Learnability and Perceived Benefits of Parallel Faceted Browsing: Two User Studies

Anthony Jameson, Adrian Spirescu, Tanja Schneeberger, Edit Kapcari, and Sven Buschbeck
German Research Center for Artificial Intelligence (DFKI)
Stuhlsatzenhausweg 3
Saarbrücken, Germany
firstname.lastname@dfki.de

ABSTRACT

Following up on work presented at IESD 2012 introducing the paradigm of multifocal exploration of semantic data, the present paper reports on two user studies of prototypes that instantiate parallel faceted browsing—a generalization of faceted browsing that enables multiple interrelated queries and their results to be displayed at the same time. In the first study, with the "World Design Capital Helsinki" demonstrator, 100 participants remotely tested the prototype for a few minutes each, performing simple tasks without explicit instructions about how to operate the system. The majority of participants were able to understand the system after engaging in trial and error, but even the successful ones found it rather unfamiliar-looking at first; and a feeling of unfamiliarity appears to have discouraged the less successful subjects from exploring the interface in the first place—a result that indicates a need to provide explicit explanation and motivation for the benefit of users who are less inclined to engage in trial and error. The participants spontaneously noticed a variety of benefits of parallel faceted browsing relative to existing interaction paradigms. In the second study, which involved a different instantiation of parallel faceted browsing in the domain of food and recipes, results concerning learnability and perceived benefits were generally consistent with those of the first study. Subjective ratings revealed mostly positive evaluations of the demonstrator, though a minority of participants stopped working with it before they perceived its benefits.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: User interfaces—*Evaluation / methodology*

Keywords

Parallel faceted browsing; Learnability

IESD'13, May 1, 2013, Paris, France.

1 Multifocal Exploration and Its Hypothesized Benefits

In the keynote talk at IESD 2012,[1] Jameson argued that more attention should be paid to *multifocal exploration* of semantic data. A system for multifocal exploration enables the user to explore in several directions in parallel, instead of being forced to pursue one line of exploration at a time, as is the case with almost all of even the most sophisticated data exploration systems.

Jameson postulated several general benefits of multifocality, which follow from its basic nature:

1. Multifocality enables users to deal more effectively with the uncertainty that is inherent in exploration. By definition, someone who is exploring isn't sure exactly where to go at any given moment. Going in a direction that turns out not to be rewarding typically results in backtracking and trying out a new direction, as in the computational search strategies of hill climbing and depth-first search. Note that computational strategies also include some that involve exploring multiple lines at a time, such as beam search and best-first search. Such strategies are much less common in interactive systems, presumably because pursuing more than one line at a time is inherently more demanding in terms of both screen real estate and cognitive complexity.

2. Multifocality helps users when they need to find a set of two or more items that are related in some way, each of which can be found in a different place. Being forced to identify one item and then look for another one that fits it is often less effective than looking for two suitably related items in parallel.

3. A system that supports multiple lines of exploration can yield as a side effect a structured overview of a subset of the space in question. Such an overview can serve as a useful source of information for future reference, for the users themselves or for others.

4. Multifocality makes it possible for two or more persons to explore along different lines in a coordinated way, either synchronously or asynchronously.

1.1 Previous Related Work

In areas outside of semantic data exploration, there has been some experimentation with systems that support multifocal exploration. Lunzer and Hornbæk (4) introduced the concept of *subjunctive interfaces*, which "provide mechanisms for the parallel setup, viewing and control of scenarios". This concept has been realized in various domains (see, e.g., 3) but not (to our knowledge) applied to semantic data. The benefits of multifocal exploration that Jameson argued for should in principle apply to these interfaces as well;

[1] `http://imash.leeds.ac.uk/event/keynote.html`

they are not in fact discussed explicitly in this way by the above-mentioned authors, though the first and most general advantage—that of better coping with uncertainty about where to explore next, is discussed in other terms.

More generally, support for mutifocality is hard to find in the many existing types of system for exploratory search (see, e.g., 8, 7).

1.2 Goals of the User Studies

At IESD 2012, Buschbeck et al. (2) presented an example of multifocal exploration of semantic data, introducing the paradigm of *parallel faceted browsing*: an interface for faceted browsing (see 6 for a thorough survey) that enables the user to create multiple interrelated queries on the screen at the same time and examine their results in relation to each other.[2]

Although this prototype had been subjected to iterative user testing, there had been no summative testing of its usability and usefulness with a large number of users. In the present paper, we report on a study of a later version of the PFB demonstrator which aims to fill this gap. We also report more briefly (in Section 5) on an evaluation of an independent instantiation of PFB which, although it does not access semantic data, embodies most of the same basic ideas as the semantically based prototype.

2 Research Questions

There are many conceivable user studies of a parallel faceted browsing system that could be conducted. For example, one could compare it directly with a normal faceted browsing system to see which one was more effective for particular types of task. But an even more basic question is that of whether PFB can be made reasonably comprehensible and learnable even to users who are not willing to spend much time getting to know it. Also, users ought to be able to perceive the advantages of PFB so as to be motivated to use it. If these conditions are not satisfied, then the paradigm is unlikely to gain wide acceptance.

Consequently, the evaluation studies reported on here focused on the questions of immediate learnability and perceivability of benefits. We chose a method of recruiting participants that yields a relatively large number of participants who expect to be working for only a few minutes and who can essentially stop at any time if they do not feel motivated to continue. Although more extensive contact with study participants is normally considered desirable, this sort of brief contact has the advantage of being closer to the situation of a casual user who visits a website with novel technology and must quickly decide whether it is worthwhile to figure out how to use the novel system, as opposed to clicking away quickly to visit a more conventional site.

3 Study 1: Method

3.1 Prototype

The PFB demonstrator used in the first study gives access to 899 events that are relevant to Helsinki's role as the 2012 World Design Capital. In addition to the exhibitions associated with the design capital itself, the repository contains a large number of cultural and sports events that might be of interest to visitors to the World Design Capital.

Figure 1 shows how several interrelated queries are visualized at the same time on the screen: The small pile of cards in the lower left-hand corner represents the set of 12 events that occurred during the last 22 days of December, 2012 that involve musical concerts in Helsinki; the smaller pile of cards to the right includes the analogous events in Tampere. Much larger trees of interrelated queries can be built up according to the same principle. As can be seen in Figure 2, the user can examine the results of each query by clicking on the pile of cards, and they can also "pin" descriptions of the individual events so as to be able to keep them in view.

Readers who are interested in seeing exactly how the demonstrator works (which is not necessary for the understanding of this paper) can visit a website[3] with links to (among other things) the demonstrator itself and a video in which its use is demonstrated. The prototype's user interface, implemented in the GOOGLE WEB TOOLKIT, essentially runs in any web browser. Information about how the event repository was constructed can be found in (2).

3.2 Pilot Study

As preparation for a larger-scale study, five students were observed as they performed several typical tasks with the demonstrator. In addition to suggesting several usability improvements that were realized immediately, this pilot study suggested that users do not find it particularly helpful to be given a legend explaining the visual notation and the controls in the PFB interface. As a consequence, it was decided to give the participants in the main study an already existing PFB structure to examine and explore before they began extending it themselves.

3.3 Main Study

Participants

The participants were 100 persons from the United States, Canada, and the United Kingdom who were registered with AMAZON MECHANICAL TURK and who responded to an invitation to participate in a brief on-line website evaluation. Demographic details about the participants are not available, but the fact that they were registered with AMAZON MECHANICAL TURK suggests that they are regular computer users who are familiar with the use of websites; this assumption is confirmed by the comments that they made.

Procedure

Participants were instructed to visit a URL which took them to the view shown in Figure 1. After reading the introduction shown in the right-hand side of that figure, they were to follow the task instructions shown in Table 1, which they could access by scrolling downward. Note that these task instructions include no explanation of the visual notation in the interface, including the card metaphor, or any indication of where the participant should click. Participants could figure out what to do by examining the interface itself, including the tooltips that appeared when the cursor was placed over an icon, and by engaging in trial and error.

A participant who successfully followed all of the task instructions ended up with a view like that shown in Figure 2.

[2] A more recent version is being presented in the CHI 2013 Interactivity track (1).

[3] `http://parallel-faceted-browsing.com`

Figure 1: Initial view of the World Design Capital prototype seen by study participants.
(By scrolling down in the right-hand sidebar, they could read the specific task instructions, which are shown here in Figure 1.)

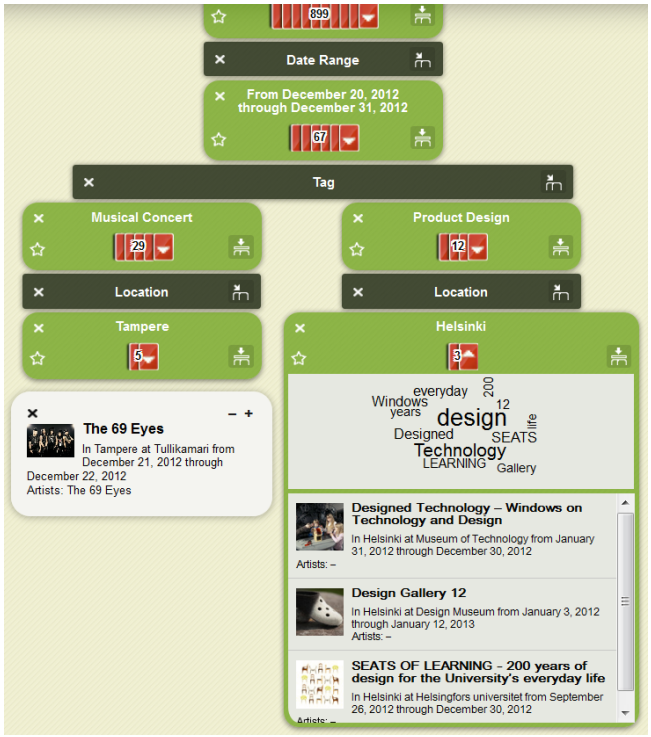

Figure 2: Final view of the World Design Capital prototype seen by study participants who successfully followed all of the task instructions.

Table 1: Tasks presented to participants in the study.
(These instructions appeared below the heading "What You Should Do" that is visible in Figure 1.)

Alice has told you that her favorite Finnish band is appearing in both Tampere and Helsinki, so please ...

1. Figure out which band that is.
2. Pull out a description of the concert by that band in Tampere, so that Alice will see it later.
3. To save space,

   hide the details about the other events in Tampere;

   get rid of the information about the concerts in Helsinki.

Alice has said she'd also like to attend exhibitions about product design, so please ...

Find the events during this period with the tag "product design".

Among these exhibitions, pull out the set that will take place in Helsinki.

Check whether some of those exhibitions will still be going on after Alice has attended the concert in Tampere.

To finish up, ...

Create a link ("bookmark") to the display that you have just created, and copy it (so that you could send it to Alice by email).

Close this tab and click "Next" in the tab with the instructions.

Table 2: Questions asked of the participants after they had completed their tasks.

1. Please paste in here the user ID that you copied before leaving the site:
2. Do you feel that you now understand how to use this method of exploring a large set of events?
3. Can you think of a website that would be more useful if it offered this method of exploring things like events or products? If so, please give its name or web address.
4. Compared to other ways of exploring things on the web, what's the main new advantage of this method?

After performing the tasks, the participant was asked to return to the AMAZON MECHANICAL TURK page and answer several questions, which are shown in Table 2.

Logging

The prototype is instrumented in such a way that every action performed by the user is logged in a MYSQL database, each action being labeled with a unique user ID. Since the study participants were asked to report the user ID that they had been assigned, it was possible to identify all of the actions performed by each participant (though in fact some participants did not follow these instructions adequately, as will be noted below).

4 Study 1: Results and Discussion

4.1 Completion of Milestones

For 11 of the 100 participants, the logging did not work for some technical reason; most of these were using Version 9 of Microsoft Internet Explorer. For another 18 participants, the initialization event was logged but no user actions were recorded. Since the verbal comments made by these participants made it clear that many of them did have some experience in using the prototype, it seems most likely that they simply did not correctly follow the instruction for copying and reporting their user ID. Consequently, these 18 participants are likewise omitted from the quantitative analyses below.

For each of the remaining 71 participants, their actions were analyzed as follows: The tasks described in Table 1 can be broken down into 8 small segments, each of which begins and ends with a "milestone". The first milestone is the trivial one of waiting for the initial view to load; achieving this milestone simply confirms that the user actually visited the site. The next milestone is the action of clicking on the pile of cards in one of the bottom nodes in Figure 1, which causes the search results to be displayed (cf. Figure 2). The remaining milestones can be seen on the x-axis of Figure 3; this figure shows, for each milestone, the number of participants who achieved it and the average amount of time taken to achieve the milestone. The longest gaps (e.g., the one until the first action "Open first stack") include the time required to read the task instructions, so they do not actually imply that participants were actively experimenting with the system for as long as one minute before figuring out how to perform the action.

Figure 4 (A) shows how many participants completed each number of milestones; the data points are shown along the y-axis in terms of the number of actions performed. It can be seen that a number of participants achieved only 2, 3, or 4 of the 9 milestones. It is unknown how many of these participants simply aimed to collect their modest financial reward as quickly as possible, believing that it would not be noticed whether they actually experimented with the system or not.[4] There are several comments from participants in that group which suggest that they found the interface initially unusual and confusing, to the point where they didn't even want to try exploring it. Since, as we will see, even the successful subjects often reported having initially been taken aback by the unfamiliarity of the interface, it seems that there is a threshold here that is too high for some participants, at least if they have no particular motivation to spend some time experimenting. The three milestones that were achieved least frequently—"Minimize Tampere", "Add Helsinki", and "Open Helsinki"—required the use of icons that were apparently not as intuitively suggestive as they should have been. On the basis of these results, additional attention was paid to the detailed design of these icons, resulting in the ones shown in Figures 1 and 2, and the formulations of the tooltips were also improved. It can therefore be hoped that the next evaluation will yield higher success rates for these actions.

It is interesting to check the relationship between the number of milestones achieved and the number of actions that each participant performed. Figure 4(A) shows this relationship with one data point for each participant. Figure 4(B) shows the same relationship after averaging of the number of actions for each number of milestones. Overall, it can be seen that participants who achieved fewer milestones also performed fewer actions. This result is not obvious; conceivably, many participants could have experimented extensively with the system but still not figured out how to operate it. In the graphs, we see that the less successful participants, with just a couple of exceptions, apparently gave up or lost interest after only a modest amount of exploration. So the bottleneck appears to lie more with the motivation of the participants to experiment than with the inherent difficulty of discovering the methods for operating the system.

Another factor that might affect participants success' is the size of the screen. Because of the parallelism, parallel faceted browsing sometimes requires more screen real estate than is needed for normal faceted browsing, even though the query results are shown only when requested. For participants with smaller screens, it is possible to zoom out so as to retain an overview of the entire tree (or large parts of it); but not all users may think of this possibility in the absence of instructions.

It can be seen in Figure 4(C) that there is a barely noticeable relationship (which does not remotely approach statistical significance) between screen size and mile stone achievement: Several users with quite small screens achieved almost all of the milestones, which suggests that scrolling and zooming are adequate if used appropriately. Note that it would be possible to include explicit hints about scrolling and zooming for the benefit of users with smaller screens.

4.2 Self-Assessments of Understanding

In response to Question 2: "Do you feel that you now understand how to use this method of exploring a large set of events?", 73% of the participants essentially answered "yes" and the others "no". Understandably, the latter participants were mainly ones who

[4]The fact that logging would occur was not announced, and in fact no participant was denied remuneration on account of having made inadequate effort.

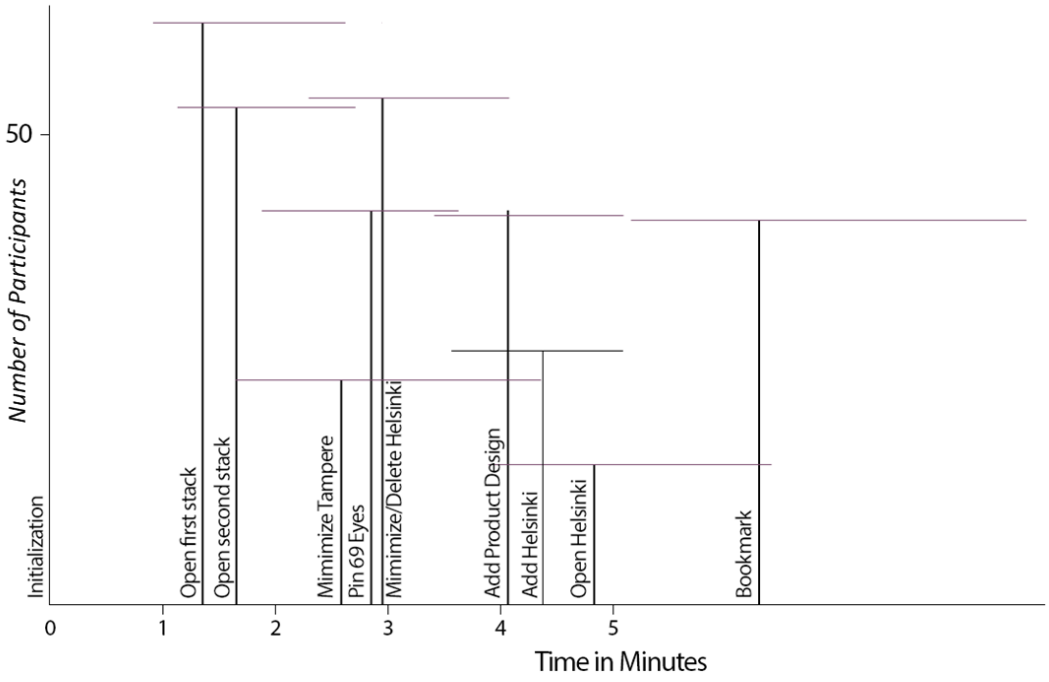

Figure 3: Number of participants who completed each milestone and the time required to do so.

(The x-axis shows the average time between completion of initialization and completion of each milestone. The height of each line shows the number of participants (out of the 71 included in the analysis) who completed it; the horizontal error bars show the range between the first quartile and the third quartile (i.e., 25% of the successful participants were faster than the shortest time and 25% were slower than the longest time. Note: The milestone "Pin 69 Eyes" logically precedes "Minimize Tampere", but the average time before completion of the latter milestone is slightly shorter; this pattern is possible because the averages are based on different subsets of participants, namely those who achieved the milestone in question.)

achieved fewer milestones. Here is a typical negative comment, from a participant who, according to the logs, performed only one action: "It was rather confusing. I feel like this could have much more easily been managed with filters, or separate webpages that were nested. (All Events $->$ Date $->$ Place $->$ etc)." It appears that some participants conclude, while simply inspecting the interface, that it will be hard to use and consequently do not seriously try it out. Note that this participant seems to think that there are no benefits relative to a normal faceted browsing interface—in contrast to the comments of the more successful participants, which are reported in 4.4.

Here are two typical responses of successful participants, which show that the initial impression of difficulty soon disappears if the participant takes the trouble to take a few actions:

- "The way to explore a large set of events was fairly self explanatory after some trial and error. I was initially not fond of the layout/setup due to the lack of a legend detailing what each symbol stood for. Also, there was not very much instruction on exactly HOW to fully utilize the website. However, after trial and error in using the website, the value became quite clear."
- "Yes, although it was intimidating when I first saw it. But following the steps that were outlined was easy. And I liked that it turned out to be easy. There was a sense of discovery which made me want to keep going."

In sum, the results are consistent with our hypotheses that the interface is quite learnable if users actually perform actions with it. But they also show that we cannot count on users being willing to try actions if they don't have immediate confidence that their actions will be successful. Combined with a number of comments of unsuccessful participants indicating a desire for explicit usage instructions, these results suggest that explicit hints should be made available for those participants who are not inclined to engage in trial and error.

4.3 Perceived Application Areas

Question 3 asked "Can you think of a website that would be more useful if it offered this method of exploring things like events or products? ..." The most frequently mentioned application area was event sites (e.g. eventful[5]), which is understandable in that the demonstrator's repository included only events. But several participants also saw applicability for online shops, social networks, and travel sites, respectively, indicating an ability to perceive the general benefits offered by the system.

4.4 Perceived Benefits

In response to the Question 4, "Compared to other ways of exploring things on the web, what's the main new advantage of this method?", some participants expressed insights that correspond fairly closely to the theoretically derived benefits mentioned in Section 1.

For example, the advantages for dealing with uncertainty about what direction to take next were summarized by one participant as follows: "It is easy to access a lot of different information and

[5] `http://eventful.com`

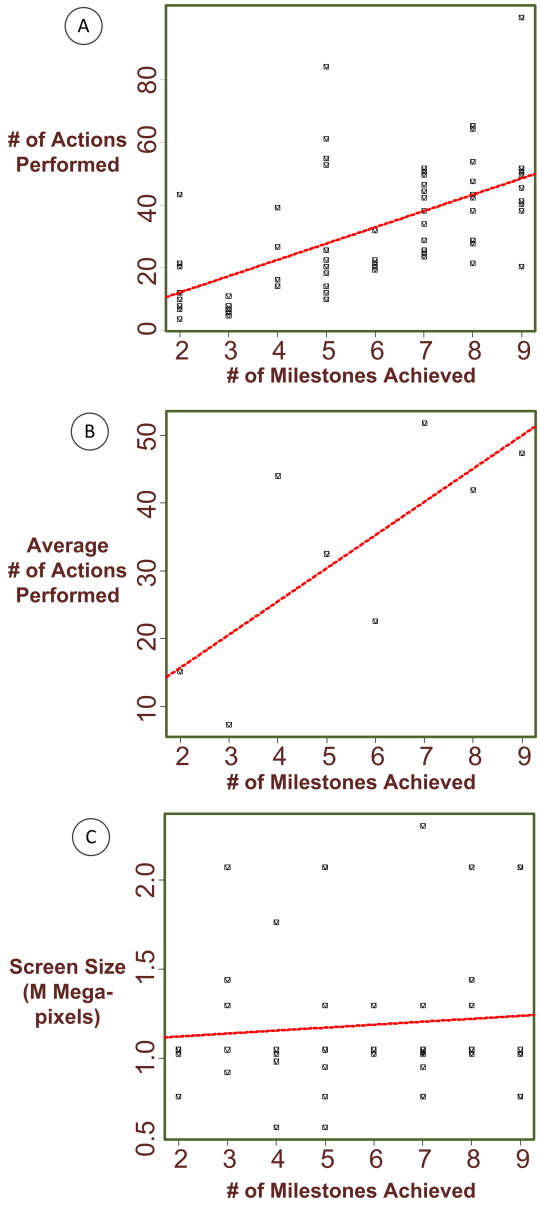

Figure 4: A and B: Differences among more and less successful users of the World Design Capital prototype in terms of their amount of exploration; C: relationship between success and screen size.
(Further explanation in text.)

retrieve it again if need be. No need to hit 'back' button or to search in one's 'history' folder. It all folds up or out all on one page." Two other participants likewise described this general benefit in terms of no longer having to use clumsy workarounds: "I like this method much more than trying to open up several websites and cross referencing myself"; "It's all on one page so you don't have to have many tabs open at once".

With regard to the possibility of looking for two or more related items, one participant wrote "... you can look at multiple things in different places and helps give you a way of figuring out how to line up multiple events ...".

With regard to the side effect of producing a structured overview of items, one participant wrote: "I like the filtering down and pinning of events .... Also then being able to share this chart with my friends to be able to collaborate together about the events that we want to attend. Very Nice!". Note that the collaboration benefit mentioned here concerns the ability to save and share a particular view, not the ability of two or more persons to work on such a view simultaneously; the participants in this study did not have an opportunity to experience this benefit. Other comments in the same vein are: "It can be shared with friends who want to attend"; and "I could easily send the link to a friend or she/he could come back to check the events she/he might like better".

Some of the other comments do not identify any of the listed benefits explicitly but do show an appreciation of the basic features of the prototype that enable them: "It can quickly overview and compare events in different places and time periods"; "All the info is in one place, I like the way you can categorize and pull up results side by side".

Many of the other comments about benefits do not explicitly identify a benefit of PFB that distinguishes it from related methods such as normal faceted browsing. For example, one participant wrote "The main advantage, for me, of using this method is how fast it is to systematically eliminate things to find exactly what you are looking for"—a benefit that, as formulated here, might also be found with normal faceted browsing.

Other participants expressed appreciation for the way in which the event repository includes events of different types, ranging from World Design Capital events to sports events: "This shows all of the events taking place in one city. It is very helpful for people who like to experience variety of things." This benefit is due to the semantic web technology supplied by the colleagues from Aalto University and EURECOM who were responsible for creating the event repository.

In retrospect, it is understandable that simply asking participants to compare the current system with "other ways of exploring things on the web" was not the ideal way to elicit very sharp comments identifying the distinguishing features of multifocal exploration; but still some of the comments do identify such features, and many of the others can be interpreted as being consistent with the postulated benefits.

5 Study 2

To give an idea of the generalizability of the results reported for the World Design Capital prototype, we will briefly summarize the results of a similar study conducted with an independent instantiation of parallel faceted browsing.

5.1 Prototype

The German-language website MY MIRACLE[6] provides information related to foods and recipes with the goal of supporting healthy eating. The site's owner, Spirescu (5), implemented and tested an experimental PFB interface[7] (see Figure 6) that supplements the more conventional methods that the main part of the site offers for searching for foods and recipes. The implementation is technically completely different from that found in the World Design Capital prototype: The database is a MYSQL database, and the user interface is implemented in PHP.

[6] `http://www.my-miracle.de`
[7] `http://parallel-faceted-browsing.com`

Figure 6: Screenshot of the MY MIRACLE PFB prototype.
(The English translations are not present in the prototype. The icon at the right-hand end of the node "Kategorie" and two other nodes enables the user to "copy and paste" an entire subtree so as to avoid having to construct similar subtrees manually.)

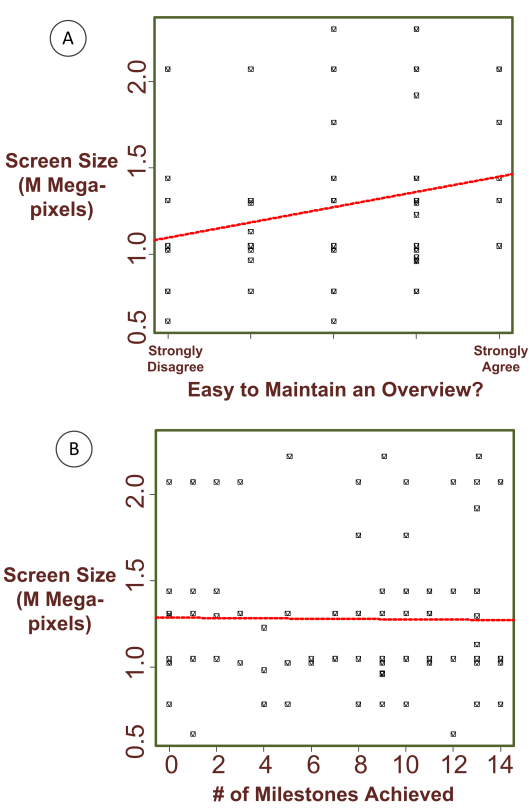

Figure 5: Subjective and objective effects of screen size in the MY MIRACLE study.
(A: Effect on agreement with the statement "The system makes it easy to maintain an overview"; German: "Ich finde das System übersichtlich"; B: effect on the number of milestones achieved.)

5.2 Participants

The participants were 116 regular users (97% female) of MY MIRACLE. They mostly lacked specialized knowledge of computers, but they were familiar with the domain.

5.3 Method

Participants typically devoted 20–30 minutes to the study, all of whose instructions were embedded within the MY MIRACLE site. They first saw an introductory explanation of the PFB prototype that comprised only a compact legend pointing to the most important functions of the prototype. They were then given step-by-step task-level instructions—roughly comparable to those in the World Design Capital study—that enabled them to try out all of the main functions of the prototype. Subsequently, they performed three less structured tasks with goals such as selecting a set of ingredients for the food to be prepared for a picnic.

5.4 Learnability

When asked whether a more detailed introduction to the prototype would be desirable, only 18% of the participants indicated that they would prefer to have more introductory information than the legend that was provided.

When expressing on a Likert scale their degrees of agreement with a number of statements about the interface, the participants generally indicated that they found the prototype easy to understand and use—though with all questions there were some participants who gave negative responses. For example, 65% agreed moderately or strongly with the statement "The system is easy to use", but 19% expressed strong or moderate disagreement.

5.5 Benefits

In an open question, the participants were asked what benefits they saw in the PFB interface. An advantage spontaneously mentioned by 27% of the participants was the ability to compare products and find suitable combinations of products. This point corresponds to the first two benefits mentioned in Section 1. A more specific benefit, mentioned by 16%, is that the prototype would be useful for planning meals—a specific instantiation of the more general benefit of being able to take into account relationships among retrieved items.

Appreciation was expressed by 23% the participants for the resulting overview of the retrieved items, which corresponds to the third benefit listed in Section 1.

5.6 Effects of Screen Size

As with the World Design Capital demonstrator, no reliable correlation was found between the number of megapixels in the participant's screen and the number of milestones that they successfully completed (see Figure 5B). When asked with a Likert scale whether they agreed with the statement "The system makes it easy to maintain an overview", the participants' answers did show a significant correlation with screen size (Spearman's rho $= .25$, $p = 0.013$; cf. Figure 5A). This discrepancy between the objective and subjective correlations is understandable: A participant who finds it difficult to maintain an overview can nonetheless perform the specified tasks, for example by scrolling and zooming where necessary. Still, it seems worthwhile to adapt the MY MIRACLE prototype's graphical representations to make them more subjectively satisfying for users with smaller screens. In contrast to the World Design Capital prototype, the MY MIRACLE prototype currently represents the query results in tables, each of which comprises several columns and consequently takes up more space than the more compact result representations of the World Design Capital prototype.

6 Conclusions

These user studies show that, despite the inevitable novelty of a user interface for multifocal exploration of semantic data, more or less experienced web users can quickly learn to use such an interface if they engage in trial and error, even if they are given little or no explicit explanation or usage instructions. On the other hand, in both studies a minority of the participants were insufficiently inclined to engage in trial and error because of initially being put off by the unfamiliarity of the new interface and/or being unconvinced that the system was worth trying out. Therefore, the most important strategies for increasing user acceptance appear to be (a) to find ways of encouraging these remaining users to try the interface out (so that they can see that it is easy to learn) and (b) to make it easier to recognize the benefits of PFB (though many users do manage to recognize these benefits), possibly through graphical representations or feedback that more directly suggest the benefits.

The benefits and application areas for multifocal exploration spontaneously mentioned by the study participants were generally consistent with those derived from a theoretical analysis, but the participants' comments mostly did not sharply identify distinguishing features of multifocal exploration—understandably, given that the participants were not in a position to compare it directly with alternative paradigms.

To the extent to which explicit recognition of the differences between paradigms is desirable, it will evidently be necessary to offer users an opportunity for side-by-side comparison.

7 Acknowledgments

The World Design Capital demonstrator was developed in collaboration with Raphaël Troncy and Houda Khrouf of EURECOM and Osma Suominen and Eero Hyvönen of Aalto University. The research and development for this demonstrator were supported by the European Union's FP7 via the projects GLOCAL (GA 248984), LinkedTV (GA 287911), ALIAS (AAL-2009-2-049); by the Finnish Funding Agency for Technology and Innovation (Tekes) via the projects SUBI and FinnONTO; and by the EIT ITC Labs via the activity EVENTMAP (TFMC 12116). Contributions to earlier versions of the interaction design of the World Design Capital demonstrator were made by New Grit, Berlin (www.new-grit.com) and by students in the 2012 course on User Interface Design at Saarland University.

8 References

[1] Sven Buschbeck, Anthony Jameson, Adrian Spirescu, Tanja Schneeberger, Raphaël Troncy, Houda Khrouf, Osma Suominen, and Eero Hyvönen. Parallel faceted browsing. In *Extended Abstracts of CHI 2013, the Conference on Human Factors in Computing Systems (Interactivity Track)*, 2013.

[2] Sven Buschbeck, Anthony Jameson, Raphaël Troncy, Houda Khrouf, Osma Suominen, and Adrian Spirescu. A demonstrator for parallel faceted browsing. In *Proceedings of the EKAW 2012 Workshop on Intelligent Exploration of Semantic Data*, Galway, Ireland, 2012.

[3] Waqas Javed, Sohaib Ghani, and Niklas Elmqvist. PolyZoom: Multiscale and multifocus exploration in 2D visual spaces. In Joseph A. Konstan, Ed H. Chi, Kristina Höök, Susanne Bødker, and Dan Olsen, editors, *Human Factors in Computing Systems: CHI 2012 Conference Proceedings*. ACM, New York, 2012.

[4] Aran Lunzer and Kasper Hornbæk. Subjunctive interfaces: Extending applications to support parallel setup, viewing and control of alternative scenarios. *ACM Transactions on Computer-Human Interaction*, 14(4):17, 2008.

[5] Adrian Spirescu. Parallel faceted browsing: Design, implementation, and evaluation, 2013. Bachelor's thesis, Department of Computer Science, Saarland University (in German).

[6] Daniel Tunkelang. *Faceted Search*. Morgan & Claypool, Palo Alto, CA, 2009.

[7] Ryen W. White and Resa A. Roth. *Exploratory Search: Beyond the Query-Response Paradigm*. Morgan & Claypool, San Francisco, CA, 2009.

[8] Max L. Wilson, Bill Kules, m. c. schraefel, and Ben Shneiderman. From keyword search to exploration: Designing future search interfaces for the web. *Foundations and Trends in Web Science*, 2(1), 2010.