Influencing Participant Behavior Through a Notification-Based Recommendation System

Venkata Reddy[1], Brian Bushree[2], Marcus Chong[2], Matthew Law[3], Mayank Thirani[2], Mark Yan[2], Sami Rollins[2(✉)], Nilanjan Banerjee[1], and Alark Joshi[2]

[1] University of Maryland, 1000 Hilltop Cir, Baltimore, MD, USA
[2] University of San Francisco, 2130 Fulton Street, San Francisco, CA, USA
srollins@cs.usfca.edu
[3] Cornell University, Ithaca, NY, USA

Abstract. Behavioral recommendations for achieving energy savings in the home are extremely common, however how to effectively influence users to adopt such recommendations is not well understood. In this work, we present the results of a feasibility study, conducted over a 4-week period, that deployed a phone-based recommendation system designed to encourage participants to follow the popular utility-company recommendation: *Consider dimmer switches to adjust the light to the lowest level necessary for an activity*. We found that the system did influence participants to follow the recommendation and some even realized that they preferred dimmer lighting, suggesting that recommendation systems can serve to demonstrate to participants that they can maintain comfort even with lower energy consumption levels.

1 Introduction

Recommendation-based approaches for influencing user behavior toward energy savings are common, however most are either manual—relying on the user to apply a static recommendation—or automatic—for example, the Nest thermostat. Manual solutions may inconvenience the user, while automated systems often cause frustration. In this work, we report results of a feasibility study of a hybrid, in-situ recommendation system to provide cues to help participants follow a popular utility company recommendation - *Consider dimmer switches to adjust the light to the lowest level necessary for an activity.* Our system measures the brightness of a light bulb and uses a heuristic to determine whether the bulb is brighter than necessary. If so, a recommendation to dim the bulb will be delivered via a phone notification. Our study was conducted over a four-week period during which we collected data on the light bulb usage and, after which, we conducted in-person interviews with the participants from our study.

Our results show that our recommendations *did* influence participant behavior and increased their awareness towards the use of a smart home automation systems. Additionally, we show that participant-in-the-loop systems can

J. Ham et al. (Eds.): PERSUASIVE 2018, LNCS 10809, pp. 113–119, 2018.
https://doi.org/10.1007/978-3-319-78978-1_9

avoid frustrations caused by automated smart home and energy saving systems. Finally, we demonstrate that the content of the recommendations and their frequency of delivery are two key attributes that must be properly designed for a recommendation system to be useful in the context of home energy management systems.

2 Related Work

Our paper builds on previous work on home energy systems. Several papers study which recommendation attributes are effective for long term adoption. For instance, Abrahamse et al. [1] study several types of interventions: goal-setting, information tailoring, modeling, and feedback. Allcott [2] study personalized recommendations based on historical usage patterns and demographics. Castelli et al. [3] prototype a context-based recommendation system on a smartphone, and Costa and Kahn [4] and Gonzales et al. [5] discuss the effectiveness of nudging and social cognition and persuasion in increasing the effectiveness of recommendation systems. Most of the research in this field is based on mock prototypes and does not involve any real end-to-end system deployment and evaluation.

A group of previous work focuses on designing automatic energy management systems. The approach taken is to model energy consumption behavior using machine learning techniques and predict future energy usage [6,7]. Based on our results, we find that such automatic systems can frustrate the participants and, while human-in-the-loop system are more cumbersome to use, they are more effective in changing energy use behavior.

3 Study Setup

Our goal in this work is to explore whether a lighting recommendation-based system is a plausible approach for influencing energy decisions. We focus on one appliance—a dimmable light bulb—and our findings focus on the benefits and limitations of this type of system. Our goal is *not to quantify energy savings*, but to better understand whether this type of system can influence behavior.

Our system measures the brightness of a dimmable light bulb and calculates an ideal level based on the time of day and local weather—other contextual cues like activity performed by the user are not considered in this study. The system then sends a phone notification recommending that the participant reduce the brightness of the bulb. As shown in Fig. 1, the system uses the Philips Hue infrastructure to measure and control a light bulb, which interfaces with a Node.js API through a Java client running on a Raspberry Pi 3. For the purposes of our study, the system records the brightness level of the bulb every 30 s, which allows us to determine whether the participant accepted or overrode any given recommendation.

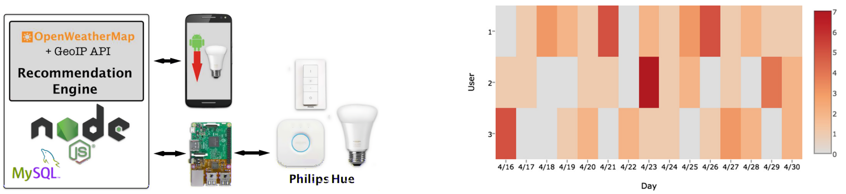

(a) Client side System (b) Brightness Override Frequency

Fig. 1. (a) The client-side system consists of a Raspberry Pi 3 interfaced with a Philips Hue light bulb via a wireless Hue bridge. (b) This heat map represents the brightness override frequency for 'automatic' participants in Phase 1. Each row represents one participant and the intensity of the color for a segment the number of times the participant overrode the system on that day. (Color figure online)

Study Design: We performed a small-scale deployment over four weeks with eight participants (three females) who received a Philips Hue bulb, a physical dimmer switch, a Raspberry Pi and, if necessary, an Android phone. To study interactions between recommendations and automation through our system we designed the following four treatments: (1) **Manual** - no recommendations are communicated to the participant. Participant uses a physical dimmer switch to manually adjust the brightness of the bulb. (2) **Automatic** - recommended light levels are sent directly to the participant's lighting system, which *automatically* adjusts the brightness of the bulb. Participants may override the setting using the physical dimmer switch. (3) **Recommendation+Manual** - recommendations are sent to the participant's phone. To apply a recommendation, the participant must manually adjust the light using the physical dimmer switch. (4) **Recommendation+Automatic** - recommendations are sent to the participant's phone, and s/he may use the phone to accept the recommendation. If accepted, the system will automatically adjust the brightness of the bulb to the recommended level.

The study was organized into two-treatment within-participant tracks (automatic and manual), each of which introduced recommendations in the *second* phase. Thus, one group successively underwent treatments (1) and (3) mentioned above, and the other (2) and (4), as described below.

Phase 1: In the *manual* group, five participants were given the equipment and encouraged for the first phase to use the physical dimmer switch to adjust the brightness of the lamp to a comfortable level. No further instructions or recommendations were provided. The remaining three participants (in the *automatic* group) were informed that their lamp would automatically adjust to the recommended level, which they could override at any time.

Phase 2: In this phase, participants installed either the *manual* (Treatment 3) or *automatic* (Treatment 4) version of the recommendation application on the phone. At the conclusion of Phase 2, each participant participated in a 15-min in-person qualitative interview during which we collected feedback on aspects of the system such as instances when they followed or did not follow a recommendation, whether they would continue to use such as system, and if they had any recommendations for us.

4 Findings

Insight 1: Recommendations may influence participant behavior either directly or indirectly by increasing their contextual awareness.

We observed changes in participant behavior due to the introduction of recommendations in **both** manual and automatic groups, however we observed a fair amount of variance. Table 1 displays the mean light bulb brightness level (from 0 to 255) for each participant during Phase 1 (no phone recommendations) and Phase 2 (recommendations). We exclude readings below light level 10 as one participant appeared to have turned her light to the lowest brightness setting rather than completely off in Phase 1. We observed that over 78% of the dim events occurred within 20 min of a notification, implying a probable causation between the participants dimming the bulb and receiving a notification.

Table 1. Mean bulb brightness and number of participant dim events during Phase 1 and Phase 2 of the study. **Dim Events** are when a participant adjusted their light by following a recommendation or manual. Participant #7 encountered technical difficulties and was unable to complete Phase 2.

Group	ID	P1 mean level	P2 mean level	P1 dim events	P2 dim events
Automatic	1	156	216	n/a	3
	2	183	198	n/a	2
	3	194	199	n/a	8
Manual	4	157	96	4	2
	5	254	254	0	0
	6	252	234	0	2
	7	210	n/a	4	n/a
	8	198	184	7	11

The observed changes in light levels suggest that recommendations could affect participant behavior in ways that could either reduce *or increase* energy usage. In the manual group, participants almost uniformly exhibited a *reduced mean light level* in Phase 2. In the automatic group, surprisingly, we observe an increase in mean light level for all three participants. This is particularly interesting, considering participants had the choice to override automatic light levels at any time. Nonetheless, frequent participant overrides (Fig. 1(b)) hint at the possible influence of an over-aggressive dimming algorithm, with at least one participant expressing gratitude for the control afforded by the recommendations. It is worth noting that, even with the increase, light levels were still below the maximum (255) for all three participants in Phase 2 (when they received recommendations). Overall, these two seemingly discordant trends could suggest that recommendations offer participants opportunities to discover ways to reduce consumption, while giving them the freedom to adapt consumption more closely to their needs than an automatic system.

In the post-study interviews, two participants cited the potential of the system to provide contextually meaningful nudges. One participant remarked, *"The chiming sound of the recommendation otherwise was useful because it reminded me that oh maybe I don't need it to be this bright."* Another participant observed that, when he changes context, i.e. finishes studying and begins relaxing, the brightness of the bulb may not be on his mind but when he receives the notification it reminds him *"... I don't need that much light at that time"*.

Three of the eight participants reported preferring dimmer lights as a result of their experience in the study, suggesting that timely recommendations can serve to demonstrate possibilities for lower consumption without discomfort. One participant observed, *"After using this device I prefer using a dimmer switch so I can set it to a [...] low setting, especially at night."*. Another participant noted, *"I kinda got used to it, and once your eyes adjust to it, it really doesn't make too much of a difference."*. Despite the possibility of some initial discomfort, this participant continued, *"I wish that I had dimmers on more of my lights, just because when you are aware of the fact [...] a little less power getting consumed and if I could do this with all of my lights I definitely would."*.

Insight 2: Participant-in-the-loop systems can help to avoid participant frustration caused by automatic solutions, but must be carefully designed to avoid inefficiencies.

Our follow up interviews confirmed that two of three automatic participants did prefer the automatic system, however one participant's comments suggested that the automatic system may not be widely accepted. He expressed frustration with the automated system, noting it was not *"what [he] wanted"* and further explained, *"I think I liked phone recommendation more 'cause [...] I have access to say yes or no."* Automatic participants' override behavior (Fig. 1(b)) supports this observation. Even though one participant reported infrequently overriding our settings, there were several overrides most days for all three participants. This raises the concern of long-term frustration with the system.

Though a participant-in-the-loop system may overcome many of the challenges of an automated system, care must be taken that it does not enable self-defeating usage decisions. In our study, most reported behavioral changes were positive, however some participants reported behavioral changes that could increase energy usage. One participant noted, *"I probably had the light on more or later in the night"* due to the ability to choose a lower setting when he didn't want full brightness. Without a baseline for typical usage, it is difficult to quantify if this participant's overall consumption dropped or increased as a result, but it is reasonable to imagine how increased usage could offset brightness reductions.

Insight 3: The recommendation content, frequency and delivery mechanism are critical system design components.
The content of the recommendation itself is critical for behavioral change. One participant reported that our simple notification—We recommend you dim your light—*"...didn't really motivate me to want to override my choice."* and another stated, *"I want to know more information, like how much I need to dim the light."* Another participant reported feelings of guilt when he chose not to follow the recommendations noting that he was, *"nervous when the notification came... It made me feel bad."* These observations point to the need to provide motivation, such as potential savings in terms of energy or cost, as well as the potential for using positive reinforcement rather than negative reinforcement to suggest behavioral change.

The frequency and timing of notifications should be carefully tuned to avoid participant frustration. Participants complained that recommendations were made too frequently (every 30 min) or too quickly after the light was switched on. One participant requested a "busy mode" to pause recommendations for a longer period of time. Several participants also commented on the fact that a recommendation would happen immediately after the light was switched on. One participant even began to distrust the system due to the immediate notifications, declaring that *"...I don't know whether that is, that is more smarter decision or it just popped out no matter, every time you just open the light."* There is a tradeoff between frequent notifications that may frustrate the participant, and infrequent notifications that may miss opportunities to save energy.

A feature-rich phone application is a good choice for delivering recommendations, however variance in participant preferences suggests the need for a customizable delivery mechanism. Five of eight participants responded positively to a question asking them whether phone-based notifications were their preferred option. One participant would have preferred to view and interact with his data on the phone app, while others wanted to have more control of the light from the app. The majority of the six participants that were loaned Android phones asserted that they would have preferred to use their personal phone. Participants had other suggestions for recommendation delivery such as an ambient sound, fixed device near the light, and an Xbox. Overall, the diversity of desires articulated by participants in response to our system point to the need to provide multiple options for participants and allow personalization based on preference.

5 Conclusions

In this paper, we present the results of a deployment of an energy-based recommendation system and its feasibility for changing energy usage behavior. We show that human-in-the-loop energy saving systems can influence human behavior. However, it is important to design feature rich recommendations and carefully control the frequency of generating these recommendations. Future directions for this work include studying how personal characteristics, for example vision, affect how the system is used; how the system would influence behavior over time; and the effect the system has on energy use.

References

1. Abrahamse, W., Steg, L., Vlek, C., Rothengatter, T.: A review of intervention studies aimed at household energy conservation. J. Environ. Psychol. **25**(3), 273–291 (2005)
2. Allcott, H.: Social norms and energy conservation. J. Public Econ. **95**(9), 1082–1095 (2011)
3. Castelli, N., Stevens, G., Jakobi, T., Schönau, N.: Switch off the light in the living room, please!-Making eco-feedback meaningful through room context information. In: EnviroInfo, pp. 589–596 (2014)
4. Costa, D.L., Kahn, M.E.: Energy conservation "nudges" and environmentalist ideology: evidence from a randomized residential electricity field experiment. J. Eur. Econ. Assoc. **11**(3), 680–702 (2013)
5. Gonzales, M.H., Aronson, E., Costanzo, M.A.: Using social cognition and persuasion to promote energy conservation: a quasi-experiment. J. Appl. Soc. Psychol. **18**(12), 1049–1066 (1988)
6. Ozturk, Y., Senthilkumar, D., Kumar, S., Lee, G.: An intelligent home energy management system to improve demand response. IEEE Trans. Smart Grid **4**(2), 694–701 (2013)
7. Choi, J., Shin, D., Shin, D.: Research and implementation of the context-aware middleware for controlling home appliances. IEEE Trans. Consum. Electron. **51**(1), 301–306 (2005)