Toward Building an Individual Preference Model for Personalizing Settings in the Vehicle

Olav Laudy
Causality Link
Sandy, UT. USA
olav@laudy.net

Johann P. Prenninger
BMW Group
Munich, Germany
Johann.Prenninger@bmw.de

Alvin Chin
BMW Technology Corporation
Chicago, IL, USA
alvin.chin@bmwna.com

Jilei Tian
BMW Technology Corporation
Chicago, IL, USA
jilei.tian@bmwna.com

***Abstract*—It is highly expected that your car learns your individual preferences and context so that when it is cold, it automatically switches on the seat heating and when it is warm it switches on the air conditioning. This paper describes a simple computational method for learning individual preferences that has the advantage that it does not require to be trained in the vehicle on-board for every customer separately, while maintaining its ability to tune to every customer individually rather than defaulting to the average customer's preference. Our results show that our method outperforms the non-individual and clustering models, and demonstrates the feasibility of integration into the vehicle.**

Keywords—seat heating, automatic settings, individual preferences, personalization

I. Introduction

Would it not be fantastic if your vehicle knows your individual preferences and context so that it can recommend or automatically switch on the seat heating or air conditioning? Many vehicles have a feature in the head unit where the seat heating can be turned on when the temperature reaches a particular temperature. However, very few customers may know this and in addition, it may be cumbersome to set up these personalized preferences among several other settings. We could use an average model where everyone gets the same preference at the beginning, then the users adjust accordingly. Nonetheless, this is not desirable as users may feel uncomfortable with those beginning settings already.

Thus, an average model is not sufficient and a predictive model needs to be created that learns from your observed behavior and starts applying that in an automated fashion to eliminate having to set up manual rules. Predictive modeling is usually done offline or online after the data has been collected. In the case for the vehicle, the predictive modeling has to be performed in real time using real-time streaming vehicle data with personal and vehicle context. Due to performance constraints and providing an optimum user experience, the naïve implementation of such would require the predictive model to be trained on-board in the vehicle head unit for each individual customer's vehicle separately, which has consequences for the computational power and capacity of the onboard CPUs or for frequent (and potentially undesirable) data transfers [1, 2]. Therefore, the problem is how to build an individual customer preference model that can be learned quickly and can run in the vehicle on-board with limited resources.

In this paper, we demonstrate how to build an individual customer preference model using basic techniques such as cluster analysis, PCA and artificial neural networks, and apply them to the example of seat heating. We use data that is based on a real example, but is simulated in order to not disclose certain details. We evaluate the individual preference model with the non-individual model and the cluster-based approach using AUC (Area Under the Curve), and discover that the individual preference model outperforms the other models ($AUC_{individual} = 0.98$, $AUC_{non\text{-}individual} = 0.92$, $AUC_{cluster} = 0.95$ to 0.98).

Our contributions are the following. First, we create a model for learning a customer's individual preferences automatically and show that it outperforms other state-of-the-art methods. Second, we apply this model to a case study of seat heating in the vehicle. Third, we show the viability of our model for direct implementation in the vehicle.

The paper is organized as follows. Section 2 describes about related work in learning user preferences and individual behavior. Section 3 introduces the case study, setup and data collection for learning seat heating preferences. Section 4 presents previous models that can be used for learning individual preference. Section 5 describes our individual model and the results from using the simulated data. Finally, Section 6 concludes the paper and references are outlined at the end.

II. Related Work

Learning user preferences from existing data is not new. Many attempts have been reported where personal data is used as input for predicting and modeling human behavior [3, 4, 5]. For example, in [3], the authors use reinforcement learning to learn human behavior from playing a virtual game. In [4], adaptive system behavior (which is human behavior) is formally defined. In [5], the authors discuss the importance of considering the concurrency of automated functions, as these

978-1-5386-6358-5/18/$31.00 ©2018 IEEE

alter the user behavior in a way which cause less regular and predictable interactions.

In addition, there exists work that takes a user's comfort and ergonomic considerations into account and extends the classical behavior-based approach [7]. To do so, specific Key Performance Indicators are studied to assess and optimize individual performance leading to satisfy ergonomic goals of ISO-7730 for thermal comfort [8].

In particular, seat heating seems to be a typical case to learn people's habits to control the temperature. The easiest solution today is to allow users to set a certain but static trigger temperature to automatically initiate seat heating [9]. These user settings obviously depend on how the system reaction is understood and perceived. The user acceptance can later be used as a starting point for learning user's behavior.

Our work differs from prior art in that we do not use the classical behavior-based approach or use the average behavior model for all users. We create a fully individual behavior model that learns the appropriate seat heating switches (on or off) automatically without the user having to apply any pre-settings in advance.

III. Case Study, Setup and Data Collection

Before explaining the models that we use for learning individual preferences, we first describe the setup that we use for offline learning and prediction, the implementation considerations, and how we collect the data. We use seat heating as our case study throughout the paper, but this approach can apparently be extended to other scenarios.

A. *The simplest solution is not sufficient*

The simplest method for predicting the seat heating could be to derive the optimal switch point as the midpoint between the maximum observed temperature when the seat heating was switched on and the minimum temperature when the seat heating was not switched on. However, it may also depend on the inside and outside temperature of the vehicle. Understanding how to combine the inside and outside temperature without a machine learning model is to make a guess about the decision boundary for an individual user rather than learning from the data [3, 4]. In essence, this is a simple matched rule using thresholds as the matching conditions that basically follows the If-This-Then-That (IFTTT) model or trigger-action rule-based model [10].

B. *Implementation considerations*

To start with, for simplicity, assume there is a single driver per vehicle. Since one tries to learn individual customer behavior, it seems logical to build one model per customer. Since one observes a customer using seat heating only after a vehicle is in use, it seems necessary to put the learning algorithm itself in the vehicle which is complex because it requires iterative optimization, rather than only the scoring code which in contrast is easy because it has a fixed ruleset or formula.

So, one could go ahead and build a chip that contains some machine learning algorithm and apply some online learning algorithm. However, what if the model itself contains a collection of decision boundaries and the scoring consists of classifying a customer's past observations to its most optimal decision boundary? Then, no online learning is needed, and one could build a model across customers to learn the collection of decision boundaries, rather than building one model per customer [5, 6]. This is the approach that we take in our work.

C. *The data*

We obtained a training dataset from simulating a sample of 1,000 vehicles. For each vehicle, two true switching temperatures for seat heating are sampled, *temp_in_True* (inside temperature) and *temp_out_True* (outside temperature). We sample using a normal distribution of mean=X, std.dev=Y which we denote as Normal(X,Y). The outside temperature (*temp_out*) is sampled from a Normal(5,5) and the inside temperature (*temp_in*) is sampled from a Normal(10,5). Then, for each vehicle, 100 days of observations are sampled from a Normal(5,10) and a Normal(10,10) for the inside and outside temperature, respectively. Finally, a small noise term is sampled from a Normal(0,1). To determine whether or not the seat heating is switched on (Target = 1 for seat heating on), the following logic in Eq. (1) is applied:

$$\begin{gathered}\textit{If (temp\_in+noise)} < \textit{temp\_in\_True, then Target} = 1 \\ \textit{Else Target} = 0 \\ \textit{If (temp\_out+noise)} < \textit{temp\_out\_True, then Target} = 1 \\ \textit{Else Target} = 0\end{gathered} \tag{1}$$

Table 1 shows an example of data simulated for a vehicle, using the above generation.

CAR ID	DAY	TEMP_IN	TEMP_OUT	TARGET
1	1	6.7	-5.9	1
1	2	-1.0	-7.3	1
1	3	24.3	-8.9	1
1	4	15.0	-0.1	1
1	5	0.5	9.0	1
1	6	-1.9	13.9	0
1	7	17.9	7.8	1

Table 1. Example of data simulated for a vehicle.

From the data, it appears that individual users are consistent when it comes to switching on the seat heating, while between users there seem to be large differences (i.e. at certain temperatures half the users do switch on seat heating, while the other half does not). Hence, this means that we need to have an individual preference model.

IV. STATE-OF-THE-ART MODELS

We present 2 models for modeling and predicting the personal seat heating preference from state-of-the-art: the non-individual preference model and the cluster-based preference model.

A. The non-individual preference model

This model deals with the general case and is not tailored personally to the driver. It is the simplest model and takes the inside and outside temperature as predictors and predicts whether or not the seat heating will be switched on, based on all the drivers' data. We use a simple neural network with a multilayer perceptron with a tanh activation function for the hidden layer and a softmax for the last layer and a cross-entropy loss function [13].

A classification model is typically evaluated by the AUC (Area Under the Curve), which was 0.92 for the training data as well as for the hold-out validation data. The model has an excellent fit and is capable of generalizing to new data. Most of the cases of percentage of correct 'guesses' of the model per vehicle owner, fall between 80% and 100%, which indicates that for most vehicle owners, if they get into their vehicle, at least 8 out of 10 times, the model guesses correctly whether or not the seat heating should be switched on. However, for some vehicle owners, the model only guesses right in 50% of the cases, which is not acceptable for a production grade model.

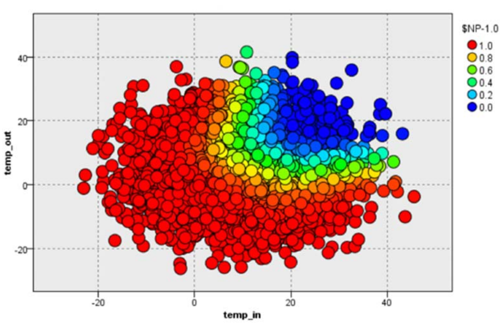

Fig. 1. Cluster-based decision boundaries for non-individual preference model. Red indicates actual seat heating = off and predicted seat heating = on, light blue indicates actual seat heating = off and predicted seat heating = off, dark blue indicates actual seat heating = on and predicted seat heating = off, green indicates actual seat heating = on and predicted seat heating = on, and yellow indicates the decision boundary.

From Figure 1, summarizing the model in a generic decision boundary suggests the following rule: Switch on the seat heating when either the inside or the outside temperature is below 10 degrees Celsius.

B. Cluster-based preference model

The previous model made no differences between the different owners of the vehicle. To learn the preference of a vehicle owner, we cluster vehicle owners together with similar behavior and build a model per cluster or take the cluster membership in as an additional predictor. Then, we create deciles from the temperatures across subjects, and per vehicle owner, compute the percentage of times that per decile seat heating was switched on (with the base of the percentage, the total number of times a temperature was observed in the decile at hand for that vehicle owner).

Inspecting the data shows that deciles is probably a fair number. With less categories, there would be less resolution. With more categories, many of them would have insufficient observations. Note that since the transformation is taken per variable, we implicitly assume that the resulting decision boundaries are parallel to the temperature. Alternatively, one could create categories of the inside and outside temperatures together and bring those to the columns. That decision becomes more relevant when there are more observations per vehicle owner available.

The resulting 20 columns are input to a TwoStep agglomerative hierarchical clustering method [12] which resulted in 8 clusters. The data for the seat heating model now has one additional variable: per vehicle id, a cluster membership is available. Note that when using this approach in score mode, a few observations are needed as input to the cluster model to determine the cluster the vehicle owner belongs to. There are two options for modeling: adding the resulting cluster variable in the model as predictor (Model A) or building a model per cluster (Model B). Ideally, adding the cluster variable as predictor is the best thing to do: if the modeling technique is flexible enough that it is capable of finding and fitting variable interactions, then the data will indicate to what extent the cluster variable will need to be incorporated. If the cluster variable is very important, the model will include the cluster variable as a full interaction with the other predictors, which mathematically is equivalent to building a separate model per cluster.

The AUCs of the models are 0.95 for Model A and 0.98 for Model B for both training and validation. Compared to the non-individual model (AUC = 0.92), the AUC has increased significantly. Also, the model per cluster (Model B) is significantly stronger than the cluster as predictor model (Model A). The one hidden layer neural network, although capable of automatically fitting interaction effects, does not have enough flexibility to incorporate full interaction between the clusters and the temperature variables, which is available in data, and hence the slight worse performance. Note that rather than one 8 cluster solution, a sequence of cluster models was run, ranging from 1 to 15 clusters and each cluster solution in turn is added to the seat heating model. A cluster model with 8 clusters leads to the best results on the hold-out data. The percentage correct per vehicle owner has increased, as well as shown by the AUC.

Figure 2 shows observed temperatures for two individual vehicles along with the true state and their prediction, as well as the decision boundary. The two vehicles have different decision boundaries, as the two vehicles are chosen from different clusters. There are only 8 different decision boundaries, as there are only 8 clusters. The figure only shows 2 of the 8 to meet space constraints.

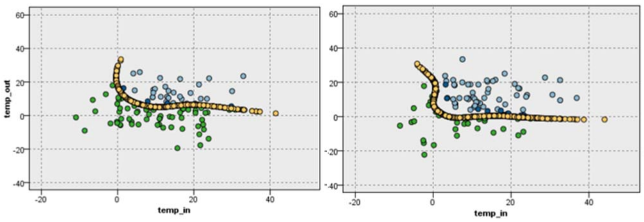

(a) (b)

Fig. 2. Cluster-based decision boundaries for 2 vehicles: (a) vehicle 1 and (b) vehicle 2. The colors are per the description in Figure 1.

V. A Fully Individual Model

The previous model used clusters to create decision boundaries for users with similar behavior. A cluster model creates latent or unobserved groups in the data. Would it be possible to project the histogram predictors into a continuous space rather than in a categorical space? There are many techniques for this, but one obvious one is principal component analysis (PCA). PCA and cluster analysis are mathematically closely related to each other (they can both be described as low rank matrix factorization methods).

For the fully individual model, we use PCA. The AUC of the fully individual model (0.98) shows the best performance compared to the previous models so far. Since the cluster analysis and the PCA contain so much of the same information, one wonders why the neural network would not be learned enough from the PCA factors alone. We attribute this to the fact that, similar to fitting a model per cluster rather than adding the cluster membership as predictor, the neural network is not able to fully incorporate the information. Adding the PCA factors per cluster, gives the neural network the ability to, within a cluster, apply a finer grained control to determine the exact shape of the individual decision boundary.

The histogram of percentage correct of switched seat heating decisions per vehicle owner, now shows that no vehicle owner has scores below 80% correct, with the majority scoring between 90% and 100%. In practice, the incorrect guesses will lie close to the individual decision boundaries, i.e. the vehicle owner was very close to switching on the seat heating but did not or vice versa.

To demonstrate the true individual decision boundaries, a small experiment is conducted. Figure 2 shows various decision boundaries while carefully manipulating the true switching variables. Figure 3 (a) shows the decision boundaries for cluster 7 where the true inside temperature is kept between 11 and 12 degrees Celsius, and Figure 3 (b) is for cluster 7 where the outside temperature is between 11 and 12 degrees Celsius but varies the inside temperature. The graph shows the decision boundaries for various outside temperatures indicated as colors and per legend. As can be seen, truly each combination of inside and outside temperature corresponds to a slightly different decision boundary and hence, this is a fully individual model.

The PCA factors cannot be directly interpreted as the true switching temperatures, as the scale is different, and the factors are linear transformations of the true switching temperatures. Yet, it can be assumed that the full information of the true switching temperatures is in there, as can be seen from the high performance of the model.

One may wonder: why not build a model with all the original histogram predictors in there? Although the full information is available there, the neural network has only one hidden layer and cannot capture the complexity of the required transformation. Basically, this is asking the neural network to perform a PCA or similar analysis. Potentially a neural network with more layers is capable of doing this, yet, helping the model to find the right relation by preparing the data with a cluster analysis and a PCA, certainly works as well. The AUC of this model is around 0.83.

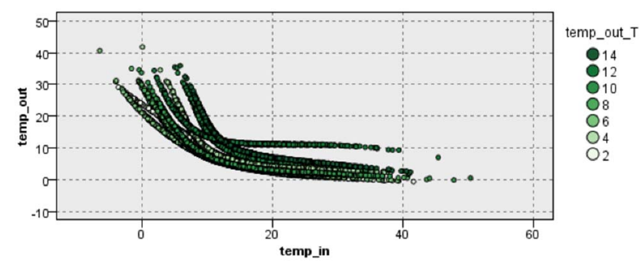

(a)

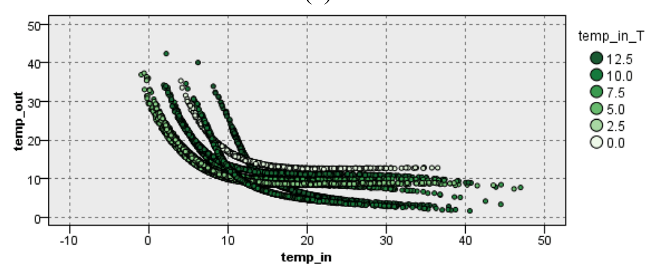

(b)

Fig. 3. Individual-based decision boundaries for (a) cluster 7 where the true inside temperature is kept between 11 and 12 degrees Celsius and (b) cluster 7 where the true outside temperature is kept between 11 and 12 degrees Celsius.

The scoring procedure comes down to a set of simple formulas: first, every new observation is classified in the corresponding decile and updates the percentage in there. Next, the factor analysis takes this as input and produces two factor scores and similarly, the cluster model produces a cluster score. Finally, the current inside and outside temperatures, together with the cluster and factor scores feed into the neural network to produce a probability, which, with a chosen cutoff, can be used to decide whether or not to turn the seat heating. Thus, there is no online learning, and there is no use for one model per vehicle. In fact, a small demonstration can show that one model per vehicle is likely not a good idea.

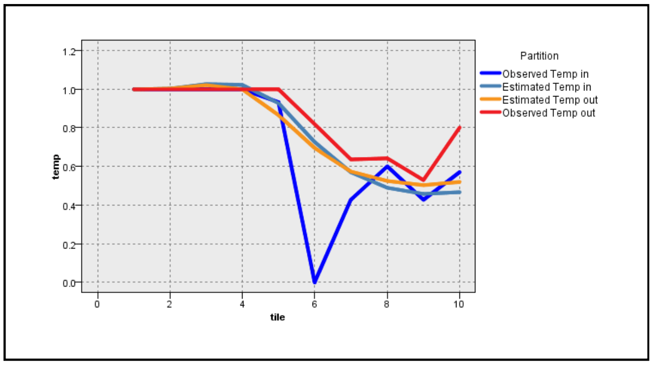

Fig. 3. Smoothing of individual vehicle owners.

The non-smoothed curves in Figure 3 show the observed percentages per decile for one vehicle. The dark blue line (observed inside temperature) shows a dip in decile 6, as there were no observations in that interval. The dark red line (observed outside temperature) goes up towards decile 10, which may be due to cold inside temperatures that were observed while the outside temperature was in decile 10. If one would create a model per car owner, this non-smooth behavior will likely be picked up on, with the result that the subsequent scoring of the model (i.e. the automatic switching of the seat heating) shows undesirable behavior. Instead, the PCA looks at the correlations between the deciles and transforms the data to a lower dimensional space. This can be seen as a smoothing operation: the random variations will be captured in the later factors rather than in the first two. Since those factors are discarded, the result is a smoothed curve. Indeed, this is the case: for the given vehicle, the two-dimensional point can be transformed back to the original decile data, and the lighter blue (estimated inside temperature) and orange curve (estimated outside temperature) show the smoothed curve. Also, the higher the decile, the lower the probability that the seat heating is switched on. We conclude that placing the data in a lower dimensional space, smoothens the volatility out of the data, and the availability of many vehicles together in the analysis helps this fact, as the PCA learns from all the vehicles together.

VI. Conclusions And Future Work

In this paper, we have demonstrated how to derive individual preferences or traits from data in order to better serve the customer. This was illustrated using the example of seat heating in a vehicle to make the life of a customer easier by learning from her behavior and applying the learning automatically. The initial idea might be to build a model per customer and use an online learning method to update the model when data becomes available, however an approach using the fully individual preference model was presented where 1) modeling was done on the whole set of customers so that individual predictions improve by learning from neighboring data and 2) new data arriving changes the predictions as part of the scoring procedure due to the fact that the model contains all the decision boundaries after learning is done. Our results showed that the fully individual preference model outperformed the non-individual and cluster-based preference models based on the AUC.

For future work, we plan to focus on the user's experience when such a predictive model is actually applied onboard. First focus will be how to signal when an automation is done to achieve a maximum of user acceptance in practice. A second focus will be to understand what adaption speed is expected by customers and realize that the model is actually doing helpful and appreciated support. The third focus will be around the change of the driver. Finally, the architecture and computational requirements to deploy such a model in the vehicle and connected vehicle backend need to be evaluated.

References

[1] S. Russell and P. Norvig, *Artificial Intelligence A Modern Approach*, 3rd ed. New Jersey, USA: Prentice Hall, 2009.

[2] B. Leite Dantas Bezerra and F. D. A. Tenorio de Carvalho, "Symbolic data analysis tools for recommendation systems," in *Knowledge and Information Systems*, 2010, vol. 26, no. 3, pp. 385-418.

[3] C. Bauckhage, C. Gorman, B Thurau, and M. Humphrys, "Learning Human Behavior from Analyzing Activities in Virtual Environments," in *MMI-Interaktiv*, 2007, No. 12, pp. 3-17.

[4] M. Broy, C. Leuxner, W. Sitou, B. Spanfelner, and C. Winter, "Formalizing the notion of adaptive system behavior," in *Proc. SAC '09*, Honolulu, Hawaii, 2009, pp. 1029-1033.

[5] H. Wolf, R. Zöllner, and H. Bubb, "Ergonomic solutions for the concurrent feedback of several driver assistance systems to the driver," in *Proc. Active Safety Conference*, Munich, Germany, 2006.

[6] X. Xu and C. Science, "Bachelor's thesis: 'An Innovative SRL Algorithm: Joint-Bayesian-Network'," Simon Fraser University and Zhejiang University, 2010.

[7] A. Schumann, N. Wilson, and M. Burillo, "Learning User Preferences to Maximise Occupant Comfort in Office Buildings," in *Trends in Applied Intelligent Systems*, 2010, pp. 681-690.

[8] ISO Standard, "EN ISO 7730 - Ergonomics of the thermal environment - Analytical determination and interpretation of thermal comfort using calculation of the PMV and PPD indices and local thermal comfort criteria," 2005.

[9] BMW of Lincoln. (2017) How to Program your BMW's Automatic Seat Heating and Ventilation. [Online]. Available: https://www.bmwoflincoln.com/blog/2017/march/1/how-to-program-your-bmw-s-automatic-seat-heating-and-ventilation.htm.

[10] B. Ur, M. P. Y. Ho, S. Brawner, J. Lee, S. Mennicken, N. Picard, D. Schulze, and M. L. Littman. "Trigger-Action Programming in the Wild: An Analysis of 200,000 IFTTT Recipes". in *Proc. 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. San Jose, CA, May 7-12, 2016, pp. 3227-3231.

[11] I.T. Jolliffe. "Principal Component Analysis and Factor Analysis". in *Principal Component Analysis. Springer Series in Statistics*. New York, NY: Springer, 1986.

[12] SPSS. The SPSS TwoStep Cluster Component. [Online]. Available: https://www.spss.ch/upload/1122644952_The%20SPSS%20TwoStep%20Cluster%20Component.pdf.

[13] IBM Knowledge Center. Activation Function (multilayer perceptron algorithms). [Online]. Available: https://www.ibm.com/support/knowledgecenter/en/SSLVMB_22.0.0/com.ibm.spss.statistics.algorithms/alg_mlp_architecture_activation-functions.htm.