Wattom: a Energy Aware Smart Plug with Mid-air Controls

ABSTRACT

This paper presents Wattom, a highly interactive ambient eco-feedback smart plug that aims to support a more sustainable use of electricity by being tightly coupled to users' energy-related activities. We describe three use cases of the system: using Wattom to power connected appliances and understand the environmental impact of their use in real time; scheduling these power events; and presenting users with personal consumption data desegregated by device. We conclude with a user study in which the effectiveness of the plug's novel interactive capabilities is assessed (mid-air, hand-based motion matching). The study explores the effectiveness of Wattom and motion matching input in a realistic setup, where the user is not always directly ahead of the interface, and not always willing to point straight at the device (e.g., when the plug is at an uncomfortable angle). Despite not using a graphical display, our results demonstrate that our motion matching implementation was effective in line with previous work, and that participants' pointing angle did not significantly affect their performance. On the other hand, participants were more effective while pointing straight at Wattom, but reported not to finding this significantly more strenuating then when pointing to a comfortable position of their choice.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous.

Author Keywords

Smart plug; motion matching; sustainability; smart watches; mid-air interaction; eco-feedback; ambient display.

INTRODUCTION AND RELATED WORK

The rising energy consumption and its impact on the environment is one of the biggest issues of our times; one that is currently being tackled across several fields such as economics, engineering or psychology. The end-user or energy consumer is an important element of this widespread sustainability effort. If we consider that 13% of the total energy consumption in the world happens in a domestic environment, savings made here can have a big impact in our sustainability efforts, especially since this value is expected to continue to increase [2]. Approaches that deal with the energy consumption of end-users fall under the umbrella of eco-feedback technologies, defined as technology that provides feedback on individual or group behaviors with a goal of reducing their impact on the environment[18]. In general, these technologies present consumption information in terms of monetary cost or Watts, either numerically, in charts [14]), or even using metaphors, such as a digital fish tank whose representation improves under the user's sustainable behaviors [19]. However, some have argued that this types of information are often disconnected from the actual behavior driving the consumption, or from the resources being consumed [11]. Pierce *et al.* [25] was one of the first to address this issue by materializing energy through a set of tangible artifacts.

This sub-field of eco-feedback has been described as tangible or ambient eco-feedback, and generally, work in this area relies on physical artifacts that change their appearance to inform users on their energy usage. One of the early examples is the Power-Aware Cord [20], a power extension that displays the electricity consumed through a colored animation on the cord itself. Heller *et al.* [21, 22] explored a similar approach, in which the consumption information was provided not on a power cord, but on a wall socket (through colors and animation). Other examples describe ambient displays as bespoke physical artifacts [1, 24, 26], or embedded into a home's smart lighting system [13]. And although tangible and ambient eco-feedback displays have been shown to be very effective at raising awareness of users' energy consumption [6], these tend to present one-dimensional data, support very little interaction, and ultimately, limit how users act on their consumption [27].

We present Wattom, a smart plug that builds on the Power-Socket [21] by supporting direct and mid-air interaction with the device itself, and the consumption information it provides. Interaction with Wattom is built around motion matching [17, 29, 31], a novel interaction paradigm where users interact with the system by tracking its different animations with their hands (tracked through the motion sensors embedded on most wrist-worn wearable devices). Compared to other mid-air and embodied interaction techniques [4, 12, 23], motion matching does not require gesture discovery and memorization [10]; and our implementation is not affected by the limitations of optical tracking, such as a limited field-of-view (FOV), occlusion and changing lighting conditions, and well-known privacy concerns when used in the context of the home [7]. This paper provides the following contributions. First, we describe the

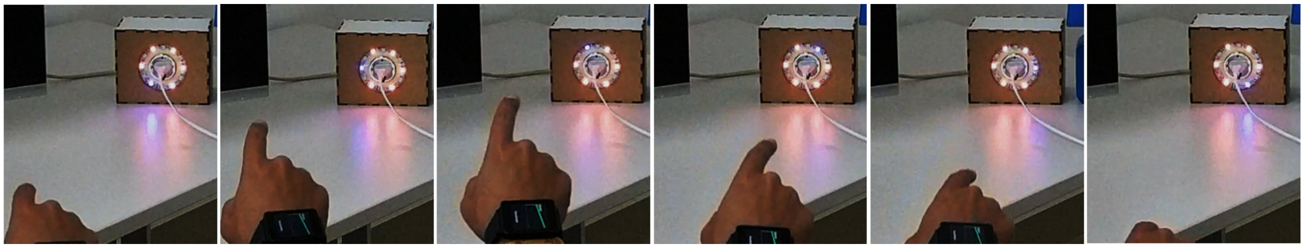

Figure 1. Wattom supports mid-air input through a motion matching approach, in which interface elements (a blue light) move in distinct trajectories, and the user interacts with these by tracking their motion using their arms (e.g., to power the connected appliance). Wattom can also adjust its backdrop color based on the amount of renewable energy present in the grid at the moment of interaction (red indicates a renewable quota close to 0%).

implementation and affordances of Wattom, a highly interactive ambient eco-feedback device that is tightly coupled to users' energy-related activities. And second, we study the effectiveness of motion matching in the context of a smart home, where users not always perceive the moving stimulus directly ahead of them (as in a traditional lab setting [8, 9, 30]), or are not able to comfortably track interface motions by pointing directly at the device (e.g., while laying on the sofa).

WATTOM

Inspired by tangible and ambient eco-feedback devices such as the Power-Aware Cord [20] and the PowerSocket [21], Wattom is a energy aware smart plug with direct and touchless interactive capabilities. Wattom allows users to be informed on the energy usage of any electrical device connected to it, to control the power to said devices, and to be informed on the amount of electricity being generated by renewable sources at any given moment — which in turn can be used to facilitate or hinder access to the plug's control capabilities. In this section, we describe Wattom's implementation and functionality, and three example uses of the system.

Implementation

Wattom employs a Schuko plug[1], that for the sake of convenience was housed in a 16.3 (w) × 12.5 (h) × 8.7cm wooden case that connects to a standard wall socket. Surrounding the Schuko plug are 12 RGB LEDs (24-bit color), equally spaced around a 7cm circle (diameter). Wattom houses two micro-controllers: an Arduino Nano[2] responsible for controlling the LEDs; and an Intel Edison[3] that manages the program logic, server-side access, and energy measurement and control (together with an ADC, 10A relay, and ASC712 current sensor). In addition to energy measurement, the Intel Edison is also responsible for two additional core features of Wattom. It records *power events*, defined as sudden power changes, and which are normally associated with an appliance being turned on or off, or a change in state (e.g., a fridge door being opened). Afterwards, it employs a non-intrusive load monitoring (NILM) algorithm [32], enabling Wattom to identify which electrical appliance is being used simply by its power event signature. Finally, the micro-controllers communicate using the I^2C protocol. Full details and diagrams can be found at *hidden for anonymous submission*.

[1]https://en.wikipedia.org/wiki/Schuko
[2]https://store.arduino.cc/arduino-nano
[3]https://software.intel.com/en-us/iot/hardware/edison

Server-side

Wattom communicates wirelessly with a Node.js[4] web server that logs the state of every LED in real-time. The web server also provides an API that enables third-party applications to have full control over the LEDs, to power or disable the device(s) connected, and to issue any queries related to the energy consumption of these devices. The server is also capable to accessing other service providers for relevant information, such as real time data on the renewable energy quota. Finally, while only one Wattom plug was built, the web server was designed to handle multiple plugs at once with minimal configuration.

Interaction

Wattom supports mid-air interaction by implementing a motion matching approach. In motion matching [29], interface elements move in continuous and distinct trajectories, and the user interacts with these by tracking their motion as closely as possible using any body movement (see Figure 1). In Wattom, user input is captured using an approach similar to WaveTrace [30], a motion matching implementation where user input is captured through motion sensors embedded in most wrist-worn devices such as smart watches, and represented as *yaw* and *pitch* values. While visually inspired by [15], the lack of a graphical displays means that Wattom employs a novel approach of displaying interface elements as moving lights using the LEDs (the user perceives the moving lights as moving targets). Wattom can display up to six concurrent targets (so that there is at least one LED turned off in-between targets), and these can move in clockwise or counter-clockwise motions. Targets share the same speed to minimize overlapping.

Target selection is supported as follows. A bespoke Android Wear 1.0 application runs in both the user's wrist-worn device (*wearable*) and smart phone/tablet (*mobile*). As in WaveTrace, the execution starts after the user performs a flick of the wrist (a standard Android gesture), which triggers the target movements. At the start of this process, and using the server's API, the mobile device receives all target positions (i.e., active LEDs, a value from 0 to 11) and directions (0 or 1). The mobile device is then responsible for simulating this movement in a Cartesian plane (*x*- and *y*-positions) using a predefined target speed, effectively increasing the resolution of the Wattom display. The mobile device is also responsible for receiving yaw and pitch values from the user's wearable device. These activities take place at 25Hz. At 300ms intervals, the mobile device

[4]https://nodejs.org/en/

verifies the target positions on the web server, and updates its simulation if necessary. Finally, the mobile device runs various Pearson's correlations [30] between the yaw and x data points of all targets, and between the pitch and y data points. If four consecutive correlation results between user input and a target movement are over a developer-defined threshold in both axis, the target is selected (and the user receives a haptic confirmation on their wearable device. These correlations take place every 60ms, using the latest 40 data points (rolling window). Target movement is disabled after the user rests its arm by its side (to minimize visual overload).

We now describe three use cases built as third-party applications to the Wattom's web server API.

Powering devices and environmental awareness
The simplest functionality of Wattom enables users to remotely power or disable a connected appliance. This is done by displaying a moving blue target that triggers this functionality. Inspired by [20, 22], the speed of this target and backdrop colors (the other LEDs) change based on the amount of renewable energy present in the grid at the moment of interaction. Our implementation averaged the percentage of electricity from renewable sources in the grid over the last three years in *hidden for anonymous submission*. This resulted in a value close to 25%. As such, during normal operation, if the real time renewable quota is close to this baseline, Wattom displays a yellow color as backdrop to the blue target, which moves at 140°/s. On the other hand, if the real time renewable quota is below or above this baseline (from 0 to 50%), Wattom shines redder (see Figure 4) or greener, and has the target move faster (max. 180°/s) or slower (min. 100°/s), respectively. This not only informs the user of the impact of his activity at the time of the interaction, but by having the target move at a faster speed, actively makes it harder to power on a device that will draw from electricity generated from a large amount of non-renewable sources.

Simple improvements to this application could be easily implemented. For example, the renewable quota baseline could be averaged on a seasonal basis (e.g., there is more electricity generated from solar panels in the Summer months). Otherwise, instead of looking at renewable quotas, Wattom could apply the same color/speed metaphor to the household's or neighbourhood's average electrical consumption (if this data is publicly available). This way, users could become more aware of positive or negative behaviors compared to their own past activities, or that of their neighbours.

Scheduling power events
Our second application directly expands on the application above by enabling users to schedule when to power or disable an appliance. This allows for a better control of their electrical consumption; to plan certain activities when the electricity is cheaper (e.g., night plans) [3]; or for when the user anticipates a higher electrical output from their local photo-voltaic installation [28]. To achieve this, users start a bespoke Android Wear application on their smart watches, and select a start and end period for Wattom (see Figure 2). This trigger all plugs connected to a web server (e.g., representing a household) to sequentially display the start and end times using the 12 LEDs as pointers on an analog clock. Afterwards, each Wattom plug displays a single moving target, enabling users to quickly assign this schedule to one or more plugs.

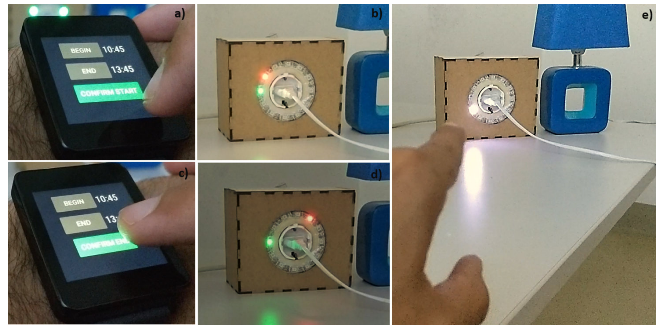

Figure 2. Wattom allows users to schedule when to power or disable an appliance through their smart watches (left). After confirming the schedule on a plug (middle), each Wattom plug displays a target that enables users to quickly assign their schedule to one or more plugs (right).

Figure 3. Wattom's NILM capabilities enables the distinction between various appliances connected simultaneously. Wattom displays a colored target per appliance (left), which upon selection sends the energy consumption of that device to the user's smart watch (right).

Disaggregating consumption
Our final example application relies on Wattom's non-intrusive load monitoring (NILM) capabilities to distinguish between multiple appliances connected simultaneously to a single Wattom plug. Wattom displays a individually-colored target per appliance (up to six appliances), which upon selection sends the latest energy consumption of that device directly to the user's smart watch (see Figure 3). Furthermore, this data reflects only the consumption of that particular user, as the data has been previously tagged with their smart watch information when used to power the appliance using our two first applications. While not implemented, other forms of data visualization could have been promptly explored, including highlighting how much of the user-appliance's consumption was powered by renewable energy sources; or comparing consumption data to an average of all other users of the system — e.g., how much does the use of the electric kettle by a user compares to its roommates.

USER STUDY
We designed a lab study to reflect the challenges of embodied, mid-air interaction with smart environments, in which the user is not always able to point at the interface comfortably (e.g.,

while laying down on the sofa), or more importantly, is not always directly facing it. This was done with 20 participants (10F) aged between 20 and 50 (M = 27.82, SD = 8.57), who rated their experience with smart watches at 3.18 (SD = 2.26) using a 7-point Likert scale (higher is better).

Experimental Setup and Design

The experiment was conducted in a quiet room, following a within-subjects design with two independent variables (2 × 6): pointing *direction* and *angle*. In the direction condition, participants were asked to point *towards* and *away* from Wattom; the latter in the most comfortable position they could find. In the angle condition, participants interacted with the plug from two meters away (2°s of visual angle), at 15° intervals between 15 and 90°s (inclusive) — participants in the 90° condition were directly facing the interface. The order in which participants interacted with these conditions was balanced using a Latin Square (2 × 2): participants moved from 15 to 90°s or vice-versa, pointing towards or away from Wattom. Participants completed 21 sequential trials in each condition, with the first trial being regarded as practice and ignored from analysis. 4800 trials were recorded in total: 2400 trials per pointing direction (20 participants × 20 trials × 6 angles), and 800 trials per angle (20 participants × 20 trials × 2 directions).

In each trial, participants were instructed to select a red target among five other white targets. The starting position and direction of these were randomized at the start of each trial, and all moved at a constant speed of 150°/sec. After a target selection, Wattom would blink twice with red or white colors to provide feedback user feedback. A timeout of seven seconds was included to minimize fatigue and frustration, and logged as a failed selection. Finally, Wattom's web server ran on a early version of a PINE64[5] from a local network.

Procedure and Metrics

Each session started with a brief explanation of the study, and by capturing participants' demographics. Afterwards, participants were handed the smart watch (an LG G Watch, placed on the wrist of their dominant hands), shown how the Wattom works, and allowed to interact informally for a short time. A chair was placed at the starting position (at either 15 or 90°s from the interface), and moved by the researcher after the trials for each angle position were completed. Participants were allowed to rest for four seconds between trials, and for ten seconds while changing angle positions. This information was displayed on an 10″ tablet placed beneath Wattom, which also produced auditory feedback after a target selection or timeout. After completing the selection task in all angles in a pointing condition (towards or away), participants were asked to complete the Borg CR10 [5] scale of perceived exertion. This captured participants' perceived effort (shoulder, arms, hands or fingers) in a scale of 0 to 10 (0.5 increments, higher is harder). Finally, data on participants' success rates and selection times were recorded for further analysis.

[5]https://www.pine64.org/

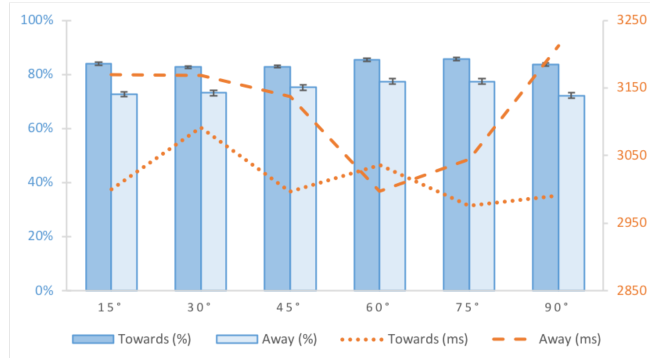

Figure 4. The results from the user study, in which the success rates and acquisition times were captured across two conditions: pointing direction (towards or away from Wattom) and angle (15 to 90°s).

Results

The results for participant performance can be seen in Figure 4. Both success rates and selection times were analyzed using a two-way repeated measures ANOVA. These which incorporated Greenhouse-Geisser corrections when Mauchly's test showed sphericity violations, and were followed by Bonferroni-corrected post-hoc t-tests. Regarding *success*, no statistically significant two-way interaction was found between pointing direction and angle, $F(5, 95) = .210$, $p = .958$ — the assumption of sphericity was met, $\chi^2(14) = 15.078$, $p = .378$. There were also no statistically significant differences between pointing *angles*, $F(5, 95) = .677$, $p = .642$ — the assumption of sphericity was met, $\chi^2(14) = 21.030$, $p = .104$; but there was a significant effect of pointing *direction* on participants' success rates: $F(1, 19) = 11.911$, $p = .003$ — the assumption of sphericity was violated, $p < .001$. Two outliers were identified prior to analysis (studentized residual values of -3.06 and -3.83), but were found to not affect these results.

The results above were mirrored for *selection times*, in which no statistically significant two-way interaction was found between our independent variables, $F(5, 90) = .678$, $p = .641$ — the assumption of sphericity was met, $\chi^2(14) = 12.638$, $p = .560$. There were also no statistically significant differences between pointing *angles*, $F(5, 90) = .612$, $p = .691$ — the assumption of sphericity was met, $\chi^2(14) = 9.202$, $p = .820$; but there was a significant effect of pointing *direction* on participants' selection times: $F(1, 18) = 4.542$, $p = .047$ — the assumption of sphericity was violated, $p < .001$. P10 was removed from the selection times analysis, as it lacked a single data point for the pointing away at a 15° angle condition (0% success rate, and one of the outliers described above).

Finally, a Wilcoxon signed-rank test found no significant differences between participants' perceived effort while pointing towards (M = 2.33, SD = 1.77) or away (M = 2.28, SD = 1.94) from Wattom, z = -.267, p = .787.

Discussion

The results above are fairly positive, and provide novel insights into the use of motion matching to interact with smart devices and environments. First, participants were more effective while pointing towards Wattom, which could potentially limit its use in real-world scenarios when the user cannot point comfortably at the plug (e.g., laying down on the sofa). That being said, participants did not find the task of pointing straight at the Wattom particularly more strenuous then pointing in a comfortable position of their choice (e.g., pointing at the floor). Furthermore, the average success rate of 84% (*SD* = 1.25) while pointing towards Wattom was very much in line with previous work based on motion sensors [16, 30], particularly taking into account the small visual angle (2°) and low resolution of the moving stimulus (12 LEDs/pixels). We argue these results can be improved by implementing Wattom with a higher resolution LED ring (e.g., 24 LEDs[6]).

Acquisition times were fairly consistent across conditions (approx. 3000ms), and in line with previous work using wrist-worn motion sensors [30]. This value can be further explained by our study design and hardware choice. That is, a target selection could only take place after an initial 40 data points (rolling window) had been collected, and this took place at a very low rate of 25Hz — at the start of each new trial, it would take Wattom at least 1600ms to be ready to respond to user input again. That being said, this selection time did not seem to aggravate participants' physical fatigue, as they reported an average Borg CR10 result indicating *weak (light)* physical exertion across conditions. Finally, and more importantly, we were able to demonstrate that participants could effectively interact with Wattom from various angles. This further suggests that the plug does not require a specific interaction 'sweet spot', potentially increasing and facilitating its real-world use. More generally, these results suggests that participants were likely able to mentally transform the visual stimulus at extreme angles (ellipse) to the ideal trajectory (circle). This information should be of value to other researchers exploring real-uses for motion matching interfaces, but further work should explore if this is the case when the target trajectory cannot be easily anticipated (e.g., random or complex target trajectories [8]).

CONCLUSION AND FUTURE WORK

This paper presented Wattom, smart plug made interactive by the use of motion matching input. This work was motivated by a need to provide users with energy consumption information that was tightly coupled with their energy use, and the energy source. Furthermore, Wattom was designed to not only facilitate control over electric appliances for a-far, but also invite users to further inquire about their energy use: either through personal data visualizations, or by hindering their access to the Wattom's control functionality. Finally, we also demonstrate that motion matching input can be provided effectively from very extreme angles, which should further facilitate the plug's use in real-world settings. Because Wattom was built as an open platform with a public API, many more applications could be easily explored in the future, including its use as a non-interactive ambient display, or outside of the scope of sustainability (e.g., the LED ring could display a connected device's battery level). More importantly, the more pressing work should look at studying Wattom in real-life deployments over longer periods of time (e.g., home, office), with the aim of assessing its effectiveness (from an awareness, interaction, and sustainable perspective), social acceptability in shared environments, and user satisfaction. We have provided three uses cases and external resources that should inspire and simplify the implementation of Wattom by other researchers, with the aim of accelerating research in this area, and expanding the uses cases of the platform.

[6]https://www.adafruit.com/product/2863

REFERENCES

1. 2007. THE ENERGY ORB: Visualize Electricity Consumption! (2007). https://inhabitat.com/the-energy-orb-monitor-your-electricity-bill/ [Online; accessed 2018-07-25].
2. 2016. *International Energy Outlook 2016*. US Energy Information Administration.
3. Alper T. Alan, Enrico Costanza, Sarvapali D. Ramchurn, Joel Fischer, Tom Rodden, and Nicholas R. Jennings. 2016. Tariff Agent: Interacting with a Future Smart Energy System at Home. *ACM Trans. Comput.-Hum. Interact.* 23, 4, Article 25 (Aug. 2016), 28 pages. DOI: http://dx.doi.org/10.1145/2943770
4. Richard A. Bolt. 1980. 'Put-that-there': Voice and Gesture at the Graphics Interface. In *Proceedings of the 7th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '80)*. ACM, New York, NY, USA, 262–270. DOI: http://dx.doi.org/10.1145/800250.807503
5. Gunnar Borg. 1998. *Borg's perceived exertion and pain scales.* Human kinetics.
6. Loove Broms, Cecilia Katzeff, Magnus Bång, Åsa Nyblom, Sara Ilstedt Hjelm, and Karin Ehrnberger. 2010. Coffee Maker Patterns and the Design of Energy Feedback Artefacts. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems (DIS '10)*. ACM, New York, NY, USA, 93–102. DOI: http://dx.doi.org/10.1145/1858171.1858191
7. A.J. Bernheim Brush, Bongshin Lee, Ratul Mahajan, Sharad Agarwal, Stefan Saroiu, and Colin Dixon. 2011. Home Automation in the Wild: Challenges and Opportunities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 2115–2124. DOI: http://dx.doi.org/10.1145/1978942.1979249
8. Marcus Carter, Eduardo Velloso, John Downs, Abigail Sellen, Kenton O'Hara, and Frank Vetere. 2016. PathSync: Multi-User Gestural Interaction with Touchless Rhythmic Path Mimicry. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3415–3427. DOI: http://dx.doi.org/10.1145/2858036.2858284
9. Christopher Clarke, Alessio Bellino, Augusto Esteves, and Hans Gellersen. 2017. Remote Control by Body Movement in Synchrony with Orbiting Widgets: An Evaluation of TraceMatch. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3 (Sept. 2017), 45:1–45:22. DOI: http://dx.doi.org/10.1145/3130910

10. Christopher Clarke, Alessio Bellino, Augusto Esteves, Eduardo Velloso, and Hans Gellersen. 2016. TraceMatch: a computer vision technique for user input by tracing of animated controls. In *UbiComp '16: Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM Press, 298–303. DOI: http://dx.doi.org/10.1145/2971648.2971714

11. S. Darby. 2006. The Effectiveness of Feedback on Energy Consumption: A Review of the Literature on Metering, Billing and Direct Displays. (2006), 1–21. http://www.google.com/search?q=Valuation+of+Non-Energy+Benefits+to+Determine+Cost-Effectiveness+of+Whole-House+Retrofit+Programs:+A+Literature+Review+Amann

12. Paul Dourish. 2001. *Where the Action is: The Foundations of Embodied Interaction*. MIT Press, Cambridge, MA, USA.

13. Eiman Y. Elbanhawy, Andrew F. G. Smith, and John Moore. 2016. Towards an ambient awareness interface for home battery storage system. ACM Press, 1608–1613. DOI:http://dx.doi.org/10.1145/2968219.2968557 [Online; accessed 2018-07-25].

14. Thomas Erickson, Ming Li, Younghun Kim, Ajay Deshpande, Sambit Sahu, Tian Chao, Piyawadee Sukaviriya, and Milind Naphade. 2013. The Dubuque Electricity Portal: Evaluation of a City-scale Residential Electricity Consumption Feedback System. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 1203–1212. DOI: http://dx.doi.org/10.1145/2470654.2466155

15. Augusto Esteves, Eduardo Velloso, Andreas Bulling, and Hans Gellersen. 2015. Orbits: Gaze Interaction for Smart Watches using Smooth Pursuit Eye Movements. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. ACM Press, 457–466. DOI:http://dx.doi.org/10.1145/2807442.2807499

16. Augusto Esteves, David Verweij, Liza Suraiya, Rasel Islam, Youryang Lee, and Ian Oakley. 2017. SmoothMoves: Smooth Pursuits Head Movements for Augmented Reality. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (UIST '17)*. ACM, New York, NY, USA, 167–178. DOI: http://dx.doi.org/10.1145/3126594.3126616

17. Jean-Daniel Fekete, Niklas Elmqvist, and Yves Guiard. 2009. Motion-pointing: Target Selection Using Elliptical Motions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 289–298. DOI: http://dx.doi.org/10.1145/1518701.1518748

18. Jon Froehlich, Leah Findlater, and James Landay. 2010. The Design of Eco-feedback Technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1999–2008. DOI: http://dx.doi.org/10.1145/1753326.1753629

19. Jon Froehlich, Leah Findlater, Marilyn Ostergren, Solai Ramanathan, Josh Peterson, Inness Wragg, Eric Larson, Fabia Fu, Mazhengmin Bai, Shwetak Patel, and James A. Landay. 2012. The Design and Evaluation of Prototype Eco-feedback Displays for Fixture-level Water Usage Data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 2367–2376. DOI: http://dx.doi.org/10.1145/2207676.2208397

20. Anton Gustafsson and Magnus Gyllenswärd. 2005. The Power-aware Cord: Energy Awareness Through Ambient Information Display. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems (CHI EA '05)*. ACM, New York, NY, USA, 1423–1426. DOI: http://dx.doi.org/10.1145/1056808.1056932

21. Florian Heller and Jan Borchers. 2011. PowerSocket: Towards On-outlet Power Consumption Visualization. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems (CHI EA '11)*. ACM, New York, NY, USA, 1981–1986. DOI: http://dx.doi.org/10.1145/1979742.1979901

22. Florian Heller and Jan Borchers. 2012. Physical Prototyping of an On-outlet Power-consumption Display. *interactions* 19, 1 (Jan. 2012), 14–17. DOI: http://dx.doi.org/10.1145/2065327.2065332

23. Mingyu Liu, Mathieu Nancel, and Daniel Vogel. 2015. Gunslinger: Subtle Arms-down Mid-air Interaction. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 63–71. DOI: http://dx.doi.org/10.1145/2807442.2807489

24. Heiko Müller, Jutta Fortmann, Martin Pielot, Tobias Hesselmann, Benjamin Poppinga, Wilko Heuten, Niels Henze, and Susanne Boll. 2012. AmbiX : Designing Ambient Light Information Displays.

25. James Pierce and Eric Paulos. 2010. Materializing energy. *Proceedings of the 8th ACM Conference on Designing Interactive Systems* (2010), 113–122. DOI: http://dx.doi.org/10.1145/1858171.1858193 [Online; accessed 2013-09-16].

26. Matevz Pogacnik and Jože Guna. 2013. Ambient visualization of energy consumption information. International Workshop on Semantic Ambient Media.

27. J. Rodgers and L. Bartram. 2011. Exploring Ambient and Artistic Visualization for Residential Energy Use Feedback. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2489–2497. DOI: http://dx.doi.org/10.1109/TVCG.2011.196

28. Johann Schrammel, Cornelia Gerdenitsch, Astrid Weiss, Patricia M. Kluckner, and Manfred Tscheligi. 2011. FORE-Watch – The Clock That Tells You When to Use: Persuading Users to Align Their Energy Consumption

with Green Power Availability. In *Ambient Intelligence*, David V. Keyson, Mary Lou Maher, Norbert Streitz, Adrian Cheok, Juan Carlos Augusto, Reiner Wichert, Gwenn Englebienne, Hamid Aghajan, and Ben J. A. Kröse (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 157–166.

29. Eduardo Velloso, Marcus Carter, Joshua Newn, Augusto Esteves, Christopher Clarke, and Hans Gellersen. 2017. Motion Correlation: Selecting Objects by Matching Their Movement. *ACM Trans. Comput.-Hum. Interact.* 24, 3 (April 2017), 22:1–22:35. DOI: http://dx.doi.org/10.1145/3064937

30. David Verweij, Augusto Esteves, Vassilis-Javed Khan, and Saskia Bakker. 2017. WaveTrace: Motion Matching Input Using Wrist-Worn Motion Sensors. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. ACM, New York, NY, USA, 2180–2186. DOI: http://dx.doi.org/10.1145/3027063.3053161

31. John Williamson and Roderick Murray-Smith. 2004. Pointing Without a Pointer. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems (CHI EA '04)*. ACM, New York, NY, USA, 1407–1410. DOI: http://dx.doi.org/10.1145/985921.986076

32. M. Zeifman and K. Roth. 2011. Nonintrusive appliance load monitoring: Review and outlook. *IEEE Transactions on Consumer Electronics* 57, 1 (February 2011), 76–84. DOI:http://dx.doi.org/10.1109/TCE.2011.5735484