Search or Explore:
Do you know what you're looking for?

Jon Pearce, Shanton Chang, Basil Alzougool,
Gregor Kennedy, Mary Ainley
The University of Melbourne
j.pearce@unimelb.edu.au

Susan Rodrigues
University of Northumbria, UK
susan.rodrigues@northumbria.ac.uk

ABSTRACT

This paper explores the distinctions between *searching* and *exploring* when looking for information. We propose that, while traditional search engines work well in supporting search behaviour, they are more limited in assisting those who are looking to explore new information, especially when the exploration task is ill-defined. We ran a pilot study using two systems: one based on a traditional database search engine, and the other – a highly innovative, engaging and playful system called iFISH – that we designed specifically to support exploration through the use of user preferences. We looked for evidence to support the concept that exploration requires a different kind of interaction. The initial results report a positive response to our exploration system and indicate the differences in preferences amongst users for systems that match their searching or exploring behaviours.

Author Keywords

Exploration engine; engagement; flow; preference-based search; recommender system; search engine

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

We live in a world with an abundance of opportunities to search for information – type 'London Olympics' into Google and you will quickly find information about the games. However, try to find a cozy restaurant for that romantic evening: one that has old-world charm, a fusion menu with plenty of spicy foods. It's not so easy – and you would probably want to explore several with similar qualities before you made your choice.

Search engines are very good if you use specific terms that bring to the top of their output ranking a set of highly relevant choices, but they are not well equipped to let you express your feelings about the kind of places you would like to explore – to let you express your taste. They are especially ineffective when exploring unfamiliar domains – domains in which you know the kind of things you like, but may not be able to articulate them using terms that would return useful results in a traditional search engine (Pu, Faltings, & Torrens, 2003).

This paper seeks to explore some of the differences between "searching" and "exploring" and addresses the issue of helping people explore constrained data sets that lend themselves to exploration and immersion, as opposed to a quick search and then moving on.

We discuss the design of an 'exploration engine' – an online system that allows people to explore domains rather than traditional keyword dependent search engines. In such an exploration engine the goal for the user is one of exploring rather than searching; that is, exploratory search rather than data retrieval (Marchionini, 2006). While search engines aim to find the "best" literal match item in the shortest time, one aim of our proposed system is to extend the time so that the user can develop a sense of engagement, and an opportunity for reflection.

We utilise the idea that the drivers for a user's exploration may often be based on the user's personal preferences for attributes that do not necessarily relate directly to the concept being explored. We play with the idea of users' exploration being driven by their expression of their preferences using non-judgemental dichotomies (e.g. 'casual vs formal', rather than 'high quality vs low quality'). We refer to these as 'user-preferences' or 'tastes'. We invoke the concept of playfulness into such a system through the application of a 'dynamic query' interface (Shneiderman & Plaisant, 2010) using direct-action sliders and animated motion on the screen.

The design of our research tool was driven by a desire to gain a better understanding of two main issues not well catered for by current search environments. The first of these issues relates to user engagement in situations in which the user may not be able to draw on the appropriate language or search terms to adequately narrow a search – where traditional search techniques may be inadequate and a more exploratory approach appreciated. This is particularly pertinent when users do not have a clear idea of what they want (Kalinov, Stantic & Sattar, 2010) or have ill-defined search requirements (MacMullin & Taylor, 1984). Our approach to this has been to construct an environment in which users can explore based on preferences that provide a meta-level (indirect) relationship to the data.

The second issue relates to ways of enticing the user to explore a set of data. The literature on 'flow' (Csikszentmihalyi, 1975) and 'play' (Rieber & Matzko, 2001) suggest that the notions of exploration, engagement and enjoyment can be utilised to produce a powerful exploratory environment. We have tried to capture these traits in an environment that is fun and easy to explore in a playful manner.

The paper is structured as follows. The next section provides a brief literature review on some key differences between searching and exploratory behaviours. We then briefly describe the two systems developed (one based on traditional search engine and the other on an explaoratory engine) to show differences in the behaviours.

PREVIOUS WORK

Whereas there is a vast quantity of literature on techniques and software to support the *search* process, there has been little research that examines the way people look for information in an *exploratory* manner.

Often reflection and reformulation can become a significant part of an exploratory process, especially in cases in which the needs of the user are not well specified. For example, whereas a web database of restaurants might respond well to queries about location, prices, opening hours and menus, it is unlikely to be able to handle the vague probings of someone who also desires a romantic night out in a quiet, yet funky, environment! This can be described as an ill-defined, or ill-structured, question and is defined as one that cannot be resolved through strictly analytical means (MacMullin & Taylor, 1984).

One approach to such problems is to give the user significant opportunity to interact and reflect as part of the search process. Kalinov et al. also note the limitations of traditional web search engines commenting on their use for research as opposed to search (Kalinov, Stantic, & Sattar, 2010). They acknowledge that whilst search engines satisfy the information locating need, they do not satisfy the information discovery need. They have built a 'web exploration engine' to facilitate both searching and exploring that uses a web spider to collect data which is then manipulated based on a combination of statistical learning and human classification. This is supported by the Alzougool, Gray and Chang's (2009) findings in their work on information needs that people have both recognised needs, which are well articulated and unrecognised needs which are either undefined or just a feeling that some information is missing.

An interesting case of exploratory behaviour is when the user is exploring based on aspects of taste (in the generic sense of the word, not the culinary sense!). One approach is to use a recommender system. These systems are generally classified by three categories of filtering: content-based, collaborative, and hybrid (Adomavicius & Tuzhilin, 2005). Collaborative filtering has been shown to work well in the area of accommodating taste, for example, for music and movies (Shardanand & Maes, 1995). However, whilst collaborative filtering sites such as MovieLens (www.movielens.org) are very successful at helping people find a list of movies that they are likely to enjoy, they do not facilitate someone who wishes to explore a range of movies and discover ones to investigate further (e.g. "I want to see movies that are deeply romantic and intellectually lightweight, but let me also see what happens if I am prepared to consider movies with increasing levels of intellect."). Collaborative sites like MovieLens tend to focus on one dimension ('like') rather than accommodating a set of

tastes (such as 'romance', 'intellect', 'pace', etc.). They commonly provide the user with the results of a search query, rather than being immersed in an exploration.

Some suggest recommender systems cannot effectively search for possible solutions without an accurate model of users' preferences (Pu et al., 2003). Not surprisingly, many researchers are now identifying preference elicitation as a key and complex issue, and are generating a body of research on various recommender system developments. However, there still remains the issues identified by Pu et al. (2003) regarding difficulties faced by users who suggest a need for "mixed initiative" systems that ensure adequate turn-taking with shared control between the system and user. Some such issues, in particular direct manipulation, shared control and affordances for collaboration, are ones that we address in our proposed system.

In summary, whilst there has been much research in the past two or three decades into search, exploration and recommender systems, we still do not have a good understanding of systems that support users in carrying out more exploratory behaviours.

METHODOLOGY

In a pilot study, we used the context of searching for a restaurant to visit. We developed two systems: one based on traditional database search engines that rely on a Boolean match of particular search parameters, and another system that was more exploratory and carried out a pattern matching process of the user's preferences for the same parameters. Figure 1 is a screen-shot of the search engine (a FileMaker database) and Figure 2 is the exploratory engine (iFISH, developed in-house).
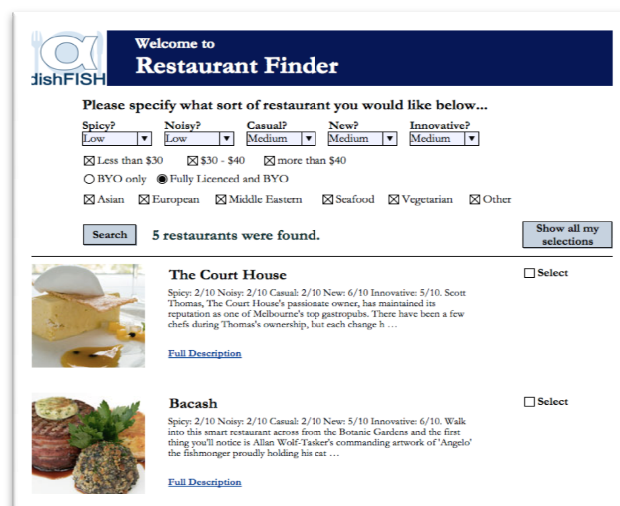

Figure 1. Restaurant Finder Search Engine

The systems behaved as follows. Each system allowed the user to enter five terms relating to their desired experience (spicy, noisy, casual, new, innovative) as well as some other checkbox filter terms relating to food types, wine, and cost. The search engine returned a list of restaurants that matched the criteria. In order to ensure sufficient hits, each criterion offered three selections (low, medium, high) representing fairly wide ranges of

numeric values. In contrast, the exploratory system ('iFISH') allowed users to make adjustments to five sliders relating to the same terms. This caused an immediate animated shuffling of restaurant images on the left of the screen, resulting in a display of restaurants ranked by how well each restaurant matched the pattern of slider positions.

Figure 2. iFISH Exploratory Engine

A fuller description of the precursor to the iFISH system can be found elsewhere (Pearce & Pardo, 2009).

We set two main tasks for a group of 23 postgraduate students, aged 24 to 34 (av. 30); 12 male, 11 female.

In Task 1, the user has a clear idea of what they are looking for. It specified: "*I am looking for a restaurant that is reasonably well-established, definitely quiet and romantic, not too formal, not too casual either. I would like it to have a choice of traditional, yet pretty hot and spicy food. I like all kinds of food, but I don't want it to be too expensive!*" We argue that Task 1 is more 'search oriented' where the user knows what they are looking for.

In Task 2, the user is less clear about the type of restaurant he/she is looking for and therefore needs to explore the options in more depth: "*You're looking to organise a night out for a small group of your friends who believe you are a bit of an expert when it comes to food. So, they have left the choice up to you*"

We let the participants spend some time using both the systems to look for 3 restaurants they might recommend within each of the tasks we set them. We then asked them about their experience and satisfaction with the results each of the systems provided in assisting them with their recommendation. The initial interview results after using the systems were then analysed for common themes.

INITIAL FINDINGS
The initial findings from the interview data showed that overall, for both Task 1 and Task 2, there was a preference for using the iFISH. However, the preference was deemed to be stronger in Task 2 than in Task 1. For example, there was a clearer preference for using the exploratory engine for Task 2 due to the task's less well-defined nature, as indicated by one participant:

"*This time (in Task 2) because it was a bit more indecisive and you try to please a general audience, so (I) spend more time to browse longer, but it's still much better than (using) the database (search engine).*"

In Task 2, participants indicated that the browsing behavior enabled by the exploratory engine met undefined needs better:

"*Browsing was more important. I did not have a specific thing in mind. So, I could not return the specific hit I needed. It is really more a case of window shop to find something.*"

On the other hand, the reactions to using either of the systems for Task 1 (with a clearer criteria) were mixed, as indicated by the following participant:

"*So for Task 1, I found it was just fine and was probably the same (using either system) in terms of finding the places giving the set criteria.*"

Moving away from the nature of the tasks, participants also discussed some of the advantages of each of the systems more generally. In particular, the 'playfulness' and interactivity of the exploratory engine appealed to many participants:

"*I like the dynamic side of things. You change the sliders (to) give you the results straight away. You can see these ones (the restaurants) are coming up to the top. Yes, I think the dynamic is what I liked the most.*"

One advantage that the search engine did have over iFISH was the sense that the search engine was something that was more familiar to users:

"*It looks familiar, like list(ing) things down the page and presume the highest match at the top, so it is very familiar to me...*" and also:

"*I would say, given my previous experience, I enjoyed using the database more because I can achieve the same meaningful outcome...*"

DISCUSSION
We live in a time where we frequently seek information and our first response is often "Google it!". In this paper we have reported on some preliminary results of a research project that examines the behaviour of 'searching' in comparison with 'exploring'. For a few participants in the study, given the choice of two search/explore systems, the familiarity with a conventional search engine dominated their preference, whether the task they were performing was essentially a searching one or an exploratory one. The entry of a few key words, skimming through the resulting list of 'hits', reading about each one and them, and then choosing one, made this approach to the task the preferable one.

However, for most of the participants, a traditional search engine was a less than satisfactory experience, especially for the less-well defined task. We had given the participants a real task of the kind we often have when required to make a decision. The vagueness of the task raises many questions: what sort of foods do my friends like? do we want a jovial environment? should we go

somewhere new? will they want to dress-up or be casual? Search engines don't handle these kinds of enquiry well, partly because the information is not encoded in their data. But even if it were, a Boolean 'AND' search *excludes* findings as you add more constraints; and an 'OR' search adds more and more unsuitable ones. Our participants discovered the limitations of the simple 'AND' search as they tried to enter requirements for these tasks in the search engine – it tended to return few, and sometimes no, suitable restaurants.

Our approach to this problem was to design an exploration engine that did not *exclude* findings based on constraints, but *ranked* findings based on personal preferences; it always returned the entire ranked list of restaurants to explore, rather than just a sub-set. We wanted participants not to enter a few quick terms and be returned several 'hits', but rather to take a little longer and explore a range of possible choices that align well with most of their terms, but not necessarily all of them. This is a true iterative exploratory process that requires manipulation of inputs, reflection and examination of the results that are returned, and then a decision as to whether you have found a suitable result or need to make adjustments and explore again. Interestingly, even for the more tightly specified search-oriented task (Task 1), this exploration approach was (slightly) preferred.

An important and powerful aspect of this exploration process was the role of *affect*. We designed the exploration engine to be highly animated and fun to use. As soon as you move a slider a small amount, the coloured images of restaurants on the left of the screen (Figure 2) burst into life and shuffle fluidly around the screen. There is a strong affordance to play; but there is also a strong affordance to probe a little deeper and try to understand the data. We have not yet analysed our log data to look for such exploratory patterns, but will be doing so in due course. However, our casual observations of users discovering these strategies suggest to us that there is a potential educative role in environments such as this one – an encouragement to become a little more engaged and ask questions "why?".

The research from this study is not yet complete, as we are still analysing participant and data logs. The exploration system is a pilot one and, of course, has its limitations. A significant one is that it is the design of the system that determines which exploration criteria (user preferences or tastes) are to be presented for use. This means that for a given context (e.g. restaurants) research must be carried out to determine the best set of criteria. Associated with this is the fact that most web data, and even specialised databases, do not have items tagged with the appropriate data required for this approach. There are various approaches to address this issue, and we are exploring some of these in further research.

We are also studying the use of these techniques in other exploratory areas, such as books, nutrition, wines and university courses.

CONCLUSIONS

This paper has shown that there are differences in searching and exploring behaviours. Understanding the differences between online searching and exploring behaviours can have very practical implications building enhanced online search and explore systems. The differences could have an impact on the ways we choose which search/exploratory engines to use in the future depending on the nature of our information needs and the specificity of what we are looking for.

REFERENCES

Adomavicius, G., & Tuzhilin, A. (2005). Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Transactions on Knowledge and Data Engineering*, *17*(6), 734–749.

Alzougool, B, Gray K & Chang, S (2009), An In-depth Look at an Informal Carer's Information Needs: A Case Study of a Carer of a Diabetic Child. *Electronic Journal of Health Informatics*. 4(1):e2.

Csikszentmihalyi, M. (1975). *Beyond Boredom and Anxiety.* San Francisco: Jossey-Bass Publishers.

Kalinov, P., Stantic, B., & Sattar, A. (2010). Let's Trust Users It is Their Search. *Web Intelligence and Intelligent Agent Technology,* 176-179.

Kalinov, P., Stantic, B., & Sattar, A. (2010). Building a dynamic classifier for large text data collections. *Proceedings of the Twenty-First Australasian Databse Conference*, 113-122.

Marchionini, G. (2006). Exploratory search: from finding to understanding. *Comms of the ACM*, *49*(4), 46.

MacMullin, S. E., & Taylor, R. S. (1984). Problem dimensions and information traits. *The information society*, *3*(1), 91-111.

Pearce, J. M. & Pardo, S. (2009). So now you're ready to play – but with what? A system to encourage playful exploration. Paper presented at the *HCI Educators 2009 Conference*, University of Abertay, Dundee, UK.

Pu, P., Faltings, B., & Torrens, M. (2003). *User-involved preference elicitation*. Proceedings from IJCAI Workshop on Configuration.

Rieber, L. P., & Matzko, M. J. (2001). Serious Design for Serious Play in Physics. *Educational Technology*, *41*(1), 14-24.

Shardanand, U., & Maes, P. (1995). *Social information filtering: algorithms for automating "word of mouth"*. Proceedings from Proceedings of the SIGCHI conference on Human factors in computing systems.

Shneiderman, B., & Plaisant, C. (2010). *Designing the user interface.* Addison-Wesley.