Search from Personal to Social Context: Progress and Challenges

Paul N. Bennett
Microsoft Research
Redmond, USA

ABSTRACT

User and behavioral modeling plays a critical role in a variety of online services such as web search, advertising, e-commerce, and news recommendation. For example, our ability to accurately interpret the intent of a web search can be informed by knowledge of the web pages a searcher was viewing when initiating the search or recent actions of the searcher such as queries issued, results clicked, and pages viewed. In this talk, I will describe a recent framework for personalized search which improves the quality of search results by enabling a representation of a broad variety of context including the searcher's long-term interests, recent activity, current focus, and other user characteristics. Then, I will review a variety of related work that extends these approaches from signals focused on the individual to social signals such as likes, cohorts, and affiliation networks. Finally, I'll speculate on how social signals and networks can provide directions for relatively unexplored directions in social personalized retrieval.

Keywords

Social Personalization and Search, Social signals

1. BIO

Paul Bennett is a Senior Researcher in the Context, Learning & User Experience for Search (CLUES) group at Microsoft Research where he focuses on the development, improvement, and analysis of machine learning and data mining methods as components of real-world, large-scale adaptive systems. His research has advanced techniques for ensemble methods and the combination of information sources, calibration, consensus methods for noisy supervision labels, active learning and evaluation, supervised classification (with an emphasis on hierarchical classification) and ranking with applications to information retrieval, crowdsourcing, behavioral modeling and analysis, and personalization. His recent work has been recognized with a SIGIR 2012 Best Paper Honorable Mention and a SIGIR 2013 Best Student Paper award. He completed his dissertation on combining text classifiers using reliability indicators in 2006 at Carnegie Mellon where he was advised by Profs. Jaime Carbonell and John Lafferty.