

ARCH: An Adaptive Agent for Retrieval Based on Concept Hierarchies

B. Mobasher

We present a client-side agent, named ARCH, for assisting users in one of the most difficult information retrieval tasks, i.e., that of formulating an effective search query. The agent utilizes a hierarchically-organized semantic knowledge base in aggregate form, as well as an automatically learned user profile, to enhance user queries. In contrast to traditional methods based on relevance feedback, ARCH assists users in query modification prior to the search task. The initial user query is (semi-)automatically modified based on the user's interaction with an embedded, but modular, concept hierarchy. The modular design of the agent allows users to switch among the representations of different domain-specific hierarchies depending on the goals of the search. ARCH passively learns a user profile by observing the user's past browsing behavior. The profiles are used to provide additional context to the user's information need represented by the initial query. The full system also incorporates mechanisms for categorizing and filtering the search results, and using these categories for performing refined searches in the background. Preliminary experiments have shown that the agent can substantially improve the effectiveness of information retrieval both in the general context of the Web, as well as for search against domain-specific document indexes.