



XPLODIV: An exploitation-exploration aware diversification approach for Recommender Systems

Title	XPLODIV: An exploitation-exploration aware diversification approach for Recommender Systems
Author(s)	Barraza-Urbina, Andrea;Heitmann, Benjamin;Hayes, Conor;Carrillo-Ramos, Angela
Publication Date	2015-07
Publisher	AAAI Press

XPLODIV: An Exploitation-Exploration Aware Diversification Approach for Recommender Systems

Andrea Barraza-Urbina

Insight Centre for Data Analytics,
NUI Galway
andrea.barraza@insight-centre.org

Benjamin Heitmann

Insight Centre for Data Analytics,
NUI Galway
benjamin.heitmann@insight-centre.org

Conor Hayes

College of Engineering and
Informatics, NUI Galway
conor.hayes@insight-centre.org

Angela Carrillo-Ramos

Pontificia Universidad
Javeriana, Bogotá
angela.carrillo@javeriana.edu.co

Abstract

Recommender Systems (*RS*) have emerged to guide users in the task of efficiently browsing/exploring a large product space, helping users to quickly identify interesting products. However, suggestions generated with traditional *RS* usually do not produce diverse results though it has been argued that diversity is a desirable feature. The study of diversity-aware *RS* has become an important research challenge in recent years, drawing inspiration from diversification solutions for Information Retrieval (*IR*). However, we argue it is not enough to adapt *IR* techniques to *RS* as they do not place the necessary importance to factors such as serendipity, novelty and discovery which are imperative to *RS*. In this work, we propose a diversification technique for *RS* that generates a diversified list of results which not only balances the trade-off between quality (in terms of accuracy) and diversity, but also considers the trade-off between exploitation of the user profile and exploration of novel products. Our experimental evaluation shows that the proposed approach has comparable results to state of the art approaches. In addition, through control parameters, our approach can be tuned towards more explorative or exploitative recommendations.

Introduction

Recommendation Systems (*RS*) help users find in less time useful products. Traditionally, *RS* are assessed with accuracy metrics. But, accuracy alone is not a clear indicator of *RS* quality (McNee *et al.* 2006). Other characteristics such as novelty and diversity should also be evaluated. Nevertheless, diversity-awareness is a challenge for *RS*, which tend to offer users lists composed of similar items (*e.g.*, the user loves Star Trek so receives suggestions of only Star Trek related items (McNee *et al.* 2006)). A diversity-aware *RS* aims to reduce the redundancy in recommendation lists by offering users a range of options, not a homogeneous set of alternatives. With diversity, *RS*

can: (a) encourage product discovery by incentivizing users to explore unknown sections of the catalog, (b) cover a wider spectrum of user preferences, and (c) respond to ambiguous user preferences with a list of varied items, thus increasing the chance the user will like at least one item.

Traditionally, the diversification problem is defined as a bi-criterion optimization problem where the aim is to select k items from a broader set that maximizes both relevance and diversity. We introduce the notion that in *RS*, in addition to the diversity *vs.* relevance trade-off, there is also a trade-off between exploitation of the user profile and exploration of novel products. Because diversification techniques for *RS* are inspired from Information Retrieval (*IR*) techniques, exploration of novel products is ignored to focus solely on exploitation of the user profile. Nonetheless, exploration is essential to the notion of discovery which is the most important feature of a *RS*. To solve this problem, we propose a diversification technique that can be tuned towards either more explorative or more exploitative recommendations.

In this paper, we will first define the diversification problem for both *IR* and *RS*. We study *IR* diversification techniques given that they serve as inspiration for *RS* techniques. Next, a comparative analysis of related work is presented for both *IR* and *RS*. Subsequently, our exploitation-exploration diversification approach is introduced. After this, we analyze results from the experimental validation. Finally, we conclude and highlight future work.

The Diversification Problem

Diversity is a concept that has been applied in many fields; mostly with the goal of obtaining a set of objects that have a high level of *dissimilarity* between them, and that as a group, maximize a quality criterion. However, there is usually a trade-off between diversity and quality; hence, the *diversification problem* is how to choose k elements from a set that maximizes diversity at a low quality sacrifice.

In *IR*, diversity is a highly desirable feature that helps to: (i) remove redundancy from retrieved results and (ii) respond to query ambiguity by offering varied options, and in this manner, increase the chance of satisfying a user with a random intent (Zheng *et al.* 2012). The study of diversity as it has been applied in *IR* serves as a strong foundation for work on diversity in *RS*. In *IR*, the goal of diversification has been defined as selecting documents that are not only relevant to the target query but that also cover as many query interpretations or sub-topics as possible. However, there is a trade-off between selecting items that are of higher relevance (which tend to be similar to each other) and obtaining diverse results (Gollapudi *et al.* 2009). Therefore, the diversification problem in *IR* is usually modelled as a bi-criteria optimization problem that aims to find the appropriate balance between two competing objectives: diversity and relevance (Gollapudi *et al.* 2009).

In *RS*, diversity is also a highly desirable feature. On one hand, diversity is important to deal with the uncertainty surrounding the user profile. The only evidence of user tastes/likes a *RS* has is encapsulated within the user profile. However, much like a user query in *IR*, the user profile could be incomplete and ambiguous. This can be explained by: the large size of item spaces and the unfeasibility of obtaining explicit rating information on all products from users, the unreliability of interpreting implicit information to understand user likes, and the dynamic nature of user preferences. In face of user profile uncertainty, *RS* should offer users a diverse set of suggestions representative of the variety of the user's tastes so as to increase the chances the user finds useful items in recommendations (Vargas 2012)(Zheng *et al.* 2012).

On the other hand, diversity is essential to the concept of novelty, which is directly related to the idea of discovery and essential to the purpose of *RS*. We define that different levels of novelty can be achieved depending on how far or *diverse* an item is from the user's past experience. In addition to aiding user discovery, novel recommendations help increase the information flow between the user and the system. It is to be expected that discovering new products would lead to an information gain for the user, but this is also true for the *RS* itself. User discovery of new items leads to user feedback on diverse/novel items. This feedback generates larger information gain for the user profile than feedback of non-novel items, broadening the knowledge over the user preferences (Lemire *et al.* 2008).

Even though diversity is a desirable feature, *RS* do not offer diverse recommendations naturally. This is due to: (a) *the heuristics that lay foundation to RS techniques are based on similarity measures*: traditional techniques that are centered on similarity-based heuristics suffer problems like overspecialization, bias towards popular items, and bias towards items which are similar to highly-rated items from the user profile; (b) *traditional evaluation metrics encourage accuracy but penalize diversity*: with traditional *RS* techniques novel products tend to receive lower

predicted ratings compared to products similar to those the user always consumes, as a consequence, accuracy metrics penalize recommending novel products; and (c) *recommendation list evaluation is performed as an aggregate of the individual scores of items, disregarding the real value of items in the context of the list*: recommendation list metrics do not evaluate each product within the context of the list and cannot determine if the list offers items that are both of high quality and sufficiently diverse to cover the spectrum of the user's interests (McNee *et al.* 2006).

It can be seen that the diversification problem in *RS* is similar to that in *IR*, where there is a trade-off between the individual accuracy of an item and the overall diversity of the recommendation list. In this manner, the diversification goal in *RS* would be to generate a list of suggested items that maximize both the predicted rating for items and coverage over the wide spectrum of user preferences. However, *RS* must also account for novel products, which by definition are not directly related to the identified user preferences. This brings up an additional trade-off between how much the *RS* wants to *exploit* the known information about the user by covering the preferences in the user profile, and how much the *RS* wants to *explore* what other preferences the user could have by offering novel products.

The trade-off between exploitation and exploration could depend on many factors, such as the maturity of the user profile and the user's openness to experience. For example, for a new user, the *RS* might want to offer more novel/explorative products in order to gain information about the user's interests. In contrast, for users that are not very open to new experiences, they may possibly prefer to receive recommendations of products similar to those they have liked in the past (*i.e.*, exploitative items). As a result, it is important for the diversification technique in *RS* to be tunable, so in this way it can be adapted to the *RS* requirements of diversity, exploitation and exploration.

In synthesis, the diversification problem in *RS* is to choose k items from a broader set, that together balance the trade-off between relevance and diversity considering the trade-off between *exploitation* of the user profile and *exploration* of novel products. The additional trade-off (exploitation *vs.* exploration) marks an important difference between the diversification problem in *IR* and *RS*.

In this section, the diversification problem has been defined for *IR* and *RS*. In the following section we will offer a comparative analysis of current work for both *IR* and *RS*.

Related Work

Diversification approaches for both *RS* and *IR* can be classified as *implicit* or *explicit*.

In *IR*, implicit approaches infer that by selecting dissimilar documents the diverse query aspects will be indirectly covered. The method *MMR* (Carbonell *et al.* 1998) is a classic example that aims to maximize "*relevant novelty*": weighted linear combination

of *relevance* and *novelty* (novelty is defined as dissimilarity from previously selected documents). In contrast, explicit approaches directly attempt to cover different query aspects or sub-topics. *IA-Select* (Agrawal *et al.* 2009) and *xQuAD* (Santos *et al.* 2010) are examples of explicit approaches. In addition, (Zheng *et al.* 2012) propose three strategies to specify coverage functions of query sub-topics that serve as a basis for their diversification solution.

RS diversification techniques are frequently inspired by *IR* approaches. In this case, explicit approaches are those that attempt to cover as many of the user preferences as possible. As examples of implicit approaches, (Smyth *et al.* 2001) propose a number of objective functions based solely on relevance and diversity, and (Ziegler *et al.* 2005) propose the Topic Diversification technique. Both draw inspiration from *MMR*. In contrast, (Vargas 2012) proposes the concept of aspect-space as a mean to translate notions from *IR* to *RS*. As a result, explicit techniques, such as *xQuAD* and *IA-Select*, can be adapted for *RS*.

Alternatively, (Adomavicius *et al.* 2009) propose several re-ranking methods (e.g., ranking by item popularity) to increase diversity but maintain accuracy. Their approach is neither implicit nor explicit.

In *Table 1*, the reviewed approaches are compared with the following criteria: (a) *greedy optimization*: is the proposed solution a greedy optimization approach?, (b) *explicit approach*: does the proposed solution directly attempt to cover the diverse aspects of the query/user profile?, (c) *implicit approach*: does the proposed solution explicitly prevent redundancy within the results?, (d) *control of diversity vs. relevance trade-off*: is there a control parameter that can tune the diversity vs. relevance trade-off?, (e) *encourages discovery*: does the proposed approach not penalize novel/serendipitous items?, and (f) *control of exploitation vs. exploration trade-off*: is there a control parameter that can tune the exploitation vs. exploration trade-off?.

	1	2	3	4	5	6	7	8
Greedy Optimization	+	+	+	+	+	+	+	-
Explicit Approach	-	+	+	+	-	-	+	-
Implicit Approach	+	-	-	-	+	+	-	-
Control of diversity vs. relevance trade-off	+	-	+	+	+	+	?	?
Encourages Discovery	?	-	-	-	?	?	-	?
Control of exploitation vs. exploration trade-off	-	-	-	-	-	-	-	-
1 - (Carbonell <i>et al.</i> 1998)	4 - (Zheng <i>et al.</i> 2012)		7 - (Vargas 2012)					
2 - (Agrawal <i>et al.</i> 2009)	5 - (Smyth <i>et al.</i> 2001)		8 - (Adomavicius <i>et al.</i> 2009).					
3 - (Santos <i>et al.</i> 2010)	6 - (Ziegler <i>et al.</i> 2005)							

Table 1. Comparison of Related Works

From *Table 1* it can be concluded that: (i) most approaches are based on greedy optimization, which performs very well when the underlying objective function is sub-modular; (ii) none of the approaches explicitly considers not penalizing novel products: implicit

approaches would add novel products by chance and explicit approaches penalize adding novel products; finally but most importantly, (iii) none of the approaches consider the trade-off between *exploitation* and *exploration*. This can be explained because most approaches for *RS* tend to be an adaptation from approaches in *IR* and therefore share the same characteristics. In *IR*, encouraging discovery and exploration are not important factors.

In this section we have analyzed diversification techniques for both *IR* and *RS*. We found that current works for *RS* are pure adaptations of solutions from *IR*, and as such, miss important properties essential to *RS*. Even though *IR* approaches serve as a foundation for advances in *RS*, the ultimate goal of both fields is different. Consequently, it is not enough to accommodate *IR* ideas towards *RS*, they must also be augmented to reflect characteristics fundamental to *RS*, such as novelty and discovery. In view of this goal, we propose the Exploitation-Exploration Diversification Technique, which we will present in the following section.

Exploitation-Exploration Diversification

In this section we will introduce the Exploitation-Exploration Diversification technique named *XPLoDIV*, which is a post-filtering approach that receives as input candidate items generated from a traditional *RS* and selects a subset of diversified relevant items. We formulate our approach as a greedy optimization problem that aims to retrieve an ordered subset of items *R*, by iteratively adding to *R* one item i^* from the set of candidate items *C*, where i^* maximizes the function *XPLoDIV* – in Equation 1 – at a given iteration step, as shown in *Algorithm 1*. The technique receives as input the user profile *U* (i.e., set of items $u \in U$ that the user has rated), a set of candidate items *C* and the desired size for *R* defined as *k*. As output, *XPLoDIV* produces an ordered set of diversified items $R \subseteq C$, where *R* has *k* items (i.e., $|R| = k$) that are shown as the final recommendations.

Input: output size *k*, set of user profile items *U*, set of candidate items *C*
Output: set of diversified items *R*

- 1: $R \leftarrow \emptyset$
- 2: **while** $|R| < k \wedge C \neq \emptyset$ **do**
- 3: $i^* \leftarrow \arg \max_{i \in C \setminus R} XPLoDIV(i, U, R)$
- 4: $C \leftarrow C \setminus \{i^*\}$
- 5: $R \leftarrow R \cup \{i^*\}$
- 6: **end while**
- 7: **return** *R*

Algorithm 1. Greedy optimization *XPLoDIV*

$$XPLoDIV(i, U, R) = \alpha \cdot rel(i) + (1 - \alpha) \cdot div(i, R) \cdot XPLO(i, U)$$

$$XPLO(i, U) = \beta \cdot exploit(i, U) + (1 - \beta) \cdot xplore(i, U)$$

Equation 1. *XPLoDIV* optimization function

As has been argued, the goal of *RS* diversification is to balance the trade-off between *relevance* and *diversity*, considering the trade-off between *exploitation* of the user profile and *exploration* of novel products. To achieve this, *XPLoDIV* has four core dimensions: relevance $rel(i)$,

diversity $div(i, R)$, exploitation $xploit(i, U)$ and exploration $xplore(i, U)$. Moreover, the approach has two control parameters, α and β , which respectively tune the trade-offs between relevance vs. diversity and exploitation vs. exploration. In addition, the diversity of selected items in *XPLODIV* is directly linked to the *exploitation vs. exploration* trade-off. As a result, the approach can be set towards more diverse exploitative items or more diverse explorative items.

The remainder of this section will discuss in detail the dimensions that compose *XPLODIV*. Each of these must be normalized to return a value in the range $[0,1]$, where one is the highest desirable value.

Relevance Dimension

The relevance dimension gives priority to items that have high predicted rating. In this fashion, the relevance value of an item i is given by *Equation 2*. In this equation, *MaxRating* is the maximum possible rating that a user can give to an item and *predicted_rating(i)* is the predicted rating for item i obtained from a traditional *RS*.

$$rel(i) = \frac{predicted_rating(i)}{MaxRating}$$

Equation 2. Relevance Dimension

Diversity Dimension

The diversity dimension $div(i, R)$ measures how diverse an item i is in relation to a set of items R . This section highlights approaches to determine $div(i, R)$.

Weitzman (Weitzman *et al.* 1992) proposes that the diversity of an element in relation to a set can be measured in terms of the amount of diversity gained for the set if the element is added. With this formulation, the diversity of an item to a set can be measured founded on popular measures of set diversity such as the Gini-Simpson index (Stirling 2007), Stirling diversity metric (Stirling 2007), among others. An example can be viewed in *Equation 3*.

$$div(i, R) = \frac{diversity(R \cup \{i\}) - diversity(R)}{diversity(R)} = \frac{GiniSimpson(R \cup \{i\}) - GiniSimpson(R)}{GiniSimpson(R)}$$

Equation 3. Diversity dimension as diversity gain

Alternatively, Weitzman (Weitzman *et al.* 1992) argues that the diversity an item would add to a set can be measured as the distance of the item to the set. Weitzman proposes the minimum distance to be one measure of item-to-set diversity, as in *Equation 4*, where distance between two items can be measured as the inverse of their similarity.

$$div(i, R) = \min_{r \in R} d(r, i)$$

$$d(r, i) = 1 - similarity(r, i)$$

Equation 4. Diversity dimension as minimum distance

Another way to measure item-to-set distance can be the average pairwise distance of the item i to each of the items in the set, as in *Equation 5*.

$$div(i, R) = \frac{1}{|R|} \sum_{r \in R} d(r, i)$$

Equation 5. Diversity dimension as average distance

Exploitation Dimension

The exploitation dimension reinforces those items that exploit known user preferences. Items that represent previously identified user preferences could turn out to be promising recommendations, following the content-based *RS* heuristic that assumes users will continue to have the same preferences they have had in the past. To achieve this, the exploitation dimension aims to determine how representative item i is of the user's preferences found in the user profile U . To determine the exploitation value of an item we propose to measure the probability that similar items within U have a high rating, as in *Equation 6*. In *Equation 6* the function $rating(u, U)$ returns the rating the user assigned to the item u . An extension to this approach is to determine the probability that the nearest neighbors of the item i from the user profile, have a high rating.

$$xploit(i, U) = \frac{\sum_{u \in U} sim(u, r) \cdot rating(u, U)}{\sum_{u \in U} rating(u, U)}$$

Equation 6. Exploitation dimension

Exploration Dimension

The exploration dimension reinforces those items that incentivize the user to explore the unknown. In other words, this dimension gives priority to novel/serendipitous items that are outside of the user's past preferences. Given that the user profile can be ambiguous and incomplete, it is not smart to always exploit known information and possibly stay stuck in a sub-optimal item space. By offering user's novel products, the *RS* is also attempting to retrieve information on unknown user preferences, hence preventing overspecialization and encouraging discovery.

We determine that the novelty an item offers to a user can be measured by how diverse an item is from the user's past experiences. As a result, we can use one of the specified measures for the diversity dimension; but instead of measuring the diversity of the item in relation to the set of selected items, we measure the diversity of item i in relation to the user's past experiences. The most clear indication of the user's past experiences is encapsulated within the user profile U . Thus, exploitation can be measured as $div(i, U)$.

Another way to determine a user's past experiences could be by considering the experiences of similar users; assuming similar users have similar experiences. For example, if an item is well known or popular among users that are similar to a target user, it is probable that the target user already knows about this product even though he/she has not rated it. In this case, novelty is measured with respect to the neighborhood of similar users, in addition to the user profile, as in *Equation 7*. In this equation, the profiles of the k nearest neighbors of the user are aggregated in to a set N .

Then, the exploration dimension is determined as the diversity of item i to the set of items formed by the union of the user profile and the profiles of the nearest neighbors.

$$N = \{N_1 \cup N_2 \dots \cup N_k\}$$

$$xplore(i, \{N \cup U\}) = div(i, \{N \cup U\})$$

Equation 7. Neighborhood exploration dimension

In this section we have specified the Exploitation-Exploration Diversification technique *XPLODIV*. This technique not only considers the *relevance vs. diversity* trade-off but also allows control over the *exploitation vs. exploration* trade-off. In the following section, we present experimental validation of the proposed technique.

Experimental Validation

We evaluate our approach from four different perspectives showing that it not only provides comparable results to state of the art techniques, but in addition, it can be tuned towards more exploitative diversity or more explorative diversity. The analyzed perspectives measure the relevance, diversity, exploitation and exploration aspects of obtained results.

We use the MovieLens 100k¹ dataset and a traditional user-user collaborative filtering *RS* (with neighborhood size of 50) to produce candidate items which serve as input for all the diversification techniques under evaluation. Candidate items are ordered according to relevance (in terms of predicted rating). Our approach is compared to the following baselines and state of the art technique: (a) *no diversity*: top k of candidate items, (b) *random diversity*: random selection of k items from candidate items, and (c) *MMR* (Carbonell *et al.* 1998)($\alpha = 0.5$): representative technique of implicit *RS* diversification, *MMR* has served as foundation for many related more recent approaches. Explicit diversification approaches are purposely omitted as by definition are biased towards only exploitative items ignoring novel items. For all experiments, an ordered set of diversified items of size fifteen ($k = 15$) was selected from a set of candidate items of size 100.

Results are evaluated using metrics defined to analyze each of the four perspectives. These metrics are defined in Table 2 as follows: (a) *relevance metric*: the normalized discounted cumulative gain (*i.e.* *nDCG*) is used, where the ideal *DCG* is obtained from the ordered top k results from the candidate items, (b) *diversity metric*: the diversity of a list is measured as the pairwise inter-list dissimilarity (*i.e.*, *PILD*) of its elements, (c) *exploitation metric*: the user profile exploitation (*i.e.*, *UPE*) metric is defined as the average of how well represented each item from the user profile U is by items in the set R (it is determined that each $u \in U$ is represented by the item in R that is most similar to the item u), and (d) *exploration metric*: percentage of novel items within the set R , where novel items are those that have

a dissimilarity from the user profile U larger or equal to the threshold τ . Through experimental observations we found that for the MovieLens dataset the value $\tau = 0.9$ was sufficiently large to observe the exploration aspect and to omit more exploitative items. Each metric returns a value in the range $[0, 1]$, where one is the highest desirable value.

Relevance Metric	
$nDCG(R) = \frac{DCG(R)}{DCG(TopK(C))}$	$DCG(R) = rel(R_1) + \sum_{i=2}^{ R } \frac{rel(R_i)}{\log_2 i}$
Diversity Metric	
$PILD(R) = \frac{2}{ R (R - 1)} \left(\sum_{\substack{i=1 \\ r_i \in R}}^{ R -1} \sum_{\substack{j=i+1, j \neq i \\ r_j \in R}}^{ R } d(r_i, r_j) \right)$	
Exploitation Metric	
$UPE(R, U) = \frac{1}{ U } \cdot \sum_{u \in U} \max_{r \in R} sim(r, u)$	
Exploration Metric	
$NDT(R, U, \tau) = \frac{1}{ R } \cdot \sum_{r \in R} DT(r, U, \tau)$	
$DT(r, U, \tau) = \begin{cases} 1, & UPd(r, U) \geq \tau \\ 0, & otherwise \end{cases}$ $UPd(r, U) = \frac{1}{ U } \cdot \sum_{u \in U} d(r, u)$	
Where:	
R_i	Item in position i within the ordered set R .
$sim(r, u)$	Similarity of item r to item u . We used Jaccard similarity coefficient between movie genres.
$d(r, u)$	Distance of item r to item u . $d(r, u) = 1 - sim(r, u)$

Table 2. Evaluation Metrics

For *XPLODIV* we used the following set-up: diversity as Equation 4, exploitation as Equation 6 and exploration as $div(i, U)$. In addition, we ran experiments for different values of the *XPLODIV* control parameters as follows: (a) *relevance bias*: $\alpha = 0.8, \beta = 0.5$; (b) *exploitation bias*: $\alpha = 0.2, \beta = 0.7$; (c) *exploration bias*: $\alpha = 0.2, \beta = 0.3$; (d) *no bias*: $\alpha = 0.5, \beta = 0.5$; (e) *pure exploitation*: $\alpha = 0.0, \beta = 1.0$; and (f) *pure exploration*: $\alpha = 0.0, \beta = 0.0$. A detailed description of the evaluation set-up can be viewed in (Barraza 2014).

Experiment results are shown in Figure 1. In addition, Figure 2 illustrates, for selected approaches, the percentage of gain or loss for each perspective relative to results obtained from the *no diversity* approach. From Figure 2, we can make observations such as: The largest exploration gain was obtained by “*pure exploration*” with exploration gain of 113.65% over “*no diversity*” results.

Specifically, from Figure 2 we can observe that all diversification approaches present a relevance loss with respect to results without diversity. However, most approaches present a significant diversity gain given the relevance sacrifice. We can also observe that our approach

¹ MovieLens Datasets: <http://grouplens.org/datasets/movielens/>

can be tuned towards more explorative results given the exploration gain in the “*exploration bias*” and “*pure exploration*” versions of *XPLODIV*. Similarly, we can observe an exploitation gain in the “*pure exploitation*” *XPLODIV*. However, there is an exploitation loss in the “*exploitation bias*” configuration which could be due to a possible indirect trade-off between exploitation and relevance which can be seen in *Figure 1*. All in all, experimental validation shows that adding diversity using *XPLODIV* allows control over the trade-offs between exploitation vs. exploration and diversity vs. relevance at minimum relevance loss. A deeper analysis of results can be viewed in (Barraza 2014).

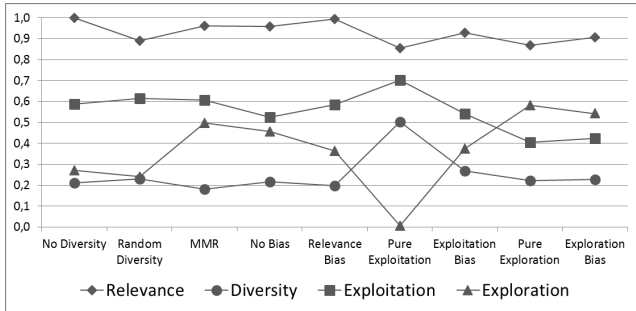


Figure 1. Evaluation Results




Figure 2. Evaluation Results relative to No Diversity

Conclusion and Future Work

We have presented the Exploitation-Exploration diversification technique *XPLODIV*, which not only considers the trade-off between diversity vs. relevance but also the trade-off between exploitation of the user profile and exploration of novel products. Through experimental validation we have demonstrated that the approach can be tuned towards more exploitative diverse results or more explorative diverse results with controlled sacrifice over relevance. As future work, we plan to define mechanisms to dynamically learn the values for the control parameters and in this way adapt *XPLODIV* to different user profile and dataset characteristics. Also, we plan to carry out further

experimentation to observe the performance of different *XPLODIV* set-ups in different contexts. Furthermore, we wish to see how we can use *XPLODIV* in other features of *RS*. For example, we could adapt *XPLODIV* to select a diverse set of neighbors in a collaborative filtering *RS*. In this case, instead of simply selecting the most similar neighbors, a diverse set of neighbors would be chosen according to their explorative or exploitative values. Finally, we would like to explore the effect of using our diversification approach to aggregate results produced from different *RS* algorithms, and in this way, enhance hybrid recommenders.

Acknowledgements: This research was made possible by funding from Science Foundation Ireland under grant number SFI/12/RC/2289 (Insight) and by the Master's Program of the Computer Science Department at the Pontificia Universidad Javeriana, Bogotá.

References

- Adomavicius, G.; Kwon, Y. 2009. Toward more diverse recommendations: Item re-ranking methods for recommender systems. In *Proc. WITS'09*.
- Agrawal, R.; Gollapudi, S.; Halverson, A.; Ieong, S. 2009. Diversifying Search Results. In *Proc. WSDM'09*, 5–14. NY, USA: ACM.
- Barraza, A. 2014. Diversification Technique for Recommender Systems that Controls the trade-off between Exploitation of the User Profile and Exploration of Novel Products. MSc. diss., Facultad de Ingeniería, Pontificia Universidad Javeriana, Colombia. Available at: [http://pegasus.javeriana.edu.co/~PI131-01-XploDiv/entregables/\[TG AndreaBarraza\]-Memorias%20Integrada.pdf](http://pegasus.javeriana.edu.co/~PI131-01-XploDiv/entregables/[TG AndreaBarraza]-Memorias%20Integrada.pdf)
- Carbonell, J.; Goldstein, J. 1998. The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries. In *Proc. of SIGIR'98*, 335-336. NY, USA: ACM.
- Gollapudi, S.; Sharma, A. 2009. An Axiomatic Approach for Result Diversification. In *Proc. of WWW'09*, 381–390. NY, USA: ACM.
- Lemire, D.; Downes, S.; Paquet, S. 2008. Diversity in open social networks. Technical report, University of Quebec, Montreal.
- McNee, S.M.; Riedl, J.; Konstan, J.A. 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *Proc. of CHI'06*, 1097-1101. NY, USA: ACM.
- Santos, R. L. T.; Macdonald, C.; Ounis, I. 2010. Exploiting Query Reformulations for Web Search Result Diversification. In *Proc. WWW'10*, 881–890. NY, USA: ACM.
- Smyth, B.; McClave, P. 2001. Similarity vs. Diversity. In *Proc. ICCBR'01*, 347–361. Berlin-Heidelberg: Springer.
- Stirling, A. 2007. A General Framework for Analysing Diversity in Science, Technology and Society. *J. R. Soc. Interface* 4(15):707-719.
- Vargas, S. 2012. Novelty and Diversity Enhancement and Evaluation in Recommender Systems. MSc. diss., Department of Ingeniería Informática, Universidad Autónoma de Madrid, Spain.
- Weitzman, M. L. 1992. On Diversity. *The Quarterly Journal of Economics* 107(2): 363–405.
- Zheng, W.; Wang, X.; Fang, H.; Cheng, H. 2012. Coverage-based search result diversification. *Information Retrieval* 15(5): 433-457.
- Ziegler, C.N.; McNee, S.M.; Konstan, J.A.; Lausen, G. 2005. Improving recommendation lists through topic diversification. In *Proc. WWW'05*, 22-32. NY, USA: ACM.