

ABSTRACT

We leverage a dataset of Dutch job postings, applicants' CVs and the corresponding hiring decisions made by real recruiters to improve the ranking of candidates with respect to search queries representative of real-world job offers. For this purpose we first propose a field relevance model of CVs, which can represent implicit domain knowledge available in a large collection of CVs without or with minimal supervision. We show in a query expansion experiment that such a model alone can improve the recall of the retrieval of candidates. In a second step we learn a (re-)ranking model which operates on the initial ranking of a search engine. We demonstrate that this model can be obtained through a learning-to-rank framework, for which we define a number of features, including features which can be computed based on the field relevance model. Hiring decisions, i. e. whether a candidate was hired, interviewed or rejected, can be used as relevance labels in the learning scheme. In our experiments we show that this supervised training procedure is able to produce a reranking model that improves significantly over the search ranking on common information retrieval performance metrics.