Recommender Systems

An Introduction

In this age of information overload, people use a variety of strategies to make choices about what to buy, how to spend their leisure time, and even whom to date. Recommender systems automate some of these strategies with the goal of providing affordable, personal, and high-quality recommendations. This book offers an overview of approaches to developing state-of-the-art recommender systems. The authors present current algorithmic approaches for generating personalized buying proposals, such as collaborative and content-based filtering, as well as more interactive and knowledge-based approaches. They also discuss how to measure the effectiveness of recommender systems and illustrate the methods with practical case studies. The authors also cover emerging topics such as recommender systems in the social web and consumer buying behavior theory. Suitable for computer science researchers and students interested in getting an overview of the field, this book will also be useful for professionals looking for the right technology to build real-world recommender systems.

DIETMAR JANNACH is a Chaired Professor of computer science at Technische Universität Dortmund, Germany. The author of more than one hundred scientific papers, he is a member of the editorial board of the *Applied Intelligence* journal and the review board of the *International Journal of Electronic Commerce*.

MARKUS ZANKER is an Assistant Professor in the Department for Applied Informatics and the director of the study program Information Management at Alpen-Adria Universität Klagenfurt, Austria. He is an associate editor of the *International Journal of Human-Computer Studies* and cofounder and director of ConfigWorks GmbH.

ALEXANDER FELFERNIG is University Professor at Technische Universität Graz, Austria. His research in recommender and configuration systems was honored in 2009 with the *Heinz Zemanek Award*. Felfernig has published more than 130 scientific papers, is a review board member of the *International Journal of Electronic Commerce*, and is a cofounder of ConfigWorks GmbH.

GERHARD FRIEDRICH is a Chaired Professor at Alpen-Adria Universität Klagenfurt, Austria, where he is head of the Institute of Applied Informatics and directs the Intelligent Systems and Business Informatics research group. He is an editor of *AI Communications* and an associate editor of the *International Journal of Mass Customisation*.

Recommender Systems

An Introduction

DIETMAR JANNACH
Technische Universität Dortmund

MARKUS ZANKER
Alpen-Adria Universität Klagenfurt

ALEXANDER FELFERNIG
Technische Universität Graz

GERHARD FRIEDRICH
Alpen-Adria Universität Klagenfurt

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore,
São Paulo, Delhi, Dubai, Tokyo, Mexico City

Cambridge University Press
32 Avenue of the Americas, New York, NY 10013-2473, USA

www.cambridge.org
Information on this title: www.cambridge.org/9780521493369

First published 2011

Printed in the United States of America

A catalog record for this publication is available from the British Library.

Library of Congress Cataloging in Publication data

Recommender systems : an introduction / Dietmar Jannach . . . [et al.].
p. cm.
Includes bibliographical references and index.
ISBN 978-0-521-49336-9 (hardback)
1. Personal communication service systems. 2. Recommender systems
(Information filtering) I. Jannach, Dietmar, 1973– II. Title.
TK5103.485.R43 2010
006.3′3 – dc22 2010021870

ISBN 978-0-521-49336-9 Hardback

Additional resources for this publication at http:www.recommenderbook.net/

Contents

Foreword

It was a seductively simple idea that emerged in the early 1990s – to harness the opinions of millions of people online in an effort to help all of us find more useful and interesting content. And, indeed, in various domains and in various forms, this simple idea proved effective. The PARC Tapestry system (Goldberg et al. 1992) introduced the idea (and terminology) of collaborative filtering and showed how both explicit annotation data and implicit behavioral data could be collected into a queryable database and tapped by users to produce personal filters. Less than two years later, the GroupLens system (Resnick et al. 1994) showed that the collaborative filtering approach could be both distributed across a network and automated. Whereas GroupLens performed automated collaborative filtering to Usenet news messages, the Ringo system at Massachusetts Institute of Technology (MIT) (Shardanand and Maes 1995) did the same for music albums and artists and the Bellcore Video Recommender (Hill et al. 1995) did the same for movies. Each of these systems used similar automation techniques – algorithms that identified other users with similar tastes and combined their ratings together into a personalized, weighted average. This simple "k-nearest neighbor" algorithm proved so effective that it quickly became the gold standard against which all collaborative filtering algorithms were compared.

Systems-oriented exploration. With hindsight, it is now clear that these early collaborative filtering systems were important examples from the first of four overlapping phases of recommender systems advances. This systems-oriented exploration stage – through not only collaborative filtering but also knowledge-based systems such as the FindMe systems (Burke et al. 1996) – demonstrated the feasibility and efficacy of recommender systems and generated substantial excitement to move the field forward, in both research and commercial practice. (I do not mean to imply that these early research efforts did not also explore

algorithms and design alternatives, but to a great extent we were so excited that "the dog sang" that we did not worry too much about whether it was perfectly in tune.)

A key event in this phase was the Collaborative Filtering Workshop at Berkeley in March 1996. This gathering helped coalesce the community, bringing together people working on personalized and nonpersonalized systems, on divergent algorithmic approaches (from statistical summaries to k-nearest neighbor to Bayesian clustering), and on different domains. By the end of the day, there was a consensus that these were all aspects of one larger problem – a problem that quickly became known as recommender systems, thanks in part to a special issue of *Communications of the ACM* that grew out of the workshop (Resnick and Varian 1997).

Rapid commercialization – the challenges of scale and value. Recommender systems emerged into a rapidly expanding Internet business climate, and commercialization was almost immediate. Pattie Maes's group at MIT founded Agents, Inc., in 1995 (later renamed Firefly Networks). Our GroupLens group at Minnesota founded Net Perceptions in 1996. Many other companies emerged as well. Quickly, we started to face real-world challenges largely unknown in the research lab. To succeed, companies had to move beyond demonstrating accurate predictions. We had to show that we could provide valuable recommendations – usually in the form of selecting a few particular products to recommend that would yield additional purchases – and that we could do so without slowing down existing web sites. These systems had to work at greater-than-research scales – handling millions of users and items and hundreds or thousands of transactions per second. It is perhaps no surprise that the first book on recommender systems, John Riedl's and my *Word of Mouse*, was targeted not at researchers but at marketing professionals.

Research at the time moved forward to address many of these technological challenges. New algorithms were developed to reduce online computation time, including item-based correlation algorithms and dimensionality-reduction approaches, both of which are still used today. Researchers became more interested in evaluating recommenders based on "top-n" recommendation list metrics. A wide set of research explored issues related to implicit ratings, startup issues for new users and new items, and issues related to user experience, including trust, explanation, and transparency.

Research explosion – recommenders go mainstream. Somewhere between 2000 and 2005, many of the recommender systems companies dried up,

imploding with the Internet bubble or simply unable to compete with more mainstream companies that integrated recommendation into a broader set of business tools. As a technology, however, recommender systems were here to stay, with wide usage in e-commerce, broader retail, and a variety of knowledge management applications.

At the same time, research in recommender systems exploded with an infusion of people and approaches from many disciplines. From across the spectrum of artificial intelligence, information retrieval, data mining, security and privacy, and business and marketing research emerged new analyses and approaches to recommender systems. The algorithmic research was fueled by the availability of large datasets and then ignited by the 2006 announcement of the $1 million Netflix Prize for a 10 percent improvement in prediction accuracy.

Moving forward – recommenders in context. The excitement of the Netflix Prize brought many researchers together in heroic efforts to improve prediction accuracy. But even as these researchers closed in on success, a wave of researchers and practitioners were arguing for a step back toward the values of exploration and value. In 2006, MyStrands organized Recommenders06, a summer school on the present and future of recommender systems. In 2007, I organized the first ACM Recommender Systems Conference – a conference that has grown from 120 people to more than 300 in 2009. A look at the programs of these events shows increased interest in viewing recommendation in context, retooling research to ground it in an understanding of how people interact with organizations or businesses, and how recommendations can facilitate those interactions. Indeed, even though the field was nearly unanimously excited by the success of Netflix in bringing in new ideas, most of us also realized that an elaborate algorithm that improved predictions of just how much a user would dislike a set of bad movies did not help the user or Netflix. It is telling that the 2009 best-paper award went to a paper that demonstrated the flaws in the field's traditional "hold some data back" method of evaluating algorithms (Marlin and Zemel 2009), and that the most cited recent research paper on recommender systems is one that lays out how to match evaluation to user needs (Herlocker et al. 2004).

That brings us to this book. Behind the modest subtitle of "an introduction" lies the type of work the field needs to do to consolidate its learnings and move forward to address new challenges. Across the chapters that follow lies both a tour of what the field knows well – a diverse collection of algorithms and approaches to recommendation – and a snapshot of where the field is today, as new approaches derived from social computing and the semantic web find their

place in the recommender systems toolbox. Let us all hope that this worthy effort spurs yet more creativity and innovation to help recommender systems move forward to new heights.

Joseph A. Konstan
Distinguished McKnight Professor
Department of Computer Science and Engineering
University of Minnesota

Preface

"Which digital camera should I buy? What is the best holiday for me and my family? Which is the best investment for supporting the education of my children? Which movie should I rent? Which web sites will I find interesting? Which book should I buy for my next vacation? Which degree and university are the best for my future?"

It is easy to expand this list with many examples in which people have to make decisions about how they want to spend their money or, on a broader level, about their future.

Traditionally, people have used a variety of strategies to solve such decision-making problems: conversations with friends, obtaining information from a trusted third party, hiring an expert team, consulting the Internet, using various methods from decision theory (if one tries to be rational), making a gut decision, or simply following the crowd.

However, almost everyone has experienced a situation in which the advice of a friendly sales rep was not really useful, in which the gut decision to follow the investments of our rich neighbor was not really in our interest, or in which spending endless hours on the Internet led to confusion rather than to quick and good decisions. To sum up, good advice is difficult to receive, is in most cases time-consuming or costly, and even then is often of questionable quality.

Wouldn't it be great to have an affordable personal advisor who helps us make good decisions efficiently?

The construction of systems that support users in their (online) decision making is the main goal of the field of recommender systems. In particular, the goal of recommender systems is to provide easily accessible, high-quality recommendations for a large user community.

This focus on volume and easy accessibility makes the technology very powerful. Although recommender systems aim at the individual decisions of users, these systems have a significant impact in a larger sense because of their

mass application – as, for instance, Amazon.com's recommendation engines. Because of the far reach of the Internet market, this issue must not be underestimated, as the control of recommender systems allows markets themselves to be controlled in a broader sense. Consider, for example, a department store in which all the sales clerks follow orders to push only certain products.

One can argue that recommender systems are for the masses who cannot afford or are not willing to pay for high-quality advice provided by experts. This is partially true in some domains, such as financial services or medicine; however, the goal of making good decisions includes the aim of outperforming domain experts. Although this is clearly not possible and also not necessary in all domains, there are many cases in which the wisdom of the crowds can be exploited to improve decisions. Thus, given the huge information bases available on the Internet, can we develop systems that provide better recommendations than humans?

The challenge of providing affordable, personal, and high-quality recommendations is central to the field and generates many interesting follow-up goals on both a technical and a psychological level. Although, on the technical level, we are concerned with finding methods that exploit the available information and knowledge as effectively and efficiently as possible, psychological factors must be considered when designing the end-user interaction processes. The design of these communication processes greatly influences the trust in the subsequent recommendations and ultimately in the decisions themselves. Users rarely act as rational agents who know exactly what they want. Even the way a recommender agent asks for a customer's preferences or which decision options are presented will affect a customer's choice. Therefore, recommender systems cannot be reduced to simple decision theoretical concepts.

More than fifteen years have passed since the software systems that are now called "recommender systems" were first developed. Since then, researchers have continuously developed new approaches for implementing recommender systems, and today most of us are used to being supported by recommendation services such as the one found on Amazon.com. Historically, recommender systems have gained much attention by applying methods from artificial intelligence to information filtering – that is, to recommend web sites or to filter and rank news items. In fact, recommendation methods such as case-based or rule-based techniques have their roots in the expert systems of the 1980s. However, the application areas of recommender systems go far beyond pure information filtering methods, and nowadays recommendation technology is providing solutions in domains as diverse as financial products, real estate, electronic consumer products, movies, books, music, news, and web sites, just to name a few.

This book provides an introduction to the broad field of recommender systems technology, as well as an overview of recent improvements. It is aimed at both graduate students or new PhDs who are starting their own research in the field and practitioners and IT experts who are looking for a starting point for the design and implementation of real-world recommender applications. Additional advanced material can be found, for instance, in *Recommender Systems Handbook* (Ricci et al. 2010), which contains a comprehensive collection of contributions from leading researchers in the field.

This book is organized into two parts. In the first part, we start by summarizing the basic approaches to implementing recommender systems and discuss their individual advantages and shortcomings. In addition to describing how such systems are built, we focus on methods for evaluating the accuracy of recommenders and examining their effect on the behavior of online customers. The second part of the book focuses on recent developments and covers issues such as trust in recommender systems and emerging applications based on Web 2.0 and Semantic Web technologies. Teaching material to accompany the topics presented in this book is provided at the site http://www.recommenderbook.net/.

We would like to thank everyone who contributed to this book, in particular, Heather Bergman and Lauren Cowles from Cambridge University Press, who supported us throughout the editorial process. Particular thanks also go to Arthur Pitman, Kostyantyn Shchekotykhin, Carla Delgado-Battenfeld, and Fatih Gedikli for their great help in proofreading the manuscript, as well as to several scholar colleagues for their effort in reviewing and giving helpful feedback.

Dietmar Jannach
Markus Zanker
Alexander Felfernig
Gerhard Friedrich
Dortmund, Klagenfurt, and Graz, 2010