

The New Challenges when Modeling Context through Diversity over Time in Recommender Systems

Amaury l’Huillier, Sylvain Castagnos, Anne Boyer

► **To cite this version:**

Amaury l’Huillier, Sylvain Castagnos, Anne Boyer. The New Challenges when Modeling Context through Diversity over Time in Recommender Systems. Proceedings of the 24th Conference on User Modeling, Adaptation and Personalization (UMAP 2016 Doctoral Consortium), Jul 2016, Halifax, Canada. pp.341-344, 10.1145/2930238.2930370 . hal-01306790

HAL Id: hal-01306790

<https://hal.inria.fr/hal-01306790>

Submitted on 2 May 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L’archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d’enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The New Challenges when Modeling Context through Diversity over Time in Recommender Systems

Amaury L'Huillier
KIWI Team, LORIA -
Université de Lorraine
Campus Scientifique, B.P. 239
54506 Vandœuvre - France
amaury.lhuillier@loria.fr

Sylvain Castagnos
KIWI Team, LORIA -
Université de Lorraine
Campus Scientifique, B.P. 239
54506 Vandœuvre - France
sylvain.castagnos@loria.fr

Anne Boyer
KIWI Team, LORIA -
Université de Lorraine
Campus Scientifique, B.P. 239
54506 Vandœuvre - France
anne.boyer@loria.fr

ABSTRACT

The main goal of recommender systems is to help users to filter all the information available by suggesting items they may like without they had to find them by themselves. Although the rating prediction is a pretty well controlled topic, being able to make a recommendation at the right moment still remain a challenging task. To this end, most researches try to integrate contextual information (weather, mood, location of users, etc.) in the recommendation process. Even if this process increases users satisfaction, using personal information faces with users' privacy issues. In a different way, our approach is only giving credits to the evolution of diversity within the recent history of consultations, allowing us to automatically detect implicit contexts. In this paper, we will discuss the scientific challenges to be overcome to take maximum advantage of those implicit contexts in the recommendation process.

Keywords

Diversity, Context, Recommendation, User Privacy

1. INTRODUCTION AND MOTIVATIONS

Recommender systems (RS) have been studied for the last two decades. Although significant progresses were achieved to date, many challenges remain. We know from many user studies that proposing items to users solely based on the rating prediction is not enough [9]. Precision is not self sufficient to provide good recommendations and other human factors have to be integrated in the recommendation process. In our case, we will focus on the relationship between two human factors which are the context [5] and the diversity [10]. We will present why and how the notion of context has become an important part of the recommendation process and what are its limitations. Then, we will present the notion of diversity and explain the reasons why monitoring its evolution over time could be promising to overcome the limitations faced by the contextual recommender systems.

The goal of a context-aware recommender system (CARS) is to adapt the recommendation to the current characteristics of the user situation, also called his context. As an example, a recommendation should not be the same either a user is alone or with some friends even if the model of preferences used is the same in both cases. In their state-of-the-art on CARS, Adomavicius and Tuzhilin [1] present several ways to include contextual information (like the weather, the time of the day, or the user mood) into the recommendation process. All the user studies reported in their paper highlighted an improvement in term of user satisfaction, regardless the domain used (restaurant, music, movies). By integrating additional information about context, the recommendations proposed are closer to users' current needs comparatively to traditional RS. However, an important drawback of all CARS is their reliance on the collect and the exploitation of this contextual information making them intrusive. While we live in a time when user privacy is more and more debated, it seems primordial to focus on non-intrusive and privacy-preserving systems. Developing such a respectful recommender system which has identical performances (in term of user satisfaction) in comparison with the models relying on contextual factors is one of the long-term goal we intend to achieve.

Rather than using personal information to detect the context, we believe that an in-depth analysis of the evolution in the user consultation path could be a promising way to understand the user needs. In other words, to preserve privacy, we argue that we do not need to know the characteristics of the context. We should instead detect the break points within a sequence of consultations (corresponding to changes of context), and then exploit meta-data about items consulted within a context since content-based techniques are known to be less intrusive [4].

We will now present some examples of use of the diversity to explain why the dimension could be related to the context. Presenting diverse recommendations to users is very useful in some cases as it offers alternatives to users [8]. However, being able to automatically detect and adjust the diversity level over time still remains a difficult task. Even if such a model did not exist yet, the possibilities offered have already been discussed [2]. The authors explain in the case of an e-commerce website that the diversity level to use in the recommendations should be high in the beginning of a navigation session and should decrease gradually as the session ends. The main idea of our works is to extend this idea, by analysing the evolution of diversity over time in the path of

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

UMAP '16 July 13-17, 2016, Halifax, NS, Canada

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4370-1/16/07.

DOI: <http://dx.doi.org/10.1145/2930238.2930265>

the users to infer their needs by detecting ends of session, but also changes of context. Hence, it is not the awareness of characteristics of the user situation (begin or end of session) which is used to adapt recommendations and diversity level, it is the evolution of diversity itself which gives to the system the clues to make the most appropriate recommendations. Currently, we have developed a model to monitor in real time the evolution of diversity that we will present and discuss in the next part.

2. FROM DIVERSITY TO CONTEXT

As explain in the previous section, we have developed a model to monitor and exploit the evolution of diversity over time that we called DANCE [6].

2.1 General Principle

DANCE can be used to compute at each time step (each consultation) the diversity brought by an item (the target) according to the previous consulted items (the history). The size of the history used to compute the diversity is fixed by a parameter of our model. So, with the diversity value obtained, it is possible to quantify how an item is similar or different compared to the previous consulted items. All we need is the set of attributes for each type of item (webpages, music, books, ...) and the associated values. The diversity can be then computed for any type of attribute (numerical, binary, string, interval, ...) as long as it is possible to compute the similarity¹ existing between two values of this attributes. It is possible to model the evolution of diversity for one attribute in particular, for a subset of attributes, or for all the attributes available by averaging the values of diversity obtained. The full mathematical formalism of our model is available in [3].

2.2 Utility and Application Domains

DANCE can model the evolution of diversity over time and then it is possible to detect some important variations of diversity. Detecting those important variations is very useful as it means that the user is changing his interests, he does not want similar items anymore. Hence, the recommender system must quickly adapt its strategy in order to fit the recommendation to these new interests. In [7], we defined the notion of implicit context as the common characteristics shared by the consulted items during a certain time lapse. The notion of implicit context was created in opposition with explicit context, which is defined as a situation for which the contextual factors are known. An implicit context only refers to a sequence of items while an explicit context includes additional information about the users (ip address, gender, age, ...). In term of privacy, implicit contexts are clearly less intrusive, as all pieces of information are extracted from the items and not from the user. Then, our model DANCE is able to detect changes of context within users' sequences of consultations, and each subsequence of items consulted between two changes of context corresponds to an implicit context. We then aim to use these implicit contexts so as to provide relevant recommendations to users. We tested and validated the ability of our model to split sequences of consultations into implicit contexts with a large musical dataset (more than 200,000 consultations). We computed the diversity based on the musical attributes extracted

¹The similarity is the opposite value of the diversity.

from TheEchonest². We gather 7 artist-related attributes, and 6 song-related attributes. We have therefore developed diversity formulas according to the nature of those attributes: interval (band years of activities), numerical (duration, tempo), binary (mode), list of words (genre), location (coordinates of artist). We have demonstrated that there exists a strong overlap between implicit contexts and existing sessions [6]. This overlap confirms that we can give meaning to the changes of context detected by our model, and by extension that implicit contexts are good alternatives to explicit contexts. Indeed, when a user starts a new session, his explicit context changes (he was at the office and now he starts to listen some music at home). As most of changes of implicit contexts also correspond to ends of sessions, this supports that implicit contexts can be used to provide contextual recommendations. In addition to the overlap between implicit contexts and sessions, we have noticed that it could exist several implicit contexts inside a single session.

Despite the fact that we validated our approach with a musical dataset, DANCE is a generic model and can be used in many application domains. The only specific parts which have to be adapted are the diversity formulas per attribute which, in our case, were related to songs and artists.

2.3 Strengths and Limitations

Our model is highly reactive, as computing the diversity evolution and detecting changes of implicit context can be done in real time. The small size of the history used to do so both preserves user privacy as we do not retain all the user history, and ensures a linear complexity. As the diversity is computed by using several attributes, the model can easily be parameterized by adjusting the weight of each attribute according to the user characteristics. As an example, if the variation of an attribute seems important to a user, this attribute must have more impact in the value of the diversity computed. In the next section, we explain more precisely how such a parameterization could be carried out. In [3], we demonstrated that DANCE is robust up to 60% of missing data and provides good performances even if users are consulting different types of items. Those results support the genericity of our model and will allow us to deploy it in real conditions without significant modifications. Only the formulas used to compute diversity must be adapted to the attributes that characterize available items.

One limitation of our model is that it only fits to situations and application domains where we have contiguous sequences of consultations, with items consulted in relatively short time laps (at least few items must be consulted per hour). Music, news, e-commerce or professional social networks like LinkedIn³ or Yunque⁴ are domains for which our model can then be used. For all those domains, users consult many items and can provide feedbacks within just a few minutes (maybe less). Our model can then try to capture the current implicit context of users by analysing diversity evolution. Conversely, if there are not enough consultations or if the gaps of time between the consultations are too long, the model will fail at providing contextualized recommendations. As an example, users rarely watch several movies in a row.

²<http://the.echonest.com/>

³<https://www.linkedin.com/>

⁴<https://my.yunque.com/>

3. SCIENTIFIC CHALLENGES

As stated in the previous section, we proposed a model to detect changes of implicit context. Our long-term goal is to rely on this model to provide recommendations of items in accordance with the user context, while preserving his privacy. The following subsection will introduce the scientific challenges to address on the road to such a system.

3.1 Parameterization of the model

DANCE is a generic model and can be easily parameterized. The parameterization occurs in the computation of the diversity, and while defining the detection conditions of a change of context. As the diversity is computed on the basis of a set of attributes, it is possible to weight these attributes in order to give them more impact on the value of the global diversity (which is the value actually used to detect the changes of implicit context). In [7], we used a genetic algorithm to adjust some parameters of the model. Genetic algorithms are often used to find a solution to a problem for which state spaces are too large to be explored in a reasonable time. In our case, we aim at finding the most appropriate weight for each attribute, according to a fitness function. The latter allows us to specialize our model, and to train it to favour the detection of certain types of context changes (ends of session, skipped items, thematic changes, strong or slight variations, ...). The results showed that by adapting the weight of some attributes, the percentage of overlap between changes of implicit context and ends of session can be increased. In some cases, the distribution of the weights can give meaning to the changes of context. As an example, the tempo and the energy of songs were often highly weighted when applying our model to our music dataset, while the weight of the location of artist was low. The weights are computed for each user, since we all have different habits of consultations (even if they could also be computed for all the users put together). However, if some attributes are always lowly weighted for almost all users, it can indicate that these attributes are not relevant for the type of context changes we aim to detect or for this application domain. On the opposite, if they are always highly weighted, they are mandatory when modelling usages. We also proved that the parameterization of our model was stable because the configurations of the weights obtained were very similar from one execution of our genetic algorithm to another.

One difficulty remains for this parameterization process. We should provide a way to dynamically adapt the fitness function to match user's feedback and expectations all over their consultations so as to improve the performances of our model. This could lead the model to recompute the weights of attributes for each user periodically or at some specific moments/situations. We can also imagine that user behaviors varies from an explicit context to another. Recomputing weights will thus allow us to build a map of relations between implicit and explicit contexts (assuming that we have the information about the contextual factors, even if collected *a posteriori*).

3.2 Context-based recommendations

Currently, DANCE can be used to split a sequence of consultations into implicit contexts on the basis of the evolution of diversity. All the items contained in a same implicit context share common characteristics that we could use to pro-

vide recommendations. Previously, we have explained that using contextual information leads to better recommendations and we want to use our implicit contexts to achieve the same goal. We will now present the different required steps so as to use implicit context:

- detect in real time and as soon as possible the current context (CC). The faster the system identify the implicit context, the better the recommendations will fit to the context. For the moment, we validated our approach *a posteriori* on an offline dataset. In real situation, the model will have to analyse sequences while they are not over;
- use the history of the active user, or all the users to find all the implicit similar contexts to CC;
- mine those similar contexts to find the most relevant items to recommend. The goal here is to predict the user interest for the items found in the similar contexts. As the number of items to be recommended is limited, all the items found previously have to be ranked according to their relevance to CC. The relevance can be estimated from the history of the user. For instance, if a user usually consults several times some items, he might have few tolerance toward novelty and serendipity. Otherwise, we can find relevant items from other users' history. Analysing the way the attributes of the current context evolves could be another solution to rank the items. For example, in a context where the price of the consulted items is always within a fixed interval, selecting items that respect that condition could be an efficient way to find relevant items.

This process of recommendation offers a major advantage: only the items directly related to the active user's short-term needs are used in the recommendation process. Indeed, unlike traditional systems which use all the items of the current session to provide recommendations, our approach ensures to select relevant items.

DANCE could also be used to predict the end of an explicit context and even to anticipate the next one. The history of a user, or of all the users, can be used as a train set and several rules can be extracted and used to improve the recommendation. Pattern mining techniques or association rules algorithms could be used to increase the prediction of transitions between contexts.

3.3 Characterization of implicit contexts

Providing recommendations adapted to the current context is very important but not self sufficient. Whatever how well a recommendation fits to the context according to the system, if the user does not accept it, this recommendation is a failure. That is why an important part of the recommendation process is to ensure that the user will understand the reasons why the system proposes those items in particular. To maximize the acceptance rate of the recommendations based on the implicit contexts, we have to find ways to automatically characterize the current context. Such a characterization could be used to explain to a user both how the system perceives his current context and why the recommendations made are adapted to it.

Extracting the characteristics could be done by analysing the differences of diversity between attributes. As we currently do not know what is the best way to characterise an

implicit context, we aim at developing an automatic machine learning model able to detect and extract the representative characteristics of each implicit context. Such a model could determine the subset of attributes for which the evolution of diversity is remarkable: very low or very high variations, significant correlation between several attributes... The goal is not to find a unique characterisation for each context, which would not only make the calculations of similar contexts difficult, but would complicate the user explanation task. The system must find the best compromise between the mathematical identification of principal components within the context and the understanding/acceptance by users, as the optimal classification of the characteristics can be different from users' perception. As an example, there could exist some contexts for which the diversity of all the attributes is pretty low. Conversely, there can exist some contexts for which the user expects a high specialisation for some attributes and a high diversity for all the others. For instance, in a professional social network, a possible explanation of an implicit context detected in a consultation sequence of job offers could be: "According to your previous consultations we realized that you are looking for a job in the health informatics domain with a minimum wage of 2000 Euros per month but the location seems regardless to you. That is why we recommend you those items".

The characterization task we want to develop is close to the task performed by a critique-based recommender. Indeed, after each cycle of recommendation, a critique-based recommender tries to discover the right combinations of attributes in order to help the user to filter the items.

3.4 User perception of implicit contexts

Our model DANCE has been evaluated offline and our recommendation systems based on implicit contexts will be evaluated as well. However, conducting user studies and confronting our model to real users in a real situation of recommendation is necessary in order to evaluate the user satisfaction, the acceptance and the adoption rates [2].

In a first time, we want to measure the perception of the notion of implicit context. We would like to ask users to split contiguous sequences of consultation into shorter sequences and ask them to explain why they split them in this manner. The main idea is to compare the contexts determined by the users with the implicit contexts computed by our model. By comparing these two kinds of contexts, 4 possibilities could be considered:

- Users' contexts are the same, or closely similar in comparison to the implicit contexts;
- Users' contexts are partially the same compared to implicit contexts;
- Users' contexts are different from implicit contexts, but users would also agree with the implicit contexts;
- Users' contexts are different from implicit contexts and users do not understand them.

In the first case, the implicit contexts are perceived by users and this supports the relevance of the notion of implicit context. In the second case, since there is an overlap between the implicit contexts and the contexts given by the users, we can suppose that a parameterization of our model could reduce this difference. The explanations of the users

could be used to set the fitness function. The third case could mean that there are different relevant ways to split a sequence into implicit contexts. Furthermore, the two approaches may potentially be both relevant and could be used to provide recommendations. In the fourth case, it would be necessary to adapt the explanations of the recommendations as the user will not understand the implicit context used to provide the recommendations.

In a second time, we want to test our approach of recommendation according to the user satisfaction. To achieve this goal, we are developing an online music service in order to monitor the whole recommendation process. With such a tool, we will be able to collect user feedback with satisfaction survey, to test different parameterizations and ways to compute our implicit contexts (either by our current split process or by using the users' one), to make A/B testing to compare our method with several recommendation approaches.

4. CONCLUSION

As a conclusion, we had promising results to model users' context from the observation of diversity over time. We are expecting this approach to be the cornerstone of an all new privacy-preserving recommendation framework.

5. REFERENCES

- [1] G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. *Recommender Systems Handbook*, pages 217–253, 2011.
- [2] S. Castagnos, A. Brun, and A. Boyer. When diversity is needed... but not expected! In *IMMM*, pages 44–50, 2013.
- [3] S. Castagnos, A. L'huillier, and A. Boyer. Toward a Robust Diversity-Based Model to Detect Changes of Context. In *27th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2015)*, Vietri sul Mare, Italy, Nov. 2015.
- [4] L. F. Cranor. Hey, that's personal! In L. Ardissono, P. Bruna, and A. Mitrovic, editors, *User Modeling 2005*, volume 3538 of *Lecture Notes in Computer Science*, pages 4–4. Springer Berlin Heidelberg, 2005.
- [5] A. K. Dey, G. D. Abowd, and D. Salber. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Hum.-Comput. Interact.*, 16(2):97–166, Dec. 2001.
- [6] A. L'Huillier, S. Castagnos, and A. Boyer. Understanding Usages by Modeling Diversity over Time. volume 1181 of *UMAP 2014 Extended Proceedings*, Aalborg, Denmark, July 2014.
- [7] A. L'Huillier, S. Castagnos, and A. Boyer. Modéliser la diversité au cours du temps pour détecter le contexte dans un service de musique en ligne. *Revue des Sciences et Technologies de l'Information*, 2016.
- [8] L. McGinty and B. Smyth. On the role of diversity in conversational recommender systems. In *Proceedings of the Fifth International Conference on Case-Based Reasoning*, pages 276–290. Springer, 2003.
- [9] S. M. McNee, J. Riedl, and J. A. Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. pages 1097–1101, 2006.
- [10] B. Smyth and P. McClave. Similarity vs. diversity. *ICCB '01*, pages 347–361, London, UK, UK, 2001.