
CARL: Activity-Aware Automation for Energy Efficiency

Brian L. Thomas

Washington State University
355 Spokane St. EME 102
Pullman, WA 99164 USA
bthomas@eecs.wsu.edu

Diane J. Cook

Washington State University
355 Spokane St. EME 121
Pullman, WA 99164 USA
cook@eecs.wsu.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org. *UbiComp '14*, September 13 - 17 2014, Seattle, WA, USA

Copyright 2014 ACM 978-1-4503-3047-3/14/09...\$15.00.
<http://dx.doi.org/10.1145/2638728.2641554>

Abstract

Society is becoming increasingly aware of the impact that our lifestyle choices have on energy usage and the environment. This paper explores the hypothesis that ubiquitous computing technologies can be used to understand this impact and to provide activity-aware interventions to reduce energy consumption. Specifically, we introduce a method to provide energy-efficient home automation based on the recognition of activities and their associated devices. We describe CARL (CASAS Activity-based Resource Limitation), a prototype energy-efficient smart home, and evaluate the performance of our activity-aware automation when using both historic and real-time sensor data to drive intelligent home automation.

Author Keywords

smart homes, activity recognition, home automation, energy efficiency

ACM Classification Keywords

I.2.6 [Learning]: (concept learning); H.2.8 [Database Applications]: (data mining)

Introduction

The impacts of lifestyle choices on energy usage and the environment are increasingly becoming more noticeable

and discussed in society. As a result, research resources are being directed toward green technology, environmentally-friendly building designs, and active demand response within the smart grid. In this paper, we look more closely at the user side of sustainability and ways ubiquitous computing may aid in reducing energy consumption. In particular, we design a machine learning-based intervention to promote energy-efficient, sustainable behavior through home automation.

In 2010, the United States consumed 97.722 quadrillion btu of energy, a 300% increase from 1949 [12]. The growth of energy usage is not entirely due to manufacturing plants and automobiles. In fact, the residential sector is responsible for 16-50% of energy consumed worldwide [7]. A utility bill provides insights about monthly whole-building consumption, which leaves homeowners to guess at the reasons for an unusually high or low bill. Earlier studies reveal that home residents reduce consumption by as much as 15% in response to simply viewing raw usage data [4] and indicate that widely-varying residential behavior can influence usage by as much as 100% in a single house [10].

This paper hypothesizes that activity recognition techniques can be used to provide energy-efficient automation in a smart home environment. This type of automation makes use of behavior analysis and ubiquitous computing to reduce energy consumption, while still supporting a resident's everyday activities. To validate this hypothesis, CARL (CASAS Activity-based Resource Limitation), a smart home automation technology was designed and implemented to control devices in the home. CARL utilizes a real-time activity recognition algorithm [8] to identify current activities being performed by the resident from sensor data in an instrumented smart home.

Electrical device usage distributions are estimated for each activity, and any devices not associated with the current activities may be turned off by CARL to reduce energy consumption. The utility of CARL is demonstrated in one of the smart home testbeds available as part of the CASAS (Center for Advanced Study of Adaptive Systems) smart home project at Washington State University. An evaluation of CARL is based on activity recognition accuracy, effectiveness at identifying devices to turn off, and energy consumption with, and without, activity-aware automation.

Related Work

Until recently, occupant behavior has been difficult to accurately capture. Adherence to self-reports of behavior and energy consumption is error prone for some populations [11] and whole-home meter monitoring does not capture the behaviors in the home that influence consumption.

Approaches have been utilized to explore the gap between the minimum amount of consumption that is needed for daily activities and the consumption that is actually observed [13]. Some early work has focused on linking resident activity with energy consumption. The hypothesis that providing users with knowledge about the relationship between their activities and energy consumption and automation support for energy reduction will result in substantial decreases in overall consumption is supported by an increasing body of work that links awareness of energy consumption and its impact on behavioral routines and behavioral change [1, 5, 9]. Until recently, validating this hypothesis was not possible. However, with the convergence of technologies in ubiquitous computing and machine learning, gathering data on human behavior is now automatable. Data can be collected from sensor-filled

smart homes [2] and smart phones [6] in an unobtrusive manner while individuals perform their normal daily routines. Because these sensor modalities operate in a continuous mode, feedback and interventions repeat ad infinitum, thereby maximizing the persistence effect.

Smart Home Monitoring

CARL utilizes the sensing and automation technologies available in the CASAS smart home architecture. The CASAS smart home project at Washington State University [3] defines a smart home environment as one that acquires and applies knowledge about its residents and their physical surroundings in order to improve their experience in that setting [1]. Such home environments, equipped with sensors for detecting features such as motion, light level, temperature, and energy and water consumption, are ideal testbeds for investigating the relationship between resident behavior and energy consumption and for providing activity-aware energy-efficient home automation.

The CASAS “smart home in a box” software architecture components are shown in Figure 1. During perception, sensed information flows up from the physical components through the middleware to the software applications. When taking an action, control moves down from the application layer to the physical components that automate the action. Each of the layers is lightweight, extensible and ready to use as is, without additional customization or training.

The CASAS physical layer contains hardware components including sensors and actuators, utilizing a ZigBee wireless mesh to communicate directly with the hardware components. The middleware layer is governed by a publish/subscribe manager, providing named broadcast

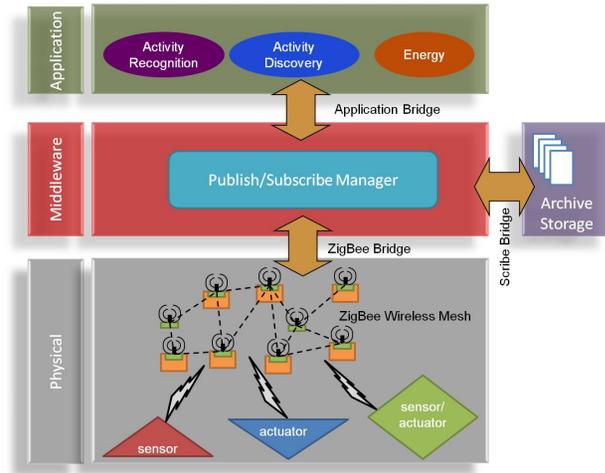


Figure 1: CASAS smart home components.

channels that allow component bridges to publish and receive messages. In addition, the middleware provides valuable services including adding time stamps to events, assigning UUIDs, and maintaining site-wide sensor state model. Every component of the CASAS architecture communicates via a customized bridge to this manager over XMPP. Examples of such bridges are the ZigBee bridge, the Scribe bridge, which archives messages in permanent storage, and bridges for each of the software components in the application layer. To date, CASAS has installed over 40 smart home testbeds.

Activity Recognition

CARL's activity recognition software, called AR (Activity Recognition) [8], provides real-time activity labeling as sensor events arrive in a stream. The problem of activity recognition is to map a sequence of sensor events, $x = \langle e_1 e_2 \dots e_n \rangle$, onto a value from a set of predefined

activity labels, $a \in A$. Activity recognition can be viewed as a type of supervised machine learning problem. The sequential nature of the input data, the ambiguous partitioning of data into distinct data points, and the common overlapping of activity classes mean that additional data processing must be performed in order to accomplish the goal of recognizing activities. The Activity Recognition (AR) algorithm learns a function that maps a feature vector, X , describing a particular sensor event sequence onto an activity label, $h : X \rightarrow A$. In this supervised machine learning problem the classes are the activity labels and the sensor events are represented by features combined into the input vector X of D dimensions. AR can use the learned function to recognize, or label, occurrences of the learned activity.

Automated Smart Home Testbed

To test our ideas for an activity-aware automated home, the Navan smart home testbed was constructed as part of the CASAS smart home project at Washington State University. The apartment houses a single resident and the floorplan and sensor layout for Navan are shown in Figure 2.

Navan utilizes the CASAS smart home architecture to sense and control the space. Passive infrared motion sensors with integrated luminosity sensors are placed throughout the residence to detect motion events and relative light levels at each location. Magnetic door closure sensors are placed on the doors and windows. Two magnetic door sensors adjacent to each other are used to identify when a sliding door or window is closed, cracked open up to 4", or open more than 4". Temperature sensors are placed throughout the smart home either as a single sensor or as a pair of temperature sensors (one sensor is placed 8" from the ceiling and one 12" from the

floor to identify temperature gradients). Power data is collected using a Ted5000 power meter, providing instantaneous wattage every few seconds. Arduino-based Wi-Fi thermostats were designed in-house and installed to control and monitor use of the baseboard heaters and log temperature setpoints.

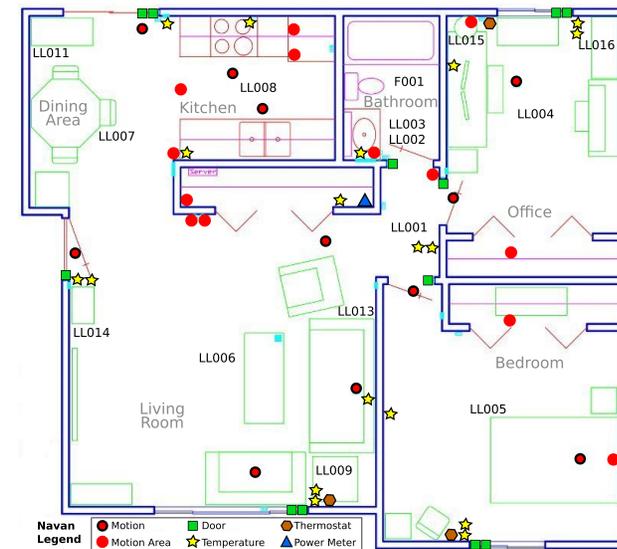


Figure 2: Navan automated home testbed.

In the Navan testbed, ZigBee light switches were attached to lights and the bathroom fan, facilitating automated control of these devices. In addition, custom electrical boxes equipped with ZigBee light switches were designed to allow monitoring and control of additional devices such as reading lamps and speakers. Each light switch reports the current state of the device upon state changes, and also reports button taps and tap counts. These taps provide a mechanism for the resident to provide feedback

to CARL. In Figure 2, the locations of devices that are controlled by the ZigBee light switches are indicated by the name of each device. All of the indicated devices are lights except for F001 (the bathroom fan) and LL014 (the television speakers). Figure 3 shows the energy consumption for each device over the monitored time period. Device voltages ranged from 30 watts (F001, bathroom fan) to 250 watts (LL003, bathroom heat lamp), while most lights were between 60 watts and 120 watts.

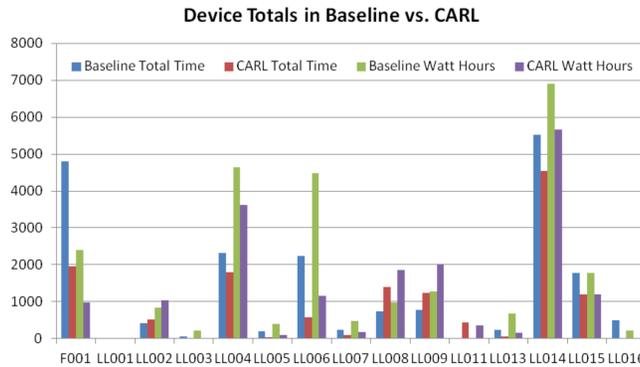


Figure 3: Time (in minutes) and energy consumption (in watt hours), listed for each device, with (CARL) and without (baseline) activity-aware home automation.

As a first step, AR was trained to learn a model of 17 routine activities based on the captured data. Ground truth activity labels were provided by an external annotator based on observing a playback of the sensor data events, the floorplan with sensor layout, and a summary of typical activity locations and times provided by the resident. The learned activities include Sleep, Watch TV, Toilet, Relax, Work on computer, Bathe, Water plants, Leave home, Enter home, Wash dishes, Eat,

Drink, Cook, Bed toilet transition, Dress, Entertain guests, and Other.

We recognize the robustness of an activity-aware automated home depends on the accuracy of the activity prediction algorithm itself. Therefore AR is first tested using 10-fold cross validation on the collected data. AR classifies the set of 17 activities in this case with 97.27% accuracy.

CARL Automation Algorithm

CARL's initial automation strategy is to turn off devices that are not needed in support of the current set of activities. To do this, first identify all of the current activities then determine the devices that are used by these activities.

Determining the set of current activities presents a particular challenge when processing data in real time. Segmentation algorithms can be used to mark the begin point and end point for any particular activity. However, these typically process historic data in offline mode. For real-time processing, a delay value is maintained for each activity class and device. The delay value $Delay_{A_i}$ represents the amount of time that has elapsed since a sensor event was observed that was labeled with activity label A_i . If the delay value for a particular activity is within a threshold number of time units then the activity is considered current. The delay value represents the amount of time that has elapsed since device D_j was observed changing state to ON. Each device D_j has an associated $DeviceThreshold_j$ that is used to prevent the device from being turned off seconds after the resident turned the device on, allowing new activities to be initiated before the device is evaluated by CARL. Additionally, turning the device on with a double tap of

```

CARL_Automation( $A, D$ ):
//  $A$  is the set of known activity classes
//  $D$  is the set of available devices
LongThreshold = DelayThreshold * 4
for  $j = 1$  to  $|D|$  : // set threshold for each device
|  $DeviceThreshold_j =$  DelayThreshold
 $t = 1$ 
while observe new sensor event  $e_t$  :
| CurrentActivities =  $\emptyset$ 
| CurrentDevices =  $\emptyset$ 
| // get activity label for sensor event
|  $A^t = AR(e_t, A)$ 
| for  $i = 1$  to  $|A|$  : // update times for each activity
| | expedite if  $A_i = A^t$  :
| | |  $Delay_{A_i} = 0$ 
| | else:
| | |  $Delay_{A_i} = Delay_{A_i} + 1$ 
| | for  $j = 1$  to  $|D|$  : // update times for each device
| | | if  $e_t == D_j$  and  $State(D_j, ON)$  :
| | | |  $Delay_{D_j} = 0$ 
| | | | if  $e_t == DoubleTapOn$  :
| | | | |  $DeviceThreshold_j =$  LongThreshold
| | | | else:
| | | | |  $DeviceThreshold_j =$  DelayThreshold
| | | else:
| | | |  $Delay_{D_j} = Delay_{D_j} + \delta_t$ 
| | for  $i = 1$  to  $|A|$  : // get set of current activities
| | | if  $Delay_{A_i} < DelayThreshold$  :
| | | | Append(CurrentActivities,  $A_i$ )
| | for  $j = 1$  to  $|D|$  : // get set of current devices
| | | if  $D_j \in Devices(CurrentActivities)$  :
| | | | Append(CurrentDevices,  $D_j$ )
| | for  $j = 1$  to  $|D|$  : // turn off not-needed devices
| | | if  $State(D_j, ON)$  and  $D_j \notin CurrentDevices$  :
| | | | if  $Delay_{D_j} > DeviceThreshold_j$  :
| | | | | ChangeState( $D_j, OFF$ )
|  $t+ = 1$ 

```

Algorithm 1: Activity-aware device control algorithm.

the light switch informs CARL that the automatic decision to turn it off was incorrect and $DeviceThreshold_j$ should be set to $LongThreshold$.

Next, for each activity A_i , the algorithm maintains a probability distribution over devices whose status is ON. Devices with a sufficiently great probability of being associated with A_i are left untouched when activity A_i is current. On the other hand, any device D_j whose status is not typically ON for any of the current activities is a candidate for CARL to turn off if is greater than $DeviceThreshold_j$. The activity-aware automation algorithm is summarized in Algorithm 1.

A modification to the algorithm was made after a brief pilot test with the resident. A guard statement was added to prevent CARL from turning off lights in the bathroom while the bathroom door was closed. There are no windows in the bathroom so unexpectedly turning off the light could present an unsafe situation.

Experimental Results

Data collection on a single resident for the baseline method occurred over 14 days in February 2014, shortly followed by 14 days of data collection with CARL running in February and March 2014. The resident was instructed on disabling CARL if the system started malfunctioning or there would be guests in the apartment, and did so for a 6 hour window when entertaining guests one evening. Data from the window was removed to ensure comparison of single resident behavior only between the baseline method and CARL. Ultimately, the primary goal is to determine how much energy savings can be achieved by using CARL. However, human factors need to be considered as well. As a result, CARL is evaluated on a number of different performance measures. These measures include: total

time, watt hours, RMSE, and NRMSE.

Total time measures the sum total time that devices are on. *Watt hours* measures the sum watt-hours of power that are used by a particular method (CARL or the baseline method of no automation) combined over all of the devices that could be controlled in the residence. Power data from other devices, refrigerator, stove, heater, etc. were not included in the calculations. The *RMSE* (Root Mean Square Error) is computed to indicate the error in CARL's automation due to being too aggressive (turning off devices when the resident needed them) or too conservative (not turning off devices when they were not needed). The too-conservative case occurred twice. In both situations, this was due to the smart home infrastructure not detecting the request to turn off a device. The too-aggressive situation was more common and was typically due to errors in recognizing the current activity. The errors were indicated by button taps by the resident and are therefore measured using this feedback, as shown in Equation 1. In this equation, D represents the number of devices, $CARL_Off(i)$ represents the number of times that CARL turned off device i , and $DoubleTap(i)$ represents the number of times that the resident indicated an automation error via a double tap.

$$RMSE = \sqrt{\frac{\sum_{i=0}^D (CARL_Off(i) - DoubleTap(i))^2}{D}} \quad (1)$$

The *NRMSE* (Normalized RMSE) is computed in order to better interpret the RMSE values. This is obtained by dividing the RMSE result by the maximum possible error. The resulting value ranges between 0.0 (no error) and 1.0 (maximum possible error).

$$NRMSE = \frac{RMSE}{\max_{i \in D} (CARL_Off(i))} \quad (2)$$

The results comparing total time and watt hours for CARL-based automation and no automation are summarized in Figure 3 and Table 1. As the figures indicate, there is consistent reduction in energy consumption and unnecessary device utilization through activity-aware automation. Most of the energy reduction is due to turning off lamps when they are not needed (e.g., when the resident is not at home) and turning off speakers when the resident is not in the living room listening to music or watching television.

	Baseline	CARL
Total Time (minutes)	19,826.04	13,867.64
Consumption (watt hours)	25,297.85	18,352.27
RMSE	0.00	11.15
NRMSE	0.00	0.15

Table 1: Reduction and accuracy results for CARL.

Table 1 summarizes the reduction and accuracy results. There is a significant ($p < 0.01$) reduction in both energy consumption and total time using activity aware automation. This automation does come at the expense of occasional overly-aggressive device control. This is reflective of the almost 3% error reported earlier for recognizing the targeted 17 activities in real time. Error is also influenced by the rate at which the resident switches between different activities, a resident with a stronger, deliberate routine will have less activity recognition errors than a resident who switches between several activities quickly and often. As activity recognition becomes more robust this error will be further reduced.

Conclusion

In this paper, a novel method for automating energy-efficient smart homes based on activity awareness

was introduced. Initial experimental results indicate that real-time activity recognition and device usage profiling can provide effective mechanisms for turning off unneeded devices and thereby reducing energy consumption. This work is just a first step in utilizing activity awareness for home automation. Reliable home automation depends heavily on reliable activity recognition, which is continuing to be refined. As this capability becomes more robust, more control may be given to CARL. This control will include turning devices on as well as off. Future work will also explore the creation of machine learning techniques to time various activities (i.e., washing clothes) for ideal conditions (e.g., when renewable or low-cost energy is available) and to automate alternative activity supporting mechanisms (e.g., open the blinds instead of turn on a light to provide sufficient light for reading).

Acknowledgments

This work is supported in part by the National Science Foundation (NSF grant CNS-0852172), the National Institute of Health (NIBIB grant R01EB009675), and the Life Sciences Discovery Fund.

References

- [1] Brounen, D., Kok, N., and Quigley, J. M. Residential energy literacy and conservation. *Chicago: 47th Annual AREUEA Meeting* (2012).
- [2] Cook, D., and Das, S. *Smart Environments: Technology, Protocols and Applications*, vol. 43. Wiley-Interscience, 2004.
- [3] Cook, D. J., Crandall, A. S., Thomas, B. L., and Krishnan, N. C. CASAS: A smart home in a box. *Computer* 46, 7 (2013).
- [4] Darby, S. The effectiveness of feedback on energy consumption. *A Review for DEFRA of the Literature on Metering, Billing and direct Displays* 486 (2006), 2006.
- [5] Darby, S. Smart metering: what potential for householder engagement? *Building Research & Information* 38, 5 (2010), 442–457.
- [6] Györbíró, N., Fábíán, A., and Hományi, G. An activity recognition system for mobile phones. *Mobile Networks and Applications* 14, 1 (Feb. 2009), 82–91.
- [7] Ichinose, T., Shimodozono, K., and Hanaki, K. Impact of anthropogenic heat on urban climate in Tokyo. *Atmospheric Environment* 33, 24 (1999), 3897–3909.
- [8] Krishnan, N. C., and Cook, D. J. Activity recognition on streaming sensor data. *Pervasive and Mobile Computing* 10, Part B, 0 (2014), 138 – 154.
- [9] Riche, Y., Dodge, J., and Metoyer, R. A. Studying always-on electricity feedback in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, ACM (New York, NY, USA, 2010), 1995–1998.
- [10] Seryak, J., and Kissock, K. Occupancy and behavioral affects on residential energy use. In *Proceedings of the Solar Conference*, ASES, AIA (2003), 717–722.
- [11] Szewczyk, S., Dwan, K., Minor, B., Swedlove, B., and Cook, D. Annotating smart environment sensor data for activity learning. *Technology and Health Care* 17, 3 (August 2009), 161–169.
- [12] U.S. Energy Information Administration (EIA). Annual energy review 2011. Tech. Rep. DOE/EIA-0384(2011), Office of Energy Statistics, U.S. Department of Energy (DOE), September 2012. Forrestal Building, Room 1E-210, 1000 Independence Avenue, S.W., Washington, D.C. 20585.
- [13] Wilson, C., and Dowlatabadi, H. Models of decision making and residential energy use. *Annu. Rev. Environ. Resour.* 32 (2007), 169–203.