Search Literacy: Learning to Search to Learn

Max L. Wilson[1], Chaoyu Ye[1], Michael B. Twidale[2], Hannah Grasse[2],
Jacob Rosenthal[2], Max McKittrick[2]

[1]Mixed Reality Lab
School of Computer Science
University of Nottingham, UK
[max.wilson, psxcy1]@nottingham.ac.uk

[2]School of Information Sciences
University of Illinois at Urbana-Champaign
USA

ABSTRACT

People can often find themselves out of their depth when they face knowledge-based problems, such as faulty technology, or medical concerns. This can also happen in everyday domains that users are simply inexperienced with, like cooking. These are common exploratory search conditions, where users don't quite know enough about the domain to know if they are submitting a good query, nor if the results directly resolve their need or can be translated to do so. In such situations, people turn to their friends for help, or to forums like StackOverflow, so that someone can *explain* things to them and translate information to their specific need. This short paper describes work-in-progress within a Google-funded project focusing on Search Literacy in these situations, where improved search skills will help users to learn as they search, to search better, and to better comprehend the results. Focusing on the technology-problem domain, we present initial results from a qualitative study of questions asked and answers given in StackOverflow, and present plans for designing search engine support to help searchers learn as they search.

CCS Concepts

• **Information systems → Information retrieval → Users and interactive retrieval → Search Interfaces.** • **Information systems → World Wide Web → Web searching and information discovery.**

Keywords

Search Literacy, Search User Interfaces, Information Seeking.

1. INTRODUCTION

While there are many facets that create different kinds of exploratory search situations [18], and even less task-oriented casual-leisure situations [6], Exploratory Search was originally characterized as occurring when users are 1) unfamiliar with their domain, 2) unsure of which words to use, and 3) unable to judge the usefulness of results [17]. This work aims to study how *Search Literacy* helps to make progress within such confusing search situations. To do this, we focus on searchers trying to solve "tech problems", where they are likely to experience all three Exploratory Search characterizations – they don't really understand the technology and may not really know what the

underlying problem is, they may not know the correct terminology to describe the problem or to search for a solution, and find it hard to understand if results are relevant. Indeed they may struggle to find a result that explains the solution in a way that they can understand without doing yet more searches.

In the "tech problems" case study domain, we see examples of both domain-novice users struggling to comprehend technical jargon and whether results will help them sort out their problems, but we also see examples of domain-expert users, who fully understand the technical jargon, but are synthesizing or diagnosing more complex or combined technical problems, and are seeking more specific specialized knowledge. Our research is driven by the observation that the behavioral difference between *techies* solving these problems, and *novices*, is that techies use search skills, associated with higher search literacy [7, 15], to resolve the situation: e.g. when they encounter something they don't understand, they resolve the new information needs with supplementary searches.

Regardless of domain expertise, research indicates that searching and learning are often closely interleaved [16]. A person may choose to learn about a technology, tinker with it, get stuck, search online for help, find a resource (such as a tutorial, blogpost or how-to video) - or ask for help in a forum. This can lead to either a solution or further learning goals. As well as searching-as-part-of-learning, a person may also be learning-as-part-of-searching: learning better search skills and information literacy, but also in technical areas, learning how to debug a problem better, how to isolate the cause of the failure they have encountered, how to do better diagnosis of a technological impasse – or of their understanding of that impasse. This project, therefore, aims to observe strong searching skills, in order to design new Search User Interface features [19] that encourage search novices to learn and improve their search literacy.

2. INITIAL STUDY

To examine experiences of solving tech problems, we looked at the questions asked and answers given in StackOverflow (SO), a social collaborative Q&A site (SQA) for technical questions [e.g. 13]. The aim was to look at a venue where technical questions were asked and answered, often quite complex ones and including technically sophisticated question askers and answerers (although also including some novices). We wanted to get a better understanding of what it takes to ask good questions and obtain good answers in a social collaborative setting. However our main focus was less on "how does SO work so well?" and more on the lessons and ideas it might inspire when thinking about how a search engine to help with searching for technical answers. Seeing how humans do it well can be informative, even if we cannot, or do not want to, directly translate the methods to a search agent.

As part of a larger Trace Ethnography [8] investigation, we first looked at our small sample of questions in SO from the perspective of the features that seem to be part of what makes a good question in this setting. We then compared them with what we see in a generic search engine such as Google. Finally we compared social question asking with the well-established and well-documented case of reference librarianship where a designated professional tries to help people with all kinds of questions. These analyses are informing ongoing design ideas for a better search engine, and we note some preliminary ideas.

2.1 Methods

We selected sixty-four questions from Stack Overflow. Special attention was paid to questions that had garnered responses from other users. The topics of the questions varied, but they all related in some way to programming in a range of languages. Topicality was limited to questions that the research team had prior knowledge about (so we could analyze the discussion).

Based on an emergent thematic analysis approach, the questions were coded from three main perspectives:

- *Aspects of the question*. These include points informally characterized as: How-do-I? Is this possible?, My main goal is X, and so I am trying to do Y, In particular, what I really want is…, Please recommend X, Why is this the case?, Here's a weird thing, Is X even possible?, etc.
- *Supplemental Information.* These include specific examples, code fragments, URLs and images, as well as what might be termed "due diligence" - what the question asker has already tried, how that failed, places looked for information, etc.
- *Tone of the question.* These include how the question was asked, the formality of phrasing, whether more of a narrative, whether particularly focussed, identifying background as a newbie or expert, or issues of question-asking etiquette.

2.2 Selected findings from question analysis

Both search engine use and SQA have certain features in common: 1) Initial query, 2) Results, ranked somehow, 3) Selection of those to attend to, 4) Assessment of quality and relevance, 5) Query refinement and iteration. With SO, however, there is of course a human in the loop. Indeed several humans and at each stage of the Information Seeking Process (ISP). This makes aspects of the ISP much more visible and helps to consider not just how it operates in this particular case, but how it might operate otherwise in other cases. Although do not have space to present our three full taxonomies in this paper, we now present several initial findings.

There are various features that seem to help make a good question – one that is easier for others to answer and indeed invites others to answer. These features are articulated in various kinds of advice given about best practices on SO, which are embodied in an established etiquette in SO usage. Examples of features that make a good question include context: the technical setup that the person was using, and the overall goal of what the person was actually trying to do that led to the particular question that was asked. Some kind of due diligence information is common. This can include what the asker tried in order to solve the problem herself (and what resulted and why it was not helpful, thereby necessitating this help request), mention of search attempts to try and find an answer and diagnostic activities to try and simplify down the initial problem to isolate the underlying cause.

One feature we found particularly interesting was not merely question refinement as an iterative activity (analogous to iterated queries in a search engine) but also question editing – both by the original asker and by other participants. The formulation of the question is considered important work in SO and not to be done in a sloppy manner. The construction of the question title and choice of tags used needs particular care – in order to catch the eye and interest of potential question answerers. These findings are similar to related work [10].

Question askers may occasionally answer their own question – and then take the time to report that back to SO. This and the previous point about question phrasing remind us that SO can also be seen in terms of knowledge management – including how answered questions can potentially be reused by others, whereas we usually think of queries into a search engine as use-once disposable activities. This insight reinforces the value of tracking user journeys of understanding [3], and might inspire developments in ideas like the retired SearchPad [5].

Taken together, these points serve to remind us that a well-posed question is a learnable skill. Learning it is desirable to increase the odds of getting a good response from SO members. It also can help in understanding your own current problem – to the extent that you may be able to solve your problem yourself while you are waiting for a solution from others. Furthermore, learning how to form good questions may help in future tech problems as part of an armory of metacognitive strategies. This is similar to academic research, where asking the right question or the question in the right way is a critical part of the research process and one to be taught to new researchers.

SO has various affordances for learning how to ask good questions. You can learn vicariously by seeing other people's questions as exemplars. Your own question can lead to follow-up clarification questions, or even having your question edited by others as a more immediate indicator of what you should have done and should think about doing next time. The voting mechanism also allows an indication of collective views of what makes a good question, including seeing your own question being up-voted as it is improved. In future work we want to think about how search interfaces (typically rather solitary places) might also facilitate such kinds of learning that occurs almost spontaneously in SQA – given that we suspect that SO was designed and developed with much more attention to giving answers than explicitly facilitating these kinds of learning.

Finally, some question askers explicitly note their level of technical sophistication as a way of indicating the kind of answer they need or would appreciate. An expert typically can manage with a far more terse, abstract and technical answer than a novice. However, expertise is not a single scale. A person may be technically adept but is currently asking about a problem with JavaScript, having never used it before. Although these levels of 'techiness' may be explicitly stated, they can also be inferred from the way the question is phrased, and it seems that this is taken into account in the answers given.

2.3 SQA vs. Reference Librarianship

We re-coded our sample of questions and answers for features identifying similarities and differences from what is seen in reference librarianship (RL). For this we used the expertise of one of the authors, who works as a reference librarian and has studied its theory and practice, e.g. [11]. There were considerably more similarities than differences (twice as many coded items that are similarities). Librarians use a combination of both hard and soft skills – and this creates an interesting lens to look at what happens in SO. Similarities included numerous instruction-like how-to

operations (rather than simply giving a factual answer – known in RL as 'ready reference'). We found considerable rapport-building communication that is a core aspect of the reference interview (RI), despite the fact that SO guidelines discourage the use of this kind of communication.

Half the threads included some kind of clarification question – a recurrent aspect of the RI where it is often about trying to understand the underlying goals of what the patron really wants to do and how that may contrast with what they have currently actually asked. This is something that basic search engines do not support – and with good reason. It is extremely hard to do well and requires considerable sensitivity on the part of the RL, because sometimes the patron may not want to answer questions like: "Why are you asking about that?", "What do you want to do with it?" or: "What are you actually trying to do?".

As noted, due diligence activities occur as ways of noting what the asker has tried so far. Other categories include whether the answer was complex or a simple fact-based answer, and whether external resources (documents or people) were pointed to. Complex answers including pointers to resources that can enable additional kinds of learning. They may include the answer to the particular question asked, but also additional insights, larger framings and generalizations. This kind of enrichment learning/teaching is a feature of reference interactions that are in that context termed 'research' rather than 'ready reference'.

Reference librarians are discouraged from offering judgments or opinions; instead they are meant to focus on 'just the facts' and providing access to information resources. Similarly SO guidelines discourage requests for opinions and recommendations. Despite this we do find a significant minority of recommendation requests, especially for choosing between alternative approaches.

3. PROPOSED DESIGNS

The aim of the first period of work, described above, was to create design implications. Although the results and design implications are still underway, we present three initial designs below.

3.1 Design 1: Elaborating the Detail

The common technique seen in 'good questions' on SO, is to make sure important contextual information is provided upfront. This is to help answers estimate Common Ground in understanding [4]. Our first initial prototype explores the possibilities of a search interface identifying the question type and prompting for detail on the right answer.
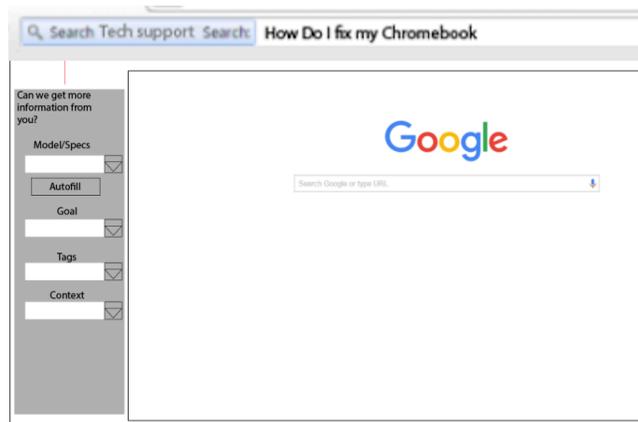

Figure 1: A key factor of strong questions in StackOverflow is to enter all the valuable detail to set the context of the problem

As indicated in Figure 1, browsers could auto-detect some of this information and offer it to the users. For example, the browser can detect the operating system and version being used. This can only be indicative since searchers may be searching for tech problems experienced on other devices. The key challenge here is to iteratively discover the correct boxes to suggest, which becomes harder once we consider supporting many different search problems. In some kinds of remote library reference, patrons are asked to fill out an online form which can help to structure the interaction, and at least help the patron to consider providing information that may help the librarian give the best advice.

3.2 Design 2: Eliciting through Dialogue

With the increasing predominance of spoken interfaces like Google Now, Siri, and Cortana, spoken search [9] presents an opportunity for providing tech support through dialogue. Although we have come a long way since automated support like Clippy, there are still many open challenges with Spoken Search, like experiencing an error midway through a multi-stage interaction [12]. However, spoken search presents a new opportunity to learn from and re-engage the ideas of Search Intermediaries [14] like Reference Librarians. Studies of Search Intermediaries led to Dialogue oriented systems in the 90s [2], which may now have new relevance in spoken search. With dialogue based interaction it becomes less critical for the question asker to provide all the contextual information up-front (desirable in SO) and indeed to need to know what that contextual information may be. Figure 2 shows a visual alternative to this scenario for non-spoken on-screen dialogues.
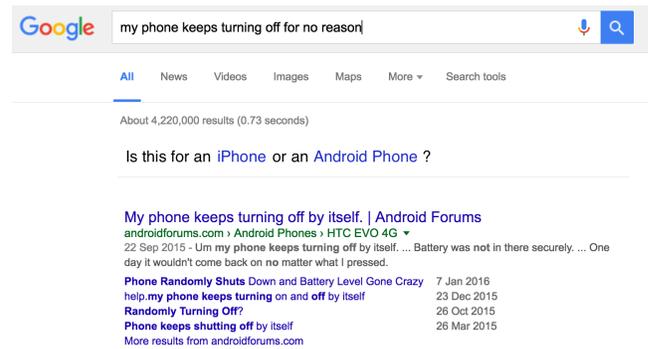

Figure 2: For some StackOverflow questions, a back and forth conversation is needed to identify the right answer
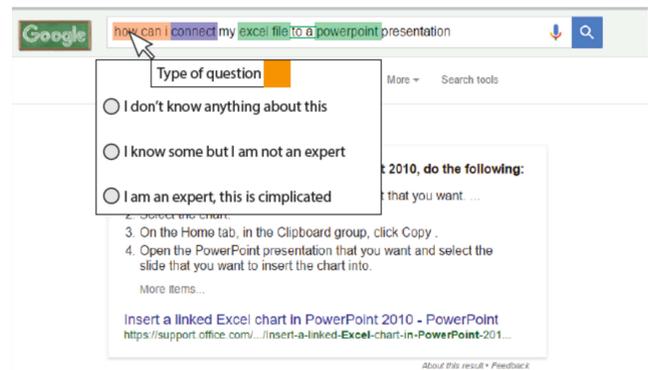

Figure 3: Sometimes users do not understand the jargon and its implications, in either the query or the results

3.3 Design 3: Exploring the Definitions

A key problem experienced by users is that they don't recognize whether items are object-specific or generic terms, and thus whether advice specifically relates to their situation, or does in principle. One way to improve search literacy would be to help users to interrogate key words or phrases. These might be to find out more about query terms, or terms in the SERPs, supporting what Bates called the TRACE tactic [1]. As shown in Figure 3, a key element of this idea might be to set of vary the existing knowledge of the user, which would help with the 'levels of techiness' problem, where those levels affect the kind of answer that you may want. Another key element of this idea is to remain within the context of the search, but be able to interrogate and explore concepts returned in the *results*, rather than queries, in situ to develop confidence in the results.

4. CONCLUSIONS & FUTURE WORK

The aim of this on-going project is to investigate Search Literacy in situations where users become easily confused within their search domain. Solving technical problems leads to many kinds of interleaved search and learning. You may be choosing to learn a technology and use search as part of your learning activities. Or you may be searching, hoping to find a simple fact-like answer, but along the way learn other things including how to search better and how to go about learning other technologies better. This might be because they are domain novices, or they may even be domain-experts who just lack the specific knowledge they need. It may be that technology related problems have features that make certain kinds of search hard, not because technology problems are unique but that they exacerbate issues seen in many other settings. These include multiple kinds of expertise, problems with not knowing terminology, confusing interactions with other technologies and large amounts of prerequisite knowledge you may not yet have, or see as related.

Our on-going work has begun to produce three taxonomies of how users overcome these barriers by asking questions within the StackOverflow community. Initial analysis of this taxonomy has led to the identification of a few initial design ideas, presented above, that might help users to improve their search literacy and make progress even in confusing circumstances.

Future work will focus on elaborating on these design ideas and producing further ideas. Refining these designs will be important, as keeping them within the familiar experiences of Google will help study their potential benefits. Initial work has also begun into building Chrome extensions to customize the appearance of Google pages, in order to deploy final prototypes with real users. Subsequent observational and experimental studies will attempt to investigate the benefits of each design idea.

5. ACKNOWLEDGMENTS

6. REFERENCES

[1] Maric J, Bates. (1979), Information search tactics. J. Am. Soc. Inf. Sci., 30: 205–214.

[2] Nicholas J. Belkin, Colleen Cool, Adelheit Stein, and Ulrich Thiel. (1995) Cases, scripts, and information-seeking strategies: On the design of interactive information retrieval systems. *Expert systems with applications* 9, no. 3: 379-395.

[3] Mikhail Bilenko and Ryen W. White. (2008) Mining the search trails of surfing crowds: identifying relevant websites from user activity. In *Proc. WWW '08*. ACM, New York, NY, USA, 51-60.

[4] Clark, Herbert H., and Susan E. Brennan. "Grounding in communication."*Perspectives on socially shared cognition* 13.1991 (1991): 127-149.

[5] Debora Donato, Francesco Bonchi, Tom Chi, and Yoelle Maarek. (2010) Do you want to take notes?: identifying research missions in Yahoo! search pad. In *Proc. WWW '10*. ACM, New York, NY, USA, 321-330.

[6] David Elsweiler, Max L. Wilson, and Brian Kirkegaard Lunn (2011), Chapter 9 Understanding Casual-Leisure Information Behaviour, in Amanda Spink, Jannica Heinström (ed.) *New Directions in Information Behaviour,* Emerald Group Publishing Limited, pp.211 – 241

[7] Susan R. Goldman. (2010). Literacy in the digital world: Comprehending and learning from multiple sources. In M.G. McKeown & L. Kucan (Eds.), Bringing reading research to life (pp. 257–284). New York: Guilford.

[8] R. Stuart Geiger and David Ribes, (2011), Trace Ethnography: Following Coordination through Documentary Practices, *System Sciences (HICSS'11)*, Kauai, HI, pp. 1-10.

[9] Larry P. Heck, Dilek Hakkani-Tür, Madhu Chinthakunta, Gökhan Tür, Rukmini Iyer, Partha Parthasarathy, Lisa Stifelman, Elizabeth Shriberg, and Ashley Fidler. (2013). Multi-Modal Conversational Search and Browse. In *SLAM@ INTERSPEECH*, pp. 96-101.

[10] Jeon, Grace YoungJoo, and Soo Young Rieh. "Social search behavior in a social Q&A service: Goals, strategies, and outcomes." *Proc ASIST 2015,* 52.1 (2015): 1-10.

[11] William A. Katz, (2002) *Introduction to Reference Work.* Boston: McGraw-Hill,.

[12] Julia Kiseleva, Kyle Williams, Jiepu Jiang, Ahmed Hassan Awadallah, Aidan C. Crook, Imed Zitouni, and Tasos Anastasakos, (2016). Understanding User Satisfaction with Intelligent Assistants. In *Proc. CHIIR '16*. 121-130.

[13] Lena Mamykina, Bella Manoim, Manas Mittal, George Hripcsak, and Björn Hartmann. (2011) Design lessons from the fastest Q&A site in the west. In *Proc CHI'11* 2857-2866.

[14] Richard S. Marcus (1983). An experimental comparison of the effectiveness of computers and humans as search intermediaries. *J. Am. Soc. Inf. Sci.*, 34: 381–404.

[15] Daniel Russell (2015). Mindtools: what does it mean to be literate in the age of Google? *J. Comput. Sci. Coll.* 30(3) 5-6.

[16] Pertti Vakkari (2016). Searching as learning: A systematization based on literature. *Journal of Information Science.* 42: 7-18.

[17] Ryen W. White, Bill Kules, Steven M. Drucker, and m.c. schraefel. (2006). Introduction. *Comm. ACM* 49 (4), 36-39.

[18] Barbara M. Wildemuth and Luanne Freund. (2012). Assigning search tasks designed to elicit exploratory search behaviors. In *Proc HCIR '12*. ACM, New York, NY, USA,

[19] Max L. Wilson (2011). "Search user interface design." *Synthesis lectures on information concepts, retrieval, and services* 3.3: 1-143.