

Information Retrieval Evaluation as Search Simulation

Chengxiang Zhai

University of Illinois at Urbana-Champaign

ABSTRACT

Due to the empirical nature of the Information Retrieval (IR) task, experimental evaluation of IR methods and systems is essential. Historically, evaluation initiatives such as TREC, CLEF, and NTCIR have made significant impacts on IR research and resulted in many test collections that can be reused by researchers to study a wide range of IR tasks in the future. However, despite its great success, the traditional Cranfield evaluation methodology using a test collection has significant limitations, especially for evaluating an interactive IR system, and it remains an open challenge how to evaluate interactive IR systems using reproducible experiments. In this talk, I will discuss how we can address this challenge by framing the problem of IR evaluation more generally as search simulation, i.e., having an IR system interact with simulated users and measuring the performance of the system based on its interaction with the simulated users. I will first present a general formal framework for evaluating IR systems based on search session simulation, discussing how the framework can not only cover the traditional Cranfield evaluation method as a special case but also reveal potential limitations of the traditional IR evaluation measures. I will then review the recent research progress in developing formal models for user simulation and evaluating user simulators. Finally, I will discuss how we may leverage the current IR test collections to support simulation-based evaluation by developing and deploying user simulators based on those existing collections. I will conclude the talk with a brief discussion of important future research directions in simulation-based IR evaluation.

BIOGRAPHY

ChengXiang Zhai is a Donald Biggar Willett Professor in Engineering of the Department of Computer Science at the University of Illinois at Urbana-Champaign, where he also holds a joint appointment at the Carl R. Woese Institute for Genomic Biology, Department of Statistics, and the School of Information Sciences. He received a Ph.D. in Computer Science from Nanjing University in 1990, and a Ph.D. in Language and Information Technologies from Carnegie Mellon University in 2002. He worked at Clairvoyance Corp. as a Research Scientist and a Senior Research Scientist from 1997 to 2000. His research interests are in the general area of intelligent information systems, including specifically intelligent information retrieval, data mining, natural language processing, machine learning, and their applications in domains such as biomedical informatics, and intelligent education systems. He has published over 300 papers in these areas and holds 6 patents. He offers two Massive Open Online Courses (MOOCs) on Coursera covering Text Retrieval and Search Engines and Text Mining and Analytics, respectively, and was a key contributor of the Lemur text retrieval and mining toolkit. He served as Associate Editors for major journals in multiple areas including information retrieval (ACM TOIS, IPM), data mining (ACM TKDD), intelligent systems (ACM TIST), and medical informatics (BMC MIDM), Program Co-Chairs of NAACL HLT'07, SIGIR'09, and WWW'15, and Conference Co-Chairs of CIKM'16, WSDM'18, and IEEE BigData'20. He is an ACM Fellow and a member of ACM SIGIR Academy. He received multiple awards, including ACM SIGIR Gerard Salton Award, ACM SIGIR Test of Time Paper Award (three times), the 2004 Presidential Early Career Award for Scientists and Engineers (PECASE), Alfred P. Sloan Research Fellowship, IBM Faculty Award, HP Innovation Research Award, Microsoft Beyond Search Research Award, UIUC Rose Award for Teaching Excellence, and UIUC Campus Award for Excellence in Graduate Student Mentoring. He has graduated 38 PhD students and over 50 MS students.