

# Research Statement

## Michael Ekstrand

I research ways to help people find, filter, track, remember, and make better use of information. Within this broad interest, I have devoted most of my research energies to date to studying recommender systems. My research activities include system building, offline experimentation with public data sets, and user studies, all bent towards one goal: understanding how algorithms can help users find the things they want or need.

My early work on recommendation focused on applications. I have published work on recommending reading lists for new researchers and on mining users' interactions with an application to improve a search engine's ability to find appropriate help resources.

I have since focused my work around enabling *recommender engineering*: attempting to discover and organize the knowledge needed to engineer recommenders from well-known principles instead of building and testing every recommender from scratch. Under this paradigm, a recommender developer will identify properties of their domain, the tasks they need to support, and the goals of their users; they will use this information find an algorithm or two that meets their requirements and fine-tune the result with user testing.

Many current experiments examine individual pairings of algorithms and tasks. In industrial settings, A/B testing has been very valuable to develop and test recommenders. But pairwise approaches do not directly produce generalizable results, particularly about mediating factors: what made algorithm A better for task B, and how can that knowledge be applied to engineer a new recommender application?

My work moves the field forward by two means: first, I have built software to support reproducible research in recommender systems and make it easier to experiment with new algorithm or evaluation ideas in both offline and online, user-facing settings. This makes it easier to explore and study a wide range of recommendation topics, and to replicate and validate the results. This work has also required novel design and development for object-oriented software configuration.

Second, I am conducting user experiments to understand how algorithms differ in their ability to meet users' needs in using a system. I am going past previous work to try to determine *how* they differ (e.g. what differing characteristics their outputs exhibit) in a psychologically robust way. To achieve this, I have been collaborating with Dr. Martijn Willemsen, a decision psychologist at the Eindhoven University of Technology.

I hope to continue my blend of data-driven and user-based experimentation, along with interdisciplinary collaboration and open-source software publication, in my future research. My recent research has focused on entertainment recommendation (movies and music) due to the availability of data and experimentation platforms. I would like to work again on recommendation, retrieval, filtering, and exploration in more 'serious' domains such as research papers, news, and social media.

I am also interested in recommendation and privacy, both developing recommendation technologies that interact well with user expectations of privacy and using recommendation to improve the usability of privacy-enhancing systems.