Engineering vs. Ambient type visualizations: Quantifying effects of different data visualizations on energy consumption

Lucas Spangher*

lucas_spangher@berkeley.edu

University of California at Berkeley, CA 94709

Alex Devonport

alex_devonport@berkeley.edu

University of California at Berkeley, CA 94709

Akaash Tawade

akaasht@berkeley.edu

University of California at Berkeley, CA 94709

Costas Spanos

spanos@berkeley.edu

University of California at Berkeley, CA 94709

ABSTRACT

Engaging consumers with energy use is a challenging endeavor; however, different forms of data visualization may help by creating relationships between a user and their data. Here we propose and conduct an experiment to test two different visualizations: an “engineering-type” visualization of a bar chart, and an “ambient-type” visualization of a cartoon tree. Participants in a competitive energy-saving game within an office building in Singapore received weekly emails visualizing their energy consumption and encouraging them to interact with their online game profile, and energy use was measured in addition to qualitative feedback through pre-experiment and post-experiment surveys. We find that the engineering-type visual did not result in statistically significant reductions whereas the ambient-type did, resulting in an average decrease of 82 Wh per day for the three days following treatment. Additionally, we find that greater engagement uniformly correlates positively with energy savings. A finite state machine simulation indicates that were this experiment scaled to 1000 people over 100 days, the total energy savings might be expected to be 1.1MWh.

*Corresponding author

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UrbSys'19, November 13–14, 2019, New York, NY, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-7014-1/19/11...\$15.00

<https://doi.org/10.1145/3363459.3363527>

CCS CONCEPTS

• **Hardware** → **Smart grid**; *Enterprise level and data centers power issues*; Impact on the environment; • **Applied computing** → Multi-criterion optimization and decision-making.

ACM Reference Format:

Lucas Spangher, Akaash Tawade, Alex Devonport, and Costas Spanos. 2019. Engineering vs. Ambient type visualizations: Quantifying effects of different data visualizations on energy consumption. In *1st ACM International Workshop on Urban Building Energy Sensing, Controls, Big Data Analysis, and Visualization (UrbSys'19)*, November 13–14, 2019, New York, NY, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3363459.3363527>

1 INTRODUCTION AND BACKGROUND LITERATURE

Identifying and eliminating plug-load waste is a challenging endeavor— not only due technoeconomic limitations of plug load sensing, but also limitations in a central planner’s ability to influence individual behavior. Perhaps because of this, the overall energy demands in buildings is increasing. In residential and commercial buildings, plug loads represent 30% of total electricity use ([1], [2]). In addition, the quantity of energy used by plugs is increasing more quickly than any other load type in both residential and commercial buildings [3].

To this end, notable work has been performed in creating and administering “Social Games” – for our purposes, defined as competitions around energy use. Social Games tend to comprise of: (1) informing each player of their energy use in a friendly and easy to consume manner and (2) an accompanying competition in which higher energy savings relative to others is rewarded. Many studies have either created Social Games themselves ([4], [5], [6], [7]) or studied existing Social Games ([8], [9]) to draw insight. To this end, the general finding has been that Social Games increase the

extent to which an individual is motivated to and able to save energy in their daily functioning within their office or home.

Data visualizations as they relate to energy have also been studied extensively ([10], [11], [12]). The literature generally taxonomizes the space of visualizations into two main types that we consider here [13]: (1) “engineering-type” visualizations composed of barcharts, line graphs, scatterplots – i.e. a formal, work-like presentation of the data in a way that would be used in scientific papers and (2) “ambient-type” visualizations, in which a linear scale is communicated by some abstract, pleasant, artistic visuals¹. Examples of an “ambient-type” visualization would include the so-called “power-aware” cord, a power cord that glows more brightly the more energy is being used [14], a “thrifty faucet” which shines a red or blue light depending on the temperature of the water [15], or an ambient battery communication system that communicates the battery’s charge with a proportional intensity of light [16].

The consensus in the literature is centered around the idea that ambient visualizations are more effective than engineering type visualizations at consistently communicating to a user their energy usage ([17], [18], [19], [20], [21]). The effect has been studied qualitatively. However, we are unaware of a study that attempts to quantify the behavioral differences that types of data visualizations engender. Therefore, we aim to reproduce a Social Game experiment and quantify the difference in energy consumption from participants exposed to different types of visualizations.

The rest of the paper is organized as follows: in Section 2, we will first describe a pre-treatment survey, which attempts to characterize the population being studied. We will then describe the experimental setup and execution. We will describe an overview of the post-treatment survey. We will then describe a Hidden Markov Model that simulates the experiment if it were extended to many more people. In Section 3, we describe the results of all four segments of this study. We will finish with concluding remarks and future directions.

2 METHODS

2.1 Pre-treatment survey

We conducted a pre-treatment survey to identify trends in the social game population as it pertains to energy literacy and climate change opinions. The survey also collected demographic information from participants, including age, ethnicity, and job position. Our objective was to explore

¹There is a third type, the (3) “natural type” visualizations, in which images closely match the natural world and map energy onto environmental impact. However, it has been generally found that this type of visualization invokes guilt in the viewer and is thus less often used.

the extent to which the social game population was heterogeneous in their demographic, energy literacy, and beliefs about climate change. The survey was emailed through a Google Form to all participants six weeks before treatments were first administered. 18 responses to the survey were collected of the 28 individuals studied, representing a 64% response rate. In order to look for relationships between question responses, we used exploratory data visualizations and multiple regressions.

2.2 The Social Game experiment

2.2.1 Experimental setup. In this section, we will briefly describe the setup for the experiment that was conducted.

We administered a Social Game previously designed and presented in the literature [4] in the Campus for Research Excellence And Technological Enterprise (CREATE) tower, a building in Singapore. This consisted of an online platform where participants could monitor their progress in saving energy relative to others, receive tips on how to change their energy saving behavior, and, crucially, to schedule times in which certain plugs at their desk would be on or off. Office workers were voluntarily enrolled in the Game at the start of August, 2018.

The Game was structured in the following manner: first, a normal distribution was fit to the historical energy data of each participant to make energy savings relative to their own baseline. Then, for five periods of two weeks each, participants competed against each other to reduce their energy consumption. Amongst the top five of every round, a random number generator picked the first, second, third and fourth place winners, all of whom received prizes of various amounts.

2.2.2 Experimental Treatment. We were interested in exploring the effect of different types of data visualizations on the user’s engagement with the system and energy savings. To this end, two types of visualizations were created: an engineering-type bar chart (please see Figure 1), and an ambient-type visualization (please see Figure 2). The ambient-type was intended to communicate a sense of linear scale without appearing too formal. Although we developed it in house, we relied heavily on the tree design presented by Froelich et al in their 2009 ACM CHI Conference on Human Factors in Computing Systems Proceedings [22]. All participants were emailed a treatment once a week, on either a Wednesday or a Thursday.² The emails consisted of Treatment A, which included the barchart and the following language:

²There was no pattern in whether a Wednesday or Thursday was chosen; it was due to technical limitations on the time of our software engineer.

Figure 1: The “engineering type” visualization sent to Social Game participants as a treatment.

Thank you for your participation thus far in the Social Game!
 Through participating in this game, your actions are being studied to learn more about how energy is used in office settings.
 Your ranking so far is:
 __ of 27.
 You progress with regards to last week and your baseline use is as shown:

Treatment B consisted of the same language and an image of the tree visualization, and finally a control email consisted only of the language.

The curious reader might wonder why we opted to include language in the control instead of simply not sending an email at all to a control group. Our experimental setup was modeled after the experimental setup noted in [23]. Here, the authors send a control postcard noting that the participants are being studied in order to deal with the *Hawthorne Effect*, a noted positive effect that occurs simply because the subject is aware of her being studied. Therefore, comparing the effects of a control whose email differs from a treated subject only by a visualization, we have in the difference of the responses the effect of the visualization itself.

The treatments are labelled as follows. Treatment A is an engineering-type visualization with the above language. Treatment B is an ambient-type visualization with the above language.

Participants were randomly selected to receive either Treatment A, Treatment B, or a control. Assignments switched

Figure 2: The “ambient type” visualization sent to Social Game participants as a treatment.

three and six weeks following a step-wedge assignment protocol.

2.3 Post experiment survey

Following the conclusion of the Social Game, all participants were sent a questionnaire seeking feedback on the energy visualizations and the competition’s impact on their energy consumption habits. Responses were compared with findings from the experimental data to analyze trends between the qualitative participant feedback and quantitative outputs from Social Game. Results from the post experiment survey were also used to frame the future direction of our experimentation.

2.4 Simulation

Our ultimate goal is to make predictions about the effect of visualization type on total energy usage at the scale of an entire building, using the energy usage data we have collected from a small number of participants.

One way to use individual energy data to make predictions about energy use at the building scale is to construct a probabilistic model for individual energy use based on the individual energy data, and “simulate” the energy use of a building’s worth of participants.

Here we introduce one such model, a Hidden Markov Model (HMM) for an individual’s energy use per unit time

Figure 3: Diagram of the HMM model for individual energy use. The hidden states E and B represent whether the individual is following their energy-saving distribution (E state) or their baseline distribution (B state). The observations, all normally-distributed in this model, correspond to the energy use of the individual per unit time.

(daily, for instance) [24]. This model is intended to be as simple as possible while capturing the effect that the visualization has on individual energy use.

Our model assumes that the energy use of an individual can be modeled at least roughly in the following way. In the absence of any kind of treatment, we assume that the individual's energy use per unit time follows some probability distribution, which is their *baseline distribution*. Some time after receiving treatment (that is, the energy visualization), we assume that the individual will modify their energy use in such a way that their energy use follows a new distribution, which is their *energy-saving distribution*. After spending some time in the energy-saving distribution, the individual eventually resumes their baseline distribution, and remains there indefinitely.

With this in mind, our model is an HMM with Gaussian observations, shown in Figure 3. The states correspond to whether the user is following their baseline distribution (state B) or their energy-saving distribution (state E). An execution of the HMM model represents a trace of the modeled individual's energy use over some time period. Since we have collected individual energy use over time for several participants, we can use this data to fit model parameters for each individual using an EM algorithm.

Under this model, how much energy does an individual save when they are influenced by the treatment? Without the treatment, the individual would remain in the baseline state the entire time, so the difference in energy use effected

by the visualization comes entirely from the time spent in the energy-saving distribution. Assuming that the individual enters the E state once per treatment, and otherwise remains in the B state, the number of time units spent in the E state n_E is distributed geometrically, that is $n_E \sim \text{Geom}(p_{EB})$. As such, the total energy used during the user's duration in the E state is a random sum of random variables, which we will denote as S_E . By Wald's equation, the expected energy use during the energy-saving state is

$$E[S_E] = \frac{1 - P_{EB}}{P_{EB}} \mu_E.$$

Had there been no treatment, the individual would have spent the same amount of time in the baseline distribution, using an amount of energy S_B , whose expected value is

$$E[S_B] = \frac{1 - P_{EB}}{P_{EB}} \mu_B.$$

Therefore, the expected energy savings effected by the visualization for an individual is

$$E[S_B - S_E] = \frac{1 - P_{EB}}{P_{EB}} (\mu_B - \mu_E).$$

For a more detailed look at the behavior of this model, in particular the effect of having models of multiple participants each with unique parameters, we turn to simulations, which we discuss in section 3.

3 RESULTS

Here we discuss the main results of the experiment.

3.1 Pre-treatment survey

Distributions of scores on the energy literacy questions (figure 4a) indicate a varied knowledge of climate related topics, though average participant scores were somewhat respectable at 68%. When pressed on whether Singapore is managing its energy correctly, participants responded across a uniform distribution centered at 6.4 points out of 10 (figure 4b). This result indicates a spread in faith in the energy management of the Singaporean government.

We completed a regression analysis on bivariate relationships in the data and present the following two takeaways. First, those who believed that the world should take collective action on climate change tended to also value scientists' portrayal of climate change as accurate ($p=0.001$) (figure 5a). Second, those who believed Singapore is managing its energy correctly also tended to view it as a responsible actor in the world's climate goals ($p=0.005$) (figure 5b).

Our results concur with other literature on Singaporeans' energy literacy and climate change opinions. Rosenthal and Ho note that most Singaporeans (84%) view climate change as somewhat or very serious [25], which agrees with our observation that respondents tended to place an importance on

(a)

(b)

Figure 4: (a) Distribution of scores on energy literacy test survey questions amongst participants (out of 11 points) (b) Distribution in opinion on whether Singapore is managing its energy correctly amongst participants (10 indicates highest agreement)

the necessity for global climate action. Interestingly, a 2016 survey conducted by Singapore's National Climate Change Secretariat (NCCS) found that most Singaporeans strongly believed that climate action should be driven by the government as opposed to individual actions [26].³ The NCCS results may explain the correlation in our survey between positive appraisals on Singapore's energy management and how

³Many attribute this trend to the so-called "nanny-state" syndrome, whereby years of state intervention have created a dependency on the government to address climate issues.

(a)

(b)

Figure 5: (a) Relationship between participant opinion on whether the world should take collective action on climate change and whether scientists' portrayal of climate change matches the phenomenon (10 indicates highest agreement) (b) Relationship between participant opinion on whether Singapore is managing its energy correctly and whether Singapore is responsible actor in the world's climate goals (10 indicates highest agreement)

responsible the Singaporean government is in the world's climate goals.

3.2 Experimental treatment

This section describes the results of an extensive exploratory analysis. First, we show a display of the data in figure 6. There are several items of note herein, which we will note and

Figure 6: Summary of the main results from the experiment. Here we show averages of groups per day across the week, with the yellow section highlighting the part of the week which occurred after the treatment. The x-axis, “centered time around treatment” is the days translated such that the treatment day always occurs on day 0.

support through regression analysis: (1) working days are generally indistinguishable from each other, (2) people tend to engage with the Game less, even though energy savings persist over the duration, (3) greater logins to the online platform correlated with greater energy savings (4) numerous predictor variables and interactions significantly predict daily aggregate energy use (E), which we present in a regression output shown in Figure 7 (section 3.2.4).

3.2.1 Weekend Indicator: Uniformity of weekday vs. weekend. First, we observe that the working days of the week (i.e. Monday through Friday) are visually indistinguishable from each other with respect to E, yet significantly different than weekend usage. This was further confirmed in a regression analysis predicting E; a categorical variable treating each day of the week as an independent predictor found four of the five weekdays to have insignificant⁴ predictive power with respect to the response, after the variation from the first weekday was controlled for.

3.2.2 Treatment Number: Effect of repeated treatments. Second, we observe that participants login less after multiple treatments, but save energy consistently throughout the experiment. We tested multiple treatments by coding a “Treatment Number” (TN): an integer value that counted the number of times an individual had been treated. TN returned

⁴When we say a predictor was insignificant in predicting a response, we use conventional definitions of 5% significance in a frequentist Ordinary Least Square (OLS) framework.

as significantly negative when included in a regression predicting a users’ Login Count (LC), which was an integer value counting how many times a user logs onto their online profile. The significant relationship implied that there was a significant decrease over time with how often users interacted with the Game. However, despite this, TN was insignificant when predicting E across time, which implies that the energy habits that the Social Game encouraged persisted throughout the treatment period.

3.2.3 Login Count: Effect of interaction with the platform. Third, we observe a positive relationship between LC and E. A safeguard was put into place to ensure that LC was not unrealistically inflated; the count could increment only once every five minutes. Regressions showed that Login Count robustly predicted E when controlling for a variety of factors such as weekday vs. weekend and TN. As expected, the regression coefficient in LC was negative, implying that the more a user logged into their system, the more energy they tended to save.

3.2.4 Main treatment effect on energy savings. Finally, we observe numerous significant predictors of E, which, again, is our main response of interest, measuring daily aggregate energy from each plug load of each desk. Given that E is a periodic time series, we segmented the resulting data into three days before treatment and three days after treatment. We include the previously noted variables, testing for a multiple regression of the following form:

$$E \sim TN + WI + LC * TT + TI : TT + \text{out of office} \quad (1)$$

Here, E is predicted by TN, WI, an interaction between LC and treatment type (TT: i.e., A, B, or control), and an interaction between TT and “Treatment Indicator” (TI). TI is an indicator variable that is 0 for three days before each treatment and three days after, which measures the average energy use that persists after a few days. Note that the equation follows standard R syntax, where * is an interaction with terms considered independently, and : is an interaction with only bivariate terms considered.

The first interaction, LC:TT, attempts to elucidate an effect between people’s LC and their TT that is above the average effect of LC; i.e., does a different type of treatment engender different impact from interacting with the system? The second interaction, TI:TT attempts to provide the average treatment effect for different treatment types several days after treatment. The final term, OO, (“Out of Office”), is an indicator that keeps track of whether or not a person was physically in the office, which we guessed from looking at distributions of energy demand.

Results of this regression conform to our hypotheses. TN is insignificant, which implies that users’ energy savings

Dep. Var:	Energy	R-sq:	0.284
Model:	OLS	Method:	Least Sq
No. Obs:	1223	Df Resid:	1213
Df Model:	9		
	coef	std err	P> t
Intercept	599.27	34.03	0.00
TN	-4.43	5.54	0.42
WI	77.73	34.157	0.023
LC	-71.98	18.30	0.00
LC:TT[Ambient]	22.21	29.38	0.45
LC:TT[Control]	-64.17	56.52	0.26
TI:TT[Engineering]	34.93	27.75	0.21
TI:TT[Ambient]	-82.67	37.07	0.03
TI:TT[Control]	19.58	54.56	0.72
OO	-610.19	33.20	0.00

Figure 7: Summary of the OLS regression results.

does not majorly diminish during the extent of the game. Whether or not a day is a weekend significantly predicts E, as we expect: during the weekends, desks use 77Wh less energy than during weekdays after controlling for the other factors. LC remains significantly negative in this regression; for an average increase of one login, we expect energy use to be approximately 71 Wh less per day. The TT interacted with LC, meanwhile, do not significantly predict total energy, implying that the treatments do not effect the relationship between interacting with the platform and reducing energy use. The TT interacted show a strongly significant negative effect of treatment B: for the three days following treatment B, we can expect on average 82 Wh less energy use after controlling for other variables. When people are out of the office, we expect 610 Wh of energy use less (this variable is highly collinear with, but not entirely explained by, WI.)

3.3 Post experiment survey

Through the post experimental survey, we endeavored to explore whether the quantitative results we measured were perceived by the participants.

First, we queried the participants on the negative effect of TN on LC (see section 3.2.2): people seemed to interface less with the game through repeated treatments. When asked whether interaction lessened over time, two thirds of the respondents answered "Yes", which confirms the quantitative findings.

Second, we queried the participants on our observation that the tree visualisation seemed more engaging. Of the tree, a respondent stated: "the tree visualization provided... insight and view on energy usage." Meanwhile, of the bargraph, a respondent stated: "It was clear but I did not feel it to have any

distinctly engaging quality." We received several comments similar to the ones shown above.

Finally, we received confirmation that the format of the treatment was sound, addressing a concern that the treatment emails may have often been left unread. When asked how often they opened the weekly email (the treatment), 83% of respondents said that they opened it "all the time" or "almost all the time". This was significantly different from responses to being asked how important the email was to their interactions with the Game: responses were uniformly distributed across a five point scale from "unimportant" to "important", confirming that while the emails may not have been impactful, at least that the method of treatment delivery was sound.

3.4 Simulation

We will now argue that a scale up in simulation is appropriate. Our test population reflects the wider Singaporean office worker population enough that behavioral dynamics from the test population can be seen as generalizable. This is important because individual energy savings only significantly improve sustainability when the savings are part of a collective shift in behavior.

We claim generalizability for a number of reasons. First, the Social Game was run in a general office with standard desk equipment⁵ and a normal 40 hours per week working schedule (Monday-Friday). Second, the distribution of energy literacy's width (see section 3.1) indicates that the participants are not uniformly experts on the subject — as one might expect in a general office. Third, baseline energy consumption varied widely amongst participants, indicating a spread in energy consumption as well.

We ran our simulation with 1000 simulated actors across 100 days. In total, the energy savings were estimated to be 1.1 MWh, or approximately 1 kWh a person over these days. Each person was in the "energy saving state" for a different amount of time, so producing a estimate of savings per day would be a less meaningful task.

4 DISCUSSION AND CONCLUSION

Our experiment demonstrates significant statistical evidence that the tree visual, an ambient-type visual, is effective in curtailing energy use. We see that receiving the ambient-type visual leads to a significant decrease in energy use for multiple days following treatment.

4.1 Limitations

This experiment presents a start to the quantification of the effects that different visualizations can have. However, we

⁵Laptop or desktop computer, 1-2 monitors, desk lamp, etc.

Figure 8: Distribution of participant weekday baseline energy consumption in log(kWh)

caution against generalizing the results too widely. Here we present some limitations to our work.

The first limitation we discuss is that the bar chart may not be representative of the range of engineering-type visualizations. Indeed, we claim here that the bar chart is more representative of an average to low level of engagement that engineering-type visualizations may invoke rather than an extreme level of engagement. Modern engineering-type visuals incorporate interactivity, exciting or smooth colors, and more intentional design choices [14].

However, the ambient type visualization that we tested is likely also in the average to low level of engagement that ambient-type visualizations might invoke. It was also not professionally designed, nor was it interactive, dynamic, or shaded. To quantify exactly where either visualization falls on the level of engagement within its own type is beyond the scope of this work.

The second limitation we discuss is the unique framework of the treatment *within* the larger Social Game. Whether the treatment would be more effective outside the Social Game, or whether the population being relatively energy aware dampens the effect of the treatment is impossible to say. We caution against the generalizability of the findings for this reason.

The third limitation that we discuss is the limitations in the Hidden Markov Model (HMM) simulation. We intend it as a rough estimate to show how much energy might be saved if this were to be deployed at scale. However, a two state representation of human energy use might be entirely insufficient to capture the ways humans use energy. In addition, the dynamics of a Social Game might change drastically

when scaled to 1000 people as opposed to 27 people, and so it is unclear whether people would actually save more or less energy than the HMM implied.

4.2 Applications and future directions of these findings

There are many ways in which this finding can be applied, and future questions we are interested in studying. We discuss several of them below.

We show that insights from the pre-treatment survey can be used to explore heterogeneity in the social game population. The survey is also used as a basis for establishing generalizability. We imagine future social games to be preceded with greater pre-treatment survey depth on energy literacy and behavioral tendencies. We hope this establishes a better basis to individualize interventions by tying pre-treatment questions to desks and understanding more about what drives people's responses.

We show that weekly exposure to an email containing an ambient-type visual is effective in increasing energy savings. How would this change if the exposure were more permanent? Our visualization could be extended into an inexpensive display (either in a desktop or in a building lobby), via some simple engineering of a Raspberry Pi or through an e-ink platform. We are, indeed, interested in studying whether this would create a lasting impact on the relationship that people develop with their energy usage and data.

We show that individual weekly exposure to an ambient-type visual effectively increases energy savings. How would this change if participants were grouped together in teams? Would energy savings increase as participants pool together knowledge and strategy in order to advance their team's collective standing? How could we optimize team formation and Game scoring in order to increase all individual's collective input? We are interested in studying further the effects of team-based dynamics.

We show that a static visual is effective in increasing energy savings. How would a dynamic visual effect energy savings? We imagine several possibilities. First, we propose the creation of an AI powered avatar that interacts with users; it might encourage user's energy savings and offer suggestions for new and creative ways to save energy. Second, we propose a generative model that produces new and unseen trees and might be individually tailored to a person. This model might take the form of a Generative Adversarial Neural Network (GAN) or a simpler convolutional network (CNN).

5 ACKNOWLEDGEMENTS

We would like to acknowledge the generous support of the Singapore–Berkeley Building Efficiency and Sustainability

in the Tropics (SinBerBEST) program, which is funded by the National Research Foundation, Prime Minister's Office, Singapore. In addition, we would like to especially thank and acknowledge Ioannis Konstantakopoulos [4] and Jon Froehlich [22] for designing the framework of the Social Game and our ambient visualization, respectively.

REFERENCES

- [1] S Lanzisera, S Dawson-Haggerty, H. Y. I. Cheung, J Taneja, D Culler, and R Brown. Methods for detailed energy data collection of miscellaneous and electronic loads in a commercial office building. *Building and Environment*, 65:170–177, 2013.
- [2] R. S. Srinivasan, J Lakshmanan, E Santosa, and D Srivastav. Plug load densities for energy analysis: K-12 schools,. *Energy and Buildings*, 43:3289 – 3294, 2011.
- [3] O Comstock and K Jarzomski. Consumption and saturation trends of residential miscellaneous end-use loads. *ACEEE Summer Study on Energy Efficiency in Buildings*, Pacific Grove, CA, USA, 2012.
- [4] I. C. Konstantakopoulos, L. J. Ratliff, M. Jin, and C. J. Spanos. Leveraging correlations in utility learning. In *2017 American Control Conference (ACC)*, pages 5249–5256, May 2017.
- [5] L. J. Ratliff, M. Jin, I. C. Konstantakopoulos, C. Spanos, and S. S. Sastri. Social game for building energy efficiency: Incentive design. In *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1011–1018, Sep. 2014.
- [6] T.G. Papaioannou, N. Dimitriou, K. Vasilakis, A. Schoofs, M. Niki-forakis, F. Pursche, N. Deliyski, A. Taha, D. Kotsopoulos, C. Bardaki, S. Kotsilitis, and A. Garbi. An iot-based gamified approach for reducing occupants' energy wastage in public buildings. *Sensors (Switzerland)*, 18(2), 2018.
- [7] T.G. Papaioannou and G.D. Stamoulis. Teaming and competition for demand-side management in office buildings. volume 2018-January, pages 332–337, 2018.
- [8] Ben Cowley, Jose Luiz Moutinho, Chris Bateman, and Alvaro Oliveira. Learning principles and interaction design for “green my place”: A massively multiplayer serious game. *Entertainment Computing*, 2(2):103 – 113, 2011. Serious Games Development and Applications.
- [9] Ian Ayres, Sophie Raseman, and Alice Shih. Evidence from Two Large Field Experiments that Peer Comparison Feedback Can Reduce Residential Energy Usage. *The Journal of Law, Economics, and Organization*, 29(5):992–1022, 08 2012.
- [10] Tiffany Grace Holmes. Eco-visualization: Combining art and technology to reduce energy consumption. In *Proceedings of the 6th ACM SIGCHI Conference on Creativity & Cognition*, C&C '07, pages 153–162, New York, NY, USA, 2007. ACM.
- [11] Dirk Börner, Jeroen Storm, Marco Kalz, and Marcus Specht. Energy awareness displays: Prototype for personalised energy consumption feedback. In *Proceedings of the 7th European Conference on Technology Enhanced Learning*, EC-TEL'12, pages 471–476, Berlin, Heidelberg, 2012. Springer-Verlag.
- [12] Latha Karthigaa Murugesan, Rashina Hoda, and Zoran Salcic. Investigating visualization of energy consumption. In *Proceedings of the 14th Annual ACM SIGCHI NZ, CHINZ '13*, pages 12:1–12:1, New York, NY, USA, 2013. ACM.
- [13] John G Rodgers. *Residential resource use feedback: exploring ambient and artistic approaches*. PhD thesis, Communication, Art & Technology: School of Interactive Arts and Technology, 2011.
- [14] Anton Gustafsson and Magnus Gyllenswärd. The power-aware cord: Energy awareness through ambient information display. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '05, pages 1423–1426, New York, NY, USA, 2005. ACM.
- [15] Jonas Togler, Fabian Hemmert, and Reto Wettach. Living interfaces: The thrifty faucet. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, TEI '09, pages 43–44, New York, NY, USA, 2009. ACM.
- [16] Eiman Y. Elbanhawy, Andrew F. G. Smith, and John Moore. Towards an ambient awareness interface for home battery storage system. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, UbiComp '16, pages 1608–1613, New York, NY, USA, 2016. ACM.
- [17] Marshini Chetty, A.J. Bernheim Brush, Brian R. Meyers, and Paul Johns. It's not easy being green: Understanding home computer power management. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, pages 1033–1042, New York, NY, USA, 2009. ACM.
- [18] Lara S. G. Piccolo, Cecilia Baranauskas, Miriam Fernandez, Harith Alani, and Anna de Liddo. Energy consumption awareness in the workplace: Technical artefacts and practices. In *Proceedings of the 13th Brazilian Symposium on Human Factors in Computing Systems*, IHC '14, pages 41–50, Porto Alegre, Brazil, Brazil, 2014. Sociedade Brasileira de Computação.
- [19] Tanyoung Kim, Hwajung Hong, and Brian Magerko. Corallog: Use-aware visualization connecting human micro-activities to environmental change. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '09, pages 4303–4308, New York, NY, USA, 2009. ACM.
- [20] Filipe Quintal, Nuno J. Nunes, Adrian Ocneanu, and Mario Berges. Sinais: Home consumption package: A low-cost eco-feedback energy-monitoring research platform. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, DIS '10, pages 419–421, New York, NY, USA, 2010. ACM.
- [21] Deb Polson and Cassandra Selin. The ecos green buildings project: Data dramatization, visualization and manipulation. In *Proceedings of the Second International Conference on ICT As Key Technology Against Global Warming*, ICT-GLOW'12, pages 33–43, Berlin, Heidelberg, 2012. Springer-Verlag.
- [22] Jon Froehlich, Tawanna Dillahunt, Predrag Klasnja, Jennifer Mankoff, Sunny Consolvo, Beverly Harrison, and James A Landay. UbiGreen: Investigating a Mobile Tool for Tracking and Supporting Green Transportation Habits. 2009.
- [23] Alan S. Gerber, Donald P. Green, and Christopher W. Larimer. Social pressure and voter turnout: Evidence from a large-scale field experiment. *American Political Science Review*, 102(1):33–48, 2008.
- [24] Leonard E. Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state markov chains. *Ann. Math. Statist.*, 37(6):1554–1563, 12 1966.
- [25] Sonny Rosenthal, Edmund Lee, Shirley Ho, and Benjamin Detenber. Perceptions of climate change in Singapore and the United States. 06 2013.
- [26] Sofiah Jamil. Fighting climate change from Singapore. 2017.