The Recommender Problem Revisited

[Morning Tutorial]

Xavier Amatriain
Netflix
100 Winchester Cr.
Los Gatos, CA, USA
xavier@netflix.com

Bamshad Mobasher
De Paul University
243 South Wabash Avenue
Chicago, IL, USA
mobasher@cs.depaul.edu

ABSTRACT

In 2006, Netflix announced a $1M prize competition to advance recommendation algorithms. The recommendation problem was simplified as the accuracy in predicting a user rating measured by the Root Mean Squared Error. While that formulation helped get the attention of the research community in the area, it may have put an excessive focus on what is simply one of possible approaches to recommendations. In this tutorial we will describe different components of modern recommender systems such as: personalized ranking, similarity, explanations, context-awareness, or search as recommendation. In the first part, we will use the Netflix use case as a driving example of a prototypical industrial-scale recommender system. We will also review the usage of modern algorithmic approaches that include algorithms such as Factorization Machines, Restricted Boltzmann Machines, SimRank, Deep Neural Networks, or Listwise Learning-to-rank. In the second part, we will focus on the area of context-aware recommendations where the two dimensional user-item recommender problem is turned into an n-dimensional space.

1. WHO WHOULD ATTEND

This tutorial targets two different kinds of audience. On the one hand, the tutorial content should appeal to academic researchers who are interested in understanding the practical applications of some of the latest machine learning and data mining approaches for building recommender systems. On the other hand, it will appeal to industry practitioners who want to understand how latest algorithmic approaches can be applied.

The tutorial assumes basic understanding of supervised and non-supervised machine learning approaches. Some prior knowledge of traditional recommendation techniques such as collaborative filtering is not required but is encouraged in order to obtain the most out of the tutorial.

Just as the Netflix Prize helped bring the Recommender Problem into the spotlight of the KDD research community, we think this tutorial will help expanding the somewhat shallow view the researchers now have of recommender systems. We therefore anticipate a big interest and impact in the KDD community.

KDD'14, August 24–27, 2014, New York, NY, USA.
ACM 978-1-4503-2956-9/14/08.
http://dx.doi.org/10.1145/2623330.2630807.

2. INSTRUCTORS

Xavier Amatriain (PhD) is Director of Algorithms Engineering at Netflix. He leads a team of researchers and engineers designing and implementing the next wave of machine learning approaches to power the Netflix product. Previous to this, he was a Researcher in Recommender Systems, and neighboring areas such as Data Mining, Machine Learning, Information Retrieval, and Multimedia. He has authored more than 50 papers including book chapters, journals, and articles in international conferences. He was General Chair for the 2010 ACM Recommender Systems Conference. He is often invited speaker at conferences, universities, and companies. He has also lectured in different universities including the University of California Santa Barbara and UPF in Barcelona, Spain, where he is originally from.

Bamshad Mobasher is a Professor of Computer Science and the director of the Center for Web Intelligence at the School of Computing of DePaul University in Chicago. His research areas include Web mining, Web personalization, recommender systems, predictive user modeling, and information retrieval. He has published five edited books as well as more than 170 scientific articles, including several seminal papers in Web mining and Web personalization that are among the most cited in these areas. Most recently, has served as the program chair and steering committee member of the ACM International Conference on Recommender Systems, as a program chair for the International Conference on User Modeling, Adaptation and Personalization, and as local organizing chair for the ACM Conference on Knowledge Discovery and Data Mining. As the director of the Center for Web Intelligence, Dr. Mobasher is directing research in Web mining, predictive analytics, and recommender systems, as well as overseeing several related joint projects with the industry. Dr. Mobasher serves as an associate editor for the ACM Transactions on the Web, ACM Transactions on Intelligent Interactive Systems and the ACM Transactions on Internet Technology. He also serves on the editorial board of User Modeling and User-Adapted Interaction, The Journal of Personalization Research.