Context Relevance Assessment for Recommender Systems

Linas Baltrunas
Free University of Bozen-Bolzano
Piazza Domenicani, 3
39100 Bozen-Bolzano, Italy
lbaltrunas@unibz.it

Bernd Ludwig
Free University of Bozen-Bolzano
Piazza Domenicani, 3
39100 Bozen-Bolzano, Italy
bernd.ludwig@unibz.it

Francesco Ricci
Free University of Bozen-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. But, often attempts to show the influence of context have failed. 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 a contextual feature, and judge if this feature is 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 used for other domains (e.g., music). The experimental results show that it is possible to identify the contextual factors that are relevant for the given task and that the relevancy depends on the type of the place of interest to be included in the plan.

Author Keywords
Context-aware, recommender systems, user preferences.

ACM Classification Keywords
H.3.3 Information Search and Retrieval: Information Filtering.

General Terms
Algorithms, Experimentation, Human Factors.

INTRODUCTION
Recommender Systems (RSs) are tools providing suggestions for items to be of use to a user [2]. Generating good recommendations is hard because they are evaluated subjectively and the RS's knowledge about the user's current preferences is largely *uncertain*. Even worse, the user's decision is mostly influenced by contextual conditions that differ each time the decision is taken. As an illustrative example, take the two recommended routes in Figure 1 for visiting the city of Cles starting from Bolzano by car. Both of them are correct; however, they have different properties. For motor bikers, route 1 would be a great experience as it includes the famous Mendelpass while travelers with children in the car would prefer route 2, a more comfortable, although longer route on the highway.

As this example illustrates, often a recommendation can be more relevant if its context is known. For this reason, context-aware recommender systems (CARSs) are gaining more and more attention [3], and various approaches have been used to incorporate context knowledge, improving performance measures, such as: mean absolute error [4], or recall [1], or prediction accuracy [9]. However, to adapt to the context the dependency of the user preferences from the contextual conditions must be modeled. This requires to record explicit user evaluations (ratings) for items in alternative contexts, e.g., the rating for a movie after it was watched with the partner. Such data is difficult to obtain because it requires substantial user effort, since the user must provide ratings in several contextual conditions. Moreover, one can acquire such ratings and later discover that the considered contextual conditions where actually irrelevant, i.e., the ratings are not influenced, and the RS is not improved [6]. Hence, a major issue for the design of CARSs is assessing the contextual factors that are worth considering. This 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 be acquired must be estimated to minimize the data acquisition cost [10].

The main contribution of this paper is a methodology for the quantitative assessment of the dependency of the user preferences from a candidate set of contextual factors. It is based on a interface for acquiring context relevance judgements and a data analysis method for identifying the contextual factors that are likely to influence the user decisions. This approach can be adopted after a qualitative analysis, such as a diary study, has revealed the contextual factors that are potentially relevant for a recommendation task. This 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 is planned to offer context-dependent recommendations for touristic POIs [5] that are updated as soon as the contextual conditions change.

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

IUI'11, February 13–16, 2011, Palo Alto, California, USA.

ACQUIRING CONTEXT RELEVANCE

In order to assess the influence of some contextual factors on the user decisions we collected data describing how users **change** their inclination to visit a POI while they imagine that a contextual condition holds. For that purpose, a large set of contextual conditions (as found in the relevant literature [11]) and a (relatively small) list of categories of POIs in Bolzano (and other nearby 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, the users could indicate the influence of these contextual conditions on their decision to visit POIs belonging to a randomly selected category. The influence is measured with three values: positive, negative or neutral. Three different contextual conditions (i.e., values for contextual factors) were tested in a single page while a full questionnaire consisted of five of such pages (as in Figure 2).

We observe that [8] 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 only if a contextual factor 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 to what extend that a context factor does influence the user. So, considering only conditions with high influence we can reduce significantly the number of false positives. Hence, our method is proposed as a tool for selecting potentially relevant contextual factors; while the true evaluations/ratings of the items under alternative contextual conditions can be acquired in a classical way by asking the users to rate items when they are really experienced in a contextual condition (the next step of our future work).

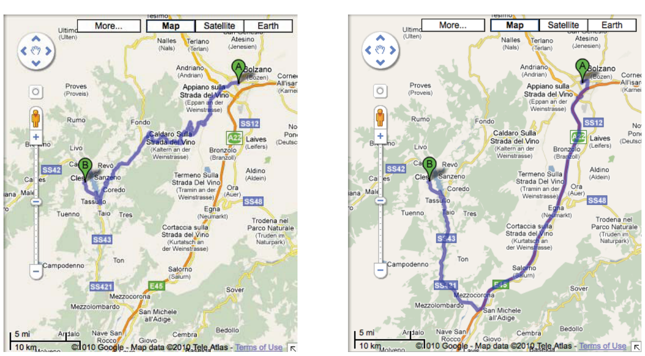

Route 1 (50.7 km/58 min) Route 2 (66.2 km/55 min)

Figure 1. Comparison of different Routes from Bolzano to Cles.

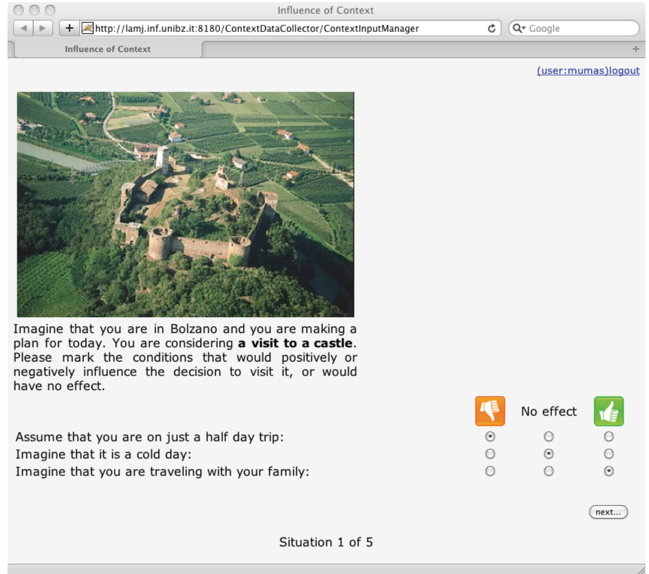

Figure 2. Web survey tool.

33 participants (mostly from our computer science faculty) took part in the web survey. Overall, they gave 1524 responses. In a single response to one of the questions shown in Figure 2, the user evaluates the influence of one contextual condition on his decision to visit an item of the given category. For the specification of the context, the factors and conditions presented in Table 1 were applied in a randomized way: for each question a category is drawn at random along with a value (condition) for a context factor. This sampling has been implemented such that a uniform distribution over the possible categories and context conditions 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 influences user decisions whether to visit or not a POI. As no information about the relationships 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 appropriate: The web survey delivered samples for the distribution

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

where I (Influence) is the response variable, taking one of the three values: positive, negative, or neutral. T is the POI category, and the $C_1, \ldots, C_N$ are the context factors that (may or may not) influence the user decision. The probabilities $P(I|T,C_i)$ model the influence of the context factors on the user's decision. The knowledge of $P(I|T,C_i)$ can drive the acquisition of context-dependent ratings for the context factors that have a large probability to increase or de-

Context Factor	Conditions	Context Factor	Conditions	Context Factor	Conditions	Context Factor	Conditions
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 to POI	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. Context factors and conditions used in the web survey

crease 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 [7]. If $P(X = x_i) = \pi_i$, the entropy of X is:

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

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 context factor of the current situation which may have an influence on the user's decision, e.g. *the 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 $(X_1|X_2)$ is 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:

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

U is 1 if the spread of $(X_1|X_2)$ is zero. This occurs 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 on 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:

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

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 $X_2 = \{y_1, \ldots, y_l\}$. $\pi_{i,\bullet} = \sum_{1 \leq j \leq l} \pi_{i,j}$ and $\pi_{\bullet,j} = \sum_{1 \leq i \leq k} \pi_{i,j}$. We note that U is the mutual information of X_1 and X_2 normalized to the $[0, 1]$ interval.

Results

Given this definition, U can be used to measure how well I (the influence) can be predicted if C_i (a context factor) is known. Therefore, in order to understand which contextual factors are more important in predicting whether the user will change his inclination to visit a POI, 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 to the POI*, *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. For lack of space we cannot fully describe the dependencies shown in this table. In fact, there are results that may appear controversial, e.g., that "distance" is not important for cycling. But, this is explained by observing that here "distance" refers to "distance to the POI", which is actually less important for a cyclist than for a pedestrian.

CONCLUSIONS AND FUTURE WORK

In this paper we have illustrated a methodology and a tool for acquiring explicit users' evaluations about the relevancy of contextual factors for item selection and recommendation. Contextual information is known to 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*, pages 217–253. Springer Verlag, 2010.

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

5. V. Bellotti, B. Begole, E. H. Chi, N. Ducheneaut, J. Fang, E. Isaacs, T. King, M. W. Newman, K. Partridge, B. Price, P. Rasmussen, M. Roberts, D. J. Schiano, and A. Walendowski. Activity-based serendipitous recommendations with the magitti mobile leisure guide. In *SIGCHI conference on Human factors in computing systems*, CHI '08, pages 1157–1166, New York, NY, USA, 2008. ACM.

6. 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.

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

8. 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.

9. 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.

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*, pages 735–767. Springer Verlag, 2010.

11. J. Swarbrooke and S. Horner. *Consumer Behaviour in Tourism, Second Edition*. Butterworth-Heinemann, 2006.

APPENDIX

Ranking of context factors by their influence on the user decision to visit an place of interest.

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