Recommending Eco-Friendly Route Plans

Efthimios Bothos
National Technical University of Athens
Athens, Greece
mpthim@mail.ntua.gr

Dimitris Apostolou
University of Piraeus
Piraeus, Greece
dapost@unipi.gr

Gregoris Mentzas
National Technical University of Athens
Athens, Greece
gmentzas@mail.ntua.gr

ABSTRACT

As personal transportation is one of the greatest contributors of CO2 emissions, means able to assist travelers in reducing their ecological impact are urgently needed. In this work we focus on travel recommenders that encourage green transportation habits among travelers who have a pre-existing interest in taking action to lessen their impact on the environment. We aim to provide urban travelers with a personalized travel recommender that will nudge them to plan routes while considering the environmentally friendliest travel modes. We present a novel, ecologically-aware approach for travel recommender systems and propose a system architecture that incorporates dimensions of recommendation information elements and profile matching methods.

Categories and Subject Descriptors

H.3.3 [**Information Search and Retrieval**]: Information filtering, Selection process; H.3.4 [**Systems and Software**]: User profiles and alert services

General Terms

Design, Human Factors, Algorithms

Keywords

Travel Recommenders, Choice Architecture, Nudging, Persuasive Technologies, Lifestyle Change

1. INTRODUCTION

Environmental issues are becoming increasingly pressing in our times and means to reduce the ecological impact of citizens' activities are needed urgently. A major source of environmental pollution from citizens' activities is carbon emissions due to traffic and mobility. It is estimated that urban transport in the European Union accounts for 15% of all greenhouse gas emissions [12]. As work and leisure life become progressively geographically distributed, a research issue of high importance pertains the development of methods and tools able to support and guide citizens towards pro-environmental behaviors with respect to their traveling habits and decisions.

Previous research has demonstrated that information regarding transport-related attributes such as travel time, travel costs and carbon emissions can lead to changes in citizens' travel behavior [3]. Nevertheless, although individuals base their choices on the attributes of the choice set (content), the presentation of information (context) has also a strong effect on travelers' behavior [4]. The presentation of choices, also known as "choice architecture" [17], refers to the design and incorporation of small features or nudges in the choice making process, which can assist individuals to overcome cognitive biases by highlighting the better choices for them, without restricting their freedom of choice. Tools available to choice architects can be divided into two categories: those used in structuring the choice task and those used in describing the choice options [9]. Recommender systems can act as tools for structuring the choice task and address the problem of what to present to travelers. Furthermore the use of information technologies incorporating feedback and personalization can be central to make lifestyle or behavioral changes [5] and, in our case, can nudge environmentally-responsible behavior.

In this work in progress we focus on recommender systems that encourage lifestyle changes towards green transportation habits among travelers who have a pre-existing interest in taking action to lessen their impact on the environment. We aim to provide urban travelers with a personalized travel recommender that will nudge them to plan multi-modal routes while considering the environmentally friendliest travel modes. We present a novel, ecologically-aware approach for travel recommender systems and propose a system architecture that incorporates dimensions of recommendation information elements and profile matching methods.

Our approach is detailed in Section 2. We synthesize concepts from multi-criteria decision making (MCDM) recommender systems and recommendations diversification to infuse the ecological dimension on travel recommenders. Namely, we focus on MCDM to infer user preferences and we balance the utility of routes with their carbon footprint in order to generate travel recommendations with ecological characteristics. In Section 3 we analyze the conceptual architecture of a system that implements the proposed approach. An illustrative scenario depicts the various user interactions with the proposed system in Section 4. We conclude with related

Lifestyle @RecSys'12, September 13, 2012, Dublin, Ireland

work and future directions.

2. APPROACH

Contrary to the vast majority of previous research on recommender systems that has focused on improving the accuracy of recommendations, i.e. better modeling user preferences to present individually preferred items, we focus on recommender systems as a tool for nudging users towards eco-friendly traveling decisions. Specifically, the recommender generates a list of suggested routes which reside within the limits of users' preferences and presents choices with low carbon emissions. With our approach we address the problem of a "filter bubble" [15] in its ecological dimension: users of existing navigation services may be trapped in a self-reinforcing cycle of emission-intensive travel modes while never being pushed to discover alternatives.

The problem an ecologically aware travel recommender system is asked to solve can be formulated as follows: Given a user u, find a subset $S \subseteq AvailableRoutes(u)$ such that $|S| = PresentedRoutes$ and the choice of S provides a good balance between the user perceived route utility and CO2 emissions. The research agenda of the above problem includes two main issues: First what is meant by user perceived route utility and how this is calculated and second what is the meaning of the term 'balance'. Both issues can be answered in a number of ways. Our approach is based on utility-based recommenders and involves a three-step process: users provide their preferences which are then transformed to a user perceived route utility value. In the final step, the utility and the CO2 emissions of a route are provided as input to a recommendation algorithm that selects $|S|$ results to be presented to the user.

2.1 User Preferences

Following [18] we adopt a utility based approach to elicit user preferences. In more details users provide their preferences over a set of criteria when planing a route. The revealed preferences are used to infer a user perceived utility per route.

First users are asked to assign themselves in one of six groups of drivers as identified by [2] - Hard driver, Complacent car addict, Malcontented motorist, Aspiring environmentalist, Car-less crusader, Reluctant rider (for a thorough description of these categories please see [2]). This information is asked only once and affects the level of nudging the user may be inclined to accept (i.e. an Aspiring environmentalist will be presented with more routes that involve public transportation and walking than a Hard driver).

Although most navigation applications provide the quickest routes as suggestions, in real life situations users are concerned with other aspects when deciding on a specific trip in a city. For example, the price of the ticket or the fare (e.g. for a taxi) of the transport mean might influence the user's decisions [18]. Moreover travelers interested in reducing their carbon footprint may be willing to walk a bit more or accept a longer trip. Based on the above arguments, in a second step users are asked to provide their preferences on a set of criteria which are then used to calculate a per route utility value. Indicative criteria are: preferred delay for arrival, preferred walking or bicycling time and preferred travel cost.

2.2 Routes and Utility Calculation

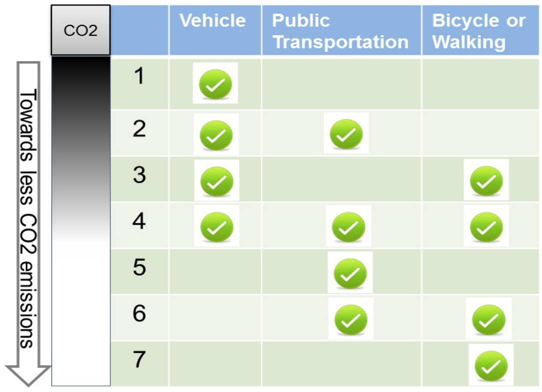

Figure 1: Travel profiles as combinations of alternative travel modes and corresponding qualitative CO2 emissions.

The alternative routes emerge from 'travel profiles' [18] which in our case are defined as the combination of one or more of the major transportation modes (personal vehicle, public transportation, walking or bicycle). In total there are $\sum_{k=1}^{3} \binom{3}{k} = 7$ travel profiles to choose from. Based on the travel mode characteristics and associated emission models of each travel profile we can infer that the use of more walking or bicycle leads to less CO2 emissions (see Figure 1), thus our aim is to nudge users into using travel profiles that include walking or bicycle.

The alternative routes are annotated with a utility value based on the submitted user preferences. To this direction Multi-Criteria Decision Making (MCDM), a set of widely studied methods in the Operations Research domain for decision making, can be employed. With MCDM a decision problem can be seen as the selection of the best alternative from a decision matrix $M \times N$ with N alternatives and M criteria. More specifically we select Multi-Attribute Utility Theory (MAUT) methods [7] which determine the utility of alternatives from user preferences on selected criteria. These methods are based on the concept that bad performing alternatives on one criterion can be compensated by good performing criteria. In our case an alternative is a route with criteria C_j. Each criterion has a weight W_j and the elements $a_{i,j}$ in the decision matrix denote the utility $U(c_{i,j})$ of criterion $c_{i,j}$. Indicative MCDM models that can be used include Weighted Sum and Weighted Product models.

In Weighted Sum Models a weighted mean over all criteria dimension for all alternatives is calculated. The result is a utility score per alternative: $U_i = \sum_{j=1}^{n} a_{ij} w_j$. Weighted Product Models multiply instead of summing up the criteria, and power instead of multiplying the weights in order to calculate the utility scores: $U_i = \prod_{j=1}^{n} a_{ij}^{w_j}$.

2.3 Recommendation Strategies

Given a set of candidate routes $AvailableRoutes(u)$ and a given threshold K of final desired number of recommendations, the optimal scenario of recommendation is finding a set of routes, that has the highest perceived utility and the

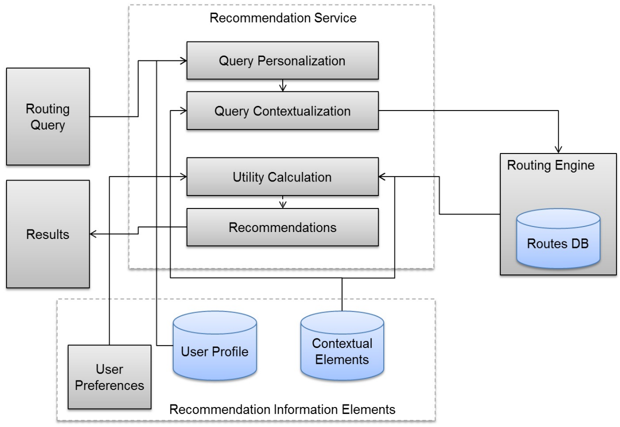

Figure 2: Proposed Architecture.

lowest CO2 emissions. However such an optimal $top-K$ answer set in general does not exist: lowering CO2 emissions typically does not correlate with the highest utility routes being selected. As a result, we have to achieve a balance between CO2 emissions and route utility. In order to generate lists of suggested eco-friendly routes, recommendation diversification algorithms can be employed following [23]. The two problems share similarities: diversification solutions attempt to identify relevant yet diversified items whereas we want to suggest relevant yet eco-friendly routes.

Two optimal algorithms are the MaxUtil which maximizes the utility of the K routes presented and the MinCO2 that minimizes the CO2 emissions of the K routes. Additional heuristic algorithms are the Swap and Greedy similarly to [20] and [22]. With algorithm Swap we begin with the K highest utility routes, and swap the route with the highest emissions with the next highest utility route among the remaining routes. A route is swapped only if the overall CO2 emissions of the displayed set is decreased. To prevent a sudden drop of the overall utility of the resulting set, a pre-defined upper-bound UB denoting how much drop in utility is tolerated has to be used. With the use of UB, swapping stops when the utility of the resulting routes becomes lower than UB. Furthermore the value of UB depends on the drivers group the user has assigned herself (see Section 2.1). With algorithm Greedy recommendation lists are formed by combining routes from different travel profiles. The list with the lowest emissions and acceptable utility is selected. Lists with acceptable utility are those whose difference with the highest utility list resides within certain limits: $HU - U_i \leq AD$ where HU is the Highest Utility, U_i is the utility of list i and AD is the Acceptable Difference which depends on the drivers group the user has assigned herself.

3. ARCHITECTURE

In this section we describe a system architecture that shows how our approach can be instantiated and extended to incorporate personal and contextual information. The proposed architecture comprises of the following components: Recommendation information elements, Recommendation service and Routing engine (see Figure 2).

3.1 Recommendation information elements

These elements incorporate the individual user profile and preferences as well as information related to the current context. In more details we identify the following information elements:

- User preferences provided by the user through a multi-criteria input interface together with the routing query before the trip planning.
- User profile configured by the user through an input interface on the first use of the system.
- Current context of the user, e.g. trip purpose (business, leisure, tourism), weather and traffic information.

3.2 Routing engine

The routing engine takes as input a set of routing options and generates a set of itineraries. It is controlled by the Recommendation service that manages the options on behalf of the user and adjusts the values based on the user's profile. Routing options to be supported include route characteristics such as travel modes. The results should include information regarding emission levels, calculated with emission models and the estimated arrival time at the destination.

3.3 Recommendation Service

This component comprises of four distinct functions responsible for personalizing and contextualizing the alternative routes to be presented to the user. The first two, query personalization and contextualization, transform the user routing query and context signals into the appropriate routing engine API parameters. Query personalization is dependent on the available transportation means the user has at her disposal i.e. car/motorcycle and bicycle and considers any disabilities the user may have. Two rules are defined for these cases:

- If the user owns a vehicle then routing results involving car/ motorcycle should be considered, similarly if the user owns a bicycle, routing results involving a bicycle should be considered.
- If the user has disabilities then bicycle and public means of transportation that do not provide amenities for persons with disabilities should be avoided.

Query contextualization considers a number of static rules to further filter the initial set of results:

- Weather data: if the day is rainy, then bike and walking time should be kept to a minimum.
- Traffic data: if there is indication of high traffic density, car time should be kept to a minimum.
- Trip purpose affects the possible delays with respect to the time of arrival. Expected delays should be minimized for business trips, can be moderately tolerable for leisure trips, and tolerable for tourism trips.

Based on the aforementioned rules, the user query is augmented and a request is sent to the routing engine for alternative itineraries.

Following query personalization and contextualization, the routing engine is triggered to generate a set of n results per travel profile given the set of personalization and contextualization parameters. Once the results are available, two

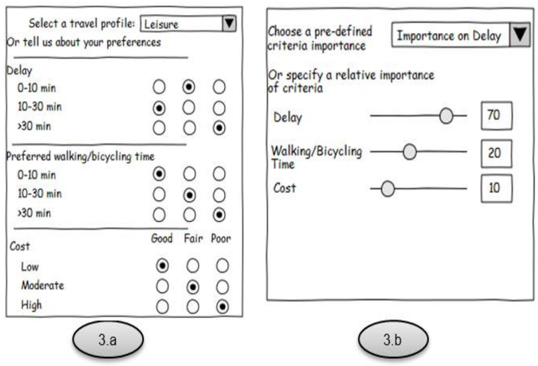

Figure 3: User input: Preferences on the criteria, and relative importance of criteria.

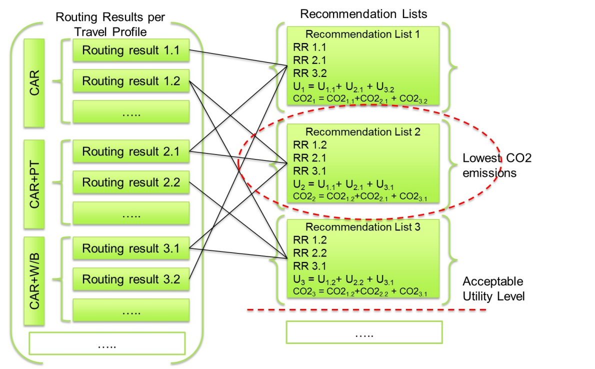

Figure 4: Recommendation lists as combinations of travel profiles. Each travel profile is a combination of one or more travel modes (CAR, PT - Public Transport, W/B - Walking or Bicycle).

more functions are triggered. The utility calculation function maps the recommendation information elements and the characteristics of the route to a perceived utility value per user and route following MCDM methods as described in Section 2.2. This step allows the projection of the user's decision strategy on the results. The final step refers to the generation of recommendations following Section 2.3.

4. ILLUSTRATIVE SCENARIO

In the following we describe an illustrative use case scenario of our approach. John is about to go out and meet his friends at a movie theater and uses his eco-friendly travel recommender to plan the route.

4.1 Query Personalization and Contextualization

The recommendation service interacts with the routing engine and retrieves a number of routes to present to John. According to the user profile, John owns a car, has no disabilities and has described himself as a 'complacent car addict'. According to the contextual information elements, the weather conditions are good, traffic is low and the trip is for leisure. A number of results are retrieved from the routing engine per travel profile.

4.2 User Preference Elicitation

John is asked to define the poor, fair and good levels of each option per criterion (Figure 3.a). Normalized scales are selected for the criteria in order to make the alternatives comparable. Similarly to [16] we employ qualitative scales which are then transformed to numerical values according to the rank order rule for further processing. The numerical mapping is 1 for poor, 2 for fair and 3 for good.

Furthermore John specifies the relative importance of criteria on a percent range, with weights summing up to a total of 100% as shown in Figure 3.b. Changes in one of the sliders in Figure 3.b adapt the values of the rest of the criteria so as to preserve the total of 100. In order to ease user input we can determine a set of predefined profiles (e.g. in the Figures we see that the 'Leisure' preferences profile has the Delay criterion set to 10-30 minutes and the 'Importance on Delay' option assigns higher weight to the 'Delay' criterion).

4.3 Utility Calculation

In this scenario we use the Ordered Weighted Averaging (OWA) MCDM method [19]. With OWA the normalized criteria values a_{ij} (numerical values of the poor, fair, good selections) are multiplied with the corresponding importance weights w_j (importance percentages). Next, rather than being aggregated, weighted criteria values $b_{ij} = a_{ij}w_j$ for each alternative i are re-ordered by descending value so that $b_{i_1} > ... > b_{i_n}$. An OWA operator is applied to the ordered criteria values that can potentially emphasize the better or the poorer values. At this preliminary phase of this work we opt for the neutral operator [19] which assigns equal weights to each criterion and the final utility scores are calculated as the weighted sum of the criterion values.

4.4 Recommendations

Using algorithm Greedy, as explained in Section 2.3, we generate lists of recommended routes by combining results from travel profiles (see Figure 4). The total utility and CO2 emissions of each list are calculated as the sum of the utilities and emissions of each element in the list. The 'Recommendation List 1' has the highest utility for John. The acceptable difference indicates that the recommendation lists one to three should be considered and from those, list 2 has the lowest emissions and is presented to John:

1. Using only his car, John can reach his destination within 30 minutes.

2. Using his car to a parking spot near his destination and then walk for 15 minutes, John can reach his destination within 40 minutes but save 20% of CO2 emissions.

3. Using his car to reach a bus stop close to his home John can reach his destination within 30 minutes and save 30% of CO2 emissions.

John decides to follow option 2 to reach his destination and save 20% of CO2 emissions.

5. RELATED WORK

Commonly, recommender systems generate prioritized lists of unseen items, e.g., music, books, by trying to predict a user's preferences based upon their profile. Travel recommender systems are designed to support travel planning decisions before travel or while on-the-move [6]. These systems capture user preferences, either explicitly or implicitly and suggest destinations to visit, points of interest (POIs), events or activities and/ or alternative routes. The main objective of a travel recommender system is to ease the information search process of the traveler and to convince her of the appropriateness of the proposed services [10].

With respect to route suggestion, certain systems consider multi-modal itineraries (i.e. routes that involve the use of more than one transportation means, for example reaching the destination with a combination of car, bus and walking). Tumas and Riccie [18] present a personalized mobile city transport advisory system that allows users to receive recommendations for personalized paths between two arbitrary points in the city of Bolzano on their mobile phone. They specify travel and user profiles which are then utilized to rank different multi-modal routes in the city and present the top ranked to users. They focus on computing suggestions according to users' travel-related preferences captured through questionnaires and based on four criteria: walking, bus changes, time of arrival at the destination and sightseeing. Zenker and Bernd [21] combine event recommendations and pedestrian navigation with (live) public transport support in order to assist passengers in finding interesting events and navigating to them.

Decision making is a central component in route planing applications [8]. In this respect, MCDM techniques have been employed to model combinations of user desires and to allow users to specify their personal decision strategies while receiving personalized alternatives adjusted to their needs. This view is similar to recent definitions of recommendation problems as MCDM problems. Multi-criteria based recommenders provide suggestions by modeling a user's utility for an item as a vector of ratings along several criteria. A comprehensive study of recommender systems based on MCDM methods was done in [11]. Other related work includes Nadi and Delavar [14] who study the use of OWA in a route planning application in order to generate personalized routes while considering user specific decision strategies in a manner that can provide multiple options regarding the user's preferred decision strategies. Additionally Rinner and Raubal [16] proposed a personalized location-based service using OWA, for finding the most preferable hotel.

Diversity in recommender systems is a research stream that tries to provide solutions beyond improving the relevance of recommendations [13]. Most approaches focus on attribute-based diversification, e.g. in [23], an order-independent intra-list similarity metric to assess the topical diversity of recommendation lists and a topic diversification approach for decreasing the intra-list similarity is introduced. Recently, explanation based approaches have been suggested [20]. They rest on the premise that for two different recommended items i and j, the closer their explanations (i.e., the sets of items that are similar to the recommended items and that are liked by the user), the more homogeneous i and j. Last but not least, in [1] a number of recommendation ranking techniques for diversity are proposed.

6. CONCLUSIONS AND FUTURE WORK

The environmental problems of our times demand new methods and applications able to provide nudges to citizens towards pro-environmental traveling choices. We suggested an approach that infuses nudges in travel recommenders. Users are presented with route alternatives that reside within the limits of their preferences and yield reduced carbon emissions. Furthermore we described a system architecture which combines multi-criteria recommendation techniques with profile matching methods.

There are various aspects that need further research. First, the field of MCDM encompasses a number of methods which could potentially fit into our problem, such as the Analytical Hierarchy Process and the Linguistic Ordered Weighted Averaging (LOWA). Our plan is to examine the applicability of these methods as well as compare them. Second, implementing and comparing the suggested algorithms for striking a good balance between routes utility and eco-friendliness will reveal which are best suited for this problem. Third we will investigate combinations of the proposed recommender with persuasive interfaces for eco-feedback. We expect that such interfaces, by informing users of their carbon footprint, can persuade them to choose the recommendations with low environmental impact. Last, we are going to evaluate our proposed approach in real life situations. To this direction, a prototype system that materializes the approach and related architecture is under development within the Peacox FP7 project. In addition two field trials in Vienna and Dublin have already been planned. In these trials a set of 65 users will use and evaluate the proposed recommendation service in their everyday life for a total duration of two weeks.

7. ACKNOWLEDGMENTS

Research reported in this paper has been partially funded by the European Commission in the Information Society Technologies (IST) project "Ecological aware navigation: usable persuasive trip advisor for reducing CO2 consumption" (Contract no.: 63761).

8. REFERENCES

[1] G. Adomavicius and Y. Kwon. Improving Aggregate Recommendation Diversity Using Ranking-Based Techniques.

[2] J. Anable. Complacent car addicts or aspiring environmentalists? identifying travel behaviour segments using attitude theory. *Transport Policy*, 12(1):65 – 78, 2005.

[3] E. Avineri and O. Waygood. Applying goal framing to enhance the effect of information on transport-related co2 emissions. 9th Biennial Conference on Environmental Psychology, 2011.

[4] M. Ben-Akiva and S. Lerman. *Discrete Choice Analysis: Theory and Application to Travel Demand.* MIT Press, 1985.

[5] P. Eslambolchilar, M. L. Wilson, and A. Komninos. Nudge: influence through mobile devices. In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*, MobileHCI '10, pages 527–530, New York, NY, USA, 2010. ACM.

[6] D. R. Fesenmaier, H. Werthner, and K. W. Wober. *Travel Destination Recommendation Systems: Behavioural Foundations and Applications (Cabi Publishing).* CAB International, 2006.

[7] J. Fülöp. Introduction to Decision Making Methods. Technical report, Hungarian Academy of Sciences, 2005.

[8] M. J. Geiger and W. Wenger. On the interactive resolution of multi-objective vehicle routing problems. In *Proceedings of the 4th international conference on Evolutionary multi-criterion optimization*, EMO'07, pages 687–699, Berlin, Heidelberg, 2007. Springer-Verlag.

[9] E. J. Johnson, S. B. Shu, B. G. C. Dellaert, C. Fox, D. G. Goldstein, G. Haubl, R. P. Larrick, J. W. Payne, E. Peters, and D. Schkade. Beyond nudges: Tools of a choice architecture. *MARKETING LETTERS*, 23(2):487 – 504, 2012.

[10] T. Mahmood, F. Ricci, A. Venturini, and W. Hopken. Adaptive recommender systems for travel planning. *Strategies*, pages 1–11, 2008.

[11] N. Manouselis and C. Costopoulou. Analysis and classification of multi-criteria recommender systems. *World Wide Web*, 10(4):415–441, Dec. 2007.

[12] A. D. May, M. Page, and A. Hull. Developing a set of decision-support tools for sustainable urban transport in the uk. *Transport Policy*, 15(6):328 – 340, 2008.

[13] S. M. McNee, J. Riedl, and J. A. Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI '06 extended abstracts on Human factors in computing systems*, CHI EA '06, pages 1097–1101, New York, NY, USA, 2006. ACM.

[14] S. Nadi and M. R. Delavar. Multi-criteria, personalized route planning using quantifier-guided ordered weighted averaging operators. *International Journal of Applied Earth Observation and Geoinformation*, 13(3):322–335, 2011.

[15] E. Pariser. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Press HC, The, May 2011.

[16] C. Rinner and M. Raubal. Personalized multi-criteria decision strategies in location-based decision support. *Journal of Geographic Information Sciences*, 10:149–156, 2004.

[17] R. H. Thaler and C. R. Sunstein. *Nudge: Improving decisions about health, wealth and happiness*. Penguin, London, Mar. 2008.

[18] G. Tumas and F. Ricci. Personalized Mobile City Transport Advisory System. In W. Höpken, U. Gretzel, and R. Law, editors, *Information and Communication Technologies in Tourism 2009*, chapter 15, pages 173–183. Springer Vienna, Vienna, 2009.

[19] R. R. Yager. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Trans. Syst. Man Cybern.*, 18(1):183–190, Jan. 1988.

[20] C. Yu, L. Lakshmanan, and S. Amer-Yahia. It takes variety to make a world: diversification in recommender systems. In *Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology*, EDBT '09, pages 368–378, New York, NY, USA, 2009. ACM.

[21] B. Zenker and B. Ludwig. Rose - an intelligent mobile assistant - discovering preferred events and finding comfortable transportation links. In *ICAART (1)*, pages 365–370, 2010.

[22] Y. Zheng and X. Xie. Learning travel recommendations from user-generated gps traces. *ACM Trans. Intell. Syst. Technol.*, 2(1):2:1–2:29, Jan. 2011.

[23] C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*, WWW '05, pages 22–32, New York, NY, USA, 2005. ACM.