Abstract

This report constitutes the proceedings of the eighth Text REtrieval Conference (TREC-8) held in Gaithersburg, Maryland, November 16–19, 1999. The conference was co-sponsored by the National Institute of Standards and Technology (NIST) and the Defense Advanced Research Projects Agency (DARPA), and was attended by 170 people. Sixty-six groups including participants from 16 different countries were represented.

The goal of the conference was to bring research groups together to discuss their work on a large test collection. The diversity of the participants meant that a wide variety of retrieval techniques were represented, including new models for retrieval and refined machine learning techniques. Results were scored using a common evaluation package, so groups were able to compare the effectiveness of different techniques, and to discuss how differences between systems affected performance. In addition to the main evaluation, seven additional evaluations, called "tracks," allowed participants to focus on particular common subproblems. Two new tracks were introduced this year. The question answering track encouraged research into systems that return answers, rather than ranked lists of documents, in response to a question. The web track investigated retrieval performance over a collection of World Wide Web pages.

The conference included paper sessions and discussion groups. This proceedings includes papers from most of the participants (some groups did not submit papers), track reports that define the problem addressed by the track plus summarize the main track results, and tables of individual group results. The TREC-8 proceedings web site also contains system descriptions that detail the timing and storage requirements of the different runs.