

CS Talk (Friday)

Title: Balancing Information Protection and Information Access

Speaker: Douglas W. Oard, University of Maryland (USA)

Abstract:

Current search engines are designed to find things, but there are many cases in which we actually don't want some things to be found. In particular, we can't yet make many potentially valuable collections available to be searched because they contain some intermixed sensitive content that requires protection, but for which reliable sensitivity labels are not available. Some prominent examples include government transparency regimes such as New Zealand's Official Information Act, the rapidly growing backlog of national security information in the United States that is awaiting declassification review, and the vast troves of email that are now accumulating in both government archives and personal collections. The scale of many of these problems is such that asking people to mark all the sensitive content in a collection would simply be impractical. If we are ever to be able to find that which is not actually sensitive, we will thus need to build systems—systems involving both people and automation—that are able to recognize and protect that which requires protection. We formulate this as a multi-objective optimization problem in which the goal is to balance information access with information protection. I'll describe two implementations of this broad idea. In the first, designed for high-stakes tasks such as topic-focused declassification review, the search is performed on behalf of the end user by a trusted intermediary (e.g., an archivist), and the system's goal is to focus that intermediary's limited time in a way that balances the risk of missing relevant content with the risk of revealing sensitive content. In the second, designed for higher-volume but lower-stakes cases such as searching archived through email, we seek to support end-user search by using a risk-averse search engine to surface some relevant content, thus allowing searchers to explore collections to find some immediately useful content, with the beneficial side effect that their final refined queries might also be used to flag the more difficult decisions for (future) human review. Both of these techniques require automated sensitivity classifiers, so I will also briefly describe three lines of work on that problem: one for email, a second for government transparency, and a third for national security information. This is joint work with Jason Baron, Mahmoud Sayed, Nate Rollings, Fabrizio Sebastiani and Jyothi Vinjumur.

About the Speaker:

Doug Oard is a Professor at the University of Maryland, with joint appointments in the College of Information Studies and the University of Maryland Institute for Advanced Computer Studies (UMIACS). He earned his Ph.D. in Electrical Engineering from the University of Maryland, and his research interests center around the use of language technologies such as speech recognition, machine translation, document image analysis, knowledge representation, processing mathematical notation, and social network analysis to support information access to support information seeking by end users. More on Doug's research can be found at <http://terpconnect.umd.edu/~oard>. He will be visiting VUW during his sabbatical leave for the week of January 15, 2024, hosted by Chern Li Liew.