

Proceedings of the

**Workshop on Novelty and Diversity
in Recommender Systems (DiveRS 2011)**

held at the

**5th ACM International Conference
on Recommender Systems (RecSys 2011)**

23 October 2011

Chicago, Illinois, USA

Edited by

Pablo Castells¹, Jun Wang², Rubén Lara³, Dell Zhang⁴

¹ Universidad Autónoma de Madrid, Spain

² University College London, UK

³ Telefónica, Investigación y Desarrollo, Spain

⁴ Birkbeck, University of London, UK

Preface

Introduction

Most research and development efforts in the Recommender Systems field have been focused on accuracy in predicting and matching user interests. However there is a growing realization that there is more than accuracy to the practical effectiveness and added-value of recommendation. In particular, novelty and diversity have been identified as key dimensions of recommendation utility in real scenarios, and a fundamental research direction to keep making progress in the field.

Novelty is indeed essential to recommendation: in many, if not most scenarios, the whole point of recommendation is inherently linked to a notion of discovery, as recommendation makes most sense when it exposes the user to a relevant experience that she would not have found, or thought of by herself –obvious, however accurate recommendations are generally of little use.

Not only does a varied recommendation provide in itself for a richer user experience. Given the inherent uncertainty in user interest prediction –since it is based on implicit, incomplete evidence of interests, where the latter are moreover subject to change–, avoiding a too narrow array of choice is generally a good approach to enhance the chances that the user is pleased by at least some recommended item. Sales diversity may enhance businesses as well, leveraging revenues from market niches.

It is easy to increase novelty and diversity by giving up on accuracy; the challenge is to enhance these aspects while still achieving a fair match of the user’s interests. The goal is thus generally to enhance the balance in this trade-off, rather than just a diversity or novelty increase.

Research contributions to this area have addressed the enhancement, evaluation, and understanding of novelty and diversity in recommendation. Businesses are accounting for these aspects in ad-hoc ways when engineering recommendation functionalities, and researchers have started to seek principled foundations for incorporating novelty and diversity in the recommendation models, algorithms, theories, and evaluation methodologies. But large room remains for further research, which motivates the DiveRS 2011 Workshop.

The 1st ACM RecSys 2011 International Workshop on Novelty and Diversity in Recommender Systems (DiveRS 2011) gathered researchers and practitioners interested in the role of novelty and diversity in recommender systems. The workshop was motivated by the importance of these topics in the field, both in practical terms, for their relevance in the development of recommender systems applications and comprehending real user needs, and for their fundamental implications in recommender systems theory and evaluation methodologies. Novelty and diversity can thus be identified as a rich area for long-haul research. The area has started to be addressed only relatively recently in the field though, and the time thus seems appropriate for an open exchange of ideas, discussion, and reflection in an informal forum.

The workshop sought to advance towards a better understanding of what novelty and diversity are, how they can improve the effectiveness of recommendation methods and the utility of their outputs. DiveRS 2011 pursued the identification of open problems, specific gaps, relevant research directions, and opportunities for innovation in the recommendation business. The workshop sought the formation of common ground and shared perspectives to foster progress in this area.

Scope and topics

Specific topics of interest for the workshop included, among others, the following:

- Modeling novelty and diversity in recommender systems.
 - Theoretical foundation for novelty and diversity.
 - Recommendation novelty and diversity models.
 - Popularity, risk, surprisal, serendipity, freshness, discovery.
 - Link to diversity models in Information Retrieval.
- Novelty and diversity enhancement.
 - Diversification methods.
 - Recommendation of long-tail and difficult items, cold-start.
 - Individual vs. global diversity.
 - Machine Learning for novelty and diversity.
- Novelty and diversity across recommendations.
 - Novelty and diversity in sequential recommendation.
 - Novelty and diversity in interactive recommendation.
 - Aggregate diversity.
 - Novelty and diversity in time and context.
 - Novelty and trust.
- Novelty and diversity evaluation.
 - Experimental methodologies and design.
 - Novelty and diversity metrics.
 - Datasets.
 - User studies.
- Business perspective on novelty and diversity.

The following questions were, among others, raised and addressed by the workshop:

- What are the different notions and dimensions of novelty and diversity? Is it possible to establish a clear definition and/or taxonomy?
- How are novelty and diversity themselves different and related?
- How can diversity, novelty, and accuracy be enhanced together?
- What important differences arise between the end-user point of view and the system or the business perspective?
- How can novelty and diversity be measured and evaluated?
- What are the potential implications of novelty on user trust, and how can they be properly cared for?
- What are the differences and unexplored connections between diversity as researched in Recommender Systems and Information Retrieval?
- Is there a relevant relation between novelty, diversity and context in recommendation?
- To what extent are novelty and diversity procured by or missing from state-of-the-art technologies?

- Do the different state-of-the-art recommendation algorithms (content-based, nearest-neighbors, matrix factorization, social, hybrid, ensembles, etc.) perform differently to each other in terms of novelty and diversity?
- To what extent are novelty and diversity a concern in the development of real-world recommender system applications, and how are they being addressed? What is the business value in novelty and diversity enhancement?
- What are the scenarios where novelty and diversity are most/least valuable or necessary? Are there situations in which novelty and diversity are not a desirable feature?

Submissions and program

The workshop received 13 submissions, of which 7 were accepted (54%). The first three papers –in the order included in the present proceedings– were selected for long presentation, the next four having a slightly shorter slot in the workshop schedule. The workshop opened with a keynote talk by Neil Hurley, and included an open discussion after the paper presentations. We briefly summarize here the presented works and held discussions.

The keynote talk, entitled “Towards Diverse Recommendation”, provided an overview of the area, its development and main proposals over the past decade, and current perspectives. The papers presented after this covered a wide spectrum of topics, encompassing most of the aspects put forward in the intended workshop scope. The first three papers present approaches to enhance recommendation diversity and novelty, introducing or revising metrics to capture specific aspects of these dimensions. G. Adomavicius and Y. Kwon present a graph-based approach to enhance the global diversity of recommendation, understood as the total set of distinct items that are recommended to the set of all users as a whole. P. Adamopoulos and A. Tuzhilin revise and formalize the notion of unexpectedness as a particular case of user-specific novelty, and propose a method to maximize it. Also in the scope of novelty, K. Oku and F. Hattori propose a method to produce serendipitous recommendations by combining the features of different items of interest.

F. Mourão et al introduce a new angle on novelty by considering the effect that the passing of time may have on known items, which may regain some of their novelty value as past user experience fades away and is to some degree forgotten by users. They explore the positive effect that exploiting such oblivion processes may have on the diversity of recommendation. J. Golbeck and D. L. Hansen explore a new view on set-oriented recommendation by explicitly considering the recommendation of collections of items, where diversity arises as a natural quality dimension. R. Hu and P. Pu address the workshop theme from the point of view of real users and their perceptions. Their paper presents a user study of how user interface aspects influence the practical effectiveness of recommendation diversity and overall user satisfaction. Finally, S. Santini and P. Castells propose new formulations of novelty and diversity models based on fuzzy relevance, as an alternative to probabilistic formalizations based on binary relevance.

Different specific notions of novelty and diversity, distinctions and nuances between them, were identified along the presentations and discussion. The contextual nature of these dimensions –and the need for context-awareness in tackling them– were also underlined: novelty and diversity are relative to users, systems, time, viewpoint (e.g. user vs. business), tasks, session state, and other contextual variables. An issue that received particular attention in the discussion was the elucidation of when, to what extent, and in which scenarios, novelty and diversity are really appropriate in practice, from the understanding that their use should not be indiscriminate. This was a starting point for the open discussion, from which the session progressed towards further workshop topics.

While the usefulness of diversity and novelty is obvious in –or actually inherent to– many well-known applications, examples were mentioned of recommendation functionalities in commercial systems in which novelty and diversity seem to be disregarded, hinting that perhaps navigational recommendation or even the recommendation of known items might be useful in some contexts. Two points of view were distinguished to this respect: users and businesses. Regarding the latter, monetization was pointed out as a main effectiveness metric for commercial applications of recommendation technologies. It was noted to this respect that assessing the business value of novelty and diversity should require a distinction between short vs. longer-term –and direct vs. indirect– benefits. It was also noted that current commercial recommender technologies, such the ones used by Netflix, include novelty and diversity as features in some among their wide array of recommendation algorithms. There was general agreement that business studies in this area would be highly useful in shedding further light on these issues.

Regarding the end-user side, there was also a general call for user studies in order to properly understand and drive the introduction of novelty and diversity dimensions, as well as their precise need. The contribution by R. Hu and P. Pu was highlighted as an example of the studies that would be useful to this respect. User personality and attitude were indicated in this context as key aspects that should be taken into account when procuring novelty and diversity, since the attitude towards new and/or diverse experience varies considerably among users.

Conclusion

The contributions, presentations and discussions held at the workshop provided a good overview of the current progress in the area, where we stand today, and where further work is needed. The importance of the workshop theme was underlined, beyond the DiveRS workshop itself, by the recurrent references to novelty and diversity in the main RecSys conference track. The need for further work and discussion in this area was clear, as well as the interest for future initiatives in line with the present workshop. To this respect, the organizers announced the forthcoming publication of a special issue of the ACM Transactions on Intelligent Systems and Technology in the scope of the workshop.

Acknowledgments

The organizers would like to thank the Program Committee members for their high-quality and timely evaluation of the submissions; the RecSys 2011 organizers and workshop chairs for their support in the organization of this workshop; the keynote speaker, all the authors and presenters, for their contribution to a high-quality workshop program; and all participants for such fruitful discussions and valuable ideas as were exchanged during the workshop. Thanks are due to all such contributions which made DiveRS 2011 a successful venue.

Pablo Castells

Jun Wang

Rubén Lara

Dell Zhang

Organizing Committee

Pablo Castells	Universidad Autónoma de Madrid, Spain
Jun Wang	University College London, UK
Rubén Lara	Telefónica, Investigación y Desarrollo, Spain
Dell Zhang	Birkbeck, University of London, UK

Program Committee

Xavier Amatriain	Netflix, USA
Leif Azzopardi	University of Glasgow, UK
Iván Cantador	Universidad Autónoma de Madrid, Spain
Licia Capra	University College London, UK
Òscar Celma	Gracenote, USA
Charles Clarke	University of Waterloo, Canada
Sreenivas Gollapudi	Microsoft Research, USA
Neil Hurley	University College Dublin, Ireland
Oren Kurland	Technion, Israel
Neal Lathia	University College London, UK
Hao Ma	Microsoft Research, USA
Qiaozhu Mei	University of Michigan, USA
Jérôme Picault	Bell Labs, Alcatel-Lucent, France
Filip Radlinski	Microsoft, Canada
Davood Rafiei	University of Alberta, Canada
Francesco Ricci	Free University of Bozen-Bolzano, Italy
David Vallet	Universidad Autónoma de Madrid, Spain
Paulo Villegas	Telefónica, Investigación y Desarrollo, Spain
ChengXiang Zhai	University of Illinois at Urbana-Champaign, USA
Tao Zhou	University of Electronic Science and Technology of China, China
Jianhan Zhu	True Knowledge, UK

Table of Contents

Keynote talk

Neil J. Hurley <i>Towards Diverse Recommendation</i>	1
---	---

Papers

Gediminas Adomavicius and Youngok Kwon <i>Maximizing Aggregate Recommendation Diversity: A Graph-Theoretic Approach</i>	3
Panagiotis Adamopoulos and Alexander Tuzhilin <i>On Unexpectedness in Recommender Systems: Or How to Expect the Unexpected</i>	11
Kenta Oku and Fumio Hattori <i>Fusion-based Recommender System for Improving Serendipity</i>	19
Fernando Mourão, Claudiane Fonseca, Camila Araújo and Wagner Meira Jr. <i>The Oblivion Problem: Exploiting Forgotten Items to Improve Recommendation Diversity</i>	27
Jennifer Golbeck and Derek L. Hansen <i>A Framework for Recommending Collections</i>	35
Rong Hu and Pearl Pu <i>Helping Users Perceive Recommendation Diversity</i>	43
Simone Santini and Pablo Castells <i>An Evaluation of Novelty and Diversity Based on Fuzzy Logic</i>	51