

The Neural Network House: An Environment that Adapts to its Inhabitants

Michael C. Mozer

Department of Computer Science and
Institute of Cognitive Science
University of Colorado
Boulder, CO 80309-0430
mozer@colorado.edu

From: AAAI Technical Report SS-98-02. Compilation copyright © 1998, AAAI (www.aaai.org). All rights reserved.

Abstract

Although the prospect of computerized homes has a long history, home automation has never become terribly popular because the benefits are seldom seen to outweigh the costs. One significant cost of an automated home is that someone has to program it to behave appropriately. Typical inhabitants do not want to program simple devices such as VCRs, let alone a much broader range of electronic devices, appliances, and comfort systems that have even greater functionality. We describe an alternative approach in which the goal is for the home to essentially *program itself* by observing the lifestyle and desires of the inhabitants, and learning to anticipate and accommodate their needs. The system we have developed controls basic residential comfort systems—air heating, lighting, ventilation, and water heating. We have constructed a prototype system in an actual residence, and describe initial results and the current state of the project.

Introduction

Since the mid 1940s, the home automation industry has promised to revolutionize our living environments. The so-called “smart home” has been hyped in the popular press. The vision of the industry is that household devices—appliances, entertainment centers, utilities, thermostats, lights, etc.—will be endowed with microprocessors that allow the devices to communicate with one another and thereby behave intelligently. The dishwasher can ask the hot water heater whether it has sufficient capacity to operate; inhabitants can telephone home and remotely instruct the VCR to record a favorite show; the TV might lower its volume when the phone rings; or the clothes dryer might make an announcement over an intercom system when it has completed its cycle.

As attractive as this scenario is, the software required to achieve the intelligence is highly complex and unwieldy, and worse, the software must be tailored to a particular home and family, and updated as the family’s lifestyle changes. Tackling the programming task is far beyond the capabilities and interest of typical home inhabitants.

Indeed, even rudimentary forms of regulation, such as operating a set back thermostat, which allows different temperature settings depending on the time of day, are inordinately difficult for people (Gregorek, 1991). The alternative of hiring professional technicians to update programs as necessary is used in some commercial systems, but is costly and inconvenient. Partly due to these difficulties in programming, home automation has never become a widely available and accepted technology.

In contrast to standard computerized homes that can be programmed to perform various functions, the crux of our project is to develop a home that essentially *programs itself* by observing the lifestyle and desires of the inhabitants, and learning to anticipate and accommodate their needs. The system we have developed controls basic residential comfort systems—air heating, lighting, ventilation, and water heating.

ACHE

We call the system ACHE, which stands for adaptive control of home environments. ACHE monitors the environment, observes the actions taken by occupants (e.g., adjusting the thermostat; turning on a particular configuration of lights), and attempts to infer patterns in the environment that predict these actions.

ACHE has two objectives. One is anticipation of inhabitants’ needs. Lighting, air temperature, and ventilation should be maintained to the inhabitants’ comfort; hot water should be available on demand. When inhabitants manually adjust environmental setpoints, it is an indication that their needs have not been satisfied and will serve as a training signal for ACHE. If ACHE can learn to anticipate needs, manual control of the environment will be avoided. The second objective of ACHE is energy conservation. Lights should be set to the minimum intensity required; hot water should be maintained at the minimum temperature needed to satisfy the demand; only rooms that are likely to be occu-




Figure 1. The Neural Network House, circa 1926.

pied in the near future should be heated; when several options exist to heat a room (e.g., furnace, ceiling fans forcing hot air down, opening blinds to admit sunlight), the alternative minimizing expected energy consumption should be selected.

Achieving either one of these objectives in isolation is fairly straightforward. If ACHE were concerned only with appeasing the inhabitants, the air temperature could be maintained at a comfortable 70° at all times. If ACHE were concerned only with energy conservation, all devices could be turned off. ACHE's challenge is to achieve both objectives simultaneously. This requires the ability to anticipate inhabitant activities, occupancy patterns, and tolerances.

Optimal Control

In what sort of framework can the two objectives—appeasing the inhabitants and conserving energy—be integrated? Supervised learning will not do: If a temperature setpoint chosen by the inhabitant serve as the target for a supervised learning system, energy costs will not be considered. Instead, we have adopted an *optimal control* framework in which failing to satisfy each objective has an associated cost. A *discomfort cost* is incurred if inhabitant preferences are not met, i.e., if the inhabitant is not happy with the settings determined by ACHE, as indicated by manual control of the environment. An *energy cost* is incurred based on the use of electricity or gas resources. The *expected average cost*, $J(t_0)$, starting at time t_0 can then be expressed as

$$J(t_0) = E \left[\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{t=t_0+1}^{t_0+K} d(x_t) + e(u_t) \right]$$

where $d(x_t)$ is the discomfort cost associated with the environmental state x at time t , and $e(u_t)$ is the energy cost associated with the control decision u at time t . The goal is

to find an optimal control *policy*—a mapping from states x_t to decisions u_t —that minimizes the expected average cost.

This framework requires that discomfort and energy costs be expressed in the same currency. We have chosen dollars as this currency, which makes a characterization of energy costs straightforward. Relative discomfort is indicated by overriding the choices of ACHE, and this relative discomfort is translated to a dollar amount by means of a misery-to-dollars conversion factor. One technique we have explored for determining this factor, based on an economic analysis, depends on the loss in productivity that occurs when ACHE ignores the inhabitants' desires. Another technique adjusts the conversion factor over a several month period based on how much inhabitants are willing to pay for gas and electricity.

Implementation

We have implemented ACHE in an actual residence. The residence is a former three-room school house built in 1905 near Boulder, Colorado, originally serving children of the mining town of Marshall (Figure 1). The school was closed in 1956 and was completely renovated in 1992, at which time the infrastructure needed for the ACHE project was incorporated into the house, including nearly five miles of low-voltage conductor for collecting sensor data and a power-line communication system for controlling lighting, fans, and electric outlets. The residence is an ideal candidate for intelligent energy management because of its age, 13-25 foot ceilings, and exposed south and west faces that hold potential for passive solar heating.

ACHE is equipped with sensors that report the state of the environment. The sensory state includes the following for each room in the home:

- status of lights (on or off, and if on, intensity level)
- status of fans (speed)
- status of temperature control user interface (a fancy digital thermostat that specifies the current setpoint)




Figure 2. A floor plan of the adaptive house, including locations of sensors and actuators.

temperature for the room, and can be adjusted by the inhabitant)

- ambient illumination
- room temperature
- sound level
- motion detector activity (motion or no motion)
- status of all doors and windows (open or closed).

In addition, the system receives the following global information:

- water heater temperature
- water heater energy usage
- water heater outflow
- furnace energy usage
- outdoor temperature
- outdoor insolation (sunlight)
- gas and electricity costs
- time of day, day of week, date.

At present, ACHE has the ability to control the following actuators:

- on/off status and intensity of light banks (22 total)
- on/off status and speed of ceiling fans (6 total)
- on/off status of water heater
- on/off status of gas furnace
- on/off status of electric space heaters (2 total)
- on/off status of speakers in each room through which computer can communicate (12 total)

Figure 2 shows a floor plan of the residence, as well as the approximate location of selected sensors and actuators.

ACHE Architecture

Adaptive control of building energy systems is difficult. We have incomplete models of the environment and controlled devices. The environment, including the behavior of the inhabitants, is nonstationary and stochastic. Controlled devices are nonlinear. Multiple interacting devices must be controlled simultaneously. Under such circumstances, traditional techniques from control theory and artificial intelligence have great difficulty (Dean & Wellman, 1991).

The basic system architecture of ACHE is presented in Figure 3. This architecture is replicated for each control domain—lighting, air heating, water heating, and ventilation. The instantaneous environmental state is fed through a *state transformation* that computes statistics such as averages, minima, maxima, and variances in a given temporal window. The result is a state representation that provides more information about the environment than the instantaneous values. The instantaneous state is also given to an *occupancy model* that determines for each *zone* of the house—usually corresponding to a room—whether or not it is occupied. The occupancy model relies on motion detector signals, but it includes rules that say, essentially, “a zone remains occupied, even when there is no motion, unless there is motion in an adjacent zone that was previous unoccupied.” Consequently, the occupancy model maintains occupancy status even when there is no motion.

The three adaptive components of ACHE are shown in the top of Figure 3. Various *predictors* attempt to take the current state and forecast future states. Examples of predictions include: expected occupancy patterns in the house over the next few hours, expected hot water usage, likeli-

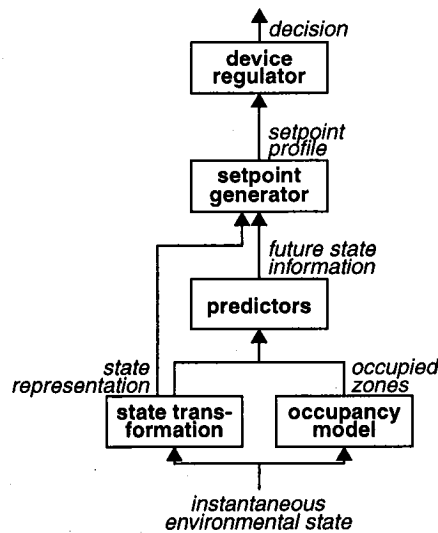


Figure 3. System architecture of ACHE

hood that a zone will be entered in the next few seconds. The predictors are implemented as feedforward neural networks trained with back propagation, or as a combination of a neural net and a look up table.

Given the predictions of future states, control decisions need to be made concerning the energy devices in the home. The decision making process is split into two stages. The *setpoint generator* determines a setpoint profile specifying the target value of some environmental variable (lighting level, air temperature, water temperature, etc.) over a window of time. The *device regulator* controls physical devices to achieve the setpoint. The device regulator may have many alternative devices at its disposal. It must determine which one or which subset to use.

The reason for dividing control between the setpoint generator and device regulator is to *encapsulate knowledge*. The setpoint generator requires knowledge about inhabitant preferences, while the device regulator has knowledge about the physical layout and characteristics of the environment and controlled devices. If the inhabitants or their preferences change over time, only the setpoint generator need relearn.

The setpoint generator and device regulator in each domain are based on one of two approaches to control: indirect control using dynamic programming and models of the environment and inhabitant, or direct control using reinforcement learning. For example, the device regulator for indoor air temperature uses a predictive model of the indoor air temperature, as a function of the current indoor temperature, outdoor temperature, and the states of the furnace and electric space heaters. This model is based on a simple RC thermal model of the house and furnace, with a neural network that learns deviations from this simple model and the

actual behavior of the house. Given this model, achieving a particular setpoint temperature involves little more than exhaustively searching through the space of heating device actions and finding an action sequence that achieves the setpoint. In contrast to this indirect approach, the setpoint generator for the lighting controller uses a direct approach with reinforcement learning because it would be difficult to learn an explicit model of inhabitant preferences.

Current Implementation Status

We have conducted simulation studies of the heating control system (Mozer, Vidmar, & Dodier, 1997), using actual occupancy data and outdoor temperature profiles, evaluating various control policies. ACHE robustly outperforms three alternative policies, showing a lower total (discomfort plus energy) cost across a range of values for the relative cost of inhabitant discomfort and the degree of nondeterminism in occupancy patterns.

We have also implemented and tested a lighting controller in the house (Mozer & Miller, in press). To give the flavor of its operation, we describe a sample scenario of its behavior. The first time that the inhabitant enters a zone (we'll refer to this as a *trial*), ACHE decides to leave the light off, based on the initialization assumption that the inhabitant has no preference with regard to light settings. If the inhabitant overrides this decision by turning on the light, ACHE immediately learns that leaving the light off will incur a higher cost (the discomfort cost) than turning on the light to some intensity (the energy cost). On the next trial, ACHE decides to turn on the light, but has no reason to believe that one intensity setting will be preferred over

another. Consequently, the lowest intensity setting is selected. On any trial in which the inhabitant adjusts the light intensity upward, the decision chosen by ACHE will incur a discomfort cost, and on the following trial, a higher intensity will be selected. Training thus requires just three or four trials, and explores the space of decisions to find the lowest acceptable intensity. ACHE also attempts to conserve energy by occasionally "testing" the inhabitant, selecting an intensity setting lower than the setting believed to be optimal. If the inhabitant does not complain, the cost of the decision is updated to reflect this fact, and eventually the lower setting will be evaluated as optimal.

Evaluating ACHE

It is our conviction that intelligent control techniques for complex systems in dynamic environments must be developed and evaluated in naturalistic settings such as the Neural Network House. While there are numerous examples illustrating the potential of neural nets for control of building energy systems (e.g., Curtiss, Kreider, & Brandemuehl, 1994; Miller & Seem, 1991; Seem & Braun, 1991; Scott, Shavlik, & Ray, 1992), this research focuses on narrowly defined problems and is generally confined to computer simulations. The research that does involve control of actual equipment makes simplifying assumptions about operating conditions and the environment. We intend to show that adaptive control will yield benefits in natural environments under realistic operating conditions.

The research program hinges on a careful evaluation phase. In the long term, the primary empirical question we must answer is whether there are sufficiently robust regularities in the inhabitants' behavior that ACHE can benefit from them. On first consideration, most people conclude that their daily schedules are not "regular"; they sometimes come home at 5 p.m., sometimes at 6 p.m., sometimes not until 8 p.m. However, even subtle statistical patterns in behavior—such as the fact that if one is not home at 3 a.m., one is unlikely to be home at 4 a.m.—are useful to ACHE. These are patterns that people are not likely to consider when they discuss the irregularities of their daily lives. These patterns are certainly present, and we believe that they can be usefully exploited in adaptive control of living environments.

Acknowledgements

We are grateful to Marc Anderson and Robert Dodier, who helped develop the software infrastructure for the Neural Network House. The Neural Network House is supported by the Sensory Home Automation Research Project (SHARP) of Sensory Inc., as well as a CRCW grant-in-aid from the University of Colorado, NSF award IRI-9058450, McDonnell-Pew award 97-18.

References

- Curtiss, P. S., Kreider, J. F., & Brandemuehl, M. J. (1994). Local and global control of commercial building HVAC systems using artificial neural networks. *Proceedings of the 1994 American Control Conference ACC'94*. New York: IEEE Press.
- Dean, T. L., & Wellman, M. P. (1991). *Planning and control*. San Mateo, CA: Morgan Kaufmann.
- Gregorek, T. (1991). The energy revolution. *Electronic House*, 6, 10–15.
- Miller, R. C., & Seem, J. E. (1991). Comparison of artificial neural networks with traditional methods of predicting return time from night or weekend setback. *ASHRAE Transactions*, 97.
- Mozer, M. C., Dodier, R., & Vidmar, S. (1997). The Neurothermostat: Adaptive control of residential heating systems. In M. C. Mozer, M. I. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems 9* (pp. 953–959). Cambridge, MA: MIT Press.
- Mozer, M. C., & Miller, D. J. (in press). Parsing the stream of time: The value of event-based segmentation in a complex, real-world control problem. In M. Gori (Ed.), *Adaptive processing of temporal information*. Springer Verlag.
- Scott, G. M., Shavlik, J. W., & Ray, W. H. (1992). Refining PID controllers using neural nets. In J. E. Moody, S. J. Hanson, & R. P. Lippmann (Eds.), *Advances in Neural Information Processing Systems 4* (pp. 555–562). San Mateo, CA: Morgan Kaufmann.
- Seem, J. E., & Braun, J. E. (1991). Adaptive methods for real-time forecasting of building electrical demand. *ASHRAE Transactions* 97.