Green Multimedia: Informing People of their Carbon Footprint through Two Simple Sensors

Aiden R. Doherty, Zhengwei Qiu, Colum Foley, Hyowon Lee, Cathal Gurrin, Alan F. Smeaton
CLARITY: Centre for Sensor Web Technologies
Dublin City University, Ireland
aiden.doherty@dcu.ie

ABSTRACT

In this work we discuss a new, but highly relevant, topic to the multimedia community; systems to inform individuals of their carbon footprint, which could ultimately effect change in community carbon footprint-related activities. The reduction of carbon emissions is now an important policy driver of many governments, and one of the major areas of focus is in reducing the energy demand from the consumers i.e. all of us individually. In terms of CO_2 generated from energy consumption, there are three predominant factors, namely electricity usage, thermal related costs, and transport usage. Standard home electricity and heating sensors can be used to measure the former two aspects, and in this paper we evaluate a novel technique to estimate an individual's transport-related carbon emissions through the use of a simple wearable accelerometer. We investigate how providing this novel estimation of transport-related carbon emissions through an interactive web site and mobile phone app engages a set of users in becoming more aware of their carbon emissions. Our evaluations involve a group of 6 users collecting 25 million accelerometer readings and 12.5 million power readings vs. a control group of 16 users collecting 29.7 million power readings.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; H.5 [**Information Interfaces and Presentation**]: General; J.2 [**Pysical Sciences and Engineering**]: Engineering

General Terms

Algorithms, Measurement

Keywords

CO_2, home energy monitoring, activity monitoring, multimodal feedback, machine learning application

MM'10, October 25–29, 2010, Firenze, Italy.

1. INTRODUCTION

One of the topical areas of discussion in the last decade is that of climate change. One of the most commonly referred to elements within this discussion is the role of carbon emissions, and in particular the role of Carbon Dioxide (CO_2). Many governments have produced policy documents and targets related to reducing CO_2 emissions over the coming years, with one of the most ambitious targets set by the European Union 20-20-20 target whereby the aim is to reduce the 2020 CO_2 emissions by 20% from 1990 levels, through a 20% reduction of energy consumption and a 20% increase in renewable energy sources [9]. As a result of this many resources are now allocated to energy research to contribute to meeting these ambitious targets.

Figure 1: The mobile carbon foot-printing website

In relation to CO_2 there are three main sources of energy that offload high quantities of CO_2 to the environment, namely electricity, thermal, and transport. The form of energy we immediately think of is that of electricity which was responsible for approximately 31% of CO_2 emissions in one country in 2008, however thermal forms of energy were responsible for approximately 33% of CO_2 emissions. Of most concern is the fact that the largest contributor to energy-related CO_2 emissions was from the transport domain, with 36% of CO_2 emissions, representing a growth of 177% over the period of 1990-2008 [12].

In terms of helping to reduce CO_2 emissions a broad spectrum of expertise is required to meet this challenge. For electricity generation, physicists and engineers are concentrating on creating more efficient power plants [5] so as to alter the fact that currently 47% of energy input into power generation plants is lost to heat. For heat generation, plasma researchers are investigating techniques to improve the efficiency of solar panels (currently only 14% of possible energy is converted into usable output), as well as material scientists exploring the possibilities of new forms of insulation for buildings meaning less heating is required. In terms of transport, vehicle manufacturers are now aggressively marketing new models based on improved energy consumptions ratings.

This research on the supply side of energy is now quite well established and focused. However it is becoming increasingly evident that a subtle change in end user behaviour can reduce overall demand placed on the energy suppliers. Even though some claim changing end user behaviour is not necessary [10], this demand shift management has now gained increased attention from the major energy suppliers who actually want to better model the demand from users so as to better predict when peak energy demands are, and how much energy will be required at those times. A key factor in understanding end user behaviour is in firstly being able to measure their current consumption to a fine-grained level of granularity.

Buildings (both domestic and commercial) can be retrofitted with an array of commercial sensors to measure the electricity and thermal usage patterns of a multitude of appliances in the relevant settings. Cheap, simple (and relatively accurate) energy and thermal monitors can be fitted to a domestic fusebox and boiler respectively to capture overall electricity and thermal consumption in the home. Also devices could be easily manufactured so that cars can upload fuel consumption information to a central server. However while these solutions may be commonplace in 20 years time, they are prohibitively expensive at the present. To our knowledge no ambient means of giving individuals' feedback on their transport CO_2 emissions has become established as of yet.

In this paper we consider the prevalence of wearable sensors, in particular accelerometers which are present in many mobile phones, as a possible means of estimating the transport related CO_2 emissions of an individual. This is achieved through building a classifier to detect when an individual is driving, based on x/y/z motion sensor values alone. From this we estimate how long they are driving for, from which an estimate can be made on how many litres of gasoline are consumed, which finally can be converted to a CO_2 emission — all automated with no data gathering or logging overhead placed on the user.

We believe the multimedia community can contribute to the CO_2 reduction effort through heavy involvement on the demand shift management, by creating engaging systems to create long-term sustainable reductions in end-user energy usage behaviour. In this paper we present two systems we used to inform 20 users of their driving CO_2 emissions and electricity costs, one a web-based system, and another a mobile application.

This paper is organised as follows: Section 2 provides a context to this work via a review of the relevant literature and policy documents. In Section 3 we describe algorithms to estimate an individual's CO_2 emissions based on 2 simple sensors, one in the home and one tied to one's keyring to detect driving. Section 4 describes the systems we built to engage our 20 participants with their energy usage through the form of 1) a desktop application, 2) a web page and 3) a mobile application. Section 5 describes our experimental setup while in Section 6 we supply results on the accuracy of our driving CO_2 estimator and on how people engaged with the various applications. In Section 7 we discuss our experiences thus far and how we believe the multimedia community can move forward in contributing to the CO_2 reduction agenda. Finally Section 8 provides some of the major conclusions from this piece of work.

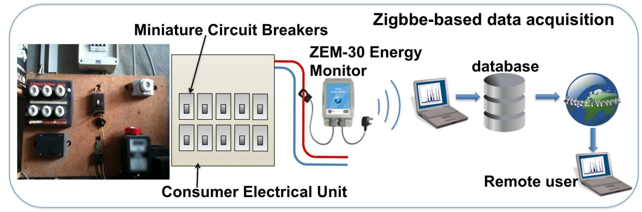

Figure 2: Overview of capturing home electricity usage data and uploading to a central server

2. RELATED WORK

Various projects have investigated the demand reduction benefits of providing electricity consumption feedback to users, through over-instrumenting residences [20]. One of the best examples is the Olympic Peninsula study on 112 residential homes which were fitted with smart appliances (which could be switched on/off remotely), including thermostats, water heaters and clothes dryers. Data was uploaded to a central server every 5 minutes with price-plans changing in real-time based on the energy demand requested by the 112 residences. On average the consumers saved approximately 10% on their bills from the previous year. Importantly from a power generation point of view the peak distribution load requests were significantly reduced [20]. Other "smart meter" trials have similarly shown that increased awareness of energy use, and helping consumers interpret that data, consistently leads to reductions and flattening of that peak demand curve [23].

The results of these findings have convinced many major electricity providers to roll out SmartMeter programs, e.g. the main Italian energy provider, ENEL, has already installed about 30 million metering units [19]. There are a number of start-ups and big, established companies like GE building energy management tools including the smart meter, management software, and energy dashboards. One of the earliest example of web-based system is GreenQuest[1], an energy efficiency internet software tool, which considers data from utility bills, data associated to the area where the house is located, type of building, weather data, and more — to create a custom energy report card. The user is also able to compare personal usage to peer groups, visualize savings over the past months, or years, and track greenhouse gas usage, while businesses can also submit to get an energy star rating. However this still requires effort from the user to enter in data from the utility bills.

Of late both Microsoft and Google have joined the home-energy management market, with quite different approaches.

[1]http://www.MyGreenQuest.com

The recently released Microsoft's Hohm [2] is a free web-based expert system that uses advanced analytics licensed from the Lawrence Berkeley National Laboratory and the US Department of Energy to help people track, understand and manage their personal energy usage. A homeowner can type in a Web ID and zip code to see average energy use in their region or nationally. Microsoft Hohm provides savings recommendations, which can range from placing new caulking on windows to removing air leaks to installing a programmable thermostat. These recommendations are tailored based on specific circumstances in the consumer's home including house features (e.g. square footage), usage patterns and appliances (e.g. brand of water heater). If consumers don't provide their data, Microsoft Hohm will base its recommendations on local and national averages, however we believe a drawback of this system is that it places an excessive time investment from users to receive personalised feedback.

We believe that Google are currently a little more consumer friendly as they rely on data *passively sensed* in the home from standard electricity "smartmeters" . Google is already working with a number of utilities on this program as well as smart meter maker Itron. Using Google PowerMeter[3], a person can view details, such as real-time electricity use and weekly trends from a Web browser or using a smart phone running iGoogle. Data is pushed to the Google servers every 10 minutes, the same as the work described in this paper (Section 3.1).

As can be seen there has been much recent activity in informing users of their electricity consumption and achieving savings through this. However considering that approximately 36% of CO_2 energy related emissions are transport related, we are not aware of any centralised efforts to inform users of their transport related behaviour. Newer models of cars do provide CO_2 emissions feedback to drivers in the car, and this is the most accurate means of calculating CO_2 emissions from driving. However there is no easy method of gaining access to this data and then integrating it with other sources to give users a more complete picture of their overall carbon footprint. In this paper we utilise wearable sensors through the form of the Actigraph GT3x accelerometer [14] attached to the keyrings of participants. This work was motivated by the growing prevalence of wearable sensors and even lifelogging devices (e.g. the Vicon Revue[4]) that are now available to both the research community and the public at large. There has been associated research on predicting the movements of people using wearable sensors [2], efforts to make these wearable sensors even less obtrusive [6], etc. To our knowledge this paper presents the first approach in estimating one's CO_2 emissions from electricity plus transport.

3. CALCULATING CO_2 EMISSIONS

In this section we describe the approaches used to estimate one's CO_2 emissions from their home and from inferred driving activity. The values inferred from these approaches can then be added together to provide an individual with a more complete CO_2 picture, all from just two cheap, simple sensors.

[2]http://www.microsoft.com/presspass/press/2009/jun09/06-24EnergyUsagePR.mspx

[3]http://www.google.org/powermeter/

[4]http://www.omg3d.com/html/IPLicenseagreement.html

3.1 CO_2 from Home

The calculation of CO_2 emissions based on residential electricity consumption is quite a straightforward process. Electricity meters are placed on the main fusebox of each domestic home, and the power consumed by the household is measured in Watts (joules/second). One kilowatt hour (KW/h) equates to a single unit of energy, which is what electricity providers use to bill consumers. Based on the amount of CO_2 released by the service providers to produce and deliver the necessary electricity to meet the demand of the customers (plus a safety buffer of extra load), the supplier can calculate the CO_2 emission per unit of electricity given to the relevant consumers.

In this work, to monitor electricity consumption of an individual home, we make use of the EpiSensor ZEM-30 data logging unit illustrated in Figure 2. The ZEM-30 includes a plug-in power supply, waterproof enclosure and a CT clip which is attached around the live wire running into a given household's main fusebox. This measures 11 different electrical parameters: *RMS/Peak/SAG current/voltage, real/apparent power* and of most relevance *Watt hours*. This information is relayed across a local Zigbee network to a local PC/laptop which logs the data onto a local relational database. This PC/laptop runs 24/7 and records the electrical parameters every minute. The machines are also connected to the Internet and data is uploaded every 10 minutes to a central web server, thus backing-up the data and also making it available to be accessed from anywhere.

Each unit of electricity consumed can be correlated with a related CO_2 emission value, with one unit generally producing 0.5kgs CO_2 [8].

3.2 Driving CO_2 from Wearable Accelerometer

Given the growing prevalence of wearable sensors, lifelogging, accelerometers built into mobile phones, etc. we now discuss a possible approach to estimate an individual's CO_2 output based on wearable accelerometers passively recording data which is then classified based on how long the user was driving for.

Feature Selection & Training: Each accelerometer sensor consists of tri-axial readings along the X/Y/Z axes. From each axis we derive a set of features to model the range of motion over a given time window to better indicate what activity a user may be involved in. Given the low sampling rate (1 Hz) of the devices we use in our work, were were restrained from converting value into the frequency domain as has been used by others [17, 24, 21], therefore we extract the following additional features for each X/Y/Z reading: *mean of previous 5/20/120/300 readings, magnitude between maximum and minimum values of previous 5/20/120/300 readings, & standard deviation of previous 5/20/120/300 reading.* As a result of this, in addition to the raw X/Y/Z values we have 39 extra metrics on which to describe the data. In order to train our driving classifier, we used the SVM-light [13] implementation of Support Vector Machines and optimised the parameters using cross-fold validation. We used the RBF kernel, and optimised parameters C and γ (gamma). Given the relatively low number of positive driving examples, we based our training on the precision metric of the driving class. As a form of post-processing we perform median smoothing on the classified output, to remove isolated incorrectly classified instances, thus exploiting the temporal re-occurrence nature of the data i.e. it's unlikely that someone would be not driving for 5 minutes, then suddenly drive for 1 second, followed by no driving for another 5 minutes.

After feature selection, training, and then classification on new data, we have a knowledge on when a given user has been driving. From this it is then possible to make an estimation on how long the user has been driving. With an element of domain knowledge we can estimate the average speed the user has been driving at to estimate the number of kilometres covered. Knowing the number of kilometres covered and the typical fuel consumption of the user's car, we then are in a position to estimate the number of litres of gasoline consumed. There is a direct correlation between the number of litres consumed and an associated CO_2 emission. In this paper we investigate how accurate our estimated output is, based on one litre of fuel generally producing 2.2kgs CO_2 [8]. $Timedriving \propto km \propto litresgasoline \propto CO_2$

4. SYSTEMS TO INFORM USERS OF THEIR CARBON ESTIMATION

Here we describe three systems to inform users of their electricity consumption and CO_2 emissions. One system (*Always-On-Desktop*) was provided to give users detailed feedback on their home electricity consumption, another one (*web page*) was provided to give users both detailed feedback on their home emissions and also their driving CO_2 emissions (where relevant). A mobile application was also provided to inform users of home & driving CO_2 emissions on the go. We now describe these systems and later we will discuss elements of change in user behaviour they may have induced.

4.1 Always-On In-Home-Display: Electricity

The interaction platform considered ideal for the home deployment of our technology is an in-home display with touch-screen capability, standing on a table or a shelf in a living room, kitchen or bedroom. We chose ASUS Eee Top Touch-screen PC as the device that exhibits the interaction quality we were targeting, though nowadays an Apple iPad would do just as well.

Figure 3: The In-Home Display as designed, implemented and deployed in 22 homes

As more non-desktop PC platforms such as interactive TableTops, Multi-Touch walls, and various embedded appliances (such as interactive digital picture frames, refrigerator door that displays data on the contents and wall-embedded displays) are becoming more widespread, one of the challenges in designing for such platforms is the need to understand the characteristics of the interaction each of these platforms affords and the situations in which it occurs [16]. In-home displays as we deal with in this paper, are no exception. Designing for a touch-screen, 24/7-on information device requires special consideration on the characteristics of the interaction and the situation where the interaction occurs [16]. The identified characteristics of the device and the afforded interaction in concern are briefly summarised here.

Simple and Easy to use - Unlike a Website viewed with a Web browser on a PC or other stand-alone interface, the In-Home Display we targeted for is less for an intensive or lengthy interaction with the user but more for all-time display of useful information such as a clock on the wall. Thus what is displayed on the screen should in most cases satisfy the user's needs simply by being glanced at, with options for interaction only if the user wishes to approach and interact. Its touch-screen front-end without keyboard or mouse and the home usage context means the interaction should be extremely simple and easy.

Dark background - the always-on nature of the In-Home Display means it should be easy to read its displayed screen contents in home, whether brightly-/dimly-lit or totally dark. A bright background screen will probably be reasonably easy to read when the room is lit, but at night time it brightens a room or blinds a user's eyes when passing by in a dark corridor. Due to this, making the colour orientation of the screen having a dark background with bright information content, similar to the digits of a bedside clockwork radio, should suit. Alternatively, the screen colours and brightness could automatically adjust itself depending on the time of the day or the lighting condition in the room.

No "main menu" - The "main menu" screens that we often see on museum kiosks or mobile phones are deemed not suitable for the In-Home Display since a main menu (e.g. containing 5-10 buttons each leading to a different screen) does not contain any useful information in itself. When a home user is passing by and glances at the screen, such an administrative screen without any content is of no use to him/her as he/she did not gain any useful information with that glance; you might as well keep the device turned off. Thus the device should collect most useful set of information and display upfront on the default screen as the main content at all times. Other options or buttons required to enable the interaction with the user should also be on the screen but not as the primary object but taken to side as secondary.

Information in context - In order for a piece of information displayed on the screen (e.g. today's electricity consumption in kW/h) to be useful and meaningful, it needs to be displayed with at least one other piece of information that is of same type thus set into a context. Even an interpreted figure that is more meaningful to the user (e.g. today's electricity consumption in kW/h converted into a figure in a monetary currency such as Euro) does not provide the sense of whether that figure means high or low in itself. Juxtaposing the figure with yesterday's figure, last month's average figure, a figure from this time last year, or a figure averaged from the whole block of neighbours, for example, makes the displayed figure more meaningful. Thus graphical representation of the temporal trends or a time-based comparison of monthly readings can condense the displayed information in such a way as to provide richer contextual information surrounding the main reading and assists making sense of the displayed information and its implication.

Based on the above identification of the interaction characteristics, an In-Home Display interface for power usage monitoring was designed. Figure 3 shows the display sitting in a dimply-lit kitchen in one of our home users.

The designed interface uses two shades of dark background (dark grey and black) with bright yellow-green and orange colour information contents in order to help its reading in a bright light and at the same time not to blind the eye or brighten the room when the lighting is dim or turned off. The layout of the information organisation emphasises two sets of bar-chart style graphs: the upper one shows daily power consumption reading in the context of a week; while the lower one shows hourly reading in the context of a 24-hour duration. The screen shown in Figure 3 is the main screen which displays such daily and hourly readings, thus a home user can simply glance over to the screen and see the information. Interposed in the readings is the average usage data, and when today's reading turns out higher than the average usage value the bar colour changes to orange. In the figure, the home used less electricity on Wednesday than the average daily usage, but between 8am–9am used more than the average hourly usage. A thin yellow-green bar at the top of the screen and the current time at the top-right both turn to orange colour when the user starts over-using compared to the average value, quietly but clearly indicating such a state. Also displayed left of the reading bars are information on the current power consumption in numeric figure and its monetary equivalent; expected usage in the current time scale; and the number of times each of the home devices has been turned on.

There are other interaction possibilities with circular buttons on the top-left and bottom-right corners. The top-left corner shows four circular buttons (*Day View*, *Week View*, *Month View* and *Year View*) with the default chosen as *Day View*. Touching these buttons changes the two main reading panels to the chosen time scales. The bottom-right buttons provide user settings (e.g. reading units, currency and time/date).

4.2 Web Interaction: Electricity & Driving CO_2

For those who do not have a touch-screen device at home as described in the previous section, more conventional Web access through a desktop would still be a useful means to monitor their power consumption. Taking similar two-panel layout as the touch-screen UI but taking advantage of the more fine-grained mouse control and more elaborate navigation possible under a desktop PC setting, a Silverlight-based Web interface was developed and deployed. As evidenced in Figure 4 users have the ability to view summarised hourly/daily/weekly/monthly versions of their home electricity consumption. Individual elements of the chart can be clicked on to view more details for the selected hour/day/week/month. Again the user is provided with feedback on the number of units they use (i.e. an indicator of CO_2 emitted), and importantly these values are provided in context (as against personalised baseline/average/expected consumption values based on the user's consumption history).

Figure 4: Version of web page displayed to users logging both home & driving CO_2

For users carrying wearable accelerometers using our algorithm to automatically estimate their driving CO_2 emissions, a pie chart is provided showing their relative home vs. driving CO_2 for the given day/week/month. This chart is illustrated on the right hand side of Figure 4 thus providing the user with both home and driving CO_2 indicators.

4.3 Mobile App: Home & Driving CO_2

In addition to touch-screen and desktop Web interfaces, the mobile access is provided as an alternative means of enhancing the power monitoring activity. The small screen, most often an awkward input mechanism, and the context of usage where distraction during use is expected make the mobile interaction quite different from other more conventional platforms. The mobile website, shown in Figure 1, allows users to access their CO_2 Consumption from anywhere using their mobile phone's web browser. Each user in the experiment was provided with a unique URL which they were encouraged to bookmark in their phone's browser for easy access during the experiment. When users accessed their mobile webpage they were presented with a page similar to that shown in Figure 1. This page showed, side by side, the user's CO_2 consumption (in terms of kilograms of carbon dioxide produced) for their home and while driving for today, yesterday and their overall usage. The page also uses visual columns whose height change in proportion to the carbon usage for both home and driving.

A simple PHP file accepts as input a unique user ID, and then produces a simple webpage showing the usage data. The webpage produced conforms to the XHTML Mobile profile which allows it to render correctly across the widest range of mobile phone browsers. Furthermore, given that users could access this page from across a cellular network, we also needed to ensure that the page was lightweight to ensure that it was both low cost in terms of data delivery charges and could render reasonably quickly on a mobile phone. Using the Ready.mobi [5] tester and analysis application, the mobile Carbon site scored a maximum of 5 out of 5.

[5] *Ready.mobi site check:* http://ready.mobi

5. EXPERIMENTAL SETUP

To evaluate the effectiveness of our proposed techniques we now introduce the datasets gathered to investigate our driving detection algorithm and driving CO_2 estimation technique. Afterwards we discuss the data collected by a group of 22 households, 6 of whom logged driving activities, to investigate potential changes in energy consumption behaviour.

5.1 Wearable Accelerometer Driving Detection

Our first experiment involved the collection of one week's wearable accelerometer data from one user. Given the traditional problem of accurately labelling training data for human activity detection, we use the SenseCam which has onboard tri-axial accelerometers [11]. The SenseCam is a valuable context reinstatement tool and the images that it takes can be used by the wearer to retrospectively annotate exactly when he was driving due to the powerful memory cues of visual imagery [3]. In total 132,247 accelerometer readings were logged over a 6 day period, with 6,371 readings being positive examples of driving. 40% of this data was used in cross fold validation to train our SVM parameters.

5.2 Driving CO_2 Estimation

To investigate the validity of our driving CO_2 estimation algorithm, we supplied one of our users with a SenseCam for a period of one year to log all their activity data. In total 9,370,647 accelerometer readings were captured over a 58 week period between February 2009 and April 2010. This user maintained a "driving CO_2" blog over this period of time, recording the number of kilometres travelled, average speed (km/h), and fuel consumption (l/100 km) every week. Therefore the blog CO_2 values were taken as the groundtruth on which to attempt to best match the classification outputs on the 9.3M accelerometer readings using the technique referred to in Section 5.1. The first 20 weeks were used as a training set to optimise the various parameters to map time spent driving to CO_2 emissions, with the remaining 38 weeks used as the test set to report our results.

5.3 Effects on Energy Consumption

Here we describe the data collected to investigate whether being returned additional passively sensed information to provide a more complete estimation of one's carbon footprint, affects change in energy behaviour habits. 6 users were provided with wearable accelerometers over a 6 week period of time, to capture their activity data to identify instances of driving data, using the techniques developed in the aforementioned subsections. Due to the lack of availability of SenseCams, and also the privacy concerns surrounding personal image capture [1], we provided these 6 users with an Actigraph GT3x unobtrusive tri-axis accelerometer [14] where the users were instructed to place it on their keyrings as illustrated in Figure 5, so as not to miss possible driving events. An issue with the capture of this data was the lack of real-time information, given that data had to be manually downloaded from the wearable accelerometers via USB cable. To simulate as near as possible to real-time captured driving data, one of the authors of this paper went to each of the 6 users every morning to swap their accelerometer with a new fully charged one, and afterwards uploading the previous day of data to our databases, so that the users could view their driving CO_2 emissions within an hour, on both the web page and the mobile phone app. In total 14,257,036 accelerometer readings were captured by our 6 subjects, 11% of which was classified as driving data, using techniques investigated in the prior two subsections in this paper, as detailed in Table 2.

These 6 users were also supplied with Episensor ZEM-30 home electricity monitoring kits as illustrated earlier in Figure 2. Here energy usage was logged every minute and automatically updated to a central repository every 10 minutes, essentially providing real-time access to their domestic electricity consumption. This information was gathered for a period of almost one year, with a total of 12,568,215 electricity parameter readings logged across the 6 users as detailed in Table 2, 1,142,565 of which were kW/h records. To validate these readings, user 6 had a Smartmeter installed in his home as part of a national trial by an electricity supply company. Our user's bill for Feb-09 (647.28 units) was almost identical to the data that was logged by our Episensor ZEM-30 home electricity monitor (652.81 units).

To contextualise any potential savings our control group may make, we also provided Episensor ZEM-30 home electricity monitoring kits to a "control group" of 16 participants who logged 29,741,184 electricity parameter records as detailed in Table 3. 2,703,744 of these readings were watt hour records and this group of users exhibit home electricity consumption habits similar to the broader (Irish) population

Figure 5: Red Actigraph GT3x given to users to attach to their keyrings

User	Start	End	Num Acc Readings	Classified Driving
1	05-Mar	06-Apr	2,450,462	168,898 (7%)
2	23-Feb	29-Mar	2,326,638	325,384 (14%)
3	23-Feb	07-Apr	1,872,093	163,717 (9%)
4	23-Feb	06-Apr	2,584,743	63,767 (2%)
5	23-Feb	03-Apr	2,991,693	459,553 (15%)
6	23-Feb	20-Mar	2,031,407	374,052 (18%)
Sum	23-Feb	07-Apr	14,257,036	1,555,371 (11%)

Table 1: Acc. data captured by "test group" users.

User	Start Time	End Time	Num Sensor Readings
1	Sep-09	Apr-10	1,249,996
2	Sep-09	Apr-10	1,986,380
3	Sep-09	Apr-10	2,666,103
4	Feb-10	Apr-10	913,891
5	Sep-09	Apr-10	2,177,505
6	May-09	Apr-10	3,574,340
Overall	May-09	Apr-10	12,568,215

Table 2: Electricity data captured by "test group" of users.

User	Start Time	End Time	Num Sensor Readings
1	May-09	Apr-10	3,875,641
2	Oct-09	Apr-10	1,943,370
3	Oct-09	Apr-10	2,194,819
4	Aug-09	Apr-10	2,314,114
5	Nov-09	Apr-10	1,888,436
6	Oct-09	Apr-10	2,310,847
7	Oct-09	Apr-10	2,192,047
8	Oct-09	Apr-10	2,184,578
9	Oct-09	Apr-10	875,996
10	Oct-09	Apr-10	1,586,981
11	Aug-09	Mar-10	1,586,948
12	Feb-10	Mar-10	433,136
13	Nov-09	Apr-10	1,996,269
14	Nov-09	Apr-10	1,494,053
15	Sep-09	Mar-10	1,333,035
16	Aug-09	Apr-10	1,530,914
Overall	May-09	Apr-10	29,741,184

Table 3: Electricity data captured by "control group" of users.

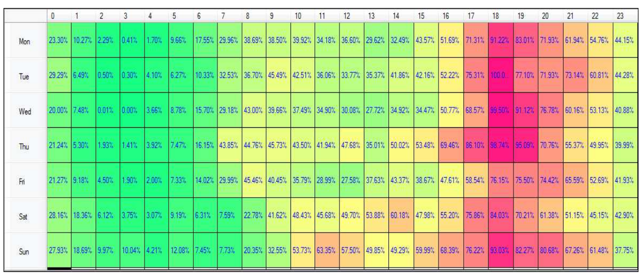

	0	1	2	3	4	5	6	7	8	9	10	11
Mon	23.30%	10.27%	2.29%	0.41%	1.70%	9.66%	17.55%	29.96%	38.69%	38.50%	39.92%	34.18%
Tue	29.29%	6.49%	0.50%	0.30%	4.10%	6.27%	10.33%	32.53%	36.70%	45.49%	42.51%	36.06%
Wed	20.00%	7.48%	0.01%	0.00%	3.66%	8.78%	15.70%	29.18%	43.00%	39.66%	37.49%	34.90%
Thu	21.24%	5.30%	1.93%	1.41%	3.92%	7.47%	16.15%	43.85%	44.76%	45.73%	43.50%	41.94%
Fri	21.27%	9.18%	4.50%	1.90%	2.00%	7.33%	14.02%	29.99%	45.46%	40.45%	35.79%	28.99%
Sat	28.16%	18.36%	6.12%	3.75%	3.07%	9.19%	6.31%	7.59%	22.78%	41.62%	48.43%	45.68%
Sun	27.93%	18.69%	9.97%	10.04%	4.21%	12.08%	7.45%	7.73%	20.35%	32.55%	53.73%	63.35%

	12	13	14	15	16	17	18	19	20	21	22	23
Mon	36.60%	29.62%	32.49%	43.57%	51.69%	71.31%	91.22%	83.01%	71.93%	61.94%	54.76%	44.15%
Tue	33.77%	35.37%	41.86%	42.16%	52.22%	75.31%	100.0...	77.10%	71.93%	73.14%	60.81%	44.28%
Wed	30.08%	27.72%	34.92%	34.47%	50.77%	68.57%	99.50%	91.12%	76.78%	60.16%	53.13%	40.88%
Thu	47.68%	35.01%	50.02%	53.48%	69.46%	86.10%	98.74%	95.09%	70.76%	55.37%	49.95%	39.99%
Fri	27.58%	37.63%	43.37%	38.67%	47.61%	58.54%	76.15%	75.50%	74.42%	65.59%	52.69%	41.93%
Sat	49.70%	53.88%	60.18%	47.98%	55.20%	75.86%	84.03%	70.21%	61.38%	51.15%	45.15%	42.90%
Sun	57.50%	49.85%	49.29%	59.99%	68.39%	76.22%	93.03%	82.27%	80.68%	67.26%	61.48%	37.75%

Figure 6: The energy consumption patterns of our control group during an average week represent the broader Irish population with the typical 5pm-7pm peak in electricity consumption.

with the usual 5pm-7pm spike in electricity consumption as shown in Figure 6 where the rows represent the days in the week and the columns represent the hour in each day. The values in each cell represent the average consumption of that particular hour over the entire dataset, normalised against the most intense consumption of the 168 hours in the week.

6. RESULTS

We now discuss the results obtained after setting up the experiments outlined in Section 5. Firstly we report the accuracy of our driving detector applied on wearable accelerometer data, and afterwards we evaluate the effectiveness of our driving CO_2 estimation algorithm. Finally we provide analysis on the changes in home energy consumption and driving activity behaviour in our test group of 6 individuals vs. a control group of 22 people.

6.1 Accuracy of our driving classifier

After optimising our SVM C and γ parameters on the training set of 40% of the 132,247 accelerometer readings, we achieved a "driving precision" score of just 0.4424. However by considering the inherent temporal re-occurrence nature of human activity data, and removing isolated positive/negative classified outputs through using median smoothing, we were able to boost the "driving precision" score to 0.8203. It is interesting to consider that a smoothing window of size of as little as 2 nearest neighbours significantly boosts the precision score, while thereafter increasing the window size doesn't have any significant effect. This more than likely indicates that the wrongly classified instances are isolated noise, rather than occurring in grouped blocks or chunks.

6.2 Accuracy of our driving CO_2 estimator

One of our subjects recorded his fuel consumption every week for a period of over 1 year. Using our driving detector we then applied this to the 58 week dataset of 9,370,647 logged accelerometer readings. After applying our SVM classification and post-processing smoothing, we were able to automatically predict the amount of time engaged in driving-like activities. Using a small subset of this data to tune parameters to map a correlation between the amount of time driving over a week and the amount of human recorded fuel consumption, where we found that almost 20 minutes of driving related to 1kg CO_2 on a Ford Focus 1.6l car. We were thus able to automatically estimate the user's CO_2 output from driving using the tri-axis accelerometer alone. After optimising our driving:CO_2 parameters on the 20 week training set, we now illustrate the results on our 38 week test set in Figure 7.

A striking feature is the two highlighted 2 week periods (2009 weeks 40-41, 2010 weeks 1-2) where our subject was on extensive trips abroad, where his car had no recorded CO_2 emissions, but he was in taxis which were picked up by our driving detector. This provides an estimation with a greater degree of entropy than the manually logged blog in these instances. Indeed on initial visual inspection the results appear quite promising with just a median difference of only 0.991 kgs/week (0.022 standard deviations) between the automatically estimated CO_2 values and the manual groundtruth. However we recognise that the deviation of differences is very large (85.38), and as illustrated in Figure 8 there is a large degree of variability between the individual weekly results, with some estimates being incorrect by almost over 3 standard deviations. Currently 64.71% of estimations had a degree of error to within 1 standard deviation, thus highlighting that many future challenges lie ahead in optimising the estimation of driving CO_2 from just wearable accelerometers. However we feel that this work also shows that there is much promise in exploring this avenue of research too.

6.3 Behavioural changes

Building upon the work above, we then supplied 6 users with a wearable accelerometer to put on their keyrings to capture their driving activity. As a different accelerometer was used for this task (Actigraph GT3x vs. SenseCam) we carried out the same training procedure on a small subset of data collected by the Actigraphs to build a relevant driving estimation model. In total 14,257,036 accelerometer readings were captured by our 6 subjects, 11% of which was classified as driving data. In addition this "test group" of 6 users also logged their home energy data over a period of approximately 9 months, which we will now compare to a "control group" of 16 users logging home energy data over the same period of time - thus seasonal effects of changes in energy behaviour are accounted for across both groups.

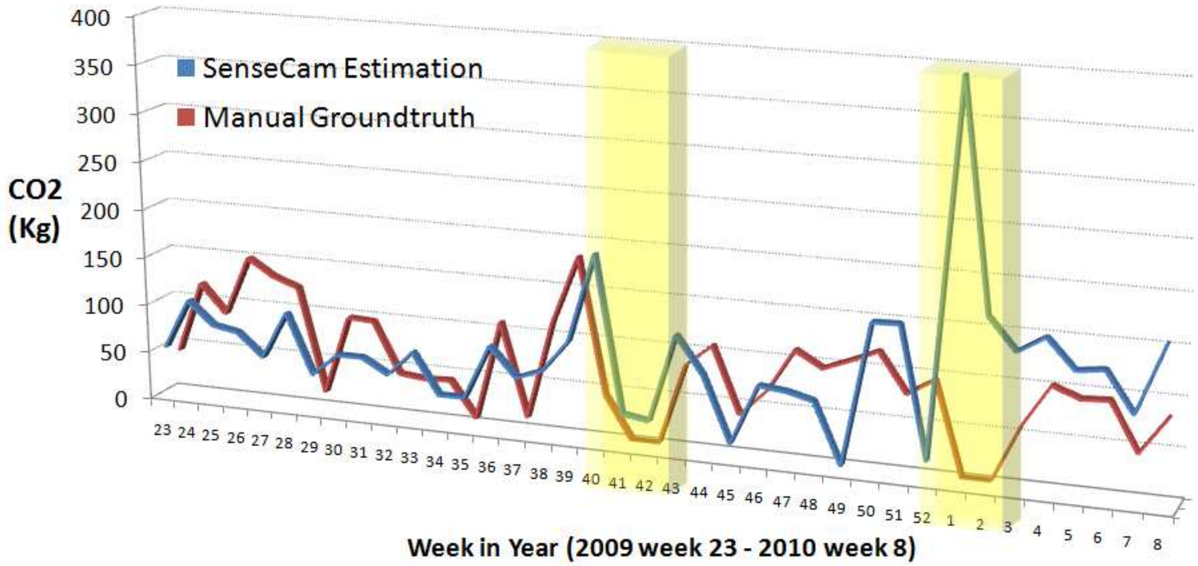

Figure 7: Comparison of our driving CO_2 estimation algorithm vs. a manual groundtruth

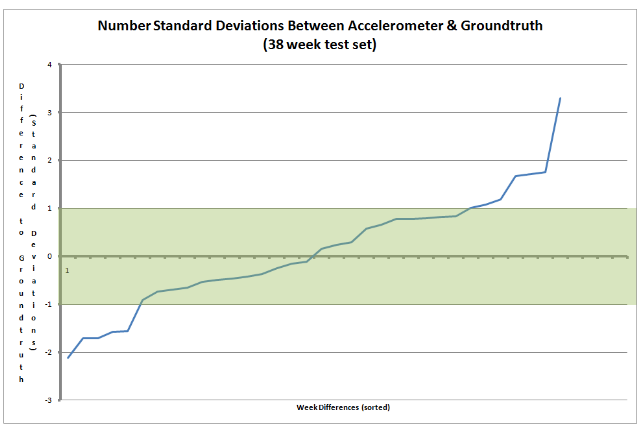

Figure 8: Overall accuracy of driving CO_2 estimation algorithm on 38 week test set, with close to y=0 being the ideal value

6.3.1 Driving Behaviour

Firstly we consider the amount of driving activity of the last 2 weeks vs. the first 2 weeks in our 6 week trial period[6]. As illustrated in Figure 9 we classified less of the logged accelerometer data as driving, thus indicating a possible change in user behaviour. A reason for this may be that the charts in our web and mobile applications displayed quite strikingly that CO_2 emissions from driving are very high, in relation to home electricity consumption. However we feel that this trial should be run over a longer period of time before making strongly asserted claims on changes in transport usage behaviour, to adjust for any effects such as the weather or public holidays like Easter, etc. Also more in-depth driving activity detection experiments will be carried out on the Actigraph GT3x accelerometer which may have different characteristics to the SenseCam. One possible reason for this is that the SenseCam is a wearable device and therefore is at rest when the wearer is driving. The Actigraph accelerometer on the keyring of the car key is more likely to be in motion as the car is moving.

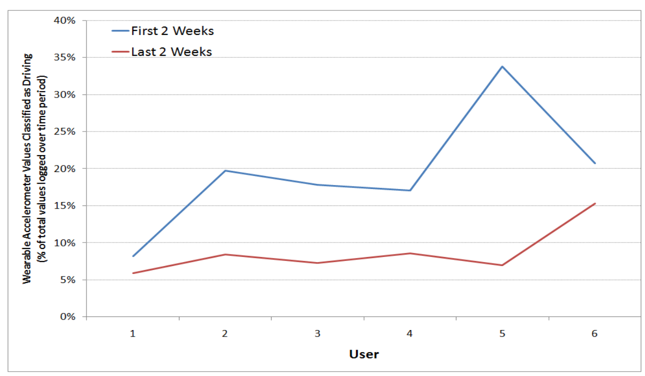

Figure 9: Indication of driving activity detected over 6 week trial period.

6.3.2 Home Electricity Behaviour

We now consider the home electricity consumption of our "test group" of 6 drivers over the 6 week trial, to see if they were influenced in any form by the feedback given by our web and mobile applications. As illustrated in Figure 12 there was an average saving of 0.656 kgs/day of CO_2 by our users in the last 2 weeks of the trial, as compared to the first 2 weeks. To contextualise this perceived saving, so as to allow for seasonal effects, the saving of 0.656 kgs/day (8.37%) in CO_2 home electricity made by our "test group" compared very favourably to the overall 0.077 kgs/day saving (1.35%) made by our "control group" (see Figure 11). This saving of

[6]Owing to the smaller amount of time that user 6 collected data, some days may overlap between "first 2 week" and "last 2 week" segments

8.37% over a short 6 week period of time suggests that our participants found the mobile and web applications to offer useful information, possibly through the additional feature of being offered an estimated figure for their driving CO_2 footprint.

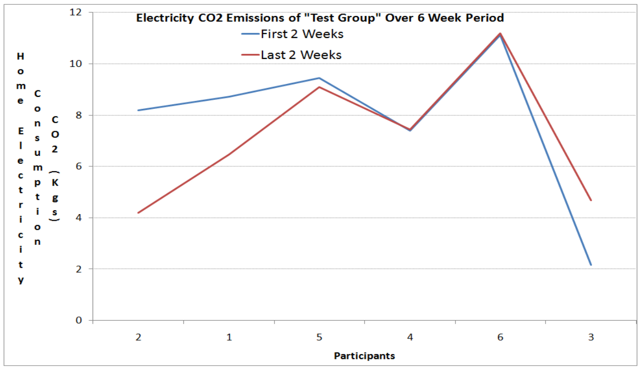

Figure 10: Electricity consumption over 6 week trial period of our "test group" of drivers

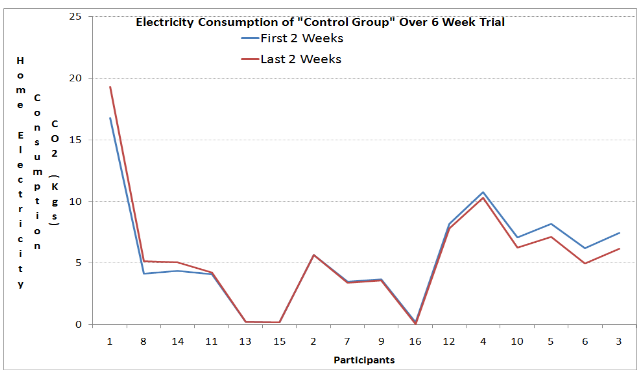

Figure 11: Electricity consumption over 6 week trial period of our "control group"

In terms of the effectiveness of the "Always-On-Display" given to both the control and test groups we now consider the home electricity consumption made by both the control and test group taken together. To achieve this we look at the average daily consumption over the last 4 weeks of the 6 week trial period (6.185 kgs/day CO_2) and compare it to the 4 weeks of electricity consumption immediately prior to the 6 week trial period (6.435 kgs/day CO_2). On investigation of Figure 12 we can see a modest saving of 3.88% here, albeit achieved in a very short time period, but already showing a change in energy consumption behaviour e.g. 6 users saved over 10%. Again any number of seasonal effects, public holidays, users going on holidays, etc. may come into play here, but our results do mirror other similar studies [7, 20].

7. DISCUSSION & FUTURE WORK

We view this paper as an initial, first step in a new research opportunity that the multimedia community can become engaged in, and make a positive contribution to. Based on our experiences thus far we have learned a few valuable lessons to guide the field in future endeavours.

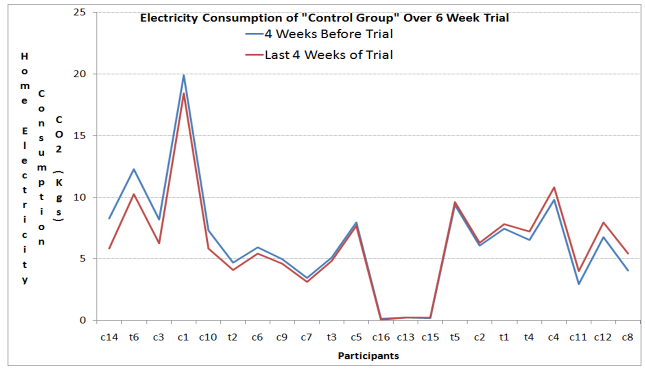

Figure 12: Electricity consumption at end of trial period vs. before trial period of all our participants

Visual lifelogging devices such as the SenseCam can be used to capture good quality annotations of human activity data, due to their inherent context re-instatement nature. For this work we tried asking individuals to record the time that they were driving in a written diary, but the users were unable to provide a good quality of information, a problem found in other studies [15], and therefore we couldn't rely on this annotation data to train our classifiers to learn activities. In contrast, the lifelogging device is always operational and does not suffer from any information quality issues. In future we will seek to integrate the accelerometer classifiers with image feature classifiers [4] for SenseCam captured visual images, so as to offer improved detection of driving events, an area the multimedia community is in the strongest position to contribute towards.

Currently our algorithm does not distinguish between a user being in a car or on a bus. A possible future mechanism to offer assistance here would be in passively logging the co-present Bluetooth devices in the area, which may offer a clue as to whether being on a car or bus, i.e. many other people/BT devices will be present when in a bus, while there may be no other people/BT devices while in a car. Alternatively the stop-start nature of a typical bus journey could be used for the same purpose. Longer term however it is quite likely that our wearable accelerometer (plus possibly BT logging) will be superseded by car manufacturers creating in-car-networks which will relay fuel and CO_2 emission information to any phone in the vicinity. Such an approach would provide the most accurate means of estimating CO_2 emissions, however for the interim and unto such standards are defined, we believe the wearable accelerometer route provides interesting research and commercial opportunities.

Thus far we have only passively engaged our users with regards to their electricity consumption. However we have noted that providing them startling personalised information has an effect on their energy usage, i.e. the "test group" appeared to drive less and also saved 8.37% in home energy consumption vs. 1.35% for our "control group". In future we believe the multimedia community can make a positive contribution to the global energy reduction agenda through designing thought-provoking systems that entice users to reduce their energy consumption habits. This may be achieved through social networking, thus offering an element of competition to act as a catalyst in sustainably reducing energy consumption [22]. Other options could include the use of smart textiles to ambiently offer feedback to constantly remind people of their energy usage e.g. a tablecloth that changes colour based on current home electricity consumption [18].

8. CONCLUSIONS

In this paper we have introduced a novel technique to automatically estimate an individual's CO_2 emissions from driving, solely though a cheap, simple wearable accelerometer. By offering this in conjunction with a home energy monitoring kit, we have moved towards the ability to offer individuals a more complete picture of their personalised carbon footprint which is ambiently monitored through 2 simple sensors. We have noted that individuals offered this additional functionality appeared to be more engaged with their energy usage, than against a control group of individuals who "only" had home electricity monitoring.

Acknowledgments

This material is based upon works supported by the Science Foundation Ireland under Grant No. 07/CE/I1147. We are also grateful to all the participants in the 22 locations logging electricity usage.

9. REFERENCES

[1] A. L. Allen. Dredging-up the past: Lifelogging, memory and surveillance. *New Yorker*, pages 38–44, 2007.

[2] M. Blum, A. S. Pentland, and G. Troster. Interest-based life logging. *IEEE Multimedia*, 13(4):40–48, 2006.

[3] W. F. Brewer. *Practical Aspects of Memory: Current Research and Issues*, chapter Qualitative analysis of the recalls of randomly samples autobiographical events, pages 263–268. Wiley, 1988.

[4] D. Byrne, A. R. Doherty, C. G. M. Snoek, G. J. Jones, and A. F. Smeaton. Everyday concept detection in visual lifelogs: Validation, relationships and trends. *Multimedia Tools and Applications*, October 2009.

[5] R. Church and C. Gellings. Reinventing renewable energy - toward a technology strategy for improving security, creating jobs & reducing emissions. Technical report, Electric Power Research Institute (EPRI) & American Council on Renewable Energy (ACORE), 2009.

[6] S. Coyle and D. Diamond. Smart nanotextiles: Materials and their application. *Encyclopedia of Materials: Science and Technology*, pages 1 – 5, 2010. published online.

[7] S. Darby. The effectiveness of feedback on energy consumption. Technical report, Environmental Change Institute, University of Oxford, 2006.

[8] EPA. 2010 u.s. greenhouse gas inventory report. Technical report, U.S. Environmental Protection Agency, 2010.

[9] EU. Eu action against climate change. leading global action to 2020 and beyond. Technical report, Luxembourg: Office for Official Publications of the European Communities, 2007.

[10] H. Herring. Energy efficiency–a critical view. *Energy*, 31(1):10 – 20, 2006. The Second Biennial International Workshop "Advances in Energy Studies".

[11] S. Hodges, L. Williams, E. Berry, S. Izadi, J. Srinivasan, A. Butler, G. Smyth, N. Kapur, and K. Wood. Sensecam: A retrospective memory aid. In *UbiComp: 8th International Conference on Ubiquitous Computing*, *LNCS* 4602, pp 177–193, 2006.

[12] M. Howley, B. O. Gallachoir, and E. Dennehy. energy in ireland 1990-2008. Technical report, sustainable energy ireland energy policy statistical support unit, 2009.

[13] T. Joachims. Svm light, http://svmlight.joachims.org, 2002.

[14] D. John, B. Tyo, and D. Bassett. Comparison of four actigraph accelerometers during walking and running. *Medicine & Science in Sports & Exercise*, 42(2):368, 2010.

[15] P. Kelly and C. Foster. Using photographs to understand health behaiours. Global Positioning Systems in Health Research Network, December 2009.

[16] H. Lee and A. F. Smeaton. Interaction platform-orientated perspective in designing novel applications. In *Create 2009 - Creative Inventions and Innovations for Everyday HCI*, 2009.

[17] A. Mannini and A. M. Sabatini. Machine learning methods for classifying human physical activity from on-body accelerometers. *Sensors*, 10:1154–1175, 2010.

[18] E. Mitchell, S. Coyle, A. R. Doherty, N. E. O'Connor, A. F. Smeaton, and D. Diamond. Smart tablecloths - ambient feedback of domestic electricity consumption. In *Irish Human Computer Interaction (iHCI 2010) Conference*.

[19] P. Petroni, M. Cotti, and O. Bono. The new edge for the enel telegestore: An integrated solution for the remote management of electricity and gas distribution allowing a total management of the energy consumptions. In *Electricity Distribution, 2009 20th International Conference and Exhibition on*, pages 1–5, June 2009.

[20] T. O. P. Project. Gridwise demonstration project fast facts. Technical report, Pacific Northwest National Laboratory, 2007.

[21] N. Ravi, N. Dandekar, P. Mysore, and M. L. Littman. Activity recognition from accelerometer data. In *Proceedings of the National Conference on Artificial Intelligence*, volume 20, page 1541. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2005.

[22] P. W. Schultz, J. M. Nolan, R. B. Cialdini, N. J. Goldstein, and V. Griskevicius. The constructive, destructive, and reconstructive power of social norms. *Pyschological Science*, 18(5):429–434, 2007.

[23] D. Velte and A. Martin. Energy services: Reducing the energy consumption of residents by behavioural changes. Technical report, BewareE, 2009.

[24] J. Yang. Toward physical activity diary: motion recognition using simple acceleration features with mobile phones. In *IMCE '09: Proceedings of the 1st international workshop on Interactive multimedia for consumer electronics*, pages 1–10, New York, NY, USA, 2009. ACM.