mSpace: What do Numbers and Totals Mean in a Flexible Semantic Browser

Max L. Wilson, m.c. schraefel

IAM Research Group
School of Electronics and Computer Science
University of Southampton, UK
{mlw05r, mc}@ecs.soton.ac.uk

Abstract. With the Semantic Web community's growing interest in Human Computer Interaction, this paper addresses a challenge for user interface design and future shifts in search paradigms. Where browsers using current search paradigms often use numeric values to indicate volumes of sub-hierarchies, future semantic browsers will not be limited to fixed hierarchical datasets, but allow flexible exploration through multiple intersecting domains. With the future use of similar numeric indicators uncertain, research here suggests that the inclusion of such indicators should be based around focal data objects within each information domain. Further research is required, as a significant number of contradicting participant expectations were present. It is the concern of the Semantic Web community to make sure that future browsers of alternative semantic search paradigms can best support their users.

1 Introduction

For the first time in the ISWC conference's 5 year history, there are two explicit Human Computer Interaction topics listed[1] where authors assign categories to their papers for review: 'Applications of Semantic Web technologies with lessons learned' and 'User-centered Semantic Web applications and/or interaction design'. These explicit categories demonstrate the Semantic Web research community's growing interest in and understanding of the role of the interface – what is becoming the implicit Presentation layer at the top of the Semantic Web Layer Cake[2] - in investigating the potential for Semantic Web technologies to deliver new models for exploring information on the Web. Indeed, the opportunity for information on the (Semantic) Web to be explored across associations between graphs, enabling one to move from exploring topics in Classical Music to say the History around a particular Composer, challenges the current mode of exploration within one data set only. Apple's Music Store, for instance will only show someone links to Beethoven recordings; its fixed database does not enable one to explore information about Beethoven or about political activity in Europe at the time of Beethoven. The Semantic Web can enable such

[1] http://iswc2006.semanticweb.org/submissions/res_ac_track.htm
[2] http://www.w3.org/2002/Talks/04-sweb/slide12-0.html

associative connections. The question arises, therefore, as to how well current design heuristics for Web-based navigation support the new search possibilities afforded by the Semantic Web. One such example of a heuristic challenge posed by the Semantic Web is the use of numbers, which follow categories returned by search results.

Current search paradigms, such as category-based or facetted browsers, are targeted towards specific data. In Figure 1 we see an example e-commerce site: eBuyer.com[3]. A user selects an intended target, such as USB Flash Drives. Once the target is chosen, the interface presents various levels of a hierarchy that are orientated towards that target. With the target as 'USB Flash Drives' in Figure 1, optional subcategories appear, showing various manufacturers. Selecting a manufacturer then reveals the various models the manufacturer produces. After each category listed is a number. In this paper such numbers are called Numeric Volume Indicators (NVIs). NVIs are based upon the initial selection of a target. NVIs in the USB Flash Drive case reflect how many of these objects can be found through each category. It has become simply accepted as good design practice to use NVIs to let people know how many objects, which match a given search, can be found in each category. Besides stores, we find this approach used in facetted browsers of digital library collections like those modeled in Flamenco [4] and Longwell [9]

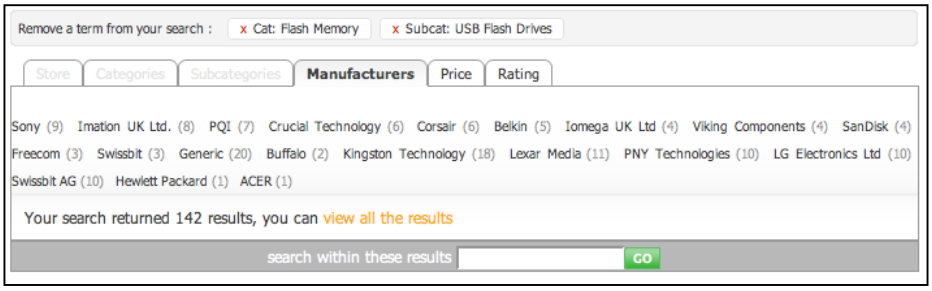

Figure 1: The search interface from eBuyer.com, showing NVIs for each manufacturer that represent the number of USB Flash Drives.

NVIs work well when one initiates a search for a specific thing, whether this search is started by typing in a keyword, as in the above choice, or via direct manipulation [1][6], as in Flamenco's UC Berkeley Architecture Slide Library Search demo[4] where one picks a specific object, like 18th Century Books. In the latter case NVIs show how many 18th C books are in each of the following categories. Each of these search examples, however, assumes that the search is initiated with a specific target in mind. There are other forms of search, however, that do not assume that one starts with such a specific goal. Exploratory search paradigms [10] suggest that one may be interested in a domain rather than an entity within a domain, or that one may be interested in exploring relations within and among domains rather than looking for specific instances within a domain. The Semantic Web has the potential to support exploration of such richly associated relationships within and among domains. In a movie domain, for instance, which may let people explore the information space by genre,

[3] http://www.ebuyer.com/

[4] http://orange.sims.berkeley.edu/cgi-bin/flamenco.cgi/spiro/Flamenco

period, director, actors, country (of actor, director, film, etc) the typical target approach to NVIs would be for to NVIs represent number of movies. But people exploring a domain may want to know how many Directors were producing Film Noir projects during World War 1, rather than the more typical – how many Film Noirs were made in a given period. Similarly with the possibility of connecting other domains such as History for example with Film domains, one might want to know more about how many films were part of propaganda campaigns used by Germany or the US during WWII. Once representations of information spaces support exploration as well as specific, targeted search, a question becomes what do NVIs represent if one's approach to exploration does not start with either a specific target, or where the target may vary? Clearly, the interface could be more explicit about what NVIs are counting: 18 Movies, but which columns should be counted and how many columns should be represented? It is also a good design heuristic to allow expert users to control which column is counted, but the question here refers to what novice users expect the software is trying to convey. As far as we have been able to discover, such target differentiation is ignored by interfaces that present NVIs: such interfaces clearly have the Target/Number of Targets in Category and Subcategory approach firmly established as the exploration paradigm for their collections. Flamenco, Endeca, Longwell, and various online stores are examples of this.

Over the past several years we have been working with the mSpace browser – a flexible Semantic Web oriented facetted browser, described in Section 3 below, that enables certain manipulations of categories that preclude easy assignment of NVIs to elements in its categories because they can be rearranged from one or more hierarchies or even polyarchies [5]. Since it is obvious that NVIs provide valuable information for making decisions about both the domain in general, and where people may wish to explore next in particular, we wanted to look at what assumptions people using such a flexible browser like mSpace might make about what NVIs represent were we to include them. In other words, we wanted to understand the "mental model" people constructed about the information space when NVIs were present. To this end, we ran a lightweight paper prototype study described below, to explore how NVIs are interpreted in a flexible exploratory search interface. We follow a discussion of findings with interim design guidelines for Semantic Web exploratory search interface designers.

2 Related Work

A number of projects have tried to break away from the traditional category-based hierarchy. eBuyer.com, along with projects like eBay Express[5], are types of facetted browsers, which do not follow a traditional hierarchy but allow the user to filter the results by categories in any order. Marti Hearst's work on the Flamenco browser is at the forefront of this research [4], but it is still based on the targeted search paradigm; it also does not use semantic technologies. Realistically, most projects in the Semantic

[5] http://express.ebay.com/

Web community are still targeted at finding particular items of information. Endeca[6] is a commercial project that uses semantic technologies to support facetted browsing, yet its aim is still to find particular target objects; as a result it is becoming popular in e-commerce. Longwell is MIT's Semantic Web browser, but is also target oriented and provides NVIs to support existing search paradigms [9].

Research such as that by Quan and Karger [7] considered how an arbitrary semantic structure might be accessed. The aims of such projects are to support the exploration from topics such as film to alternative topics such as history, music and architecture. It is this kind of vision for search paradigms that require simple design heuristics to be re-addressed.

3 The mSpace Interface

mSpace is, in part, a flexible browser that makes use of affordances provided by the Semantic Web and is used in this study as a test-bed interface for presenting potential alternative NVI implementations. This interface is particularly suitable for this study as it supports the flexibility described above in future exploratory search paradigms. mSpace is defined in detail in previous publications [2][8], but a core element of the interface is the columnar browser highlighted in Figure 2. Suffice it to say that the data storage is in RDF and access to it is provided by 3store [3]; information can, therefore, permit tangents into potentially any related domain through the directed-graph RDF relationships.

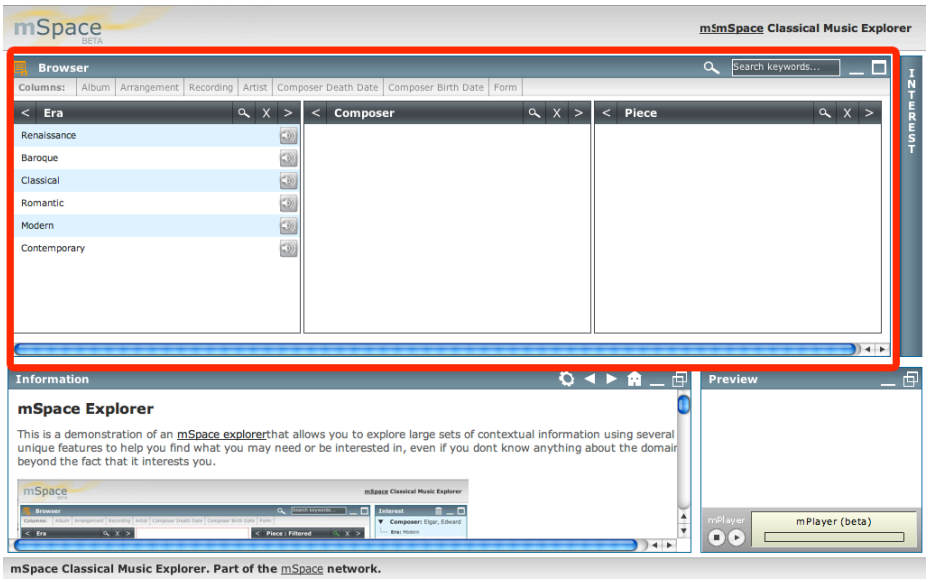

Figure 2: The mSpace browser, with the columnar section highlighted.

[6] http://www.endeca.com/

The goal of mSpace is to make it possible for people to explore rich (or high dimensional) spaces without having to be domain experts or even familiar with the domain lexicon. Part of the problem to support such a goal is to be able to represent n-dimensional spaces for easy access. High dimensional spaces are hard to visualize on a 2D screen, so mSpace imagines a projection of these dimensions onto the space, flattening it to create a slice. These slices create, effectively, a temporary hierarchy of the subset of dimensions in the domain. The dimensional slice is represented as a multi-column viewer. These columns can be rearranged, removed and new columns/dimensions added at any point in the slice to allow users to have this freedom of exploration described above. Essentially, the user can focus the exploration towards any item of information that they currently care about.

Interaction with this slice is as follows. Selecting an item from the list in the first column, with a single click, produces a filtered list in the second column. In Figure 2, selecting an Era produces a list of Composers, *from that Era*, in the second column. Similarly, selecting a Composer in the second column then produces a list of pieces in the third column (Piece), *from the selected Era and by the Selected Composer*. These columns can be rearranged freely and so to see all of Bach's Cello pieces, the column 'Arrangement' can be added between the second and third columns. Then to see all of the composers who produced Cello music in the Baroque Era, the Arrangement column can be moved to be before the Composer column. The interface will then show a list of composers*, from the Baroque Era who have produced Cello Music*.

mSpaces are currently domain specific. That is, an mSpace will cover Classical Music, Film or Computer Science Research in the UK. Our goal is a more generic browser where Composer in a Classical Music context can shift seamlessly to Historical Figure in a Political context.

4 Study

With the mSpace approach, there are 3 attributes NVIs could represent: 1) number of items in the next column; 2) number of items in the column that the information is oriented around, such as Movies; and 3) the number of items in the last column, if the last column is understood as a "goal" for the exploration. These approaches are explained more below in detail, and inform what we postulate as the Next Column, Domain Focus and Dynamic Target hypotheses, respectively.

With the Next Column hypothesis, the number indicates the number of items in the next column to the right. As there is no fixed target and no foreseeable end-point to browse towards in these new search paradigms, the furthest guaranteed point that one could base NVIs upon is potentially only the next column. NVIs are then indicating the broadness of the connections from each node in the graph.

In the Domain Focus hypothesis, the number would indicate a total of some object to which the majority of information is oriented towards. Considering an IMDB[7] dataset, where the domain focus is centered on movies, the figures would then always represent the number of Movies in each category. Although the user may be con-

[7] Internet Movie Database – http://www.imdb.com/

cerned with Actors, their search is then routed in the focus of the domain. This has a further interesting question – should these numbers be (a) filtered with the columns or (b) always be the total number of movies for that category. In situation (a), the Country column might indicate the number of Film-Noir movies produced in, say, France. In situation (b) the NVI shows the number of films in France and ignores the previously selected Genre.

Finally, the third competing Dynamic Target hypothesis would base the figures on the furthest right column. In mSpace, the furthest right column shows the results of all the previously selected filters. If Director is the furthest right column, then it could be potentially used to decide that Director is the user's current target, rather than Movie. The key point of this case is to dynamically determine the user's intentions, based upon their customized layout. However, this assumes that the user is purposefully using the entirety of a slice, and also *requires* the user to remove, rather than ignore, uninteresting columns in order to get the NVI that they require.

4.1 Study Design

From these three options, a study was designed to gather responses to two questions. First, the study was designed to gather immediate expectations as to the information represented by NVIs in an exploratory search interface with flexible targets. Second, the study was designed to tease out the impact of various cues on participant rationale for interpreting NVI meaning. To this end we ran a between-groups, with matched groups, study with each group exposed to one of 4 interface conditions, each condition presenting NVIs in distinct configurations. We used this between-groups approach in order to ascertain both first impressions of what the NVI represented and subsequent rationalization for interpreting the NVIs in that way. We could then correlate if and how similar or different mental models were at play across conditions.

20 people participated, 5 people in each condition. The groups were matched for computing experience and previous experience with mSpace.

4.2 Interfaces

Each interface was given the same slice of columns: Genre, Country, Movie, Director so that the order supported the Dynamic Target hypothesis, as the slice was not oriented towards Movie; the participant was not told it was a movie domain, but this column order was specifically mentioned, including that it ended with Director. To support the Domain Focus and Next Column hypotheses, the NVIs were designed so that they alone could not lead participants to decide meaning. For example, the number shown in the Genre column could as easily be the number of Countries (the Next Column hypothesis) or the number of Movies (the Domain Focus hypothesis).

The base condition (Figure 3) is the simplest of the four and like current category browsers, simply provides a number beside each element in a list; this interface provides no cues as to what the NVIs could represent. This was designed to grasp the most instinctive answer from participants. The remaining three interfaces are designed

to provide slightly more information to the user to see if simple cues affect their instinctive responses. Essentially, they are designed to test the rigidity of the users' expectations. Interface condition 2 (Figure 4) is designed to provide a simple cue that removes the possibility of NVIs representing the number of items in the next column. As the entire list is visible in the Country column the user can clearly see that it is less than 86. We have named this the *Visible Size Cue*.

Figure 3: The base condition, no selection made and no cues provided as to the value of the NVI.

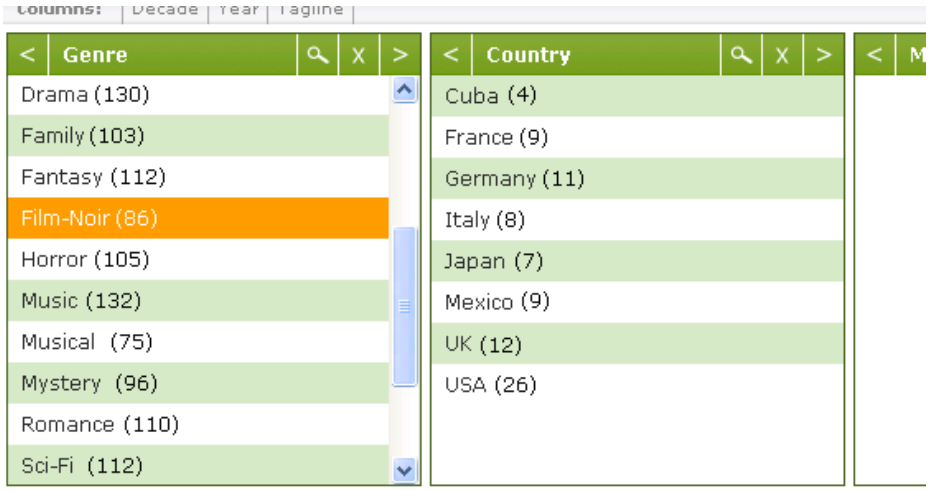

Figure 4: Interface Condition 2 has a single selection in the first column and a completely visible list in the second column; this includes a useful Visual Size Cue.

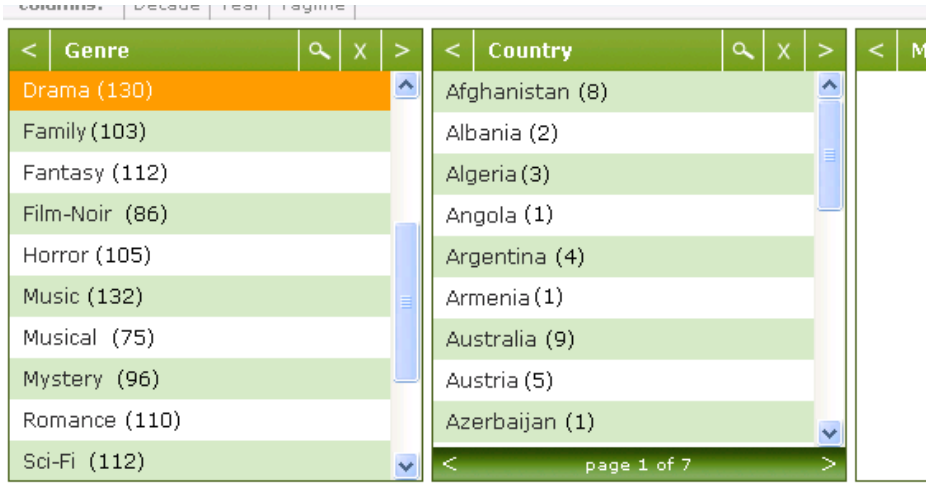

Figure 5: Interface Condition 3 has a single selection in the first column and then a partially viewable list in the second; this includes a useless Visual Size Cue.

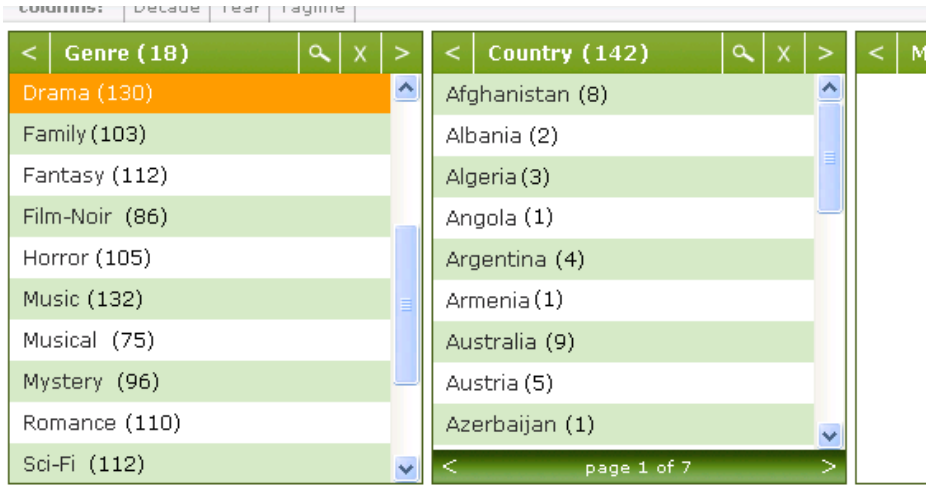

Figure 6: Interface Condition 4 is the same as Interface Condition 3 but with NVIs that indicate the number of the items in each populated column; this includes a useless Visual Size Cue and a useful Numeric Size Cue.

Interface condition 3 (Figure 5), however, is designed to emphasize the ambiguity of the NVIs; the Visible Size Cue is no longer useful as the list in the Countries column is potentially as large as the NVI for the selected category: Drama. Finally, Interface Condition 4 (Figure 6) is designed to show a different cue, indicating that the NVI does not represent the number of Countries. This interface also includes totals at the top of each column, showing the number of items found in its list. This column total is inconsistent with the selected NVI for Drama, and so the NVI cannot mean that the

number of Countries that have a Drama movie; this cue provided by the column totals we have called the Numeric Size Cue.

4.3 Procedure

Participants were presented with an example mSpace for exploring a conference schedule and were walked through how slices and selections work. This interface had no numbers assigned to any of its attributes. Participants were then told they would be considering a similar mSpace, which allowed them to explore information about Movies, Actors and Directors. We characterized the domain space with these three possible foci rather than saying the domain is about movies so that participants would not be led to that one particular attribute by our biasing the way the space was presented. As the columns were presented, the specific order, with director last, was stated to make sure each participant had addressed the exact sequence. Once ready, participants were shown one of the four paper prototype interfaces and asked to suggest what the NVIs shown should represent. Once a specific answer had been given, the user was engaged in discussion with the three options we had identified and invited to suggest further alternatives.

4.4 Results

Figure 7 shows the number of answers across the different participant groups (interface conditions). All participants shown the Base Condition chose the Domain Focus hypothesis; similarly for Interface Condition 2, where the ambiguity was reduced. Where the ambiguity was increased, through a useless Visual Size Cue (Interface Condition 3), responses were more mixed with 60% of participants choosing the Next Column hypothesis. In Interface Condition 4, which has the ambiguity of condition 3 slightly removed through the Numeric Size Cue; only 1 of the 5 participants chose Country for the NVIs to represent (the Next Column hypothesis).

No participant chose the NVI to represent Director, even though Director was the last column in the slice; the Dynamic Target hypothesis also was not chosen. Of the remaining responses, across all interface conditions, 80% immediately expected the NVIs to represent the number of films (the Domain Focus hypothesis). 20% of participants, across all conditions, assumed the NVI to be the number of items in the next column (the Next Column hypothesis); each of these participants rated themselves as having little or no previous experience with mSpace.

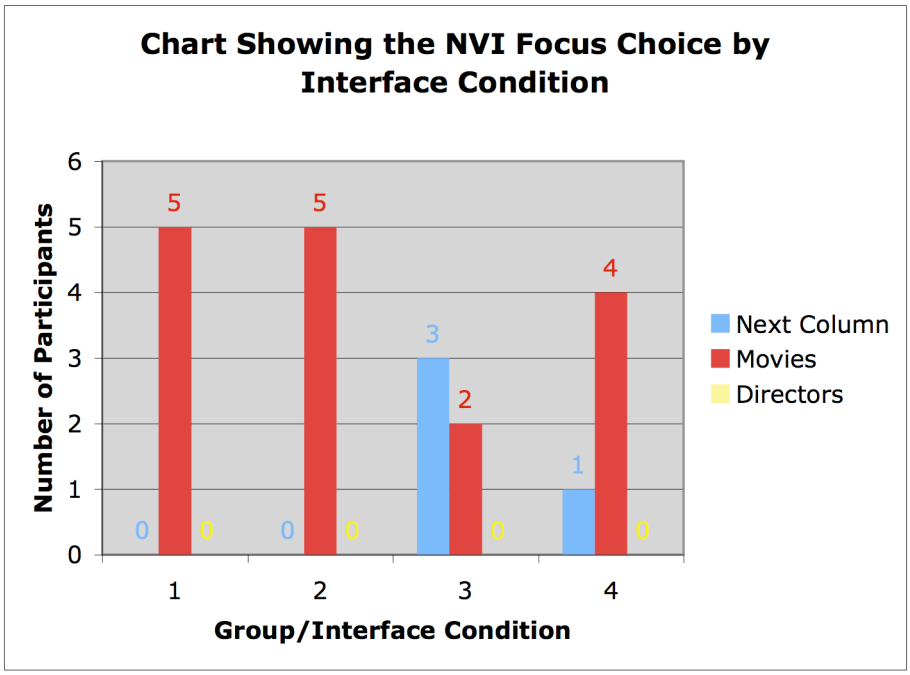

Figure 7: Chart Showing the NVI Focus Choice by Group/Interface Condition.

4.5 Observations

The results above indicate the first responses given by participants, however three participants, from various conditions, entertained all three hypotheses noted above in follow-up discussions; no fourth option was presented by a participant. Although not the immediate expectation, few participants rejected the potential of the alternative options for what the NVI might represent. Interestingly, most were happy to accept the Dynamic Target option, but if they had selected Domain Focus for the NVIs to represent, they were less willing to accept the Next Column hypothesis, and *visa versa*. Further, those who had selected the Next Column hypothesis could rarely verbalize reasoning for their own expectation. In contrast, those who selected 'Movie' as the item represented by the NVIs were most adamant that, as the focus of the dataset, browsing other information about Directors, for example, was still oriented around the Movies themselves. Some more concise reasons noted that all of the possible columns could be attributed to movie, but not all to any other columns. For example, while a Director may have a Genre and a Country, a Country does not have either a Genre or a Director. Although Director does appear to be quite atomic within the data, the most consistent root of all the information was a Movie.

Perhaps the most concise argument given by a participant against the Next Column hypothesis was that it provides no useful information to the user. In the Next Column

hypothesis, the NVIs do not indicate the total of some target object but represent the subsequent breakdown of those targets at the next level.

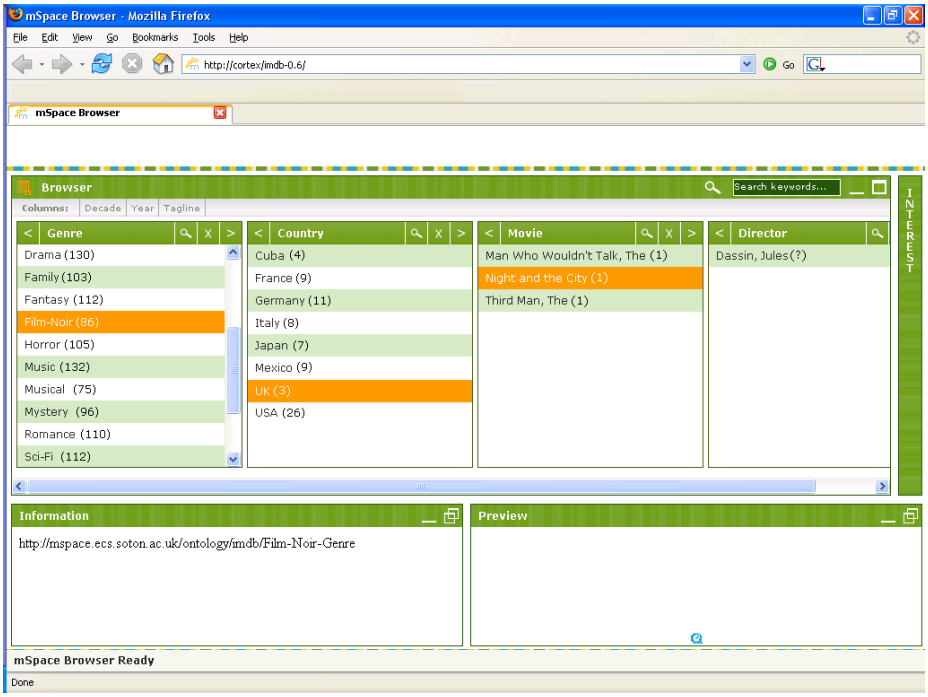

Figure 8: A conceptual slice with NVIs, representing the problem participants had with the Domain Focus hypothesis when the domain focus is not the most right column.

Finally, it is important to note that, although 'Movie' was the most popular answer for what the NVI represented, participants were unable to calculate the value of the NVI in columns that are right of the NVI focus. In respect to the conceptual image shown in Figure 8, the NVIs in the Movie column are all of value 1, but no user was able to imagine the NVIs for the Director column. Logically these values should be of value 1, as expressed verbally, the query represented by these column selections represents "The number of films directed by A, with the movie name B, from the country C and in the genre D." All NVIs right of the NVI focus column would therefore be 1, unless NVIs were not filtered by previous column selections; although this unfiltered possibility was not entertained by participants.

5 Discussion

With no participants choosing the Dynamic Target hypothesis, it is reasonable to suggest that NVIs should not be dynamically based upon what the final column in the slice may be. There seems to be no correlation between that column position and user

assumptions about how that column may or may not affect the NVI. We may need to consider, however, that these initial impressions of what these single interface shots represent as presented to the participants in the experiment may still be based upon familiarity with current search paradigms in which categories represent a fixed target. Our study only presented one state in an interaction. A more detailed study which enabled participants to step through a sequence of interactions to build up an exploration that specifically asked for information among categories (how many Russians were making movies in the US during the McCarthy era, for instance) might push on participants' mental models about the space in ways that would more explicitly register a desire to have NVIs be selectable. This notion is supported by the participants' acceptance of the Dynamic Target hypothesis after more detailed discussions within the debriefing sessions.

One of the most pertinent arguments given against the Next Column hypothesis was that it does not indicate the volume of a space, but instead the broadness of its subcategories. This point notes that the Next Column hypothesis effectively contradicts the original NVI heuristic to indicate some volume of subsequent search space, rather than representing the number of ways in which that top category is divided into subsequent categories. Dismissing this possibility for previewing information, however, may ignore the shift in search paradigm that is being afforded by the Semantic Web enabled browsers like mSpace, where different kinds of queries about the types of information that is available are possible. It may be valuable for a searcher to know that Romantic music breaks down into 60 subdivisions while Baroque has only 3. That said, it was still the majority expectation, across interface conditions that the NVIs represent movie, supporting the Domain Focus hypothesis.

Two cues were tested in the above experiment, and it is clear that both can be used to reduce the ambiguity of situations. Although the Visual Size Cue can be very clear, as in Interface Condition 2, the results from Interface Condition 3 show that it can also cause greater ambiguity for the user. Participants were unanimous in condition 1, where the NVIs could have represented either column, we expect that the user guessed 'the most likely answer'. We believe that the Visual Size Cue supported the user in Interface Condition 2, but this cannot be proven from our results. However, with Interface Condition 3, participants were clearly put in doubt. In Interface Condition 4, however, the combination of Visual Size Cue and Numeric Size Cue removed this ambiguity. Interestingly, though, we see that one of the participants was still drawn to the Next Column hypothesis in Interface Condition 4, indicating that the Numeric Size Cue is weaker than when the Visual Size Cue, alone, if useful.

6 Conclusions and Future Work

There are three possible conclusions to be drawn from this work. First, it is clear that the majority expectation on first encountering an interface with NVIs is that NVIs represent (1) a number of specific artifacts (not categories) and (2) that these artifacts represent a domain focus. Thus where there is a *known* domain focus, most participants use that focus to determine what the NVIs represent. Second, when ambiguity

about exploratory situations is introduced – that is, when the query of the exploration becomes less domain artefact specific (queries about directors rather than about movies in a movie domain) participants clearly entertain the possibility of different representations. This finding suggests that interfaces may need to implement user-controlled NVI switching. 20% of the sample population in our study, for instance, had a contradicting view to the majority about NVIs. Through cognitive walkthroughs of the interface beyond the initial exposure condition, we saw that some users wanted to have a Dynamic Target NVI. Third, combining both Visual and Numeric Size Cues seems to provide additional benefit beyond that of than NVIs alone. Indeed, based on these findings, it is our intention to develop an interactive extension for the mSpace interface that includes Visual Size Cues and also allows the user to choose the dimension/target focus for the NVIs. We also intend to further explore the acceptance of a Dynamic Focus approach, compared to a Domain Focus approach.

This work also shows that the mental model for including NVIs within flexible target browsers is not completely understood. Although we know a majority expectation, we do not know how expectations about NVIs will change during longer interaction with browsers that take advantage of Semantic Web associations within and among domain information. Consequently, it becomes clear that the simple heuristic of NVIs = Domain-Focus-as-Target may not be sufficient to support richer interactions that Semantic Web explorers may afford; indeed, that heuristic may unduly restrict richer mental models about exploratory possibilities that these interfaces support. As such, it will be important for Semantic Web researchers and developers designing interfaces for their services to keep the issue of affording multiple NVI possibilities in mind. Indeed, the question of new heuristics for NVIs in Semantic Web explorers is likely just the tip of the iceberg of design heuristics, which will need to be reconsidered for optimal effectiveness within the Semantic Web community. These questions provide an interesting opportunity for interdisciplinary work between semantic Web systems and Semantic Web interface researchers. They also serve to underline the current ISWC's acknowledgement of the role of UI research within the Semantic Web community itself, and suggest it may be time to make explicit in the Semantic Web Layer Cake the once invisible and now implicit Presentation Layer.

7 Acknowledgements

Thanks to all participants who contributed to this study, along with Daniel Smith and Alistair Russell who have provided great support within the mSpace project.

8 References

1. Chimera, R., and Shneiderman B. An exploratory evaluation of three interfaces for browsing large hierarchical tables of contents. ACM Transactions on Information Systems (TOIS), 1994.

2. Harris, C., Owens, A., Russel, A. and Smith, D. A. (2004) mSpace: Exploring The Semantic Web. A Technical Report in Support of the mSpace software framework. Technical Report, ECS, University of Southampton.
3. Harris, S. and Gibbins, N. (2003) 3store: Efficient Bulk RDF Storage. In Proceedings of 1st International Workshop on Practical and Scalable Semantic Systems (PSSS'03), pp. 1-15, Sanibel Island, Florida.
4. Hearst, M., Elliott, A., English, K., Sinha, R., Swearingen, K., and Yee, K. 2002. Finding the flow in web site search. *Commun. ACM* 45, 9 (Sep. 2002), 42-49.
5. McGuffin, M. J. and schraefel, m. c. (2004) A Comparison of Hyperstructures: Zzstructures, mSpaces, and Polyarchies. In Proceedings of ACM Conference on Hypertext and Hypermedia, 2004, pp. 153-162, Santa Cruz, California, USA.
6. Plaisant, C., Shneiderman, B., Doan, K., and Bruns, T. Interface and data architecture for query preview in networked information systems. ACM Transactions on Information Systems (TOIS), 17(3):320 - 341, 1999.
7. Quan, D. A. and Karger, D. R. 2004. How to make a semantic web browser. In Proceedings of the 13th international Conference on World Wide Web (New York, NY, USA, May 17 - 20, 2004). WWW '04. ACM Press, New York, NY, 255-265. DOI= http://doi.acm.org/10.1145/988672.988707
8. schraefel, m.c., Smith, D.A., Owens, A., Russell, A., Harris, C. and Wilson, M.L., The evolving mSpace platform: leveraging the Semantic Web on the Trail of the Memex. Hypertext (HT05), (2005). 174-183
9. SIMILE: Longwell RDF Browser (2003-2005) http://simile.mit.edu/longwell/.
10. White, R. W., Kules, B., Drucker, S. M., and schraefel, m.c. 2006. Introduction. Commun. ACM 49, 4 (Apr. 2006), 36-39.