



LTI COLLOQUIUM



Wei-Ying Ma
Microsoft Research, Asia

Friday, November 22.
2:30 PM, Doherty 2315

Advancing Web Search with New Capabilities from Knowledge Mining, Deep Learning, and Crowdsourcing

In recent years, we have seen dramatic improvements in machine learning, knowledge mining, and crowdsourcing that are providing search engines with new capabilities to perform deeper natural language processing and understanding on queries and documents. For example, deep learning offers "bottom-up" capability to learn (feature) representation from big data, and there has been exciting progress in using deep learning for the semantic embedding of words, queries, and documents into a vector space for advanced text processing. Knowledge mining, which refers to a variety of techniques that data-mine and extract knowledge such as entities and relationships from the Web, has made significant advances in the past few years, and these techniques are being used by search engines to create a comprehensive knowledge graph that offers "top-down" capability for understanding data. Combining the bottom-up and top-down capabilities with crowdsourcing that cleverly includes human computation in the loop to receive implicit or explicit feedback and verification, we have seen many new ideas and systems developed to assist automatic algorithms for higher accuracy and performance. In this talk, I will introduce some of the exciting progress in these areas and show how search engines might be able to leverage these new capabilities to enable new search experiences and scenarios.

Bio: Dr. Wei-Ying Ma is an Assistant Managing Director at Microsoft Research Asia, where he oversees multiple research groups, including Web Search and Data Management, Natural Language Computing, Knowledge Mining, Machine Learning, and Internet Economics and Computational Advertising. He and his team of researchers have developed many key technologies that have been transferred to Microsoft's Online Services Division, including Bing Search Engine and Microsoft Advertising. He has published more than 250 papers at international conferences and in journals. He is a Fellow of the IEEE and a Distinguished Scientist of the ACM. He served on the editorial boards of ACM Transactions on Information System (TOIS) and is a member of the International World Wide Web (WWW) Conferences Steering Committee. In recent years, he has served as program co-chair of WWW 2008 and as general co-chair of ACM SIGIR 2011.