

A STUDY OF NAVIGATION STRATEGIES IN SPATIAL-SEMANTIC VISUALISATIONS

Timothy Cribbin and Chaomei Chen

The Vivid Research Centre, Department of Information Systems and Computing,
Brunel University, Uxbridge, United Kingdom. UB8 3PH

ABSTRACT

Visualisations of abstract data are believed to assist the searcher by providing an overview of the semantic structure of a document collection whereby semantically similar items tend to cluster in space. Cribbin and Chen (2001) found that similarity data represented using minimum spanning tree (MST) graphs provided greater levels of support to users when conducting a range of information seeking tasks, in comparison to simple scatter graphs. MST graphs emphasise the most salient relationships between nodes by means of connecting links. This paper is based on the premise that it is the provision of these links that facilitated search performance. Using a combination of visual observations and existing theory, hypotheses predicting navigational strategies afforded by the MST link structure are presented and tested. The utility, in terms of navigational efficiency and retrieval success, of these and other observed strategies is then examined.

1. INTRODUCTION

The representation of unstructured document collections by means of node based graphing methods is an approach that has been widely adopted by the information visualisation community. These visualisation methods aim to support information seeking by providing the user with a spatial map representing an existing model of the semantic structure of a collection. Documents are generally represented as nodes plotted in visual space whereby inter-document proximity implies their degree of similarity.

Visual overviews should guide the information seeker to important information whilst minimising the amount of reading and scanning required. The potential value of overviews, therefore, tends to increase in line with document collection size. Unfortunately as thematic diversity also tends to increase with collection size an overall drop in the accuracy of inter-document proximities can often result. When the reliability of spatial proximity cues are low, the user will be more frequently misled, reducing their navigational efficiency and causing frustration. Furthermore, large maps can appear quite amorphous and cluttered with few obvious landmarks to guide orientation.

1.1. Using visual links to communicate salient structure

One approach to alleviating this problem is to use other cues to further emphasise the most important or strongest inter-document relationships. Previous empirical work (Hascoet, 1998) has found that users preferred to search using visualisations that made clusters more identifiable and that provided good clues as to which direction to move next once good items had been found. By connecting highly similar documents with visual links, minimum-spanning tree (MST) graphs hold the potential to satisfy both these criteria. MST graphs are always composed of N-1 links, the minimum required to form a continuous tree structure. The example in Figure 1 shows how the branching structure of the MST graph provides additional landmarks and is less cluttered in overall appearance.

1.2. Information seeking performance and visual links

(Cribbin & Chen, 2001) reported results from a study that compared information-seeking performance using a range of graphing algorithms. Principle components analysis (PCA) was used to create a traditional scatter-graph solution from similarity data generated by Latent Semantic Indexing (Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990). MST visualisations were also created from the same data. Participants completed information-seeking tasks,

according to a number of given criteria, using each of these visualisations for navigation. Use of MSTs resulted in general and significant performance improvements in comparison to PCA visualisations. This paper examines navigational strategies employed by participants when using the MST visualisation, focusing in particular on the early stages of exploration. It is proposed that the link structures provide a number of navigational cues that, when exploited, lead to superior information-seeking performance.




Figure1: PCA and MST graphs representing the same semantic proximity data

1.3. A classification of node types in tree networks

Using MST graphs to represent a document collection provides the user with important cues that are known to support navigation in the real world, most notably pathways, edges and nodes² (Lynch, 1960). The continuous link structure provides pathways that support route following and learning. A document radiating a single link clearly indicates the limits of a cluster or pathway. In contrast, a document possessing many links suggests an important document that may form a useful focal point.

In this paper nodes are classified into three types according to the number of direct links to other nodes. An *extremity* possesses a single link, a *thread* node possesses two links and *branch* nodes link to three or more other nodes. From a combination of existing theory and observational evidence acquired from visual replays of task event data a range of hypotheses are put forward.

1.4. Hypotheses and rationale

The analysis is split into two sections. Firstly, an investigation into navigational strategies will test three hypotheses relating to landmark and pathway usage. Branch and extremity nodes are both presented as natural landmark candidates. Secondly, performance benefits in terms of navigation and retrieval performances will be explored in within the context of these strategies.

1.4.1. Using landmarks to guide early orientation

On first presentation, the spatial layout was completely novel to participants. Furthermore, the semantic properties of visible regions were also unknown. When navigating the real world, people tend to identify the most salient landmarks and use them as anchor points from which they can form a working cognitive model of the space (Couclelis, Golledge, Gale, & Tobler, 1987). Our observations also show that landmarks, extremities in particular, seem to be used regularly during the early orientation stage. *H1 therefore predicts that visits to extremity and branch nodes will be comparatively more frequent during the early stages of the task.*

1.4.2. Using pathways to maximise return

The inclusion of a visual link connecting two nodes implies a strong semantic relationship between the two associated documents. Given that participants were aware of this feature it can be expected that, on discovery of a relevant document, users will attempt to capitalise on the pathways radiating from the current node, an activity

² Lynch use of the term node differs from that of this paper, referring instead to a focal point within a city such as a town square.

referred to hereon as chaining. *H2 therefore predicts that discovery of a relevant document will result in an increase in chaining activity.*

1.4.3. Using landmarks to escape from sparse regions

In contrast, users will only spend time browsing information in a particular area if they believe that further useful discoveries are likely and imminent (Chalmers, 2000). Discovering a non-relevant document within a spatial-semantic environment would suggest to the information-seeker that proximal documents are also unlikely to be useful. In such circumstances, it is predicted that users will be more inclined to jump to a different region of the space. Furthermore, they will seek to reorientate themselves by jumping to a landmark node. *H3 predicts that visits to branch and extremity nodes will be comparatively more frequent when the current document node is non-relevant to the task criteria.*

1.4.4. Utility of identified strategies

Correlational analysis will then examine the relationship between the extent to which these strategies were employed and overall task performance. Criteria for success were search recall, precision, general efficiency, lostness and time taken to retrieve the first relevant document. Efficiency is calculated by combining recall and precision using the F-measure formula (see van Rijsbergen, 1979). Lostness is also created with the same formula using a combination of a backtracking measure (time spent re-visiting non-relevant nodes) and the proportion of time spent examining relevant nodes.

2. METHOD

The complete methodology used in this study can be found in a previous paper (Cribbin & Chen, 2001). In brief, four document collections each comprising 200 documents were subjected to Latent Semantic Indexing (LSI) and similarity data computed. Documents were retrieved from the TREC Los Angeles Times database using single keywords. These were alcohol, endanger, gaming, storm. In each case, a set of two-dimensional coordinates and inter-document links were computed using an MST algorithm. The computer display was split into two panes. In the left pane an interactive VRML world, representing the MST model. The right pane displayed full document text when users clicked on a node.

The task examined here required participants to search the visualisation for documents that conformed to a broad set of criteria. For example, for the 'alcohol' collection users were asked to find any document that mentioned an incident of drinking and driving. For each collection approximately 12% (22-26 documents) of all documents were classified as relevant.

Sixteen participants completed the task. All were postgraduate students at Brunel University. Each collection was administered an equal number of times across the sample. Participants received no prior experience of either the spatial structure or the documents contained within. They were instructed to locate and mark all documents relevant to the given criteria. Five minutes was allowed to complete the task although participants could terminate the task earlier if they wished.

3. RESULTS

3.1. Observed uptake of predicted strategies

Early activity was defined as the first five nodes visited by the participant. The proportion of all events (visits, clicks and markings) for both the early and whole task period, was computed by node type. These values were then normalised against chance levels by subtracting the actual proportions of each node type existing within the space.

Figure 2 shows the comparison of early events against events recorded across the whole task period. ANOVA showed a significant quadratic interaction between node type and period, $F(1,15) = 5.01$, $p < .05$. There was a clear bias during early activity towards branch nodes. This was not significantly different from chance, $t(15) = 1.22$ or from the whole task period, $t(15) = 1.03$. There was also a tendency for participants to avoid thread nodes and this was significant with respect to chance, $t(15) = 1.94$, $p(1) < .05$, and whole period events, $t(15) = 2.24$, $p(1) < .025$. Extremity visits occurred at approximately chance level for both the early, $t(15) = .05$, and whole period, $t = 1.24$ and

there was no significant difference between early and whole period event likelihood, $t(15)=.33$. H1 is therefore partially supported, although the data suggests that extremities were not viewed as useful landmarks during the early stages of navigation.

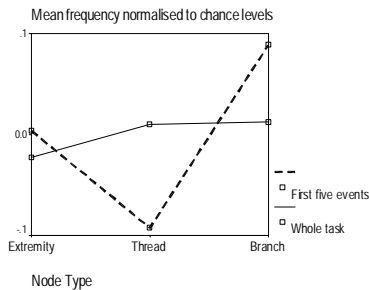


Figure 2: Likelihood of transition to node type as a function of task period

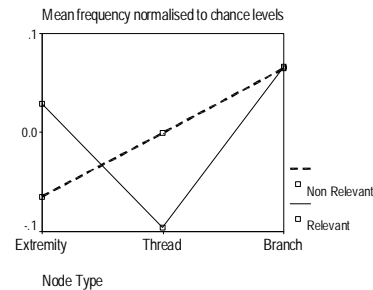


Figure 3: Likelihood of transition to node type as a function of current node relevance

With regards to chaining activity, the likelihood of following a link to an adjacent node was not found to be significantly higher when the source node was relevant, $t(15)=.441$. H2 is therefore rejected.

To test H3, ANOVA was computed to determine the relative likelihoods of jump destinations from non-relevant documents. Figure 3 shows a linear pattern of means suggesting a positive relationship between the number of radiated links and the attractiveness of a document node. Branch visits were clearly most likely in this context, with the likelihood of a visit being significantly greater than chance, $t(15)=1.87$, $p(1)<.05$. Thread visits occurred roughly at chance level and extremity visits occurred at a below chance level, suggesting they were intentionally avoided, although this was not significant, $t(15)=1.58$.

In comparison, transitions from relevant nodes show the kind of pattern originally expected from non-relevant nodes with means forming a U-shape graph (see Figure 3). The observed difference in distribution was confirmed from a general interaction, $F(1,15)=3.33$, $p<.05$. Thread visits fell significantly below the chance level, $t(15)=2.64$, $p(2)<.02$. There was an almost identical tendency towards branch nodes visits and a greater than chance tendency towards extremities, although this was not significant, $t(15) = .50$. At a relative level, extremity visits were more common when the source node was a relevant document. H3 is therefore rejected.

3.2. Utility of observed strategies

The previous analysis showed an overall preference for using branches as landmarks during the early stages of search. Use of this strategy, however, only weakly predicted the time taken to mark the first relevant document ($r=-.25$, ns). Virtually no relationship was observed between extremity visits and time to first marking was found ($r=.09$).

Given that a link between two documents implies high similarity, H2 predicted that chaining would be more common from relevant documents. This was not found to be the case. Furthermore, correlations show that it is the degree to which participants chained from non-relevant documents that was most predictive of overall performance. Significant correlations were found for all performance criteria with the exception of time to mark first relevant document. Significant coefficients were in the range $r=.52$ to $.61$ ($p<.05$). Although weakly positive, none of coefficients computed with respect to chaining from relevant documents approached significance.

H3 predicted that jumping to landmarks would be seen as an efficient means of escaping a sparse region of the information space. The strategy of jumping to extremities from non-relevant documents was negatively predictive of performance across all criteria. Participants who used this strategy tended to be less precise with respect to the relevance of documents they marked ($r=-.58$, $p=.02$) experienced more lostness ($r=-.48$, $p=.06$) and read (clicked on) fewer relevant documents ($r=-.45$, $p=.08$).

Despite being a common strategy, the extent to which participants jumped from non-relevant documents to branch nodes did not predict performance in any respect. In contrast, the frequency of transitions to branches from relevant documents was indicative of efficient navigational, but not retrieval performance. For instance the F-measure based

just on nodes visited correlated significantly with usage of this strategy ($r=.50$, $p=.05$). In contrast, jumping to extremities from relevant documents produced the opposite effect on navigational efficiency ($r=-.47$, $p=.06$).

As a final note of interest, the correlational data shows that participants could be split into two more, unexpected styles. It was found in all cases that frequency of transitions to extremities tended to correlate negatively with those to branches ($r=-.79$, $r=-.70$, $p<.01$). This suggests that some users preferred to either pick away at the ends of branches whilst others navigated more extensively around the branch structures. Evidence in support of this assertion comes from negative correlations between thread visits and extremity visits ($r=-.55$, $r=-.51$, $p<.05$). No significant correlations were found between thread and branch visits. Given our data, it seems like the latter strategy was the most successful one.

4. DISCUSSION AND CONCLUSIONS

The analysis of navigation patterns conducted here has provided some interesting insights into the strategies adopted by information seekers when navigating document collections using MST graph visualisations.

When first confronted by a novel document space, users seem to rely on branch, rather than extremity, nodes as landmarks to guide orientation. This strategy, however, did not reliably predict the time taken to identify the first relevant document. When the session data is considered as a whole, however, it seems reliance on branch nodes does become a slightly more successful strategy although this was only in terms of navigational rather than retrieval effectiveness. High reliance on extremity nodes, on the other hand, was generally counter-beneficial to good performance. This was unexpected but may be because searchers did not use extremities in the expected sense. That is, they did not follow up the visit by travelling down the available pathway towards the source branch.

The prediction that chaining activity would be used on discovery of a relevant document as an efficient means of locating others did not seem to be the case. This could be an artifact resulting from the broadness of the relevance criteria given in this task. Clustering of relevant nodes may not have been as coherent as expected. This could have lead to users losing faith in the reliability of the link structure for these purposes. Chaining was, in fact, a more successful strategy when used to navigate from non-relevant documents. Perhaps this is indicative of the utility of taking a generally more methodical approach, particularly when lost. Further analysis is necessary to clarify this issue.

MSTs clearly provide valuable support to the information navigation process. This paper, although only a snapshot, shows that the way these features are exploited is variable and that chosen strategy can have a significant impact on overall performance. Future research and analysis should seek to understand how these features can be best exploited and what further perceptual cues are necessary in order to maximise value of these structures.

5. REFERENCES

- Chalmers, R. (2000, 11 November). Surf like a bushman. *New Scientist*, 39-41.
- Couclelis, H., Golledge, R., Gale, N., & Tobler, W. (1987). Exploring the anchor-point hypothesis of spatial cognition. *Journal of Environmental Psychology*, 7(2), 99-122.
- Cribbin, T., & Chen, C. (2001, January 21-26). *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., Furnas, G., Landauer, T., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6), 391-407.
- Hascoet, M. (1998, 18 - 23 April). *Analytical versus empirical evaluation of spatial displays*. Paper presented at CHI 98: Human Factors in Computing Systems, Los Angeles, CA.
- Lynch, K. (1960). *The Image of the City*. Cambridge, MA: MIT Press.
- van Rijsbergen, C. (1979). *Information Retrieval*. London: Butterworths.