Credibility in Search Systems via Information Retrieval theory

Mihai Lupu
Vienna University of Technology
Favoriten Strasse 9-11/188
Vienna 1040, Austria
<http://www.mihailupu.net>
lupu@ifs.tuwien.ac.at

Information Retrieval methods, and search systems in general, are now a significant part of our interaction with the world outside of our direct reach. As such, they influence the way we perceive the reality we cannot directly experience. From here, the issue of credibility of IA systems is raised, and in this poster we look at the work which has already been done, observe its limits and propose some directions for future work in this area.

Credibility, Trust, Information Retrieval, Information Access

1. INTRODUCTION

A number of studies deal with credibility of the data—web pages, answers, tags. This is a considerable and challenging problem, and systems and methods have been proposed to help the users in assessing the credibility of the information they receive (Schwarz and Morris (2011)). Such studies assume the IR method and search system in general to be an independent, impartial, *trustworthy* intermediary. Even if such were the case, the user may still be entitled to mistrust a, for the lack of a better word, incompetent system. In our experience, this comes up first in professional search scenarios (Patent, Legal, Medical), but there is no reason to stop here, since we all take search systems for granted in many of our daily interactions with knowledge.

This problem is not specific to information access systems, nor even only to computer systems, and as such, the issue of credibility of computer systems is not new in itself. In the following we will very briefly summarize the work which has already been done and propose some directions for future research in IR models and practice.

2. PRIOR WORK

The issues of how a human can trust a system have been studied for other types of computer systems (see for instance Galletta et al. (2005) and Lai et al. (2011) for spell-checkers and internet-based inter

organizational systems, respectively) and in general in the more humanistic literature (Kiran and Verbeek (2010); Taddeo (2010)), but less so for information retrieval engines. However, before proceeding, we should provide a definition for our understanding of *credibility*.

The vast majority of researchers identify two components of credibility: *trustworthiness* and *expertise* (Fogg and Tseng (1999)). In a general context, *trustworthiness* is unbiased, truthful, well intentioned, while *expertise* is knowledgeable, experienced, or competent. For IR engines and systems, *trustworthiness* reflects the perception of the user that the search system is not filtering out potentially desirable results (e.g. censorship) or biasing the results according to an unknown agenda (e.g. hidden advertisement). *Expertise* for IR systems is, on one hand, market popularity (less interesting for us) and, on the other hand, effectiveness and efficiency evaluations to the extent to which these help build confidence in the quality of the search system.

In fact, a lot of the work already done in IR can be casted as a conveyor of credibility in the performance of the system. Here are a few IR research areas and how they can be viewed in terms of credibility. To save space and because many of these are well-known research areas in our field, we have refrained from using citations, except in cases where a specific point was to be made.

Probabilistic retrieval together with methods to re-generate probabilities of relevance from retrieval status values (e.g. Nottelmann and Fuhr (2003)) are methods to convey to the user more than the set of most relevant documents, but also *how relevant* these most relevant documents are.

Automatic explanations have appeared mostly in Question-Answering systems, but recent work has also looked at recommender systems and attempted to provide explanations based on text similarity (Blanco et al. (2012)).

Diversity can be viewed as the answer to the need of the user to explore the entire knowledge space before making a decision.

Findability looks at the core ability of an engine to retrieve documents and as such can be viewed as a measure of expertise, as well as trustworthiness (if a set of documents is intentionally not retrievable).

Evaluation campaigns are a direct measure of expertise of an IR method at a specific task.

Human-computer interfaces study the objective and subjective ways in which the presentation of the results affects the, ultimately subjective, perception of credibility.

All this work can be cast as credibility, but is only a series of proxies, each independently looking at a different aspect, providing a fragmented image of credibility of IA systems.

3. PROPOSED DIRECTIONS

There are many directions starting from here. Here are some that come first to mind.

Back to probabilities Use the now pervasive big data to improve the probabilistic model. A number of assumptions and simplifications have been made in the original models, which have subsequently been further refined in countless articles, but ultimately still with the goal of providing the set of most relevant documents, rather than a precise probability of how relevant each one is.

New benchmarks While the evaluation campaign benchmarks are, as we said, an indicator of expertise and therefore credibility, developments of new IR results are linked to the existence of specific test collections and metrics.

Automatic explanations are already present in some form (e.g. text snippets) or for specific (linked) data, but what is now left to the user to assess, could be done also automatically by the system.

Consistency assessment refers to the experience that users tend to adapt to a system and consider it reliable even if it has known and considerable weaknesses, as long as those weaknesses are consistent. Furthermore, an inconsistent result may be an indicator of censorship or hidden advertisement.

User studies are ultimately needed because credibility is essentially a subjective assessment. This will have to include, or in some way factor out, the human-computer interface.

4. CONCLUSION

The assessment of the expertise (quality of results) and trustworthiness (impartiality of results) of a search system is performed constantly by each user, either consciously or unconsciously. The question we have is to what extent, and in which way the underlying IR method can be used to assist in this evaluation. There is no answer at the time, but hopefully, the set of research directions proposed here would bring us closer to one.

REFERENCES

- Blanco, R., Ceccarelli, D., Lucchese, C., Perego, R. and Silvestri, F. (2012), You should read this! let me explain you why: explaining news recommendations to users, *in* 'Proc. of CIKM'.
- Fogg, B. J. and Tseng, H. (1999), The elements of computer credibility, *in* 'Proc. of SIGCHI'.
- Galletta, D. F., Durcikova, A., Everard, A. and Jones, B. M. (2005), 'Does spell-checking software need a warning label?', *Commun. ACM* **48**(7).
- Kiran, A. and Verbeek, P.-P. (2010), 'Trusting Our Selves to Technology', *Knowledge, Technology & Policy* **23**(3-4).
- Lai, I. K. W., Tong, V. W. L. and Lai, D. C. F. (2011), 'Trust factors influencing the adoption of internet-based interorganizational systems', *Electron. Commer. Rec. Appl.* **10**(1).
- Nottelmann, H. and Fuhr, N. (2003), 'From retrieval status values to probabilities of relevance for advanced ir applications', *Information Retrieval* **6**.
- Schwarz, J. and Morris, M. (2011), Augmenting web pages and search results to support credibility assessment, *in* 'Proc. of SIGCHI'.
- Taddeo, M. (2010), 'Trust in Technology: A Distinctive and a Problematic Relation', *Knowledge, Technology & Policy* **24**(3-4).