

Advice Provision for Energy Saving in Automobile Climate Control Systems

Amos Azaria^{1,*}, Ariel Rosenfeld², Sarit Kraus², Claudia V. Goldman³ and Omer Tsimhoni³

¹ Machine Learning department, Carnegie Mellon University, Pittsburgh, PA, USA

² Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel

³ General Motors Advanced Technical Center, Herzliya 46725, Israel

{azariaa1,rosenfa5,sarit}@cs.biu.ac.il, {claudia.goldman,omer.tsimhoni}@gm.com

Abstract

Reducing energy consumption of climate control systems is important in order to reduce human environmental footprint. The need to save energy becomes even greater when considering an electric car, since heavy use of the climate control system may exhaust the battery. In this paper we consider a method for an automated agent to provide advice to drivers which will motivate them to reduce the energy consumption of their climate control unit.

Our approach takes into account both the energy consumption of the climate control system and the expected comfort level of the driver. We therefore build two models, one for assessing the energy consumption of the climate control system as a function of the system's settings, and the other, models human comfort level as a function of the climate control system's settings. Using these models, the agent provides advice to the driver considering how to set the climate control system. The agent advises settings which try to preserve a high level of comfort while consuming as little energy as possible. We empirically show that drivers equipped with our agent which provides them with advice significantly save energy as compared to drivers not equipped with our agent.

Introduction

There is a growing interest in electrical cars. Since 2012 there has been an increase of 170% in electrical cars worldwide (as of July 2014) (Trigg and Telleen 2013; Pontes 2014). Yet, one of the most reported reasons for refraining exchanging to electrical cars is the relatively limited travel range it has (w.r.t petrol cars). Extending the travel range is of course desirable, it's economically beneficial and environmentally friendly - totally gas-free and tailpipe-emissions-free.

Auto experts from Edmunds.com and the Society of Automotive Engineers (SAE) have reported that the air conditioner reduces the cars fuel efficiency by up to 10 percent (Reed and Romans 2009; Hill et al. 2004). Thus, we propose an automated agent that advises the driver on how to set the car's climate control system, in a way that would reduce energy consumption while keeping the driver comfortable. We conducted this research in the summer time using

*This work was done while affiliated with Bar-Ilan University. Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Chevrolet GM Volt car interior.

a Chevrolet GM Volt car. During that time temperatures varied between 30-36°C.

Unfortunately, the agent and human user do not share the exact same goal. While the agent may care mostly about the car's energy consumption, the driver, on the other hand, is usually more interested in his own comfort level while less interested in the car's energy consumption. Thus, the agent faces the challenge of providing advice that will reduce energy consumption while taking into consideration the driver's comfort level, i.e., advice that will persuade the driver to set the system settings such that he reduces the energy consumption of the system.

The agent has to overcome two sources of uncertainty. First, it should try to model the preferences of the driver, estimating his comfort level in a given climate control setting. Second, it should estimate the energy consumption of a given setting. Both the drivers' preferences and the car's energy consumption are very noisy and difficult to estimate. Both models were built using data collected by running experiments in the Chevrolet Volt. Based on the constructed models we formalized the optimization problem of the agent, which wishes to minimize the energy consumption while maintaining a reasonable level of estimated comfort. We also designed a GUI that allows the agent to provide the advice in a convenient and attractive way for the driver.

We have conducted an extensive three different advice provision methods, with 49 human users who were required to set the climate control parameters of the Chevrolet Volt when it was very hot outside.

The proposed agent can be deployed in petrol and electrical cars, and as the results will show, has great long-term and short-term benefits in Chevrolet GM Volt car.

Related Work

Agents for the improvement of energy efficiency are a challenge for researchers and practitioners alike. Many works on the subject have been put forward, an example is Koehler et al. (2013), where the authors present automated approaches that can better match heating control to users routines and preferences. Al Mahmud et al. (2007) investigate the design and evaluation of the iParrot, an intelligent agent that helps to persuade family members to conserve energy in their home. Froehlich et al. (2009) suggest a mobile application that senses and reveals information about transportation behavior, in an attempt to persuade people to increase their use of green transportation.

Attempts to persuade people to change their behavior "for the better" is not restricted to energy saving. An example of such work is Consolvo et al. (2008), which by activity modeling peoples activities throughout everyday life they try to encourage physical activity. In his book (Fogg 2002), Fogg surveys many technologies trying to persuade humans, and analyzes the main properties required for such persuasion technologies to be successful. One example (pg. 50) is an exercise bicycle connected to a TV ("Telecycle"). In this system, as you pedal at a higher rate the image on the TV becomes clearer. This way the Telecycle encourages humans to exercise at higher rates. However, in most of the works, not only is the goal clear (exercise more or consume less energy), but so is the suggested way to achieve it. Therefore, the system is not required to provide *advice* as to how to achieve this goal, but merely persuade the user into doing so.

When considering advice providing systems (recommendation systems), they have been focused on predicting rating behavior by users in order to best provide them with recommendations. (See Ricci et al. (2011) for a review). Most works in this realm have only considered the utility of the users and minimize prediction error with respect to users' choices. Other works do explicitly consider the utility of the system (Chen et al. 2008; Azaria et al. 2013). These works build a user model, which allows the prediction of the probability that a user will accept a recommendation (or a set of recommendations). Using this prediction, they solve the optimization problem for the system in order to maximize its expected outcome. In all of these works, the user may either accept or reject the advice. Similarly, Azaria et al. (2012) model the long-term effect of advice given by a self-interested system on the users in route selection problems. Sarne et al. have shown that users' performance can be substantially improved through manipulating the input (e.g., the information concerning the different choices) that they receive (Sarne et al. 2011). Elmalech et al. suggest an approach according to which the decision regarding the advice

to be provided should not be made merely based on the encapsulated utility, but rather also based on the likelihood of its acceptance by the user (Elmalech et al. 2015). Das et al. (2009) theoretically analyze a recommender system trying to maximize its own expected utility. They assume the existence of some threshold in which, if the recommendations' quality is within the assumed threshold, the acceptance rate for the users remains the same. They analyze the benefit that the system may gain from providing recommendations, which are sub-optimal to the user but are close enough in quality and within the assumed threshold. Inspired by their work, we also set a threshold and assume that if the advice is above this threshold, the users are not likely to ignore it, but will be influenced by the advice received.

In our work, however, the challenge is doubled; the agent has to figure out it's own utility from every action as well as the human drivers' utility. We need to combine both these utility function into an advice provision model in which the human driver could be persuaded to save energy. To the best of our knowledge, no persuasive work thus far has focused on automobile climate control systems.

The Volt Climate Control System

The study in this paper was based on the Volt's climate control system. In this system the drivers can control the settings S as described in this tuple (T, F, D, M) where:

- **Temperature** (T) is associated with a temperature in Celsius and can receive values between 16 and 35 degrees;
- **Fan strength** (F) is associated with the fan blower and can receive values between 1 and 6;
- **Air delivery** (D) may either be set to face (in which D is set to 0) or face and feet (in which D is set to 1);
- and **Mode** (M) may either be set to "eco" (when M is set to 0) or to "comfort" (when M is set to 1). According to the Volt's user manual, the 'eco' mode tries to reduce energy consumption, while the "comfort" mode aims at maximizing the user's comfort level.

Given a setting s we use subscript s_T to refer to the temperature in that setting, s_F to refer to the fan strength, s_D for the air delivery and s_M for the mode of the setting.

Figure 2 presents the original climate control system panel (from the user manual). Figure 3 provides a short description for each of the variables (from the user manual).

CARE

In this section we present our Climate control Adviser for Reducing Energy consumption (CARE). CARE requires the composition of two models, one for modeling the climate control's energy consumption as a function of its settings and the other for modeling human comfort level as a function of the climate control's settings. CARE uses these models in order to provide a driver with advice regarding the settings of the climate control system, taking into account both the expected energy consumption and the expected comfort level. The comfort level is captured by a number between 1 to 10 where:




Figure 2: Original Volt climate control panel, taken from the user manual.

| Climate Touch Screen Controls | |
|---|--|
| A. Driver and Passenger Auto Heated Seats | F. Climate Modes: Fan Only, ECO, Comfort |
| B. Air Delivery Mode Controls | G. Climate Power Gauge |
| C. Auto Recirculation | H. Manual Fan Control |
| D. Manual Recirculation | I. Temperature Setting Display |
| E. Outside Air Temperature Display | J. Air Conditioning Indicator |
| | K. Heat Status Indicator |

Figure 3: Description of the Volt’s climate control variables.

- 1: "I’m very uncomfortable; I would not be willing to drive under these conditions.”;
- 3: "I’m uncomfortable, but I might be willing to compromise.”;
- 5: "Reasonable, I would be willing to drive under these conditions.”;
- 7: "I’m comfortable; I would like to drive under these conditions.”; and
- 10: "I’m most comfortable, I would be happy to drive under these conditions.”

CARE Training Data

Constructing CARE requires two sets of training data: ψ_e and ψ_c . ψ_e is used to train the parameters for the energy consumption model. (see Section). It is composed of a tuple with the following format for every instance i : $\psi_e^i = (e, T, F, D, M, E, I)$ where e is the energy consumption level, given the other parameters; T, F, D and M are the variables set on the climate control system; E is the external temperature as displayed in the dashboard; and I is the internal temperature as measured with a manual thermometer located between the 2 front seats. Both the external and internal temperatures could be viewed by the drivers.

ψ_c is used to train the parameters for the comfort model. (see Section). It is composed of a tuple with the following format for every instance i : $\psi_c^i = (c, T, F, D, C, E, I)$ where c is the comfort level reported by the subject, given the other parameters; C is the initial comfort level, i.e. the comfort level reported when the driver enters the car; and all

other parameters are as described in ψ_e ¹.

Energy Consumption Model

We model the energy consumption of the climate control system based on the following equation:

$$e(T, F, D, M, E, I) = (w_1 \cdot (-T) + w_2 \cdot F + w_3 \cdot D + w_4 \cdot E + w_5 \cdot I) \cdot ((1 + w_6) \cdot M) \quad (1)$$

where w_1, w_2, \dots, w_6 are parameters learned by the model. This form of function assumes that all variables except the climate mode have a linear impact on the final energy consumption. The climate mode is assumed to have a multiplicative impact on the total energy consumption, since in the “comfort” climate mode, all of the climate control components seem to work harder and thus consume more energy. This form of function was compared to other forms and yielded the best fit to the data collected². All parameters are assumed to be positive, except w_3 which models the impact of air delivery on energy consumption. w_3 was allowed to obtain negative values and in fact it did end up with a negative value. We use the training data, ψ_e , and search for the parameters w_1, w_2, \dots, w_6 which maximize the likelihood of the training data (maximum likelihood estimation). We use interior point methods (Nesterov, Nemirovskii, and Ye 1994) to search these parameters.

Human Comfort Level Model

CARE’s model for the human comfort level is based on the following equation:

$$c(T, F, D, C, E, I) = v_0 - v_1 \cdot T + v_2 \cdot F - v_3 \cdot F^2 - v_4 \cdot D + v_5 \cdot C - v_6 \cdot E - v_7 \cdot I \quad (2)$$

where v_0, v_1, \dots, v_7 are parameters learned by the model. F^2 tries to capture the effect of the noise created by the fan, which is super-linear in the fan’s level. The human comfort level model assumes that the human comfort level is a linear combination of all of the parameters that the human faces (assuming that F^2 models the noise effect). This assumption is common in the literature (Nguyen et al. 2013; Azaria et al. 2011). According to the car’s user manual, the ‘eco’ mode is supposed to save energy, therefore, CARE never recommended to set the mode to “comfort”, and we only gathered data on subjects’ comfort level when using the ‘eco’ mode. For that reason, the human model does not take the mode into account, and only tries to predict the comfort level for when the mode is set to “eco”. We use the training data, ψ_c , and search for the parameters v_0, v_1, \dots, v_7 which maximize the likelihood of the training data (maximum likelihood estimation). We use again the interior point method to search these parameters similar to the search performed in Section for finding the suitable parameters for the energy

¹Notice that the mode, M , does not appear in the comfort level model; this attribute will be explained later.

²Some of the other functions that were tested included one or more of the following modifications to the above function: the use of M as an additive variable; F as having a multiplicative impact or T as having an impact depending on its offset from I or E .

consumption model. Note that the initial comfort level (C) may change from person to person. This will cause the expected comfort level to vary among people, and thus also the advice provided by CARE may vary among different people. This causes the advice to be personalized, i.e. different drivers may receive different advice. However, it is possible that people reporting the same comfort level in a given setting will have slightly different preferences that can be used for further energy saving. Furthermore, it will be preferred if the driver's preferences will be learned without her need to explicitly report her comfort level. These two improvements (among others) can be done only when the system interact repeatedly with the user. See section Current Work for our work on the subject.

CARE Method for Advice Provision

Given both the energy consumption model and the human comfort level model, CARE provides the driver with advice regarding the settings of the climate control system. Given the external temperature (E), the internal temperature (I) and the initial comfort level (C), CARE provides the driver with advice, $a(E, I, C) \in S$, that yields an expected comfort level of at least 7 while minimizing the expected energy consumption of the climate control system. CARE only considers advice in which the mode is set to "eco" (i.e. M is set to 0). Comfort level 7 was chosen as the minimal target comfort level since a comfort level of 7 means that the driver is comfortable. More formally, CARE provides advice such that:

$$a(E, I, C) = \arg \min_{s \in S} e(s_T, s_F, s_D, M, E, I) \text{ s.t.} \quad (3)$$

$$s_M = 0; c(s_T, s_F, s_D, C, E, I) \geq 7$$

where $e(\cdot)$ is obtained from Equation 1, and $c(\cdot)$ is obtained from Equation 2. Since the search space is small ($|S|$ is much smaller than 1000), we perform an exhaustive search to find the optimal advice. However, in a climate control system with additional variables, CARE may consider a more efficient method of search.

Algorithm 1 presents an overview for the entire procedure of the construction and usage of CARE.

Data Collection

In order to train the Energy Consumption and Human Comfort Level models we used a Chevrolet GM Volt car parked (idle) in a closed parking lot at General Motors Advanced Technical Center in Herzliya, Israel (GM ATCI). The parking lot was chosen due to its stable temperature, and the fact that it is shaded at all times. These conditions were repeated in the actual experiment described in the Experimental Evaluation section.

Data Collection for Modeling Energy Consumption

The energy consumption of setting s is the sum of two factors: Energy consumed by the blower (the fan) and the energy consumed by the compressor. The data was collected directly from the car's feedback using 120 measurements, each measurement was a 10-minute episode in which the

```

Find  $w_1, w_2, \dots, w_6$  for Equation 1 which maximize the
likelihood of the training set  $\psi_e$ .
Find  $v_1, v_2, \dots, v_7$  for Equation 2 which maximize the
likelihood of the training set  $\psi_c$ .
Obtain initial comfort level ( $C$ ) from driver.
Obtain external ( $E$ ) and internal ( $I$ ) temperatures.
 $min \leftarrow \infty$ 
for each  $s \in S$  where  $s_M = 0$  do
  if  $e(s_T, s_F, s_D, s_M, E, I) < min$  and
   $c(s_T, s_F, s_D, C, E, I) \geq 7$  then
     $advice \leftarrow s$ 
     $min \leftarrow e(s_T, s_F, s_L, M, E, I)$ 
  end if
end for
return  $advice$ 

```

Algorithm 1: Construction and use of CARE.

climate control system was on (resulting in a total of 20 hours). We were interested in the total energy consumption in each of these 10-minute episodes (and not momentary energy consumption which varied a lot). To maintain the integrity of the measurements, we've let the car warm up (and the compressor cool down) for 10 minutes between consecutive measurements. The measurements were for various temperatures, starting at $T = 16$ and up to $T = 26$, and various fan speeds, starting at $F = 1$ up to $F = 5$. The measurements were used to train ψ_e .

It was encouraging to observe that there are settings where a large percentage of energy can be saved. For example, when the temperature in the car and outside the car is $26^\circ C$, then the energy consumption when setting the climate control system temperature to $16^\circ C$, the fan to 5 and the mode to "comfort" is 75% higher than when setting the temperature to $22^\circ C$, the fan to 1 and the mode to "eco". Of course, this is an extreme case.

According to our measurements, when all other settings are identical, 'eco' indeed consumes less energy than "comfort".

The final function obtained was: $e(T, F, D, M, E, I) = (-0.0095T + 0.016F - 0.003D + 0.005E + 0.005I) \cdot (1.17M)$.

Data Collection for Modeling Human Users

Collecting data to train the human model (ψ_c) is far more difficult than the data collection for the energy consumption model. We had to find subjects which were willing to enter the parked car (several times) in a hot day, and experience CCS levels which were not necessarily those that were the most convenient for them. As we will describe here-under, the subjects were also required to wait outside the car between measurements, so the car could warm-up while the car warmed up between.

We want to get as many instances as possible, due to the high cost of recruiting subjects. These instances should preferably be in the range of plausible settings. We recruited 15 subjects for training the Human Model, out of which 4 subjects were females and 11 were males. The subjects'

ages ranged from 21 to 73, with a mean of 30 and a median of 27. All subjects live in Israel. The subjects were first asked to fill out a questionnaire collecting demographic information. Then the comfort level scale was explained to them (See the CARE Training Data section).

The subjects were asked to enter the car and sit in the driver's seat with their hands on the steering wheel. While the climate control system was still off, the subjects were asked to set the vents to point in their direction. At this point, the subjects were asked to rate their comfort level. Then the subjects were told how to operate the climate control and were asked to set it so that they would feel most comfortable. The selected settings were left on for 4 minutes. The subjects were then asked for their comfort level and were required to explain why they had chosen that level. The subjects were then asked to exit the car and the car was left to warm up for an additional 4 minutes.

In order to test as many settings as possible, the subjects then returned to the car and the experiment operator set additional 8 settings for them, each was left on for 4 minutes followed by the subjects rating their comfort level. Between one setting to the next the subject was asked to stay outside the car for 4 minutes while the car warms up. and waited 4 minutes. This process resulted in 120 instances – 15 subjects, each provided 8 instances.

The subjects' comfort levels seem to have been mostly influenced by the temperature that was set on the climate control system, T . The fan, F , also had an impact on the comfort level, though not as strong as the temperature. Recall, that the opposite phenomena occurred when modeling the energy consumption level. This result motivated CARE to advise settings with the fan set to low values. Most subjects reported a reduced comfort level when the fan was too strong, some reported that the noise was what bothered them. The other parameters seemed to have a milder impact on the subject's comfort level. The final formula for the Human Comfort Level model is:

$$c(T, F, D, C, E, I) = 16.608 - 1.2995T + 0.9841F - 0.0642F^2 - 1.2188D + 0.3238C - 0.1727E - 0.4817I$$

Experimental Evaluation

To see if CARE can reduce energy consumption we implemented a panel based on the original climate control panel in the VOLT car (Figure 2), with additional add-ons and functionality. We tested CARE against two agents; a "Silent" agent – which does not offer any advice yet records the subjects choices and energy consumption, and CAREless, which only provides information on the energy consumption the current setting produces.

- **Silent:** As shown in Figure 4, the GUI of this agent is based to the original climate control panel in the VOLT car. The silent agent merely presents the current climate control settings and records the drivers actions (changes in AC settings) and energy consumption (as seen from the car's data). The driver neither receives any information nor any advice.

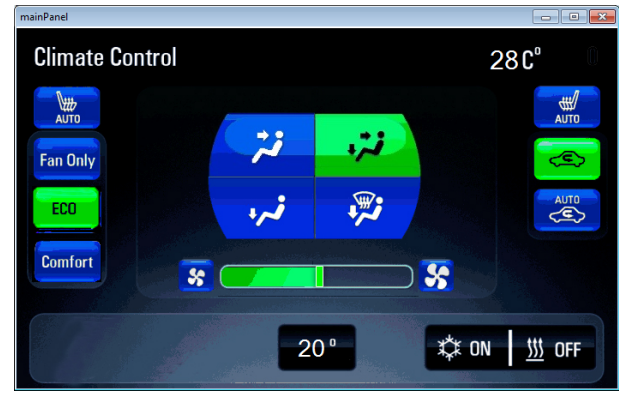


Figure 4: A screen-shot of the GUI when no advice is provided.



Figure 5: A screen-shot of the GUI with additional energy consumption information provided by CAREless (the circle in the bottom left corner).

- **Careless:** As shown in Figure 5, on top of the silent agent functionality, CAREless has an additional information circle, presented in the bottom left, which supplies the driver with an estimation of the current energy consumption level. This information appears as the percent of the current energy consumption from the maximum energy consumption obtained in the training data (the lower the better). Note that CAREless does not provide any active advice. In Figure 5 we can see an example where the current consumption is 40% of the maximum.
- **CARE:** On top of CAREless functionality, CARE provides advice. As soon as the driver gets into the car s/he is presented with an advice indicated in purple (See Figure 7). The driver can set the climate control in any settings he/she desires, which is not necessarily the advice provided by CARE. Nevertheless, s/he is presented with an estimate of the current energy consumption (exactly as provided by CAREless) as well as the advice with its projected energy consumption indicated in purple in Figure 6. These add-ons are not present in CAREless, as there is no advice. Figure 8 shows a screen-shot of a case in which the driver set the climate control system to match

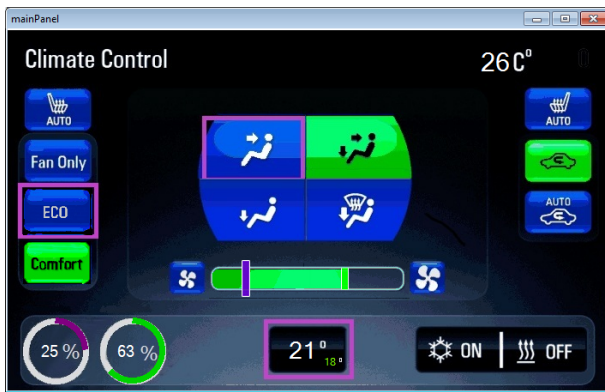


Figure 6: A screen-shot of the GUI when the driver set the climate control system differently from CARE's advice. In this example the driver decided to set the temperature to 18°C (rather than 21°C as advised by CARE), the fan to 4 (rather than 1), the air delivery to both face and feet (rather than only to the face) and the mode to 'comfort' (rather than 'eco'). This resulted in an energy consumption level of 63% of the maximal energy consumption level, rather than only 25% if the driver would have followed CARE's advice.

the advice.

Methodology

People have different preferences when it comes to climate control; they vary in their preferred temperature, fan strength, etc. Some of these differences are physiological, larger people tend to gain and lose heat more slowly than smaller ones, by virtue of their smaller surface-area-to-volume ratios. Also, it has been shown in Wunderlich's studies in the early 19th century that women tend to have lower core temperatures than men (Wunderlich 1870)³, regardless of differences of height and weight. These mentioned differences, as well as others that were studied such as occupation, place of birth and others, impose a big concern when testing agents between-subjects. Not only that, external factors in our experimental environment tend to change, i.e., the external temperature (E). Although the study was conducted in the summer time temperatures were not constant, varying from 30 to 36°C. Therefore, in order to control this variance, we chose an experimental design that examined the effect of advice as a within-subject variable rather than a between-subject variable, thus overcoming the mentioned challenges.

First, we recruited 49 Israeli subjects, 33 males and 16 females, aging from 21 to 73 (mean 35, median 31)⁴. Each subject was asked to fill out forms and demographic data. Then, the subject was led to a Chevrolet GM Volt car parked in GM ATCI. We had each subject run the experiment twice, once with the silent agent (as a baseline) and once with either CARE or CAREless. We counterbalanced the order among the type of experiments, i.e. approximately half of the sub-

³ Available online as Google eBook

⁴ All experiments with human subjects were approved by the corresponding IRB.



Figure 7: A screen-shot of the GUI with advice from CARE. Note that the advice is shown in purple. The left purple circle represents the expected energy consumption from following the advice. The right circle represents the current energy consumption is empty, since the driver did not yet turn on the climate control system.



Figure 8: A screen-shot of the GUI when the driver set the climate control system to match CARE's advice. The right green circle represents the current user energy consumption and the left purple circle represents the advice's expected energy consumption.




Figure 9: The mean energy consumption level of the subjects who were assigned to CARE and CAREless agents, compared to the mean energy consumption levels of the subjects when they did not receive any advice.

jects first ran the experiment with no advice. 24 subjects were assigned to receive advice from CARE, while 25 subjects were assigned to receive the information provided by CAREless (randomly).

At each phase, the experiment operator asked the subject’s initial comfort level (denoted C in our model). Then the subject was given 10 minutes to be in the Volt car (parked in ATCI) and he or she was free to tell the experimenter what settings to set in the climate control system. The GUI of the designated agent was displayed on a laptop, while the built-in car display was covered to avoid distractions. The experimenter updated the climate control of the car as many times as requested by the driver. While in the car, the subject was given a smart phone with a driving simulator “Bus Simulator 3D”⁵ to be played while the experiment goes on. The motivation was to set the conditions similar to regular driving and give the subjects something to do. After 10 minutes, the subject had to go outside the car and wait until inside of the car gets warm again to simulate initial conditions. To that aim, we left the doors and the trunk open for 10 minutes while the car was turned off. Then the second stage was examined for another 10 minutes, in the exact same fashion.

The process took about 40 minutes per subject (including the paperwork and instructions), for which we paid each subject 100 NIS (approximately 27\$). In real terms, 100 NIS is the price of a fancy lunch in Israel.

Results

The results were analyzed using repeated measures of ANOVA with total energy consumption as a dependent variable, silent (True/False) as a within-subject variable, type of agent (CARE/CAREless), gender of the subject and order of presentation (baseline, first or second) as between-subject variables. Thus, the statistical model had one within-subject factor and three between-subject factors. The statistical analysis revealed no significant findings, except a trend suggesting that the effect of the agent depended on the

type (either CARE or CAREless). We therefore ran separate analyses for each of the two advice types.

When subjects were given advice by the CARE algorithm, their total energy consumption significantly decreased from 0.24 KWH to 0.20 KWH, an improvement of 17% ($F(1, 21) = 7.6, p < 0.05$). We also corrected for multiple comparisons, and after the Bonferroni correction, the type-I error remains < 0.05 . This improvement amounted to a mean energy savings described in the 95% confidence interval: $[-24\%, -5\%]$. The effect of presentation order and its interaction with the effect of advice were both not significant. A similar analysis for the CAREless advice did not show any improvement in total energy consumption ($F(1, 23) = 0.12$). Figure 9 presents the mean energy consumption level of the climate control system, which was obtained by the subjects who were assigned to CARE or CAREless, compared to the mean energy consumption level of the same subjects when they did not receive any advice at all.

Figure 10 shows the energy consumption level of the climate control system of each subject when receiving advice from CARE compared to the baseline of that same subject when not receiving any advice. As illustrated by the figure, 19 out of the 24 subjects have shown an improvement over their baseline when receiving advice from CARE (their associated points appear under the 45° diagonal). The figure also shows that for three subjects, CARE reduced energy consumption by approximately 50% (from approximately 0.25 KWH to approximately 0.12 KWH).

When comparing men to women in no advice condition, it turns out that females tend to consume less energy than males, 0.201KWH vs 0.242KWH, which fits the common myth that women like the air conditioner weaker than men.

In order to ensure that the advice provided to the user is easy to understand, we asked the subjects the following question: “Was the information on the screen clear?” and asked them to specify a number between 1 and 10. The average answer was 9.15, indicating that the GUI is very understandable.

Discussion

As shown by the results, CARE significantly decreases the energy consumption of its users while CAREless did not do as well (w.r.t no advice condition). The conclusive statistical finding at 95% confidence that Care was better than the baseline is based on a within subject experiment and model. Astute readers might notice that the baseline values for the Care and careless conditions were not the same and might ask themselves if this difference might have affected the results. The beauty of within subject analyses is that even when baseline values are different, the comparisons of conditions are done per each individual and therefore are robust to variations among subjects. In our case, we found a statistical difference between Care and the baseline and did not find such differences for the other condition. This statistical analysis also captures any other effects which might have been in play, such as the psychological effect which the displayed internal and external temperatures may have had on the subjects.

⁵Available free in Google Play store.




Figure 10: The energy consumption level of the climate control system of each subject when receiving advice from CARE compared to the baseline of that same subject when not receiving any advice.

We notice that the advice provided by CARE is not in a "take it or leave it" fashion, rather than it is a reference point, an anchor if you will, for the users to select their settings by. For example, one of the subjects, when receiving no advice, set the climate control system to a temperature of 23°C and the fan to 4. However, when that same subject received the advice from CARE to set the temperature to 24°C and the fan to 1, she set the temperature to 24°C as suggested, but set the fan to 2. Later, when she became a little too warm, she set the fan to 3, and the temperature setting at 24°C . The effect seen here can in part be attributed to the anchoring effect. This effect proposes that when people do not know the exact value of a product (or answer to a question), if they are first shown a possible value (or answer) which was randomly generated, then later, when they need to evaluate the product (or give their own answer to the question), their evaluation (or answer) is relatively close to the original evaluation (or answer) (Ariely, Loewenstein, and Prelec 2006; Johnson and Schkade 1989). Obviously, though, a system may not rely solely on the anchoring effect, since a system will lose all its credibility by offering an unreasonable value (such as setting the temperature to 30°C).

Notice, that our study has focused on summer conditions, and can be easily expanded to winter conditions as well. Providing advice which depends on the actual weather seems essential for people to treat the advice seriously. Also, the CARE methodology can be implemented on both electrical and petrol cars alike, yet measurements in each car are needed to estimate the models correctly.

Clearly, CARE has reduced her energy consumption, which in fact translate naturally to longer travel distance and lower electricity cost. This also suggests a lower fuel consumption levels for a petrol car implementation.

The methodology presented in this paper can be extended

to include other automotive systems such as the navigation or adaptive cruise control (ACC) systems. These systems share a common characteristic with the climate control system studied in this work—the system and human user do not share the exact same goal. For example, while the ACC system may care mostly about the car's energy consumption and driver's safety, the driver, on the other hand, is usually more interested in reducing travel time. Thus, the agent faces the challenge of providing advice that will be beneficiary for the system while accounting for what the driver considers acceptable. Rosenfeld et al. (Rosenfeld et al. 2012) have showed that learning drivers' behavior can improve the use of the ACC system. Thus, combining their model with a model quantifying the system's goals can bring about an even better use of the ACC system.

Current Work

We have recently finished testing agents that interact repeatedly with their users to save energy in a platform similar to the one used for the one-time advice provision. Among the challenges we faced when extending the work to the repeated interactions case, is the notion of building trust between the agent and the driver, personalizing the advice even when the comfort level is not provided, taking into account long term effect of an advice and a new graphical user interface. Experiments run so far show that these agents were able to compute personalized advice even when the users were not asked for their initial comfort level and were adaptive to their drivers' actions. For example, when a driver rejected an advice, the agent would adapt its computations so that whenever it provided a new advice, it would be better suited for that driver. And furthermore, when a driver accepted an advice, the agent tries to improve the energy saving of the car by providing a new advice in the next interaction given the past behavior of the driver. The results from testing our repeated-interaction agent suggest that although tested in a hotter environment than CARE (external temperatures were $35 - 37^{\circ}\text{C}$), the average energy consumption at each 10 minute episode is 33% lower than the "no advice" case tested. Moreover, the new agent outperformed CARE on average 0.17KWH vs. 0.2KWH. See Rosenfeld et al. (2015) for a report about these new models and results.

Conclusions

In this article, we presented a method to persuade a driver to reduce the energy consumption of the climate control system of his electrical car. Via experiments, we showed that the proposed methodology leads to a significant reduction of energy consumption. The methodology requires the collection of data on the energy consumption of the climate control system and on the drivers' behavior, but is effective even with a small number of examples (15 drivers in our experiment). We designed a GUI for presenting the advice that facilitates understanding of the advice. The reported work is the first step in the process of the deployment of a persuasive agent in electrical and petrol-fueled cars alike. As discussed above, the results and insights obtained in this study were used once again to accomplish an additional step

towards real deployment. We hope that this work will inspire researchers and practitioners to translate the proposed methodology into persuasive and advice-providing agents in other fields.

Acknowledgments

We would like to thank Amihai Gottlieb, Jonathan Azaria and Shira Abuhatzera for their help in this research.

References

- Al Mahmud, A.; Dadlani, P.; Mubin, O.; Shahid, S.; Midden, C.; and Moran, O. 2007. iparrot: towards designing a persuasive agent for energy conservation. In *Persuasive Technology*. Springer. 64–67.
- Ariely, D.; Loewenstein, G.; and Prelec, D. 2006. Coherent arbitrariness: Stable demand curves without stable preferences. *Journal of Business Research* 59(10-11):1053–1062.
- Azaria, A.; Rabinovich, Z.; Kraus, S.; and Goldman, C. V. 2011. Strategic information disclosure to people with multiple alternatives. In *AAAI*.
- Azaria, A.; Rabinovich, Z.; Kraus, S.; Goldman, C. V.; and Gal, Y. 2012. Strategic advice provision in repeated human-agent interactions. In *AAAI*.
- Azaria, A.; Hassidim, A.; Kraus, S.; Eshkol, A.; Weintraub, O.; and Netanel, I. 2013. Movie recommender system for profit maximization. In *RecSys*.
- Chen, L.; Hsu, F.; Chen, M.; and Hsu, Y. 2008. Developing recommender systems with the consideration of product profitability for sellers. *Information Sciences* 178(4):1032–1048.
- Consolvo, S.; McDonald, D. W.; Toscos, T.; Chen, M. Y.; Froehlich, J.; Harrison, B.; Klasnja, P.; LaMarca, A.; LeGrand, L.; Libby, R.; et al. 2008. Activity sensing in the wild: a field trial of ubifit garden. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1797–1806. ACM.
- Das, A.; Mathieu, C.; and Ricketts, D. 2009. Maximizing profit using recommender systems. *arXiv preprint arXiv:0908.3633*.
- Elmalech, A.; Sarne, D.; Rosenfeld, A.; and Erez, E. S. 2015. When suboptimal rules. In *AAAI*.
- Fogg, B. 2002. *Persuasive Technology: Using Computers to Change What We Think and Do (Interactive Technologies)*. Morgan Kaufmann.
- Froehlich, J.; Dillahunt, T.; Klasnja, P.; Mankoff, J.; Consolvo, S.; Harrison, B.; and Landay, J. A. 2009. Ubigreen: investigating a mobile tool for tracking and supporting green transportation habits. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1043–1052. ACM.
- Hill, W.; Lebut, D.; Major, G.; and Schenkel, F. 2004. Affect of windows down on vehicle fuel economy as compared to ac load. Technical report, IEA.
- Johnson, E. J., and Schkade, D. A. 1989. Bias in utility assessments: Further evidence and explanations. *Management science* 35(4):406–424.
- Koehler, C.; Ziebart, B. D.; Mankoff, J.; and Dey, A. K. 2013. Therml: occupancy prediction for thermostat control. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, 103–112. ACM.
- Nesterov, Y.; Nemirovskii, A. S.; and Ye, Y. 1994. *Interior-point polynomial algorithms in convex programming*, volume 13. SIAM.
- Nguyen, T.; Yang, R.; Azaria, A.; Kraus, S.; and Tambe, M. 2013. Analyzing the effectiveness of adversary modeling in security games. In *AAAI*.
- Pontes, J. 2014. World Top 10, May 2014. Technical report, EVSales.com.
- Reed, P., and Romans, B. 2009. Affect of windows down on vehicle fuel economy as compared to ac load. Technical report, edmunds.com.
- Ricci, F.; Rokach, L.; Shapira, B.; and Kantor, P., eds. 2011. *Recommender Systems Handbook*. Springer.
- Rosenfeld, A.; Bareket, Z.; Goldman, C. V.; Kraus, S.; LeBlanc, D. J.; and Tsimhoni, O. 2012. Learning driver's behavior to improve the acceptance of adaptive cruise control. In *IAAI*.
- Rosenfeld, A.; Azaria, A.; Kraus, S.; Goldman, C. V.; and Tsimhoni, O. 2015. Adaptive advice in automobile climate control systems. In *Proc. of The AAI Workshop on Artificial Intelligence for Transportation WAIT-15*.
- Sarne, D.; Elmalech, A.; Grosz, B. J.; and Geva, M. 2011. Less is more: restructuring decisions to improve agent search. In *AAMAS*, 431–438.
- Trigg, T., and Telleen, P. 2013. Global ev outlook 2013 - understanding the electric vehicle landscape to 2020. Technical report, International Energy Agency.
- Wunderlich, C. A., ed. 1870. *Das Verhalten der Eigenwärme in Krankheiten*. Otto Wigand.