A Smarter Smart Home:

Case Studies of Ambient Intelligence

Stephen Makonin, Lyn Bartram, and Fred Popowich
Simon Fraser University

Abstract—Research has shown that small changes in behavior in how we use our homes can result in substantial energy and water savings. Home automation and the integration of computational intelligence capabilities in the "smart home" are often cited as promising advances in the design and renovation of efficient buildings. However, the design and implementation of such technologies are largely based on energy-use simulations, smart automation of the building systems, and components for optimal performance rather than on effectively supporting how people use their homes. Additional factors, including system complexity and awkward automation, can discourage acceptance of smart home technologies. In this article, the authors propose that technological support for sustainable home use lies in more subtle and contextually appropriate interventions that integrate more informative models of occupant behavior, provide hybrid levels of automated control, and use ambient sensing for localized decisions. They discuss several cases from their experience in designing sustainable home systems and describe two current design cases for ambient intelligence in home control.

Keywords: smart home; ambient intelligence; adaptive lighting; adaptive HVAC, pervasive computing

A sustainable home is more than a green building; it's a living experience that encourages occupants to use resources more effectively. Research has shown that small changes in behavior—such as turning off lights, changing the thermostat, or shortening showers—can result in substantial energy and water savings.[1] But changing how occupants use these resources is challenging.

Combining ubiquitous computing and computational intelligence offers an opportunity to help occupants dynamically interact with building technologies for feedback and control regarding performance and atmosphere while empowering them as agents of behavioral change. These technologies and patterns of use are the focus of recent design research,[2–6] but sustainable home design is still in its infancy. We know very little about how to design, situate, and integrate these various technologies to support occupants in making more efficient resource-use decisions.

Our goal is to help occupants optimize home use through computational interventions without imposing undue technological complexity, effort, or inconvenience. These interventions comprise a combination of information and interaction design;[2,3,6] automation and control; adaptive intelligent agents; and distributed, ambient sensing. We seek to achieve a level of support with what Peter Tolmie and his colleagues refer to as *unremarkable computing*[7]—that is, technology that assists without imposing and that remains largely in the background.

Critical questions to address in this research are[8]

- What are the right conditions to affect automated or intelligent decisions related to controlling the residents' environment?
- What are the recovery and override efforts in correcting an automated system action that occupants have deemed wrong?
- How do we accommodate and adapt to different user profiles and tasks?

Here, we provide a framework for human-computer interaction in a smart home and present case studies based on our experience with designing two sustainable homes.

DOI: 10.1109/MPRV.2012.58

Beyond the Typical Smart Home

Architects and designers of sustainable buildings are becoming more aware of the importance of *occupant intelligence*—that is, how occupants of a building engage with its operation.[9,10] Buildings designed around occupant intelligence should provide flexible, adaptive task environments, refined control zones, and technologies that maximize occupants' access to adaptive opportunities.[9]

Architects, engineers, and system designers must reframe design strategies, given the co-evolution of human and building intelligence, to encourage and reinforce sustainable use. This requires new models of design that go beyond the typical smart home, encompassing occupant behavior and motivational strategies and exploring how automation can affect occupants' daily rituals and sense of comfort. This last issue is particularly challenging.

The Automation Challenge

Smart homes are often difficult to manage.[11] Even technologically passionate residents who equip their homes with varying degrees of control often find the control systems_cumbersome.[5] Managing such systems is even more difficult when the automation moves from user control to some form of rule-based behavior, because the rules might insufficiently reflect what the resident actually wants.[5] In a recent study of home automation,[5] most users expressed a desire for such systems to help with conservation actions—reducing heat and turning off lights, appliances, and other home devices—yet current tools are still small, piecemeal devices that aren't integrated into the larger home systems.

Humans have an uneasy relationship with automation,[12] as we discovered in our first net-zero home project.[14] At the same time, automation is well suited to reducing energy use in circumstances involving vigilant attention to simple, distributed controls, such as thermostats and electrical switches. Computational intelligence offers the promise of simple, adaptive controls that can minimize user effort and optimize energy use—but the issue is how to best provide this assistance without causing discomfort and incurring override reactions, which are problematic in overly intrusive systems.

Automated home systems provide two types of control:[5] user control (an aggregation of lower-level controls into more zone-specific groups) and rule-based control, where the system makes decisions—usually based on a schedule. These types of controls are difficult to manage because they force the static definition of a complex, a priori configuration that lacks a dynamic and holistic view of the home.[5,8,11] According to a recent study, occupants overwhelmingly view rule-based systems—the programmable thermostat being the simplest example—as problematic and error-prone.[5] There are two reasons for such systems' failure.

First, the thermostat interface is nonstandard and invariably complex, so few users configure them successfully. More important, however, is the functional design. Most peoples' lives are more complex and variable than the simple schedules the rules accept; thus, people end up setting simple schedules and needlessly running house systems when they're not home—notably, heating and cooling systems.

Design Considerations

In line with Stephen Intille[14] and others, we advocate for technology that helps humans behave appropriately rather than relieving them of any operational involvement with their homes.[15] Issues of trust and customization are important.[16] The advantages of smart homes to date have tended to be outweighed by complexity, but it's clear that context-aware systems,[15] distributed smart sensor or agent networks,[17,18] and adaptive behavior hold substantial promise. However, given the simplistic conditions that might guide behavior, we posit that two factors influence the efficacy of these approaches and they must be expressly modeled in the design.

DOI: 10.1109/MPRV.2012.58

First, we need to consider *the cost of being wrong*. For example, what happens if all the lights turn off because you fall asleep on the couch? The cost of being wrong might simply be the effort required to recover a desired state, but it could also implicate additional technological interventions, such as the need to include more information in the interface about the performance effects when running in a nonoptimal energy-efficiency mode.

Second, we should recognize that *the appropriateness of smart intervention* can be highly contextual. Consider adaptive approaches to night lighting, which could provide adequate light for navigation when someone gets out of bed during the night without disturbing his or her Circadian sleep rhythms. The same manipulation of light levels during the day could prove intrusive. Designers who wish to explore the affordances and potential of these systems must be able to simulate them with a variable degree of automation. Coupled with feedback systems, the challenging design question is to balance the appropriate responsibility between prompting occupants for action and assisting them by carrying out that action automatically.

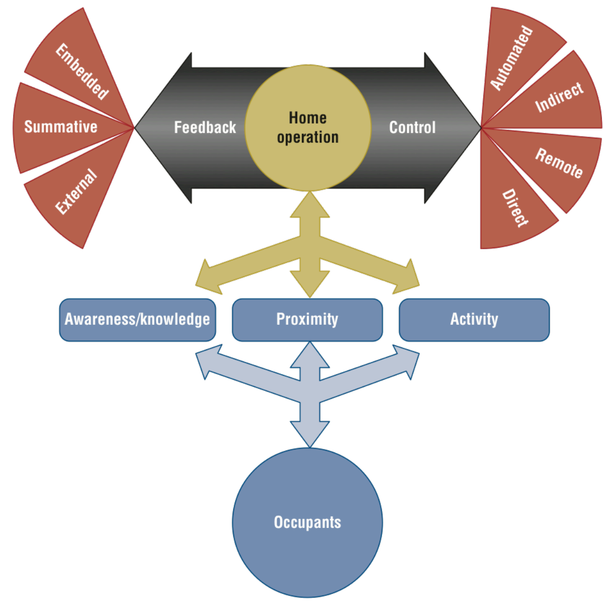

Figure 1. A depiction of human-home interaction. How humans interact with the home is central to any automation strategy and depends on awareness or knowledge, proximity, and activity.

So, what factors should be considered when examining appropriate technological intervention with the occupants of a smart home?

Smart Intervention Framework

We conceptually model the smart home as a complex ecosystem involving the occupants, physical and operational components of the home, external and internal context, and dependencies between these. The operational components are the technologies that let the occupant run things in the home. Along with the standard idea of smart home technologies—such as monitoring displays and thermostats—we point out that windows, doors, appliance displays, fans, and other aspects of the home not traditionally considered technology are in fact devices that allow some form of feedback and control.

How humans interact with the home is central to any automation strategy and depends on awareness or knowledge, proximity, and activity (see Figure 1).

The home is a complex system whose current state encompasses numerous attributes. Its state is also dependent, in a hierarchical manner, on its occupants, household activities (animate and inanimate), and the configuration of its components. Figure 2a provides a high-level view of some of these dependencies.

DOI: 10.1109/MPRV.2012.58

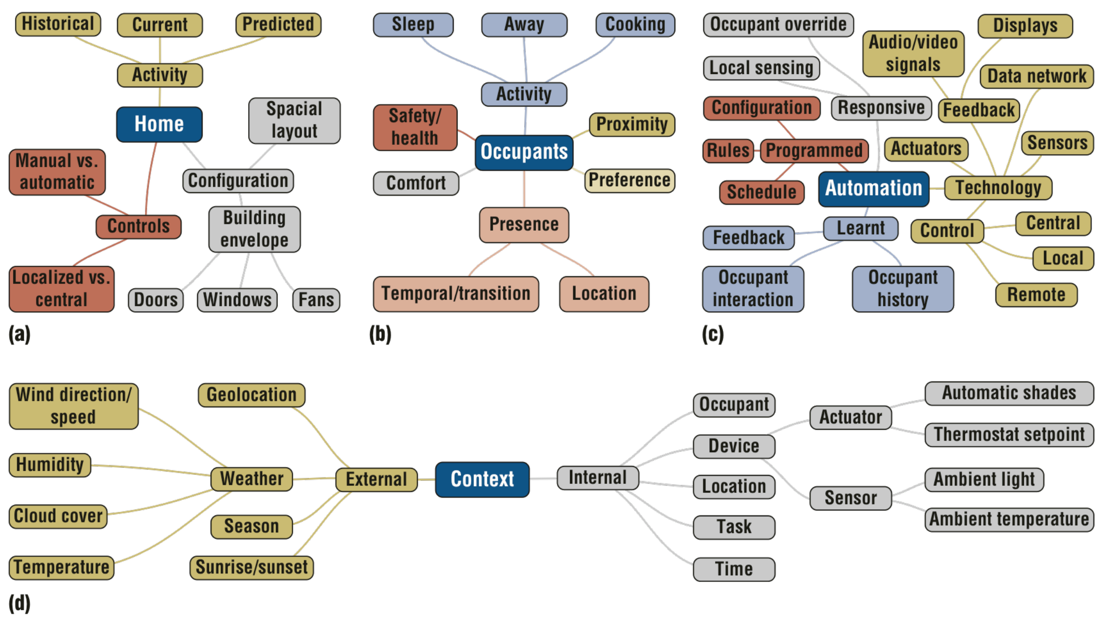

Figure 2. Concept maps of the different factors relevant for our smart intervention framework: the (a) home, (b) occupants, (c) automation, and (d) context.

Occupant use of the home involves numerous factors. These are computationally difficult to capture, but even simple models of activity, presence, proximity, health, and comfort are critical to designing intervention strategies (see Figure 2b). How the occupant controls and is aware of the home state will depend on these as well as the affordances of home operation detailed in Figure 1. For example, the occupant who wants detailed control over lighting in the evening might benefit from automated light activation in the middle of the night when getting out of bed to check on a child—especially if the controls for the lights are across the room. Similarly, a display that shows energy use for the home is only useful when the occupant is close enough to read it.

So, different classes of automation should exist in a home, ranging from learned to programmed automation and from fully automatic to human-initiated actions (see Figure 2c). The technology used in the automation (sensors and actuators) and how they communicate (the data network) has a great influence on the automation's functionality and ability to perform with an acceptable level of accuracy.

Furthermore, when considering a smart intervention for a home, we need to provide a context for the intervention that includes input values obtained not only inside the home but also outside the home, as Figure 2d shows (this is an extension of Wolfgang Kaltz and his colleagues' work[19]). Note that the input from animate agents (occupants) can be considered in addition to inanimate devices.

Consequently, we can formally view a smart intervention (*SI*) as a function to produce a new home state:

$$H'_S = SI(H_S, H, O, Ctx, Aut, E),$$

where H_S is the previous home state, H is the home, O is the occupants, Ctx is the context, Aut is the automation, and E is an evaluation of the appropriateness function. We introduce E to measure the cost of being wrong. H_S is an instantiation of H, O, Ctx, and Aut at a point along a given time sequence.

E would be a function of normalized measurements (accuracy, intuitiveness, and overrideability) where we maximize over the function to balance out the goals of energy efficiency and comfort. Note that E can vary from household to household and can even change over time. Furthermore, the smart intervention

DOI: 10.1109/MPRV.2012.58

function can be defined such that an intervention won't occur if a minimum *E* threshold isn't met.

Smart interventions are realized when we consider the previous state of the home (Figure 2a), the type and activity of the occupants (Figure 2b), the type of automation (Figure 2c), the context of the home (Figure 2d), and a measure of how appropriate the action is for the occupants. Obviously, an intervention having low accuracy, requiring the occupants to read a how-to manual, or not letting the occupants override the system when it's wrong wouldn't be smart or appropriate.

This *appropriateness* measure is extremely important yet challenging. In some cases, it relates to user preferences (such as comfort); in other cases, it depends on a more complex balance between previously known best practices (such as light levels for circadian rhythms), user activity, and external context. Furthermore, appropriate interventions aren't noticed as much as inappropriate ones and thus provoke less annoyance. Consequently, typical HCI evaluation methods that rely on self-reporting, usability, or performance metrics don't capture this measure of appropriate intervention at the necessary systemic level. This is why we instead derived our framework from a series of case studies, drawn from our experience in building sustainable home systems.

Case Study: Occupancy and Sleep Detection

A basic issue is determining whether someone is home. Although accurately performing this task can be difficult, it has several exploitable benefits in terms of energy conservation, sustainability, and cost savings for the occupant. Consider an intelligent agent (IA) that can turn down the heat when no one is home. Such automatic behavior could save money by lowering heating costs. It could also save the occupant time, because he or she wouldn't need to (remember to) override the thermostat setting.

Detecting if the occupants are sleeping can provide additional savings by automatically turning off lights that the occupant forgot about or setting the home automation system to a *sleep mode*. Our work on the Home Occupancy Agent (HOA) investigated such automated system behaviors.[17,18] We first created rules for detecting nightly sleep patterns and then added rules for daytime nap detection from the activities we observed in our test house. Our results were mixed.

Our initial examination of sleep detection used only ambient light sensors to determine the occupants' nightly sleeping period. However, such an algorithm is brittle because sleeping patterns from one home to another can be extremely different—even within a given household, sleeping patterns change over time.

So, we decided to look for a lux spike in certain ambient light sensors and create a general rule for the HOA to follow. To simplify things, we assumed that the occupants went to sleep before midnight and awoke before noon. Over the eight-month period, we collected data at 15-minute intervals, which the system used to achieve an f-score (the harmonic mean of precision and recall) of 0.89, indicating that the algorithm worked well in our test home.

However, the performance of our sleep-detection algorithms could have been better. Polling the ambient light sensors every 15 minutes was too infrequent. If the occupants went to bed between two interval periods and took less than 15 minutes to go to bed, then the lighting trigger never fired, causing inaccuracies. Also, there were anomalies of 1.2 lux spikes during summer that caused the algorithm to think that the sleep period ended prematurely.

When determining a napping period, we looked for light-level changes in a room designated for naps. However, the light levels in this nap room were greatly affected by opening and closing the curtains, which caused a breakdown in our rules. After running the algorithm over the eight-month period, the resulting f-score was only 0.10. So what went wrong?

We wrote rules based on a subset (one month) of the overall eight months of data. Environmental lighting and the variability of the occupants napping were the two main problems. Our findings suggest that online learning—such as reinforcement learning—could help make the rules-based agents more adaptable.

The HOA system aimed to replace the complicated programmable thermostat schedule using automated, smart intervention to reduce heating costs. Although we designed the HOA with the appropriateness of smart intervention in mind, the cost of being wrong was significant. Having the HOA raise and lower the heating system temperature setpoint (when it was wrong about the state of occupancy and sleep) could cause occupant discomfort. Furthermore, the occupant would be inconvenienced by (repeatedly) having to override the thermostat settings. Owing to the lack of accuracy, the system failed to balance occupancy comfort with optimal energy efficiency, leading us to conclude that the lower the accuracy, the greater the cost of being wrong.

DOI: 10.1109/MPRV.2012.58

Case Study: Optimal Energy Conservation

North House is a small solar-powered home designed to achieve net-zero performance (producing at least as much energy as it consumes) in the challenging Canadian climate. It incorporates sophisticated custom energy systems, adaptive intelligent building envelope technologies, specialized lighting and climate systems, and automated optimization behavior. For 10 days during the 2009 Solar Decathlon, North House saw more than 60,000 visitors, placing fourth (out of 20 entries) in the competition.

The control system in North House employed several optimized subsystems with intelligent behavior—notably, external shades that tracked the sun for efficient heating and cooling. The North House architects used an ESP–r simulation to tune the behavior of the intelligent shading in their original model,[13] which proved insufficient in the later design stages. We (the interaction design team) came in relatively late and didn't work on the final controls specification or system deployment in the house until approximately eight months before the competition.

We immediately identified a problem with the shade automation: What if the resident wanted to alter the external shades for comfort, privacy, and natural light? To the system, this potentially put North House into a *nonoptimal* mode. Interface modes that indicated the shades' mode and a time-out function to return to optimization were required. Our anecdotal experience in North House was that visitors (our potential users) struggled to understand how the system worked, what the optimal and nonoptimal modes represented, and how they might balance their needs with the optimized system state.

Placement of the interactive interface controls was equally constrained by the building envelope and materials. Because the facades of North House are almost entirely glass, there were few places to embed controls for lights, thermostats, or other devices. A digital touchscreen panel provided the only means for the resident to control, track, and manage energy performance in North House, and the only place to put it was over a deep kitchen counter. High levels of natural light during the day made it perceptually difficult to see. In addition, an iPhone application served as a remote control for lighting and thermal controls. This turned out to be more than just an interesting design piece; many visitors noted the need for a remote control so that they didn't have to move to a central location simply to turn on a light.

This case highlights the problems with designing an energy-efficient home without being mindful of the appropriateness of smart intervention and the cost of being wrong. The North House engineers created a technological complex system that was smart in terms of optimal energy efficiency, but the house didn't balance occupancy comfort with optimal energy efficiency. In fact, it didn't even know that the occupants existed. The effort for an occupant to override the system was considerable and created a major inconvenience. We can conclude that a lack of considering the appropriateness of smart intervention immediately creates a high cost of being wrong, so much so that the occupant might disable the entire system.

Both of our case studies reveal that rigid rule-based energy conservation must be dynamic in terms of co-existing with occupants. However, this doesn't mean that we should throw away the rules on which they operate and replace them with a new machine-learning technique. To make rules less rigid, we should make them dynamic so that they can adapt to occupant behaviors.

Design Case: Adaptive Lighting for Circadian Rhythms

We're currently exploring adaptive approaches to night lighting. When someone gets out of bed, we avoid disturbing their circadian sleep rhythms by providing adequate (but not full) light to navigate.[20] Use of the same light-level implementation during the day would be intrusive and inappropriate. Designers who wish to explore the affordances and potential of these systems will need to be able to simulate them with a variable degree of automation. Coupled with feedback systems, the challenging design question is determining when to prompt for occupant action and when to carry out that action automatically.

Figure 3 demonstrates one possible implementation. The system must be able to infer that the occupants are sleeping. Localized sensors in the room need to coordinate action and feedback to raise the ambient light level in the room until it has reached a lighting level below what would disturb the occupant's circadian sleep rhythms.

DOI: 10.1109/MPRV.2012.58

```
HOME_STATE (variables):
  home_sleeping = yes;
  bedroom_light_level = 0;
  bedroom_light_switch = off;

RULES:
  WHEN bedroom_light_switch IS pressed_once AND
    WHEN bedroom_light_switch WAS off AND
    WHEN home_sleeping IS yes
    SET bedroom_light_level = CIRCADIAN_LIGHT_LEVEL;
  WHEN bedroom_light_switch IS pressed_twice
    SET bedroom_light_level = MAX_LIGHT_LEVEL;
  WHEN bedroom_light_switch IS pressed_once AND
    WHEN home_sleeping IS no
    SET bedroom_light_level = MAX_LIGHT_LEVEL;
  WHEN bedroom_light_switch IS pressed_once AND
    WHEN bedroom_light_switch WAS on
    SET bedroom_light_level = MIN_LIGHT_LEVEL;
```

Figure 3. Circadian-rhythm-aware lighting at night. The system must be able to infer that the occupants are sleeping. Localized sensors in the room need to coordinate action and feedback to raise the ambient light level in the room until it has reached a lighting level below what would disturb the occupant's circadian sleep rhythms.

The rules in Figure 3 are triggered when an event happens. The *home sleeping* state is determined by an online learning algorithm. The first rule will set the ambient light in a room to a lighting level that won't disturb the occupant's circadian sleep rhythms if the light switch is pressed once when it's off and the house is in *sleeping mode*. The second rule overrides the first if the occupant presses the light switch twice. In case the occupant wants to wake up and fully turn on the lights, the third rule says that if the light switch is pressed once when it was off and the house isn't in *sleeping mode*, then we should have full lighting. The fourth rule handles turning off the lights in a room.

The proposed smart intervention operates much the same way as most digital dimmer switches do and eliminates the inconvenience of having the occupant's circadian sleep rhythms interrupted while keeping the light switch's intuitive operation. The cost of being wrong is eliminated because of the easy override with a second press. A balance has been struck between occupant comfort and the technology.

Design Case: Adaptive HVAC Acclimatization

Monitoring home occupancy using power monitoring and ambient light sensors is a first step toward achieving an *adaptive HVAC system*. We've further modified the system to use the arming and disarming of the alarm system to determine home occupancy.

When a home is unoccupied, it can be put into a *power saving mode* by changing the thermostat setpoint and turning off equipment such as ambient displays. We can do by sending specific commands to different equipment around the house or by broadcasting a message over a data network of the *home state* changes. When the occupant returns, the *power saving mode* is reversed. In the case of heating (HVAC), we might want to infer when the occupant will most likely be home and preheat the home. Figure 4 demonstrates how this might look.

DOI: 10.1109/MPRV.2012.58

```
HOME_STATE (variables):
  home_occupied = yes;
  alarm_system = disarmed;
  temp_setpoint = 21.0;
  TEMP_INCREMENT = VACANT_TEMP_SETPOINT / 2;

RULES:
  WHEN alarm_system IS armed_away
    SET home_occupied = no AND
    SET temp_setpoint = VACANT_TEMP_SETPOINT AND
    SET ambient_displays.send_cmd = 'turn off' AND
    SET automation.send_cmd = 'set away profile';
  WHEN alarm_system IS armed_away AND
    WHEN home_occupied 95% yes WITHIN_TIME - 30 MIN
    SET temp_setpoint += TEMP_INCREMENT
  WHEN alarm_system IS disarmed
    SET home_occupied = yes AND
    SET temp_setpoint += TEMP_INCREMENT AND
    SET ambient_displays.send_cmd = 'turn on' AND
    SET automation.send_cmd = 'set athome profile';
```

Figure 4. Electricity and heating cost-savings with occupancy detection. When a home is unoccupied, it can be put into power-saving mode by changing the thermostat setpoint. When there's a high chance (in this case, 95 percent) that the occupant will be home within 30 minutes, the system slowly increases the heat.

This example doesn't address every situation in which the user might want to override the *temperature setpoint*. Instead, it highlights a couple of key points that demonstrate how rule-based systems with online learning could be used. The first rule says that when the security system is armed in the *away mode*, the following actions apply: consider the house to be unoccupied, change the thermostat setpoint to save energy, turn off all ambient displays, and set the home automation system to an *away profile*. We could have added rules to check for thermostat overrides, but that would complicate our simple example.

The second rule specifies that when the house is unoccupied and there's a high chance (in this case, 95 percent) that the occupant will be home within 30 minutes, then we slowly increase the heat (or cooling if TEMP_INCREMENT is negative) so that the occupant comes back to a more pleasant house temperature. This inference comes from a learning algorithm that makes the rules more adaptable.

The third rule says that the occupant is home when the security system is disarmed. Once the user is home, we change the thermostat setpoint to its original value, turn on the ambient displays, and set the home automation system to an *at home profile*.

Using historical occupancy usage data to control the setpoint in an HVAC system is a goal similar to that of the first case study; however, we reduce significantly the cost of being wrong. If the smart intervention of the adaptive HVAC system is wrong, it would be when the occupant either returns home earlier or later than usual. If the occupant returns home earlier than usual, the heating/cooling setpoint would have been changed at the same time the occupant enters the home (worst-case scenario). The occupant's comfort wouldn't be optimal, but the HVAC system is optimal in terms of energy efficiency. This wouldn't be the case if the occupant used the programmable option of the thermostat. So, the cost of being wrong is still less than without the automated system. Additionally, the occupant doesn't have the inconvenience of setting the thermostat after entering the home.

If the occupant returns home later than usual, his or her comfort would be optimal but not the energy efficiency (meaning the occupant would have spent money heating an empty home). In this case, the occupant might be inconvenienced by a slightly higher energy bill, but comfort is maintained, and there's the added convenience of not having to worry about programming the thermostat.

Having dynamic rules that adapt to occupant behavior strikes a balance between occupant comfort and energy efficiency. The fact that there are rules means that users can understand and modify them to increase satisfaction with the system. However, the same can't be said about having a system that relies on machine-learning techniques, such as artificial neural networks and support vector machines, which rely on statistical probability and the recomposition of that problem into linearly separable classifications.

A New Design Model

Understanding how residents use energy in the home is an emerging area of active design research. Designing for humans requires a user-centered approach.[8] A key concept in *ambient intelligence* is that its operation should have to be explicitly learned or managed by the occupant. This has been a problem for smart homes in general, where poor usability and intrusive and inappropriate operation have overwhelmed users.[12] Explicitly modeling the cost of inaccurate operation with the cost to the user to recover or override the decision provides a mechanism to evaluate how appropriate and effective a reasoned intervention might be and (hopefully) discourages a simplistic reliance on how the home should operate.

Given the state of the art, it's clearly infeasible to design a rule-based (or learning) system that fully captures the complexity of daily life, but even simple models can go a long way toward appropriate interventions. Perhaps more importantly, these models and their exploration in use can help us avoid inappropriate interventions that evoke frustration and technological resistance in occupants.

The smart intervention need not be full automation; in fact, it might resemble a "power-steering" assistive approach rather than a fully automated "self-driving" approach, as in the lighting example. As we pointed out, a low-cost intervention is one that lets the occupant easily and efficiently override the intervention without undue effort (as in the lighting case study). We propose that the first step to providing more effective ambient intelligence in the home is to identify scenarios and contexts in which these low-cost interventions can be usefully deployed: in other words, rather than trying to solve the entire problem, we're interested in discovering the subproblems that are most tractable to these kinds of interventions.

Taking this approach has the advantage that smart interventions can be localized and distributed without relying on whole-home systems. It also means that hybrid approaches (where the intelligent operation exists in conjunction with user overrides) can be more easily implemented. Distributing rules with learning that solve subproblems should reside in sensors and other equipment around the home. This distribution of intelligence is needed to reduce system latency and create modularity.

Reducing system latency increases system responsiveness to human-home interactions—which is important when responsiveness must be immediate, as in the case of adaptive lighting for circadian rhythms (responsiveness to light-switch presses must be in the microseconds). Addressing subproblems in a modularized fashion creates a type of plug-and-play system. This benefits the smart home by allowing functionality to be extensible, malleable, and upgradeable.

Conclusion

We've derived this model from our emerging experience in designing systems for sustainable homes. This emphasizes the fact that none of the approaches we advocate can be theoretically tested in isolation but must be actually deployed and studied in situ. We're currently implementing and exploring ambient intelligence interventions in two sustainable homes, and plan to extend the deployment to more. Our current and future research will rely heavily on evaluation with occupants in a variety of home contexts.

Acknowledgments

This work was supported in part by a grant from the National Sciences and Engineering Research Council (NSERC) of Canada, along with the Graphics, Animation, and New Media Network of Centres of Excellence (GRAND NCE) of Canada.

DOI: 10.1109/MPRV.2012.58

References

1. R. Woodbury et al., *Buildings and Climate Solutions*, Pacific Inst. for Climate Solutions, Nov. 2008.
2. M. Chetty, D. Tran, and R. Grinter, "Getting to Green: Understanding Resource Consumption in the Home," *Proc. 10th Int'l Conf. Ubiquitous Computing*, ACM, 2008, pp. 242–251.
3. E. Blevis, "Sustainable Interaction Design: Invention and Disposal, Renewal & Reuse," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, ACM, 2007, pp. 503–512.
4. A. Woodruff, J. Hasbrouck, and S. Augustin, "A Bright Green Perspective on Sustainable Choices," *Proc. 26th Ann. SIGCHI Conf. Human Factors in Computing Systems*, ACM, 2008, pp. 313–322.
5. A. Brush et al., "Home Automation in the Wild: Challenges and Opportunities," *ACM CHI*, 2011, pp. 2115–2124.
6. L. Bartram, J. Rodgers, and K. Muise, "Chasing the Negawatt: Visualization for Sustainable Living," *IEEE Computer Graphics and Applications*, vol. 30, no. 3 2010, pp. 8–14.
7. P. Tolmie et al., "Unremarkable Computing," Proc. SIGCHI Conf. Human Factors in Computing Systems: Changing Our World, Changing Ourselves, ACM, 2002, pp. 399–406.
8. L. Bartram and R. Woodbury, "Smart Homes or Smart Occupants? Reframing Computational Design Models for the Green Home," *Proc. 2011 AAAI Spring Symp. Series*, Assoc. for the Advancement of Artificial Intelligence (AAAI), 2011; www.aaai.org/ocs/index.php/SSS/SSS11/paper/view/2500.
9. R. Cole and Z. Brown, "Reconciling Human and Automated Intelligence in the Provision of Occupant Comfort," *Intelligent Buildings Int'l*, vol. 1, no. 1, 2009, pp. 39–55.
10. A. Leaman and B. Bordass, "Assessing Building Performance in Use 4: The Probe Occupant Surveys and Their Implications," *Building Research & Information*, vol. 29, no. 2, 2001, pp. 129–143.
11. R. Eckl and A. MacWilliams, "Smart Home Challenges and Approaches to Solve Them: A Practical Industrial Perspective," *Intelligent Interactive Assistance and Mobile Multimedia Computing*, Springer-Verlag, 2009, pp. 119–130.
12. D. Woods, "Decomposing Automation: Apparent Simplicity, Real Complexity," *Automation and Human Performance: Theory and Applications*, CRC Press, 1996, pp. 3–17.
13. K. Velikov and L. Bartram, "North House: Developing Intelligent Building Technology and User Interface in Energy Independent Domestic Environments," *Proc. 26th Conf. on Passive and Low Energy Architecture* (PLEA 09), Les Presses de l'Université Laval 2009.
14. S. Intille, "The Goal: Smart People, Not Smart Homes," *Proc. 4th Int'l Conf. Smart Homes and Health Telematics* (Icost 06), IOS Press, 2006, p. 3.
15. E. Tapia, S. Intille, and K. Larson, "Activity Recognition in the Home Using Simple and Ubiquitous Sensors," *Proc. 2nd Int'l Conf. Pervasive Computing*, LNCS 3001, Springer, 2004, pp. 158–175.
16. P. Noy et al., "Design Issues in Personalizing Intelligent Buildings," *Proc. 2nd IET Int'l Conf. Intelligent Environments* (IE 06), IEEE, vol. 1, 2006, pp. 143–149.
17. S. Makonin and F. Popowich, "An Intelligent Agent for Determining Home Occupancy Using Power Monitors and Light Sensors," *Toward Useful Services for Elderly and People with Disabilities*, Springer-Verlag, 2011, pp. 236–240.
18. S. Makonin and F. Popowich, "Home Occupancy Agent: Occupancy and Sleep Detection," *GSTF J. on Computing*, vol. 2, no. 1, 2012, pp. 182–186.
19. J. Kaltz, J. Ziegler, and S. Lohmann, "Context-Aware Web Engineering: Modeling and Applications," *Revue d'Intelligence Artificielle*, vol. 19, no. 3, 2005, pp. 439–458.
20. J. Veitch, "Principles of Healthy Lighting: Highlights of CIE TC 6-11's Forthcoming Report," *Final Report—The 5th Int'l LRO Lighting Research Symp.—Light and Human Health*, NRC, 2009, pp. 1–8.

Stephen Makonin *is a PhD candidate in the School of Computing Science at Simon Fraser University. He's also an applied researcher in the Technology Centre at the British Columbia Institute of Technology. His research interests focus on intelligent systems in smart meters, smart homes, and a smart grid for energy conservation and sustainability. He blogs about his research at http://eco-sustain.org. Contact him at stephen@makonin.com.*

Lyn Bartram *is an associate professor in the School of Interactive Art and Technology at Simon Fraser University, where she's a cofounder and director of the Human-Centered Systems for Sustainable Living research group. Her theoretical research interests span perception, cognition, visualization, and human-computer interaction, particularly with respect to how humans interact with technologies off the desktop. Her current work explores the intersecting potential of information technologies, ubiquitous computing, computational media, and sustainable building design in encouraging conservation and reducing our ecological footprint in our homes and personal activities. Bartram received her PhD in computer science from Simon Fraser University. Contact her at lyn@sfu.ca.*

Fred Popowich *is a professor of computing science at Simon Fraser University, director of the Vancouver Institute for Visual Analytics, and member of the Institute for Computing, Information and Cognitive Systems at the University of British Columbia. His research involves several aspects of artificial intelligence with an overall objective to make it easier for human beings to interact with computers, or to make it easier for human beings to interact with each other. He's also active in technology commercialization. Popowich received his PhD in cognitive science from the University of Edinburgh. Contact him at popowich@sfu.ca.*