Unsupervised learning and control provide ambient intelligence to smart buildings

Buildings are changing their nature from static structures of bricks and mortar to dynamic work and living environments that actively support and assist their inhabitants. These new buildings are expected to behave intelligently. In addition, to satisfying the needs of its inhabitants, a building is an active, autonomous entity that pursues its own goals (energy consumption, security). To fulfill this goal, a building must continually take decisions: but specifying rules that describe which actions to take, at which time, and because of which conditions, is complex and time consuming. In addition, these rules have to be changed constantly, because preferences and needs of users change. The building thus needs to learn its own rules of behavior, and continually adapt them, based on feedback from its occupants.

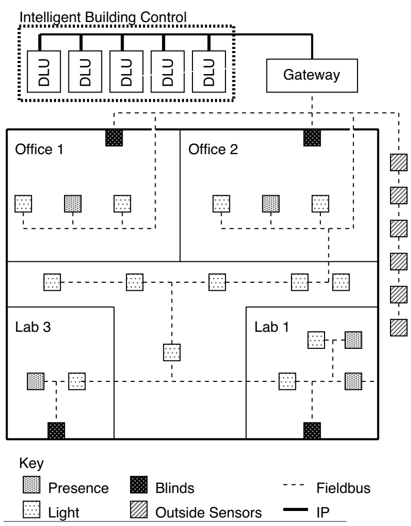

Figure 1. The floor of a typical building structured into rooms. All sensors and effectors are wired to a common fieldbus network. A gateway allows the agents to access the fieldbus network.

Engineering such a system poses a number of challenges. Decisions must be made in near-real time. The system must have a way to interact with its users to obtain feedback. On the other hand, it should not intrude on the user. In this article we particularly highlight the problem[1] from a machine-learning perspective.

Unsupervised learning and control

We tackle the learning problem in several stages. All knowledge that the system has about its goals and users is stored as fuzzy logic[2] rules (in a rulebase). These rules are continually adapted based on feedback from the environment. A number of agents,[3] each responsible for a small part of the whole decision space, use individual rulebases for decision-making. Using fuzzy rules has the advantage that the rulebase, although completely automatically constructed by the learning algorithm, is easily read and modifiable by humans.

Two-stage memory and decision process

The inherent non-stationarity and noisiness of user interactions with the building make it difficult to acquire stable, long-term knowledge. On one hand, the building must be able to retain and recognize long-term overall patterns of behavior; but on the other hand it must also react to short-term changes in requirements without destroying such long-term knowledge. Motivated by biological systems, we use a two-stage memory process: short-term (STM) and long-term (LTM) memory.[4] Newly-acquired knowledge remains in STM and does not become part of LTM before it is confirmed and generalized. STM changes rapidly based on user demand and always gets precedence over LTM during day-to-day operation. However, if the former is discarded after a short time if it does not get incorporated into latter. Knowledge transformation from STM to LTM, as in natural systems, takes place transparently: in parallel to decision making and learning. For efficient generalization of knowledge, a genetic algorithm[5] is used (see Figure 2). This serves as an efficient mechanism to retain long-term knowledge in spite of inherent contradictions in feedback that occur in the short-term. This is a two-way process: knowledge in LTM, as in STM, is discarded if it is not needed anymore, or if it contradicts new facts. This property is essential because of the non-stationarity of the environment.

We have been running various versions of this system in a real building to evaluate its performance. Initial results are promising and show that the system is able to improve the inhabitants' comfort[1] and minimize energy consumption.

Ueli Rutishauser*†, Jonas Trindler† and Rodney Douglas†
*Computation and Neural Systems
California Institute of Technology
†Institute of Neuroinformatics
ETH/University of Zurich
E-mail: urut@caltech.edu
and trindler@ini.phys.ethz.ch
http://www.ini.unizh.ch/~trindler/abi

References

1. U. Rutishauser, J. Joller, R.J. Douglas, *Control and learning of ambience by an intelligent building,* **IEEE T. on Systems, Man and Cybernetics,** submitted, 2004.
2. R. Fuller, **Neural fuzzy systems,** Springer, Berlin, 2000.
3. G. Weiss, **Multiagent systems,** The MIT Press, Cambridge, 2001.
4. H. Zhou, *CSM: A Computational Model of Cumulative Learning,* **Machine Learning. 5,** pp 383-406, 1990.
5. M. Mitchell, **An introduction to Genetic Algorithms,** The MIT Press, Cambridge, 1998.

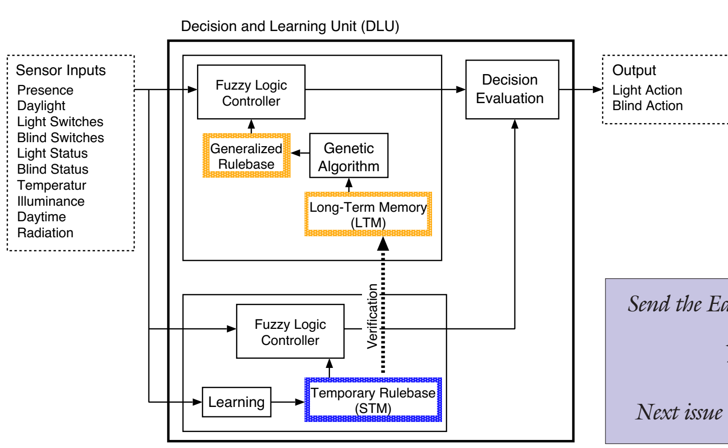

Figure 2. Decisions are taken based on fuzzy logic rules, which are stored in either long- or short-term memory. Feedback from the environment is used to continually adapt the rules.

Send the Editor ideas for the next issue by:

19 July 2004

Next issue copy due 17 September 2004.