

Users typically search the web by entering a small number of keywords in a search engine. The search engine uses an inverted index to identify all documents containing the keywords and ranks the results given some relevance function that depends on many different factors, including the query context, the application domain, location, temporal criteria, and more. Nevertheless, natural language is inherently ambiguous hence, to complicate matters further, expressing a query using a few keywords only, results in even larger inherent ambiguity and loss of relevant context. If a search engine does not consider the Diversity of query results during ranking, especially for queries with large ambiguity, the results produced will be unsatisfactory with high probability. A simple example is one of the most popular queries in YellowPages.com, Yahoo! Local and Google Maps: “Restaurant”. Depending on the query location, a typical search engine might rank a large number of Italian or Chinese restaurants before the first Thai restaurant appears. Clearly, the user intent for this particular query is not known, but it is safe to assume that since the user did not explicitly specify a particular cuisine, he/she might be interested in all options available in the vicinity of the query. It can be argued that diversification of results will invariably result in increased user satisfaction.

Query result diversity has been extensively studied the past few years and has gained traction in the industrial sector as well (for example, all major search engines implement some form of diversity ranking). Moreover, diversity is not limited to information retrieval. In general, it can be applied to any query that can return a very large number of results. Diversity based queries on traditional database systems, diversity in recommender systems, diversity through user preferences, diversity in publish/subscribe systems, diversity on structured data, and diverse skyline computation using preference functions are some other application domains where diversification plays a big role in improving the end-user experience. This special issue of the *Data Engineering Bulletin* presents a snapshot of the current state-of-art research on diversity for the aforementioned topics.

The first article [Gollapudi and Sharma] presents an axiomatic framework for diversification, formalizing the problem and setting the stage for what follows. The second article [Haritsa] presents a diversity aware algorithm for k-NN processing in traditional database systems. The third [Ziegler and Lausen] and fourth [Amer-Yahia, Lakshmanan, Vassilvitskii and Yu] articles discuss diversification in recommender systems, and more specifically collaborative filtering. The article by [Chen and Li] focuses on efficient query evaluation based on pre-defined user preferences (as a means of expressing diversity). The article by [Drosou and Pitoura] discusses content diversity in the case of continuous data delivery (in the context of publish/subscribe systems). The next article [Vee, Shanmugasundaram, Amer-Yahia] talks about diversification of query results on structured data in the presence of a preference hierarchy of attributes. The final article [Tao] discusses diversification of skyline query results, based on user preference functions.

The goal of this special issue is to instigate further research into the essential problem of diversification. Even though these past few years there has been an impressive number of contributions of high quality into this field, we feel that there is still a lot of room for growth and improvement. It is our belief that current research has only scratched the surface of diversification in the respective domains, and the applications of diversification will keep growing. We would like to thank all authors for the time and effort they put into editing their excellent articles for this special issue. We would also like to extend a special thanks to Marcos Vieira at UC Riverside for editorial assistance with the issue.