

Modelling Search and Stopping in Interactive Information Retrieval

David Maxwell
University of Glasgow, Scotland
maxwelld90@acm.org

Abstract

Searching for information when using a computerised retrieval system is a complex and inherently interactive process. Individuals during a search session may issue multiple queries, and examine a varying number of result summaries and documents per query. Searchers must also decide when to stop assessing content for relevance – or decide when to stop their search session altogether. Despite being such a fundamental activity, only a limited number of studies have explored stopping behaviours in detail, with a majority reporting that searchers stop because they decide that what they have found feels “*good enough*”. Notwithstanding the limited exploration of stopping during search, the phenomenon is central to the study of Information Retrieval, playing a role in the models and measures that we employ. However, the current *de facto* assumption considers that searchers will examine k documents – examining up to a *fixed depth*.

In this thesis, we examine searcher stopping behaviours under a number of different search contexts. We conduct and report on two user studies, examining how result summary lengths and a variation of search tasks and goals affect such behaviours. Interaction data from these studies are then used to ground extensive *simulations of interaction*, exploring a number of different *stopping heuristics* (operationalised as twelve *stopping strategies*). We consider how well the proposed strategies perform and match up with real-world stopping behaviours. As part of our contribution, we also propose the *Complex Searcher Model*, a high-level conceptual searcher model that encodes stopping behaviours at different points throughout the search process (see Figure 1 below). Within the Complex Searcher Model, we also propose a new results page stopping decision point. From this new stopping decision point, searchers can obtain an impression of the page before deciding to *enter* or *abandon* it.

Results presented and discussed demonstrate that searchers employ a range of different stopping strategies, with no strategy standing out in terms of performance and approximations offered. Stopping behaviours are clearly not fixed, but are rather *adaptive* in nature. This complex picture reinforces the idea that modelling stopping behaviour is difficult. However, simplistic stopping strategies do offer good performance and approximations, such as the *frustration*-based stopping strategy. This strategy considers a searcher’s tolerance to non-relevance. We also find that *combination* strategies – such as those combining a searcher’s *satisfaction* with finding relevant material, and their frustration towards observing non-relevant material – also consistently offer good approximations and performance. In addition, we also demonstrate that the inclusion of the additional stopping decision point

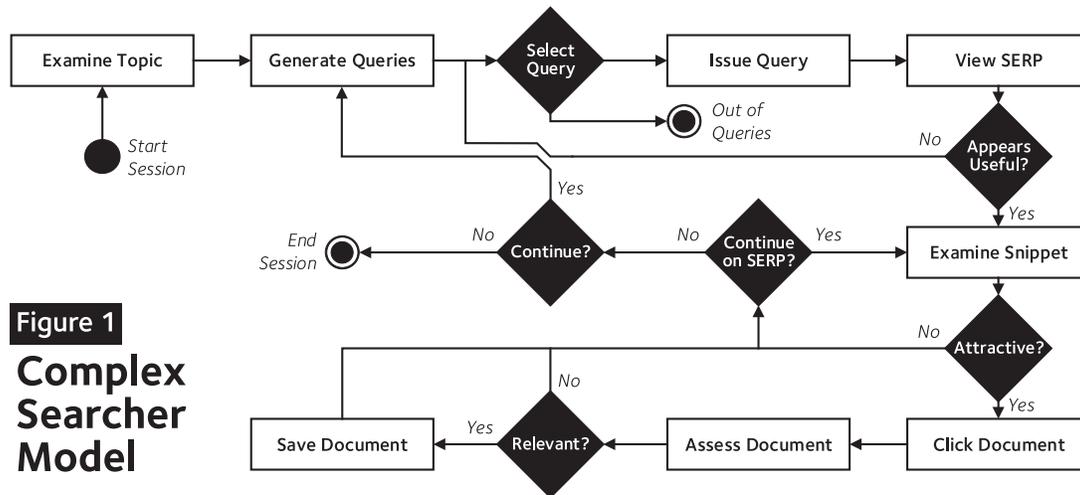


Figure 1
Complex Searcher Model

within the Complex Searcher Model provides significant improvements to performance over our baseline implementation. It also offers improvements to the approximations of real-world searcher stopping behaviours.

This work motivates a revision of how we currently model the search process and demonstrates that different stopping heuristics need to be considered within the models and measures that we use in Information Retrieval. Measures should be reformed according to the stopping behaviours of searchers. A number of potential avenues for future exploration can also be considered, such as modelling the stopping behaviours of searchers individually (rather than as a population), and to explore and consider a wider variety of different stopping heuristics under different search contexts. Despite the inherently difficult task that understanding and modelling the stopping behaviours of searchers represents, potential benefits of further exploration in this area will undoubtedly aid the searchers of future retrieval systems – with further work bringing about improved interfaces and experiences.

Doctoral Supervisor Dr Leif Azzopardi (University of Strathclyde, Scotland)

Examination Committee Professor Iadh Ounis (University of Glasgow, Scotland) and Dr Suzan Verberne (Leiden University, The Netherlands). *Thanks to both of you for your insightful and fair questioning during the defence!*

Availability This thesis is available to download from <http://www.dmax.org.uk/thesis/>, or the University of Glasgow’s *Enlighten* repository – see <http://theses.gla.ac.uk/41132/>.

A Quick Thank You Five years of hard work has got me to the point at which I can now submit the abstract of my doctoral thesis to the *SIGIR Forum*. There have been plenty of ups and downs, but I’m super pleased with the result! Even though there is only a single name on the front cover of this thesis, there are many people who have helped me get to where I am today. You all know who you are – from my friends and family, those who granted me so many fantastic opportunities to travel and see the world – and of course, to Leif. *Thanks to all of you for confiding your belief and trust in me, even when I may have momentarily lost that belief and trust in myself. This thesis is for you all.*