

A Personalized Location Aware Multi-Criteria Recommender System Based on Context-Aware User Preference Models

Salvador Valencia Rodríguez and Herna Lydia Viktor

School of Electrical Engineering and Computer Science, University of Ottawa
800 King Edward Road, Ottawa, Ontario, Canada
{svale054,hviktor}@uottawa.ca

Abstract. Recommender Systems have been applied in a large number of domains. However, current approaches rarely consider multiple criteria or the level of mobility and location of a user. In this paper, we introduce a novel algorithm to construct personalized multi-criteria Recommender Systems. Our algorithm incorporates the user's current context, and techniques from the Multiple Criteria Decision Analysis field of study to model user preferences. The obtained preference model is used to assess the utility of each item, to then recommend the items with the highest utility. The criteria considered when creating preference models are the user location, mobility level and user profile. The latter is obtained considering the user requirements, and generalizing the user data from a large-scale demographic database. The evaluation of our algorithm shows that our system accurately identifies the demographic groups where a user may belong, and generates highly accurate recommendations that match his/her preference value scale.

Keywords: Recommender Systems, Location Aware, Multi-Criteria, Preference Models, Personalization.

1 Introduction

Consumers often find themselves in situations where they have to choose one item over others. For instance, they may wish to decide which movie to view, which book to read, what items to buy, and so forth. However, due to the advance in technology and the large amount of information available in databases, our options have dramatically increased. We have now reached a point where we have thousands of options at our fingertips, through the use of web-based systems. Consequently, in order to address this information overload, and aid consumers through these daily decision-processes, Recommender Systems have been developed. These systems aim to provide personalized services to each user, showing them only the information that they are most likely to be interested into [1].

A number of techniques to predict the best items have been proposed. While some of these Recommender System implementations have been successful in many domains, a number of challenges still remain. Most implementations consider only

single criteria ratings, and consequently are unable to identify *why* a user prefers an item over others. Some systems classify the user into one single group or cluster, an approach that has limited use, since real world users share commonalities in different degrees with diverse types of users. Finally, other systems require a large amount of previously gathered data about users' interactions and preferences, in order to be successfully applied.

This work introduces an algorithm to overcome these previously mentioned disadvantages. Our algorithm, as presented in this paper, builds user preference models considering multiple criteria. That is, we include the users' special needs, and context in the decision making process. This enables our system to clearly identify most important criteria, for a specific user when selecting an item over the others, and correspondingly create accurate recommendations that match his/her preference's value scale. Moreover, by including the user context as part of the recommendation process, the system is able to produce different types of recommendations to the same user, depending on his/her current context. Additionally, our algorithm exploits the information contained in a large-scale demographic database to generalize the information as provided by the user. This is done by clustering the user into one or more demographic groups. This aspect allows our system to leverage commonalities between similar user types, and create richer user profiles without the need of previously stored data about other users' interactions.

This paper is organized as follows. Section 2 details our recommender system algorithm. Section 3 describes our case study and Section 4 concludes.

2 Personalized Location-Aware System

We introduce a Personalized Multi-Criteria Context-Aware Recommender-System, which has been designed to achieve the following goals. Firstly, we aim to identify the user preferences, as based on multiple criteria that lead him/her to select an item over others, and correspondingly potentially generate more accurate recommendations. Further, our objective is to model and consider the user context during the recommendation process, instead of using it only as a filter. That is, we follow a context-aware approach, where the context is used as one of multiple criteria for the creation of preference models [6, 7]. Finally, our ultimate goal is to produce accurate recommendations by generalizing user profiles without gathering data from previous users' interactions. We use a demographic generalization technique to exploit the information contained in large-scale demographic databases that encompass interests and preferences of similar people [8].

Our algorithm is based on the creation of a multiple criteria user preference model considering the following four criteria. These are 1) the probability that a user prefers an item, given the probabilities that the user belongs to each of the demographic clusters identified in a large-scale demographic database; 2) the weight of each item, based on its attributes and the weights the user assigns for each of them; 3) the distance between the user's current location and the item's location; and, 4) the time the user would require to reach each item, based on his/her mobility level, together with




Fig. 1. High-level schematic overview of our algorithm

other geo-spatial constraints (such as a river) and average route times. The user preference model is then used to assess the utility of each item for the user. Finally, the system will recommend to the user the items with the highest estimated utility. The following figure presents the high-level schematic overview of our algorithm.

2.1 Create the User Profile

We create a user profile consisting of two dimensions. The first is the set of all possible demographic clusters where the user potentially belongs to. The second aspect concerns how much importance a specific user assigns to an item. These dimensions may be expressed as the following two vectors. $DemogClusters(u)$, which represents the possible demographic clusters where a user may belong, along with the probability of belonging to each of them. The second is $AttributeWeights(u)$, which contains the possible items' attributes along with the user-assigned weight for each of them.

The set of demographic clusters where a user may belong (k clusters), that form the $DemogClusters(u)$ vector is obtained by applying a demographic generalization technique. First, the algorithm select the m demographic variables that best differentiate each cluster from the others using feature selection techniques. Second, the user is queried to obtain his/her information for the selected demographic variables. Third, our system obtains the probability of a user belonging to each demographic cluster, based on the user given information and the users' distribution (included in the database) for each possible value for each demographic variable. Finally, the set of possible demographic clusters where the user may belong, are those in which the product of the users' distributions for the user given information is higher than a selected threshold. The following equation presents how these clusters are obtained.

$$PossibleClusters(u) = \left\{ C_x : \prod_{j=1}^m P_{u,C_x}(DV_j = Val(DV_j, u)) \geq Threshold \right\} \forall 1 \leq x \leq N \quad (1)$$

Here, DV stands for *Demographic Variable*; $Val(DV_j, u)$ represents the value of the demographic variable j given by user u ; and finally, $P_{u, C_x}(DV_j = Y) = Z\%$ expresses that there is a $Z\%$ probability (obtain from the users' distributions) that user u belongs to cluster x , given that the value of the demographic variable j for that user is Y .

Once the demographic clusters a user could belong to have been selected, the probabilities already obtained from equation (1) may be considered as the probabilities of belonging to each of them. However, this would imply classifying the user based only on his/her demographic information, and not considering other data that might also be included the database (e.g. information about leisure, shopping, media preferences, hobbies, etc). Therefore, for databases where this type of information is available, we propose the following technique to obtain the probabilities of belonging to each of the selected demographic clusters. First, we obtain the Cartesian product of the non-demographic information included in the database for each of the previously selected clusters (e.g. Leisure X Shopping X Media). Second, we apply the k-means data mining clustering algorithm over the Cartesian product dataset, to learn the demographic clusters and create the same number of groups as previously selected clusters (k groups).

The results of the k-means algorithm are k groups, each labeled by the most representative non-demographic information it contains. The k-means algorithm also creates a table, with the distribution of instances from each demographic cluster classified into the newly created groups, as shown in Table 1.

Table 1. Instances from demographic clusters classified in the newly created groups

Total instances	Group 1	Group 2	...	Group k	
Inst(C ₁)	Inst(C ₁ ,G ₁)	Inst(C ₁ ,G ₂)		Inst(C ₁ ,G _k)	Dem. Cluster 1
Inst(C ₂)	Inst(C ₂ ,G ₁)	Inst(C ₂ ,G ₂)		Inst(C ₂ ,G _k)	Dem. Cluster 2
					...
Inst(C _k)	Inst(C _k ,G ₁)	Inst(C _k ,G ₂)		Inst(C _k ,G _k)	Dem. Cluster k
	Inst(G ₁)	Inst(G ₂)		Inst(G _k)	Total instances

Diagram annotations:

- Red arrow pointing down to the table: Total instances per demographic clusters
- Red arrow pointing down to the table: Demographic clusters
- Red arrow pointing right to the table: Newly created groups
- Red arrow pointing right to the table: Total instances per newly created groups

In Table 1, $Inst(C_x)$ represents the total number of instances belonging to the demographic cluster x , $Inst(G_y)$ represents the total number of instances assigned to the newly created group y , and $Inst(C_x, G_y)$ represents the number of instances from the demographic cluster x that were assigned to the newly created group y .

Third, we present to the user the obtained newly created groups, so he/she can select the one (or more) group(s) that best define him/her (p selected groups), based on

the non-demographic information that labels each of them. Fourth, we obtain probabilities for each of the k demographic clusters, using the following equation:

$$P(u, C_x) \forall 1 \leq x \leq k = \begin{cases} \text{If } \sum_{j=1}^p \text{Inst}(C_x, G_j) = 0 : 0 \\ \text{Otherwise} : \frac{1}{p} \sum_{j=1}^p \left[\frac{\text{Inst}(C_x, G_j) * 100}{\text{Inst}(G_j)} \right] \end{cases} \quad (2)$$

Once we have obtained the possible demographic clusters where a user may belong and computed the probability for each of them, the $\text{DemogClusters}(u)$ vector is complete. Equation (3) shows the structure of this vector, formed by pairs of the form $\langle C_x, P(u, C_x) \rangle$, where C_x represents the demographic cluster x , and $P(u, C_x)$ the probability that the user u belongs to cluster x .

$$\text{DemogClusters}(u) = \{ \langle C_1, P(u, C_1) \rangle, \langle C_2, P(u, C_2) \rangle, \dots, \langle C_k, P(u, C_k) \rangle \} \quad (3)$$

The second vector included in the user profile ($\text{AttributeWeights}(u)$) consists of pairs of the form $\langle A_x, W(u, A_x) \rangle$, where A_x represents the item's attribute x , and $W(u, A_x)$ is the weight that user u assigns to attribute x . Equation (4) shows the structure of this vector, where r represents the selected number of items' attributes that best differentiate the items between them:

$$\text{AttributeWeights}(u) = \{ \langle A_1, W(u, A_1) \rangle, \langle A_2, W(u, A_2) \rangle, \dots, \langle A_r, W(u, A_r) \rangle \} \quad (4)$$

2.2 Location Awareness and Mobility Level

The second phase of the algorithm consists of obtaining the user's mobility level and current location. Consequently, the user is asked to choose the means of transport (i.e. mobility level) that best describes his/her current situation (e.g. driving a car, riding a bike, wheelchair, or walking). Based on the user selection, the routes to be considered when predicting the distance and time to each item are selected. Subsequently the user's current latitude and longitude spatial coordinates are obtained from Google Maps®, through its application programming interface (API), using the current user address.

2.3 Assessing the Items' Utilities Considering the User Profile

The third phase of our algorithm is divided into two stages. The first stage entails obtaining an utility for each item ($\text{Demographic Utility (DU)}$), based on the previously created vector $\text{DemogClusters}(u)$. However, in order to do so, all the items in the database need to be mapped to one or more demographic clusters. Consequently, we assign each item to one or more categories, which in turn are linked to one or more demographic clusters.

Once all the items are related to one or more demographic clusters through one or more categories, we compute the probability that a user might prefer each of these categories. As shown in equation (5), these probabilities are calculated based on the probabilities of belonging to each of the k previously selected demographic cluster.

$$P(u, ItemCat_x) = \sum_{i=1}^k \begin{cases} \text{If } ItemCat_x \in C_i : P(u, C_i) \\ \text{Otherwise} : 0 \end{cases} \quad (5)$$

Here, $P(u, ItemCat_x)$ represents the probability that user u prefers the item category x . Subsequently, these probabilities are used to obtain the demographic utility (DU) for each item. This DU is computed as the average of the probabilities that the user prefers each category where the item belongs to. The DU of an item is obtained by using equation (6), where q represents the number of identified items' categories.

$$DU(I_x, u) = \frac{1}{Num\ of\ Cat\ of\ I_x} \sum_{i=1}^q \begin{cases} \text{If } I_x \in ItemCat_i : P(u, ItemCat_i) \\ \text{Otherwise} : 0 \end{cases} \quad (6)$$

The second stage involves assessing a utility for each item (*Weighted Utility (WU)*), considering now the previously created *AttributeWeights(u)* vector. The WU for each item is calculated as the sum of each user given weight for all the attributes of an item, as shown in equation (7).

$$WU(I_x, u) = \sum_{i=1}^r \begin{cases} \text{If } A_i \in I_x : W(u, A_i) \\ \text{Otherwise} : 0 \end{cases} \quad (7)$$

2.4 Assessing the Items' Utilities Considering the User Location

The fourth step in the algorithm involves assessing a utility for each item, based on the user's current location and mobility level. In this phase, each item is assessed under two additional criteria: the distance from the user (*Distance Utility (DisU)*), and the time it would take the user to reach the item (*Time Utility (TU)*), considering the user's mobility level. It is important to note that these new utilities, expressed in kilometers and minutes from the user, are inversely proportional to the item utility for the user. That is, the closer an item is the higher its utility for the user. Finally it is noteworthy to mention that while the TU is obtained from Google Maps®; the $DisU$ is obtained by means of SQL spatial queries, performed over a spatial database using the user coordinates.

2.5 Creating the User Preference Model

The utilities that were obtained during the third and fourth phases represent the four criteria considered to create the user preference model, which in turn represents the importance to the user. To obtain the user preference model we apply the UTA* algorithm. This algorithm is a regression-based technique that infers preference models from given global preferences (e.g. previous user choices) [2, 3]. For more information on this algorithm, we recommend [2-5].

In order to apply this algorithm we first need to obtain the user's weak preference order, which represents the user preferences for some items, after considering the four criterions for each of them. Consequently, the system randomly selects ten items, along with their four utilities, to present to the user, who is then asked to rate them in a non-descending order. The result of the UTA* algorithm is a vector of four values,

which represent the weights of each of these four criteria for that user. These weights are calculated using linear programming techniques. The resulting vector from the UTA* algorithm has the following form:

$$PM(u) = \{W_{DU}, W_{WU}, W_{DisU}, W_{TU}\} \quad (8)$$

2.6 Recommending the Items That Best Match the User Preferences

The last phase of the algorithm involves assessing an integrated, final utility for each item, using the previously obtained user preference model. It follows that, to apply the previously obtained weights from the model, the values of the four different utilities must be normalized so they are in the same range $\{0, \dots, 1\}$, otherwise the final utility would be biased toward the criterion with the highest scale. Therefore, each utility is divided by its highest value among all the items. Moreover, since the distance and time utilities are inversely proportional to the user utility, after dividing their values by their highest utility the resulting value is subtracted from 1, and the remainder is the utility to be considered. The final integrated item's utilities are obtained by applying the following equation:

$$U(I_x, u) = W_{DU} \frac{DU(I_x, u)}{\max(DU(u))} + W_{WU} \frac{WU(I_x, u)}{\max(WU(u))} + W_{DisU} \left(1 - \frac{DisU(I_x, u)}{\max(DisU(u))} \right) + W_{TU} \left(1 - \frac{TU(I_x, u)}{\max(TU(u))} \right) \quad (9)$$

3 Experimental Evaluation

As a case study, we created PeRS, a Personal Recommender System that recommends events to attend to consumers. The source datasets used in this case study are the following. The Ottawa Open Data Events database is our items dataset, which contains the events that will take place within the National Capital Region of Canada. By events we refer to festivals/fairs, film/new media, galleries, music, theater, and so on [9]. Our large-scale demographic database is the PRIZM C2 Database that classifies Canada's neighborhoods into 66 unique lifestyle types, providing insights into the behavior and mindsets of residents [10]. The Ottawa Open Data Geospatial databases contain spatial information to geographically locate different elements of interest (i.e. rivers, buildings, museums, municipalities, wards, roads, cities and country areas) [11].

3.1 Experimental Design

We evaluated our algorithm using an offline experiment. This type of experiment was chosen because it focuses on the recommendation technique that is being used, rather than the system interface, which makes it highly suitable for our environment [1]. Since the design and implementation of a Recommender System depends on the specific requirements, goals, and sources of information, the properties used to evaluate a system should be selected according to each application domain. The selected properties to be evaluated in our offline experiment were the system's prediction accuracy, item and user coverage, and confidence. (For more information on other properties to evaluate and compare Recommender Systems, we recommend [1]).

We performed a pilot study with 30 human subjects to obtain, according to the central limit theorem, a punctual estimation of the standard deviation of the population [11]. Table 2 presents the structure of the questionnaire answered by the human subjects that took part in the experiment. This questionnaire was designed to gather the required data to create the users' preference models, and to evaluate the recommendations.

Table 2. Structure of the questionnaire applied to the users

Section	Gathered Information
Personal Information	The user is asked 9 personal questions, along with his/her mobility level
Preferences	The user is asked to provide weights to the identified events' attributes
Activities	The user is asked to select the groups of activities where he/she identifies the most.
Events	The user is asked to rate 10 randomly selected events, considering the information for all the criteria.

3.2 Experimental Results

Table 3 presents a summary of the results obtained. Our analysis is centered on four aspects. Firstly, our goal was to analyze the distribution of the subjects considered in the experiment. Secondly, we aimed to evaluate the system accuracy and the ability to cluster the subjects into the most adequate demographic groups. Thirdly, we learned the system accuracy from the user profiles. Finally, we evaluated the system coverage, in terms of user and item space coverage.

From the results, as shown in Table 3, we may conclude that, based on their demographic information, 93% of the subjects that partook in experiment have differences in their profiles. This allowed us to test the system for a wide-spread distribution of users, and therefore potentially draw more reliable conclusions for the general population. Additionally, we concluded that the system was able to narrow down the possible matching clusters where a subject may belong, using only the selected demographic variables asked to the subjects. Furthermore, the subjects identified themselves within 86% of the clusters selected by the system, considering the characteristic activities of the people within each group. This indicates that the system accurately classified the subjects into the correct demographic groups.

The system has a predictive accuracy between 75% and 82% with a confidence of 95%. This accuracy represents the effectiveness of the system to identify the user preferences and rank the items in the same way he/she would approach this task. Moreover, for 80% of the subjects considered in our experiment the system was able to rank the items with at least 70% match with the way they would have proceeded (i.e. user space coverage). In addition, the system recommended 81% of the items in the database to at least one user (i.e. item space coverage).

Finally, from the results obtained from the performed classification and statistical analyzes, we concluded that the accuracy of the produced recommendations does not solely depend on the user profile. Therefore, it could potentially be used for a wide range of the population and produce equally accurate recommendations. We will explore the validity of this observation in our future work.

Table 3. Experimental results

Property	Measured Value	Interpretation
Users Distribution	93.33% of the users in the experiment have different profiles (based on their personal information).	The users present an equal distribution in most of the considered personal information. Correspondingly this wide spread distribution of users allows the system to be tested for different user types, and therefore draw reliable conclusions for a wide range of the population.
Narrowing Percentage	The system identified an average of 4 out of 66 possible clusters where a user may belong	The system is able to narrow in a 95.38% the possible demographic groups where a user may belong, during the demographic generalization process.
Accuracy to identify the matching clusters	26 out of 30 users identified themselves, in every one of the selected demographic groups for him/her.	86.66% of the users identified themselves with the activities included in the demographic groups selected by the system.
Prediction Accuracy and Confidence	Accuracy: {75.36 % – 82.96%} Confidence: 95%	The system recommends items that match the user preferences with an average of 79.16%, for 95% of the population.
Learning the System Accuracy	Based on the results obtained from applying classification and statistical regression analyzes, using a decision tree, we conclude that the system accuracy doesn't depend on the user profile. Therefore it follows that the system can produce equally accurate recommendations for a wide range of the population.	
User Space Coverage	For 80% of the users the system created models that can rank the events with at least a 70% match with how the users would have ranked them.	
Item Space Coverage	81.97% of the items in the database were recommended to a user at least once, with at least 50% match of his/her preferences.	

4 Conclusion

This paper presents a Recommender System that utilizes an algorithm to generate user models which includes the user context and preferences in the decision model. Our algorithm, as shown in the experiments, accurately identifies the demographic groups where a user may belong. Our results indicate that we are able to produce highly accurate recommendations for a wide range of the population. Further, we are able to construct accurate user profiles. Our algorithm has the following noteworthy characteristics. It uses more than one decision criteria and is therefore able to identify what it is most important for a user when selecting an item over the others. It creates user preference models, consequently the system is capable to understand the user interests and preferences, and is not biased from previous user actions. It generalizes the user profile, thus the system is capable to reason from the commonalities between demographic groups, and it doesn't require previous gathered information about other users' interactions. Importantly, it is location aware and as a result the system varies the recommendations as the user location or level of mobility changes. Our future work will focus on producing accurate preference models in an imbalanced preference

setting. This would imply that our algorithm should distinguish user preferences, even when they may be very similar. Additionally, it would be interesting to evaluate the prediction accuracy of our system considering a larger number of users and different locations. Finally, it would be worthwhile to compare the obtained results from our system, against other linear and non-linear recommendation techniques [12].

References

1. Ricci, F., Rokach, L., Shapira, B., Kantor, P.B.: *Recommender Systems Handbook*. Springer (2011)
2. Lakiotaki, K., Matsatsinis, N.F., Tsoukias, A.: Multi-Criteria User Modeling in Recommender Systems. *Intelligent Systems* 26(2), 64–76 (2011)
3. Siskos, Y., Grigoroudis, E., Matsatsinis, N.: UTA Methods. In: *Multiple Criteria Decision Analysis: State of the Art Surveys*, pp. 297–344. Springer, New York (2005)
4. Siskos, Y., Yannacopoulos, D.: UTASTAR: An ordinal regression method for building additive value functions. *Investigação Operacional* 5(1), 39–53 (1985)
5. Jacquet-Lagrèze, E., Siskos, J.: Assessing a set of additive utility functions for multicriteria decision making: The UTA method. *European Journal of Operational Research* 10(2), 151–164 (1982)
6. Lee, J., Mateo, R.M.A., Gerardo, B.D., Go, S.-H.: Location-Aware Agent Using Data Mining for the Distributed Location-Based Services. In: Gavrilova, M.L., Gervasi, O., Kumar, V., Tan, C.J.K., Taniar, D., Laganá, A., Mun, Y., Choo, H. (eds.) *ICCSA 2006*. LNCS, vol. 3984, pp. 867–876. Springer, Heidelberg (2006)
7. Schilit, B., Theimer, M.: Disseminating active map information to mobile hosts. *IEEE Network* 8(5), 22–32 (1994)
8. Krulwich, B.: Lifestyle Finder. *Intelligent User Profiling Using Large-Scale Demographic Data*. *AI Magazine* 18(2), 37–45 (1997)
9. City of Ottawa, Ottawa Open Data, <http://ottawa.ca/en/open-data-ottawa>
10. PRIZM C2 Segmentation System, <http://www.environmentalanalytics.ca/>
11. Berenson, M., Levine, D.: *Basic Business Statistics: Concepts and Applications*. Prentice Hall (1995)
12. Klačnja-Milićević, A., Ivanović, M., Nanopoulos, A.: The use of Nonlinear Manifold Learning in Recommender Systems. In: *Proc. of the International Conference on Information Technology (ICIT 2009)*, Amman, Jordan, pp. 2–6 (2009)