

# Research Statement

Michael Ekstrand

I apply the lens of human-computer interaction research to study recommender systems and other information retrieval, filtering, navigation, and decision support tools. Two key questions drive my research agenda:

1. How can these systems be designed and adapted to better meet human needs in real applications?
2. How do these systems interact with their users individually, the user base as a group, and society more broadly?

Recommender systems are a ubiquitous part of the Internet experience, helping users to find the information and resources they want and improving the sales of their proprietors. However, there is a great deal we still do not yet know about how to best design these systems to support their users' needs and decision-making processes, and how the recommender and its sociotechnical context support and affect each other.

The high-level aim of my research is to maximize the benefit of recommenders and other intelligent information systems to their users and to society, while identifying and mitigating possible risks that can arise from the use and deployment of these systems.

My earliest recommendation work was application-focused, building a recommender system for research papers [5] and a context-sensitive search tool for software help [6]. As my research program matured, I engaged in two major projects: building the LensKit toolkit to support reproducible recommender systems research [3] and studying the ways in which recommender algorithms differ in their ability to meet the information needs of different users [1, 4]. This work is motivated by the idea that it is unlikely that there is a one-size-fits-all recommendation technique, and that different users will be better served by different algorithms; the goal was to map out better the ways in which different algorithms produce different results, and how those differences correspond to the user's satisfaction with the system.

I have followed up on this work with another study on how users select recommender algorithms [2], and am supervising students in ongoing work aimed at improving the reliability and efficiency of tuning recommender algorithms and understanding additional dimensions along which the output of various algorithms may differ in useful ways.

My work building LensKit has raised a number of technical challenges, one of which is the issue of configuring recommender algorithms in a flexible and composable fashion; our solution to this problem, developed with Michael Ludwig, has resulted in another paper currently under review for the *Journal of Object Technology*.

I employ a variety of methods in my research, including offline experiments with public data sets, online user studies and surveys, and building software to support reproducible

research. My research dissemination strategy relies heavily on open-source software publication, enabling others to more easily reproduce, validate, and build on my results.

My research also benefits from strong interdisciplinary collaborations. Researchers from the university libraries contributed to the research paper recommendation work, and our work on user-perceptible differences in recommender algorithms is part of an ongoing collaboration with Martijn Willemsen, a decision psychologist at the Eindhoven University of Technology. As I extend my research into new applications and questions, I hope to continue collaborating with researchers with both domain expertise and skills useful to studying the behavior of recommender systems and their users.

To date, my research has focused on the interaction of the recommender system with a single user. I am developing a new research agenda that will broaden this lens to study the ways in which recommenders (and other intelligent information systems) respond to and affect their user population as a whole, as well as the broader social structures in which the systems are embedded. I currently have a proposal under review at the NSF and a student working on preliminary results for my first project in this program.

I have another proposal under review to develop interactive recommenders to help users in the process of building organized collections of items such as research bibliographies. This work will look at how to integrate multiple recommender algorithms, each with a unique perspective on what items might be relevant, into a compelling interactive system that makes it easier for users to understand how they can build better collections.

Recommender systems particularly, and intelligent information systems in general, are powerful tools that help their users achieve great things for themselves and, perhaps, for society. It's the goal of my research to understand how we can maximize this good.

## References

- [1] Ekstrand, M.D., Harper, F.M., Willemsen, M.C. and Konstan, J.A. 2014. User Perception of Differences in Recommender Algorithms. In *RecSys '14* (Oct. 2014).
- [2] Ekstrand, M.D., Kluver, D., Harper, F.M. and Konstan, J.A. 2015. Letting Users Choose Recommender Algorithms: An Experimental Study. In *RecSys '15*, 11–18.
- [3] Ekstrand, M.D., Ludwig, M., Konstan, J.A. and Riedl, J.T. 2011. Rethinking the Recommender Research Ecosystem: Reproducibility, Openness, and LensKit. In *RecSys '11* (2011), 133–140.
- [4] Ekstrand, M.D. and Riedl, J.T. 2012. When recommenders fail: predicting recommender failure for algorithm selection and combination. In *RecSys '12*, 233–236.
- [5] Ekstrand, M., Kannan, P., Stemper, J., Butler, J., Konstan, J.A. and Riedl, J. 2010. Automatically Building Research Reading Lists. In *RecSys '10* (2010), 159–166.
- [6] Ekstrand, M., Li, W., Grossman, T., Matejka, J. and Fitzmaurice, G. 2011. Searching for Software Learning Resources Using Application Context. In *UIST '11* (2011), 195–204.