Joint DELOS-NSF Workshop on Personalisation and Recommender Systems in Digital Libraries

Alan Smeaton
School of Computer Applications
Dublin City University
Dublin 9, IRELAND
(asmeaton@compapp.dcu.ie)

Jamie Callan
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA
(callan@cs.cmu.edu)

Introduction

One of the important ways for users to feel comfortable with and become productive using information technology is to personalise or tailor systems to individuals or groups of users. This covers both explicit personalisation directly by the user, and implicit tailoring by systems that track users usage patterns and preferences and adapt systems and interfaces accordingly. The concept of personalisation thus is about making systems different for individual people, but the concept of personalization itself can mean different things.

One type of personalisation that is growing in use is recommender systems. Such systems take input directly or indirectly from users and based on user needs, preferences and usage patterns, recommender systems will then make personalised recommendations of products or services. These vary from recommending books to buy or TV programs to watch, to suggesting web pages to visit. The ultimate goal of such recommender systems is to reduce the amount of explicit user input and to operate, effectively, based on usage patterns alone, thus giving users what they want without them having to ask.

Personalisation and recommender systems may have different characteristics in digital library environments, because individual user behavior and aggregate traffic patterns may differ significantly from those of Web and E-commerce environments. For example, few digital libraries will see millions of transactions within a short period of time. Digital library resources may be structured (e.g., by digital librarians), and more stable (e.g., as compared with commercial sites). Digital library characteristics and usage patterns may provide opportunities for types of long-term learning that would be difficult or impractical in other environments, for example, learning about the relationships among resources in a single digital library or among resources in multiple digital libraries.

The Joint DELOS-NSF Workshop on Personalisation and Recommender Systems in Digital Libraries held at Dublin City University June 18-20, 2001, brought together 57 researchers and practitioners from 14 countries to discuss the development of personalisation and recommender systems and techniques, particularly as they apply to digital libraries. Fifteen papers and three invited talks were presented, covering a range of topics from basic personalisation technologies to studies of user interfaces. Moderated discussion sessions at the end of each day provided an opportunity for participants to reflect on the day's talks and to address recurring themes or issues missing from the talks.

Presentations

Below we provide very brief descriptions of the workshop presentations, to give a sense of the range of themes and topics covered. The complete workshop proceedings are available in paper and electronic formats from the European Research Consortium for Informatics and Mathematics (ERCIM). The URL for the electronic format is http://www.ercim.org/publication/ws-proceedings/DelNoe02/.

Susan Dumais opened the workshop with an invited talk giving an overview of personalisation research, viewed through a framework of "Assay, Isolate, and Accommodate" actions. Her talk discussed advantages and disadvantages of personalisation, and provided examples from commercial and research systems. The talk raised the question of whether there is an important difference between personalising system behavior, which sometimes annoys people, and personalising content, which seems to have been more successful.

Most of the papers at the workshop addressed either techniques or algorithms for personalisation and recommendation. Clodagh Moriarty spoke about using collaborative filtering techniques to predict student answers to exam questions. The advantage of this approach is that it needs no domain theory or complex user model, which enables new systems to be created quickly. Lillian Cassel and Ursula Wolz discussed a research Web search interface that would use contextual information from past searches to automatically expand search queries and to generate multiple query variants, to improve search effectiveness. Miles Efron and Gary Geisler presented a movie recommendation system that uses Singular Value Decomposition (SVD) for dimensionality reduction and generalization. A question raised by this talk is whether SVD was essentially learning movie genres, or something more specific (e.g., preferred genres, directors, actors).

An alternative technique to recommender systems was reported by Matthew Chalmers who presented a path-based view in which people are represented not by individual decisions, but by sequences of decisions. The path-based approach has the advantage that it enables more localized or context-specific recommendations. These ideas are being developed in two large-scaled research projects. In related work, Robert Villa described a framework for implicitly tracking user behavior over time, and then correlating syntactic information (actions such as moving the mouse or clicking a button) with semantic information (user goals at the time). The talk argued that the use of computer systems can be studied from the point of view of non-verbal communication, drawing upon past work in semiotics (the study of signs).

Aidan Finn described a system that automatically identifies the "content" portion of a Web page (as opposed to header, footer, and navigation information) based on vocabulary growth rates, and that classifies the content into "fact" or "opinion" categories based upon vocabulary and part-of-speech patterns. Saranya Maneeroj also addressed implementation issues, specifically the problem of recommendation sparsity during the startup phase of a recommender system by combining historical ratings (recommender systems approach) and content-based descriptions (information filtering approach).

Marcos André Gonçalves presented a specification-based framework for automatically constructing Digital Libraries customized for different user communities. The framework is based upon the PIPE personalisation methodology and the 5S theory (Streams, Structures, Spaces, Scenarios, and Societies) of Digital Libraries.

Two of the presentations looked at one implementation technique in particular, namely Language Modeling. Djoerd Hiemstra reviewed his recent work on the Language Modeling approach to Information Retrieval, demonstrating that term-specific mixing parameters enable some terms to be "required" in a document (similar to a Boolean AND operator), and showing that relevance feedback can be accomplished within the Language Modeling approach. Bruce Croft reviewed the successes and failure of prior research on relevance feedback and query expansion, then showed how these problems might be addressed in the Language Modeling approach to IR using relevance models. Relevance models can represent contexts of varying granularity, providing a method of integrating short- and long-term personalisation. Bruce also discussed an approach to representing query "clarity" that might enable ambiguous queries to be identified automatically.

The rest of the presentations were mostly descriptions of actual systems. Barry Smyth gave a lively invited talk on personalisation of television program listings when hundreds of television channels are available. Barry provided examples, case studies and a demonstration from his recent work in the commercial sector, arguing that in this environment there is insufficient context for content-based systems to succeed. Dmitry Briukhov presented a mediator architecture for Digital Libraries that encompasses a wide range of resource types, including text search and relational databases. The architecture includes a personalisation level that enables local modifications of resources such as ontologies. David Hicks discussed personalisation in a distributed Digital Library for knowledge workers in automobile manufacturing. Here, personalisation is accomplished through the use of fine-grained metadata, and can be done at different layers of the organization (individual, department, plant, etc) in a system which is in daily use. Mariella Di Giacomo described Web-based personal Digital Libraries used in a large research laboratory. Librarians select, organize and disseminate resources initially, based upon general knowledge of user information needs, but individuals can add/delete their own information, organize it however they wish, and share it in a system which is also in daily use. Jim French described a recommender system that selects songs for personalised Internet radio stations. In this domain it is important to select good songs, to provide local coherence in sequences of songs, and to model the way interest in an artist varies over time.

Apart from the papers that described either techniques or systems, there was one paper that looked at evaluation. Kirsten Swearingen discussed a user study comparing recommendations given by friends to recommendations given by six commercial or widely available recommender systems. The study also investigated the degree to which assessments of overall recommendation system quality were correlated with the quality of the recommendations or the quality of the user interface.

Clifford Lynch ended the workshop with the third invited talk, which was from a Library Science and Digital Library point of view. He compared system-specific (one system, many users) and user-specific (one user, many systems) personalisation, arguing that the former is often done more for the benefit of the vendor/service-provider than for the user, and discussing the implications of each approach. Privacy issues, lessons learned from traditional libraries, and the social dynamics of information dissemination were also discussed. As a whole, this talk encouraged researchers to step back from studying individual algorithms, to think about the larger issues involved in providing many different types of information services to individuals, and to become more imaginative and creative about the services that our Digital Libraries provide. As such, it was the perfect closing talk for the workshop.

Conclusions

Several themes recurred in workshop discussions.

Although the workshop included talks on software architectures, several deployed systems, and a user study, most of the presentations focused on algorithms for making recommendations. In spite of this emphasis on algorithms by authors, the papers were generally consistent with past research in this area, which has been adaptation of existing technologies, mostly from Information Retrieval or Machine Learning, often combined with creative methods of acquiring training data. The area is not yet mature enough to have spawned its own specific and tailored techniques, or the theoretical models that would support them, which is a sign that there is much more work yet to be done.

Workshop participants naturally felt that personalisation is an important research area, and that it will receive greater attention in the coming years. There was less agreement on how it will be applied in the context of digital libraries, and indeed personalisation research in digital library environments was underrepresented at the workshop. Clifford Lynch's talk suggested several reasons for this under-representation, including a strong tradition of privacy in the Library Science community, which discourages long-term storage of personal information, and the development history of recommender systems, which has been driven by the commercial sector. Most of the "success" stories are systems that encourage additional consumption, for example of movies and books. It is harder to find systems that provide other types of improvements that might be more important in a digital library environment, for example productivity improvements.

As is often the case, research has been influenced strongly by the data available to the research community. Movie recommender systems are a popular research vehicle because movie databases are freely available on the Web, but only a narrow range of research can be done in this context. Some researchers have access to more interesting datasets, but these are often proprietary, which makes research results difficult to evaluate and reproduce, and which raises a barrier to entry by new researchers. Workshop participants agreed that there is a strong need for a more diverse set of generally-available

data resources for personalisation research, and encourage the funding community to consider this in their funding decisions.

Evaluation is a related problem. Most of the systems discussed in the workshop were evaluated in some manner, but few of the evaluations could be called rigorous. It *might* be clear how to evaluate a movie recommender system, but it is less clear how to evaluate more complex systems. Researchers in this area need to begin considering a broad approach to evaluation that embraces not just quantitative evaluation, but also methods and tools from disciplines such as sociology. Researchers in this area recognize the importance of evaluation, but worry that an overemphasis on evaluation will stifle new ideas in this fledgling field. Finding the right balance between creative development of new ideas and scientific evaluation remains an issue in this research area.

Prior research on personalisation and recommender systems has focused on relatively short periods of time. Systems are beginning to be deployed that are intended for daily use over long periods of time (two were discussed in this workshop), but little is known about how such systems and their users might evolve over time. It is important to begin studying long-term personalisation issues, for example following a particular group of users over a several year period. Workshop participants felt that funding bodies need to give greater attention to longer-term projects.

Finally, as Clifford Lynch pointed out in his talk, most usage information is specific to particular systems, not to individuals. There is a need for research that would enable users to "take their histories and preferences with them" when they switch vendors or service providers, otherwise personalisation becomes a barrier to competition.

Acknowledgements

We thank the other program committee members, Nick Belkin, William Cohen, Tom Fawcett, Yannis Ioannidis, Maurice Mulvenna, John Riedl, and Barry Smyth, for their help in organizing the workshop.

The Joint DELOS-NSF Workshop on Personalisation and Recommender Systems in Digital Libraries was funded by the DELOS Network of Excellence on Digital Libraries and by the National Science Foundation under grant IIS-011855. Any opinions, findings, conclusions, or recommendations expressed in this report are the authors', and do not necessarily reflect those of the sponsors.