Search has become a daily activity. Due to its popularity, most of us expect to be able to use search everywhere and, when we do, want it to work, that is, return good results regardless of the complex nature of the data or the task we are looking to achieve. Despite its popularity, search is yet to be improved. This issue contains several messages about how we should go about making search better. The nature of the data and its richness, structured objects, XML data, user-contributed content, require various semantic interpretations of search. The user intention, whether searching personal information, looking to explore topics or finding product information, is another dimension affecting the effectiveness of search results. Finally, the advent of the social Web is changing users' expectations and their perceptions of data quality, and hence, the requirements of search.

The first article, by Kumar and Tomkins, gives us insights on how people search based on a large-scale analysis of search and toolbar logs. It observes, among other things, that more than half of search queries are about structured objects. The authors then study the relationship between search and e-commerce queries as they represent a monetizable user intent.

The second article, by Anand Rajaraman, describes Kosmix, a general-purpose topic exploration engine which goes beyond Web search to extract Deep Web information using taxonomies.

The following two articles, by Marian and Wang, and, Dittrich, Vas Sallez and Blunschi, discuss the state of the art in searching personal information. The first article describes a method which uses query approximation to search personal content across file boundaries. The second article describes the interplay between search and data integration. Users can not only search their data collections but also semantically integrate it as they search.

The fourth article reports on experiences and challenges in building an XML search engine. one main distinction of searching XML data is the ability to determine which portion of a document is most relevant to a user's request, as opposed to traditional Web search, where entire documents are returned.

Searching user-contributed and shared content is tackled in the last two articles by Abiteboul and Polyzotis and, Agichtein, Gabrilovich and Zha. The first article describes data rings, an infrastructure, which enables users to declaratively share and consume other users' content. The second article applies machine learning techniques to finding appropriate ranking functions when searching social content.

I would like to thank the authors who graciously volunteered their time and effort in putting together this special issue.

Happy reading!