Keynote: Search and Recommendation in the Enterprise

Paul Bennett
Microsoft Research
Paul.N.Bennett@microsoft.com

Abstract

Advances in information retrieval have mainly been tied to shared collections such as the web where massive amounts of click data drives signals of relevance. However, advances have been slower over enterprise search and personal information collections where click data is sparse and privacy restricts cross-user generalization. Unlike in web search where we have developed a solid understanding of the types of tasks people undertake and how those relate to query intents, we lack a complete understanding of how peoples tasks relate to actions they perform in a productivity setting, e.g. interacting with emails, finding files, etc. We review work that seeks to understand personal tasks and how this relates to information finding as well as task planning. We present arguments more generally that information assistance should happen contextually within an application. Specifically we focus on several projects which provide contextual intelligence by: (1) presenting the information you want before you ask at the right time, place, and context; (2) and addressing your information needs in a context-specific way to the task you are doing. We present a summary of work that aims to identify the underlying goals which drive peoples behavior in memory recall, task planning, and meeting preparation and describe prototypes which leverage these insights to better support people through intelligent context-aware interfaces. Finally we conclude with a discussion of research challenges in the enterprise and personal information space.

This talk presents joint work with Adam Fourney, Ryen White, Eric Horvitz, David Graus, Xin Rong, Qian Zhao, Susan Dumais, Adam Troy, Shane Williams and Anne Loomis Thompson as well as Ahmed Awadallah, Horaiu Bota, Robin Brewer, Nirupama Chandrasekaran, Fernando Diaz, Cristina Garbacea, Nick Ghotbi, Marcello Hasegawa, Richard Hughes, Abhishek Jha, Ece Kamar, John Krumm, Merrie Morris, Rev Rameshkumar and likely many others.

Bio

Paul Bennett is the Principal Research Manager of the Information & Data Sciences group in Microsoft Research AI. His published research has focused on a variety of topics surrounding the use of machine learning in information retrieval - including ensemble methods and the combination of information sources, calibration, consensus methods for noisy supervision labels, active learning and evaluation, supervised classification and ranking, crowdsourcing, behavioral modeling and analysis, and personalization. Some of his work has been recognized with awards at SIGIR, CHI, and ACM UMAP as well as an ECIR Test of Time Honorable Mention award. Prior to joining MSR in 2006, he completed his dissertation in the Computer Science Department at Carnegie Mellon with Jaime Carbonell and John Lafferty. While at CMU, he also acted as the Chief Learning Architect on the RADAR project from 2005-2006 while a postdoctoral fellow in the Language Technologies Institute.

In: Joint Proceedings of the First International Workshop on Professional Search (ProfS2018); the Second Workshop on Knowledge Graphs and Semantics for Text Retrieval, Analysis, and Understanding (KG4IR); and the International Workshop on Data Search (DATA:SEARCH18). Co-located with SIGIR 2018, Ann Arbor, Michigan, USA – 12 July 2018, published at http://ceur-ws.org