Context Relevance Assessment for Recommender Systems

Linas Baltrunas
Centre for Database and Information Systems
Free University of Bolzano
Piazza Domenicani, 3,
39100 Bozen-Bolzano, Italy
lbaltrunas@unibz.it

Bernd Ludwig
Centre for Database and Information Systems
Free University of Bolzano
Piazza Domenicani, 3,
39100 Bozen-Bolzano, Italy
bernd.ludwig@unibz.it

Francesco Ricci
Centre for Database and Information Systems
Free University of Bolzano
Piazza Domenicani, 3,
39100 Bozen-Bolzano, Italy
fricci@unibz.it

ABSTRACT
Research on context aware recommender systems is taking for granted that context matters. On the contrary, often attempts to show the influence of certain contextual conditions have failed, or succeeded only using context features that are obviously correlated with the rating to be predicted. In this paper we consider the problem of quantitatively assessing context relevance. For this purpose we are assuming that users can: imagine a situation described by contextual features, and judge if these features are relevant for their decision making task. We have designed a UI suited for acquiring such information in a travel planning scenario. In fact this interface is generic and can also be applied to other domains (e.g., music). The experimental results clearly show that it is possible to identify the contextual dimensions that are relevant for the given task and this relevancy depends on the typology of the point of interest to be included in the plan.

INTRODUCTION
Recommender Systems (RSs) are software tools and techniques providing suggestions for items to be of use to a user [2]. Often, recommendations can be improved if the context of the recommendation is known, e.g., in a travel recommender, the means of transportation or the season of the travel. For this reason, context-aware recommender systems (CARSs) are gaining more and more attention [3]. Various approaches have been used to incorporate contextual information into recommender systems, improving performance measures, such as: mean absolute error [4], or recall [1], or prediction accuracy [8].

However, to adapt the recommendations to the context the dependency of the user preferences from the contextual conditions must be modeled and example data must be acquired. This requires, for instance in the Collaborative Filtering approach [2], to record explicit user evaluations (ratings) for items in alternative contexts (e.g., the rating for a movie to be watched with the partner). Such data is difficult to obtain because it requires a substantial user effort, since the user must provide items' evaluations (ratings) in several different contextual conditions. Moreover, one can set up a process for acquiring such ratings and discover a posteriori that the selected contextual conditions where actually irrelevant (i.e., the rating is not influenced), or simply they do not help in improving the recommender system effectiveness [5, 9].

Selecting a recommendation is an example of a difficult decision that the user may take while performing an activity. Generating good recommendations is hard because: they are evaluated subjectively, there is no *correct* decision, and the recommender system's knowledge about the user's current attitude is largely *uncertain*. Even worse, the user's decision is mostly influenced by contextual conditions that are different every time the decision has to be taken. As an illustrative example, take the two recommended routes in Fig. 1 for visiting the city of Cles starting from Bolzano by car. Both of them are correct; however, they have different properties: for users with a motor bike, route 2 would be a great experience as it includes the famous Mendelpass while users with children in the car would prefer route 1, a more comfortable, although longer route on the highway. Even such preferences may change depending on weather and traffic conditions, for example. Hence, a major initial issue for the correct design of a CARS is the assessment of the contextual factors that are worth considering when generating recommendations. This is not an easy problem: it requires to formulate informed conjectures about the influence of some data, before collecting the real data. It is a kind of active learning problem, where the relevance of the data to acquire must be estimated to minimize the cost of real data acquisition [10].

The main contribution of this paper is a methodology for the quantitative assessment of the dependency of user preferences from a set of contextual dimensions. This methodology is based on a tool for acquiring context relevance judgements and a statistical data analysis method for identifying the contextual dimensions that are more likely to influence the user decisions for different items' types. This approach can be adopted after a qualitative study, such as a diary study, has revealed the contextual dimensions that

Submitted for review to IUI 2011.

are potentially relevant for a users' population. The proposed methodology has been tested on a travel planning application aimed at recommending points of interests (POIs) to mobile users[1]. The mobile assistant we are developing in this scenario is planned to offer two main functionalities. Firstly, context-dependent and personalized recommendation of touristic POI. Secondly, assistance in the preparation of an complete itinerary and the modification of the itinerary according to circumstances and eventualities that occur during the itinerary.

ACQUIRING CONTEXT RELEVANCE

In order to assess the influence of alternative contextual conditions on user decisions, we collected data describing how users **change** their inclination to visit a POI while they imagine that certain contextual circumstances hold. For that purpose, we designed an online survey. A large set of contextual conditions (as found in the relevant literature [11]) and a (relatively small) list of categories for POIs in Bolzano (and other cities) have been incorporated in a web form (see Figure 2). POIs were aggregated into categories in order to avoid sparseness of the collected data. We defined eleven categories: castle; nature wonder; cycling and mountain biking; theater event; folk festival, arts and crafts event; church or monastery; museum; spa and pampering; music event; walking path. In the web application, users could indicate the influence of these conditions on their decision to visit POIs belonging to a randomly selected item category. The influence is measured with three values: positive, negative or neutral. Three different contextual conditions were tested in a single page while a full questionnaire consisted of five of such pages (as in Figure 2).

We observe that [7] already tried to estimate the impact of contextual conditions on the user evaluations by asking the user to imagine a given contextual condition. They have shown that this method must be used with care as users rate differently in real and supposed contexts. When the context is just supposed there is a tendency of the users to exaggerate its importance. In fact, in our case we are trying to measure

[1]It is also being tested on a in car music recommendations scenario (not illustrated here for lack of space).

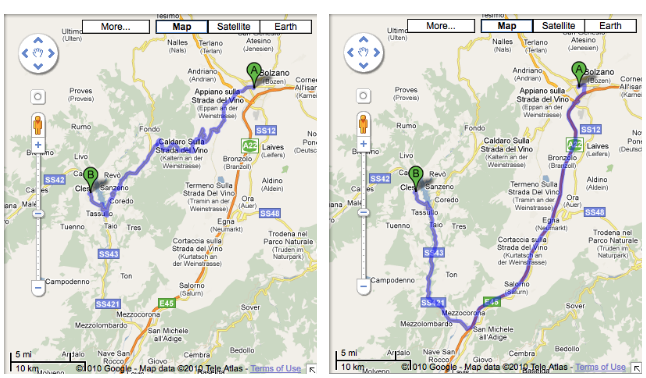

Route 1 (50.7 km/58 min) Route 2 (66.2 km/55 min)

Figure 1. Comparison of different Routes from Bolzano to Cles

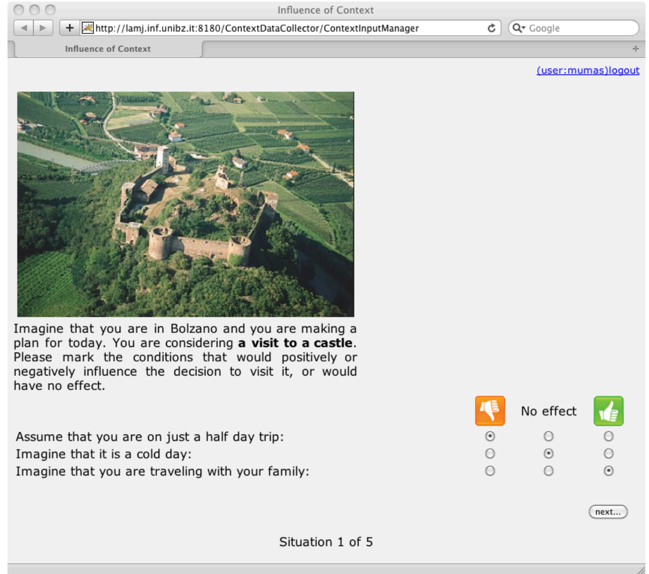

Figure 2. Interview conducted with the web survey tool

only if a contextual condition has an influence (positive or negative) on the user's decisions and not the real value of the user's ratings. For instance, we want to understand if the proximity to a POI is influential and not how the rating for a precise POI changes as a function of the user proximity. Moreover, as it is shown later, our statistical approach can predict that a context dimension does influence the user with a given reliability measure. So, considering only conditions with high reliability we can reduce significantly the number of false positives. In addition, our method is proposed as a tool for selecting potentially relevant contextual conditions; while the true evaluations/ratings of the items under the selected contextual conditions can be acquired in a classical way by asking the users to rate items when they are really experienced in a contextual situation (which is one of the next steps of our future work).

33 participants (mostly from computer science departments) took part in the web survey. Overall, they gave 1524 responses. In a single response to one of the questions shown in Fig. 2, the user tells which influence one contextual condition has on his decision to visit an item of a given category. For the specification of the context, the factors presented in Tab. 1 were applied in a randomized way: for each question a category is drawn at random along with a value for a context factor. This sampling has been implemented such that a uniform distribution over the possible categories and the possible values is achieved. A different sampling is also applicable if a prior distribution is known.

ANALYSIS

With the web survey we aimed at finding indications about which context factors influence user decisions whether to visit or not a POI. As no information about the relationship between response variable and context was available, parametric tests such as χ^2 were not applicable. Therefore, a non-parametric statistical analysis seemed to be more appro-

Context Factor	**Values**	**Context Factor**	**Values**	**Context Factor**	**Values**	**Context Factor**	**Values**
budget	budget traveler	crowdedness	not crowded	companion	with girl-/boy-friend	season	spring
	high spender		crowded		with family		summer
	price for quality		empty		with children		autumn
time of the day	morning time		health care		alone		winter
	afternoon	travel goal	cultural experience		with friends	transport	public transport
	night time		scenic/landscape	weather	snowing		no means of transport
day of the week	weekend		education		clear sky		bicycle
	working day		hedonistic/fun		sunny		car
distance	near by		social event		rainy	temperature	warm
	far away		religion		cloudy		cold
knowledge	new to city		activity/sport	mood	happy		hot
about about area	citizen of the city		visiting friends		active	time available	half day
	returning visitor		business		sad		more than a day
							one day

Table 1. Overview of the context factors used in the web survey

priate: The web survey delivered samples for the distribution

$$P(I|T, C_1, \ldots, C_N) = \frac{P(I, T, C_1, \ldots, C_N)}{P(T, C_1, \ldots, C_N)} \approx \left(\prod_{i=1}^{N} \frac{P(I|C_i, T)}{P(I|T)} \right) \cdot P(I|T)$$

where I (Influence) is the response variable, taking one of the three values: positive, negative, or neutral. T is a POI category, and the C_1, ..., C_N are the context factors that (may or may not) influence the user decisions. The probabilities $P(I|C_i, T)$ model the influence of the context factors on the user's decision. The knowledge of $P(I|C_i, T)$ can drive the acquisition of context-dependent ratings for the context factors that have a large probability to increase or decrease the user evaluation for the items in a given category T. Hence, it is interesting to understand which C_i have impact on I, or in other words, which C_i explain I better than other context factors.

Statistical Methodology

The spread of a categorical variable $X = \{x_1, \ldots, x_n\}$ can be measured by looking at the entropy of the random variable [6]. If $P(X = x_i) = \pi_i$, the entropy of X is:

$$E(X) = -\sum_{1 \le i \le n} \pi_i \cdot \log \pi_i$$

This measure of the spread can be used to estimate the association of two variables X_1 and X_2, i.e., how well one variable explains the other. In the considered tourist recommendation scenario X_1 is the variable *Inclination of the user to visit an item*, while X_2 is a factor in the context of the current situation which may have an influence on the user's decision, e.g. *Current weather condition*. Informally, this influence is strong if the knowledge about the weather reduces the spread of X_1, and it is weak if the spread of X_1 remains unchanged even if one knows the weather. Therefore, the difference between the spread of X_1 and the expected spread of $(X_1|X_2)$ if a measure for the association of X_1 and X_2. As the spread of $(X_1|X_2)$ should not be larger than that of X_1 alone we can normalize the difference to the interval $[0, 1]$ by:

$$U = \frac{E(X_1) - \mathcal{E}(E(X_1|X_2))}{E(X_1)}$$

$\mathcal{E}(X)$ denotes the expected value of the random variable X. U is 1 if the spread of $(X_1|X_2)$ is zero. This occurrs if for each value of X_2 the value X_1 is certain (i.e. X_1 is a deterministic function of X_2). U is zero, however, if X_2 does not have any influence of X_1, in which case the spread of $(X_1|X_2)$ is not different from that of X_1. Using entropy to measure spread, we get the following formula:

$$U = -\frac{\sum_{1 \le i \le k} \sum_{1 \le j \le l} \pi_{i,j} \cdot \log \left(\frac{\pi_{i,j}}{\pi_{i,\bullet} \pi_{\bullet,j}} \right)}{\sum_{1 \le j \le l} \pi_{\bullet,j} \cdot \log \pi_{\bullet,j}}$$

where, $\pi_{i,j} = P(X_1 = x_i, X_2 = y_j)$. X_1 and X_2 are categorical variables with $X_1 = \{x_1, \ldots, x_k\}$ and $X_2 = \{y_1, \ldots, y_l\}$. $\pi_{i,\bullet} = \sum_{1 \le j \le l} \pi_{i,j}$ and $\pi_{\bullet,j} = \sum_{1 \le i \le k} \pi_{i,j}$. U can be seen as the mutual information of X_1 and X_2 normalized to the interval $[0, 1]$.

Results

Given this definition, U can be used to measure how good I – the influence of context on the user's decision – can be predicted if C_i – one of the relevant context factors – is known. Therefore, in order to understand which context factors help most to decrease the uncertainty about I, we have computed U for all factors and POI categories. Ordering the factors in descending value of U, one gets the results reported in the Appendix of this paper. That table indicates that there are some factors that indeed seem to be relevant for all the categories, among them *distance*, *time available*, *crowdedness*, and *knowledge of the surroundings*. Others often appear to be less relevant: *transport*, *travel goal*, *day of the week*. Finally some factors appear to have a different relevance depending on the category.

CONCLUSIONS AND FUTURE WORK

In this short paper we have illustrated a methodology and a tool for acquiring explicit users' evaluations about the relevancy of contextual factors for item selection in a recommender system. Contextual information is known to have a large impact on user decision making but often the relationship between context and decision is largely unknown and uncertain. Which contextual factor is relevant, in a specific decision making situation, is hard to predict and wrong assumptions may lead to unnecessary and misleading reasoning models. The proposed methodology tackles these problems and has been applied to a travel planning scenario. It has been shown that tourists' preferences are strongly influenced and vary significantly with respect to context and item category. The proposed methodology provides quantitative measures of context relevancy, complementing other qualitative approaches and results coming from consumer behavior literature [11]. The collected data are now being used in a mobile tourist assistant that pushes new recommendations to tourists when contextual conditions changes.

In conclusion, we have shown how the uncertain relationships between context and decision can be explored and measured. We are applying the proposed approach in a different decision making scenario, namely music recommendation for a group of passengers in a car, to understand to what extend the approach can be generalized to other tasks.

REFERENCES

1. G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin. Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Trans. Inf. Syst.*, 23(1):103–145, 2005.

2. G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: a survey of the state-of- the-art and possible extensions. *Knowledge and Data Engineering, IEEE Transactions on*, 17:734–749, 2005.

3. G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. In F. Ricci, L. Rokach, B. Shapira, and P. Kantor, editors, *Recommender Systems Handbook*. Springer Verlag, 2010. in press.

4. L. Baltrunas and F. Ricci. Context-based splitting of item ratings in collaborative filtering. In R. Burke, A. Felfernig, and L. Schmidt-Thieme, editors, *RecSys '09: Proceedings of the 2009 ACM conference on Recommender systems*, October 22-25, New York, USA, 2009. ACM Press.

5. S. Jeong, S. Kalasapur, D. Cheng, H. Song, and H. Cho. Clustering and naive bayesian approaches for situation-aware recommendation on mobile devices. In *International Conference on Machine Learning and Applications, ICMLA 2009, Miami Beach, Florida, USA, December 13-15, 2009*, pages 353–358, 2009.

6. C. J. Lloyd. *Statistical Analysis of Categorical Data*. Wiley-Interscience, 1999.

7. C. Ono, Y. Takishima, Y. Motomura, and H. Asoh. Context-aware preference model based on a study of difference between real and supposed situation data. In *User Modeling, Adaptation, and Personalization, 17th International Conference, UMAP 2009, Trento, Italy, June 22-26, 2009*, pages 102–113, 2009.

8. C. Palmisano, A. Tuzhilin, and M. Gorgoglione. Using context to improve predictive modeling of customers in personalization applications. *IEEE Transactions on Knowledge and Data Engineering*, 20(11):1535–1549, Nov. 2008.

9. K. Partridge and B. Price. Enhancing mobile recommender systems with activity inference. In *User Modeling, Adaptation, and Personalization, 17th International Conference, UMAP 2009, Trento, Italy, June 22-26, 2009*, pages 307–318, 2009.

10. N. Rubens, D. Kaplan, and M. Sugiyama. Active learning in recommender systems. In F. Ricci, L. Rokach, B. Shapira, and P. Kantor, editors, *Recommender Systems Handbook*. Springer Verlag, 2010. in press.

11. J. Swarbrooke and S. Horner. *Consumer Behaviour in Tourism*. Butterworth-Heinemann, 2nd edition, 2006.

APPENDIX

Ranking of context factors to their association with the user responses on the influence of a factor on their decision to visit an item:

castle	**church or monastery**	**cycling or mountain biking**	**folk festival, arts and crafts event**	**museum**	**music event**	**nature wonder**	**spa**	**theater event**	**walking**
distance	distance	time available	distance	distance	crowdedness	day week	distance	distance	time available
day week	time available	budget	temperature	temperature	day week	distance	knowledge about area	time available	distance
crowdedness	companion	season	knowledge about area	budget	time available	temperature	crowdedness	day time	budget
knowledge about area	day time	crowdedness	season	crowdedness	mood	season	time available	budget	temperature
season	transport	temperature	weather	knowledge about area	budget	crowdedness	budget	temperature	crowdedness
budget	travel goal	weather	day time	time available	companion	weather	season	day week	knowledge about area
time available	mood	day time	budget	weather	distance	time available	temperature	mood	season
day time	temperature	mood	time available	companion	knowledge about area	companion	travel goal	knowledge of surrounding	weather
companion travel goal	weather crowdedness	day week knowledge of surroundings	companion crowdedness	mood travel goal	day time transport	day time travel goal	weather companion	travel goal season	day time transport
mood transport temperature	season budget knowledge about area	transport travel goal companion	day week travel goal transport	season transport day week	temperature travel goal season	mood transport budget	mood transport day time	weather companion crowdedness	mood day week companion
weather	day week	distance	mood	day time	weather	knowledge about area	day week	transport	travel goal