

## The Neural Network House: An Overview

Michael C. Mozer<sup>†\*</sup>  
Robert H. Dodier<sup>†</sup>  
Marc Anderson<sup>†</sup>  
Lucky Vidmar<sup>†</sup>  
Robert F. Cruickshank III<sup>‡</sup>  
Debra Miller<sup>†</sup>

<sup>†</sup>*Department of Computer Science*

<sup>\*</sup>*Institute of Cognitive Science*

<sup>‡</sup>*Department of Civil, Environmental, and Architectural Engineering  
University of Colorado  
Boulder, CO 80309 USA*

### Abstract

Typical home comfort systems utilize only rudimentary forms of energy management and conservation. The most sophisticated technology in common use today is an automatic setback thermostat. Tremendous potential remains for improving the efficiency of electric and gas usage. However, home residents who are ignorant of the physics of energy utilization cannot design environmental control strategies, but neither can energy management experts who are ignorant of the behavior patterns of the inhabitants. Adaptive control seems the only alternative. We have begun building an adaptive control system that can infer appropriate rules of operation for home comfort systems based on the lifestyle of the inhabitants and energy conservation goals. Recent research has demonstrated the potential of neural networks for intelligent control. We are constructing a prototype control system in an actual residence using neural network reinforcement learning and prediction techniques. The residence is equipped with sensors to provide information about environmental conditions (e.g., temperatures, ambient lighting level, sound and motion in each room) and actuators to control the gas furnace, electric space heaters, gas hot water heater, lighting, motorized blinds, ceiling fans, and dampers in the heating ducts. This paper presents an overview of the project as it now stands.

## 1 Introduction

The burgeoning home automation industry will likely have a tremendous impact on the residence of the future. One vision is that household devices—appliances, entertainment centers, utilities, temperature and lighting control units—will be endowed with microprocessors that will be able to communicate with one another. 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 stereo can switch on automatically when the inhabitants return home.

While publicity surrounding this technology focuses on the visible and easily realized operations, the greatest—and relatively unexplored—potential afforded by the technology lies behind the scenes, in the difficult tasks of energy management and conservation. Consider the problem of controlling the setpoint temperature of a water heater. Residential water heaters in the US are inefficient; they are set to a fixed temperature at all times. The heater could conserve energy if it shut itself down at times of the day when there is no demand for hot water. Of course, the demand will be different from household to household, one month to the next, on weekends versus weekdays. And the presence of house guests should be taken into account to anticipate both how much and when hot water is required. The demand could also depend on weather: Cool showers are more likely on a hot, humid morning than in the dead of winter. Patterns of utilization of appliances such as dishwashers and washing machines must be considered; these patterns of utilization in turn depend on the schedules of the inhabitants. And if time-of-use utility pricing becomes reality, it may be worthwhile to schedule water heating at times when energy is cheapest, not when a demand is likely to occur. Finally, it may be important to consider that running the water heater could influence other aspects of the home environment such as peak electricity demand or room temperature.

Optimal control of a water heater is certainly nontrivial. What is really needed is a system whose rules of operation are based on the lifestyle of the inhabitants and knowledge of energy costs. The system would be fitted with sensors to monitor the state of the environment, e.g., sound and motion detectors in the spare bedroom to determine whether house guests were visiting. These sensors allow the system to detect patterns in the environment that reliably predict demands on the device being controlled, which in turn will lead to an appropriate control strategy.

A major challenge to building such a system is programming it to behave appropriately. Technophiles may find it fun to program their personal computers to control energy and electronic devices in their home. This program would consist of lists of rules like “if the mean value of motion detector 7 over the past hour is greater than .5 and if it’s later than 10 p.m., then turn on the water heater 2 hours earlier than usual.”

Tackling this programming task is far beyond the capabilities and interest of typical home inhabitants, though. 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). Inhabitants cannot be expected to design efficient control strategies given their ignorance of technology and the physics of energy utilization, but neither can energy management experts

who are ignorant of the dynamic behavior patterns of individuals.

What seems required here is an *adaptive* system that can infer appropriate rules of operation by observing the home's inhabitants and by exploring alternative control strategies. Essentially, the system would *program itself* rather than requiring human expertise to specify the rules of operation.

## 2 ACHE

We are constructing a system that will manage all comfort aspects of an actual residence, including air and water temperature, ventilation, and lighting. We call the system ACHE, which stands for *adaptive control of home environments*.

ACHE has two objectives. One is anticipation of inhabitants' needs. Lighting, air temperature, and ventilation should then be maintained to the inhabitants' comfort; hot water should be available on demand. When inhabitants manually regulate the environment, 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 the minimum temperature needed to satisfy the demand; only rooms that are likely to be occupied 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. Note that achieving the conservation objective requires the ability to anticipate inhabitant activities and tolerances.

## 3 Hardware Implementation

ACHE is provided 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 temperature control user interface (a fancy digital thermostat that specifies the current setpoint 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

Figure 1: Floor plan of residence and approximate location of selected sensors and actuators

- 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 1 shows a floor plan of the residence, as well as the approximate location of selected sensors and actuators. The residence is a former three-room schoolhouse built in 1905 near Boulder, Colorado, originally serving

Figure 2: Hardware underlying ACHE.

children of the mining town of Marshall. The school was closed in 1956 and was completely renovated in 1990, at which time the infrastructure needed for the ACHE project was incorporated into the house, including nearly five miles of low voltage wire for collecting sensor data and a powerline 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.

Figure 2 depicts the hardware underlying ACHE and the flow of information. Data from the roughly seventy-five sensors is fed into a PC via three data acquisition boards. The PC detects changes in the sensor values and sends them, via a serial line, to a DEC 3000/600 UNIX workstation. On the workstation, a *server* is running that maintains the current state of the house, computes statistics of the state, watches for critical sensory events, and communicates with one or more *clients* that perform the actual operations of ACHE. When clients wish to control a device, they pass a message to the PC, via the server, which either communicates with the device directly (for low voltage devices) or passes a message to a dedicated microcontroller which sends out a command superimposed over the 60 Hz powerline signal, directly to the device being controlled (e.g., shut off an electric outlet, turn on a light). The powerlines are also used to communicate messages to the PC, via the microcontroller, when the inhabitant changes the state of a device (e.g., dims a light, turns off a fan).

## 4 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 primary 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, likelihood that a zone will be entered in the next few seconds. The predictors are implemented as feedforward neural networks trained with back propagation, as look up tables, 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 environ-

Figure 3: System architecture of ACHE.

ment and controlled devices. If the inhabitants or their preferences change over time, only the setpoint generator need relearn. Both the setpoint generator and device regulator are built as a combination of lookup tables and neural networks, and are trained with Q learning (Watkins & Dayan, 1992), a reinforcement learning technique.

## 5 Current Implementation Status

We are currently implementing various components of ACHE. We started with the predictors, and are now moving on to control tasks. We summarize several of the current projects here.

- *Occupancy Predictor.* This is a module that predicts expected fraction of the time in the coming 30, 60, or 90 minutes that inhabitants will be home. Its predictions are based on the current occupancy status of the house, time of day, day of week, the occupancy pattern from the past three weeks on the same day of the week at the current time, and the occupancy pattern from the past three days at the current time. This module is being expanded to predict occupancy of individual zones, as well as the whole house.
- *Zone Anticipator.* This module predicts, for every currently unoccupied zone, whether the zone will become occupied in the coming two seconds. Its predictions are based on motion patterns, sound levels, door status, and time of day. Its usage is in lighting control: If lights are turned on only when a motion detector responds, one can occasionally enter well into a zone before the lights can respond.

By anticipating zone occupancy, lights are turned on prior to a zone being entered.

- *Lighting Setpoint Generator.* This module determines the lighting setpoints in a zone. The setpoint is specified in terms of the level of light sensors (photoresistors) in the zone. Most zones contain a single sensor; the sensor reading corresponds roughly to perceived brightness. This module uses the output of the zone anticipator, output from a module that predicts how long before a zone is again occupied, the occupancy model, the position of the sun, and an outside light sensor reading. It is trained via reinforcement learning, being punished whenever the inhabitants change the settings of the lights from the automatically determined settings, and is also punished for consuming energy.
- *Lighting Device Controller.* We have implemented this module for the great room zone (see Figure 1), which has seven independently controlled banks of lights and four light sensor readings (Dodier, Lukianow, Ries, & Mozer, 1994). Each bank of lights can be set to one of 16 possible intensity levels. The task of this module is to determine the appropriate light intensity settings to achieve a particular pattern of sensor readings.

## 6 Evaluation of 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, 1992a, b; 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 neural nets will yield benefits in natural environments under realistic operating conditions.

The research program hinges on a careful evaluation phase. We are not yet at the stage where energy savings or convenience benefits can be quantified. 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.

## 7 Acknowledgements

This project owes its existence to the heroic efforts of students at the University of Colorado, particularly Josh Anderson, Dan Bertini, Matt Bronder, Brian Daugherty, Mark Fontenot, Paul Kooros, Diane Lukianow, Tom Moyer, Charles Myers, James Ries, Erik Skorpen, and Joel Sloss. Support for the project to date comes from a CRCW grant-in-aid from the University of Colorado, the Computer Science Department at the University of Colorado, NSF Presidential Young Investigator award IRI-9058450, and grant 90-21 from the James S. McDonnell Foundation.

## References

- Curtiss, P., Kreider, J., & Brandemuehl, M. (1992a). *Adaptive control of HVAC processes using predictive neural networks*. Manuscript submitted for publication.
- Curtiss, P., Kreider, J., & Brandemuehl, M. (1992b). *Energy management in central HVAC plants using neural networks*. Manuscript submitted for publication.
- Dodier, R. H., Lukianow, D., Ries, J., & Mozer, M. C. (1994). Comparison of neural net and conventional techniques for lighting control. *Applied Mathematics and Computer Science*. In press.
- 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.
- Dean, T. L., & Wellman, M. P. (1991). *Planning and control*. San Mateo, CA: Morgan Kaufmann.
- 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.
- Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8, 279-292.