Exploring cognitive issues in visual information retrieval

Timothy Cribbin & Chaomei Chen

VIVID Research Centre, Brunel University, Uxbridge, UK

Timothy.Cribbin@brunel.ac.uk

Abstract: A study was conducted that compared user performance across a range of search tasks supported by both a textual and a visual information retrieval interface (VIRI). Test scores representing seven distinct cognitive abilities were examined in relation to user performance. Results indicate that, when using VIRIs, visual-perceptual abilities account for significant amounts of within-subjects variance, particularly when the relevance criteria were highly specific. Visualisation ability also seemed to be a critical factor when users were required to change topical perspective within the visualisation. Suggestions are made for navigational cues that may help to reduce the effects of these individual differences.

Keywords: information visualisation, information retrieval, cognitive ability

1. Introduction

Given the increasing accessibility and volume of on-line information, there has been much research and development effort towards the implementation of visualisations as a means of assisting users in their comprehension of large datasets. It is argued that by representing the semantic properties of information using spatial dimensions or clustering techniques, one can reduce the gulf between the interface and the users conceptual model. Furthermore, access to a graphical overview may reduce the cognitive demand experienced when exploring large unfamiliar information spaces, by reducing memory load, maintaining orientation and capitalising on the relative ease with which spatial patterns can be visually perceived. A good visualisation should assist the user in their comprehension of a corpus by imposing a logical semantic structure that they may otherwise find it difficult to create by themselves. Cognitive, particularly spatial, ability is a major source of variance in this respect as it is with many areas of human-computer interaction (HCI). Egan (1988) commented that, whilst the importance of spatial ability as a predictor of contemporary HCI performance was undoubted, what was unclear was *which* cognitive factors, generally, were most influential. This is still true today, particularly with respect to visual information retrieval. A review of the current literature still shows a stark inconsistency between evaluation studies, not just in terms of psychometric measures but also with respect to task, data and population types used.

This paper will examine the importance of a range of different cognitive factors, spanning three cognitive domains. Our data show that visualisations created using the minimum spanning tree (MST) algorithm consistently provide higher levels of support to users across a range of information seeking tasks when compared to the more traditional simple scatter graph (Cribbin & Chen, 2001). Performance scores, however, show large amounts of within-subjects variance with respect to the MST based interface. This paper aims to ascertain how much of this variance can be accounted for by cognitive ability. Focusing on a comparison between the MST based interface and a more traditional textual interface, it is hoped that the results of our regression analysis will provide clues as to the nature of most cognitively demanding activities involved when conducting a visual information search.

2. Individual differences in information retrieval

Individual differences in interface usability have long been of interest to the HCI community. This is particularly true for the field of information retrieval. Previous studies have generally focussed on a fairly narrow range of cognitive factors. In an effort to provide good coverage, tests representing seven factors were selected for this study from across three generally recognised cognitive domains: Visual-perception, learning and memory and idea production. What follows is a brief description of each domain, their associated factors and their known relevance to visual and non-visual information retrieval. All tests used in this study were drawn from the ETS Kit of Factor Referenced Tests (Ekstrom, French et al., 1976).

2.1. Visual-perception

'Spatial ability' is a general term that actually describes a range of statistically distinguishable factors belonging to the broader cognitive domain of visual-perception (see Carroll, 1993). The visual-perceptual domain breaks down roughly into five factors – Visualisation, spatial relations, closure speed, closure flexibility and perceptual speed (P) (Carroll, 1993). Little attempt has been made to deconstruct spatial ability within the context of a consistent experimental paradigm (although see Hook et al., 1996).

Visual information retrieval interface (VIRI) research has historically tended to focus on spatial visualisation (VZ) ability (Westerman & Cribbin, 2000b). VZ is a cognitive factor that relates to a person's ability to mentally manipulate objects within space. It seems particularly predictive of the ease with which users are able to construct a mental model of an information (Stanney & Salvendy, 1995). In our experiment this was measured using the paper folding/punching task (VZ-2).

Tests of spatial relations or 'spatial orientation' are described in the ETS factor kit as measure "the ability to perceive spatial patterns or to maintain orientation with respect to objects in space" (Ekstrom, French, Harman, & Derman, 1976). S-1, the card rotations task, is used here.

Scanning large amounts of text is an intensive perceptual task. Allen (1992) found those with low perceptual speed produced lower quality search results because they were slower at drawing out the key terms. Grouping documents together conceptually enables a pan and zoom strategy that drastically reduces scanning demands. Once the user has located one or two interesting documents, they can simply concentrate their browsing efforts on that locality, rather than scanning through the entire list. Studies show that that by placing documents within some meaningful context, significant benefits are returned with respect to retrieval speed and that magnitude of these benefits are dependent on both the fidelity (Westerman & Cribbin, 2000b) and type of contextual information (Chen & Dumais, 2000).

2.2. Learning and Memory

To navigate a visual information space effectively, the user must not only retain a good overview of the structure, but must also remember which themes each of the various visual features represent. Tests of visual memory (MV) and associative memory (MA) are both intuitively relevant in this respect.

The MA test used here, MA-2, measures the testee's ability to remember imposed relationships between pairs of unrelated data items, a number and an object. (Westerman & Cribbin, 2000a) found that users who were high in this ability were able to capitalise on spatial-semantic visualisations when thematic clustering was clear and found target items (animal names) faster than the low group. Interestingly, when spatial location was arbitrary, this pattern was reversed with the high group performing more poorly than the low group. User comments suggested this could be attributed, in part, to differences in search strategy between the two groups with the high ability users relying more on cognitive maps of the space. The low group, on the other hand, seemed to opt for a more systematic strategy, perhaps better suited to an arbitrary layout.

MV is a psychometrically difficult factor to classify. Some experts place it within visual-perception whereas others assign it to learning and memory (Carroll, 1993). This study uses the test MV-1. Whilst having face validity, little is known about it's relevance to VIRI navigation.

2.3. Idea Production

The two final measures included in our battery were both fluency measures. associational or verbal fluency (FA) is "the ability to produce words from a restricted area of meanings". In this case, test FA-1 requires the testee to think of synonyms for a number of sample words. Swan and Allen (1998) found that verbal ability, although not strongly predictive, was still likely to be a primary factor as far as cognitive abilities are concerned.

Finally, ideational fluency (FI), which measures "the facility to call up ideas wherein quantity and not quality of ideas is emphasised". FI-1 is used here, which requires the testee to generate ideas associated with a given topic. This factor is untested by the IR community, but has good face validity and is sufficiently distinct from the other included factors to warrant investigation.

3. Accommodating individual differences

Designing an interface that both minimises the effects of dominant cognitive factors whilst still providing incremental performance benefits for all users is a highly valuable objective. The 'robust' interface is diametrically opposed to the training approach. (Egan, 1988, p.558) Three steps are involved. First one must ascertain which user characteristics predict the biggest differences in performance. Second, one must isolate the sources of variation at the task or interface design component level. Finally, the offending task or interface features can be redesigned in order to maximise the benefits for all groups.

3.1. Hypotheses

This paper will focus on step one of this approach. The effects of cognitive ability will be studied at the by task and by interface levels. We found that the benefits of using the MST interface, over the textual interface, emerged as the structured tasks progressed (Cribbin & Chen, 2001). This suggests that users were gradually understanding the structure and using it to maintain focus within the space. Due to large amounts of within-subjects variance, however, few of the observed differences were found to be significant. *H1 therefore is that much of this variance can be explained in terms of individual differences in cognitive ability.*

It is expected that the cognitive processes involved in visual information retrieval will be very different to those associated with retrieval using textual lists. Visualisation ability is known to be a good predictor of retrieval performance when navigating non-spatial interfaces, principally because of its relationship to mental model construction. As the main promise of spatial-semantic VIRIs is to reduce this load, *H2 is that visualisation ability will become less important.*

Although spatial-semantic structures should assist users in their comprehension of the information space, the medium of interaction is still inherently visual-perceptual. For this reason *H3 predicts that other spatial abilities, such as spatial orientation and possibly visual memory will play an important role.*

It was apparent (Cribbin & Chen, 2001) that the main benefits of using the VIRI emerge over time, when conducting zoom and filter type tasks (see section 4.1.). When using text with no contextual structure, performance decreases in-line with the size of the relevance pool, probably because the common strategy is a sequential scanning one, and discovery is dependent on more on serendipity that directed navigation. In contrast, when given a spatial-semantic structure, although some unguided exploration is necessary en-route to locating the thematic 'hotspot'. Once this is achieved, scanning demands are reduced considerably and navigation performance remains fairly constant. O*n this basis, H4 is that perceptual speed will have a strong effect on text-based search performance but play a less important role in the VIRI search condition, particularly during the later phases of the zoom/filter tasks.*

4. Method

4.1. Design

What follows is a brief overview of the design and methodology. A more complete description can be found in (Cribbin & Chen, 2001).

A formal testbed comprising four information spaces and four interfaces was initially tested. Each information space comprised 200 newspaper articles drawn from the LA Times 1989-90 (TREC v5) database, in each case using a single keyword query. To minimise the effects of reading speed, only documents between the lengths of 250 and 750 words were retained. Keywords used were 'alcohol', 'endangered', 'storm' and 'gambling'.

For the textual interface condition (TEXT), article titles were presented in typical search engine format, as a single HTML page in rank order (determined by keyword hits). To create the visualisation, the data set was subjected to Latent Semantic Indexing (Deerwester, Dumais, Landauer, Furnas, & Beck, 1988). The resulting similarity matrix was then scaled to two dimensions using the minimum-spanning tree (MST) algorithm. Visualisations were realised as VRML models, with each article being represented by a spherical node (see Figure 1.). Titles were viewable by holding the mouse over a node. For all conditions, including TEXT, articles could be selected for reading by clicking on the appropriate sphere or link. The full text could then be read from the adjacent frame (see figure 1).

A four-task design was used to test each interface. This was based on Shneiderman's visual information seeking taxonomy (Shneiderman, 1996).

Tasks A to C formed what will be referred to as structured tasks. Each consecutive task required the user to zoom progressively more deeply into a particular topic within the information space. Task A was a exploratory browse task requiring retrieval of articles relating to a fairly general theme e.g. “Find and mark all articles relating to drink driving”. These articles numbered approximately 25 documents out of the total 200. The criteria for task B required users to find and mark a sub-set of the previous relevance pool (around 6-10 documents). Finally, task C required the location of a specific ‘thread’ or two directly related articles. Again these were articles that were also relevant to the previous task.

Task D was a single unstructured task that required the user to go straight to a two article thread on a new topic i.e. the articles were not relevant to any of the previous tasks.

4.2. Users

Twenty-one people attended the psychometric testing sessions. All were volunteers from the Brunel University student population. Thirteen people (7 males, 6 females) then returned, and completed both the TEXT and MST VIRI interface conditions. Most (10) ranged between the ages of 20 and 29 and the remainder were between 30 and 34 years of age.

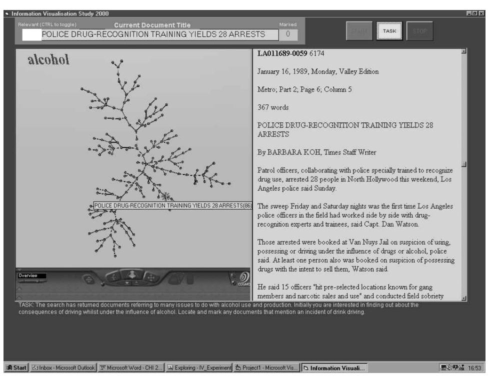

Figure 1: Visual Information Retrieval interface incorporating an MST graph showing the ‘Alcohol’ space.

4.3. Procedure

Interface order and interface-information space combination was counter-balanced across the sample. Users were given five minutes to complete each task and were instructed to mark articles they deemed relevant by pressing the CTRL key. Visited and marked nodes did not overtly change in appearance, although users were able to tell if they had already marked a node from an asterisk that would appear on the interface when they held the mouse cursor over the node. At the beginning of each new task, all previously marked documents were unmarked (i.e. users had to remark any articles that were still relevant). Tasks were always presented in alphabetical order.

Regression analysis was used to identify the most influential factors at each stage of the task treatments. Five objective performance measures were examined. Two of the measures combine precision and recall scores in to the well known F-measure (Figure 2a). The first, Fmk combines recall and precision scores based on number of documents marked and serves as an overall measure of *retrieval success*. The second, Fno, is a measure of *navigational accuracy*, being a combination of recall and precision based on nodes simply visited (titles viewed). LOST, is a navigational efficiency measure (figure 2b) and is known to correlate well with subjective perceptions of lostness (Cribbin & Chen, 2001). Finally, the time taken to locate the first relevant node (Tno) will act as a measure of orientation speed.

$$F_{pr} = \frac{\mathbf{2}pr}{p + r} \qquad LOST = \frac{\mathbf{2}(D/T)\ (R/E)}{D/T + R/E}$$

p = precision
r = recall

D = different irrelevant nodes visited
T = total irrelevant nodes visited
R = number of relevant events logged
E = total events logged.

Figure 2a, b: Formulae for F_{pr} and LOST

Test scores were subjected to principal components analysis, with varimax rotation being used to achieve the final solution. Split-half reliabilities, corrected using the Spearman-Brown formula were computed to check for good internal reliability.

Backward regression was then used to examine the relative predictive value of cognitive factors in terms of search performance by task and by interface. All seven factors were included at the start of each procedure, and factors iteratively rejected until only those that made a significant difference to the predictive power of the model remained. In this case, the factor rejection level was set at alpha >.10.

5. Results

Split half correlations, showed reasonable to good reliabilities for most measures ($.68 < r_{sb} < .92$). Results of the factor analysis show that the cognitive factors split into three components that generally matched the described domains. Visual memory loaded most heavily on the same component as the visual-perceptual factors, rather than associative memory.

5.1. Cognitive predictors of TEXT interface performance

The results of the regression analyses can be seen in Table 1. The first feature that emerged here is that cognitive ability did not tend to predict search performance during the structured tasks (i.e. A, B and C). Only orientation time (Tno) and retrieval performance (Fmk) were reliably predicted by cognitive ability.

Task	D.V.	*TEXT (n=13)* -	+	r^2	*VIRI (n=13)* -	+	r^2
A	*Fmk*						
	Fno						
	LOST						
	Tno	**FI**		.29*	**MV**	**S**	.59**
B	*Fmk*	**MV**	**P**	.42*	**P**	**S**	.35*
	Fno				P	**S**	.41*
	LOST				**P,FA**	**S,MA**	.72**
	Tno		**MA**	.25*	FI	**S,MA**	.55*
C	*Fmk*		**P**	.26*			
	Fno				**P**	**S**	.63**
	LOST				**P**	**S**	.45*
	Tno		**VZ**	.29*	**P**	**S**	.63**
D	*Fmk*		**FI**	.26*	**MV**	**VZ,P, FI**	.72**
	Fno		**FI, VZ, FA**	.74**	**MV**	VZ,P	.48*
	LOST	**MV**	**P,S, MA**	.80**	**MV**	**VZ, P**, MA	.78**
	Tno		**FI,S**	.81**	**MV**	**VZ**	.59**

*Bold characters indicate a significant beta coefficient (p<.05) for a particular factor. R^2 significance refers to the overall model (*p<.05; **p<.01).*

Table 1: Cognitive predictors of retrieval and navigational performance.

Perceptual speed (P) seems to be the important predictor of retrieval performance, but only for the later tasks, B and C. In contrast, Tno was predicted reliably on all three tasks, but each time this was by a different factor. Ideational fluency (FI), rather counter-intuitively was a negative predictor of Tno on task A, predicting 29% of the variance. High associative memory was associated with faster orientation times on task B, predicting 25% of the total variance. Finally on task C, visualisation ability (VZ) was the only important factor, positively predicting 29% of all variance.

The unstructured task (D) sees a more complex and important role for cognitive ability emerging. particularly in the case of the navigation measures where between 74% and 81% of all variance could be accounted for by cognitive ability. This task is perhaps the most difficult task as the user experiences a track change and must locate two unfamiliar articles. The picture is dominated most clearly by ideational fluency. This factor was the strongest predictor of Fmk, Fno and Tno.

VZ again played only a minor role, predicting only Fno significantly, but only as part of a more complex model that included ideational and associational fluency.

Spatial orientation (S) was also predictive of navigational efficiency (LOST) and orientation time (Tno).

5.2. Cognitive ability as a predictor of VIRI performance

As expected, a completely different picture emerges when users were searching with the VIRI. Firstly, cognitive ability plays a much stronger role during the structured tasks, with cognitive models being predictive of Tno from task A. From task B onwards, cognitive ability accounted for significant amounts of variance on all tasks and all performance measures (with the exception of Fmk on task C). As with TEXT, the unstructured task seemed to be the most cognitively demanding with cognitive ability models predicting half and three-quarters of all observed variance. *These findings lend good support for H1.*

5.2.1. Hypothesis 2: The role of visualisation ability

Given the known relationship of visualisation ability to information seeking performance, H2 predicted that VZ would be a less important factor when searching using spatial-semantic spaces. Indeed VZ does not feature at all until task D. This is particularly encouraging with respect to task C, suggesting that given time, low visualisers were comprehending the space equally well as highs. In contrast, the importance of VZ is clear to see on task D. In fact, in terms of reorientation time, VZ was the strongest single predictive factor (r=-.62, p=.025). Hence high VZ ability advantaged people in their reorientation following the change of topic track. That said, although VZ was not tremendously important on TEXT, it certainly featured more consistently during the structured tasks. *Hence H2 is* *partially supported, but VZ still plays a part in predicting the speed with which users are able to switch topical perspective within an information space.*

5.2.2. Hypothesis 3: The role of other spatial factors

H3 predicted that given the visual nature of VIRIs, spatial ability would still have an important role to play in terms of information seeking performance.

For the structured tasks (A to C), spatial orientation seems to dominate the picture, with high scorers tending to perform better. Furthermore, the results show that visual memory (MV) may also be important, although this relationship was, rather counter-intuitively, consistently negative.

As users move through the structured tasks (B & C) MV no longer plays a significant role. This is surprising given that, on task B at least, there were still moderate to strong correlations between visual memory and the navigation measures ($.44 < |r| < .63$). MV, however, is often considered to be a minor spatial factor that relates more to the efficiency of the higher order ability factors such as orientation and visualisation. As the visual-perceptual factors do share a lot of variance, it was conjectured that the predictive power of visual memory, in this context, is principally a result of variance shared with one of the stronger factors. Partial correlations show that spatial orientation alone accounted for most of the variance, with coefficients for all three navigation metrics falling well below significance (see Table 2).

Model	Fno	LOST	Tno
MV	.44	.48*	-.63*
S	.53*	.59*	-.58*
MV– S	.12	.11	-.38

p (1-tailed) <.05, df = 12

Table 2: Full and partial correlations of MV and S with three navigational performance measures

It can therefore be concluded that spatial ability is an important factor in visual information search, with spatial orientation ability being the critical factor. *H3 is therefore accepted.*

5.2.3. Hypothesis 4: Scanning demands

H4 predicted that use of a VIRI would result in lower scanning demands, due to the zoom-focus strategy afforded by the spatial-semantic structure.

Table 1 shows that, indeed, perceptual speed (P) did positively predict performance on TEXT, particularly with respect to retrieval performance on tasks B and C. Counter to H4, however, P seemed to play a more important role on these tasks, when using the VIRI. Furthermore, it is negatively predictive of performance during these tasks. It is unclear why this should be although there is some evidence suggest a less directed strategy from fast scanners. A closer look at the Task C log data reveals P was positively correlated with total nodes visited ($r=.32$). High Ps were also more likely to backtrack over non-relevant nodes ($r=-.35$) and spent less time overall looking at relevant nodes ($r=-.46$). None of these correlations were significant, however.

On task D, the predictive power of P becomes positive probably due to the panning/scanning demands required by the sudden change of topic track.

Overall, the impact of perceptual speed seems somewhat more critical, albeit different in nature, when searching using a VIRI. For this reason H4 is rejected.

5.2.4. Other effects of topical track change

As with TEXT, cognitive ability becomes highly predictive of performance on all measures on task D and so warrants special attention. P is positively predictive on all measures except Tno, but is not the major feature here. Nor is spatial orientation, which hardly features at all. It is MV and, as already described VZ that figure most prominently in the models.

In contrast to VZ, MV is negatively predictive of performance here. This is strange given that direct correlations between this factor and the performance measures were generally weak ($-.05 \leq r \leq -.29$) and non-significant. Partial correlations, however, suggest a suppressed correlation with coefficients rising to between $r=-.48$ ($p=.11$, $df=10$) and $r=-.73$ ($p<.01$, $df=10$) when the effects of P are partialled out. The sample was split into high and low groups on both these measures and the means examined. This revealed an interactive effect which was supported by ANOVA analysis for both Tno, $F(1,9)=7.1$, $p<.05$ and LOST, $F(1,9)=5.9$, $p<.05$. It seems that high MV was counter-beneficial for those of high perceptual speed and made little difference to low Ps. Likewise, only when MV was low did perceptual speed positively predict performance. Overall, a combination of low MV and high P seemed to result in the best performance. Whilst the group sizes are small, the interactive effect does fit well with the correlation and regression data.

6. Discussion

6.1. Cognitive issues in visual information retrieval

Information seeking on the VIRI clearly requires a different set of strategies and as such the cognitive models that predict performance are very different. Spatial orientation dominates all the way through the structured tasks from A to C. Interestingly this factor was far more important than visual memory in this respect and any predictive variance possessed by visual memory was accounted for, in the large part, by this factor. Given the nature of these tasks, it seems that the critical ability here was to maintain and, in the later tasks, reorientate to a previously held perspective. On task D, however, the track change meant a panning and scanning strategy became more important and hence spatial orientation was superseded by other factors such as perceptual speed, visualisation ability and visual memory.

The observed negative predictive effect of perceptual speed during task B was unexpected, but was consistent across all measures except Tno. This feature continues over to task C and for the navigational performance measures at least, became more powerful. This was most likely to be due to differences in strategy. Weak associations between perceptual speed and some of the more low-level log metrics suggest that fast scanners visited more nodes, spent a lower proportion of their time visiting relevant nodes and tended to back-track over non-relevant nodes more frequently. All this points to a more random navigation strategy suggesting they were paying less attention to the structure. The exception of Tno on task B, however, shows that this factor did not seem to affect reorientation time. Based on this, it may be that fast scanners were also more prone to search over a wider area after initially locating the 'hotspot' relatively efficiently.

Task D is an extremely difficult task in comparison to the structured tasks. Users were required to reorientate and zoom in to retrieve just two documents in a previously non-relevant area of the visualisation. Visual memory, visualisation and perceptual speed seem to take over from spatial orientation. Perceptual speed changes from being a negative to a positive predictor. This factor has an interesting interactive effect with visual memory. Visual memory predicted performance in a negative direction, but only when perceptual speed was high. Likewise, the positive predictive value of perceptual speed only held if visual memory was low. Overall, high perceptuals with poor visual memory were the best performers suggesting a simple, fast scanning strategy was optimum here. The key need here, therefore, seems to be to minimise scanning demands during the overview to focus process.

Visualisation ability was also strongly predictive. High visualisers were particularly good at gaining the new focus following the track change and were more navigationally efficient in doing so. This could well relate to the richness of mental models constructed during earlier tasks. Possibly, high visualisers did a better job of comparing and integrating other, non-relevant themes as they explored during the structured tasks.

The final factor of examined was ideational fluency (FI). FI was an important predictor of TEXT search performance on task D. Being able to spontaneously generate related ideas on a new topic is clearly an asset when performing this type of task. Interestingly, this factor was far less influential in predicting VIRI performance. This was particularly true for the navigation metrics. This may be misleading, however, as when searching using TEXT, users did not have 'visit' a node in order to read the title information. They were therefore able to assess document relevance prior to interacting directly with a node. In fact, when one considers the recall and precision levels for marking, on the other hand, this factor does appear in the model, although it's importance is still relatively low in comparison to TEXT.

6.2. Improving usability through navigational cues

The interfaces used in this study were intentionally poor of navigational cues. For this reason, the cognitive issues raised here are fundamental ones relating to visual database navigation and provide clear indicators to the nature of the support that particular user groups may require.

Further visual cues are clearly necessary to facilitate comprehension and support faster, more efficient navigation. Our data allow several informed suggestions in this respect:

1. During the continuous zoom/filter (structured) tasks maintaining orientation within the space seemed to be a key factor. Navigational history should be represented explicitly, preferably in a way that implies recency.
2. The provision of high-level context descriptors has shown to be highly beneficial when navigating unfamiliar (Chen & Dumais, 2000). Such descriptors placed within the appropriate

spatial contexts could enrich VIRIs and support users who may experience trouble in forming/applying a mental model to a large information structure. Thematic ‘landmarks’ could be defined and located either automatically, using text co-occurrence analysis, or manually (e.g. user-annotations). Computer generated landmarks would be particularly beneficial during the early stages of navigation.

3. Perceptual speed seemed counter-productive to good performance on some tasks. Our data suggests this may have due to fast scanners adopting a more random navigation strategy (i.e. less reliant on the node structure). In this respect it may be useful to explicitly ‘zone’ the visualisation by representing the spatial extents of all or certain salient contexts, in order to prevent users accidentally straying too far from areas of interest.

7. Conclusions

This study found that the visualisation approach to supporting information search may, without careful, informed design, raise more human factors issues than it solves. There seems to be a complex interaction between cognitive ability and task type. One must bear in mind, however, that the study sample was relatively small and users’ exposure to the visual interface was brief. Future studies must address these and other related issues.

Acknowledgements

This study forms part of a research project funded by the Library and Information Commission.

References

Allen, B. (1992). *Cognitive differences in end user searching of a CD-ROM index.* Paper presented at the Fifteenth Annual International ACM SIGIR conference on Research and development in information retrieval, Copenhagen Denmark.

Carroll, J. (1993). *Human cognitive abilities: A survey of factor-analytic studies*. Cambridge, UK: Cambridge University Press.

Chen, H., & Dumais, S. (2000). *Bringing order to the Web: automatically categorizing search results.* Paper presented at the CHI 2000 conference on Human factors in computing systems, The Hague Netherlands.

Cribbin, T., & Chen, C. (2001). *Visual-Spatial Exploration of Thematic Spaces: A Comparative Study of Three Visualisation Models.* Paper presented at Electronic Imaging 2001: Visual Data Exploration and Analysis VIII, San Jose, CA.

Deerwester, S., Dumais, S., Landauer, T., Furnas, G., & Beck, L. (1988). *Improving Information Retrieval with Latent Semantic Indexing.* Paper presented at the 51st Annual Meeting of the American Society for Information Science, ASIS '88, Atlanta GA.

Egan, D. (1988). Individual Differences in Human-Computer Interaction. In M. Helander (Ed.), *Handbook of Human-Computer Interaction*. North-Holland: Elsevier.

Ekstrom, R., French, J., Harman, H., & Derman, D. (1976). *Manual for kit of factor-referenced cognitive tests*. Princeton, NJ: Educational Testing Service.

Hook, K., Sjölinder, M., & Dahlback, N. (1996). *Individual Differences and Navigation in Hypermedia.* Paper presented at the Eigth European Conference on Cognitive Ergonomics (ECCE-8), Spain.

Shneiderman, B. (1996). *The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations.* Paper presented at the IEEE Conference on Visual Languages, Boulder, CO.

Stanney, K., & Salvendy, G. (1995). Information Visualization: Assisting low-spatial individuals with information access tasks through the use of visual mediators. *Ergonomics, 38*(6), 1184-1198.

Swan, R., & Allan, J. (1998). *Aspect Windows, 3-D Visualizations and indirect comparisons of Information Retrieval system.* Paper presented at the 21st annual international ACM SIGIR conference on Research and development in information retrieval, Melbourne, Australia.

Westerman, S. J., & Cribbin, T. (2000a). Cognitive ability and information retrieval: When less is more. *Virtual Reality, 5*(1), 1-7.

Westerman, S. J., & Cribbin, T. (2000b). Mapping semantic information in virtual space: Dimensions, variance, and individual differences. *International Journal of Human-Computer Studies, 53*(5), 765-788.