

Spatial Consistency and Contextual Cues for Incidental Learning in Browser Design

First Author Name (Blank if Blind Review)

Affiliation (Blank if Blind Review)

Address (Blank if Blind Review)

e-mail address (Blank if Blind Review)

Optional phone number (Blank if Blind Review)

Second Author Name (Blank if Blind Review)

Affiliation (Blank if Blind Review)

Address (Blank if Blind Review)

e-mail address (Blank if Blind Review)

Optional phone number (Blank if Blind Review)

ABSTRACT

This paper introduces the Backward Highlighting technique for mitigating an identified flaw in directional column-faceted browsers like iTunes. Further, the technique significantly enhances the information that can be learned from the columns and encourages further interaction with facet items that were previously restricted from use. After giving a detailed overview of faceted browsing approaches, the Backward Highlighting technique is described along with possible implementations. Two of these possible implementations are compared to a control condition to statistically prove the value of Backward Highlighting. The analysis produces design recommendations for implementing the Backward Highlighting technique within faceted browsers that choose the directional column approach. The paper concludes with future work on how to further improve on the statistically proven advantages provided by the Backward Highlighting technique.

Author Keywords

Faceted, search, columns, highlights, context, spatial.

ACM Classification Keywords

H5.2. User Interfaces: Graphical user interfaces. H.3.3. Information Search and Retrieval: Information Filtering.

INTRODUCTION

Visualising and manipulating multi-dimensional information spaces is a known hard problem. Finding solutions is increasingly important, since the information available to us is growing exponentially, and creating novel and effective ways to explore this data is critical. An example approach to visualising and interacting with these kinds of information spaces is faceted browsing, where there are strengths and weakness with each implementation.

In this paper we present a lightweight technique for improving the interactivity, learnability and explorability of faceted representations for large information spaces.

In directional column-faceted browsers like iTunes¹ and mSpace [17], for example, if a user makes their first selection in the second column-facet, or those towards the right of the browser, and the browser only filters columns to the right, then there will be no associated information indicated in the first column, and all those to the left of the selected item. They will only clearly see associated information in the filtered results of the columns to the right of their selection. For instance, in iTunes shown in Figure 1, there are three column-facets, which are directionally filtered to the right: Genre, Artist and Album. If a user selects an Artist (in the second column), they can easily see the associated Albums in the filtered column to the right, but no associations with the selected Artist are conveyed in the Genre column to the left. Subsequently, the Genre column is also excluded from contributing to the search, if the user does not already know which Genres the selected Artist has worked in. Without conveying useful information, the Genre column simply becomes a long list that is potentially wasting prime screen real estate, when it could be utilised to present more interesting information about the selected Artist and used to further refine their search.

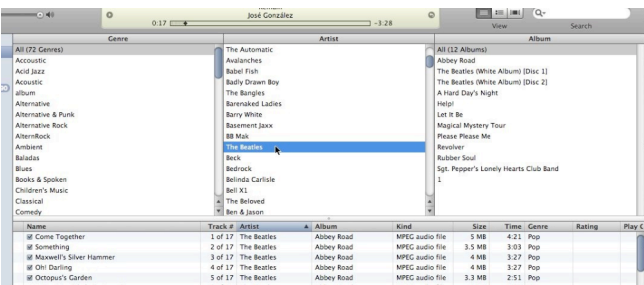


Figure 1: iTunes showing The Beatles selected, with their Albums showing in the Album column to the right, and no associations shown in the Genre column to the left.

¹ <http://www.apple.com/itunes/> - Apple – iPod + iTunes

We propose a technique for revealing associations in columns to the left of a selection in a directional browser, which we call Backward Highlighting (BH). In the example above, with Backward Highlighting, the missing associations in the Genre column would be highlighted so that the potential selectable relationships to the left, which would lead to the current selection, are clearly revealed. The BH technique, therefore, converts previously passive and uninformative columns into interactive visualisations revealing rich associations, while simultaneously removing a problem in directional browsers. While BH is a fairly intuitive suggestion for an interface improvement in directional faceted browsers, the optimal implementation is not so clear. In this paper we investigate two possible designs representations of BH. Within representation, how to assess the specific effects that the BH technique has on the user in terms of performance is also challenging. Speed of result finding may not be the optimal metric if the goal is to facilitate exploration and discovery of the domain. Therefore it may be more appropriate to consider how the technique affects: a) their increasing domain knowledge and b) their successive patterns of interaction with the column-facets.

In the rest of this paper, we first describe related work on benefits and flaws of different approaches to faceted browsers. Additionally, this related work is supported by search theory and research into incidental learning from a psychology background. This is followed by a detailed description of the BH technique in the context of related work. The two candidate BH implementations are described and evaluated through a user study, investigating the benefits and costs associated with each implementation in terms of learning and its effects on user interaction patterns. We conclude with an analysis of the results, design recommendations, and proposed future research into this area.

RELATED WORK

Since the turn of the century, a popular interactive alternative to keyword search has been faceted browsing [9], and has been shown to improve over other techniques such as basic hierarchical categories and clustering [10]. The premise of faceted browsing is to apply the theory of facets, which is well described by Capra et al. [4], to document collections so that a whole range of their metadata can be used as constraints during search. For example, Classical Music may have facets such as Song Titles, Composers, Periods, Instruments, and Arrangements, where songs have one or more values in each facet. While looking for classical music, users may want to narrow their results by items in any or all of these facets, depending on their existing knowledge. Figure 2 shows an example of a faceted browser for classical music, where three of the facets are presented in a row across the top of the interface. The user may make selections in these columns to narrow down their search for classical music. Faceted browsing has

become popular and is used in many commonly used interfaces: iTunes, eBay², and even Wal-Mart's website³.

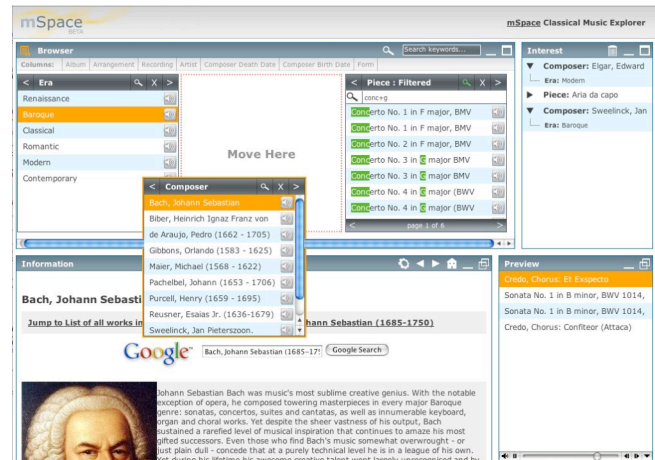


Figure 2: The mSpace directional column-faceted browser, showing three column-facets on classical music: Era, Composer and Piece Title.

In terms of search theory, the benefits of faceted browsing techniques over popular keyword searching systems alone is that they support a wider range of interaction with information [1] and a wider range of user contexts [2]. For example, it may be difficult for a user to express their interest in, for example, classical music with keyword terms alone, especially if their starting knowledge is low. In faceted search, multiple sets of search terms are listed in facets for the user to choose from, which permits recognition as well as recall. Further, having both options and selections in view allows a user to review their previous decisions and weigh up their current options all in one interactive visualisation. Both of these issues are explored in more detail by Wilson et al. [21], but the specific aim of the research is to work towards formalising the potential advantages in combining tools to improve synergy in design.

There is also strong support for faceted representations in the psychology of learning, as it encourages the association of concepts and subsequently builds knowledge of the information space. Marsick and Watkins [15, 20] define informal and incidental learning by their contrast with formal learning, as a “byproduct of some other activity, such as task accomplishment” or “trial and error experimentation.” Similarly, Kerka [13] suggests incidental learning is unintentional or unplanned and results from

² <http://express.ebay.com> - eBay Express | An exciting new way to shop

³ <http://www.walmart.com/> - Walmart.com – Always Low Prices!

other activities, often in the workplace and using computers, such as in the process of completing tasks [3]. In terms of faceted search, through the selection and subsequent filtering of facets and results, users may incidentally build knowledge on a topic, even though it was not their aim to do so. This is further supported by the following quote from Jones on incidental learning [12]:

"One of the major assets of such a system is the ability to make explicit links between concepts. Such explicit links may lead to an increase in knowledge, as the users gain an understanding of the information's underlying structure. This notion of connectivity is one of the main selling points of hypertext, since it corresponds to cognitive models proposed by research psychologists."

Further work by Jones in 1988 [12] into incidental learning during information retrieval involved a hypertext experiment examining browsing versus index use. Though no significant results were found at the overall level of performance, other insight was gained from the discussion of individual questions. Relevant to the faceted model, Jones reports that users performance and their construction of goals and plans is tied to their understanding of underlying relationships in the data.

Given all these benefits provided by faceted search, there are some significant variations in the way that they are implemented. Most notable is the difference between traditional faceted searching, such as Flamengo [7] and Exhibit [11] and directional column-faceted browsers like iTunes and mSpace [17]. In the traditional design of faceted browsers, a user makes a selection in one facet and all of the facets are filtered to show only the information that is relevant to the item selected, including the facet with the selection. This design, while effective for quickly narrowing down a search, provides less support for incidental learning as it removes all the information from the screen that is unrelated to the selection. In directed column-faceted search there is a specific order of facets (usually in columns) and the filtering usually goes from left to right; this is true of both mSpace and iTunes. That is a selection in a left column filters the columns to the right to show related information.

Let us again consider a traditional faceted system that allows users to browse classical music. Facets might include the Eras, Composers, Arrangements, and Pieces of music. Selecting a Composer in a typical faceted system would remove all of the Eras, Arrangements, and Pieces not related to that Composer. A further selection of a type of Arrangement, removes all of the Eras and Pieces unrelated to Arrangements by the previously selected Composer. The user is now unable to see all of the Arrangements that the composer produced, as all of the other Arrangements, and indeed the other Composers, are not related to the selected Arrangement by the selected Composer. In directional column-faceted browsers, shown in Figure 3, selecting a Composer would filter the Arrangement and Piece columns,

but leave all the composers in view. Then selecting one of the remaining Arrangements would filter the Piece column only, leaving all the Arrangements used by the Composer still in view. There is no benefit in removing the Composer's contemporaries and the other types of arrangements, unless screen real-estate is an issue. In the remainder of this paper, we call this type of information that is usually lost in traditional faceted search, but kept in directed faceted search, 'Added Facts'.




Figure 3: A simple directional column-faceted browser with the Arrangement column filtered by the Beethoven and the Piece column filtered by Beethoven and by Violin Arrangements. This interface is used as the control condition in the user study described below.

While it is possible that under certain conditions the traditional method may be more appropriate, search and psychological research would suggest that the directed method has significant advantages. For example, the directed method allows a user to check their decisions in terms of both what is selected and what was not selected. This directly supports one of the 32 recognised search tasks noted by Bates [1]. Previous research has also shown the importance of spatially consistent interfaces [5, 16], and where traditional faceted techniques filter information, column browsers usually maintain more spatial consistency in both the layout and content of columns. Further support for keeping these Added Facts is provided by the memory model produced by Cseh et al [6] who noted that context cues are particularly important for learning; similar emphasis has been noted in research into Focus+Context interfaces [8]. Finally, another quote from Jones [12] also supports the column-facet model:

"If knowledge is represented cognitively as an associative network or hierarchy of information, then systems which are organized along similar lines would seem to be desirable."

BACKWARD HIGHLIGHTING

Although directional column-facet browsers have the advantages described above, there is a cost as we suggest in the introduction that can be mitigated by the addition of backward highlights. This cost, given in the iTunes example, is that users may not start in the first column, and so no associations are shown in the Genre column that

would otherwise be shown by the traditional faceted search technique. The intuitive proposition is to highlight associations in columns to the left of a selection to indicate the possible paths a user could take to get to where they chose to start. As this highlighting happens backwards up the path from the top of the represented hierarchy, we call it Backward Highlighting (BH).

In our more detailed classical music example, where we have the Era, Composer, Arrangement, and Piece facets, should a user select a certain Arrangement, the potentially multiple associated Composers and Eras would simply highlight and the associated Pieces would remain in final column. The user can now see all of the composers who used the selected type of arrangement and could select any one of them to further refine their search. The ability to make a further selection to the left of the initial selection was previously the cost of moving to directional column-facet browsers but was possible in traditional faceted search. Now able to do this in directional browsers, selecting a Composer would leave only the Era of the selected Composer's Arrangements highlighted. The Piece column, as always, would show only the Pieces of that Arrangement by the selected Composer. The visible Added Facts are: 1) the Eras where the Composer did not compose any pieces in the selected Arrangement, and 2) all of the Arrangements used by the Composer, including the selected one.

The overall effect of adding the BH technique is that the user gets the benefit of both traditional and directed faceted browsing, increasing flexibility interaction but persisting some structure and sense for the user. In terms of increasing domain knowledge, however, the expected effect was that the user will see further associated information, along with their selections and filtered lists, and result in increased incidental learning. In terms of their future search behaviour, it was expected that the BH technique allows users to convert their increased domain knowledge and use previously highlighted items as filters.

Given that the BH technique is an intuitive concept, there are interesting variables in its possible implementation that are not so straightforward. For example, the related work above would suggest that items in parent facets (those to the left with BH rather than child facets to the right, which are filtered) should be simply highlighted in place to maintain spatial consistency. However, any items highlighted in this way would be embedded in potentially very long lists, where the same information in traditional faceted browsing would be collected together. The possible result is that items are highlighted but not seen, which would have a negative effect on incidental learning. An evaluation of these two possibilities is presented, along with the results, in the next section.

EVALUATION

The evaluation of the Backward Highlighting technique was designed to evaluate whether conveying this additional

information has benefits for users by encouraging incidental learning. Subsequently, the conditions and measures are specifically designed to examine the effect of the additional information conveyed by backward highlighting on incidental learning. First, to attribute any incidental learning directly to information that is highlighted or filtered in the facets, a cut down simple set of directional column facets were created, as in our control condition shown in Figure 3 in the related work section above, which shows a Classical Music dataset. The user never leaves this view, and no further information is given except for what is shown, highlighted, and filtered in the columns. The chosen columns may not be rearranged, removed, or added to.

Our first experimental condition, and our initial conceptual design for the BH technique, highlights related items in facets to the left of selections Figure 4. Highlights are differentiated to selections by colour, where selections are yellow and highlights are green. In addition, a count of highlighted items is displayed at the top to indicate whether there are any highlights that are currently out of view and can be seen by scrolling. After a selection is made, any parent facets automatically scroll until a highlight is visible; this may involve no scrolling if one is already in view. In this paper, we refer to this first experimental condition as the Backward Highlighting condition.



Figure 4: The first experimental condition called 'Backward Highlighting'. In the columns, the Cello Arrangement is selected and the Piece column shows all the Pieces with a Cello Arrangement. The Era and Composer columns show the items associated with Cello music with highlights.

The second experimental condition, shown in Figure 5 represents an alternative design, which is identical to the first experimental condition, but also copies all highlighted items into a separate 'bucket' at the top of each column. This separation is much like the dual menu designs presented by Sears and Shneiderman [18]. Then regardless of the position of the scrolled list, the highlighted items persist in view and can be scrolled independently of the column if needed. While there is a necessary trade-off for screen real-estate, this condition is designed to identify any significant advantages for using this space to group highlighted items. The items in the 'bucket' can be selected by the user, as in the column itself, and causes the same

behaviour. In this paper, we refer to this second experimental condition as the Bucket Highlighting condition.

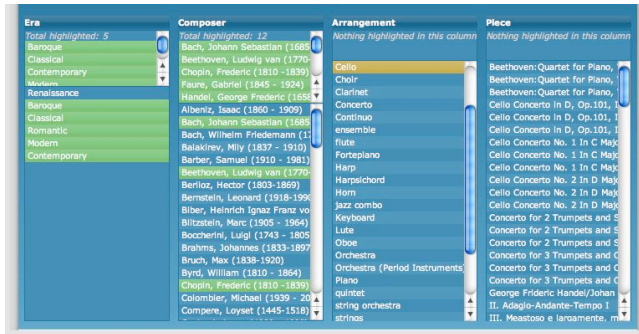


Figure 5: The second experimental condition called 'Bucket Highlighting'. The same information is shown as in Figure 4, except that the highlights in the Era and Composer columns are copied and collected into a separate scrolling list above the column. The user may now scroll through only the highlighted items.

Hypotheses

The study was carried out with three hypotheses regarding the effect of BH on incidental learning:

H1: We hypothesised that having BH would permit incidental learning on parent facets. Parent facets are those presented to the left of facets that contain the subject of learning and thus are parents in the left-to-right hierarchy. While we expected it to be true that using the BH will allow users to learn about parent facets where they could not before, it is important to check that this is true and that simply highlighting has the same effect as filtering. In particular we expected to see a significant effect between the two experimental conditions and the control condition.

H2: We hypothesised that grouping any backward highlighted items, as in the Bucket Highlighting condition, would aid incidental learning and make the technique more effective. Given that highlights are as effective for incidental learning as filtering, we suspected that collecting them together will emphasise this ability to learn incidentally from the columns. In particular we expected to see a significant increase in incidental learning between the two experimental conditions.

H3: We hypothesised that the result of increased incidental learning on parent facets would lead to an altered interaction behaviour in subsequent interactions with the columns. That is to say, rather than scrolling through a long list to find an item in the middle of a path through the columns, participants will begin to alter their patterns of interaction to start at the beginning of a path and find information using filtering. We believe that, being a more efficient method of accessing the data, users will be able to improve as their knowledge of the information space increases. In the remainder of this paper we will refer to this

Altered Pattern of Behaviour as APB. In particular, we expected to see a significant rise in the number of re-finding tasks carried out with APB between the control condition and, potentially, both control conditions.

Study Design

The study was a within-groups repeated measures design, where participants were exposed to each of the three conditions in a counter-balanced order. Each participant began by filling out an agreement and demographic survey. There were 18 participants and so each condition was trialed 18 times, and each order of exposure was trialed 3 times. The participants ranged between 18 and 65 years old, 12 were male and 6 were female, and they had varying levels of education and subject backgrounds. With each condition, the user carried out 2 task-sets, which are described below. After using all 3 conditions, participants were asked to carry out some memory tasks, also described below. Finally, each participant was debriefed with a semi-structured interview regarding the tasks and their experiences with the conditions. Participants were given a music voucher in appreciation of giving their valuable time.

To reduce learning effect between the conditions, three datasets were used in the study: a classical music dataset, a historic news film dataset, and a BBC TV programming catalogue dataset; the association of dataset and condition was also rotated. Each dataset had four facets. The facets of the classical music dataset were Era, Composer, Arrangement, and Piece. The facets of the news film dataset were: Theme, Subject, Topic and Story Title. The facets of the TV programming dataset were: Subject, Series, Contributor, and Title. Having multiple datasets also provides opportunity to evaluate any subtle differences if a significant pattern is found between them in analysis.

Two task-sets were carried out with each condition, which involved one learning task and three re-finding tasks. One task-set focused on an item in the second column and the other focused on an item in the third column. This separation of task made sure that there was always one filtered child facet and one highlighted parent facet, while supporting a comparison of the two in later analysis. Finally, every participant carried out a memory task for each task-set and each condition, leading to a total of six task-sets and six corresponding memory tasks.

Tasks and Measurements

Learning Tasks. In line with the hypothesis of the evaluation, the tasks were exploratory and involved the user learning about items in the middle two facet columns. As we are measuring learning effect, the tasks have to be very carefully planned so that they do not have an effect on what is learnt. Klauer [14] presented a meta-analysis of a decade of incidental learning studies, concluding that incidental learning is impeded by giving behavioural objectives, learning directions or questions before an instructional text. Subsequently, questions were very clear and simple, with no

added context. An example task, given simply a facet and an item within it, was: "Please learn about the Arrangement: Orchestra".

Re-finding Tasks. Based on what was discovered during their learning tasks, participants were asked to carry out three re-finding tasks. Each time the condition was reset to its starting position, so that the behaviour within each re-finding task had the same baseline and did not vary depending on the previous task.

Memory Tasks. After both task-sets had been completed on each condition, the participants were asked to recall as much information as possible about what they had learned during the initial learning tasks. These were carried out in the same order as the task-sets within each condition, so that there was a gap between learning and recall in all cases.

From there we intended to measure the amount of incidental learning in three ways. First, the user was asked to write facts about each object of the learning tasks as they interacted with the browser, in a similar manner to the study carried out by Todd on the learning skills of children [19]. Within this written account, the number of statements made about items in the parent facets to the left of the learning objects were counted. Second, we measured incidental learning through their interactions, by asking them to re-find information from their written account. Through these re-finding tasks, we measured the amount of interaction with facets to the left of the facets containing the objects of the tasks. We hypothesised that greater and more accurate interaction with these facets on the left indicate that the user has learned that they can use this metadata to filter down towards the requested information. This expected behaviour is that of users who know what they are looking for and how to find it, which is in line with the expectations of many existing search techniques. Finally, the users were asked to complete a memory task, where they try to write down as much as they remember about each of their previous learning tasks. Counting the balance of statements from this second written account will measure the amount of retained information. We hypothesised that backward highlighted items have a similar retainment in memory as filtered items.

RESULTS

In this section, we consider each of our hypotheses and the results that support or disprove them. We finish with a discussion of additional interesting results found during the study.

H1: Highlights Support Incidental Learning

As hypothesized, there was a statistically proven increase in the number of parent facts learned and remembered by participants in the experimental conditions over the control condition and there is no significant effect on child facts. Figure 6 shows the number of parent facts and children facts written down during learning and the number of

parent and child facts written down during the memory tasks.

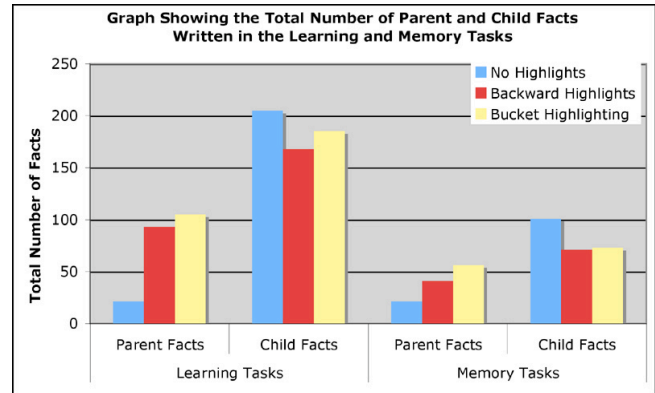


Figure 6: Shows the total number of facts written at different stages of the study, by condition, where Parent Facts are those written about facets to the left of a selection, and Child Facts are those written about filtered facets to the right. As expected, no significant difference was found in the Child Facts between any of the conditions. The differences seen in the Learning Parent facts is significant (ANOVA, $F=14.97$, $p<0.0001$), as are the Memory Parent Facts (ANOVA, $F=5.292$, $p<0.01$).

In the learning tasks, the trend in written parent facts is highly significant (ANOVA, $F=14.97$, $p<0.0001$). The significance, however, is between the control condition and the two experimental conditions, while the increase in parent facts written in the Bucket Highlighting condition, over Backward Highlighting, only has a significance of $p<0.5$ (t-test, $t=0.6972$).

In the comparison of all three conditions during memory tasks, the difference is still highly significant (ANOVA, $F=5.292$, $p<0.01$), but the amount remembered in the Bucket Highlighting condition over Backward Highlighting is more significant (t-test, $t=1.5357$, $p=0.1336$) than in the learning tasks.

While the significant effect of having highlights was expected and supports our first hypothesis, these statistics alone do not tell us about any significant difference between Backward Highlighting and Bucket Highlighting. These two specific conditions are investigated further below.

H2: There is a benefit to grouping highlights

Our second hypothesis was that grouping highlights, as in our Bucket Highlighting condition, had a significant learning advantage to the user over simply highlighting items as they appear in the columns. Although there was no significant difference in the number of facts written down in the learning and memory tasks, post-study questionnaires revealed that Bucket Highlighting was the preferred condition for 16 out of 18 participants. Further to this, the significant trend (ANOVA, $F=5.715$, $p<0.005$) shown in Figure 7 shows that users found the Bucket Highlighting

condition slightly easier than Backward, and Backward Highlighting condition easier than the control. This trend of decreasing difficulty matches the increasing trend in the number of Parent facts written in the Learning and Memory tasks in Figure 6. In Figure 7 we can also see that there was no significant difference in previous knowledge participants had within each condition.




Figure 7: Average score from a likert scale for both difficulty and previous knowledge, by condition. No significant difference was found in the previous knowledge scores, but an ANOVA showed that Bucket Highlighting was perceived as easier to use than Backwards

During further discussion in the structured interviews users consistently explained that they preferred Bucket Highlighting for when the list was longer, because then they didn't have to scroll to find information. To investigate this further, an analysis was carried out by condition and dataset. Figure 8 shows the number of parent facts written down during memory tasks for both Backward and Bucket Highlighting conditions. While there is little difference seen for the Classical and News film datasets, a significant difference (t-test, $t=2.7269$, $p<0.05$) is seen between the two conditions on the BBC TV scheduling dataset. Analysing the three datasets, we see that both the Classical Music and News film datasets have 6 and 15 items in the first column, respectively. In comparison the BBC TV scheduling dataset has over 200 items in the first column. Subsequently, we can deduce that Bucket Highlighting shows specific advantage as the size of the parent columns increase.




Figure 8: The total number of facts written in the memory task phase of the study, by condition and then dataset. The only significant difference is between the Backward and Bucket conditions with the BBC TV Scheduling dataset (t-test, $t=2.7269$, $p<0.05$). Analysis of the dataset shows that it has significantly longer left-hand columns than the other datasets.

H3: Highlights will improve future interactions

Our third hypothesis, suggesting that highlights would allow people to improve their interaction patterns in re-finding tasks, which we called APB above, was not supported by the results. Figure 9 shows the average number of tasks with APB in each condition. With a maximum average of three, we can see that users very rarely altered their pattern of interactions, opting instead to repeat their previous patterns. Most users, who did alter their patterns of interaction, did so consistently, but no significance could be obtained from this graph or from any other analysis by dataset or by demographics as to what causes this effect. There is a slight gender imbalance, as shown by Figure 10, but this was only significant to $p=0.35$ (t-test, $t=0.9221$). Interestingly, this pattern is exaggerated for younger females and older males. The most significant trend, although still low, was found by education (ANOVA, $F=1.627$, $p=0.2$), as shown in Figure 11, where we see a large bias towards users who have a Bachelors degree, and reduces with further education. This does not relate to age, as a separate analysis shows age effect to be insignificant. Subsequently we must reject our third hypothesis and leave the investigation of APB to future work.




Figure 9: The average number (out of 3) of re-finding tasks with APB, by condition. No significant difference was found by condition.




Figure 10: The total number of re-finding tasks with APB, by gender. A larger trend is seen, but is still not significant.

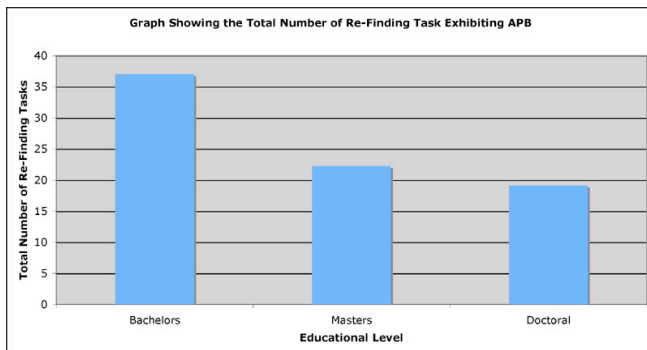


Figure 11: Total number of re-finding tasks with APB, by educational level. An even stronger trend is seen towards less educated people, but is still not significant. No significant trends were found across the collected data, and requires further investigation.

These results are not conclusive as to the cause of APB, but we have shown that it is not significantly affected by the

addition of highlights that might guide some users towards optional parent filters. It would be interesting in future work to investigate these behavioural differences more, but within a wider focus than one possible highlighting technique.

Further Discussion

One additional factor in the motivation for this research was that previous work indicates that spatial layouts are important for users [16], and here we have seen that having a grouped space provides both a significant advantage to incidental learning and a preferred user experience. However, our analysis showed that this was only significantly so when users were presented with a long list of items that might require lots of scrolling.

Subsequently, the results were analysed further to investigate whether there are any disadvantages to copying items into separate lists instead of simply keeping them in place. The results of this investigation showed that Backward Highlighting is specifically more useful for finding what we have called 'Added Facts'. Added Facts are those that can be found through directional faceted browsers (like column browsers including mSpace and iTunes), but not through more traditional style faceted browsers. Around 10% of the parent facts written during learning and memory tasks were these 'Added Facts', and, as shown in Figure 12, significantly more of these were found during Backward Highlighting (ANOVA, $F=3.869$, $p<0.05$).




Figure 12: Total number of Added Facts written in the learning and memory tasks, by condition. The increase in learned and remembered Added Facts in the Backward Highlighting condition is significant (ANOVA, $F=3.869$, $p<0.05$).

This significant difference is interesting, and by analysing interaction patterns and the semi-structured interviews, it appears that it is related to its spatial position, as related work suggested. First, having a separate list means that people do not look at the main list and therefore miss these added facts. Further, having them in a separate space means

that they do not interact with them to find other 'Added Facts' such as sibling information (i.e. Beethoven's contemporaries). Subsequently, it appears that while collected highlights, as in the Bucket Highlighting condition, is good for learning direct associations with a given facet item, but bridged information (information related by a second constraint) is more often found through highlights that are spatially consistent.

DISCUSSION

In our results above we have shown significance for the first two of our three hypotheses and rejected our third. A clear thought process would naturally lead to the assumption that highlights are more useful than not having highlights, and we have shown this to be significant. Similar results show that having these highlights grouped is preferable for users, and improves incidental learning, specifically in long list conditions. This added benefit, however, is at the cost of spatial consistency, which has been shown to be significantly important for deeper exploration, where users interact with highlights to find further related information. Finally, we have shown that although these highlights give added value for incidental learning, this does not have a significant effect on the interaction patterns of users. Investigating changes in such patterns requires a more thorough investigation in the future.

Design Recommendations

Given the results summarised above, we suggest the following dynamic implementation of the BH technique for future inclusion into directional faceted browsers. Based on given research, the importance of spatially consistent highlighting for deeper exploration, and the conditional importance of Bucket Highlighting, we suggest that column-facets should default to the first experimental condition. However, for when a list become longer than the given screen space, a small bucket should be fore-grounded that can be expanded and scrolled as desired. Although less useful for short lists, the bucket should still be available, but not fore-grounded.

Future Work

Aside from investigating the causes and factors in Altered Patterns of Behaviour in the future, there are also remaining questions in the area of Backward Highlights. For example, after a selecting an Artist in iTunes, the filtered list to the right would show related Albums and highlighted items in the left would show related Genres, but the relationships between them are not represented. For example, if an artist has worked in multiple Genres, the it is highly likely that different albums will be connected to different Genres, and it may be advantageous to show which Genre each Album is in. This is by no means a trivial problem, as indicators such as colour do not scale to potentially hundreds of highlighted items.

It may be that an interactive brushing action, which would highlight these relationships on demand, would be sufficient, but then the added information would only seen if requested. Another approach is to group the filtered list by the highlighted items in columns to the left. This approach also has limitations, however, when filtered items are associated with more than one of the highlighted parent items. This same problem is found when a filtered list is the union of two selections from the same facet, in that there is a desire to indicate which selection each filtered item relates to. No obvious solution to this problem has been discovered so far.

Finally, given the design recommendations produced from our investigation, a longitudinal study of the BH technique would most likely reveal interesting results of its every day use within a complete directional column-faceted browser.

CONCLUSIONS

In this paper we have introduced and evaluated a technique we term backward highlighting - a contextual cue for directional column-based faceted browsers that combines the benefits of both directed browsers and traditional non-directed faceted browsers. By highlighting associations in columns to the left of a selection, the BH technique converts a previously static unassociated list in two ways. First, the highlighted items in lists convey new information that is not seen without the use of Backward Highlighting. Second, the highlighted items can be used to advance or refine an existing set of selections, where a user was previously constrained to column-facets to the right only.

As well as defining the benefits of directional faceted approach over the traditional models, the user study results have let to clear design recommendations for the optimal implementation of Backward Highlighting within column browsers such as iTunes and mSpace. The recommended design is both spatially consistent to encourage further interaction with the facets, while collecting the highlights on demand for occasions when highlighted items might otherwise be lost in a long list. Finally, the investigation has identified future research in this area for further improving the quality of information conveyed by the Backward Highlighting technique.

In summary, the Backward Highlighting technique clearly improves directional column-faceted browsers by mitigating an identified flaw, providing the user with additional information, providing the user with further means of interacting with the facets, and maintaining the known benefits and interaction model. While the benefit is focused and clear for faceted browsers, the principles echo towards improving the interactivity, explorability, and learnability of large information spaces, through the maintenance of spatial consistency and maximising the potential of the data and its internal relationships.

ACKNOWLEDGMENTS

Thanks to Steven Day for work on implementation, and the participants of the studies.

REFERENCES

1. Bates, M.J. Information search tactics. *Journal of the American Society for Information Science*, 30. 205-214.
2. Belkin, N.J., Cool, C., Stein, A. and Thiel, U. Cases, scripts, and information-seeking strategies: on the design of interactive information retrieval systems. *Expert Systems with Applications*, 9 (3). 379-395.
3. Cahoon, B.B. Computer skill learning in the workplace: A comparative case study (Doctoral dissertation, University of Georgia, 1995). *Dissertation Abstracts International*.
4. Capra, R., Marchionini, G., Oh, J.S., Stutzman, F. and Zhang, Y. Effects of structure and interaction style on distinct search tasks. *Proceedings of the 2007 conference on Digital libraries*. 442-451.
5. Chun, M.M. and Jiang, Y. Contextual Cueing: Implicit Learning and Memory of Visual Context Guides Spatial Attention. *Cognitive Psychology*, 36 (1). 28-71.
6. Cseh, M., Marsick, V.J. and Watkins, K.E., Re-conceptualizing Marsick and Watkins' Model of Informal and Incidental Learning in the Workplace. *Proceedings of the Academy of Human Resource Development Conference*, (1999). 349--356
7. Elliott, A., Flamenco image browser: using metadata to improve image search during architectural design. *CHI extended abstracts on Human factors in computing systems*, (New York, NY, USA, 2001), ACM Press. 69-70
8. Furnas, G.W. Generalized fisheye views. *Proceedings of the SIGCHI conference on Human factors in computing systems*. 16-23.
9. Hearst, M.A. Next generation web search: setting our sites. *IEEE Data Engineering Bulletin: Special Issue on Next Generation Web Search*, 23 (3). 38-48.
10. Hearst, M.A. Clustering versus faceted categories for information exploration. *Commun. ACM*, 49 (4). 59-61.
11. Huynh, D.F., Miller, R. and Karger, D., Exhibit: Lightweight Structured Data Publishing. *Proceedings of the World Wide Web Conference*, (2007).
12. Jones, T. Incidental Learning During Information Retrieval: {A} Hypertext Experiment, 1989, 235-253.
13. Kerka, S. Incidental learning. Trends and alert, 2000.
14. Klauer, K.J. ntentional and Incidental Learning with Instructional Texts: A Meta-Analysis for 1970-1980. *American Educational Research Journal*, 21 (2). 323-339.
15. Marsick, V.J. and Watkins, K.E. Informal and Incidental Learning. *New Directions for Adult and Continuing Education*, 2001 (89). 25.
16. schraefel, m.c., Karam, M. and Zhao, S., Listen to the Music: Audio Preview Cues for the Exploration of Online Music. *Interact*, (Zurich, Switzerland, 2003).
17. schraefel, m.c., Wilson, M.L., Russell, A. and Smith, D.A. mSpace: improving information access to multimedia domains with multimodal exploratory search. *Communications of the ACM*, 49 (4). 47-49.
18. Sears, A. and Shneiderman, B. Split menus: effectively using selection frequency to organize menus. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 1 (1). 27-51.
19. Todd, R.J. From information to knowledge: charting and measuring changes in students' knowledge of a curriculum topic. *Information Research*, 11 (4). 2005-2006.
20. Watkins, K.E. and Marsick, V.J. Informal and Incidental Learning in the Workplace. *New York*.
21. Wilson, M.L. and schraefel, m.c., Bridging the Gap: Using IR Models for Evaluating Exploratory Search Interfaces. *First Workshop on Exploratory Search and HCI at SIGCHI07*, (2007).