

Novelty and Diversity Enhancement and Evaluation in Recommender Systems and Information Retrieval

Saúl Vargas
Universidad Autónoma de Madrid, Spain
saul.vargas@uam.es

Categories and subject descriptors: H3.3 [Information Search & Retrieval]: Information Filtering

Keywords: Recommender Systems; Novelty; Diversity

ABSTRACT

The development and evaluation of Information Retrieval and Recommender Systems has traditionally focused on the relevance and accuracy of retrieved documents and recommendations, respectively. However, there is an increasing realization that accuracy alone might be a sub-optimal strategy for a successful user experience [7, 3]. Properties such as novelty and diversity have been explored in both fields for assessing and enhancing the usefulness of search results and recommendations. In this doctoral research we study the assessment and enhancement of both properties in the confluence of Information Retrieval and Recommender Systems.

In Information Retrieval, diversity is posited a quality of result lists that helps cope with the ambiguity and under-specification in users queries, whereas novelty is understood as the quality of a system that avoids redundancy. In the last years there have been proposals and significant advances for both assessing [2] and enhancing [3] novelty and diversity in search results. Recommender Systems can be seen as a particular case of personalized Information Retrieval where there is no explicit query, but just implicit information about the user's interests. Researchers have realized that improving recommendations' usefulness and user satisfaction may require more than being accurate. Recommending novel items helps avoiding recommending too obvious or popular items [1]. Moreover, the effectiveness of recommendations can be enhanced by acknowledging the user's diverse tastes [7].

The novelty and diversity problem has nonetheless been approached under different views and formulations in both fields, giving rise to different models, methodologies, and metrics, with little convergence between both fields in this particular area. Our research addresses the problem of modeling and enhancing novelty and diversity in the context of

recommender systems, seeking connections to diversity research in Information Retrieval.

First, we propose the definition of a new framework of novelty and diversity metrics for Recommender Systems [4]. This framework adapts and unifies several metrics from the state of the art under a common probabilistic formulation. This formulation allows us to consider rank and relevance when assessing the novelty and diversity of recommendation lists, which has not been considered previously.

Second, we research the connection of recommendation diversity to the diversity principles and techniques developed in Information Retrieval for diversity [5]. We explore diversification techniques for recommender systems based on feature spaces and the extraction of sub-profiles, establishing an analogy with query interpretations or aspects for diversity in Information Retrieval. We study the relationship between the choice of feature spaces and their effectiveness in diversification as well.

Furthermore, the thesis delves into the definition of new methods to both assess and enhance novelty and diversity in recommender systems. For instance, we analyze the formalization of diversity and the explicit role of relevance under a probabilistic perspective [6]. We also consider the role of coverage, redundancy and size-awareness in diversity for recommender systems. Finally, we also study the popularity bias observed in many recommendation algorithms and its effect in terms of novelty and sales diversity, proposing ways to alleviate it.

We conduct empirical validation and evaluation of metrics and methods in offline experiments with publicly available datasets. Furthermore, we envision to incorporate online and crowdsourced experiments to gain further insights on the effects of novel and diverse recommendations.

1. REFERENCES

- [1] O. Ceña and P. Herrera. A new approach to evaluating novel recommendations. *RecSys*, 2008, pp. 1279–186
- [2] C. L. Clarke, M. Kolla, G. V. Cormack, O. Vechtomova, A. Ashkan, S. Büttcher, and I. MacKinnon. Novelty and diversity in information retrieval evaluation. *SIGIR*, 2008, pp. 659–666.
- [3] R. L. Santos, C. Macdonald, and I. Ounis. Exploiting query reformulations for web search result diversification. *WWW*, 2010, pp. 881–890.
- [4] S. Vargas and P. Castells. Rank and relevance in novelty and diversity metrics for recommender systems. *RecSys*, 2011, pp. 109–116.
- [5] S. Vargas, P. Castells, and D. Vallet. Intent-oriented diversity in recommender systems. *RecSys*, 2011, pp. 1211–1212.
- [6] S. Vargas, P. Castells, and D. Vallet. Explicit relevance models in intent-oriented information retrieval diversification. *SIGIR*, 2012, pp. 75–84.
- [7] C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. *WWW*, 2005, pp. 22–32.

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). Copyright is held by the author/owner(s).

SIGIR'14, July 6–11, 2014, Gold Coast, Queensland, Australia.

ACM 978-1-4503-2257-7/14/07.

<http://dx.doi.org/10.1145/2600428.2610382>.