Application of Machine Learning to Building Control

Jamieson Schulte and Sebastian Thrun

Problem: The design of buildings with sophisticated sensing and control systems is a growing trend that promises increased energy efficiency and personal comfort. However, the performance of such buildings is often hindered by reactive control systems based on simple thresholding and scheduling, and cannot effectively coordinate sequences of actions by multiple controllers. Furthermore, the failure or blockage of a single sensor can result in a completely broken control scheme, even if sufficient sensor information for adequate control is present.

Modern building controllers also generally fail to adapt to the preferences of individual users. In many cases it is arguable that this is more important than meeting a specific local lighting or heating setpoint, which is the goal of most existing systems. For tasks such as heating and cooling, information exists that can be used to enhance performance and efficiency by predicting future states (outdoor temperature and occupancy history, for example), but is not extensively used.

Impact: We believe that our work can lead to advances in the state of the art in both building control and machine learning. Given that heating and lighting make up a considerable portion of worldwide energy expense, efficient control algorithms deserve careful study. Any broadly-applicable improvement in building control has the potential to reduce global energy consumption. Any impact on human comfort in the workplace could significantly effect the productivity and well-being of workers in an institution.

From a learning perspective, buildings generally consist of multiple rooms with similar (but rarely identical) characteristics, which provide an opportunity to apply knowledge transfer between machine learning tasks. While a basic building control algorithm may appear obvious in some situations, a statistical analysis of the problem performed by a learner can point to significant improvements. For this reason, we propose the integration of programming and learning to achieve a near-optimal algorithm with minimal exploration by the learner. This approach can assist many applications of machine learning for which an initial human bias can improve the rate of learning, and quality and generality of the resulting solution.

State of the Art: Current technology for building control is exemplified by the Intelligent Workplace, a project of the Center for Building Performance and Diagnostics at Carnegie Mellon University. This structure is equipped with many sensors and actuators, as well as a bus system for access to both. However, only a small fraction of the collected data is actually used for control, and the range of possible actions is limited due to the use of algorithms based on input thresholds and scheduling. Because all parameter tuning and scheduling is performed by hand, the system is not particularly robust to changes in the building environment that alter the mapping of inputs to outputs. The sensing and control infrastructure of the Intelligent Workplace provides an excellent resource for data collection and experimentation, and is the focus of our current research.

Approach: Our approach to the problem consists of two parts. First, we use non-intrusive sensors to infer important information about the building environment. This is an essential aspect of the building domain, since quantities such as the illuminance profile within a room cannot be directly measured in an active office space. Second, we are working to integrate programming and learning as methods for instructing the controller. This is important in domains where learning must occur with a minimum of exploration (by using some initial bias) and the resulting policy must be readible by a human.

Much of the information useful for making building control decisions is hidden or not easily measurable.

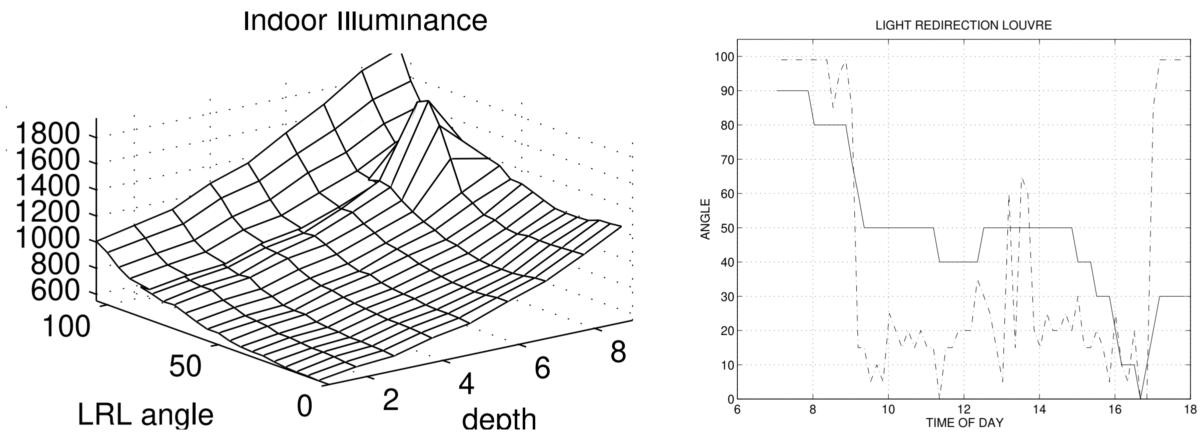

Figure 1: *Left.* The illuminance profile of an office space with different positions of a light redirection louvre (LRL) at a single point in time. *Right.* A comparison of the best LRL position for meeting an indoor lighting setpoint (dashed line) to the behavior of our reinforcement learner, which is penalized for moving more than necessary (solid line).

Among these are the illuminance profile in a space, the occupancy (and predicted future occupancy) of a room, future weather patterns. Each represents a distinct problem that we approach by incorporating data that is not traditionally used by building controllers. The output is a set of features that augments those provided by the sensors already incorporated into the Intelligent Workplace.

Given these features, our approach is to combine programming and learning in a learner that can be biased by an initial control program, and that outputs a policy in the form of such a program. This method allows a human programmer to bias learning when a sub-optimal solution to the problem is already known. This starting policy initializes the value function of a reinforcement learner such that exploration is guided away from states that are not expected to require policy refinement. The output value function of the learner is then converted into a program that is ideally both simple and similar to the initially provided algorithm. The integration of programming and learning addresses the fundamental problem of knowledge transfer in learning systems. We have so far implemented this approach with a policy encoded as a set of rules rather than a sequentially executed program and have achieved promising improvements over an existing lighting controller.

Future Work: Our work has not yet covered the on-line learning of human environmental preferences, which will be a necessary step in the automation of the control system. We have yet to extend our experiments to the control of temperature, and to examine the interrelationship this will inevitably have with lighting. The work on integrating programming and learning is still preliminary, and a suitable programming language that is both expressive enough to represent a range of policies and sufficiently learnable is still being sought.

Selected References:

[1] S. Thrun and J. Schulte. Reinforcement learning with policy constraints. 1998. internal report.