

Searching the Enterprise

Udo Kruschwitz
University of Essex
udo@essex.ac.uk

Charlie Hull
Flax
charlie@flax.co.uk

ABSTRACT

Search has become ubiquitous but that does not mean that search has been solved. Enterprise search, which is broadly speaking the use of information retrieval technology to find information within organisations, is a good example to illustrate this. It is an area that is of huge importance for businesses, yet has attracted relatively little academic interest. We introduce our recently published monograph[1] which explores the main issues involved in enterprise search both from a research as well as a practical point of view. We first plot the landscape of enterprise search and its links to related areas. This will allow us to identify key features before we survey the field in more detail and discuss what needs to be done to make enterprise search work.

KEYWORDS

Enterprise search, expertise search, email search, data silos

ACM Reference Format:

Udo Kruschwitz and Charlie Hull. 2017. Searching the Enterprise. In *Proceedings of DIR'17*. ACM, New York, NY, USA, Article 4, 1 page. https://doi.org/10.475/123_4

1 MOTIVATION

“Enterprise Search doesn’t work” summarises a view commonly held by employees trying to find information within their organisation. On the other hand, an information retrieval (IR) researcher or student might never have heard this or even thought about it.

Given the wide-spread sentiment among users and search practitioners that enterprise search does not deliver on its promises, the question arises as to what is it that causes these perceptions and do they properly reflect the reality? One of the fundamental issues underlying the overall discussion is the question of how does enterprise search relate to search in general and Web search in particular. We will provide a thorough discussion of the topic and outline implications and guidance resulting from this. We will focus on both theoretical and practical issues as well as their interplay.

Web search is a relatively recent development that has attracted much attention and for many it has become a synonym for ‘search’ in general. But Web search is just one – perhaps the most prominent – search context. There are many other application areas, enterprise search being one of them, which require fundamentally different solutions. Note that we do not simply want to reduce our discussion to a comparison of searching the Web with searching within

an organisation – other types of search include Web site search, database search, and desktop search among others. Identifying the key features of each search type allows a systematic comparison and these features can be characterised by investigating a number of dimensions that reflect, for example: collection and document properties, structural and organisational differences, different types of users and information needs, the level of support needed etc. Unpacking these differences results in a fairly complex picture with equally complex implications. To choose just one way of contrasting enterprise search with related areas, we observe that while Web search is aimed at high *precision*, in enterprise search *recall* is often at least as important as precision.

Given that our target audience are mainly researchers and PhD students with some background in information retrieval, we consider DIR’17 to be the perfect place to present our work to the large IR community in the Netherlands and Belgium (and beyond).

2 CONCLUSIONS

The key findings of this work are as follows:

- (1) Enterprise search is an area that is hugely important in industry yet has attracted relatively little academic interest.
- (2) There are substantial differences between enterprise search and other types of search such as Web search which include heterogeneous data sources, silo-based repositories, and users defined by roles in the enterprise.
- (3) Enterprise search will not work out-of-the-box, and the human in the loop is essential, e.g. for customisation, continuous relevance assessment and tuning.
- (4) People search and email search are among the dominant search types in an enterprise. Information needs are driven by business needs, e.g. people search is mainly aimed at finding experts or expertise and not for entertainment.
- (5) Evaluation in enterprise search is essential and in some way very different to other search areas.
- (6) Much progress in other search areas has not found its way yet into standard enterprise search applications.
- (7) Given the controlled environment of an enterprise there is much scope to utilise user interactions with the search system to improve search and exploration for individuals or groups of users. Little of this is currently being employed.
- (8) There are plenty of research challenges worth exploring academically to push forward the state of the art in enterprise search.
- (9) The reality is that even now enterprise search systems are falling short of expectations.

REFERENCES

- [1] Udo Kruschwitz and Charlie Hull. 2017. Searching the Enterprise. *Foundations and Trends® in Information Retrieval* 11, 1 (2017), 1–142. <https://doi.org/10.1561/15000000053>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

DIR'17, November 2017, Hilversum

© 2016 Copyright held by the owner/author(s).

ACM ISBN 123-4567-24-567/08/06...\$15.00

https://doi.org/10.475/123_4