapted from MIT's CoolVent, a natural ventilation simulation program [34]. The equations used are given below.

For the surface slice:

$$T_{1,j+1} = T_{2,j} - \left(T_{2,j} - \left(T_{1,j} + \Delta t * \frac{\left(\frac{(T_{room} - T_{1,j})}{A\left(\frac{1}{h} + \frac{1}{2kd}\right)} + Q_{incident} \right)}{\rho A d C_p} \right) \right) * e^{-\frac{2hA\Delta t}{\rho A d C_p}} \tag{13}$$

For the bottom slice:

$$T_{n,j+1} = T_{n-1,j} - \left(T_{n-1,j} - T_{n,j}\right) * e^{-\frac{2hA\Delta t}{\rho A d C_p}} \tag{14}$$

For all other slices:

$$T_{i,j+1} = \frac{1}{2}\left(\left(T_{i-1,j} + T_{i+1,j}\right) - \left(\left(T_{i-1,j} + T_{i+1,j}\right) - 2T_{i,j}\right) * e^{-\frac{4hA\Delta t}{\rho A d C_p}} \right) \tag{15}$$

Finally, to find the heat into or out of the room:

$$Q_{net} = hA\Delta t * \left(T_{room} - \frac{T_{1,j+1} + T_{1,j}}{2}\right) \tag{16}$$

For all of these, i denotes the index of the thermal mass slice (from 1 at the surface to n at the bottom), j denotes the index of the time step, d is the slice thickness, A is the exposed area, and T_{room} is the air temperature during the current time step.

The results from this change, both in terms of time as well as accuracy, are summarized below. A representative set of runs were performed as a basic demonstration of the difference between the

two calculation strategies. Time step size and spatial discretization were left the same for both schemes.

Geometry	Floors	Crank-Nicolson Time (ms)	Explicit Time (ms)	% Difference
One-sided	1	8285	7604	8.2
One-sided	3	11698	10624	9.2
Four-sided	1	33423	29829	10.8
Four-sided	3	48499	44314	8.6

Table 8: Run time for Crank-Nicolson and forward time explicit thermal mass calculations

While changing the calculation method alone leads to an appreciable time savings (roughly 10% across various types of simulations, as shown in Table 8), the accuracy loss from doing so is greater than for simply reducing the number of times per hour that the Crank-Nicolson scheme is run. This is shown below in Figure 7.

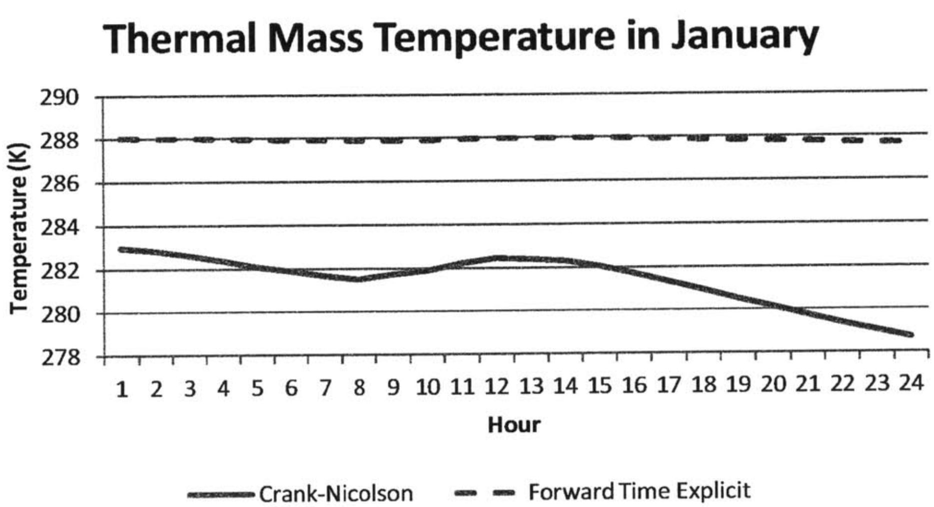

Figure 7: Comparison of thermal mass temperatures for Crank-Nicolson and forward time explicit methods

This increased error is likely due to the time step and thermal mass slice size. As discussed previously, the error in an explicit solver like this one is proportional to the time step used. When the number of slices is increased (for the "low" thermal mass setting used here, from 3 to 10), thus decreasing the slice size and the time step as per the Biot criterion, the 10% reduction in run time becomes a 25% to 50% increase in run time.

Lighting Calculations

One of the oldest pieces of code present in Design Advisor concerns the simulation of lighting in the user-defined typical room. Before the full building is formed (if applicable) or energy use is simulated, a lengthy set of calculations takes place during which the distribution of solar radiation within the room (and, concurrently, additional lighting requirements at the workplane) is predicted. This also allows for the daylighting visualizations that are a part of Design Advisor's output. The basic steps are as follows [35]:

1. Discretize the window, both side walls, the ceiling, and the floor of the shoebox room (all surfaces except for the back wall) into small rectilinear elements. Each element stores its own illumination level. The number of elements can vary greatly depending on the resolution setting chosen by the program and the dimensions of the room. For the default room, 5m wide by 5m deep by 3m high, each element could be 10cm on a side (for the crudest setting) or 0.1cm on a side (for the finest).
2. At each time step, use the calculable position of the sun relative to the window as well as the direct solar radiation data from the weather file to find the discretized surface regions that are illuminated. In all cases, the room is assumed to be empty, so there are no pieces of furniture or occupants blocking the light. If blinds are present, their effect on the initial distribution of light on the room (by absorbing incoming solar radiation or reflecting it into the room or back out into the environment) is accounted for here.
3. Add diffuse solar radiation throughout the room by treating the window as a uniform source of diffuse radiation.
4. Calculate the contribution of reflected light to the light level at each element. This is done by summing the radiosity of every other element in the room, accounting for the view factor between each pair:

$$R_i = (1 - \alpha) * \left(\sum_j^n VF_{ji} * R_j\right) \qquad (17)$$

   Where R is the radiosity of the indexed rectilinear element, α is its absorptivity, and VF is the view factor between element i and j. All of the elements are assumed to reflect light in a diffuse manner, so the angle of incoming radiation is not preserved. This step takes up a majority of the total calculation time.

5. Repeat step 4 until the change in radiosity of the room elements does not change with additional reflections. The final difference between the illumination level at the workplane and

the required light level set by the user governs the amount of lighting energy that is used during the current time step.

While this algorithm does succeed in rapidly predicting the daylighting levels in the room (this process is generally completed during the first several seconds of the simulation), it is nonetheless a significant source of potentially superfluous calculation time, especially if the goal is model optimization. Examining the code, one finds unnecessary repetition of complex calculations. There is also the question of the necessity of these calculations for the purposes of the average user.

As stated previously, this process of discretization and lighting prediction takes place before the building as a whole is created and the general energy usage simulation is performed. For a single room, this is not an issue. However, for four-sided or multi-floor simulations, this code runs multiple times. Obviously, the lighting requirements will change for different room orientations, but the map of rectilinear elements will remain the same (as the typical room is tiled around the building core). The calculations are also redone between types of floors (with and without a roof). Since the height of the building does not have a significant effect on the solar altitude, it is not included in the calculation of the lighting in the room. As a result, the lighting in a room will be the same as all those above and below it. This means that, in the most complex cases, the discretization could take place once and the lighting prediction could be done four times, once per facade direction. In those cases, as presently coded, the process described above takes place in full eight times.

On top of this, in the case in which we are most concerned with simulation speed (genetic algorithm-related runs during optimization), these calculations may not be necessary at all. The daylighting visualizations, while very informative for designers, would never be seen with the exception of the one corresponding to the optimal solution. Deactivating a portion of this code would also for additional variables or levels to be included in the initial parametric search and regression.

However, due to the structure of the Design Advisor code base, this kind of decoupling has proven to be very difficult. As discussed, the daylighting simulation leads to the lighting usage prediction during most simulations, which impacts heating and cooling energy use and is itself one of the primary outputs. Because of this, not only is the code rather complex (having been ported from MathCAD and spanning 48 files) but it is hooked into the structure of the software at a very deep level. This makes it difficult to modify but not fully replace the daylighting libraries without causing errors from dependencies elsewhere in the program.

It may instead be possible to predict required lighting energy alone with a simple correlation. The factors that affect lighting represent a somewhat reduced subset of Design Advisor's inputs:

- Solar information - solar altitude, solar-surface azimuth, incoming solar radiation (encapsulated with the selection of five cities – Boston, Miami, Anchorage, Santiago, and Kuala Lumpur - and the four cardinal directions from Design Advisor)
- Window - percent glazing area, window unit type, glass coating
- Blinds - blind angle, blind geometry, blind color
- Room - width, depth, height
- Lighting settings – required light levels, lighting control method

Recognizing that the various solar measures are themselves functions of the current time and location and choosing to disregard blinded window units (as in done below during the initial optimizer regressions) serves to shrink this list even further. This in turn allows for a parametric set of runs that can be used to relate required lighting energy to the selected variables:

Variable	Minimum	Maximum	Levels	
City	N/A	N/A	5	
Orientation	N/A	N/A	4	
% glazing	20	80	3	
Window type	Single glazed	Triple glazed	3	
Glass coating	Clear	High performance	2	
Lighting control	Single dimmer	Individual dimmers	2	
Lighting req.	300	600	2	
Room width	4	10	3	
Room depth	4	10	3	
Room height	3	5	2	
			25920	Total runs
			5	Seconds per run
			129600	Total seconds
			2160	Total minutes
			36	Total hours

Table 9: Table of variables for lighting regression

The yearly lighting energy requirement for the room was treated as the dependent variable for the regression analysis. The orientation variable was converted into a number, with 0° being north and the value increased clockwise around the cardinal directions. The city and orientation variables together were taken as signifiers of the physical configuration of the sun relative to the building façade and the incident illuminance (rather than radiation, as we are concerned with the light levels in the room). Total incident illuminance is made up of direct, diffuse, and reflected (from the ground) components, as follows:

$$I_{total} = I_{direct,incident} + I_{diffuse,vertical} + I_{reflected,incident\ from\ ground} \quad (18)$$

The typical meteorological year data files used by Design Advisor contain hourly information about direct normal radiation, defined as the "average amount of direct normal illuminance . . . received within a 5.7° field of view centered on the sun" and diffuse horizontal illuminance, defined as the "average amount of illuminance . . . received from the sky (excluding the solar disk) on a horizontal surface" [36]. These are converted into the three components mentioned above using several configuration-related factors:

$$I_{total} = I_{direct,normal} * \cos\theta + I_{diffuse,horizontal} * Y + I_{direct,normal} * (C + \sin\beta)\rho_g * \frac{(1-\cos\Sigma)}{2}$$

(19) [23]

Where θ is known as the surface-solar angle of incidence (the angle between the façade normal and the vector from the façade to the sun), C is a dimensionless factor taken to be its yearly average value of 0.118, ρ_g is the ground surface albedo (roughly 0.2), Σ is the inclination of the façade surface (assumed to be 90° here), β is the solar altitude (based on location and time of year), and Y is a dimensionless factor dependent on the surface-solar angle of incidence.

The surface-solar angle of incidence, given the stated assumption about the façade inclination, is found as follows:

$$\cos\theta = \cos\beta * \cos\gamma * \sin\Sigma + \sin\beta * \cos\Sigma = \cos\beta * \cos\gamma \quad (20)\ [23]$$

Where γ is the difference between the solar and surface azimuths (the angle between the sun and façade normal, respectively, and north).

Y, the conversion factor from diffuse illuminance falling on a horizontal surface and that falling on a vertical surface, is equal to:

$$Y = \begin{cases} 0.55 + 0.437\cos\theta + 0.313(\cos\theta)^2 & \cos\theta > -0.2 \\ 0.45 & \cos\theta \leq -0.2 \end{cases} \quad (21)\ [23]$$

These calculations can be found in full in Bryan Urban's original thesis about the MIT Design Advisor [23]. The resultant total incident illuminance levels are then averaged (only for daytime hours) for each of the four simulated cardinal directions for each city. To have a mix of different latitudes and longitudes in the simulation set, the five cities chosen were Boston, Anchorage, Santiago, Kuala Lumpur, and Miami. Note that the illuminance levels differ from east to west due to differences in solar radiation levels (as reported in the weather files used here) between the morning and afternoon hours.

Average Yearly Illuminance in Lux

	North	East	South	West
Boston	15626	36198	35529	29484
Anchorage	12559	25812	29960	22917
Santiago	27430	37014	18113	24840
Kuala Lumpur	21181	20429	22405	33098
Miami	19021	39273	34803	35205

Table 10: Average yearly illuminance levels during daylight hours for vertical surface facing in given direction

Each of the runs detailed in Table 9 were ordered by yearly required lighting energy. Performing a regression on that data set yielded the following results (note that any categorical variable that is not listed, such as double glazed window units, were selected as the default for the variable in question and have an effective coefficient of 0):

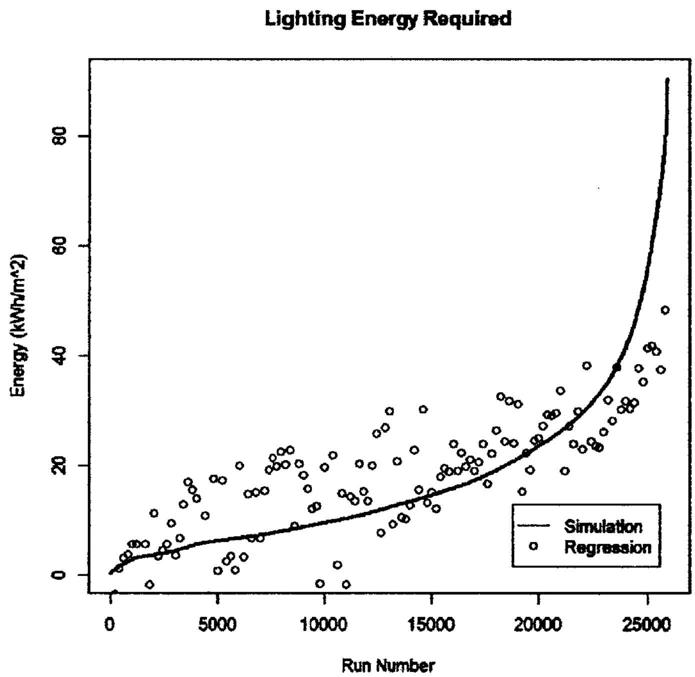

Figure 8: Calculated and predicted yearly lighting energy requirements

Coefficient	Value
Intercept	2.2
Avg. Total Ill.	0.0
Lighting control = "Single Dim"	8.9
Depth	1.9
Width	-0.8
Height	-2.8

Max Residual	41.72
Median Residual	-1.12
Min Residual	-17.643
RMSE	9.013453

Glazing %	-0.1	Adj. R^2	0.6296
Window type = "single glazed"	-1.2		
Window type = "triple glazed"	1.3		
Coating = "high performance"	2.7		
Lighting required	0.1		

Table 11: Summary of lighting regressions

While the regression yields significant residuals for scenarios corresponding to the extreme ends of the plot in Figure 8, the bulk of the predictions agree somewhat with the corresponding calculated lighting energy usage. This is especially true in the region from 20-40 kWh/m^2 per year, which covers many offices and other buildings with similar lighting needs. As these types of buildings account for a significant proportion of this country's total electricity usage, it is important that those simulations are the most accurate.

It was found during the regression analysis that when the lights are set to be on (and at full brightness) during occupied hours, even though the lighting-related energy consumption is not a function of the available daylighting, the room is still discretized as discussed before. Therefore, this is the one case where the daylighting module can be switched off entirely with no change in prediction accuracy when the daylighting visualizations are not needed. The run time savings from this change are shown below.

		Original Run Time (ms)	No Daylighting (ms)	% Difference
Single-sided	One floor	7728	6727	-12.95
Four-sided	One floor	32201	28964	-10.05
Single-sided	Three floors	12744	10812	-15.16
Four-sided	Three floors	53213	47752	-10.26

Table 12: Run time changes from removing daylighting module when lights are always on

This change serves mainly as a demonstration that there is likely more time to be saved by continuing to find and eliminate unnecessary or repeated calculations. Design Advisor's default settings assume that light fixtures respond to the lowest lighting level in the room (rather than always being on), so most simulations will not benefit from this fix.

Representative Day Formulation

Another barrier to reducing overall simulation runtime is the fact that, by default, the energy usage for a full year is calculated. Any reduction in the number of days that need to be simulated will logically reduce the total time required. As such, four methods for formulating and testing representative days were implemented. The representative days are formed from the same weather data used by Design Advisor but do not themselves necessarily correspond to any given actual day. The four methods are described below, followed by a summary of the results from each.

The first method is simply to use the 24 hourly averages of temperature and solar radiation to create a single day for each month. Each day is simulated twice, once as a training day and once to find the heating, cooling, and lighting energy usage for the month. The training day serves to "reset" the temperature of the air in the room and the temperature distribution within the thermal mass so that, for example, a day representing April is not simulated with an initial temperature distribution corresponding to the end of an average day in March. This reduces the number of days from 365 to 24.

The second method separates heating and cooling processes into two separate days per month. Heating and cooling degree hours are calculated for each hour each month (i.e.: total heating and cooling degree hours from 00:00 to 01:00 in January, etc.) relative to a standard 18.3°C setpoint. [37]–[39] From this, hourly temperatures are determined by averaging the hourly degree hour counts and finding the required temperature difference from 18.3°C for each hour. User-set loads, such as heat from equipment in the room, are split between the heating and cooling day for each month. This method produces 48 simulated days in total.

The third method attempts to capture the effects of extreme weather that are missed when using a pure average. It does this by assuming that hourly temperatures are normally distributed within each month and uses this to form three representative days. The three days represent average, low temperature (with each hourly temperature set one standard deviation below the monthly mean for that hour), and high temperature (one standard deviation above the mean) days. All three are simulated and the predicted energy usages are weighted according to the 68-95-99.7 rule for normal distributions. This rule of thumb states that, for a normally-distributed variable, roughly 68% of the values fall within one standard deviation of the mean, 95% fall within two, and 99.7% fall within three. Therefore, the energy usage predicted for the average temperature day is weighted to account for 68% of the energy usage for the month while the two extreme days account for 16% each. This is referred to in Figure 9 as "Normal Dist. Temp." The final method expands on the third by assuming that incident solar radiation follows a normal distribution as well. This version is referred to as "Normal Dist. All." For both the third and fourth methods, 72 days are simulated.

The four strategies (as well as a standard simulation) were used to simulate a standard office-type building with low thermal mass and a bitumen roof in Boston, Anchorage, Los Angeles, and

Mumbai. This was done to test the strategies across a set of disparate climates. The results of this study are presented below.

First, it is important to note that while limiting the number of days that need to be simulated does drastically reduce the overall simulation time (by 47% to 56% for the building tested here), a significant proportion of the calculations are used to set up and initialize the building. As a result, the difference in calculation time between the four representative day methods is minimal. This led in part to the effort to change the lighting model as discussed previously.

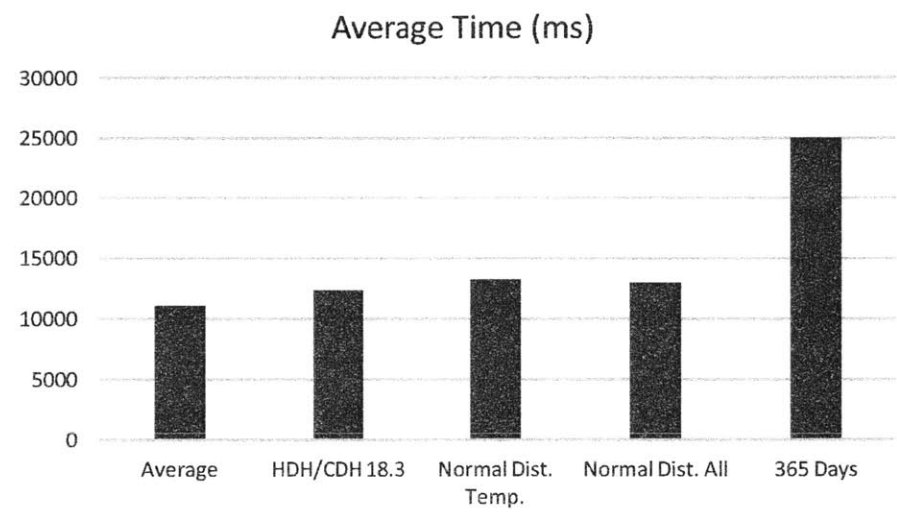

Figure 9: Average simulation time for full-year simulation and representative day methods

On a monthly basis, the average temperature and degree hour methods tend to produce the largest monthly errors. This is especially true for climates dominated by one temperature extreme, such as in the Anchorage and Mumbai examples shown below. In the case of the average temperature formulation, this is likely due to the fact that averaging hides perturbations that could be important to energy usage. In Anchorage, this leads to an underestimation of heating energy during the summer (as unseasonably cold days have been smoothed out). In Mumbai, this leads to an underestimation of cooling energy during the winter (as the unseasonably hot days are lost). It is important to know, however, that those types of days in those extreme climates represent a small percentage of yearly heating and cooling demands. This effect is seen in other fields as well; for example, crop growth simulations programs that make use of weather averaging tend to overestimate yields as they miss damaging day-to-day variation in conditions. [40] In the case of the degree hour method, it is possible that the 18.3°C setpoint is incorrect or insufficiently descriptive for extreme climates. That temperature is meant to represent a point of thermal equilibrium between a building and its environment. As this equilibrium depends on weather conditions (temperature, wind, solar gains, etc.), building construction norms, and internal loads, the accuracy of this setpoint will vary from place to place. It could be the case that the unique conditions and building insulation needs in extreme environments renders the 18.3°C setpoint too inaccurate to be used in this kind of simulation. Therefore, different equilibrium points may need to be found for these locations. Another possibility would be to use two different setpoints for heating and cooling. For example, the Pacific Energy Center uses a 65°F baseline for heating degree days and an 80°F baseline for cooling degree days. [41] This, along with finding a more accurate equilibrium point, could improve the performance of the degree hour method described here.

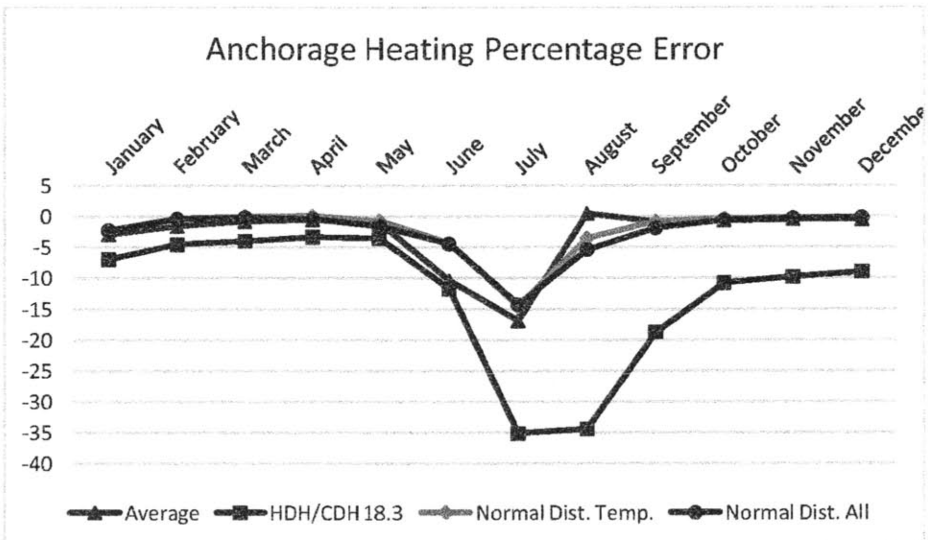

Figure 10: Monthly heating energy error for Anchorage for representative day methods

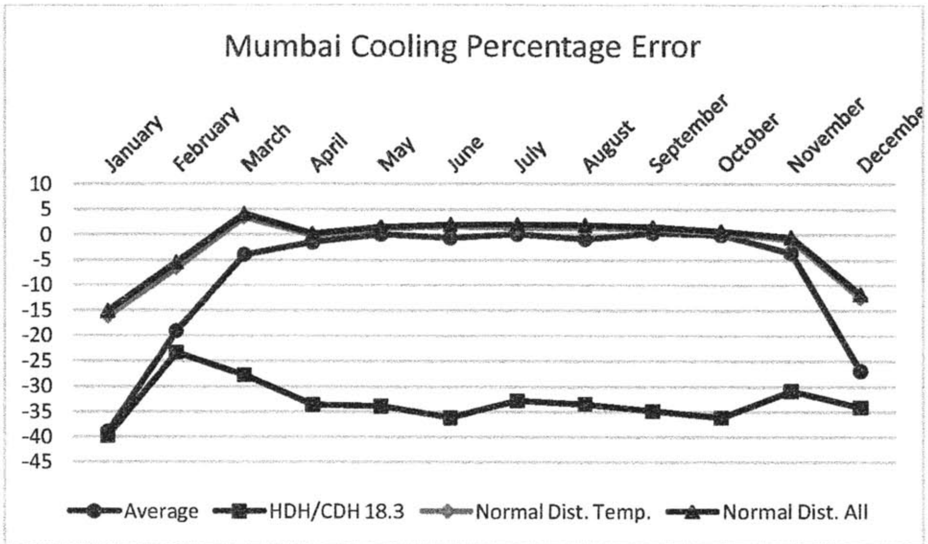

Figure 11: Monthly cooling energy error for Mumbai for representative day methods

Note: the two normal-distribution-based methods produced nearly identical results, which is why they are difficult to distinguish from one another in Figure 10 - Figure 13.

That said, the two normal distribution methods, while more accurate on average, still produce statistically significant errors. In temperate regions, such as the example in Boston shown below, both methods have their largest errors during the shoulder months of late spring and early fall. The average temperature in Boston during June and September is closer to the aforementioned 18.3°C equilibrium point (19.8°C and 18.1°C, respectively) than during any other month. As heating and cooling requirements during these months are relatively balanced and minimal (compared to other months of the year), small absolute errors can lead to large percentage differences. These errors (compared to the default full year simulation) are shown below for Boston.

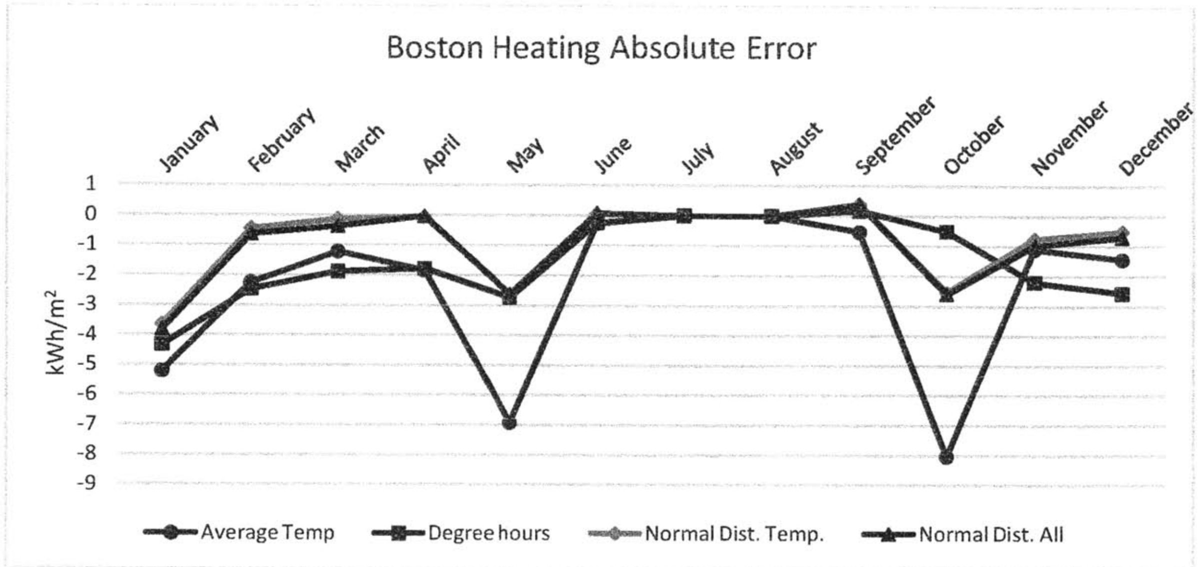

Figure 12: Monthly heating energy error for Boston for representative day methods

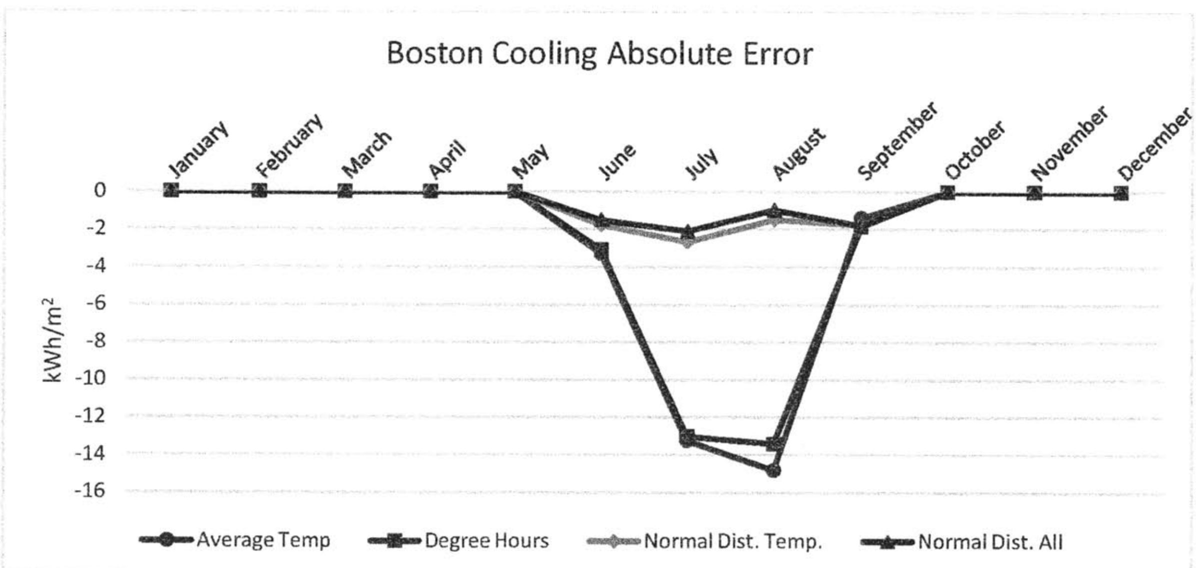

Figure 13: Monthly cooling energy error for Boston for representative day methods

The average percentage errors for whole-year energy usage for the four cities tested are presented below. Together, these measures serve to highlight the fact that, while both normal distribution methods lead to significantly more accurate results (compared to a standard Design Advisor simulation) than either the average temperature of degree hour method, the difference between the two is minor – a 2.7% difference at most across all four cities. Depending on how many simulations need to be run (such as in the case of an optimization algorithm), this slight increase in accuracy may not be worth the extra time it takes to prepare the solar radiation data (about 0.2 seconds on average). It is also interesting to note that no method overestimated whole-year heating or cooling energy usage for any of the four cities. This points back to the importance of extreme weather days and the need to capture their effects in these representative day formulations.

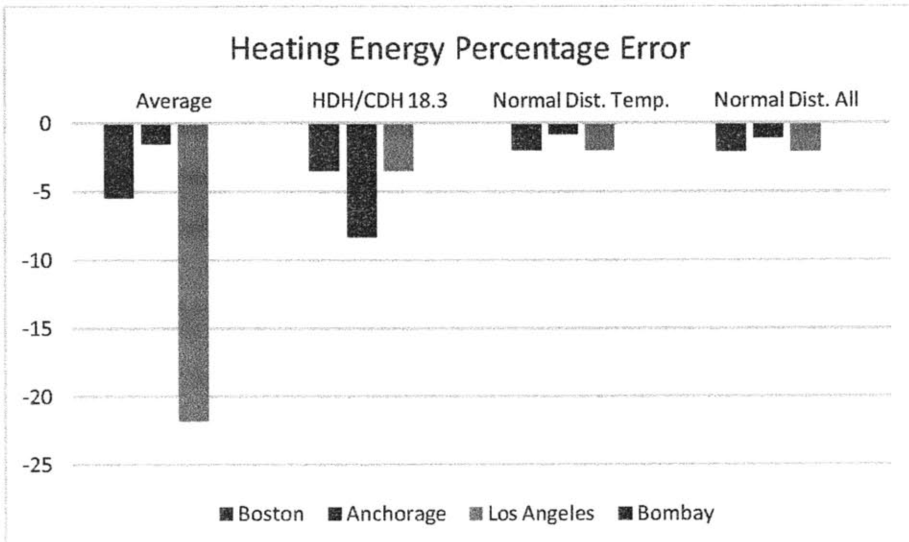

Figure 14: Total yearly heating energy error for representative day methods

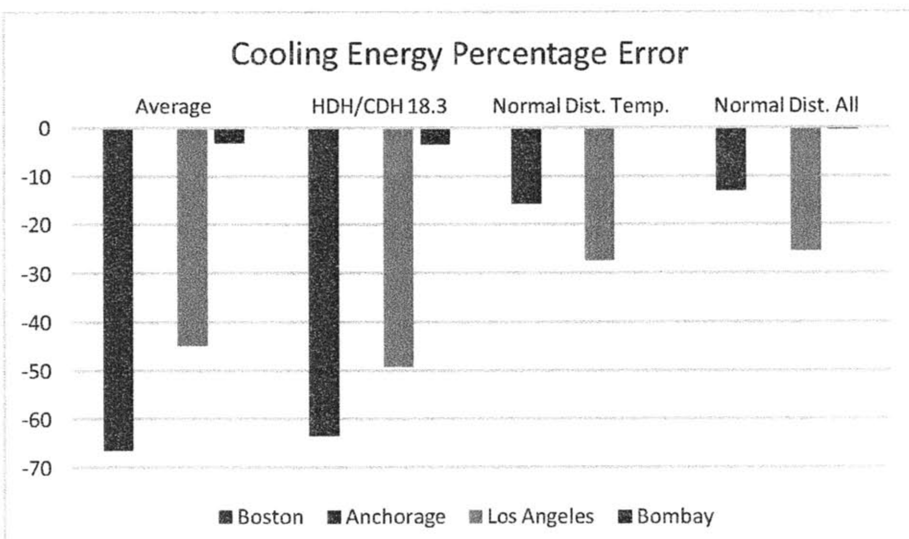

Figure 15: Total yearly cooling energy error for representative day methods

One way to accomplish this would be to move away from the assumption that hourly and daily temperatures are perfectly normally distributed within each month. A cursory examination of the weather file for any of the cities tested here shows that to be false – some monthly temperature distributions are highly skewed, some are bimodal, and many exhibit other irregularities. Incorporating these non-ideal factors without sacrificing the generalizability of these methods would require changing the weighting used (from 16% for each of these extreme days and 68% for the average day) in accordance with some additional mathematical measure such as skewness or kurtosis.

One possibility would be to assume that temperature follows a skew-normal rather than a normal distribution. First introduced by O'Hagan and Leonard in 1976, the skew-normal distribution, as the name implies, attempts to incorporate the skewness of a distribution into its probability density function. [42] A method of estimating the probability density function, from Azzalini and Valle, is given below. [43], [44]

$$f(x) = 2\phi(x)\Phi(\alpha x) \tag{22}$$

Where f is the probability density function of the skew-normal variable x and α is a parameter describing its skewness. ϕ and Φ are the probability density function and cumulative distribution function of the normal distribution, respectively, and are given below.

$$\phi(x) = \frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}} \tag{23}$$

$$\Phi(\alpha x) = \int_{-\infty}^{\alpha x} \phi(t)dt = \frac{1}{2}\left[1 + \text{erf}\left(\frac{\alpha x}{\sqrt{2}}\right)\right] \tag{24}$$

However, given the previously-discussed tradeoff between speed and accuracy (in terms of using a normal distribution for only temperature), it is unlikely that the addition of these skewness calculations would result in significant enough improvements in net yearly accuracy to be worth the extra calculation time. This is especially true if the representative days are used primarily for faster simulations during optimization rather than for individual building runs.

Design Advisor Validation

As part of the Department of Energy's Energy Efficient Buildings Hub Project (referred to as "the Hub" from this point onwards), model testing and validation was performed on three existing buildings. Two, selected by project administrators for use with all simulation engines involved, are located in the Philadelphia area. The third is located in Cambridge, Massachusetts on the campus of the Massachusetts Institute of Technology (MIT). The testing of these demonstration buildings represented the first such validation of Design Advisor using already-existing buildings. Previous validations had been performed by comparing Design Advisor's energy usage predictions to those from Energy Plus, but real buildings inevitably include inefficiencies and complexities arising from user behavior, equipment malfunctions, and other factors that are not necessarily included in simulations. The process of and results from these three buildings are summarized below.

One Montgomery Plaza

The first demonstration building to be selected was One Montgomery Plaza, a municipal office building in Norristown, Pennsylvania. It has 10 floors, approximately 250,000 square feet of gross floor area, and serves roughly 3,000 employees. It was of particular interest due to the building's ongoing need for extensive retrofits (and the controversy surrounding the final expense to Montgomery County).

A variety of documents describing the building were made available through the Hub and used to form estimates of the simulation inputs. For example, hourly utility data was used to determine the occupancy schedule of the building while a site audit documenting the repair needs mentioned above was used to find building envelope parameters [46]. Furthermore, several parameters (such as air change rates and equipment power usage) were not available in any explicit form and had to be estimated from related measures or industry guidelines. As a result, a majority of the time needed to simulate the building was spent distilling these various data sources down into the inputs required by Design Advisor. Due to the nature and format of the reported data, this would likely be true for any other simulation suite.

Using these initial inputs, a baseline run was performed. The goal was to first validate Design Advisor's stock capabilities – that is, to check the accuracy of the outputs without any additional post-simulation calculations. The results from this first simulation are shown below in Figure 16.

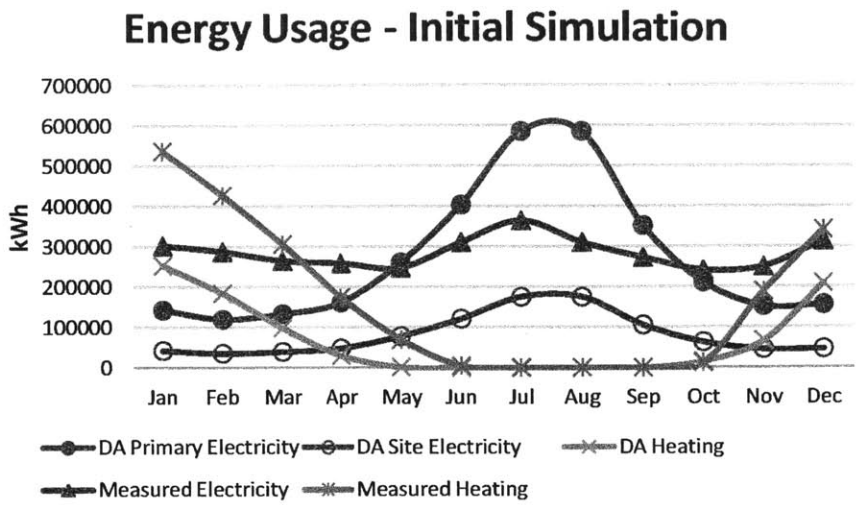

Figure 16: Initial simulation results for One Montgomery Plaza

This comparison proved to be very informative in terms of shortcomings in Design Advisor and changes that could be made for the sake of usability as well as simulation accuracy. The first is that Design Advisor's electricity (lighting and cooling) usage predictions reflect primary or source energy usage, which means that HVAC system efficiency and transmission losses are taken into account. While this is useful from a sustainability perspective, as it provides a more inclusive view of a building's impact, it can be less useful to a building manager looking at site energy usage measurements on a utility bill. It also makes the simulation seem far less accurate than it is, as seen through the three electricity curves in Figure 16.

The second is that air handler and miscellaneous equipment electricity usage is not accounted for in any of the final report categories. Equipment energy usage (in watts per unit area) is a required input, but it is treated only as a constant internal heat load. Similarly, hourly air change rate is an input, but Design Advisor does not calculate the energy required to actually facilitate that movement of air (which, in the case of One Montgomery Plaza, turned out to be fairly constant throughout the year). The energy used by the HVAC equipment was included using a simple rule of thumb: 1 CFM of airflow takes 1 W of power to move. While the output categories are clearly denoted as heating, cooling, and lighting, the lack of these other major usage categories can be confusing to users. It also paints a too-optimistic picture of the energy efficiency of the simulated building: in the case of One Montgomery Plaza, total monthly electricity consumption was underestimated by up to 250 MWh before these loads were included.

Lastly, a reexamination of the available data on One Montgomery Plaza led to a third change, the inclusion of multiple occupancy schedules. During the heating season, in order to retain occupant comfort levels, the boiler and air handlers run on what is essentially a 24-hour schedule [47]. As Design Advisor relaxes internal temperature and air change requirements during unoccupied hours, this meant that a significant use of both heating and electrical energy was being missed. To correct this, the simulation was split into two: one full-building simulation using the original occupancy schedule for the cooling season and one with a 24-hour schedule for the heating season. Alongside this, the assumed boiler efficiency was changed from the ideal case (100%) to a more realistic value (90%). The energy usage predictions were extracted for the specified months and combined to yield the final results show below in Figure 17 and Figure 18.

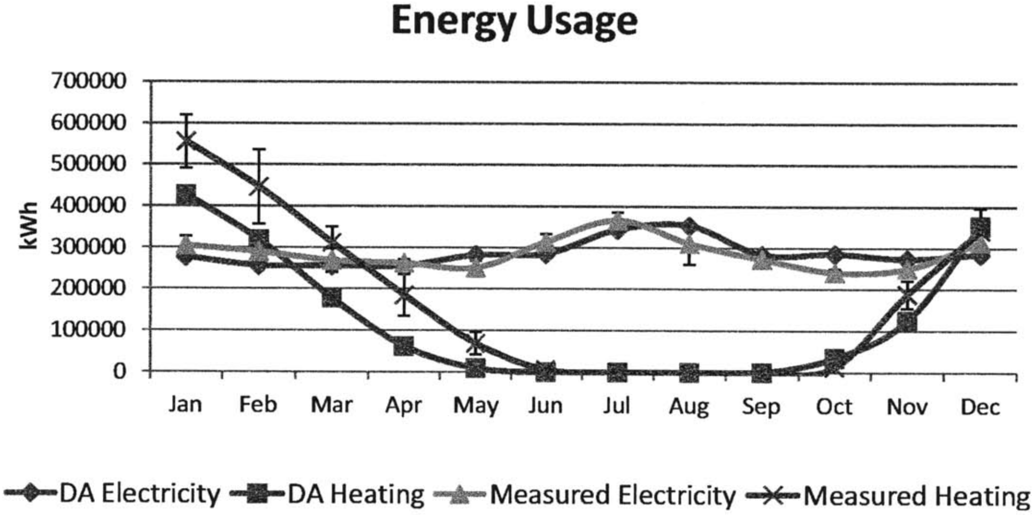

Figure 17: Heating and electricity usage for One Montgomery Plaza

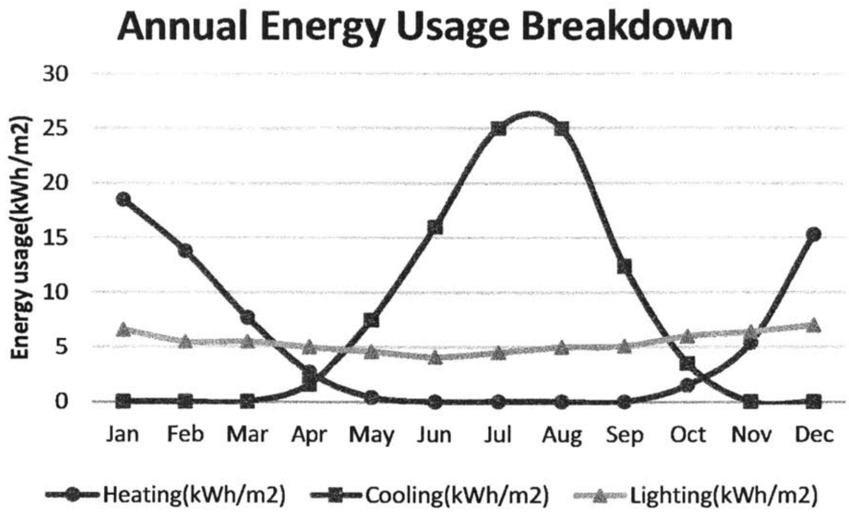

Figure 18: Breakdown of Design Advisor predictions for heating, cooling, and lighting energy usage per unit area for One Montgomery Plaza

With the aforementioned changes and additions in place, the energy usage predictions from Design Advisor agreed fairly closely with the actual utility data for the building, especially for electricity usage. Heating energy usage was still under-predicted for much of the winter and spring, most likely due to the nature of the 24-hour heating cycle. This led to an overall mean yearly bias error (for the combination of the two energy types) of -11.2% and a root-mean-square error (RMSE) of 27.1%. The absolute monthly error ranged from 2.9% in December to 29.1% in April. Additional operational details that would reduce these errors would also be required for an accurate simulation using a more comprehensive program such as EnergyPlus.

Philadelphia Navy Yard Building 101

The second case study focused on Building 101 at the Philadelphia Navy Yard, which acted as the central office for the Hub project. It is shown in map view below in Figure 19. It is a 55,000 square foot office building with three aboveground floors, an attic, and a basement. [48] It is billed as "one of the nation's most highly instrumented commercial buildings," [49] leading to a plethora of detailed data about its operation.

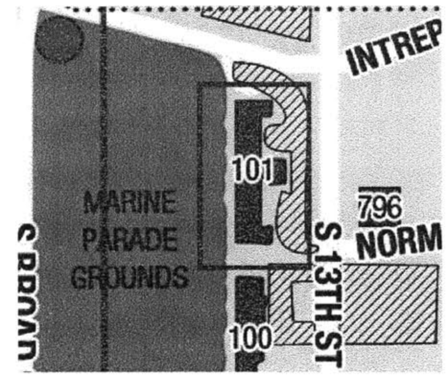

Figure 19: Map view of Building 101 [50]

The general process behind the simulation of Building 101 was the same as that for One Montgomery Plaza. First, input values were extracted or estimated from the documents made available through the Hub: for example, floor dimensions and air change rates were determined through an envelope air tightness study performed by Camroden Associates [51] and the ventilation mode (joint mechanical and natural ventilation) was chosen based on site photos that showed user-operable windows. As Design Advisor assumes a rectangular floor plan, the overall building dimensions were set such that the length and total floor area matched reality. Finally, the building was split into three separate simulated regions – basement, main floors, and attic – to allow for differences in glazing percentage, room height, and insulation.

Continuing to follow the process from the previous demonstration building, a second set of simulations was performed making use of additional knowledge about Building 101. A before, the furnace efficiency was lowered, plug loads and air handler energy usage added (using the same rule of thumb described earlier), and the occupancy schedule was modified to more accurately reflect the heating and cooling of the actual building. The most significant change involved corrected the floor plan and region definitions for the simulated building. As shown in Figure 19, Building 101 is roughly T-shaped, with a small foyer on its eastern side leading to the main floor area (running north-south). These two sections were treated and modeled separately. The airflow between them was assumed to be minimal based on the floor plans of the building. Further examination of the available information concerning the attic (specifically, that it is "used only for mechanical and storage space" [52] and appears to lack windows) led to the decision to approximate it as a highly-insulating roof. The basement was left as its own region. This led to a total of four simulated sections (basement and aboveground floors for both the lobby and main wing). The results from this final set of simulations are summarized in Figure 20 and Figure 21. The error bars on the measured data denote the standard deviation in each month's measurement.

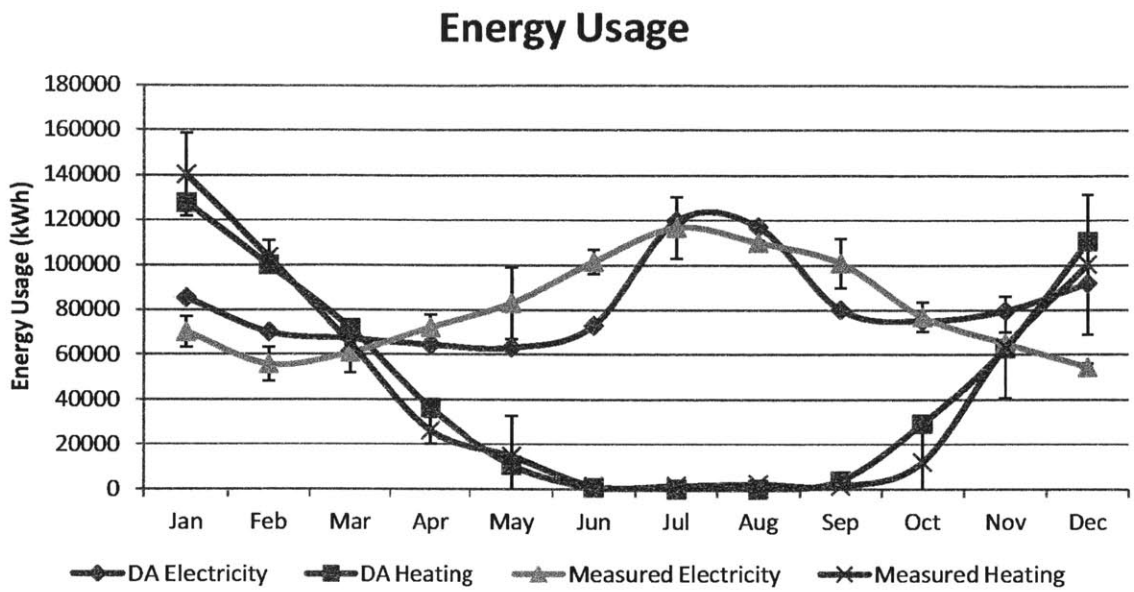

Figure 20: Heating and electricity usage for Building 101

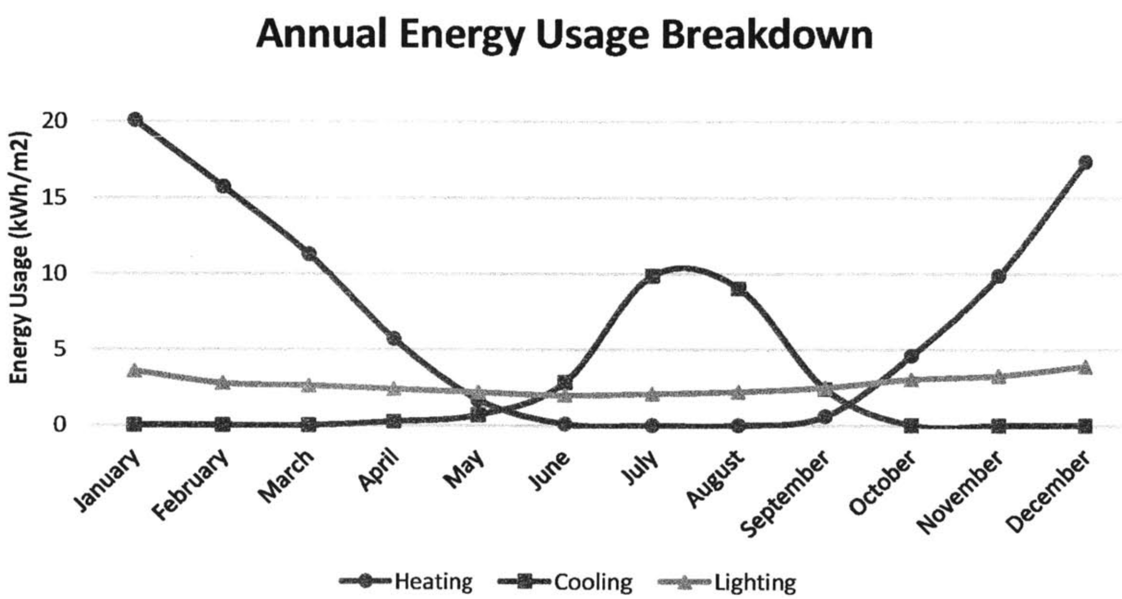

Figure 21: Breakdown of Design Advisor predictions for heating, cooling, and lighting energy usage per unit area for Building 101

Again, the predictions from Design Advisor closely match the actual energy usage of the building for much of the year. The monthly heating predictions are all within the spread of the data between the different years for which data was provided (as indicated by the bars on the measured data points). Electricity usage is overestimated in December and underestimated in the shoulder months between the heating and cooling seasons. It is difficult to determine the source of this error, as the occupancy of the building was generally very low and changed drastically over the period of time for which data was made available. This led to an overall yearly mean bias error of 2.6% and an RMSE of 31.4%. The absolute monthly error ranged from 1.2% in July to 46.3% in December. The yearly RMSE for both this building and One Montgomery Plaza compares favorably to those from the other simulation engines in the Hub Project (some of which are based on EnergyPlus), which ranged from 10% to 20%.

MIT Building E40

The third validation building was Building E40, also known as the Muckley Building, located on MIT's campus. The building was constructed in 1930 [53] and underwent major renovations in 1980 and 2000 [54]. Its usable floor area (across four above-ground floors and a basement) is primarily office space [55] (for graduate students, meeting rooms, etc.). Half of the basement is devoted to a chilled water plant that is considered separate from the building, both for modeling and utility billing purposes. It was later discovered that a portion of the remaining half of the basement houses a bank of servers, which could account for some of the error in electricity usage seen in Figure 23. An overhead map view of the building is shown in Figure 22.

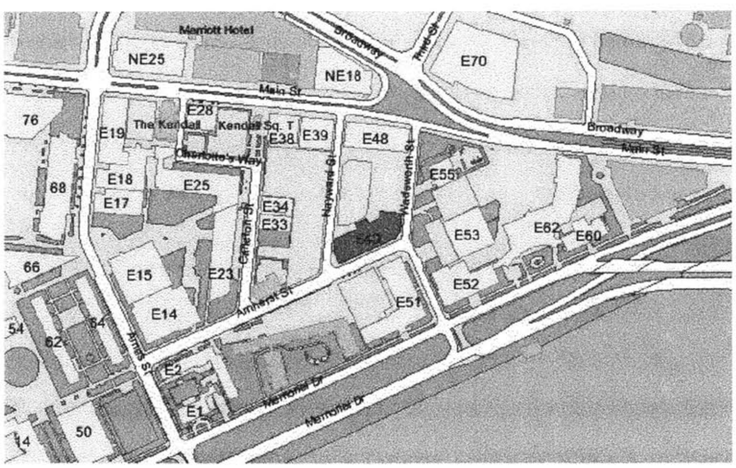

Figure 22: Map view of MIT Building E40 [56]

Again, simulation inputs were determined primarily through documents concerning the building (in this case, made available by MIT's Department of Facilities). Internal temperature and building insulation specifications were confirmed via an in-person IR survey of the site. This was done by taking corresponding measurements of air and surface temperature on both sides of various external walls of the building. Assumptions were made about convection and radiation heat transfer coefficients based on the conditions at the time. The calculations are as follows:

$T_{internal\ air} = 71°F = 21.7°C$

$T_{internal\ wall} = 66°F = 18.9°C$

$T_{outside\ air} = 37°F = 2.8°C$

$T_{external\ wall} = 41°F = 5°C$

$h_{internal} \approx 5\frac{W}{m^2K}$

$$R = \frac{\Delta T}{\dot{q}} \tag{25}$$

$$\dot{q} = h_{internal}\Delta T_{in} = (h_{internal} + \mathrm{h_R}) * (T_{internal\ air} - T_{internal\ wall}) = 10\frac{W}{m^2K} * (2.8\ K) = 28\frac{W}{m^2} \tag{26}$$

$$R = \frac{\Delta T_{wall}}{\dot{q}} = \frac{13.9\ K}{28\frac{W}{m^2}} = 0.5\frac{m^2K}{W} = 2.85\frac{hr\ ft^2\ °F}{Btu} \tag{27}$$

Using the data described above, initial energy usage simulations were performed with the weather file for Boston from the Design Advisor website. Steam and chilled water consumption were converted to kWh-equivalent to ease comparison to the cooling and heating predictions from Design Advisor (1 ton-hour of chilled water = 3.517 kWh and 1 pound of steam = 0.3523 kWh). Simulation results are summarized in Figure 23, Figure 24, and Figure 25 and are compared to average monthly utility usage from June 2010 to February 2014.

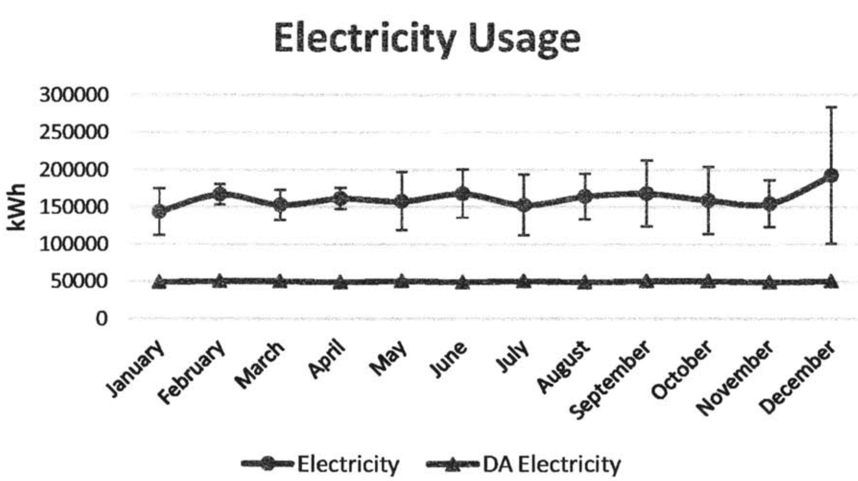

Figure 23: Electricity usage for MIT Building E40

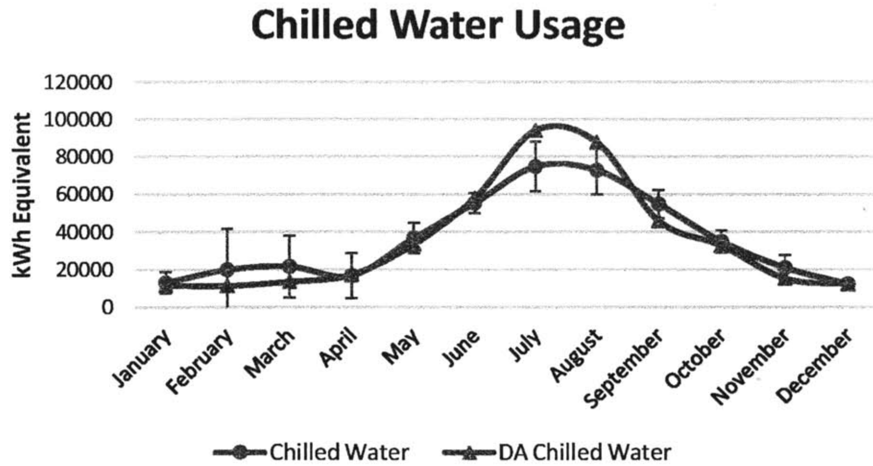

Figure 24: Chilled water usage for MIT Building E40

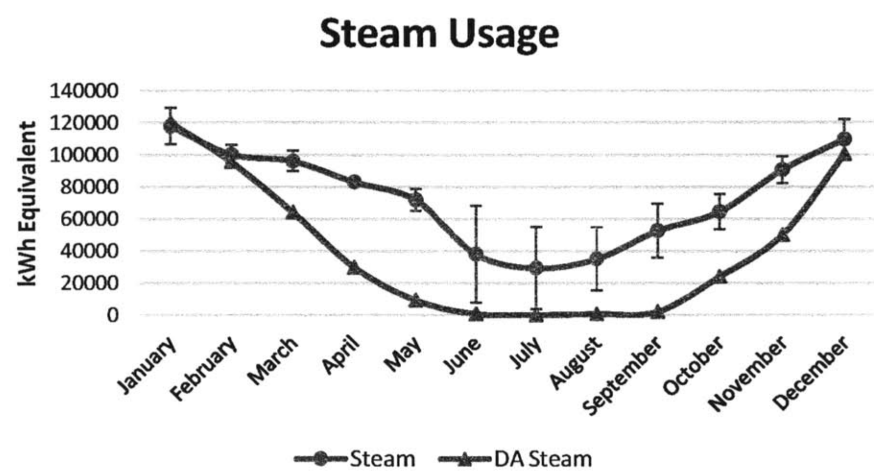

Figure 25: Steam usage for MIT Building E40

Across all three utility categories, Design Advisor's predictions were less accurate than for the previous two buildings (especially for electricity and steam usage). However, it has since been found that several important factors were either entirely unknown or otherwise unaccounted for and could explain these errors. As noted previously, a bank of servers (of unknown specification) is accounted for in Building E40's electricity usage. Given the power and cooling requirements of such devices, it is plausible that this could lead to the gap between the predicted and actual electricity consumption seen above. Without additional information, however, it is impossible to determine and will require further work.

Additionally, there is significant year-to-year variation in utility usage whose cause, as of yet, has not been found. For example, chilled water usage in March ranged from 5,050 kWh-equivalent in 2012 to 45,076 kWh-equivalent in 2011 (with a standard deviation, shown via the error bars on Figure 24, of 21,892 kWh-equivalent). Similarly, steam usage in June ranged from 8,808 kWh-equivalent in 2011 to 69,057 kWh-equivalent in 2013 (with a standard deviation of 30,233 kWh-equivalent). This implies the presence of simultaneous heating and cooling in the building. In the case of steam, the monthly averages appear to indicate that a leak developed some time in 2012 (the summer when steam usage first began to rise). For the most part, however, these abnormal usage patterns remain unexplained.

Resulting Improvements

As mentioned previously, much of the value in performing these validation simulations lay not just in quantifying Design Advisor's accuracy but also in discovering its weaknesses. As such, many of the changes discussed here were subsequently or will be added to Design Advisor to improve its usability and reliability for end users. These are briefly summarized below:

- Option to display site instead of source energy usage, allowing for comparison with actual utility bills without additional hand calculations
- Change "Lighting" category to "Electricity" to include lighting, plug loads, and an estimation of air handler and fan energy usage based on air change rate requirements
- Include more realistic estimates of furnace efficiency and chiller coefficient of performance

Other steps that were taken during the validation process have not been added to the program in order to prevent user interface clutter and limit the number of new inputs. These include but are not limited to: non-rectangular building layouts, multiple occupancy schedules per year, and allowing for different glazing percentages across floors. While these options were needed for these simulations, it would be difficult to implement them fully without overwhelming the target audience with or artificially limiting the number of possible variable combinations.

Optimization

Once a user models their proposed building design, it is not always clear what the takeaway should be. Knowing the energy use of a building does not necessarily imply knowledge of what should be changed to improve its performance. One possible solution to this is an optimizer, which would remove the guesswork by automatically finding the most energy-efficient solution based on the original design and the given constraints. As discussed previously, the challenge is not just to develop and implement any optimization algorithm (something as simple as an exhaustive parametric search would find an optimal solution) but to strike a balance between speed and accuracy to provide the maximum value to the user.

The Design Advisor website has, for some time, had a placeholder entry for an optimizer. Despite being nonfunctional, it served as the basis for the optimizer detailed here. The original user interface is shown in Figure 26.

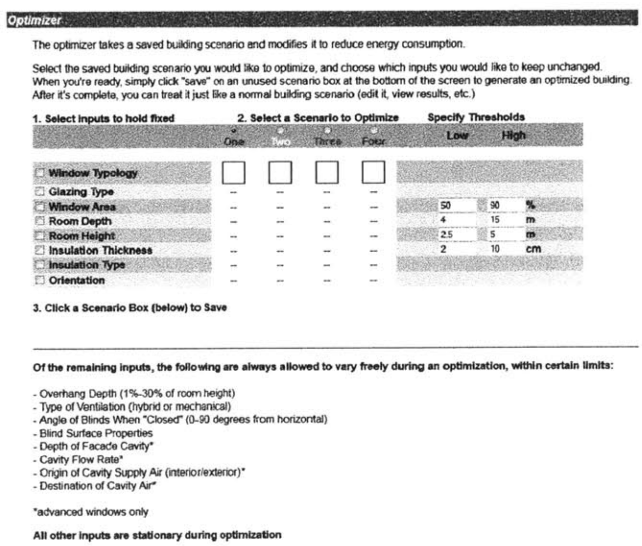

Figure 26: Original Design Advisor optimizer interface

This mockup has two key features that informed and guided the optimizer development process. The first is the integration of the optimizer into the overall scenario saving system in Design Advisor. As noted, the optimized building can be treated as a normal user-created scenario, allowing for direct comparison to the original design as well as further edits. The second is the user-selectable constraint system. Allowing the user to set not only variable limits but whether they vary at all ensures that the final design conforms to budgetary, geometric, or other constraints. Limiting the number of inputs that can change also serves to reduce the problem domain and the number of simulations that need to be performed.

A hybrid optimization technique was explored as part of this work. The optimization itself is handled using a genetic algorithm, which uses successive generations of sets of "chromosomes" (representing encoded sets of inputs) to work towards an optimal solution. To simplify the search and initially move toward an optimal solution, a multistep multivariate regression precedes the genetic algorithm, as used (with Design Advisor) by Carrie Brown, a previous graduate student at MIT [57]. It involves multiple levels of regression that allows increased accuracy compared to a single regression calculation. Both methods are detailed in the sections below.

Genetic Algorithm: Background and Usage

A genetic algorithm is a type of search algorithm that is often applied to optimization problems. Its name comes from the fact that it mimics natural selection in that the fittest solutions in each generation are the most likely to survive and pass on their traits. Numerous researchers worked independently on the application of evolutionary strategies to optimization problems throughout the 1950s and 1960s. The framework behind the modern genetic algorithm, however, was introduced by John Holland in his 1975 book *Adaptation in Natural and Artificial Systems* [58]. Holland's method defines and makes use of many of the features discussed below: binary-encoded chromosomes and genes, chromosomal crossover, mutation, etc. Previous implementations focused on individual evolutionary strategies (e.g., Fogel, Owens, and Walsh selecting the fittest solution from randomly-mutated options), lacked the scale of Holland's version (e.g., Rechenberg's use of a single pair of "parent" chromosomes rather than a larger population), or simply did not get the widespread academic attention and use that Holland's did [58].

Keeping with the evolution analogy, the solutions (or, in the case of Design Advisor, the inputs that will determine the solution) are stored in binary-encoded "chromosomes." The chromosomes can be subdivided into "genes" depending on the number of inputs that may be required. The number of bits comprising each gene (along with the range of possible values the trait or input encoded in the gene can take) determines the resolution of the search. For example, a 12-bit gene can take 4096 different values, meaning that a search on an input that ranges from 0 to 1 would have a resolution of 2.44E-4. A higher-resolution search can lead to a more precise solution but is likely to take longer, as a stricter convergence criterion will be required.

The general steps in the genetic algorithm used here are as follows:

1. Form initial chromosome set
2. If adding mutations, flip bits at random
3. Evaluate fitness of each chromosome (in this case, by simulating the building described by it)
4. Determine selection probability of each chromosome by dividing individual fitness by sum of fitness values
5. Select new set of chromosomes randomly according to calculated probabilities
6. Pair up chromosomes at random
7. Select crossover points at random and perform crossover on each pair
8. Repeat from step 2 until convergence criteria are reached

There are several ways to form the initial chromosome set (step 1), each of which addresses different concerns [59]. The chromosomes can be generated pseudo-randomly within the search space, which is a common method and adequately satisfies concerns about both the diversity of the population and speed and ease of implementation of the generator. There are also quasi-random generation schemes that space the chromosomes out on a grid or checks for some specified minimum distance between the pseudo-randomly generated points [60]. This enforces population diversity more strictly than pure pseudo-random generation. Lastly, the initial population could be chosen to be optimal in some sense, which ignores diversity but gives some direction to the algorithm. Beyond the method itself, the size of the population is also a concern [61]. A small population may lack the ability to contain enough diversity to avoid local maxima or minima. A large population would be preferable, but computation time will increase both due to the number of simulations required per generation and because the solution set will likely take longer to converge.

Step 2 is a simple programming exercises and requires no further explanation. For step 3, the fitness function being evaluated depends on the scenario in question. For Design Advisor, the fitness function is related to the difference between the energy usage of the building encoded in the chromosome and that of the user's original design. Depending on the goals of the designer, another appropriate function might be related to the total fuel cost to power the building (based on electricity and heating oil requirements). In the case of the weights for the retrocommissioning algorithm (discussed in a later chapter), the fitness function is related to the total error in predicted savings using the weights from the chromosome.

The reproduction (steps 4-7) is meant to mimic reproduction in actual creatures. For the purposes of reproduction, chromosomes are paired off into "parent" sets and have their genetic information mixed to create "offspring" chromosomes (the next generation). This can be, broadly speaking, split into two processes: parent selection and the crossover itself. There are four main parent selection strategies that are commonly used: proportionate, linear rank, tournament, and truncation selection [62], [63]. Proportionate selection is the method used here (implemented in steps 4-6). In it, chromosomes are assigned a selection probability in direct proportion to their fitness:

$$P_i = \frac{fitness_i}{\sum_i fitness} \tag{28}$$

These probabilities are then used to select chromosomes to be paired. In the algorithm described here, this is implemented by first filling a placeholder generation with copies of the original chromosomes chosen randomly as per the calculated probabilities. The placeholder copies are then paired up at random. This method is easy to implement, but has several shortcomings. It can only be applied in cases where the fitness function increases with the desirability of the solution. For example, in the case of Design Advisor, the fitness function has been written such that it increases with decreasing energy usage, which is the desired trait. Proportionate selection is also susceptible to single individuals dominating the reproduction process. This can lead to limited genetic diversity in subsequent generations, which will hinder the algorithm's ability to find the true optimum solution.

When using linear rank selection, the probability of a given chromosome being chosen for reproduction depends not on its actual fitness but on its fitness rank relative to the rest of the generation. This rank is used to assign a new fitness value as follows [62]:

$$f_{new,i} = \frac{(N-r_i)*(f_{new,max}-f_{new,min})}{N-1} + f_{new,min} \tag{29}$$

Where N is the number of chromosomes in the generation, r_i is the rank of chromosome i, and $f_{new,max}$ and $f_{new,min}$ are the predetermined maximum and minimum fitness values to be assigned. These new fitness values are then used to select parenting pairs using the same roulette wheel/probability-based method as in proportionate selection. Linear rank selection is less susceptible to dominant individuals than proportionate selection but, conversely, it can magnify the influence of weaker individuals, especially when using a small population.

Tournament selection does not rely on fitness-based probability at all. In it, a subset of chromosomes of size k is taken from the generation of size N. From this subset, the most-fit individual is chosen to become a parent. A common choice for k is 2 (a process known as binary tournament selection) [63]. In tournament selection, the only way for the least-fit chromosome to reproduce is if there are multiple individuals tied for the lowest fitness and they are the only ones chosen for the tournament. Otherwise, the presence of individual chromosomes in the parenting pairs will be inconsistent: depending on the subsets that are chosen, any given chromosome could be highly under- or over-represented.

Truncation selection is the simplest of the four methods described here. Some predetermined top-performing fraction of the chromosomes are selected to reproduce (0.5 and 0.3 are common choices [63]). Parenting pairs are then formed and each pair reproduces a number of times equal to the inverse of the truncation factor (twice each for 0.5, etc.). Truncation selection is not very common in practice as the truncation pressure it provides lacks any nuance or detailed selectivity.

Once the pairs of parents are chosen, there are also multiple strategies that can be used to perform the genetic crossover operation. Four common methods are discussed here: one point, two point, uniform, and arithmetic crossover [64]. In one point crossover, a single index point is chosen for each pair of binary-encoded chromosomes. Two offspring are created by exchanging the values from the parents from that point to the end of the chromosome. This is the crossover method that is used in the algorithm described here. Two point crossover introduces a second such swap point. In uniform crossover, each binary digit for both of the offspring is chosen from the corresponding index point in the parents according to some predetermined probability (often 50% for each parent). The two offspring will not share a parental source at any given index, so in practice the selection probability determines how intact the two parents remain in the next generation. Lastly, in arithmetic crossover, each offspring gene is formed using a weighted average of the corresponding parental genes, as follows [64]:

$$Offspring_{i,k} = a * Parent_{j=i,k} + (1-a) * Parent_{j\neq i,k} \qquad (30)$$

Where i is the number of the offspring (1 or 2), j is the number of the parent (1 or 2), and k is the index of the gene.

The appropriate convergence criteria are dependent upon many factors, including but not limited to: the complexity of the problem, the number of chromosomes per generation, the resolution of each gene, the presence of mutations, and the amount of time available for the search. One could choose to end the search based on the most or least-fit chromosome in the generation (to ensure the quality of one solution or the entire population), the average or total fitness of the population, the difference between subsequent generations, or the amount of time elapsed or number of generations simulated [65], [66]. A combination of these criteria could be employed as well, depending on the scenario.

Multistep Multivariate Regression

Multistep multivariate regression is a technique intended to take advantage of the computational speed and ease of implementation of a linear regression without the loss of accuracy that often comes with deriving said regression from a widely-varying domain. By repeatedly subdividing the data set and performing new regression analyses, we can drill down to a more accurate prediction without incurring a significant penalty in terms of computation time. The methodology used here was adapted from work done by Carrie Brown, a former graduate student in the Building Technology lab at MIT [57].

The data set is first sorted by the selected independent variable. Here, that would be heating, cooling, or other electric energy usage, depending on the needs of the user. It is then divided in two different ways – once into overlapping halves and once into overlapping quartiles. The overlap is intended to mitigate errors in the regression from edge cases. A multivariate regression analysis is performed on each of these seven tiers (full ordered data set, both halves, and all four quartiles), leading to seven regression formulas for each energy type. The percentage ranges defining the tiers [57] and the resultant counts for the data set used in the optimizer are given below in Table 13.

Tier	Min %	Max %	Min Count	Max Count
1	0	100	1	314928
2a	0	55	1	173210
2b	45	100	141717	314928
3a	0	27	1	85031
3b	23	52	72433	166912
3c	48	77	151165	242495
3d	73	100	229897	314928

Table 13: Tier limits for multistep regression

A set of graphs showing the ordering and subdivision process for the heating energy usage portion of the data set used here is given in Figure 27. The parameters defining this data set are given in Table 14, Table 15, and Table 16.

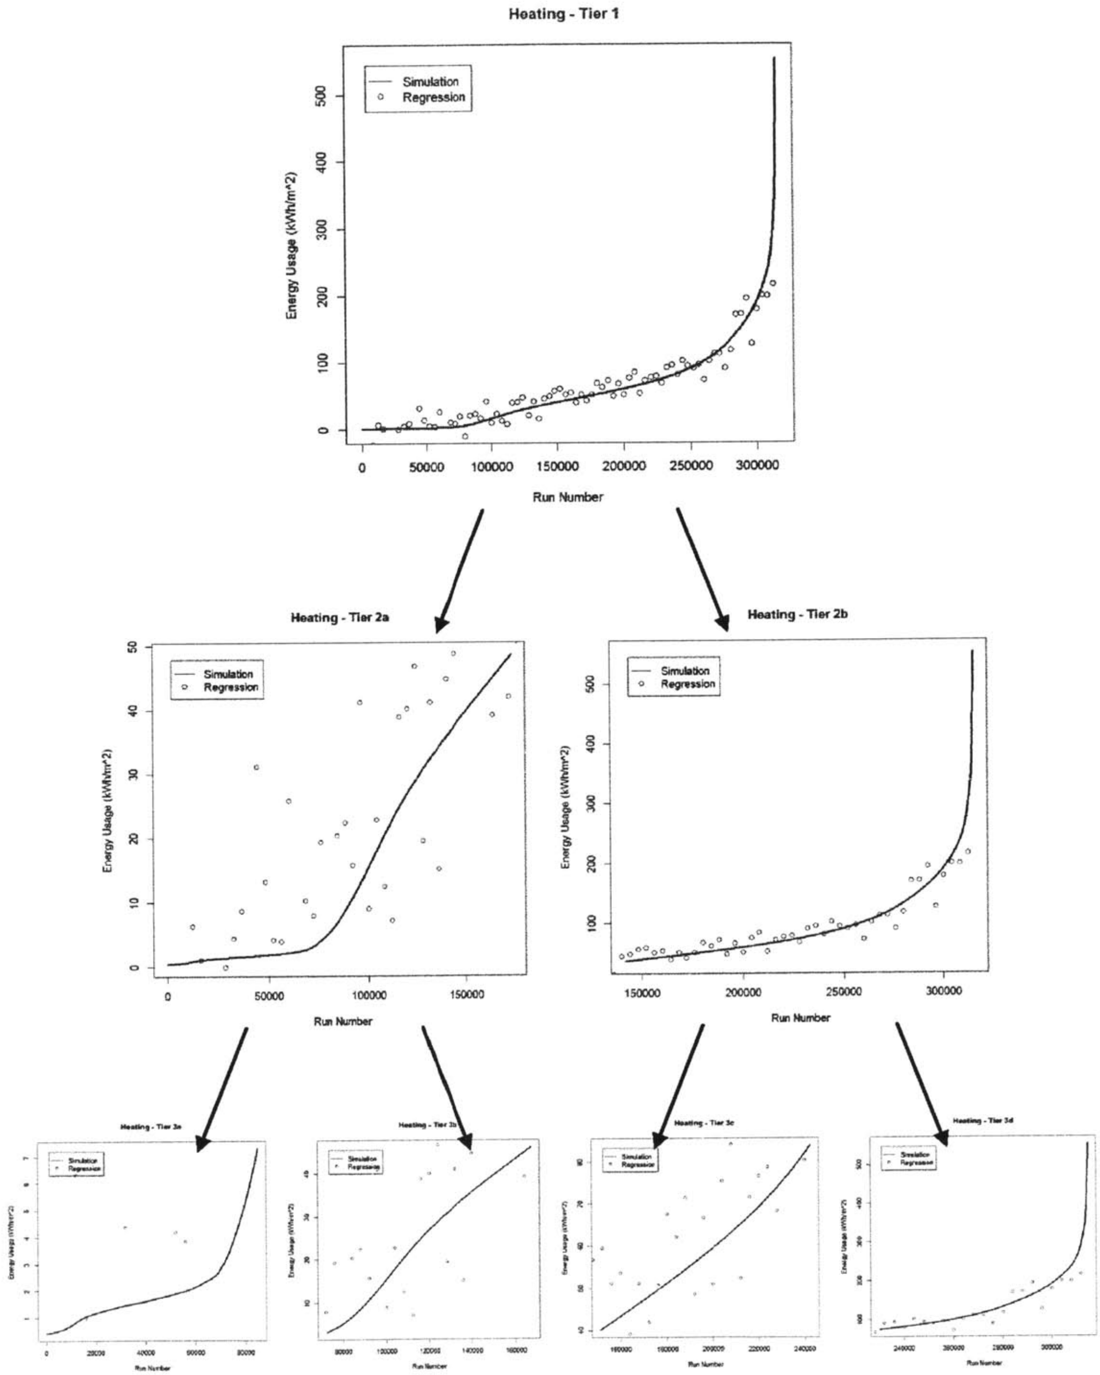

Figure 27: Subdivision of heating regressions

The equations derived from these regression analyses are used sequentially to predict the energy usage of a given target building. At each level, the input variables that define the building are used to evaluate the corresponding regression formula. Where the resultant predicted energy usage falls amidst the ordered set determines which equation will be used in the subsequent level of regression. A building found to be in the lower half of heating energy usage as per the Tier 1 regression would move on to Tier 2a, and so on. A single design can move separately through the regression subdivisions corresponding to each energy usage type, leading to a final predicted EUI that combines the three.

Optimizer

The optimizer developed here for Design Advisor makes use of both multistep multivariate regression and a genetic algorithm to determine the most efficient building design within the given constraints. The regression allows the algorithm to use a pseudo-elitist approach during the beginning of the optimization process, eliminating regions of the search domain that it knows, a priori, will not contain the desired solution. The genetic algorithm can then operate within the reduced search domain, providing greater accuracy than the regression can.

The first step in forming the multistep regression is, as discussed, performing multiple representative simulations across the relevant variable ranges. As more variables are being considered than in Carrie Brown's work, the number of steps for each variable has been reduced. A list of variables with their respective ranges and step sizes is given below in Table 14.

Input Variable Name	Input Variable Unit	Min	Max	# of Steps	
Lighting Requirements	lux	200	800	3	
Equipment Load	W/m^2	5	15	3	
Lighting Control Type	N/A			2	
Ventilation System	N/A			2	
Thermal Mass	N/A			2	
Roof Description	N/A			3	
Roof Insulation R-Value	(m^2-°C)/W				Twice wall R-value
Room Depth	m	4	10	3	
Window Area Percentage	%	20	80	3	
Window Type	N/A			3	
Glass Type	N/A			2	
Wall R-Value	(m^2-°C)/W	1	5	3	
				34992	Total runs
				5	Seconds per run
				174960	Total seconds
				2916	Total minutes
				48.6	Total hours
				2.03	Total days

Table 14: Multistep regression variables

The cities chosen are listed below. In order to increase the applicability of the regressions, the cities are converted from a single categorical variable to a set of five continuous numerical variables as in Carrie Brown's work. The five variables are latitude, heating and cooling degree days (using a 65°F baseline) [67], average relative humidity, and average global horizontal solar radiation. Latitude, relative humidity, and solar radiation were taken directly from the weather files used by Design Advisor. The solar radiation values were averaged after all null readings (implying complete darkness) were removed. One direction for future research is immediately visible once the cities are broken down in this manner: none of the nine cities are located in the Southern Hemisphere. This is due to the fact that the vast majority of the weather files included with Design Advisor are from cities in the Northern Hemisphere, but it could still be an important consideration if and when the list of available cities is expanded.

City	Latitude	HDD (65°F Baseline)	CDD (65°F Baseline)	Relative Humidity (%)	Solar Radiation (Wh/m^2)
Anchorage	61.22	10655	33	71.18	209.06
Boston	42.36	5412	903	65.72	326.02
Delhi	28.61	791	5312	56.45	455.69
Los Angeles	34.05	1537	1144	69.60	415.26
Pittsburgh	40.44	5508	1033	73.74	315.67
St. Louis	38.63	4817	1659	69.49	351.85
San Francisco	37.78	3254	215	73.74	395.69
Seattle	47.61	4789	357	73.23	282.61
Washington, D.C.	38.90	3947	1817	67.91	342.87

Table 15: Multistep regression cities and weather variables

The remaining variables are held constant:

Input Variable Name	Input Variable Unit	Value
Occupancy Start	Hours	7
Occupancy End	Hours	19
Person-desnity	People/m^2	0.1
Max Indoor Temp.	°C	26
Min. Indoor Temp.	°C	20
Max Relative Humidity	%	60
Fresh Air Rate	L/s per person	8
Air Change Rate	roomfuls/hour	1.8
Unoccupied Air Change Rate	roomfuls/hour	0
Max Temp. Unoccupied	°C	28
Min. Temp. Unoccupied	°C	18

Building Geometry	N/A	One-sided
Roof Insulation Location	N/A	Top
Number of Floors	N/A	1
Room Width	m	5
Room Height	m	3
Primary Façade Orientation	N/A	East
Window Overhang Depth	m	0
Blind Width	mm	15
Blind Schedule (Occupied Hours)	N/A	4
Blind Schedule (Unoccupied Hours)	N/A	2
Blind Angle When Closed	degrees	90
Slat Emissivity	N/A	0.2
Slat Absorptivity	N/A	0.9

Table 16: Inputs held constant during multistep regression

The predicted energy consumptions from these runs (broken into heating, cooling, and electricity) are shown below in Figure 28. They have been sorted with respect to their own usage type as for the multistep regression.

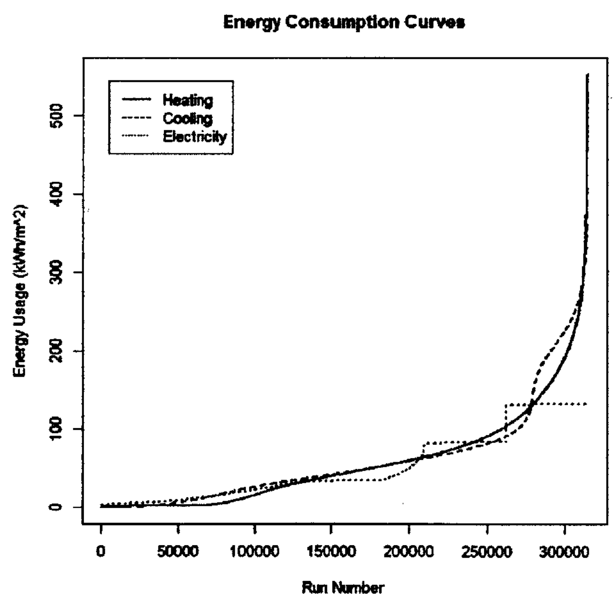

Figure 28: Predicted energy consumption from multistep regression simulations

The steps in the electricity usage plot come from the addition of HVAC and plug loads to Design Advisor's original lighting energy usage calculations. The rule of thumb used for HVAC-related energy usage depends only on room dimensions and air change rate, all of which are held constant with the exception of room depth, which varies between discrete levels rather than as a continuous numerical variable. Similarly, the energy usage attributed to miscellaneous plug loads is simply taken from the discrete "Equipment" input in Design Advisor. This means that when the lighting energy usage is constant (such as when the lights are always on, which is the least-efficient setting and would fall on the right-hand side of Figure 28), the only remaining possible variance is stepwise rather than smooth.

The regression and subdivision process was carried out as previously described using R. The coefficients for each of the 21 regression equations can be found in Appendix A. The root-mean-squared and average percentage errors for each of the tiers of each of the three energy usage types (heating, cooling, and other electricity) are given in Table 17.

Root-Mean-Square Error for Regression Subdivisions (kWh/m2)

	Tier 1	Tier 2a	Tier 2b	Tier 3a	Tier 3b	Tier 3c	Tier 3d
Heating	26.10	6.03	24.46	0.91	5.36	7.18	26.54
Cooling	14.88	5.48	15.49	2.99	4.29	5.52	16.65
Electricity	14.75	3.11	5.52	1.20	3.25	6.89	0.00

Average Percentage Error for Regression Subdivisions (kWh/m^2)

	Tier 1	Tier 2a	Tier 2b	Tier 3a	Tier 3b	Tier 3c	Tier 3d
Heating	-275.97	-42.88	-0.31	7.22	8.66	1.47	0.92
Cooling	-138.38	-83.07	0.98	27.26	2.39	0.92	0.28
Electricity	-35.08	1.44	0.86	0.27	2.05	1.76	-3.6E-16

Table 17: Root-mean-square and average percentage errors from multistep regression

The errors shown here compare favorably, both in trend and in value, to those found by Carrie Brown through her regression analysis (albeit with a different set of cities and selected variables) [57]. For heating and cooling, the primary cause of the relatively large RMSE for tiers 1, 2b, and 3d is the subset of highly-consuming or energy=inefficient buildings as seen at the right side of Figure 27. The electricity predictions are somewhat immune to this due to the limited combination of variables that directly affect them (lighting requirements, lighting control type, equipment loads, location, and window description). Conversely, percentage errors are greatest in magnitude for the buildings that consume the least energy, as even a slight error in absolute terms can give rise to a large percentage difference. Additionally, due to the nature of linear regression, highly negative energy usage is often predicted for buildings with near-zero (but obviously positive) actual energy usage of the selected type.

Using these regressions, a parametric search is performed across the variable ranges specified by the user. Each given numerical variable range is divided into five equal levels, and every combination of these levels, the varying categorical inputs, and the non-varying parameters is tested using the multistep regressions for the energy type of interest. This means that the total number of combinations depends on the number of varying parameters selected by the user. Through this, we can find the one mesh region within the variable space (as opposed to the entirety of the variable space) in which the simulation inputs for the optimal building design are most likely to be found. Note that a rough estimate of the value of those inputs can also be found, but the accuracy of this estimate is limited by the accuracy of the regression itself. A simple two-dimensional visualization of this search space reduction strategy is shown below.

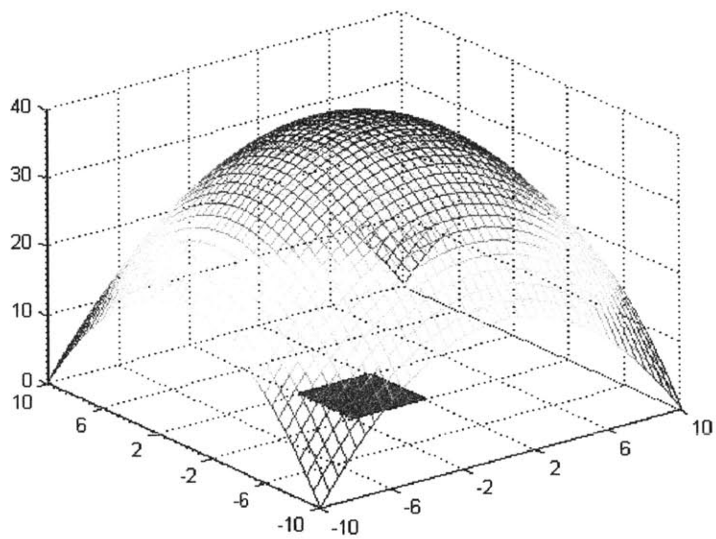

Figure 29: Simplified example of search space reduction

In this case, the local regression curve is modeled as a paraboloid of revolution. In this model, the regression was formed using data taken at the mesh points on the X-Y plane. The true maximum of the paraboloid lies at (0, 0, 40), which falls between these mesh points. By testing input values across the relevant variable ranges (-10 to 10 for both X and Y), we can identify the mesh points between which the optimal solution lies. This mesh region is highlighted on the X-Y plane in Figure 29. By doing this, the search space is reduced from 25 such mesh regions to one. The more variables/dimensions are considered, the more dramatic this reduction becomes.

Additional validation work was performed on several cities that were not included in the original regression:

City	Latitude	HDD	CDD	Rel. Hum. (%)	Avg. Sol. Rad. (Wh/m^2)
Denver	39.739167	6224	1118	52.0229	418.4259
Chicago	41.881944	6167	1128	70.7229	328.2181

Houston	29.762778	1351	3409	75.2237	299.986

Table 18: City data for multistep regression validation

HEATING	Run 1	Run 2	Run 1	Run 2	Run 1	Run 2
City	Denver	Denver	Chicago	Chicago	Houston	Houston
Light Req.	500	400	500	400	500	400
Equipment Load	5	8	5	8	5	8
Light Control	Indep. Dim	Always On	Indep. Dim	Always On	Indep. Dim	Always On
Ventilation Control	Mech	Hybrid	Mech	Hybrid	Mech	Hybrid
Thermal Mass	Low	High	Low	High	Low	High
Roof Type	Bitumen	Cool	Bitumen	Cool	Bitumen	Cool
Room Depth	5	4	5	4	5	4
Window Area %	60	40	60	40	60	40
Window Type	Single	Triple	Single	Triple	Single	Triple
Glass Coating	Clear	High perf.	Clear	High perf.	Clear	High perf.
Wall Insulation R-value	1	2	1	2	1	2
Design Advisor Result	255.92	158.07	304.96	192.00	46.97	19.45
Regression Result	178.90	48.82	209.54	81.15	96.37	45.02
% Difference	-30.10	-69.11	-31.29	-57.73	105.17	131.47

COOLING	Run 1	Run 2	Run 1	Run 2	Run 1	Run 2
City	Denver	Denver	Chicago	Chicago	Houston	Houston
Design Advisor Result	23.49	16.33	40.78	30.47	223.81	194.52
Regression Result	23.70	25.48	58.78	56.15	121.37	114.47
% Difference	0.89	56.03	44.14	84.28	-45.77	-41.15

ELECTRICITY	Run 1	Run 2	Run 1	Run 2	Run 1	Run 2
City	Denver	Denver	Chicago	Chicago	Houston	Houston
Design Advisor Result	11.64	82.87	13.17	82.87	9.16	82.87
Regression Result	12.95	66.28	15.14	65.68	10.88	61.24
% Difference	11.25	-20.02	14.96	-20.74	18.78	-26.10

TOTAL	Run 1	Run 2	Run 1	Run 2	Run 1	Run 2
City	Denver	Denver	Chicago	Chicago	Houston	Houston
Design Advisor Result	291.05	257.27	358.91	305.34	279.94	296.84
Regression Result	215.55	140.58	283.46	202.98	228.62	220.73
% Difference	-25.94	-45.36	-21.02	-33.52	-18.33	-25.64

Table 19: Comparison of results from Design Advisor and regression

These runs show markedly higher percentage error than the averages for the cities used in the regression process (see Table 17). Assuming the regressions themselves are sound, this is likely due to the selection of the subdivision points and the wide range of EUIs covered by the fourth quartile (tier 3d). As seen in Figure 28, designs with predicted EUIs (for any of the three energy types) of roughly 100 kWh/m^2 and above will fall into tier 3d. This covers a wide range of situations, from relatively-inefficient buildings in moderate climates to high-performing buildings in extreme environments. As a result (and as discussed above), residuals for individual designs using the tier 3d regression formula can be very large. There are several possible solutions to this issue. The most inefficiently-constructed buildings – for example, those with single glazed window units and meager wall insulation – could be removed from the data set, possible reducing the variety of disparate scenarios within each tier. Alternatively, all outliers with respect to EUI could be removed. This would eliminate designs with near-zero energy consumption, which often lead to negative energy usage predictions when using a regression, as well as the highly-consuming designs that cause the sharp changes in slope seen in Figure 28. Lastly, the subdivision points could be changed, effectively making the "mesh" finer where the errors are the most significant.

Once the best-performing set of parameters is found, the original variable maxima and minima are reduced to one step above and below the levels in the best combination, respectively. At this point, the optimization process is handed to the genetic algorithm. The initial chromosomes are based, when applicable, on the limits of the previously-identified mesh region (i.e., the reduced limits). For inputs that are allowed to vary but were not included in the regression simulations, the full user-defined range is employed. Each parameter has a gene length associated with it which, along with the range of possible values, defines the resolution of the variable. Each chromosome is built up by concatenating groups of binary digits corresponding to these gene lengths. Each bit is set using a uniform random variable. String-based inputs (e.g., roof composition or window type) are represented by a unique integer for each possible option. In cases where the input has fewer possible options than binary values representing it (such as thermal mass, which has three options but is represented by two bits), the extra binary values are assigned to Design Advisor's default for that variable. A simulation is performed based on the user's original inputs and the resultant energy consumption is used as the basis for comparison for the fitness function. This is true for whichever mode of energy usage the user chooses to investigate: heating, cooling, electric, or all three. Energy consumption in general is penalized, with consumption beyond the baseline punished more severely:

$$J = E_{simulated} + C * (E_{simulated} - E_{baseline}) \tag{31}$$

Where J is the cost associated with the given run and C is the additional penalty factor for excess consumption. For testing purposes, C has been set to 2. The fitness function is simply the inverse of the cost (so that it decreases with increasing energy usage):

$$f_i = \frac{1}{J_i} \tag{32}$$

As a proof of function of the genetic algorithm alone and to test the convergence criterion, a selection of runs was performed for an idealized shoebox room located in Boston. Here, the convergence criterion refers to the maximum variation between the sum of the fitness values from the previous generation and the sum from the current generation. When this falls below the set threshold and the minimum EUI from the current generation's simulations is less than that of the original design, the genetic algorithm's search is said to be complete. The best-performing design in the final generation is reported to the user. As discussed previously, it would be wrong to refer to this design as truly optimal, since there is no statistical guarantee that the genetic algorithm will find the global minimum EUI each time. As this criterion is based on a sum and not on an individual's fitness, it would, in reality, be linearly dependent upon the number of chromosomes in the generation. In this case, each generation contained 10 chromosomes. In addition, while there is some chance that the encoded simulations could change between generations without significantly altering this fitness sum, it is more likely that meeting this benchmark indicates some convergence on a set of building parameters. The results of the first set of tests, examining roof composition, roof insulation R-value, room width, window area as a percentage of wall area, and lighting requirements are shown below:

HEATING

	Roof	Roof R-Val	Width	Window Area %	Light Req.	EUI	% Diff	Gens	Time (s)	Criterion
Original	Bitumen	3	5	60	400	243.86				
Run 1	Cool	3.25	4.5625	45.625	375	223.39	-8.4	20	902.3	5.00E-06
Run 2	Green	3.4375	4.9375	49.375	425	227.13	-6.9	11	518.3	5.00E-06
Run 3	Bitumen	2.5	4.1875	42.8125	337.5	217.65	-10.7	31	1391.6	1.00E-06
Run 4	Bitumen	2.6875	4.75	48.4375	412.5	224.72	-7.8	20	903.9	1.00E-06

COOLING

	Roof	Roof R-Val	Width	Window Area %	Light Req.	EUI	% Diff	Gens	Time (s)	Criterion
Original	Bitumen	3	5	60	400	37.24				
Run 1	Cool	1	4.1875	41.875	325	28.2	-24.3	9	406.6	5.00E-06
Run 2	Green	1	4.375	43.75	350	28.2	-24.3	16	745.4	5.00E-06
Run 3	Green	3.4378	5.125	51.25	450	30.76	-17.4	27	1268.5	1.00E-06
Run 4	Green	3.25	4.375	43.75	350	28.2	-24.3	20	926.5	1.00E-06

	Roof	Roof R-Val	Width	Window Area %	Light Req.	EUI	% Diff	Gens	Time (s)	Criterion
ELECTRICITY										
Original	Bitumen	3	5	60	400	67.07				
Run 1	Cool	2.125	5.6875	57.8125	337.5	57.22	-14.7	6	273.4	5.00E-06
Run 2	Green	1.75	4	40	300	50.42	-24.8	18	839.6	5.00E-06
Run 3	Cool	2.5	4	40.9375	312.5	52.45	-21.8	6	280.3	1.00E-06
Run 4	Cool	2.5	4.375	43.75	350	58.69	-12.5	11	499.8	1.00E-06

	Roof	Roof R-Val	Width	Window Area %	Light Req.	EUI	% Diff	Gens	Time (s)	Criterion
TOTAL										
Original	Bitumen	3	5	60	400	348.17				
Run 1	Green	2.5	4.375	44.6875	362.5	315.63	-9.3	6	278.7	5.00E-06
Run 2	Bitumen	3.25	4	40	300	290.42	-16.6	14	634.9	5.00E-06
Run 3	Bitumen	3.4375	4.75	47.5	400	319.77	-8.2	15	593.2	1.00E-06
Run 4	Bitumen	3.25	4.1875	42.8125	337.5	301.72	-13.3	11	456.9	1.00E-06
						Average	-15.3	15.1	682.49	

Table 20: Genetic algorithm test results part 1

The limits for the selected variables are given in Table 21. Roof type, like other string-valued variables, does not have a set of limits associated with it.

Variable	Minimum	Maximum
Roof Type	N/A	N/A
Roof R-Value	1	4
Width	4	7
Window %	40	70
Light Req.	300	500

Table 21: Genetic algorithm test variables and ranges part 1

A second set of tests was performed using a different selection of parameters (room orientation, room depth, room width, window glass coating, and lighting control system) to ensure reproducibility over a wider variety of cases. Note that for lighting control, "Always On" refers to a system in which the lights are used to meet the lighting requirements regardless of the current natural lighting conditions, "Indep. Dim" refers to a situation in which each light fixture is individually dimmable and automatically responds to the lighting needs at its own location, and "Single Dim" refers to a situation in which all of the light fixtures in the room dim together in response to the lowest light level in the room.

HEATING

	Orientation	Depth	Height	Glass Coating	Light Control	EUI	% Diff	Gens	Time (s)	Criterion
Original	East	5	3	Clear	Always On	184.45				
Run 1	South	5	2.5	Clear	Always On	133.13	-27.8	12	488.1	5.00E-06
Run 2	Southwest	5.125	2.78	Clear	Indep. Dim	166.70	-9.6	15	619.3	5.00E-06
Run 3	Southwest	5.25	2.875	Clear	Indep. Dim	170.52	-7.6	17	700.4	1.00E-06
Run 4	South	5	2.5	Clear	Always On	133.13	-27.8	18	736.7	1.00E-06

COOLING

	Orientation	Depth	Height	Glass Coating	Light Control	EUI	% Diff	Gens	Time (s)	Criterion
Original	East	5	3	Clear	Always On	48.39				
Run 1	East	5.75	3.625	High Perf.	Single Dimt	30.58	-36.8	38	1551.3	5.00E-06
Run 2	North	5.5	3.344	High Perf.	Always On	27.26	-43.7	22	895.5	5.00E-06
Run 3	Northeast	5.5	3.344	High Perf.	Indep. Dim	26.67	-44.9	9	367.0	1.00E-06
Run 4	West	5.5	3.344	High Perf.	Single Dim	30.39	-37.2	28	1143.7	1.00E-06

ELECTRICITY

	Orientation	Depth	Height	Glass Coating	Light Control	EUI	% Diff	Gens	Time (s)	Criterion
Original	East	5	3	Clear	Always On	83.29				
Run 1	Northeast	4.125	2.688	Clear	Indep. Dim	12.86	-84.6	10	417.3	5.00E-06
Run 2	Northeast	4.25	2.969	Clear	Indep. Dim	12.90	-84.5	23	945.8	5.00E-06
Run 3	Southwest	5	2.594	Clear	Indep. Dim	13.60	-83.7	23	947.4	1.00E-06
Run 4	Southwest	4.125	2.781	Clear	Indep. Dim	12.20	-85.4	12	495.3	1.00E-06

TOTAL

	Orientation	Depth	Height	Glass Coating	Light Control	EUI	% Diff	Gens	Time (s)	Criterion
Original	East	5	3	Clear	Always On	316.13				
Run 1	East	4.125	2.688	Clear	Single Dim	250.35	-20.8	9	376.4	5.00E-06
Run 2	Northeast	5	2.5	Clear	Indep. Dim	230.98	-26.9	16	668.9	5.00E-06
Run 3	Northeast	5.125	2.688	Clear	Indep. Dim	238.00	-24.7	14	575.4	1.00E-06
Run 4	Southeast	5.375	3.063	Clear	Single Dim	236.47	-25.2	23	952.4	1.00E-06
					Average		-41.9	18.1	742.6	

Table 22: Genetic algorithm test results part 2

Variable	Minimum	Maximum
Orientation	N/A	N/A

Depth	4	6
Height	2.5	4
Coating	N/A	N/A
Light Control	N/A	N/A

Table 23: Genetic algorithm test variables and ranges part 2

The relationship between time taken (and thus number of generations simulated) and the resulting reduction in EUI (for the selected usage type) for the first set of tests is plotted in Figure 30. As shown, the only energy type for which there is a clear increase in optimizer performance with increased run time is heating. For electricity and total energy use, one of the four optimizations underperformed relative to the others. As discussed, this is a natural consequence of the stochastic nature of the genetic algorithm itself. For cooling energy, there seemed to be no relationship or even a slightly negative correlation between performance and run time at these convergence criterion levels. For an initial design that is already close to optimal, a relaxation of this test or the introduction of a limit on the total number of generations may be necessary. As a result, the proper choice for the convergence criterion remains unclear.

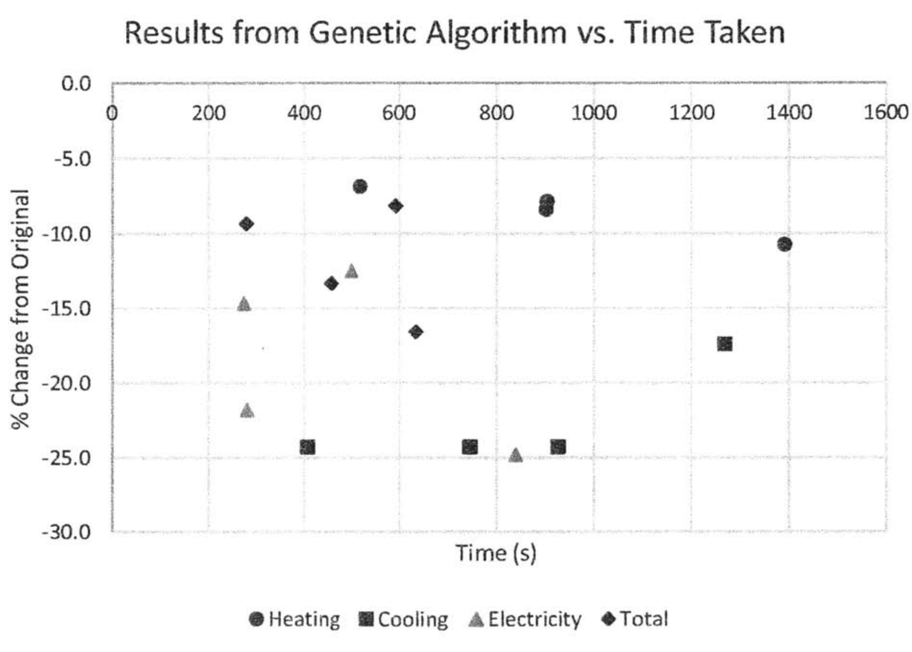

Figure 30: Percentage EUI reduction vs. time taken for first set of optimizer tests

With these separate tests done, the multistep regression and genetic algorithm were linked, as discussed, to form the final optimizer module. It was tested using the same sets of varying parameters as the genetic algorithm alone. In response to the results shown in Figure 30 and some of the more counterintuitive recommendations from the genetic algorithm (e.g., decreasing the roof insulation R-value to save heating energy), a stricter convergence criterion was included. The results are given below.

HEATING

	Roof	Roof R-Val	Width	Window Area %	Light Req.	EUI	% Diff	Gens	Time (s)	Criterion
Original	Bitumen	3	5	60	400	243.86				
Run 1	Bitumen	3	5	34.375	365	199.46	-18.2	12	378.3	5.00E-06
Run 2	Bitumen	3	5	35.5	380	201.05	-17.6	20	484.8	5.00E-06
Run 3	Bitumen	3.625	6.0625	42.25	390	210.32	-13.8	18	426.2	1.00E-06
Run 4	Bitumen	3.4375	4.9375	38.125	415	202.72	-16.9	18	453.3	1.00E-06
Run 5	Bitumen	3.25	4	34.375	365	197.14	-19.2	11	254.9	5.00E-07
Run 6	Bitumen	3.25	4.1875	34.75	370	197.77	-18.9	18	334.6	5.00E-07

COOLING

	Roof	Roof R-Val	Width	Window Area %	Light Req.	EUI	% Diff	Gens	Time (s)	Criterion
Original	Bitumen	3	5	60	400	37.24				
Run 1	Cool	3.4375	4.75	37	400	26.88	-27.8	7	193.7	5.00E-06
Run 2	Cool	3.8125	6.25	43	400	28.07	-24.6	6	148.3	5.00E-06
Run 3	Green	2.125	5.5	40	360	26.99	-27.5	8	204.9	1.00E-06
Run 4	Green	1.75	4.5625	36.25	390	26.51	-28.8	13	367.2	1.00E-06
Run 5	Cool	3.25	4.5625	36.25	390	26.68	-28.4	20	372.4	5.00E-07
Run 6	Green	3.25	4	34	360	25.84	-30.6	16	320.9	5.00E-07

ELECTRICITY

	Roof	Roof R-Val	Width	Window Area %	Light Req.	EUI	% Diff	Gens	Time (s)	Criterion
Original	Bitumen	3	5	60	400	67.07				
Run 1	Bitumen	3	4.5625	36.625	395	66.07	-1.5	22	598.7	5.00E-06
Run 2	Cool	1.375	6.0625	42.625	395	66.71	-0.5	5	143.1	5.00E-06
Run 3	Bitumen	1.375	5.6875	40.75	370	62.49	-6.8	14	379	1.00E-06
Run 4	Green	1.375	5.875	41.875	385	65.01	-3.1	18	517.5	1.00E-06
Run 5	Bitumen	1.75	4.1875	34.75	370	61.86	-7.8	9	173.2	5.00E-07
Run 6	Bitumen	3.25	4	34	360	60.16	-10.3	16	305.6	5.00E-07

TOTAL

	Roof	Roof R-Val	Width	Window Area %	Light Req.	EUI	% Diff	Gens	Time (s)	Criterion
Original	Bitumen	3	5	60	400	348.17				
Run 1	Bitumen	2.5	4	40	300	295.71	-15.1	16	435.3	5.00E-06
Run 2	Bitumen	3.4375	4.75	47.5	400	319.77	-8.2	22	592.3	5.00E-06
Run 3	Green	3.25	4.375	44.6875	362.5	310.4	-10.8	11	314.2	1.00E-06
Run 4	Bitumen	2.5	4.1875	42.8125	337.5	307.02	-11.8	14	382.3	1.00E-06
Run 5	Green	2.5	4	40.9375	312.5	300.12	-13.8	15	307.9	5.00E-07
Run 6	Cool	2.5	4	40	300	298.02	-14.4	8	155.9	5.00E-07
						Average	-15.68	14.04	343.52	

Table 24: Full optimizer test results part 1

HEATING

	Orientation	Depth	Height	Glass Coating	Light Control	EUI	% Diff	Gens	Time (s)	Criterion
Original	East	5	3	Clear	Always On	184.45				
Run 1	Southeast	4	2.59375	Clear	Single Dim	161.81	-12.3	33	960.8	5.00E-06
Run 2	Southeast	3.65	2.6875	Clear	Always On	153.98	-16.5	10	292.9	5.00E-06
Run 3	Southwest	4	2.5	Clear	Indep. Dim	162.17	-12.1	11	323.2	1.00E-06
Run 4	Southeast	4.1	2.875	Clear	Single Dim	176.47	-4.3	9	266.6	1.00E-06
Run 5	Southeast	3.6	2.5	Clear	Single Dim	154.82	-16.1	33	681.9	5.00E-07
Run 6	Southeast	4	2.59375	Clear	Single Dim	156.69	-15.1	7	151.8	5.00E-07

COOLING

	Orientation	Depth	Height	Glass Coating	Light Control	EUI	% Diff	Gens	Time (s)	Criterion
Original	East	5	3	Clear	Always On	48.39				
Run 1	Northeast	4.2	3.34375	High perf.	Indep.Dim	29.60	-38.8	15	377.4	5.00E-06
Run 2	North	4.05	2.78125	Clear	Always On	35.00	-27.7	6	130.7	5.00E-06
Run 3	East	3.9	3.625	High perf.	Single Dim	38.34	-20.8	15	344.6	1.00E-06
Run 4	West	4.2	3.25	High perf.	Single Dim	34.43	-28.8	22	498.7	1.00E-06
Run 5	Northeast	4.3	3.625	High perf.	Indep.Dim	30.36	-37.3	8	171.9	5.00E-07
Run 6	Northwest	3.8	3.25	High perf.	Single Dim	29.88	-38.3	18	383.2	5.00E-07

ELECTRICITY

	Orientation	Depth	Height	Glass Coating	Light Control	EUI	% Diff	Gens	Time (s)	Criterion
Original	East	5	3	Clear	Always On	83.29				
Run 1	East	4.35	3.90625	High perf.	Single Dim	16.68	-80.0	9	199.7	5.00E-06
Run 2	West	3.6	2.59375	Clear	Single Dim	14.97	-82.0	14	300.4	5.00E-06

Run 3	West	3.65	2.6875	Clear	Single Dim	14.88	-82.1	29	625.9	1.00E-06
Run 4	East	3.65	2.6875	Clear	Single Dim	15.24	-81.7	18	388.5	1.00E-06
Run 5	Southwest	3.85	3.53125	High perf.	Indep. Dim	13.24	-84.1	22	467.6	5.00E-07
Run 6	West	3.7	2.96785	Clear	Single Dim	14.56	-82.5	17	362.1	5.00E-07

TOTAL

	Orientation	**Depth**	**Height**	**Glass Coating**	**Light Control**	**EUI**	**% Diff**	**Gens**	**Time (s)**	**Criterion**
Original	East	5	3	Clear	Always On	316.13				
Run 1	South	5.125	2.6875	Clear	Always On	270.89	-14.3	16	341.1	5.00E-06
Run 2	East	4	2.5	Clear	Single Dim	242.25	-23.4	15	321.8	5.00E-06
Run 3	East	5.625	3.53125	High perf.	Single Dim	255.33	-19.2	18	391.0	1.00E-06
Run 4	Southeast	5.125	2.6875	Clear	Single Dim	218.18	-31.0	14	306.2	1.00E-06
Run 5	Northwest	5	2.59375	Clear	Single Dim	226.70	-28.3	33	662.3	5.00E-07
Run 6	Southeast	5	2.59375	Clear	Single Dim	205.89	-34.9	18	358.5	5.00E-07
						Average	-37.98	17.08	387.87	

Table 25: Full optimizer test results part 2

Combined, the regression and genetic algorithm find, on average, similar opportunities for reducing energy usage while taking fewer generations to do so than the genetic algorithm alone. The multistep regression calculations are completed on the order of several milliseconds, making their contribution to the overall run time negligible. Despite the fact that it is inaccurate on its own when used to predict a single EUI value, the parameter ranges obtained from the regressions appear to be broad enough to mitigate further error while still reducing the search space by a significant amount. The strictest convergence criterion option appeared to produce the most consistent (and generally lowest) final EUI value for each energy type. It is possible that there is a level that will generate only those results that are in some small local neighborhood around the true optimal design.

These initial results from the optimizer are promising. While it rarely settled on one set of parameters across multiple runs, certain trends did arise (e.g., reducing unshaded window area, reducing lighting levels, moving away from traditional bitumen roofs, etc.). For the most part, these changes conform to logical decisions that someone well-versed in building mechanics would make, acting as a basic sanity check for the optimizer. Since the optimizer converged on a solution so quickly each time (with a maximum run time of 16 minutes and an average of 6.1 across 32 tests), it would be possible for a designer to run it multiple times, noting these trends and using them to improve their original simulation.

Another key lesson is that existing designs can often be improved in a variety of different ways to reach the same end goal. As shown in the two sets of optimizer test results, multiple different parameter selections can lead to the same or similar final EUI. This allows designers to choose the most desirable option while still potentially reducing the environmental impact of their building.

Multivariate Regression for Individual Climates

Based on the results from the multistep regression on its own, an effort was also made to use the regression framework for single-city and single-climate energy usage predictions. By reducing the range of EUI results (especially for heating and cooling) and eliminating the location-dependent variables, the aim was to produce more accurate results for a single city and surrounding climatologically-similar locales. To facilitate this, additional simulations were performed for Boston using the remaining three cardinal façade orientations (north, south, and west, as east had been used in the original regressions). Regressions for each energy type were performed as before, with the addition of orientation as a variable and the removal of the five location-dependent variables (latitude, HDD, CDD, average relative humidity, and average solar radiation). The coefficients from these regressions are given in Appendix B. Tests were then performed using the same two sets of parameters as in Table 19 but for Boston and two other climatologically-similar cities (New York and Cleveland). The results from these tests are shown below.

HEATING	Run 1	Run 2	Run 1	Run 2	Run 1	Run 2
City	Boston	Boston	New York	New York	Cleveland	Cleveland
Orientation	South	North	South	North	South	North
Light Req.	500	400	500	400	500	400
Equipment Load	5	8	5	8	5	8
Light Control	Efficient	Always On	Efficient	Always On	Efficient	Always On
Ventilation Control	Mech	Hybrid	Mech	Hybrid	Mech	Hybrid
Thermal Mass	Low	High	Low	High	Low	High
Roof Type	Bitumen	Cool	Bitumen	Cool	Bitumen	Cool
Room Depth	5	4	5	4	5	4
Window Area %	60	40	60	40	60	40
Window Type	Single	Triple	Single	Triple	Single	Triple
Glass Coating	Clear	High perf.	Clear	High perf.	Clear	High perf.
Wall Insulation R-value	1	2	1	2	1	2
Design Advisor Result	152.75	80.22	122.33	60.67	169.07	85.62
Regression Result	167.39	85.73	167.39	85.73	167.39	85.73
% Difference	9.58	6.87	36.83	41.31	-0.99	0.13

COOLING	Run 1	Run 2	Run 1	Run 2	Run 1	Run 2
City	Boston	Boston	New York	New York	Cleveland	Cleveland
Design Advisor Result	31.67	21.90	46.00	34.76	27.62	22.44
Regression Result	36.37	21.30	36.37	21.30	36.37	21.30
% Difference	14.84	-2.74	-20.93	-38.72	31.68	-5.08

ELECTRICITY	Run 1	Run 2	Run 1	Run 2	Run 1	Run 2
City	Boston	Boston	New York	New York	Cleveland	Cleveland
Design Advisor Result	12.38	65.75	11.41	65.75	10.23	65.75
Regression Result	20.43	72.12	20.43	72.12	20.43	72.12
% Difference	65.02	9.69	79.05	9.69	99.71	9.69

TOTAL	Run 1	Run 2	Run 1	Run 2	Run 1	Run 2
City	Boston	Boston	New York	New York	Cleveland	Cleveland
Design Advisor Result	196.80	167.87	179.74	161.18	206.92	173.81
Regression Result	224.19	179.15	224.19	179.15	224.19	179.15
% Difference	13.92	6.72	24.73	11.15	8.35	3.07

Table 26: Comparison of results from Design Advisor and regression for single climate

The location-based parameters for these cities are summarized in Table 27.

City	Latitude	HDD (65°F Baseline)	CDD (65°F Baseline)	Relative Humidity	Avg. Solar Rad.
Boston	42.36	5412	903	65.72	326.02
New York	40.71	4640	1222	63.04	329.14
% Diff. from Bos.	-3.89	-14.26	35.33	-4.08	0.96
Cleveland	41.48	5702	1079	70.34	314.88
% Diff. from Bos.	-2.07	5.36	19.49	7.03	-3.42

Table 27: Location-based parameters for single climate test cities

Based on these preliminary results, it appears as if regression performed on a single city could be applied to cities with temperature distributions similar to the city used for the regression (as measured by latitude, heating degree days, and cooling degree days) with reasonable accuracy. To further validate this concept, additional tests will need to be done with a wider variety of climatologically-similar cities and parameter sets. The next step would be to reintroduce the location data by performing the regression analyses on a single climate rather than a single city. An end product could be a set of multivariate regressions corresponding to different climate groups around the country.

If these could be used during the genetic-algorithm-based optimization process, total processing time could be drastically reduced even with a tighter convergence criterion.

Retrocommissioning Savings Prediction Algorithm

Choice and Description of Variables

The goal of the retrocommissioning (RCx) savings prediction algorithm, much like for Design Advisor as a whole, is to generate reasonably accurate results in a transparent and tractable manner without requiring significant levels of detail or expertise from the user. As such, an effort was made to minimize both the number and complexity of the algorithm inputs. An evaluation of the key variables in Design Advisor led to the selection of three main input categories: location (which would determine climate and weatherization requirements), size (determining the technology used), and usage type (again, determining technologies used and relevant building codes). Examining the literature on RCx usage and effectiveness led to the inclusion of a fourth category, age (determining available technology and best practices at the time of construction).

To allow for the cleanest and most straightforward implementation of the k-nearest neighbors algorithm used here, these categories were then reduced to a set of numerical variables. These were heating and cooling degree days, gross square footage, energy usage intensity, and year of construction. The location category was split into heating and cooling degree days rather than a single measure of latitude due to potential issues surrounding rescaling a distance-related metric [68] as well as to allow for climatic differences across geographic regions.

Data Collection

Another goal of this algorithm, as stated previously, is to make use of historical data from real RCx projects rather than simulated savings whenever possible. As such, the process of data collection and the data set that has been built up formed the backbone of this work. Through a mix of publicly-available literature [69], [70] and partnerships with energy and RCx providers [24], [71], data on 90 individual properties have been collected, representing an average energy savings percentage (across all types of energy) of 16.75%. 44 of these are entirely new and are not contained in any publicly-available reports.

The primary issue during this process was that of completeness. Reporting on completed RCx projects, both by property owners and the companies involved, is inconsistent. Requiring enough information to find and fill in just the four categories discussed here led to significant culling of potential data points. For example, one report, "The Cost-Effectiveness of Commercial-Buildings Commissioning: A Meta-Analysis of Energy and Non-Energy Impacts in Existing Buildings and New Construction in the

United States" (Mills et al), contained information on 106 existing buildings that had undergone an RCx process. Of those, 46 entries met the standards listed here (i.e., they contain information on the variables discussed above).

There is also an understandable recalcitrance on the part of energy and commissioning providers to divulge potentially-identifying information about their clients. As such, all data used here has been made anonymous and restricted from viewing by the end user.

Each input variable is plotted against percent energy savings in the figures below. The plots can be thought of as projections of the 5-dimensional variable space onto the five mutually perpendicular planes formed by the basis vectors of the space. Note that the heating and cooling degree day plots are combined on to one set of axes to more clearly show their relation to one another.

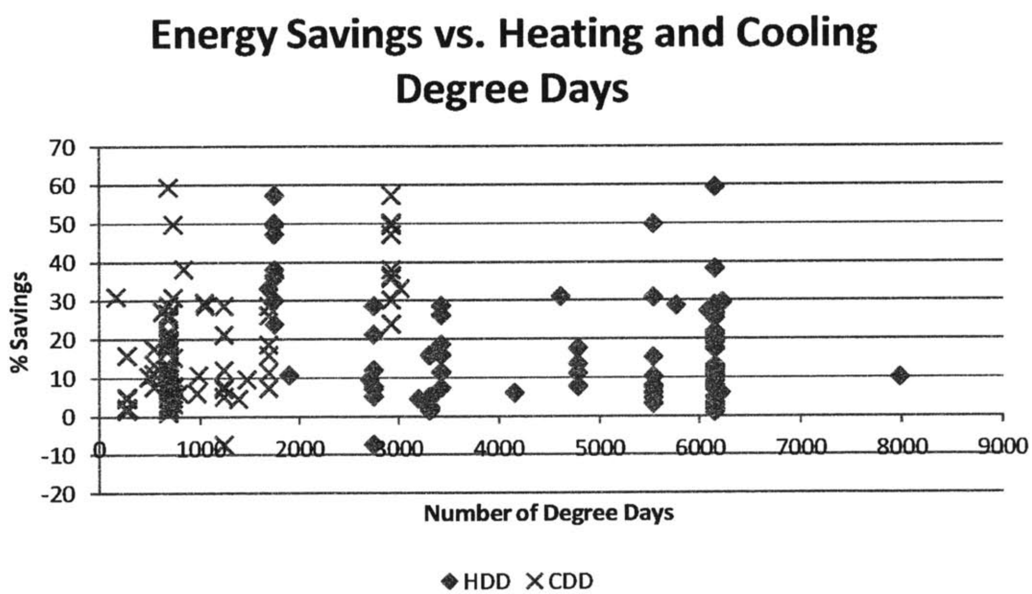

Figure 31: Percentage energy savings vs. heating and cooling degree days

The achieved savings are also compared to the total number of degree days (the sum of heating and cooling degree days):

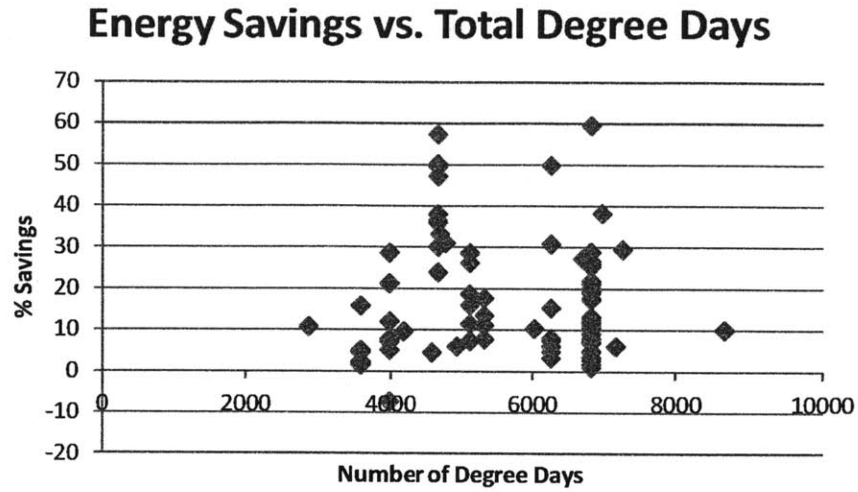

Figure 32: Percentage energy savings vs. total degree days

The columns of data points visible here denote sets of buildings in the same geographical region. Within each of those regions, a wide variety of energy savings have been realized. In most cases, the highest-resolution location information available was the city in which the building is located. For those, a variety of publicly-available sources were used to find an average number of yearly heating and cooling degree days (based on historical data) for the region in question. [72]–[75]

Overall building size, as shown below, appears to have the clearest relation to energy savings. Roughly speaking, the potential for these savings on a percentage basis diminishes with increasing building size, with a sharp change in slope between 300,000 and 400,000 square feet. Intuitively, it would make sense that a larger building would provide more opportunities for energy savings, leading to a greater absolute reduction in energy usage even if the percentage is smaller. However, given the limited number of data points at the tail end of this distribution, it is difficult to draw statistically valid conclusions.

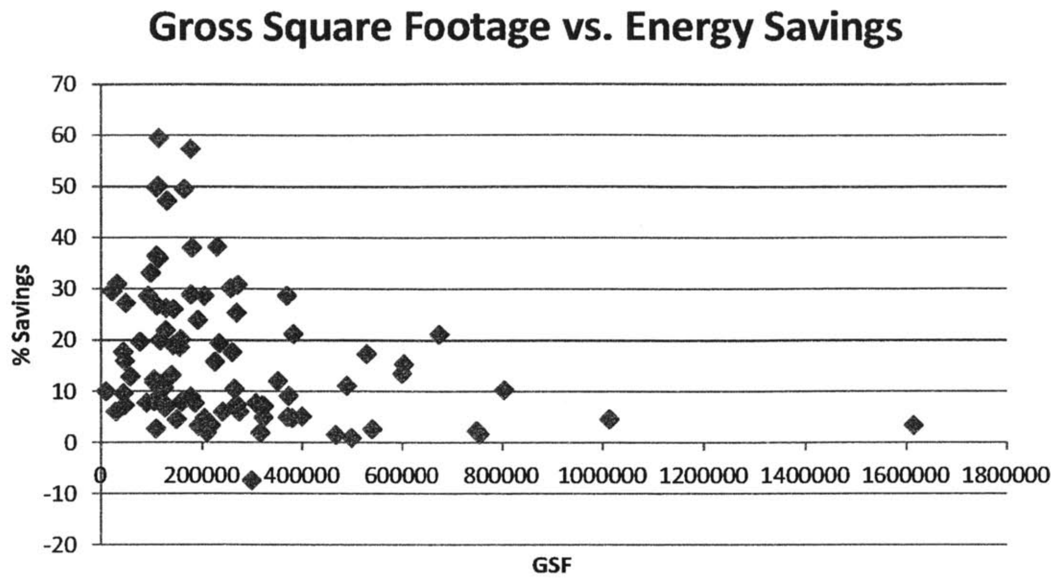

Figure 33: Building size vs. percentage energy savings

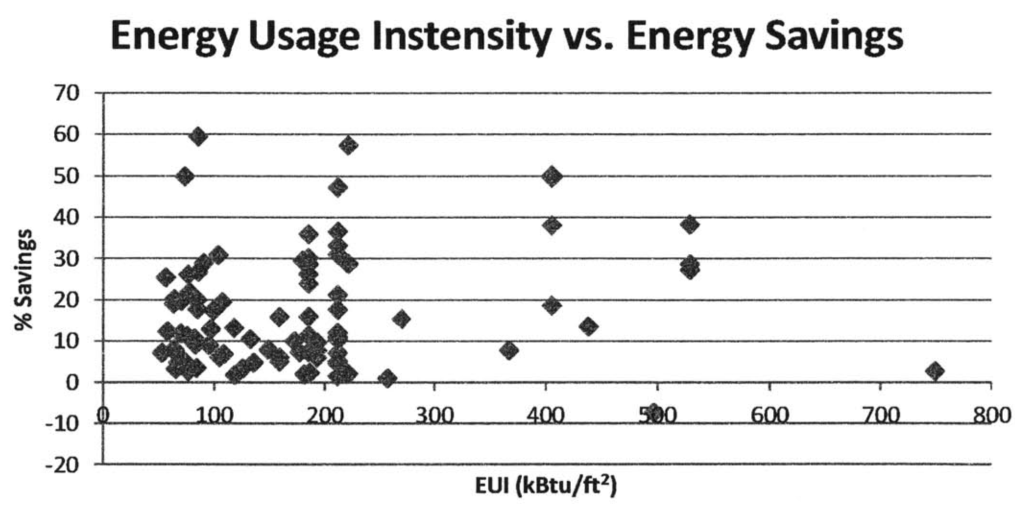

Figure 34: Energy usage intensity vs. percentage energy savings

As with building size, a higher EUI could indicate an opportunity for greater absolute energy savings, as it could denote a poorly-run or otherwise faulty building or an incorrect categorization in the data set.

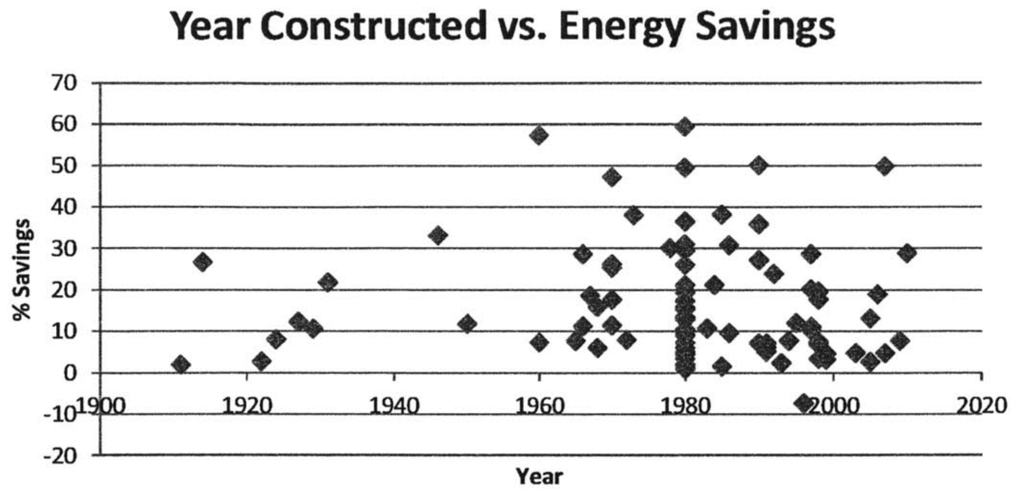

Figure 35: Year of building construction vs. percentage energy savings

There are two main reasons why building age alone is a poor predictor of energy savings. The first is that older buildings are likely to have been retrofitted or otherwise altered significantly at some point during their lifetimes. Depending on the changes that were made, the building's original construction year may no longer accurately represent the age of the primary energy-consuming technologies in the building. The second reason is that issues that arise in older buildings (superstructure degradation, subsystem failure, etc.) may well require complex fixes that lie outside the scope of most RCx investigations. More significant energy usage reductions could be attained using measures beyond RCx.

Taken as a group, these graphs reemphasize the point that no single variable is sufficient to calculate the reduction in energy usage for these buildings. As a whole, however, they could potentially characterize the conditions that make it possible to save energy through retrocommissioning and allow for a reasonably accurate prediction. For example, based on the relation apparent in Figure 33, the link between total building floor area and savings percentage was further examined:

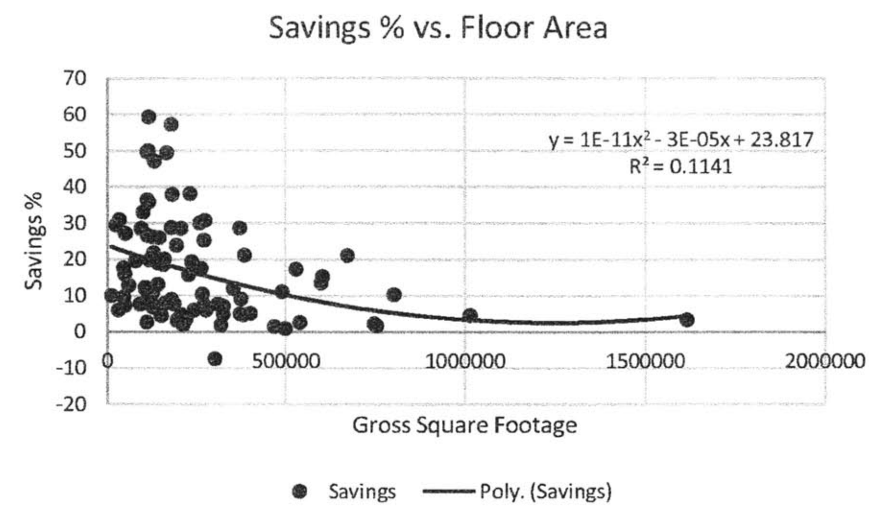

Figure 36: Building size vs. percentage energy savings with polynomial regression of order 2

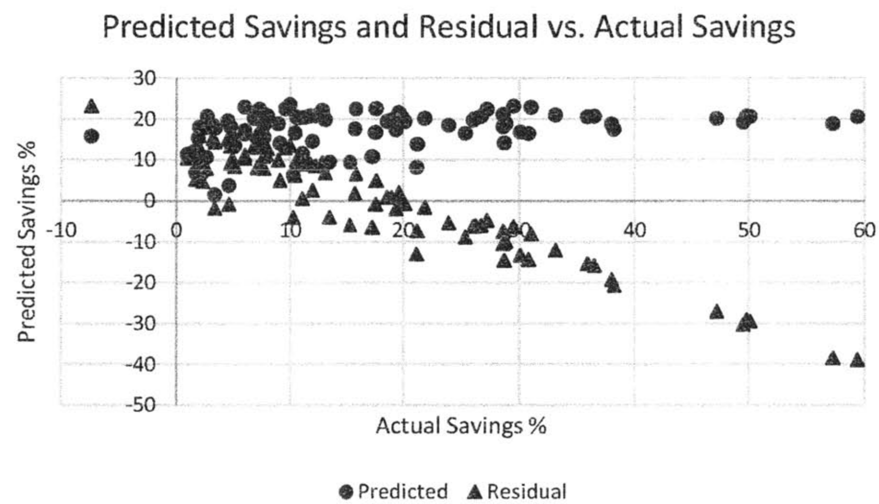

Figure 37: Predicted savings and residuals using polynomial fit

Clearly, while the available data implies some relationship, the variety of RCx outcomes for smaller buildings is far too wide to draw any strong conclusions. The limited number of very large buildings in the data set also reduces the applicability and meaningfulness of this kind of simpler analysis.

Algorithm Background

The key assumption behind the formulation of the savings prediction algorithm is that buildings belonging to the same set or cohort will develop similar problems throughout their lifetimes. That is, buildings of the same age, region, size, and usage type will be susceptible to similar modes of degradation and systems malfunctions that could be addressed through retrocommissioning. If we can categorize a building into the same "neighborhood" as other buildings on which we have data and this assumption holds, we can base a prediction of potential savings through RCx on the actual savings that were attained in the other buildings in the group.

The algorithm that was chosen to take advantage of this assumption is known, logically, as the k-nearest neighbors (KNN) algorithm. Data points can be thought of as situated in an n-dimensional space (where n is the number of similarity parameters by which new points are being evaluated). Predictions can be made about the nature of new samples based on their distance from the pre-existing data points in this space and their similarity to the nearest k of those points. The accuracy of these predictions depends on, among other things, the parameters chosen, the method of distance measurement, the presence and method of weighting, and the number of neighbors examined. A simple example of the effect of the selection of k is shown below.

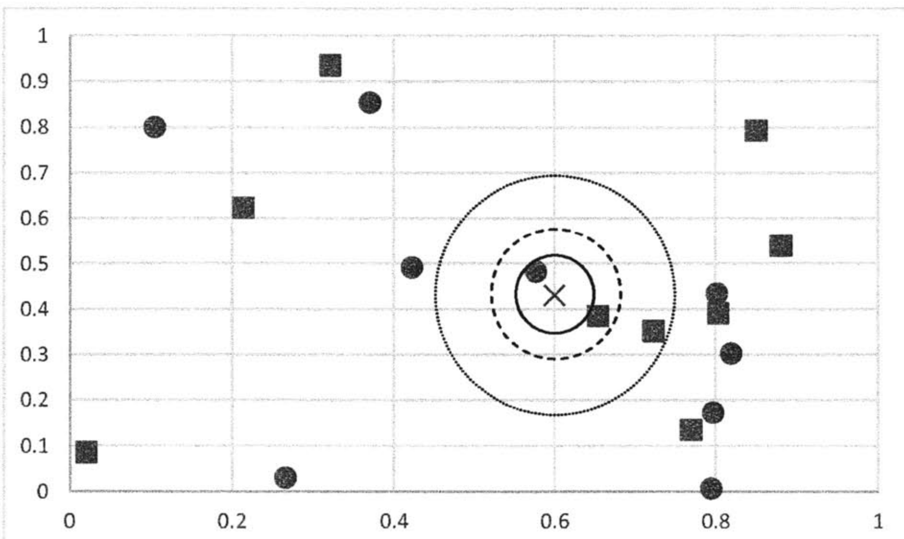

Figure 38: Effect of size of neighborhood on KNN algorithm categorizations

Here, the solid line represents a neighborhood size of 1, the dashed line represents a size of 2, and the dotted line represents a size of 3. If the goal is to categorize the new sample (denoted by an X) as either a square or a circle, these three neighborhood sizes will yield three different answers – circle, undetermined, and square, respectively. While one might be able to make an intuitive choice in this simple two-dimensional example, doing that becomes more difficult as multiple dimensions are added. This is why ensuring the completeness and integrity of the data used, as discussed in the previous section, was important – while the missing parameters could be filled in themselves using the KNN algorithm, any inaccuracy in one of the measured dimensions could lead to drastic changes in the predicted energy savings (depending on the number of neighbors surveyed).

The example above implicitly used a Euclidean measure of distance to determine the meaning of "nearest" neighbors:

$$D^2 = \sum_i (x_i - y_i)^2 \tag{33}$$

That is not the only option, especially for discrete or non-numerical data. For discrete variables with a meaningful innate definition of distance (e.g., pixels on a screen or location on a gridded map), Manhattan distance (also known as cityblock or taxicab distance) [76] may be best:

$$D = \sum_i |x_i - y_i| \tag{34}$$

Manhattan distance allows for simple decomposition of the individual parameter contributions to the total distance. It also results in an interesting change to the notion of equidistance – a circle (the locus of points equidistant from some center) becomes a square on the grid. For example, each of the points on the figure below is four units from the origin.

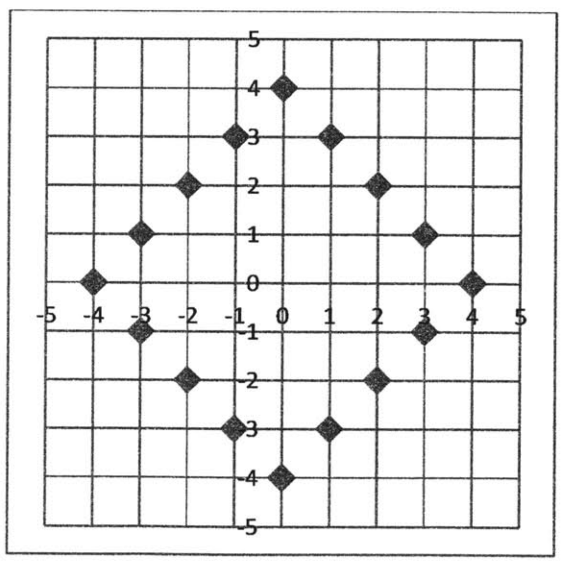

Figure 39: Circle of radius 4 using Manhattan distance

Both Euclidean distance and Manhattan distance are special cases of the more generalized Minkowski distance [68] (for q=2 and q=1, respectively):

$$D = (\sum_i |x_i - y_i|^q)^{\frac{1}{q}} \tag{35}$$

Another possibility for data that can be referred to by vectors is cosine similarity [77], [78]:

$$sim(x, y) = \frac{\bar{x} \cdot \bar{y}}{|\bar{x}||\bar{y}|} \tag{36}$$

Cosine similarity could be useful in cases where the scaling of each parameter is either inconsistent or not meaningful in terms of one of the more conventional distance measurements given above. It yields the cosine of the angle between the two vectors corresponding to the two samples being compared and thus ranges from -1 to 1.

Finally, as mentioned previously, there are also options for non-numerical variables or parameters. A common example is Hamming distance [79]:

$$D_{Hamming} = \sum_i f(x_i, y_i) \tag{37}$$

$$f(x_i, y_i) = \begin{cases} 1, & x_i = y_i \\ 0, & x_i \neq y_i \end{cases}$$

Hamming distance can be used, for example, to compare categorical similarity (where x_i and y_i are considered equal if both samples belong to the same category) or for character-wise string comparisons (where x_i and y_i are equal if the character at position i is the same for both strings).

These distance measures, in turn, raise the questions of parameter scaling and weighting. It is likely that not all of the parameters chosen for evaluation will have the same units and, furthermore, that those units will be of vastly different scale. In the case of the variables selected here, building age might be measured in tens of years, location in degrees latitude and longitude (from 0° to 90° and 0° to 180°, respectively), and EUI in hundreds of kWh/m^2. These differences in magnitude can cause one variable to dominate the distance calculation, reducing the importance of the others in judging similarity regardless of physical meaningfulness. One solution to this is to rescale and non-dimensionalize each variable such that they all range from 0 to 1, inclusive, as follows:

$$x_{i,rescaled} = \frac{x_i - x_{i,min}}{x_{i,max} - x_{i,min}} \tag{38}$$

Scaling variables in this manner allows continuous numerical and categorical variables (using the aforementioned Hamming distance) to be mixed.

Scaling, in turn, emphasizes the need for weighting. All scaled variables, unweighted, are treated as if they signify overall similarity equally. To correct this, we can augment each parameter's contribution using attribute weighting [80]:

$$D^2 = \sum_i w_i^2 (x_i - y_i)^2 \tag{39}$$

Where w_i is the individual weighting constant for the i^{th} parameter. One can also weight each neighbor's contribution to the final prediction (for continuous variables) using its distance to the test point:

$y_{test} = \sum_j f(D_j) * y_j$ (40)

These two methods of weighting allow for significant flexibility in determining the set of nearest neighbors and categorizing the test point. If no physically meaningful attribute weights are known, an evolutionary algorithm is often employed to find optimal values for the given data set.

Final Algorithm

The final prediction algorithm makes use of a scaled Euclidean distance measure as well as both attribute and neighbor contribution weighting. At first, attribute weights were set equal to the individual coefficients of determination from a linear regression of each attribute against the actual energy savings percentages. However, in an effort to improve upon those initial results, they were later chosen through multiple runs of a genetic algorithm (GA) working to minimize:

$J = \sum_i |predicted_i - actual_i|$ (41)

Each property is, in turn, isolated from the data set and a prediction of its energy savings is made using the remaining buildings. The individual errors are them summed. The chromosomes are made up of five 12-digit binary-encoded genes, each corresponding to one of the attribute weights. As each weight ranges from zero to one, this gives the genes a resolution of 1/4096 or 2E-4. The fitness function (which increases as the chromosomes converge on the optimal solution) was chosen to rescale the cost function to avoid very large values and values close to zero:

$fitness = \exp(-\sqrt[4]{J})$ (42)

The genetic algorithm used to find the attribute weights was implemented in Matlab. For ease of testing, the initial RCx prediction algorithm was also written in Matlab before being ported to Java (for compatibility with Design Advisor).

Results and Discussion

Results from the algorithm described above are summarized in the figures and table below. As shown in Figure 40, the GA-weighted predictions reproduce the actual distribution of energy savings much more closely than the R^2 weighted set. However, both fail in the same way: they fail to predict both modest (0-5%) and high (>40%) energy savings, instead over-predicting savings between 15 and 20% by a factor of four.

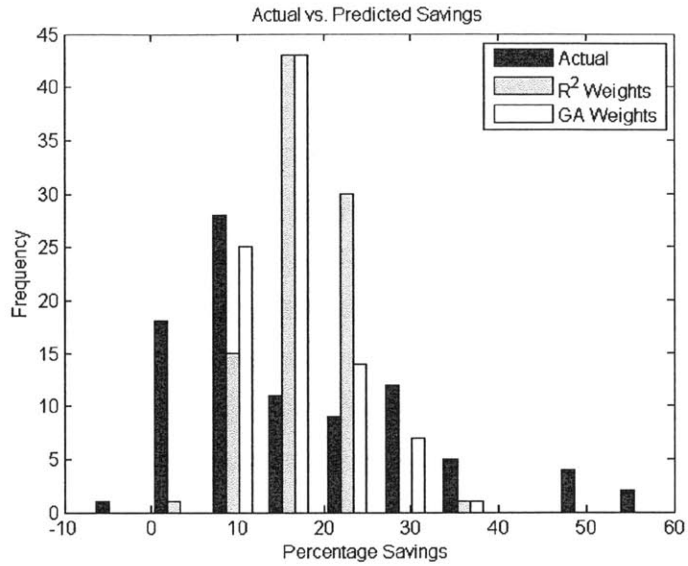

Figure 40: Actual savings vs. predictions using R^2 and genetic algorithm attribute weights

A more direct comparison of the predictions and the corresponding actual savings is shown in Figure 41. The more accurate the predictions, the closer the points will get to a straight line with a slope of unity. There is significant spread in the predictions produced using the GA attribute weights, reducing the R^2 value of the correlation to the actual savings to 0.20. The predictions using the individual R^2 attribute weights (not shown for clarity) actually correlate slightly negatively with the true savings.

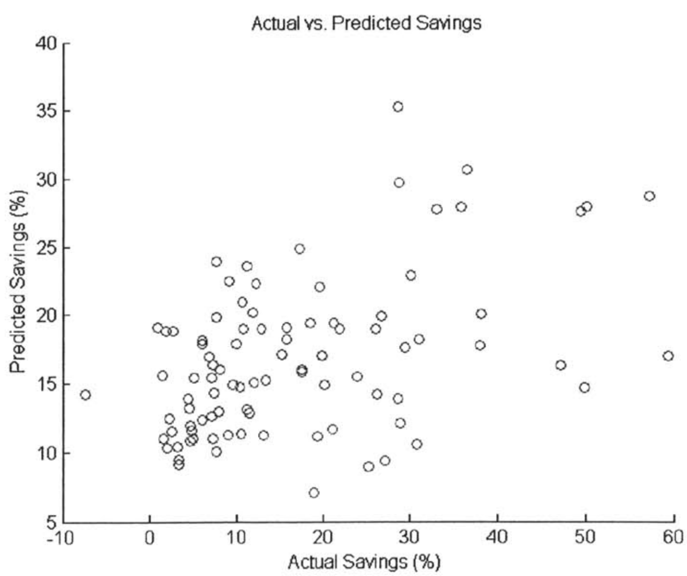

Figure 41: Actual savings vs. predictions using genetic algorithm attribute weights

Finally, three different measures of error (average error, sum of error magnitudes as used in the genetic algorithm, and root-mean-square error) are shown for the two attribute weighting schemes described here as well as no (or equal) weighting. As seen in Table 28, R^2 weighting performs worse in terms of all but average error than no weighting at all.

	Attribute Weighting		
	None	R^2 Value	GA
Average Error	-0.82	0.52	0.02
Sum of Error Magnitudes	1130.9	1132	876.4
Root-Mean-Square Error	15.5	15.7	12.4

Table 28: Error comparison for attribute weighting schemes

The nature of the algorithms used here, specifically the genetic algorithm used to determine the attribute weights, leaves this work open to future improvement in several ways. The first would be to further explore the solution space traversed by the genetic algorithm to see if a more optimal set of weights could be found. Due to the stochastic nature of this method, it is possible for the search to return a locally optimal solution (a local minimum of total error, in this case) instead of the desired global optimum. A possible solution to this problem would be to add mutations to the chromosomes at random, which would serve to perturb the algorithm enough to avoid the local minimum. This would require reexamining and resetting the threshold value used to determine when a solution has been found (as mutations will lead to an increase in generation-to-generation variance) as well as longer run times.

The cost function itself could also be changed. As seen above, the current algorithm over-predicts energy savings in the middle of the range of values (around 15-20%). This is a result of the fact that the cost function penalizes the sum of the magnitude of the prediction errors across the 90 properties tested. By predicting that many buildings will save a moderate amount of energy, the overall error can be minimized without producing individual accurate predictions. Therefore a superior cost function might penalize any error, not just their sum. This could lead to a set of weights that would yield somewhat larger (and fewer) individual errors, but that may be, for the purposes of prediction, preferable to smaller but more omnipresent errors. The best solution may be to continue working with RCx providers and utilities to obtain a larger, more robust, more detailed data set.

Conclusions

The collection of work presented in this thesis is aimed at one overarching goal: popularizing and facilitating awareness of building energy usage throughout the entire lifespan of a structure. Design Advisor in its original form was aimed at the early planning phases, when many design details would be undecided or still open to change. The retrocommissioning savings prediction algorithm was intended to expand the usability of Design Advisor to include existing buildings as well. Through the changes discussed here and the lessons learned through the real-world validation studies that were done, this position has been strengthened – simulations are faster and more accurate and the outputs relate more directly to real-world usage. The conclusions drawn from the results of these various products of research are discussed below.

Run Time Improvements

The various run time improvements discussed previously - modifying the thermal mass calculations, reducing the daylighting calculations, and implementing a representative day formulation to limit the number of days simulated per year - were integral to the other work described here, especially the optimizer. Using the original version of Design Advisor would have made it more difficult and time-consuming to plan, code, and test an optimization algorithm. For example, even the restricted set of simulations used for the initial regressions done here would have taken three times longer to compute (15 second per run compared to 5). The same is true for the genetic algorithm portion of the optimizer. Given the number of simulations per generation and the fact that hundreds or thousands of generations can pass before a suitable solution is found, the total time required can quickly balloon. For a web-based application, speed is paramount, so longer simulations might have ruled out genetic algorithms as a possible basis for the optimizer.

Additionally, by reducing the time taken for each simulation without violating the integrity of the results, we increase Design Advisor's accessibility and take another step toward near-instantaneous feedback on the energy efficiency of a proposed building design (see "Future Work" for further discussion of this goal). It reduces the chance that the user will worry that the program has stopped responding and allows them to test a wider variety of designs in the same amount of time, possibly leading to more efficient building plans.

With these improvements and positive results come some negative side effects. Introducing further approximations into the model will inevitably impact the accuracy of the predictions. As discussed in their relevant sections, the thermal mass and lighting modifications led to relatively small changes in accuracy. Simulating only a reduced subset of representative days, however, causes more problems. As demonstrated previously, the individual monthly errors for moderate climates are significant enough to prohibit the use of said monthly predictions. The average error remains small, allowing for the use of total yearly energy usage predictions (such as in the optimizer), but the user would not be able, for example, to compare heating or cooling season energy usage across designs. Further changes that could be made remain to be seen. As discussed, even given these reductions in run time, there are various required calculations that take a non-trivial amount of time to perform. This can be seen through the following breakdown of the run time of a simulation using the original version of Design Advisor:

Description	Time (ms)
Entire simulation	25481
Final data storage and manipulation	1054
Hourly energy usage calculations	14936
Supporting hourly calculations	470
Daylight solve	8915
Building creation	106

Table 29: Run time breakdown using original Design Advisor code

Even if one could remove thermal mass, daylighting, and day-to-day energy calculations in their entirety, just creating the building and formatting and storing the energy predictions would take more than a second per simulation. At a certain point, continued gains in simulation speed might require rewriting large portions of the program.

Design Advisor Validation

While Design Advisor and its constituent theoretical components (specifically the window unit solver, thermal mass calculations, and daylighting distribution prediction) have been measured against industry-standard software, the program had not previously been rigorously tested against a real-world building [23], [81]. As a part of the work done under the auspices of the Department of Energy's Energy Efficient Buildings Hub Project, this testing was done with three such buildings (two in the Philadelphia area and one on MIT's campus). This served to more concretely confirm the potential accuracy of Design Advisor in a realistic use case, at least relative to the limits of modeling as a discipline – the human element of building operation remains "very difficult (or nearly impossible) to estimate" [82], even for more complex software.

The two Philadelphia-area buildings were simulated fairly successfully, with a total mean bias error (for heating, cooling, and electricity on a kWh-equivalent basis) of -11.2% for One Montgomery Plaza and 2.6% for Navy Yard Building 101. The third building, MIT Building E40, was somewhat more problematic: the simulations resulted in a mean bias error of -76.6% for electricity, 3.7% for chilled water, and -50.4% for steam. Of the three, the most accurately predicted energy source, chilled water, was the only one for which no issues in the actual building were later discovered. The electricity underestimation can be explained, at least in part, by the previously-unreported presence of a server cluster in the basement. Without knowing the specifications of the computers, a better estimate cannot be made, but servers are known to be a significant source of electricity consumption. Similarly, when the provided data were expanded to include earlier years, the high level of steam usage observed during the cooling season was revealed to be a recent (and unexplained) development. The MIT Department of Facilities is currently investigating the site for evidence of a steam valve leak that may account for the increased consumption.

The three simulations highlight the fact that further details about building operation are needed for complete and certain knowledge of agreement between Design Advisor and actual utility bills. Errors arising from unknown or incorrectly-assumed factors both infrastructural (such as the servers in MIT E40) and operational (details about occupant activity, for example) can represent a significant proportion of the energy usage predicted by a simulation engine. This is true regardless of the software that is used.

Lastly, these studies uncovered several usability issues that had not previously been considered. Once the simulations were complete, additional calculations had to be performed to include important factors that were not accounted for in the original program. These include but are not limited to the inclusion of fan and AHU energy usage, displaying total electricity usage (including equipment loads) rather than just lighting usage, and displaying site energy usage along with primary energy usage. As discussed earlier, the lighting energy consumption calculation grew out of older daylighting simulation code, and thus was not chiefly concerned with general electricity usage. Depending on the user's selections, however, HVAC equipment and other equipment loads can be far more significant than lighting alone. If Design Advisor is intended to promote awareness of energy efficiency in early designs, however, it would be counterintuitive and potentially misleading to not include them. Similarly, while primary energy usage is a valuable measure for judging the overall impact of a building (as it includes losses during generation and transmission), it is probably not as immediately useful to a designer or property owner as site energy usage, which would be reflected on a utility bill. Including these changes will not only improve Design Advisor for the target user group but will also make it so future validation studies require fewer post-simulation calculations.

Optimizer

The multistep regression used here leveraged the various improvements summarized above to expand upon the work done previously by Carrie Brown [57]. By focusing on single-sided simulations and drastically reducing the run time of an individual simulation, the data set for the regressions grew from 7 cities to 9 and, through an increase in the number of levels for each variable, 48,383 total runs to 314,928. This allowed for a much more varied and detailed set of multistep multivariate regressions, although the run-time-reducing changes that were employed necessitated the prediction of yearly rather than monthly energy usage values.

As discussed, the initial results from the multistep regression do not suggest that it could be used (in its current form) on its own to reliably predict building energy usage. While the median prediction error is small, individual run errors can be substantial, either on a percentage or absolute basis. This seems to be especially true for buildings in cities that were not included in the original regression, despite the fact that the categorical "city" variable was replaced with five continuous numerical variables that could be extended to cover any location. It is plausible that, given the distribution of simulations that went into the regression analyses and characteristics of the designs most likely to be simulated with Design Advisor, the regression subdivision points need to be re-examined. Using smaller subsets for the more energy-intensive designs (as opposed to roughly equal quartiles) could reduce these errors and allow for standalone use of the regressions. As the regression formulas yield a result in milliseconds as opposed to seconds, this would open up many new opportunities for Design Advisor as a tool.

The results from the optimizer as a whole suggest the possibility of real usability and potential for impact in the future. Across a variety of simple test cases, the genetic algorithm alone suggested changes that would lead to substantial reductions in building EUI (upwards of 25%) despite relatively short total run time (the bulk of which is taken up, on average, by 16.6 generations of 10 simulations each, or 714 seconds in total for a single representative room). Including the search space reduction from the multistep multivariate regression further improved these design suggestions.

Future work on the optimizer would, by necessity, focus at first on finding and fixing the source of the regression errors discussed above. From there, the convergence criteria used by the genetic algorithm should be more comprehensively tested to find a set that will consistently yield design improvements without incurring unnecessary simulations due to overly-strict requirements. From a user and usability standpoint, it would also be beneficial to rewrite how parameter ranges are handled to allow for subsets of string-based variables to be selected for investigation (as opposed to the current all-or-nothing approach for variables such as thermal mass thickness and window unit type).

Retrocommissioning

The results from the retrocommissioning savings prediction algorithm are less clear. The module was intended as a proof-of-concept that would show that potential energy savings could be predicted (within a reasonable confidence interval) given only a basic and limited set of information about a building. There is some evidence that this can be done, but it is difficult to ascertain the causes of the errors and the proper way to correct them.

The most significant positive contribution of the RCx algorithm is the data set itself. As evidenced by the size and sparseness of other such collections, it can be difficult to bring together a list of even basic but not publicly-available information about a variety of properties. This is especially true for energy savings data, as commissioning providers are understandably cautious about possibly violating the privacy of their clients. Even if the prediction algorithm itself turns out to be impractical, this data set identifying 6 variables across 90 buildings (44 for which information is not otherwise available) could prove useful for other projects in the future. Otherwise, as discussed previously, the results from the algorithm itself are mixed but encouraging. While the average prediction error across all of the tested properties is very low (0.02%), individual predictions remain suspect (12.4% RMSE, which is larger than the actual savings achieved in many buildings, and an R^2 of 0.2 between the predicted and actual savings). Much of this error can be attributed to poor estimates for buildings that displayed energy savings either significantly higher or significantly lower than the average. While outliers can be difficult to work with in any analysis, it is important that the algorithm not overestimate potential energy savings, which would give the property owner a false sense of how much to invest in RCx measures.

The possible sources of error are highly varied, as are the potential solutions. Some or all of the attribute weightings could be incorrect. This could require a change to the parameters of the genetic algorithm use to determine the weights or a deeper investigation into the possibility of an analytic relationship that might determine them (although the simplest solution, using the R^2 values from individual variable regressions, was unsuccessful). Mixed-use buildings (which cannot be identified reliably using EUI alone) or facilities that are outliers for their type in terms of energy usage intensity could be confounding the results. A larger data set could solve this by allowing for individual sets of weights for each building usage type and the elimination of outliers within each category. Some of the entries in the data set might be erroneous in some way. As achieved savings are often self-reported, it would be infeasible to validate each entry. The assumptions underlying the k-nearest neighbors algorithm might not apply to this situation; each building might need to be treated as a wholly unique case. In that case, the data set would still be useful and informative, but a different machine learning strategy would need to be applied. This is also not an exhaustive list, as there could be other heretofore unconsidered issues.

Future Work

The work described here in no way represents the limits of this avenue of research (simplified building modeling and energy usage estimation). There are a wide variety of further improvements in applicability, usability, and beyond that have yet to be implemented. These can be split into two broad categories: issues for further research and programming issues. A variety of unexplored and unfinished ideas concerning both Design Advisor as well as retrocommissioning are discussed below.

One step in reducing the simulation runtime was separating the discretization of the typical room geometry from the calculation of solar gains for each building side. It may be worthwhile to take this further and separate out the daylighting visualization code as well. During optimization, it takes up processing time for a result that will never be seen by the user or used in subsequent simulations (but is needed for energy usage calculations in some cases). Even during normal use, there is no guarantee that the user's primary interest is daylighting. If it were made optional, the runtime for the average simulation would be reduced. However, given the tight integration of the visualization calculations into the overall lighting code, this would require an extensive rewrite.

Another potential change would be to expand the chiller options available to the user. At the moment, all cooling energy calculations in Design Advisor are done with an assumed chiller coefficient of performance of 3, regardless of locale, actual building needs, or other factors. A project that was started but not completed as part of the Hub was the development and implementation of a series of simple regressions that would automatically provide the user with a set of chiller options akin to those given for insulation (i.e., commercial or residential system with low, medium, or high efficiency). This was to involve performing a series of chiller simulations across various climates and given various cooling requirements and regressing the resultant COPs against the relevant inputs in Design Advisor.

This would allow the COPs displayed to the user to change realistically in response to the selected building parameters rather than remaining static.

One initial target deliverable for this work was a portable standalone/offline application to allow Design Advisor to be used more easily at prospective building sites. In its current form, Design Advisor requires web access (which may not be available in all cases) and its user interface is not optimized for mobile device touchscreens. Porting Design Advisor to iOS and/or Android would be a significant undertaking, but would go a long way towards the general goal of facilitating building energy modeling during the early stages of planning and construction. Creating a mobile-friendly version of the current site would be significantly easier but would not allow the simulations to be run entirely locally unless the multistep regressions were used. It would, however, be a good first step and would gauge the prevalence of or interest in this type of on-site usage.

For the RCx prediction algorithm, the best direction for future work is unclear. The first priority would be to expand the set of documented RCx projects. Ideally, there would be enough properties listed to be able to create separate algorithms for each building usage category (office, food service, classroom, etc.). As stated previously, the KNN algorithm relies on the assumption that data points that appear similar (as per the selected attributes) are in fact similar to one another. It is possible that mixing obviously dissimilar groups of buildings together acts as a confounding factor and reduces the efficacy of the algorithm. If, for example, the distribution of savings differed between building types, the attribute weightings would necessarily be different as well. Combining the groups (and therefore combining the distributions and weightings) would disguise that fact. The majority of the buildings in the current data set are used for education or office space, meaning that other categories would have to be populated before this change would be possible.

There are three main directions for research concerning the algorithm itself and its implementation. The first is the cost function. The current cost function penalizes the total deviation of all savings predictions from the actual achieved savings percentages. There are likely other cost functions that would more successfully reduce individual errors as well. The second is the attribute weightings. Obviously, these would change if a different cost function were used, but they could also change given more runs of the genetic algorithm if it found a more effective set. This would simply require time for more runs or for runs using a higher fitness threshold. The third direction is the algorithm itself. If, once more data was collected, it turned out that the similarity assumption underlying the KNN algorithm did not apply to this problem, a new method would have to be found.

There are numerous machine learning algorithms that could be applied to the problem, or perhaps multivariate regression could be revisited and different variable transformations could be explored.

The development of a piece of software like Design Advisor is an ongoing process. There will always be new technologies that change assumptions or require new research (such as chillers) or new practices or standards that push innovation and raise new questions (like retrocommissioning). By keeping up with these advances, we can ensure that Design Advisor makes the greatest possible impact on building energy usage.

Appendix A: Multistep Regression Coefficients

Heating

Tier	1	2a	2b	3a	3b	3c	3d
Intercept (kWh/m^2)	1.42E+02	3.62E+02	6.70E+02	2.36E+02	3.70E+02	3.74E+02	1.17E+03
Latitude	-9.69E-01	-3.64E+00	-5.58E+00	-4.62E-01	-3.80E+00	-3.09E+00	-1.01E+01
HDD	2.33E-02	9.53E-03	2.81E-02	4.34E-03	1.36E-02	1.36E-02	3.51E-02
CDD	2.35E-03	-3.32E-03	1.10E-02	-9.61E-03	-2.73E-03	2.08E-03	1.05E-02
Rel. Hum.	-1.33E+00	-1.14E+00	-2.78E+00	-3.05E+00	-1.21E+00	-1.40E+00	-4.60E+00
Avg. Sol. Rad.	-4.41E-02	-3.70E-01	-7.85E-01	NA	-3.68E-01	-3.67E-01	-1.44E+00
Light Req.	-1.22E-02	-5.22E-03	-2.06E-02	-4.81E-04	-1.01E-02	-1.08E-02	-2.29E-02
Equipment Load	-1.55E+00	-6.12E-01	-2.64E+00	-5.84E-02	-1.20E+00	-1.39E+00	-2.99E+00
Light control (Indep. Dim)	5.44E+00	2.05E+00	9.10E+00	2.22E-01	3.99E+00	4.91E+00	1.03E+01
Ventilation control (mech)	-7.70E-13	-2.06E-12	1.22E-12	8.10E-05	6.40E-13	1.49E-04	2.87E-12
Thermal mass (low)	1.75E+00	8.05E-01	2.93E+00	-4.15E-02	1.64E+00	1.83E+00	2.90E+00
Roof type (cool)	2.59E+00	7.28E-01	4.43E+00	1.65E-01	1.31E+00	2.55E+00	5.83E+00
Roof type (green)	2.02E+00	6.88E-01	3.41E+00	1.57E-01	1.22E+00	2.00E+00	4.13E+00
Room depth	-5.16E+00	-9.22E-01	-8.88E+00	-1.71E-01	-1.82E+00	-3.39E+00	-1.34E+01
Window area %	4.69E-01	9.82E-02	8.03E-01	1.70E-02	1.92E-01	3.04E-01	1.14E+00
Window type (single glazed)	2.48E+01	3.84E+00	4.01E+01	7.12E-01	7.45E+00	1.42E+01	5.73E+01
Window type (triple glazed)	-6.66E+00	-1.60E+00	-1.24E+01	-2.64E-01	-3.19E+00	-6.26E+00	-1.70E+01
Coating type (high performance)	-1.09E+01	-1.62E+00	-1.88E+01	-4.15E-01	-2.99E+00	-5.95E+00	-2.93E+01
Wall insulation R-value	-5.98E+00	-1.53E+00	-1.06E+01	-1.68E-01	-3.11E+00	-5.60E+00	-1.33E+01

Table 30: Heating energy multistep regression coefficients

Cooling

Tier	1	2a	2b	3a	3b	3c	3d
Intercept (kWh/m^2)	-1.88E+02	2.78E+01	-3.81E+02	2.03E+01	1.91E+02	-3.77E+02	-1.06E+03
Latitude	4.22E-01	-1.15E+00	-3.21E-02	-4.95E-01	-2.79E+00	1.38E+00	6.21E+00
HDD	3.49E-03	2.31E-03	1.38E-02	5.34E-04	3.77E-04	1.06E-02	1.75E-02
CDD	4.15E-02	2.53E-02	3.91E-02	1.72E-02	1.72E-02	3.35E-02	4.78E-02
Rel. Hum.	8.52E-01	4.25E-02	1.11E+00	NA	-3.10E-01	8.88E-01	3.34E+00
Avg. Sol. Rad.	2.61E-01	2.26E-02	6.97E-01	NA	-1.44E-01	6.28E-01	1.41E+00
Light Req.	1.31E-02	6.48E-03	1.78E-02	2.37E-03	8.31E-03	1.01E-02	2.41E-02
Equipment Load	1.89E+00	1.04E+00	2.50E+00	4.15E-01	1.31E+00	1.50E+00	3.36E+00
Light control (Indep. Dim)	-9.29E+00	-4.87E+00	-1.22E+01	-1.88E+00	-6.10E+00	-7.14E+00	-1.64E+01
Ventilation control (mech)	1.27E-12	2.79E-13	2.79E-12	1.63E-04	4.13E-13	3.69E-04	-1.15E-12
Thermal mass (low)	1.75E+00	1.28E+00	2.03E+00	7.70E-01	1.47E+00	1.48E+00	2.40E+00
Roof type (cool)	-7.55E+00	-4.02E+00	-9.73E+00	-1.70E+00	-4.91E+00	-5.73E+00	-1.26E+01
Roof type (green)	-7.40E+00	-3.73E+00	-9.80E+00	-1.51E+00	-4.65E+00	-5.52E+00	-1.32E+01
Room depth	-2.18E+00	-3.46E-01	-3.55E+00	1.78E-02	-6.37E-01	-1.23E+00	-5.69E+00
Window area %	2.67E-01	6.80E-02	4.15E-01	1.03E-02	1.07E-01	1.66E-01	6.46E-01
Window type (single glazed)	-5.21E+00	-4.67E+00	-4.11E+00	-2.67E+00	-4.33E+00	-4.90E+00	-7.47E-01
Window type (triple glazed)	2.76E+00	1.86E+00	2.73E+00	5.72E-01	2.35E+00	2.10E+00	2.62E+00
Coating type (high performance)	-8.37E+00	-2.15E+00	-1.28E+01	-3.28E-01	-3.32E+00	-5.06E+00	-1.95E+01
Wall insulation R-value	-1.33E-01	5.34E-01	-7.97E-01	3.31E-01	5.44E-01	1.01E-01	-1.92E+00

Table 31: Cooling energy multistep regression coefficients

Electricity

Tier	1	2a	2b	3a	3b	3c	3d
Intercept (kWh/m^2)	-1.06E+03	2.38E+01	4.89E+00	-2.04E+01	1.89E+01	-3.49E+01	-3.99E-08
Latitude	6.21E+00	1.93E-01	1.16E-01	2.79E-01	2.29E-01	5.14E-01	1.40E-14

HDD	1.75E-02	-2.19E-04	-1.07E-04	7.05E-05	-2.40E-04	-6.94E-05	6.26E-17
CDD	4.78E-02	-3.52E-04	-5.27E-04	4.30E-04	-1.89E-05	7.90E-04	-1.04E-15
Rel. Hum.	3.34E+00	-8.20E-02	-1.16E-01	3.56E-02	-2.71E-02	7.94E-02	-5.58E-13
Avg. Sol. Rad.	1.41E+00	-7.07E-03	-7.29E-03	1.66E-02	-3.12E-03	2.50E-02	-7.07E-15
Light Req.	2.41E-02	3.60E-02	2.60E-02	1.60E-01	2.66E-02	1.56E-01	1.62E-01
Equipment Load	3.36E+00	4.16E-02	1.11E-02	1.71E-02	3.63E-02	-1.98E-02	1.20E-02
Light control (Indep. Dim)	-1.64E+01	-2.80E+01	0.00E+00	-8.20E+01	-2.33E+01	-8.15E+01	-4.89E+01
Ventilation control (mech)	-1.15E-12	-6.02E-03	-1.51E-05	3.55E-04	-3.53E-03	1.50E-02	-2.61E-13
Thermal mass (low)	2.40E+00	-2.65E-03	-1.51E-05	3.55E-04	-3.53E-03	6.23E-03	-2.62E-13
Roof type (cool)	-1.26E+01	3.61E-04	-2.61E-14	-3.14E-15	5.30E-03	-4.46E-03	-1.75E-16
Roof type (green)	-1.32E+01	2.17E-03	2.26E-05	-5.33E-04	5.30E-03	-6.42E-03	3.92E-13
Room depth	-5.69E+00	1.03E+00	6.29E-01	5.20E-01	1.13E+00	7.74E-01	1.91E-01
Window area %	6.46E-01	-6.16E-02	-3.53E-02	-3.85E-02	-6.82E-02	-7.09E-02	-1.51E-17
Window type (single glazed)	-7.47E-01	-5.33E-01	-2.90E-01	-3.66E-01	-6.42E-01	-5.65E-01	4.22E-13
Window type (triple glazed)	2.62E+00	4.99E-01	3.13E-01	4.34E-01	5.52E-01	7.16E-01	1.35E-14
Coating type (high performance)	-1.95E+01	1.09E+00	6.26E-01	8.36E-01	1.23E+00	1.48E+00	-2.96E-13
Wall insulation R-value	-1.92E+00	9.90E-06	5.65E-06	-4.40E-06	1.85E-05	-1.11E-05	-1.05E-13

Table 32: Electric energy multistep regression coefficients

Note: Entries labeled "NA" are those from which R failed to determine a coefficient due to a reported singularity.

The following variables were held constant:

Input Variable Name	Input Variable Unit	Value
Occupancy Start	Hours	7
Occupancy End	Hours	19
Person-desnity	People/m^2	0.1
Max Indoor Temp.	°C	26
Min. Indoor Temp.	°C	20
Max Relative Humidity	%	60
Fresh Air Rate	L/s per person	8
Air Change Rate	roomfuls/hour	1.8
Unoccupied Air Change Rate	roomfuls/hour	0
Max Temp. Unoccupied	°C	28
Min. Temp. Unoccupied	°C	18
Building Geometry	N/A	One-sided
Roof Insulation Location	N/A	Top
Number of Floors	N/A	1
Room Width	m	5
Room Height	m	3
Primary Façade Orientation	N/A	East
Window Overhang Depth	m	0
Blind Width	mm	15
Blind Schedule (Occupied Hours)	N/A	4

Blind Schedule (Unoccupied Hours)	N/A	2
Blind Angle When Closed	degrees	90
Slat Emissivity	N/A	0.2
Slat Absorptivity	N/A	0.9

Table 33: Parameters held constant during multistep multivariate regression

Appendix B: Multistep Regression Coefficients – Boston Only

Heating

Tier	1.00E+00	2a	2b	3a	3b	3c	3d
Intercept	1.39E+02	1.02E+02	1.50E+02	8.42E+01	8.24E+01	1.04E+02	1.57E+02
Orientation - North	9.26E+00	3.60E+00	1.19E+01	2.25E+00	2.19E+00	3.39E+00	1.49E+01
Orientation - South	-1.20E+01	-7.27E+00	-1.39E+01	-5.48E+00	-3.95E+00	-4.74E+00	-1.79E+01
Orientation - West	2.06E+00	8.69E-01	2.69E+00	5.68E-01	5.63E-01	8.59E-01	3.35E+00
Light Req.	-1.63E-02	-1.35E-02	-1.64E-02	-1.26E-02	-6.04E-03	-6.55E-03	-1.71E-02
Equipment Load	-2.12E+00	-1.61E+00	-2.29E+00	-1.49E+00	-7.18E-01	-8.84E-01	-2.43E+00
Light control (Efficient)	7.86E+00	5.68E+00	8.33E+00	4.97E+00	2.71E+00	3.09E+00	9.21E+00
Ventilation control (mech)	2.61E-13	-3.73E-13	-7.48E-04	8.89E-05	4.88E-04	4.65E-04	-4.74E-14
Thermal mass (low)	2.75E+00	2.02E+00	3.03E+00	1.82E+00	9.88E-01	1.19E+00	3.13E+00
Roof type (cool)	4.04E+00	1.95E+00	5.93E+00	1.17E+00	1.24E+00	2.34E+00	7.25E+00
Roof type (green)	3.28E+00	1.78E+00	4.59E+00	1.17E+00	1.07E+00	1.82E+00	5.53E+00
Room depth	-6.65E+00	-2.55E+00	-9.03E+00	-1.78E+00	-1.44E+00	-2.37E+00	-1.33E+01
Window area %	5.94E-01	2.64E-01	7.39E-01	1.95E-01	1.46E-01	1.85E-01	9.66E-01
Window type (single glazed)	3.43E+01	1.14E+01	4.21E+01	7.25E+00	6.63E+00	1.04E+01	5.46E+01
Window type (trple glazed)	-9.32E+00	-5.56E+00	-1.15E+01	-4.03E+00	-3.18E+00	-4.11E+00	-1.49E+01
Coating type (high performance)	-1.35E+01	-3.37E+00	-2.00E+01	-1.88E+00	-2.06E+00	-3.97E+00	-3.00E+01
Wall insulation R-value	-8.13E+00	-5.62E+00	-9.83E+00	-3.33E+00	-3.25E+00	-3.75E+00	-9.00E+00

Table 34: Heating energy multistep regression coefficients - single city

Cooling

Tier	1.00E+00	2a	2b	3a	3b	3c	3d
Intercept	2.57E+01	2.08E+01	3.21E+01	1.78E+01	2.96E+01	3.44E+01	3.48E+01
Orientation - North	-1.08E+01	-8.16E+00	-1.25E+01	-6.43E+00	-8.60E+00	-1.04E+01	-1.37E+01
Orientation - South	-2.19E+00	-1.41E+00	-3.37E+00	-9.05E-01	-2.66E+00	-3.93E+00	-3.88E+00
Orientation - West	-1.38E+00	-1.53E+00	-1.33E+00	-1.30E+00	-1.57E+00	-1.53E+00	-9.83E-01
Light Req.	1.09E-02	1.31E-02	7.60E-03	1.47E-02	9.78E-03	7.36E-03	5.93E-03
Equipment Load	1.57E+00	1.75E+00	1.29E+00	1.83E+00	1.48E+00	1.22E+00	1.12E+00
Light control (Efficient)	-7.75E+00	-8.66E+00	-6.40E+00	-9.20E+00	-7.38E+00	-6.24E+00	-5.64E+00
Ventilation control (mech)	2.44E-13	-5.76E-15	1.22E-04	4.70E-04	6.75E-05	-1.07E-04	4.63E-14
Thermal mass (low)	2.59E+00	2.55E+00	2.74E+00	2.54E+00	2.83E+00	2.80E+00	2.95E+00
Roof type (cool)	-6.15E+00	-5.46E+00	-6.85E+00	-4.42E+00	-5.87E+00	-6.90E+00	-6.33E+00
Roof type (green)	-6.19E+00	-5.55E+00	-6.91E+00	-4.50E+00	-6.00E+00	-7.03E+00	-6.44E+00
Room depth	-1.03E+00	-9.45E-01	-1.44E+00	-8.03E-01	-1.42E+00	-1.83E+00	-1.80E+00
Window area %	1.41E-01	1.46E-01	1.73E-01	1.48E-01	1.85E-01	2.09E-01	2.02E-01
Window type (single glazed)	-5.95E+00	-3.20E+00	-6.16E+00	-1.89E+00	-1.77E+00	-2.78E+00	-6.59E+00
Window type (trple glazed)	2.67E+00	2.79E+00	2.46E+00	2.48E+00	2.30E+00	2.33E+00	1.49E+00
Coating type (high performance)	-4.76E+00	-4.79E+00	-5.54E+00	-4.47E+00	-5.67E+00	-6.43E+00	-6.16E+00
Wall insulation R-value	3.76E-01	4.06E-01	-5.86E-02	2.21E-01	-4.68E-01	-5.80E-01	-4.05E-01

Table 35: Cooling energy multistep regression coefficients - single city

Electricity

Tier	1.00E+00	2a	2b	3a	3b	3c	3d
Intercept	2.42E+01	2.18E+01	2.97E+01	2.21E+01	3.20E+01	3.25E+01	3.26E+01
Orientation - North	6.84E-01	1.12E+00	7.63E-01	1.27E+00	1.69E+00	1.85E+00	6.91E-01
Orientation - South	-6.32E-01	-8.71E-01	-8.86E-01	-1.37E+00	-2.33E+00	-2.19E+00	-8.43E-01
Orientation - West	-3.25E-01	-2.94E-01	-2.67E-01	-3.43E-01	-2.80E-01	-2.14E-01	-1.52E-01
Light Req.	1.07E-01	1.15E-01	9.70E-02	1.21E-01	1.07E-01	1.00E-01	9.06E-02
Equipment Load	1.20E-02	-5.29E-02	-5.97E-02	-1.34E-01	-2.98E-01	-2.54E-01	-7.73E-02

Light control (Efficient)	-6.18E+01	-6.11E+01	-6.14E+01	-6.15E+01	-5.91E+01	-5.95E+01	-6.25E+01
Ventilation control (mech)	2.52E-13	6.74E-14	5.55E-04	-2.82E-04	-8.72E-04	-1.67E-04	-8.81E-14
Thermal mass (low)	-1.51E-13	6.32E-02	6.96E-02	1.84E-01	4.14E-01	3.34E-01	1.54E-01
Roof type (cool)	-3.19E-14	1.24E-01	1.62E-01	1.51E-01	5.49E-01	6.29E-01	2.11E-01
Roof type (green)	-1.52E-13	9.23E-02	1.23E-01	1.67E-01	4.05E-01	4.55E-01	1.49E-01
Room depth	1.14E+00	1.07E+00	8.82E-01	9.36E-01	5.58E-01	4.10E-01	5.74E-01
Window area %	-7.29E-02	-6.83E-02	-5.12E-02	-5.43E-02	-1.68E-02	-1.88E-02	-2.62E-02
Window type (single glazed)	-6.11E-01	-2.74E-01	2.40E-01	9.37E-02	2.23E+00	2.40E+00	1.10E+00
Window type (trple glazed)	6.69E-01	5.33E-01	1.59E-01	2.20E-01	-4.61E-01	-6.69E-01	-1.96E-01
Coating type (high performance)	1.35E+00	1.38E+00	8.32E-01	1.32E+00	4.51E-01	1.93E-01	1.43E-01
Wall insulation R-value	8.84E-14	-2.70E-01	-3.12E-01	-5.50E-01	-1.48E+00	-1.17E+00	-2.84E-01

Table 36: Electric energy multistep regression coefficients - single city

Appendix C: Retrocommissioning Data

Entry	GSF	Year Constructed	HDD	CDD	EUI (kBtu/sf)	Savings %
1	270000	1998	5537	724	64.64	7.43
2	143190	2006	6151	677	63.57	18.98
3	373500	1980	6151	677	83.26	9.11
4	178000	1980	6151	677	95.42	8.96
5	111212	1914	6151	677	85.54	26.69
6	119000	1966	6151	677	75.66	11.22
7	59218	1980	6151	677	97.10	12.85
8	105402	1927	6151	677	57.84	12.27
9	143887	1980	6151	677	76.48	26.03
10	380814	2007	6151	677	67.53	4.75
11	106000	1924	6151	677	63.69	8.01
12	130360	1991	6151	677	53.12	7.1
13	106000	1950	6151	677	69.84	11.79
14	128270	1931	6151	677	77.63	21.85
15	117098	1929	6151	677	82.20	10.59
16	117540	1980	6151	677	84.64	19.86
17	109613	1922	6151	677	749.03	2.68
18	158000	1997	6151	677	64.04	20.12
19	90774	2009	6151	677	63.89	7.71
20	205625	2003	5537	724	135.91	4.75
21	110560	2007	5537	724	73.35	49.8
22	242000	1968	5537	724	104.74	6.02
23	160000	1972	5537	724	149.09	7.91
24	269765	1970	6151	677	56.31	25.33
25	218200	1998	6151	677	84.21	3.38
26	127000	1998	6151	677	108.86	6.83
27	44900	1998	6151	677	85.03	17.65
28	78210	1998	6151	677	70.22	19.63
29	178000	2010	6151	677	90.79	28.83
30	193772	1999	5537	724	65.18	3.2
31	272000	1986	5537	724	103.79	30.78
32	115085	1980	6151	677	85.30	59.41
33	540000	2005	6151	677	76.29	2.62
34	140000	2005	6151	677	118.26	13.13
35	1615492	1980	6151	677	126.08	3.35
36	265000	1980	6151	677	132.92	10.46
37	323408	1980	6151	677	73.15	4.91
38	676301	1980	6151	677	78.66	21.14
39	756200	1980	6151	677	118.77	1.68

40	529000	1980	6151	677	99.30	17.32
41	499000	1980	6151	677	256.78	0.93
42	603480	1980	5537	724	269.64	15.32
43	235000	1980	6151	677	107.31	19.38
44	320000	1980	6151	677	177.33	7.27
45	50000	1990	6087	622	529.25	27.2
46	156000	1967	3431	1689	405.83	18.6
47	99000	1946	1688	3016	211.7	33.1
48	165031	1980	1747.55	2924.61	405.83	49.5
49	130844	1970	1747.55	2924.61	211.7	47.2
50	113700	1990	1747.55	2924.61	405.83	50.1
51	114666	1990	1747.55	2924.61	185.35	35.9
52	257953	1978	1747.55	2924.61	185.35	30.1
53	180316	1973	1747.55	2924.61	405.83	38
54	110272	1980	1747.55	2924.61	211.7	36.5
55	177838	1960	1747.55	2924.61	221.1	57.3
56	192001	1992	1747.55	2924.61	185.35	23.9
57	205000	1966	3431	1689	185.35	28.6
58	118000	1970	3431	1689	185.35	11.4
59	129000	1970	3431	1689	185.35	26.2
60	48000	1960	3431	1689	185.35	7.3
61	48000	1968	3431	1689	185.35	15.9
62	370000	1997	5765	1047	221.1	28.7
63	30244	1991	4156	781	193.1	6
64	45372	1986	2707	1470	193.1	9.6
65	125000	1983	1896	980	211.7	10.7
66	230400	1985	6152	828	529.25	38.2
67	805000	1997	5538	489	211.7	10.3
68	261000	1970	4785	538	211.7	17.6
69	489700	1997	4785	538	211.7	11.1
70	275200	1980	6207	964	159.4	6
71	185500	1994	4785	538	193.1	7.7
72	23210	1980	6226	1040	180	29.5
73	1014133	1999	3319	273	211.7	4.6
74	371343	1980	3319	273	159	5
75	317000	1911	3319	273	181.65	1.9
76	750000	1993	3319	273	186.55	2.3
77	226383	1980	3319	273	159	15.8
78	210406	1980	3319	273	221.1	2
79	467685	1985	3319	273	211.7	1.5
80	150000	1980	3205	1383	211.7	4.5
81	383200	1984	2749	1237	211.7	21.2
82	94000	1997	2749	1237	529.25	28.6

83	300000	1996	2749	1237	497.17	-7.4
84	400000	1991	2749	1237	211.7	5.1
85	324000	1990	2749	1237	211.7	7.1
86	352000	1995	2749	1237	211.7	12
87	308360	1965	2749	1237	367.2	7.7
88	600000	1980	4785	538	438.8	13.5
89	11232	1980	7981	682	172.6	10
90	32800	1980	4611	167	211.7	31

Table 37: Retrocommissioning savings algorithm data set entries

Average energy savings = 16.75%

Works Cited

[1] "Buildings Energy Data Book - Chapter 1: Buildings Sector." U.S. Department of Energy [Online]. http://buildingsdatabook.eren.doe.gov/ChapterIntro1.aspx

[2] "Construction Spending." United States Census Bureau [Online]. http://www.census.gov/construction/c30/c30index.html

[3] T. Roberts, "Energy Modeling: Early and Often." BuildingGreen.com Environmental Building News, 01-Mar-2013 [Online]. http://www.buildinggreen.com/auth/article.cfm/2013/3/1/Energy-Modeling-Early-and-Often/

[4] "Building Energy Software Tools Directory - Whole Building Analysis: Energy Simulation." US Department of Energy [Online]. http://apps1.eere.energy.gov/buildings/tools_directory/subjects.cfm/pagename=subjects/pagename_menu=whole_building_analysis/pagename_submenu=energy_simulation

[5] K. Tupper and C. Fluhrer, "Energy Modeling at Each Design Phase: Strategies to Minimize Design Energy Use." Rocky Mountain Institute, 2010 [Online]. http://www.rmi.org/Knowledge-Center/Library/2010-27_EnergyModelingDesignPhase

[6] J. Higgins, "Modeling and Design," *ASHRAE Journal*, p. 98, Aug-2013.

[7] R. Seidl, "Introduction to ASHRE Standard 202 - Commissioning Process for Buildings and Systems," San Francisco, CA, 18-Jul-2013 [Online]. http://www.cacx.org/meetings/meetings/2013-07-18/Standard_202-Seidl.pdf

[8] "2014 NCBC," *Building Commissioning Association*. [Online]. http://www.bcxa.org/training/ncbc/

[9] J. Penny, "Retro-commissioning and Existing Building Commissioning," *Buildings*, 20-Mar-2012. [Online]. http://www.buildings.com/article-details/articleid/13743/title/retro-commissioning-and-existing-building-commissioning.aspx

[10] T. Haasl and K. Heinnemeier, "California Commissioning Guide: Existing Buildings." California Commissioning Collaborative, 2006 [Online]. http://www.cacx.org/resources/documents/CA_Commissioning_Guide_Existing.pdf

[11] "A Practical Guide for Commissioning Existing Buildings," *Oak Ridge National Laboratory*. [Online]. http://eber.ed.ornl.gov/CommercialProducts/RetroCx.htm

[12] Portland Energy Conservation, Inc., "A Retrocommissioning Guide for Building Owners." U.S. Environmental Protection Agency, 2007 [Online]. http://www.peci.org/sites/default/files/epaguide_0.pdf

[13] "What is commissioning?," *Building Commissioning: A Golden Opportunity for Reducing Energy Costs and Greenhouse-Gas Emissions*. [Online]. http://cx.lbl.gov/definition.html

[14] "LL87: Energy Audits & Retro-commissioning," *PlaNYC Green Buildings & Energy Efficiency*. [Online]. http://www.nyc.gov/html/gbee/html/plan/ll87.shtml

[15] "LL84: Benchmarking," *PlaNYC Green Buildings & Energy Efficiency*. [Online]. http://www.nyc.gov/html/gbee/html/plan/ll84.shtml

[16] "Existing Commercial Buildings Energy Performance Ordinance: Overview." SF Environment Department of the City & County of San Francisco [Online]. http://www.sfenvironment.org/sites/default/files/fliers/files/sfe_gb_ecb_ordinance_overview.pdf

[17] "AEE Certifications Recognized by NYC and San Francisco," *Association of Energy Engineers*. [Online]. http://www.aeecenter.org/i4a/pages/index.cfm?pageID=4221

[18] "Reto-Commissioning for State and Local Government," U.S. Department of Energy, May 2012 [Online]. http://www1.eere.energy.gov/seeaction/pdfs/commercialbuildings_factsheet_retrocommissioning_stateandlocal.pdf

[19] "RCx – Improving Efficiencies in Building Management System Operation," *Connecticut Light & Power.* [Online]. http://www.cl-p.com/Business/SaveEnergy/Services/Retro_Commissioning/
[20] "Programs & Incentives - Retro-Commissioning," *ComEd.* [Online]. https://www.comed.com/business-savings/programs-incentives/Pages/retro-commissioning.aspx
[21] "Retrocommissioning (RCx) Program," *PG&E.* [Online]. http://www.pge.com/en/mybusiness/save/rebates/retrocommissioning/index.page
[22] "Rates," *PG&E.* [Online]. http://www.pge.com/tariffs/rateinfo.shtml
[23] B. Urban, "The MIT Design Advisor: Simple and Rapid Energy Simulation of Early-Stage Building Designs," Massachusetts Institute of Technology, Cambridge, MA, 2007.
[24] "Building Performance Solutions," *KGS Buildings.* [Online]. http://www.kgsbuildings.com/solutions.aspx
[25] "Ongoing Building Commissioning," *Cimetrics.* [Online]. http://www.cimetrics.com/index.php/infometrics-process.html
[26] "Savings Persist with Monitoring-Based Commissioning," California Energy Commission's Public Interest Energy Research Program, California, Sep. 2008 [Online]. http://www.esource.com//system/files/files/2008-08/CEC-TB-39_MBCx.pdf
[27] "Existing Building Commissioning Toolkit," *California Commissioning Collaborative.* [Online]. http://www.cacx.org/resources/rcxtools/spreadsheet_tools.html#energy_savings_calculation_tools
[28] F. F.-H. Nah, "A Study on Tolerable Waiting Time: How Long Are Web Users Willing to Wait?," *Behav. Inf. Technol.*, vol. 23, no. 3, pp. 153–163, 2004.
[29] S. Lohr, "Impatient Web Users Flee Slow-Loading Sites," *The New York Times*, 29-Feb-2012 [Online]. http://www.nytimes.com/2012/03/01/technology/impatient-web-users-flee-slow-loading-sites.html
[30] Y. Sanyasiraju, "Parabolic partial differential equations," *Partial Diffrential Equations.* [Online]. http://mat.iitm.ac.in/home/sryedida/public_html/caimna/pde/forth/forth.html
[31] A. Geletu, "Lecture 3: Introduction to Numerical Methods for Differential and Differential Algebraic Equations," Ilmenau University of Technology, 2011 [Online]. http://www.tu-ilmenau.de/fileadmin/media/simulation/Lehre/div/Lec_Slides3.pdf
[32] R. Sureshkumar, "Forward and Backward Euler Methods," *10.001: Numerical Solution of Ordinary Differential Equations.* [Online]. http://web.mit.edu/10.001/Web/Course_Notes/Differential_Equations_Notes/node3.html
[33] G. Recktenwald, "Crank Nicolson Solution to the Heat Equation," Portland State University [Online]. http://web.cecs.pdx.edu/~gerry/class/ME448/notes/pdf/CN_slides.pdf
[34] "CoolVent," *CoolVent / The Natural Ventilation Tool by MIT.* [Online]. http://coolvent.mit.edu/
[35] M. Lehar and L. Glicksman, "Rapid Algorithm for Modeling Daylight Distributions in Office Buildings," *Build. Environ.*, vol. 42, no. 8, pp. 2908–2919, 2007 [Online]. http://www.sciencedirect.com/science/article/pii/S0360132305000673
[36] "Data Elements in the TMY2 Format," *User's Manual for TMY2s.* [Online]. http://rredc.nrel.gov/solar/pubs/tmy2/tab3-2.html
[37] V. Badescu and E. Zamfir, "Degree-days, degree-hours and ambient temperature bin data from monthly-average temperatures (Romania)," *Energy Convers. Manag.*, vol. 40, no. 8, pp. 885–900, May 1999 [Online]. http://www.sciencedirect.com/science/article/pii/S0196890498001484
[38] A. Bolattürk, "Optimum insulation thicknesses for building walls with respect to cooling and heating degree-hours in the warmest zone of Turkey," *Build. Environ.*, vol. 43, no. 6, pp. 1055–1064, Jun. 2008 [Online]. http://www.sciencedirect.com/science/article/pii/S036013230700090X

[39] K. Papakostas and N. Kyriakis, "Heating and cooling degree-hours for Athens and Thessaloniki, Greece," *Renew. Energy*, vol. 30, no. 12, pp. 1873–1880, Oct. 2005 [Online]. http://www.sciencedirect.com/science/article/pii/S0960148105000091

[40] S. Nonhebel, "The effects of use of average instead of daily weather data in crop growth simulation models," *Agric. Syst.*, vol. 44, no. 4, pp. 377–396, 1994 [Online]. http://www.sciencedirect.com/science/article/pii/0308521X9490194K

[41] "The Pacific Energy Center's Guide to: California Climate Zones and Bioclimatic Design." Pacific Energy Center, Oct-2006 [Online]. http://www.pge.com/includes/docs/pdfs/about/edusafety/training/pec/toolbox/arch/climate/california_climate_zones_01-16.pdf

[42] A. O'Hagan and T. Leonard, "Bayes Estimation Subject to Uncertainty About Parameter Constraints," *Biometrika*, vol. 63, no. 1, pp. 201–203, Apr. 1976 [Online]. http://www.jstor.org/stable/2335105

[43] A. Azzalini and A. D. Valle, "The multivariate skew-normal distribution," *Biometrika*, vol. 83, no. 4, pp. 715–726, Dec. 1996 [Online]. http://biomet.oxfordjournals.org/content/83/4/715

[44] A. Azzalini, "A very brief introduction to the skew-normal distribution." Università di Padova [Online]. http://azzalini.stat.unipd.it/SN/Intro/intro.html

[45] M. Abramowitz and I. Stegun, Eds., *Handbook of Mathematical Functions with Formulas, Graphis, and Mathematical Tables*. Washington, D.C.: National Bureau of Standards, 1972 [Online]. http://people.maths.ox.ac.uk/~macdonald/aands/abramowitz_and_stegun.pdf

[46] Klein and Hoffman, Inc., "Condition Assessment Report for One Montgomery Plaza," EEB Hub Program - Pennsylvania State University, Philadelphia, PA [Online]. https://gpicspoint.ecs.psu.edu/gpic/Shared%20Documents/BUILDING%20DATA%20-AER%20DEMONSTRATION%20PROJECTS/look%20here%20--%20Current-Active%20AER%20Projects/1%20Montgomery%20Plaza%20(1MP%20or%20OMP)/Site%20Assessment/1MP%20Condition%20Report%20-%20App%20ABFGHI.pdf

[47] "Executive Report - One Montgomery Plaza," EEB Hub Program - Pennsylvania State University [Online]. https://gpicspoint.ecs.psu.edu/gpic/Shared%20Documents/Task%202%20-%20Modeling%20and%20Simulation/Shared%20Files/One%20Montgomery%20Plaza/Site%20Assessment/Executive%20Summary%20txt.pdf

[48] "Building 101." Philadelphia Industrial Development Corporation [Online]. http://www.navyyard.org/uploads/files/buildings/rBuilding101.pdf

[49] "Navy Yard Building 101," *EEB Hub Demonstration and Deployment Team*. [Online]. http://www.eebhub.org/projects-list/navy-yard-building-101

[50] "The Navy Yard Campus Map," Philadelphia Industrial Development Corporation, Philadelphia, PA, 30-Sep-2013 [Online]. http://www.navyyard.org/userfiles/file/The_Navy_Yard_Campus_Map_11x17_9-30-13.pdf

[51] Camroden Associates, Inc., "Report on the Building Envelope Air Tightness and Pressure Relationship Mapping," Philadelphia Industrial Development Corporation, Westmoreland, NY, 2012 [Online]. https://gpicspoint.ecs.psu.edu/gpic/Shared%20Documents/BUILDING%20DATA%20-AER%20DEMONSTRATION%20PROJECTS/look%20here%20--%20Current-Active%20AER%20Projects/Building%20101/Site%20assessment/Bldg%20101%20Blower%20Door%20final%20report%203-21-12.pdf

[52] Dome-Tech, Inc., "Energy Audit Report - Philadelphia Industrial Development Corporation - Philadelphia Navy Yard, Building 101," Philadelphia Industrial Development Corporation, Edison, NJ, 2011 [Online]. https://gpicspoint.ecs.psu.edu/gpic/Shared%20Documents/BUILDING%20DATA%20-

AER%20DEMONSTRATION%20PROJECTS/look%20here%20--%20Current-Active%20AER%20Projects/Building%20101/Site%20assessment/Building%20101%20Level%202%20Energy%20Audit%20Final%20Report%2012_23_11.pdf
[53] "Buildings Owned by MIT by Year Construction Began," *MIT Department of Facilities Space Accounting*. [Online]. https://floorplans.mit.edu/cgi-bin-db-any/wdbanyscript.asp?Report=bycb&Item=MIT
[54] O. R. Simha, "North East Sector Plan," in *MIT Campus Planning, 1960-2000: An Annotated Chronology*, Cambridge, MA: Massachusetts Institute of Technology, Office of the Executive Vice President, 2001, p. 126.
[55] "Building E40," *MIT Department of Facilities Space Accounting*. [Online]. https://floorplans.mit.edu/cgi-bin-db-mit/wdbmitscript.asp?report=bs&item=E40
[56] "MIT CIS: Directions," *MIT Center for International Studies*. [Online]. http://web.mit.edu/cis/directions.html
[57] C. A. Brown, "Toward zero net energy buildings : optimized for energy use and cost," Thesis, Massachusetts Institute of Technology, 2012 [Online]. http://dspace.mit.edu/handle/1721.1/77776
[58] M. Mitchell, "Genetic Algorithms: An Overview," in *An Introduction to Genetic Algorithms*, MIT Press, 1998 [Online]. http://books.google.com/books?hl=en&lr=&id=0eznlz0TF-IC&oi=fnd&pg=PP9&dq=history+of+genetic+algorithms&ots=sfrF8055Qd&sig=qLe5tJyWH8C5K_doAeR3lCe4iTs#v=onepage
[59] C. Dyer, "Genetic Algorithm," University of Wisconsin Madison [Online]. http://pages.cs.wisc.edu/~dyer/cs540/notes/05_geneticAlgorithms.pdf
[60] H. Maaranen, K. Miettinen, and M. M. Mäkelä, "Quasi-random initial population for genetic algorithms," *Comput. Math. Appl.*, vol. 47, no. 12, pp. 1885–1895, Jun. 2004 [Online]. http://www.sciencedirect.com/science/article/pii/S0898122104840240
[61] P. Diaz-Gomez and D. Hougen, "Initial Population for Genetic Algorithms: A Metric Approach." University of Oklahoma School of Computer Science [Online]. www.researchgate.net/publication/220862320_Initial_Population_for_Genetic_Algorithms_A_Metric_Approach/file/79e4150e5dd55b727a.pdf
[62] C. Seals, "Genetic Algorithms," Auburn University [Online]. www.eng.auburn.edu/~sealscd/GAs.ppt
[63] P. Lanzi, "Lecture 2: Simple Genetic Algorithms," Politecnico di Milano [Online]. http://www.pierlucalanzi.net/wp-content/teaching/ga/GA2011-02-SimpleGeneticAlgorithm.pdf
[64] "Genetic Server/Library: Crossover," *NeuroDimension, Inc.*, 2002. [Online]. http://www.nd.com/products/genetic/crossover.htm
[65] S. N. Sivanandam and S. N. Deepa, "3.11 Search Termination (Convergence Criteria)," in *Introduction to Genetic Algorithms*, Springer, 2007, pp. 59–60 [Online]. http://books.google.com/books?id=wonrLjj2GagC&pg=PA59&lpg=PA59
[66] "Genetic Server/Library: Termination," *NeuroDimension, Inc.*, 2002. [Online]. http://www.nd.com/products/genetic/termination.htm
[67] "Custom Degree Day Data," *Heating & Cooing Degree Days*. [Online]. http://www.degreedays.net
[68] M. Kalisch, "Measuring distances - Applied multivariate statistics." Swiss Federal Institute of Technology Zurich, Spring-2012 [Online]. https://stat.ethz.ch/education/semesters/ss2012/ams/slides/v4.2.pdf
[69] E. Mills, N. Bourassa, M. A. Piette, H. Friedman, T. Haasl, T. Powell, and D. Claridge, "The Cost-Effectiveness of Commissioning New and Existing Commercial Buildings: Lessons from 224 Buildings," presented at the National Conference on Building Commissioning, New York, NY, 2005 [Online]. http://evanmills.lbl.gov/pubs/pdf/ncbc_mills_6apr05.pdf

[70] E. Mills, H. Friedman, T. Powell, N. Bourassa, D. Claridge, T. Haasl, and M. A. Piette, "The Cost-Effectiveness of Commercial-Buildings Commissioning: A Meta-Analysis of Energy and Non-Energy Impacts in Existing Buildings and New Construction in the United States." Lawrence Berkeley National Laboratory, 15-Dec-2004 [Online]. http://eetd.lbl.gov/emills/PUBS/Cx-Costs-Benefits.html
[71] "Strategic Building Systems - CL&P Energy Savings Initiatives Funded by CEEF." Connecticut Light & Power.
[72] J. Null, "Comparative Climatic Data - United States," *Golden Gate Weather Services*. [Online]. http://ggweather.com/ccd/
[73] Cornell University, "US Comparative Climate Data," *Northeast Regional Climate Center*. [Online]. http://www.nrcc.cornell.edu/page_ccd.html
[74] Desert Research Institute, "Comparative Data for the Western States," *Western Regional Climate Center*. [Online]. http://www.wrcc.dri.edu/climatedata/comparative/
[75] University of Tennessee, Knoxville, "Graphic and Tabular Climate Data," *Teaching Architecture + Energy*. [Online]. http://web.utk.edu/~archinfo/EcoDesign/escurriculum/
[76] M. Greenacre, "Chapter 5: Measures of distance between samples: non-Euclidean." Universitat Pompeu Fabra Barcelona Department of Economics and Business [Online]. http://www.econ.upf.edu/~michael/stanford/maeb5.pdf
[77] C. Manning, P. Raghavan, and H. Schütze, "Dot products," in *Introduction to Information Retrieval*, Cambridge, England: Cambridge University Press, 2008, pp. 120–123 [Online]. http://nlp.stanford.edu/IR-book/html/htmledition/dot-products-1.html
[78] C. Monz, "Machine Learning for Data Mining Week 4: kNN and Naive Bayes Classification." Queen Mary University of London School of Electronic Engineering and Computer Science [Online]. http://www.eecs.qmul.ac.uk/~christof/html/courses/ml4dm/week04-knn-naivebayes-4pp.pdf
[79] S. Sayad, "K Nearest Neighbors - Classification." [Online]. http://www.saedsayad.com/k_nearest_neighbors.htm
[80] S. Deokar, "Weighted K Nearest Neighbor," University of Minnesota Duluth, 20-Apr-2009 [Online]. http://www.d.umn.edu/~deoka001/downloads/K_Nearest_Neighbor_Algorithm.pdf
[81] B. Urban and L. Glicksman, "A Rapid Building Energy Model and Interface for Non-technical Users," presented at the Buildings X Conference, Clearwater, FL, 2007.
[82] J. Higgins, "Modeled vs. Real Utility Data," *ASHRAE J.*, vol. 56, no. 4, p. 93, Apr. 2014.