

# Advice Provision for Energy Saving in Automobile Climate Control Systems

Amos Azaria<sup>1</sup> and Sarit Kraus<sup>1</sup> and Claudia V. Goldman<sup>2</sup> and Omer Tsimhoni<sup>2</sup>

<sup>1</sup> Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel

<sup>2</sup> General Motors Advanced Technical Center, Herzliya 46725, Israel  
{azariaa1,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

In the increasingly industrialized world, various facts imply that energy consumption levels may no longer be overlooked. In addition to long term reasons, saving energy while driving electrical cars has an additional short-term benefit—it extends the range of travel. This is desirable since electric cars often have a shorter driving range than fuel-powered cars per full battery charge. Thus, in this paper, we propose an automated agent that advises a driver on saving energy. In particular, we consider a summer environment in a Chevrolet GM Volt car, in which the car is very warm and the driver would like to turn on the climate control system in order to cool down the car and drive comfortably. An agent advises the driver how to set the car's climate control system. In this scenario the computer 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.

Our 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. The data for building the drivers' model was collected from only 15 participants. 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. In order to evaluate our agent, we ran experiments with 49 human users who were required to set the climate control parameters of the Chevrolet Volt when it was very hot outside. We tested three different types of advice provision methods. The first *no advice* method did not provide any advice and presented the subjects with an interface similar to the original Volt climate control system. The second *energy info* method provided the subjects with information regarding their current energy consumption level (based on the energy consumption model we built). The third *agent* method provided the subjects with advice on how to set the climate control system, along with the energy consumption information. We show that, on average, the subjects consumed less energy when interacting with the *energy info* or *agent* method vs. the *no advice* method. However, these differences reach statistical significance only when comparing the *agent* method with the *no advice* method. We show that when using our agent, the subjects saved approximately 17% of the energy consumption of the climate control system.

## Related Work

In his book (Fogg 2002), Fogg surveys many technologies trying to persuade humans, and analyzes the main proper-

ties 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. Similarly Froehlich et al. (Froehlich, Findlater, and Landay 2010) survey many persuasive technologies with a goal of reducing environmental impact. However, in most of these 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.

Several works in recommendation systems have predicted rating behavior by users in order to best provide them with recommendations. (See Ricci et al. (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. (Azaria et al. 2012) model the long-term effect of advice given by a self-interested system on the users in route selection problems. Das et al. (Das, Mathieu, and Ricketts 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. However, all the above approaches are not applicable to our work, since in our work the user may not only accept or reject the advice, but may also be influenced by the advice and may choose an option that is on the one hand different from the advice, but on the other hand different from what the user would have chosen if he had received no advice. Therefore, building a model that would only model the acceptance rate of the user would not have been useful in our work.

### 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.

## 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:

- 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:  $\psi_e$  and  $\psi_c$ .  $\psi_e$  is used to train the parameters for the energy consumption model. 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 on the dashboard; and  $I$  is the internal temperature as we measured with a manual thermometer located between the 2 front seats.

$\psi_c$  is used to train the parameters for the comfort model. 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  $\psi_e$ .

## 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 greatest fit to the data which we collected<sup>1</sup>. 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,  $\psi_e$ , and search for the parameters  $w_1, w_2, \dots, w_6$  which maximize the likelihood of the training data (maximum likelihood estimation).

## 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,  $\psi_c$ , and search for the parameters  $v_0, v_1, \dots, v_7$  which maximize the likelihood of the training data. 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 receive different advice.

<sup>1</sup>Some 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  $E$ .

## 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:

$$a(E, I, C) = \arg \min_{s \in S} e(s_T, s_F, s_D, M, E, I) \text{ s.t.} \\ s_M = 0; c(s_T, s_F, s_D, C, E, I) \geq 7 \quad (3)$$

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.

## Training Data Collection Methods

We used the following methods for gathering the necessary data in order to train CARE’s two models.

### Data Collection for Modeling Energy Consumption

We collected data on energy consumption directly from the car in order to train the energy consumption model ( $\psi_e$ ) while the climate control system was on. We conducted a total of 120 measurements. Each measurement was a 10-minute duration. We let the car warm up (and the compressor cool down) for 10 minutes between consecutive measurements. We conducted these measurements for various temperatures, starting at  $T = 16$  and up to  $T = 26$ , and sampled different values for all of the different variables. Many of our measurements (36) were focused on the range between  $T = 20$  and  $T = 25$  and when  $M = 0$  (“eco” mode), which is the natural range for candidates for the advice, and the function is required to be most accurate at those variables.

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, and in more reasonable settings the differences are much lower.

It turns out that the fan strength,  $F$ , had a greater impact on the energy consumption level than the temperature,  $T$ . That is, increasing the fan by one unit consumed more energy than reducing the temperature by one degree Celsius. The raw data we collected strengthened this result, as we

observed that when the fan was set to a higher level, not only did the blower consume more energy, but so did the compressor. Both the external and internal temperatures ( $E$  and  $I$ ) had a milder effect on the energy consumption level. Interestingly, the air delivery,  $D$ , had a negative impact on the energy consumption level, i.e. when the air delivery was set to face and feet, the climate control system consumed less energy than when set only to face. 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

Data collection for the human model ( $\psi_c$ ) requires human subjects, and thus is difficult to gather. We therefore had to assure that as many instances as possible are in the range that is most likely to be used by CARE. Merely randomly selecting different settings may not have yielded information good enough for training the human model.

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.

The subjects entered the car and sat in the driver's seat with their hands on the steering wheel and set the vents to point in their direction. While the climate control system was still off, the subjects were asked to rate their comfort level. The subjects were told how to operate the climate control and set it so that they would feel most comfortable. These settings were left on for 4 minutes. The subjects were asked for their comfort level and were required to explain why they had chosen that level. The subjects then exited the car and the car was left to warm up for 4 minutes. The subjects then returned to the car and the experiment operator set the next setting for them and waited 4 minutes. The subjects had to report and explain their comfort level and had to wait 4 minutes outside the car. These stages were repeated for a total of 8 different settings for every subject (resulting in 120 instances for all of the 15 subjects).

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 happened 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 model is:

$$c(T, F, D, C, E, I) = 16.6 - 1.3T + 0.98F - 0.064F^2 - 1.22D + 0.32C - 0.17E - 0.48I$$

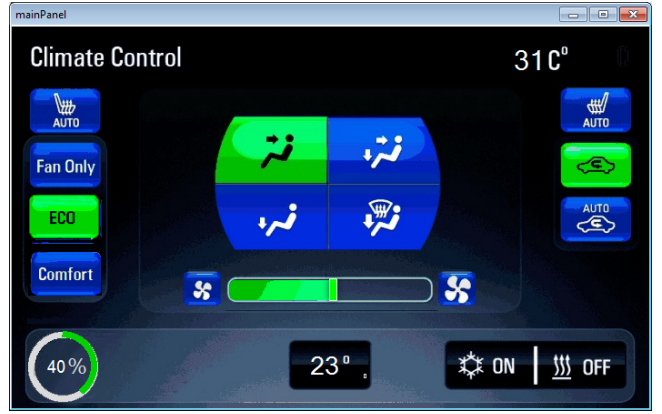


Figure 1: A screen-shot with additional energy consumption information provided by CAREless (the circle in the bottom left corner).

## Graphical User Interface

We implemented a panel based on the original climate control panel in the VOLT car, with additional add-ons. We have three different methods of advice, each with a different Graphical User Interface (GUI):

1. The first GUI is identical to the original climate control panel in the VOLT car. This option was used for the control group and is associated with a driver that does not receive any advice.
2. The second GUI has an additional information circle, which supplies the driver with an estimate 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). This GUI will be referred to as CAREless. Note that CAREless does not provide any active advice either. An example can be seen in Figure 1 (where the current consumption is 40% of the maximum.).
3. The third GUI is equipped with the full functionality of CARE. The driver is presented with both the advice provided by CARE and an estimate of the current energy consumption (similar to the information provided by CAREless). Figure 2 shows a screen-shot of a case in which the driver set the climate control differently from CARE's advice. Current user's selection is shown in green and the advice uses the purple color.

## Experimental Evaluation

In order to evaluate the performance of CARE and CAREless we recruited 49 subjects for the evaluation phase, of which 33 were males and 16 were females<sup>2</sup>. The subjects' ages ranged from 21 to 73, with a mean of 35 and median of 31. All subjects live in Israel. The subjects were paid 100

<sup>2</sup>All experiments with human subjects were approved by the corresponding IRB.

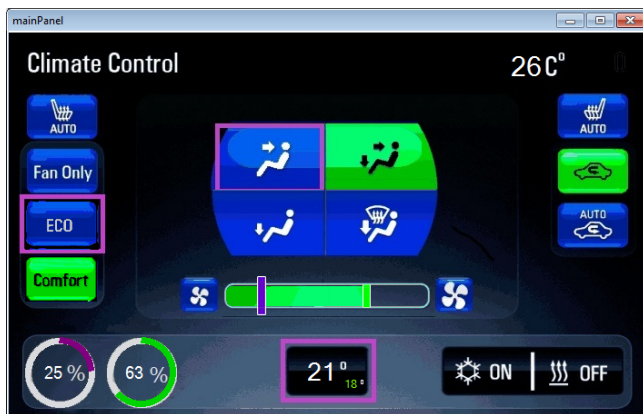


Figure 2: A screen-shot of the GUI accompanied with CARE’s advice. In this example, the driver set the temperature to  $18^{\circ}\text{C}$  (rather than  $21^{\circ}\text{C}$  as advised by CARE), the fan to 4 (rather than 1 - as shown by the purple line), the air delivery to both face and feet (rather than face-only) and the mode to “comfort” (rather than “eco”). This resulted in an energy consumption level of 63% of the maximal energy consumption level (right green circle), rather than only 25% if the driver would have followed CARE’s advice (left purple circle).

NIS each (27\$, price of fancy lunch in Israel). Each of the subjects was randomly assigned an advice provider, which was either CARE or CAREless. 24 subjects were assigned to receive advice from CARE, while 25 subjects were assigned to receive advice from CAREless. It is well known that people vary in their preferred temperature, fan, etc. settings and the outside temperature changes from time to time. Had we randomly assigned some subjects to a control group, those subjects may have had an average consumption that may have differed from the average consumption of the subjects receiving advice merely because of these differences. 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. We had each subject run the experiment twice, once with no advice (the control group) and once with advice given either from CARE or from CAREless. We counter-balanced the order among the type of experiments, i.e. approximately half of the subjects first ran the experiment with no advice, while the other half first ran the experiment with advice. Within each of these groups, approximately half of the subjects received CARE while half received CAREless.

Every subject adhered to the following procedure. First the subject was asked to fill out forms and demographic data, was then led to the car and was asked by the operator for his initial comfort level. Then, on a dedicated laptop (not on the car display) the subject was shown the GUI which, according to the experiment type, either presented CARE, CAREless or no advice. The subject then told the operator how to set the climate control system. The operator set the climate control system accordingly and updated the GUI accordingly and showed it to the subject. The subject could

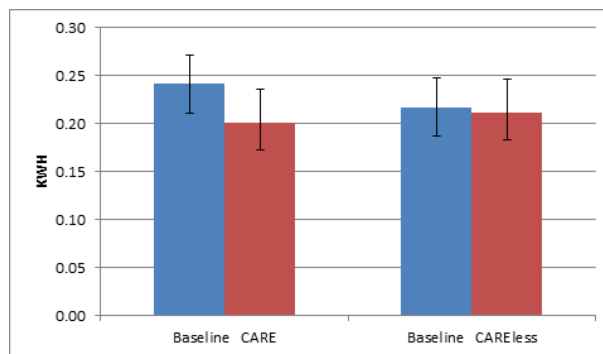


Figure 3: The mean energy consumption level of the subjects who received advice from CARE and CAREless, compared to the mean energy consumption levels of the subjects when they did not receive any advice.

then ask to modify the climate control system again. The subject remained in the car for a total of 10 minutes. The subject could ask the operator to modify the climate control system in these 10 minutes as many times as he wanted. The subjects then had to wait outside the car for 10 minutes between the experiments. The car doors and trunk were left open and the climate control system was turned off for those 10 minutes, in order to allow the temperature in the car to equalize to the outside temperature. After these 10 minutes the subject came back to the car and ran through the second experiment (which also lasted 10 minutes). The subject was then asked some final questions.

## Results

The results were analyzed using repeated measures of ANOVA with total energy consumption as a dependent variable, advice (yes/no) as a within-subject variable, type of advice (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 advice 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% ( $F(1, 21) = 7.6, p < 0.05$ )<sup>3</sup>. 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 3 presents the mean energy consumption level of the climate control system, which was obtained by the subjects when receiving advice from CARE vs. CAREless,

<sup>3</sup>We corrected for multiple comparisons, and after the Bonferroni correction, the type-I error remains  $< 0.05$ .



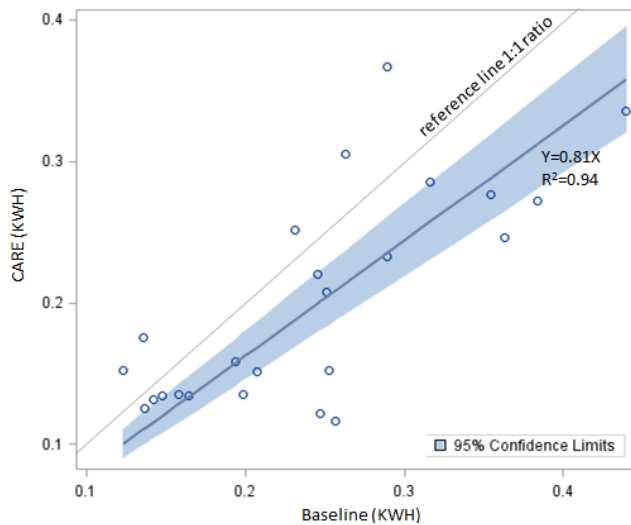


Figure 4: The energy consumption level of the climate control system of each subject when receiving advice from CARE compared to the baseline of that same subject when not receiving any advice.

compared to the mean energy consumption level of the same subjects when they did not receive any advice.

Figure 4 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).

### Discussion

CARE significantly outperformed the control group. It is possible that this occurred not only because some of the subjects actually accepted the advice and used it, but it seems that even those who did not accept the advice were influenced by it. Some subjects also used the advice as a baseline and edited it. 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 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. Clearly, CARE has reduced her energy consumption. In total, out of the 24 subjects that received CARE’s advice, only 4 followed the exact advice, however, CARE also decreased the energy consumption of the subject who did not follow its exact advice.

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. Another interesting observation is that females tend to consume less energy than males, 0.201 vs 0.242 (when looking only at the no-advice condition). This raises the idea that demographic data may be used instead of the information provided explicitly by the driver when entering the car.

### Conclusions

In this paper, we presented a method to persuade a driver to reduce the energy consumption of the climate control system of his electrical car. We showed via experiments 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 a car.

### Acknowledgments

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

### References

- 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.
- Das, A.; Mathieu, C.; and Ricketts, D. 2009. Maximizing profit using recommender systems. *ArXiv e-prints*.
- Fogg, B. J. 2002. Persuasive technology: using computers to change what we think and do. *Ubiquity* 2002(December):5.
- Froehlich, J.; Findlater, L.; and Landay, J. 2010. The design of eco-feedback technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1999–2008*. ACM.
- Nguyen, T.; Yang, R.; Azaria, A.; Kraus, S.; and Tambe, M. 2013. Analyzing the effectiveness of adversary modeling in security games. In *AAAI*.
- Ricci, F.; Rokach, L.; Shapira, B.; and Kantor, P., eds. 2011. *Recommender Systems Handbook*. Springer.