Improving Exploratory Search Interfaces: Adding Value or Information Overload?

Max L. Wilson, m.c. schraefel
School of Electronics and Computer Science
University of Southampton, UK
{mlw05r, mc}@ecs.soton.ac.uk

ABSTRACT

One method for supporting more exploratory forms of search has been to include a compound of new interface features, such as facets, previews, collection points, synchronous communication, and note-taking spaces, within a single search interface. One side effect, however, is that some compounds can be confusing, rather than supportive during search. Faceted browsing, for example, conveys domain terminology and supports rich interaction, but can potentially present an abundance of information. In this paper we focus on the faceted example and conclude with our position that Cognitive Load Theory can be used to estimate and thus manage the potential complexities of adding new features to search interfaces.

INTRODUCTION

The recent interest in supporting more exploratory forms of search [13], for when users are unfamiliar with domain terminology, information sources, or even their own goals, has spurred many new interface design ideas. One method that mSpace, Figure 1, has promoted for supporting a range of directed and exploratory search behaviours, has been to provide a gestalt of interface features [9]. Similarly, the latest version of the Relation Browser has recently extended their range of visualisations and interactions, including the addition of facet clouds [2]. Further, the recent Parallax interface to the Freebase project[1] provides a combination of faceted search, fact views, timelines, and maps to help users explore a wide range of heterogeneous data.

Both the mSpace and Relation Browser interfaces, and many others, provide a user interface with a compound of features, where the aim is for the set of features to work together in synergy in supporting users during search. Conversely, however, Schwartz has discussed the paradox of choice in that often, when users are presented with increasing numbers of options, they make poor or possibly no decisions [10]. In line with Schwartz's findings, many online faceted search websites focus on reducing decision paralysis by presenting only the key facets and their key options at each stage of the user's search [11]. This is most notable when facets, such as those presented by eBay start with a small set of values with a link to see 'more' options.

[1] http://mqlx.com/~david/parallax/ - Freebase Parallax

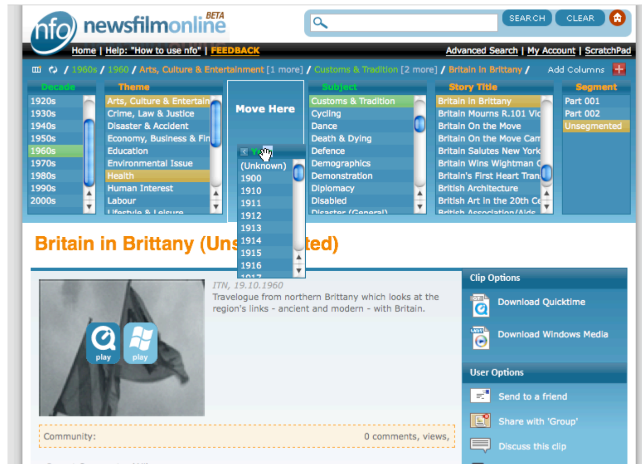

Figure 1: mSpace is a Directional Consistent Faceted Browser.

Evidently, there are two opposing forces that will affect the design of future exploratory search interfaces: 1) enriched functionality and 2) clarity in design. Unfortunately, recent work has also described the difficulties that can be faced when trying to evaluate the proposed advances in exploratory search interfaces [14].

In the next section of this paper we focus this problem by assessing the different approaches taken in providing one type of exploratory search feature: faceted browsing. We identify two dimensions that are present in the different implementations of faceted search and detail both the arguments for and against them. In the latter half of the paper, we propose that Cognitive Load Theory (CLT) [3] can be used to estimate the severity of the expected costs of the different approaches to faceted browsing. Further, the theory can be included into an existing, validated inspection framework [16] so that designs are evaluated for both synergy of features *and* complexity of design.

DEFINING THE DIMENSIONS OF AN EXAMPLE EXPLORATORY FEATURE: FACETED BROWSING

Faceted browsing [5] is an approach to supporting exploratory search that takes a set of meta-data from a corpus and presents the different attributes, and the distinct set of instances from each attribute, to the user. When shopping online for dresses, for example, users may make selections within facets such as price, colour, size, style, and material, to reduce the number of purchasable items. In general, faceted browsing has a number of expected benefits over typical keyword search [5]. One example is that faceted browsing provides users with options to choose from when searching, so that they do not have to guess keyword search terms on their own.

Although various faceted browsers are unified in their aim to provide these expected benefits to exploratory users, there is significant variation in their implementations. In particular, there are two main dimensions that vary in faceted browsers: 1) direction between facets and 2) consistency of display. These dimensions are discussed so that later, their costs can be more concisely understood, explained with CLT, and managed in the future.

Dimension 1: Direction between Facets

Apple's iTunes is an example of a faceted browser that maintains direction between facets. Selecting an Artist filters the list of Albums, but not the Genre column. Like iTunes, most directional faceted browsers present facets in a series of columns across the interface from left to right. mSpace is a directional column browser that has overcome the problem that no Genre associations are shown [15]. Most other instances of faceted browsing, like those on Google product search, Walmart, and eBay, present facets that are unanimously filtered by any selections. Selecting a price range in Google Product Search filters every facet regardless of location of facets on the screen.

The perceived benefit of keeping direction is that additional relationships between facets are clearly shown. In iTunes, selecting a Genre will filter both the Artists and the Albums. Choosing an Artist then filters the Albums, but not the Genres. Now the user sees all the Artists in the selected Genre and all the Albums from the selected Artist. One perceived '"problem"' with maintaining direction is that it can overload the users, as they would have to maintain both a notion of direction, understand the relationships between side-by-side facets, and choose which facet and value to select next to refine their search.

Dimension 2: Consistency of Display

One hypothesis, held by browsers such as Flamenco [17], is that hiding used facets and dedicating screen space to unused facets can minimize information overload. Similarly, browsers often default to show the only the most popular values in a facet to reduce the number of choices. As previous decisions, and their options, are hidden using this method, previous choices are usually placed together as a breadcrumb trail. Another benefit of this approach is that once a user's decision has been hidden, the space can be given to show sub-category options of that selection.

One potential problem with hiding used facets and making space for unmade decisions is that it can be hard to quickly compare multiple items within one facet. In order to compare one style of dress with another, users are required to make an extra step to undo their first action, before making another selection. Further, by hiding used facets, it becomes difficult for a user to make multiple selections within one facet and see the dresses in two or more styles.

The intersection of these Dimensions in Browsers

These two dimensions produce a grid, as shown in Table 1. As noted before, iTunes and mSpace are the two notable examples of faceted browsers that choose to have a direction between facets that affects which are filtered by a selection. Combined with the choice of a consistent layout, these browsers provide: a) inter-facet relationships, b) multiple selections in any facet, c) previous decisions, d) previous selections e) all unused facets and f) a result set.

The remaining browsers listed in Table 1 are all examples that do not employ a direction but allow any facet to be filtered by the facet, and value, chosen by the user. Of these remaining browsers, most also chose to hide the used facets as the users make decisions (Varying layout). As a result, the user neither has to worry about the concept of a direction can choose freely among the facets and only has to consider the facets that remain in view. This combination, however, only provides: a) previous selections b) all unused facets and c) a result set.

Table 1: Examples of Faceted Browsers categorised by Use of Direction and Consistency of Layout

	Consistent Layout	Varying Layout
Directional Filtering	**e.g.** mSpace, iTunes.	?
Universal Filtering	Exhibit, Relation Browser	Flamenco, eBay, Endeca, Google.

Exhibit is an example of a non-directional, but consistently laid out faceted browser, where used facets are not hidden. This means that the inter-facet relationships from the Genre/Artist/Album iTunes scenario can be created by the order of selections, as opposed to the order of the layout. Although this approach produces the same result set and values in each facet as a directional and consistent browser, there is yet no evidence to show that the unstructured layout makes the relationships as clear as having the three facets side-by-side. In summary, this approach provides: a) multiple selections in one facet, b) previous decisions, c) previous selections, d) all unused facets, and e) a result set.

It is worth noting here that no browser has yet attempted to provide direction in their filtering, whilst hiding previous decisions to make space for unused facets. This maybe because hiding previous decisions also removes the ability to see the inter-column relationships provided by directional browsing. Further, the combination would hide potentially unused facets (in the iTunes problem, selecting an Artist would put both the Artist and the Genre column out of view). This combination would appear to provide only a) previous selections and b) a result set.

THE COSTS ASSOCIATED WITH THESE DIMENSIONS

While the previous section indicates that some browsers have potential functional benefits over others, the opposing argument is that each additional benefit comes at a cost of interface complexity provided to the user. In the directional and consistently laid out browsers like mSpace and iTunes, the user has to comprehend the effect of direction and consider both facet-result and facet-facet relationships.

Consequently, we are left with the challenge of trying to estimate which approaches are '*better*' for the user. Certainly, the majority of examples of faceted browsers on the Web choose the less complicated non-directional and space-optimising layouts, which we consider to have less functional benefit. Alternatively, iTunes has chosen the more powerful, but perhaps more challenging approach of providing a directional and consistent layout. Wilson *et al.* have already produced an inspection-based evaluation framework that can analyse the extent of functional benefits provided by search interfaces, but consequently encourages the complicated directional and consistent designs provided by mSpace and iTunes [16]. We now discuss Cognitive Load Theory, which we believe can be integrated into the same framework to argue against complexity. The extended framework would support designers in deciding if the added benefits of new features outweigh the added complexities.

Understanding the costs using Cognitive Load Theory

Put simply, the notion of Cognitive Load Theory (CLT) is that the complexity of a learning task and any learning material both affect the users ability to gain the knowledge they seek [3]. The complexity of a learning task is called *intrinsic load*, and learning materials should aim to support users no matter how much intrinsic load their task requires. If a problem is too big for working memory, then learning material should support users in breaking it down into steps, each with lower intrinsic load. Learning materials, or the objects that support users in learning, provide *extraneous load*. The aim of learning material should also be to reduce its extraneous load on the user, so that more intrinsically loaded tasks can still be achieved. If the extraneous load is high, then only tasks with a low intrinsic load may be achieved. Ultimately, however, both need to be reduced to make space in the overall cognitive load, for *germane load*, which is required to commit anything learnt into schemas in long-term memory. According to CLT, although space for germane load can be produced by minimizing intrinsic and extraneous load, the design of learning materials can effect whether or not the space is used for germane load.

So far, CLT has been designed to understand how instruction manuals, for example, can be better designed to teach people to use machinery or computers [4]. In these scenarios, the task has been to learn how to use a computer and the material has been a book. Learning, however, is often the same task held by exploratory search users, except that the material they have to support them in achieving their goal is a search interface. Ultimately, the user is still aiming to learn something, and has resources to help them do it, and so our first position in this paper is that CLT can be applied to understand the complexity of search software. This position supported by Mu [7], who, states 'cognitive loads are closely related to the complexity of a task, the system used to operate the task, and the operators characteristics', which makes no indication that 'the system' need be instructional. Further, others have considered how CLT might help interface designers convey search result relevance [6] and explain why users rarely provide relevance feedback during search [1].

The next stage is to translate the methods that CLT has identified for reducing the complexity of instructional material, to the reduction of complexity in search interfaces. CLT presents three methods of improving instructional material: split-attention, modality, and redundancy effects.

Split Attention Effect refers to occasions when a user has to mentally integrate information from multiple sources, such as text and a diagram, in order complete their learning. Chandler and Sweller approach this problem by making sure that the text necessary to understand a diagram is embedded within the diagram [4]. Otherwise, the system places unnecessary extraneous load on users, as they have to remember textual information while interpreting the diagram, or visa versa. An example here, from mSpace, may be that previous choices are highlighted and left in place, rather than displayed as a separate list of choices in a separate location [15]. Consequently, users can see both their decision and choices in place. Conversely, it may be better to have all your choices in one breadcrumb-style place, rather than having to find them in multiple facets.

Modality Effect refers to the reduction of cognitive load, by distributing learning into the different modalities of working memory. mSpace has tried this with audio preview cues so that users may take advantage of the auditory channel when making decisions about musical domains [8]. Similarly, the Relation Browser provides graphical volume representations with each facet value, which uses a separate mode to numeric values [18].

Redundancy Effect refers to situations where the same information is displayed in multiple places, so that the user is potentially required to a) read information they have already read and b) recognize what is new or has already been seen. Chandler and Sweller further their previous diagram and text example, by removing text that simply states what is clearly demonstrated by the diagram. It would appear, for example, that reducing the redundancy effect might help protect users from decision paralysis [10].

Using CLT within an Inspection Evaluation Framework To Manage and Reduce these Costs

Most research into CLT measurement has focused on recording the actual experience of users, through physiological changes, subjective views, task performance, and secondary-task performance (where their ability to multi-task is reduced by high cognitive load). An inspection framework, however, focuses on assessment through careful estimation by some model and expected metric. Very little has been written about how to formally estimate cognitive load, but Chandler and Sweller [4] provide the following guidelines for estimating element interactivity: 'the extent to which elements interact for any given instructional material may be estimated a priori by simply counting the number of elements that must be considered simultaneously in order to learn a particular procedure.'

This process can be easily integrated into the authors' inspection framework [16], as it already counts the users 'moves' required to achieve a task. Chandler and Sweller add a caveat that this can only be applied in consideration of the user's existing capabilities. As the inspection evaluation framework also has a model of user types, this should also be easy to integrate. Further, as the framework already calculates the different interface features that allow users to carry out the same strategy, then we can also integrate measures for *split-attention* and *redundancy*.

With CLT integrated into the inspection framework, results would allow assessors to easily compare the extraneous loads produced by, in our example, different faceted browsers. This may first tell us if there is any significant cognitive load difference between the various approaches. Second, the framework would allow assessors to compare the difference between the increase in search support provided by each interface feature and the extraneous load produced. Third, the nature of the framework would allow assessors to quickly, and incrementally, consider design changes for both enriched support and reduced cognitive load. Having such a measure would complement cognitive engineering guidelines, such as the Ecological Interface Design framework [12], which encourage designs that require lower amounts of working memory.

CONCLUSIONS AND FUTURE WORK

In this paper we address the problem of a) finding the best trade-off between rich functionality and clear design, and b) discovering which combination of features best supports exploratory search. Using the inherent variation found in faceted browsers, we first discuss the root variables that cause such differences and propose that Cognitive Load Theory (CLT) may be able to provide a strong measure of clarity in design, while other existing measures push designers towards richer functionality.

The previous section has indicated that an estimate of CLT should fit nicely into an existing inspection-based evaluation framework, and so our immediate plans are to do so and validate it's findings against user studies of search interfaces. While most of the known methods of reducing CLT can be included in the framework, the *modality* effect may provide the largest challenge, as the framework currently takes no specific note of modality channels. The ultimate test, however, of using CLT this way, will be to actively improve user experiences of exploratory interfaces by providing rich functionality *and* clarity in design.

REFERENCES

1. Back, J. and Oppenheim, C. A model of cognitive load for IR: implications for user relevance feedback interaction. *Information Research 6*, 2 (2001).
2. Capra, R. and Marchionini, G. The relation browser tool for faceted exploratory search. In Proc. *JCDL08*, ACM Press (2008), 420-420.
3. Chandler, P. and Sweller, J. Cognitive Load Theory and the Format of Instruction. *Cognition and Instruction 8*, 4 (1991), 293-332.
4. Chandler, P. and Sweller, J. Cognitive load while learning to use a computer program. *Applied cognitive psychology*, *10*, 2 (1996), 151-170.
5. Hearst, M.A. Next generation web search: setting our sites. *IEEE Data Engineering Bulletin*, *23,* 3 (2000).
6. Hu, P.J.H., Ma, P.C. and Chau, P.Y.K. Evaluation of user interface designs for information retrieval systems: a computer-based experiment. *Decision Support Systems 27*, 1-2 (1999), 125-143.
7. Mu, X. Smartlinks in a Video-Based Collaborative Distance Learning System: a Cognitive Model and Evaluation Study. *School of Information and Library Science*, University of North Carolina, Chapel Hill, 2004.
8. schraefel, m.c., Karam, M. and Zhao, S., Listen to the Music: Audio Preview Cues for the Exploration of Online Music. in *Proc. Interact*, (2003).
9. schraefel, m.c., Wilson, M.L., Russell, A. and Smith, D.A. mSpace: improving information access to multimedia domains with multimodal exploratory search. *Commun. ACM 49*, 4 (2006), 47-49.
10. Schwartz, B. *The Paradox of Choice: Why More Is Less*. Harper Perennial, 2005.
11. Tunkelang, D., Guided Summarization. in *ECIR08 Industry Day Presentation*, (2008).
12. Vicente, K.J. Cognitive engineering: A theoretical framework and three case studies. *International Journal of Industrial and System Engineering 1,* 1 (2006), 168-181.
13. White, R.W., Kules, B., Drucker, S.M. and schraefel, m.c. Introduction. *Commun. ACM 49,* 4 (2006).
14. White, R.W., Marchionini, G. and Muresan, G. Evaluating exploratory search systems Introduction. *Inf. Process. Management 44*, 2 (2008)
15. Wilson, M.L., André, P. and schraefel, m.c., Backward Highlighting: Enhancing Faceted Search. in *Proc. UIST08*, ACM Press (2008).
16. Wilson, M.L., schraefel, m.c. and White, R.W. Evaluating Advanced Search Interfaces using Established Information-Seeking Models. *JASIST*. (to appear).
17. Yee, K.-P., Swearingen, K., Li, K. and Hearst, M., Faceted metadata for image search and browsing. in *Proc. CHI03*, ACM Press (2003), 401-408.
18. Zhang, J. and Marchionini, G., Evaluation and evolution of a browse and search interface: relation browser. in *Proc. National Conference. on Digital Government Research*, (2005), 179-188.