Adding Imageability Features to Information Displays

Matthew Chalmers, Robert Ingram† & Christoph Pfranger
Ubilab
Union Bank of Switzerland, Bahnhofstrasse 45, 8021 Zürich, Switzerland
chalmers@ubilab.ubs.ch

ABSTRACT
Techniques for improving the imageability of an existing data visualisation are described. The aim is to make the visualisation more easily explored, navigated and remembered. Starting from a relatively sparse landscape–like representation of a set of objects, we selectively add to the visualisation static features such as clusters, and dynamic features such as view–specific sampling of object detail. Information on past usage is used in this process, making manifest an aspect of interaction which is often neglected. Issues arising from the use of such features in a shared virtual environment are discussed.

KEYWORDS: Visualization, navigation, imageability, information design.

INTRODUCTION

Graphical displays of information attempt to make complex information more accessible and legible than textual or tabular techniques. Although profuse in today's computers, information design is by no means a trivial matter [22]. Clutter, not clarity, is too often the result. This is especially the case with regard to displays where not only are there many objects to be shown, but also each object is relatively complex.

The focus of this paper is a set of techniques intended to alleviate this problem. This work has been applied to visualisations in a 3D space of a landscape–like model of bibliographic data. This data was obtained from the HCI Bibliography Project [20]. (Other data types have also been worked on, but for the sake of clarity we will confine most of this paper's discussion to bibliographic data.) Individual documents are shown as small cubes positioned in 3D space. A mesh of triangles interconnects them, thus making a surface of a type often used in visualisation. This surface is usually flat, although some surface roughness may remain due to layout constraints. Previously the Bead system used simulated annealing for the layout task, as in [4], but more recently we have been using a more efficient technique which is the subject of a complementary paper [6].

† R. Ingram's current address: Dept. of Computer Science, University of Nottingham, NG7 2RD, U.K.

With the basic landscape model of [5], an object selected with the mouse changes colour and its detail is 'turned on' e.g. a title appears above it. This 3D text must be large enough to allow for legibility and to allow clicking on words in order to spark off searches. Alternatively, one can initiate a search by typing a word directly into the window, or using an accompanying Java widget for boolean construction of more complex queries. Searches colour matching documents, so that one can see the distribution of matches within the overall structure of the layout.

This rather sparse representation was found to be insufficiently self–disclosing and engaging. Too much user effort is required to build up knowledge of the landscape, with a tiresome amount of clicking on objects and repetition of searches being necessary. We would like the visualisation to unfold and reveal itself more naturally. Our response has been to add further objects, behaviours and colourings to aid the user in exploring and becoming familiar with the landscape. We wish to make the landscape more *imageable* (in the sense of Lynch's *Image of the City* [16]) by the addition of static and dynamic features to the information design, while avoiding self–defeating clutter.

Static features such as coloured regions or clusters of objects are added to serve as straightforward bases for orientation. Dynamic features reveal detail of the modelled objects while avoiding the occlusion problem that can occur either if all objects' detail is turned on simultaneously or if all those objects within a given distance of the viewer's eyepoint are turned on — as is often the basis of 'level of detail' mechanisms. This problem is especially likely to occur when the detail is larger than the object, as is the case here in showing legible titles. To address this, we sample objects every few seconds, so as to obtain a limited amount of detail to display before the next sampling round. Sampling is not random, but is based on such features as the set of objects within the user's field of view, the closeness of those objects to the viewer, the relative frequency of word occurrence and the history of search activity. In this way, we make displays adapt according to interaction.

The Bead system was written in a mixture of C and Java, and runs on Silicon Graphics workstations. It employs the DIVE virtual environment toolkit in offering a shared information space [7]. Some aspects of the information design relevant to synchronous and asynchronous use by multiple users are briefly touched upon before a description of our future work plans and the paper's conclusion.

BACKGROUND AND RELATED WORK

Designing effective visualisations for complex information is a difficult task, and one which has been addressed by a number of researchers. One of the essential problems is how to avoid cluttering the display, and instead concentrating detail on interesting or important regions of the display.

The fisheye lens was one of the first major techniques to address this problem [11]. Distortion like that of a wide angle lens is used to show detail at a central focus. Detail is increasingly compressed and simplified further away from the focus. As the lens is moved, what is in focus smoothly shifts to being less detailed context, while another area, previously contextual, moves into focus and gradually becomes more detailed.

In effect this process is the same as the 'level of detail' mechanisms used on many commercial and research visualisation tools, whereby greater detail is shown when an object comes within a certain (usually 3D) distance of the viewer. As was mentioned in the previous section, this can lead to legibility problems when the detail to be shown has to be made large enough to be readable.

Other researchers tried to avoid strong distortion and instead use more naturalistic perspective views, for example the UIR group at Xerox PARC which produced the Perspective Wall [17] and Cone Trees [21]. They gave a lead in richer designs where data was laid out according to abstract data structures such as hierarchies and common linear dimensions, and rightly emphasised the importance of human perception in making visualisations effective. Distant detail was considered to be less significant than that of nearby objects, and so perspective and occlusion were used to allocate screen resolution where it was most needed. Nevertheless, occlusion hindered browsing and navigation. Similar problems led to a shift from a 3D point cloud representation to a '2.1D' landscape model in an earlier version of the Bead system. The design issues behind this change are discussed in [5].

One of the fundamental works in visualising more complexly structured data is SemNet [9]. Multidimensional scaling techniques were used to create a 3D layout of Prolog modules based on their inter–calling relationships. Clusters of objects in the resulting layouts could be selectively collapsed down to smaller shapes, thus reducing occlusion effects by means of abstraction. A negative aspect of this is the loss of detail. The patterns and textures of the cluster elements may offer useful information to the observer which is abruptly lost when collapsing occurs.

Another means of handling clutter is to selectively filter out unwanted detail before zooming in to areas of interest. A good example of this is in the work of Ahlberg [1], where it was emphasised that tight coupling between filtering tools and the display helped with user engagement. A problem, however, is that filtering can remove the surrounding context of a query's matches. What might potentially be the next browsing step, and how the current query results relate to earlier queries, are harder to discover when non-matching objects disappear.

An alternative approach to selectively displaying detail was shown in the Pad system [19] and its descendant, Pad++ [2]. Here, the apparent size of an object determines the amount of detail given to it. As one zooms in on an object, further detail smoothly 'blooms' out. Previously illegible characters become legible and new objects appear, while very large objects fade out and disappear. What was background texture, giving information about the character of the region in view, smoothly becomes detail as one zooms in. Eventually it may be so large — as an abstraction it is common to all lower–level objects in the field of view — that it is no longer so important to show it. Instead, screen resolution is allocated to new, finer detail.

In Pad–like systems, however, it can be difficult to show useful or legible information on objects which are lower in the hierarchy of abstraction when there are many such objects or when their representation is complex (as is the case with bibliographic data). Given a view from one point in the information space, one must either create intermediate discretisations which describe subregions well, or one must restrict the detail shown on each low–level object to a handful of pixels.

With complex data such as documents, the latter is not an attractive option. The former brings up the difficult problems of textual analysis and interpretation. Although some attempts have been made to automatically generate labels for regions on a mapped corpus of documents [15], the best labels may be the words people actively use there (e.g. in searches) rather than words chosen automatically before work starts.

One of the Pad++ authors was involved in enriching interaction in visualisations by showing on objects traces of past use [12]. The authors maintained and exploited object–centred interaction histories, showing which subregions of a document have been the focus of earlier work either in creation or in reading. This helps answer questions such as: Which sections of a document have been read by various categories of users? Who were the last people to read a section, and when?

Remote access to bibliographic citation databases was the subject of [18]. The Butterfly's design emphasises access to remote databases where speed of access is variable. Its display of information grows as the user progresses with searches, so that pyramids and piles grow in accordance with the order of retrieval and placement i.e. the overall structure is specific to the history of search activity. Also, regions of the display are associated with different types of access e.g. searches will be done in one area but browsing on a retrieved (or a related) data set must be done in another. Focusing on a butterfly sparks off a background query process to obtain further information about that item.

Citations offer a means to link between these structures, and butterflies are linked by a linear chain of citations. Neither the more general graph of citations nor the contents of each citation are used in the structuring of the display i.e. no overall map is made based on wider semantic content.

Themescape [23] attempted to create an overall map, by first making clusters based on similarities in word occurrence, laying out cluster centroids using either principal component analysis or multidimensional scaling (MDS), then finally scattering or interpolating individual documents near to their associated centroids. Computational complexity is much reduced, but the technique tends to ignore the detail of inter–document relationships by relying on the coarse discretisation offered by the initial clustering. Also, interpolation between clusters can be ambiguous or unreliable e.g. if the centroids formed a square *ABCD*, a document placed at the centre could be there because of being an even mix of *A* and *C*, or of *B* and *D*. Lastly, choosing the number and acceptable sizes of clusters to partition the set can be a difficult, data–dependent task.

The layouts of Bead are constructed using numerical optimisation techniques similar to MDS [6]. The system approximately represents the similarities and dissimilarities between articles by their separation in the landscape i.e. the high–dimensional distances defined by similarities in word occurrence are approximated by low–dimensional spatial distances in the visualisation. At any moment, one can look at the discrepancy between the current low–dimensional distance between two articles and their high–dimensional distance, and so decide whether to add a force to push them further apart or to pull them closer together in order to reduce this error. An additional gravity–like force is used to attract the documents towards a 'ground plane'.

This system of forces incrementally lays out the landscape, with global structure emergent from the action of many attractions and repulsions between individual objects. Indirect attractions and repulsions often occur via chains of objects e.g. an object *A* may be laid out close to another, *C*, primarily because of an intervening *B* to which both *A* and *C* are similar. Consequently, the axes of the layout do not correspond to consistent increases or decreases of a small set of semantically–relevant values. Work on the efficiency of these layout algorithms has been the focus of a substantial part of the work on Bead, with algorithms developed to take the complexity of each iteration of the layout process for *N* objects down from the standard $O(N^2)$ to $O(N.logN)$ and most recently to $O(N)$.

STATIC IMAGEABILITY FEATURES

The layout program positions objects according to their relative similarity, and therefore clusters based on geometric proximity in the layout should indirectly approximate clusters in the original high–dimensional space of the data. As in [13], clusters are built using a minimal spanning tree of the objects. This tree is constructed such that the edge values are taken to be the Euclidean distance between the objects. By identifying and eliminating 'inconsistent' edges i.e. edges of the spanning tree whose values are significantly greater than the nearby edge values, clusters or 'regions' are produced. An example result is shown in Figure 1. We sometimes combined the low–dimensional distance in the layout with high-dimensional distances in defining edge values, but with little subjective improvement. Ingram also linked clusters with 'paths' which were to suggest routes for exploration. Paths are intended ultimately to evolve with use of the data so as to suggest common or interesting routes.

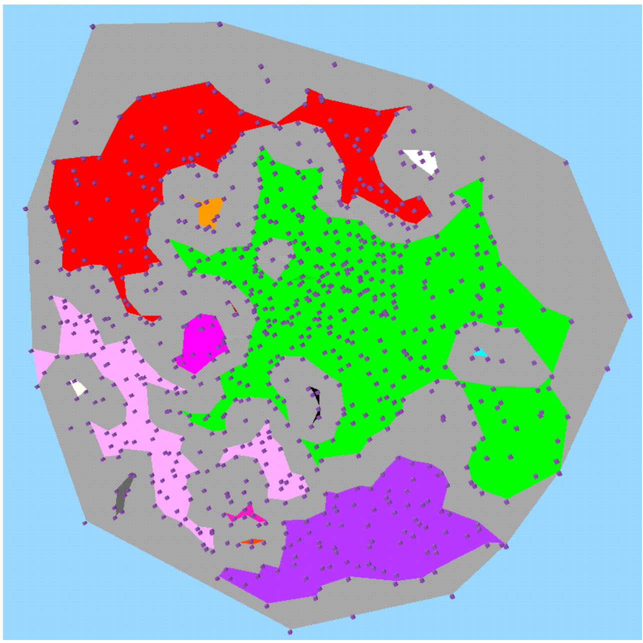

Figure 1. A layout of a set of 831 bibliography entries from CHI, CSCW and UIST conferences. Coloured clusters, based on local density in the layout, have been added. Clusters serve as static imageability features for orientation and navigation. A few large clusters dominate, although several smaller clusters exist.

In Ingram's earlier work, objects had their colour and shape changed to show membership of a cluster. The emergent effect was to show districts of similar objects. We decided not to vary shape because of the difficulty in resolving differences in shape from a distance. We also did not employ colouration of the objects themselves because of the use of object colour in keyword searches. Instead we colour the mesh of triangles linking objects, which effectively give an extra dimension for display. This option was not feasible in Ingram's earlier 3D work because of occlusion problems.

When all three vertices of a triangle are from the same cluster then the triangle is given the unique colour of the cluster. No association is currently made between a cluster's colour and any property such as density or history of activity. Remaining triangles are given a neutral colour, creating bands that reinforce the delineation of clusters. Note that long thin clusters can contain no triangles and therefore their objects sit in the neutrally–coloured bands.

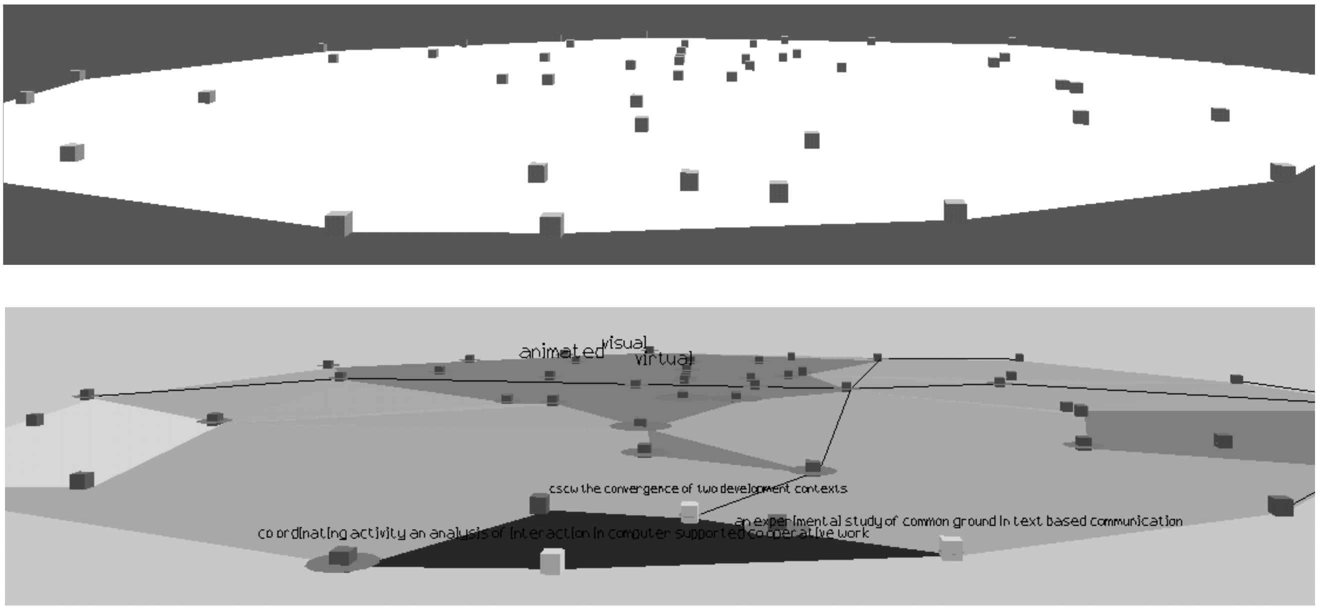

Figure 2. Adding imageability features to a layout of CHI91. The basic layout of documents is relatively sparse (above) but with the addition of imageability features (below) we obtain a richer, more accessible display. Static features such as coloured clusters and paths aid orientation and navigation, as well as emphasing global patterns in the data set. A small number of documents are dynamically but randomly chosen to be highlighted, with a bias based on nearness to the eye. Topic words are also dynamically placed in the scene, based on frequency of occurrence in the current field of view and a 'popularity' based on word usage history amongst all users. Usage discs can be made visible around each document, to more directly show how often each one has been hit by searches.

Our clustering relies on *a priori* layout positions. We consider this approach to be only provisionally acceptable, and hope to improve the 'sense' of clustering by integrating data on the ongoing activities of the users. Collecting and employing usage data are discussed in the next section.

DYNAMIC FEATURES

Dynamic imageability features are here subdivided into two types: view–independent and view–specific. The former concentrate on showing the current state of relative usage frequency of all objects. Their positions and shapes are independent of each user's field of view, and their underlying data also feed into view–specific features, described later.

As a user performs a search, and sees the pattern of hits and misses in the visualised objects, a log entry is made of the time, search key and user involved. After initially colouring objects, which tended to clash with the colouration used in searches, we opted for adding discs under the document objects whose radii conform to the relative search hit frequencies. Their size and position are intended to be most useful when seen from above i.e. in an overview, where one can more easily make relative judgements about larger numbers of objects. This information is dynamically updated within the virtual environment as activity progresses. Selecting a usage disc with the mouse prompts the system to show the history of searches that hit the object in a neighbouring window. The discs therefore show which documents have been of interest to many people. This kind of information may, for example, be of particular interest to novices in the field of the displayed documents [14].

Although we have concentrated on search data, we can also collect information on which objects are selected with the mouse and which objects are near to the user as he or she moves through the space. We wish to record as many user activities as possible: activities which express interest in, or the importance of, embodied entities in the virtual environment such as documents, words, clusters, regions, and, last but not least, other users.

We wish to be able to see detail about nearby parts of the landscape but also some detail about distant objects. Our interest ranges from individual documents to sets of documents in perhaps quite distant regions of the landscape. We wish to find which words dominate in such distant regions i.e. which topics predominate there, and therefore which regions are good candidates for further exploration and detailed view. Furthermore, we wish to do this in a way that is relatively passive for the user i.e. the landscape should be self–disclosing and unfolding, as actively clicking on (and off) every potentially interesting object so as to see its title is excessively laborious.

The DIVE system offers the ability to automatically increase the level of detail on an object when a user comes within a given distance [8]. However, since our objects have complex detail, as one moves near to a number of objects this tends to create a cluttered region of overlapping text immediately in the foreground. The text is in itself difficult

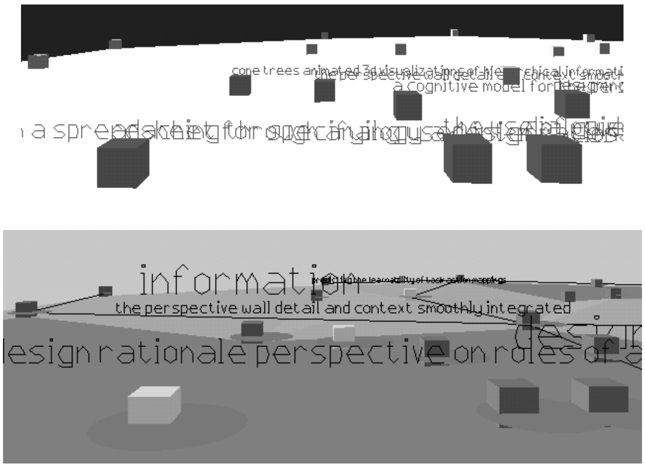

Figure 3. Titles turned on by a distance–based 'level of detail' function overlap and reduce legibility (above, from an older Bead version). The successive sampling of pop–ups (below), based on closeness to the viewer, reduces occlusion to acceptable levels.

to read, and it can also obscure more distant objects and features. An example is given in Figure 3.

Rather than rely on more complex artificial intelligence or information retrieval techniques to discover topics in text, we combine layout positions with information on the work of previously active users, and with the changing field of view of each user. We use dynamic sampling of detail to avoid cluttering the screen with an excessive number of document titles. At any one time, only a small number of objects show their titles: three if the user is at rest and one otherwise. Each second, the oldest activated document has its highlighting removed, and then we turn on a new object sampled from the current field of view. Given time looking from a single viewpoint, the viewed region will gradually disclose itself.

Sampling is based on the solid angle of each object i.e. the apparent size of the object in the perspective view. In this way, the detail of close documents is shown more often than that of distant objects. We found it best to use 3D text for titles so as to reinforce the association of title and object in the perspective view. Titles that are very far away could be so small as to be illegible, and those of extremely close objects are also often difficult to read. Therefore we use a pair of threshold distances — for 'too near' and 'too far' — to constrain the choice of objects and to set text size.

In order to tighten the coupling of word searches and the display, we make any displayed word in the environment clickable. Rather than having to always transfer one's attention to the search widget, a clicked-on word activates a search for the word (Figure 4). All matching documents are changed to have a light colour other than the 'pop-up' colour, and matching words in titles are shown. Accumulated data on searches and hits serves in a measure of the importance or popularity of words, and is used for a second type of dynamic imageability feature, topics.

A topic is intended to describe a region of the landscape currently in view. It is 3D text of a size (before perspective transformations) four times greater than title text. Each candidate word is ranked according to the following expression:

$$Popularity^2 \times \frac{n}{f} \times Bias^n$$

where *Popularity* is the number of times a word has been searched for, *n* is the number of occurrences of the word among the titles of the visible documents, *f* is the frequency of the word among all document titles (and hence *n/f* gives the proportion of occurrences in the field of view), and *Bias,* usually greater than 1, is a factor used to further favour words which occur frequently in the field of view. The top–ranked word is displayed, at a position above the centre of mass of those visible documents within which the word occurs. In this way we try to find topic words which are well-fitted or consonant with the user's model of the information. At present, popularity ratings are shared by all users i.e. the rating is given by the community and not the individual.

Topic generation can be computationally expensive and so we employ a heuristic to speed up their generation. We maintain a table of which objects have titles containing one or more of the top 1000 popularity–ranked keywords. Although topic generation uses only this subset of documents, it still makes the visualiser pause for a fraction of a second before smooth motion is recontinued.

Topics can be seen from quite far away and, being relatively unobtrusive, we show up to three simultaneously. At present we compute a new topic every four seconds and, before inserting it into the scene, remove the oldest one. This also helps in the common situation where there are several words which vie for the top ranking. A stationary view will show the three top–ranked words placed in the scene.

It might be expected that topic words would not be frequently clicked on, as doing so would apparently give little new information to the viewer. In fact, the opposite happens: people tend to strongly reinforce the popularity of topic words by clicking on interesting ones. Our suggested explanation is that the topics are easily accessible and supposedly significant words, and by clicking on them one obtains useful information in the form of the exact distribution of occurrences of the word. To constrain this feedback effect, we are looking at strategies such as showing all matching documents automatically when each new topic appears.

SHARED ENVIRONMENTS

The DIVE system is used to set Bead landscapes in a shared virtual environment. As described in [8], DIVE uses a spatial model of interaction involving volumes around people (and objects) known as aura, nimbus, and focus. This can be used for a level of detail mechanism as well as to support mutual awareness amongst users. Our imageability features have not directly been applied to users'

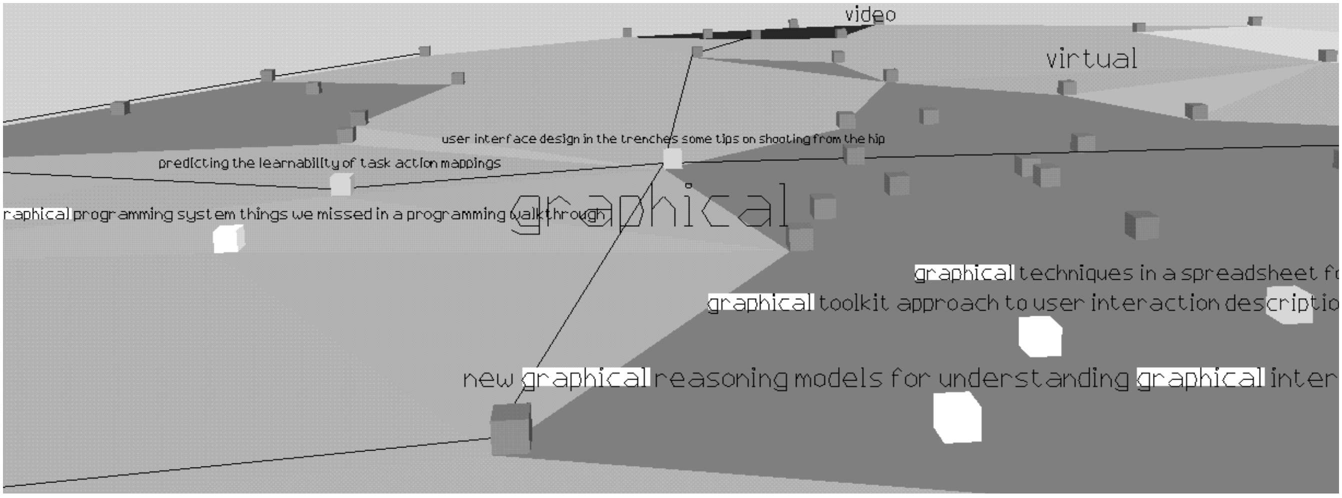

Figure 4. From a different viewpoint above the CHI91 landscape, different pop-ups and topics appear. Topic and title words (e.g. 'graphical') can be clicked on to start searches, which colour hits white. Search words can also be typed directly into the window. Each search increments the word's popularity rating, feeding back in to scene adaptation via a persistent database shared by all users of the system.

representations in the virtual environment, but there is an indirect effect worth noting, based on the way that the orientation of a pop–up title is kept perpendicular to the view of the user and aligned with the highlighted document as he or she moves. While this aids legibility for that user, it also displays the activity of the user to others. The orientation of topics is currently static. On creation they are aligned with a user's view but are often further away from the eye than titles, and therefore tend to be legible for some time without the same orientation maintenance. Searches are, at present, public actions i.e. they are visible to all users in that virtual space.

User motion and orientation is visible in DIVE, although shrinking due to perspective can make this difficult to discern from afar. Titles therefore give a clue to a distant user's position and motion, as they appear, disappear, swing and move according to that user's changing viewing frustum and sampling. Pop–ups and topics make the awareness of users and their activities richer in terms of semantics, showing information on 'what' as well as 'where'. Also, pop–up activity is reciprocal: 'you can see my pop-ups and I can see yours'. It is therefore difficult to pry without your own pop-ups giving away your presence and gaze. We suggest that such people finding themselves in such a situation is similar to two people finding out that they read the same magazine or work on the same author, informally revealing a shared interest that might lead to further interaction or collaboration.

Another aspect of user–user interaction is asynchronous. Past usage data can be accessed either indirectly e.g. by which titles appear in a region, or directly, such as when clicking on a usage disc. By doing this, one learns more about which documents and searches but also one learns about which people have worked in an area previously and what their searches and interests were. At present such data is not filtered or protected, but once done this should allow people to more comfortably make use of both past activity and raw data content.

ONGOING AND FUTURE WORK

The primary focus of our ongoing work is in making better and more varied use of usage data. To this end, we are working to make data on more of the virtual environment and the activities therein persistent.

We wish to implement policies which take into account privacy and invasiveness issues, such as those discussed in [3]. These issues will come more into the foreground as we add tools for filtering, combining and representing usage data. Such tools would allow users to select how items such as popularity assessments and topic words were made, using groups ranging from individuals to the whole community.

Another topic of current interest is the feedback of usage information into the layout and clustering process. An example is adjusting the landscape layout algorithm to weight words and documents in accordance with user activity, and not relying solely on its current criterion, the words contained in documents. The construction of features such as clusters should adaptively take into account user activity. The rankings used for topics may form a first step towards this. Following one anonymous reviewer's comments and reference to [14], we may also in future weight pop–up titles using relative frequency of search hits i.e. the same information as is shown in usage discs.

We have also been reimplementing parts of the Bead system in Java, so as to ease software development, to allow Web–based experiments and to facilitate extension to new data types. This has allowed us to lay out and visualise financial time series data such as sets of companies' stock values over time, and also some more heterogeneously structured files of customer records. This autumn we will collaboratively construct a simple pilot application with a section of our corporation concerned with foreign exchange trading. This will initially be a simple, single–user system, but we expect that having such a pilot running in the bank will offer valuable feedback and experience.

CONCLUSION

In designing displays of complex information, it is difficult to make them imageable and usable. Laying out a set of objects in an open, accessible structure is a good first step, but then showing the objects in their modelled positions with little detail often offers an insufficiently engaging and legible representation. Adding too much detail, as may occur with level of detail mechanisms, can easily lead to occlusion problems.

We have described our approach towards solving this difficulty, based on adding static and dynamic imageability features. In order to manage image complexity, techniques based in both sampling, as in pop–ups, and discretisation, as in clusters and topics, were applied. Dynamic features allow adaptation of the display in accordance with ongoing interaction. Usage history helps to guide this adaptation, tailoring it to those working with the visualisation.

More generally speaking, we have tried to take advantage of time and person. Although somewhat trite, we point out that people's activities, whether synchronous or asynchronous, have a past and a context. Most information displays do not make enough of this fact, but we consider that using this wider range of data characteristics affords circumvention of some difficult data analysis problems, and the enrichment of information designs.

ACKNOWLEDGEMENTS

Thanks go to Rob Ingram, for the work put in during his post–doc visit, and also to several student interns — especially to Christoph Pfranger for his work on the dynamic imageability features but also to Raimond Reichert and Roberto Brega.

The helpful suggestions from the anonymous reviewers are gratefully acknowledged, and I also offer thanks to Ramana Rao whose discussion and comment helped greatly.

REFERENCES

1. Ahlberg, C. & Shneiderman, B. Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays, *Proc. CHI 94*, ACM, 1994, pp. 313–317.
2. Bedersen, B. & Hollan, J. Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics, *Proc. UIST 94*, ACM, 1994, pp. 17–26.
3. Bellotti, V. & Sellen, A. Design for Privacy in Ubiquitous Computing Environments. In G. de Michelis, C. Simone and K. Schmidt (Eds.) *Proc. 3rd. European Conf. on Computer Supported Cooperative Work, (ECSCW'93)*, Kluwer, 1993, pp. 77-92.
4. Chalmers, M. & Chitson, P. Bead: Explorations in Information Visualisation, in *Proc. ACM SIGIR'92* (Copenhagen, June 1992), published as a special issue of SIGIR Forum, ACM, 1992, pp. 330–337.
5. Chalmers, M. Using a Landscape Metaphor to Represent a Corpus of Documents, *Proc. European Conf. on Spatial Information Theory*, Elba, September 1993. Published as *Spatial Information Theory*, Springer Verlag LNCS 716, A. Frank & I. Campari (eds.), 1993, pp. 377–390.
6. Chalmers, M. A Linear Iteration Time Layout Algorithm for Visualising High–Dimensional Data, to appear in *Proc. IEEE Visualization 96*, San Francisco, 1996.
7. The DIVE Home Page, http://www.sics.se/dce/dive/dive.html.
8. Fahlén, L. et al., A Space Based Model for User Interaction in Shared Synthetic Environments, *Proc. INTERCHI'93* (Amsterdam, April, 1993), ACM, 1993, pp. 43–50.
9. Fairchild, K., Poltrock, S. & Furnas, G. SemNet: Three–Dimensional Graphic Representation of Large Knowledge Bases. In R. Guindon (Ed.), *Cognitive Science and its Applications for Human–Computer Interaction*. Erlbaum, 1988.
10. Feiner, S. & Beshers, C. Worlds within Worlds: Metaphors for Exploring n-Dimensional Virtual Worlds, *Proc. UIST 90*, ACM, 1990, pp. 76-83.
11. Furnas, G. Generalised Fisheye Views, *Proc. CHI 86*, ACM, 1986, pp. 16-23.
12. Hill, W., Hollan, J., Wroblewski, D. & McCandless, T. Edit Wear and Read Wear, *Proc. CHI 92*, ACM, 1992, pp. 3–9.
13. Ingram, R. & Benford, S. Legibility Enhancement for Information Visualisation, *Proc. IEEE Visualization 95*, Atlanta, Oct. 1995, pp. 209-216.
14. Kleiboemer, A.J., Lazear, M.B. & Pedersen, J.O. Tailoring a Retrieval System for Naive Users, *Proc. 5th Symposium on Document Analysis and Information Retrieval (SDAIR96)*, Las Vegas, April 1996, pp. 209–216.
15. Lin, X., Soergel, D. & Marchionini, G. A Self–Organizing Semantic Map for Information Retrieval, in *Proc. SIGIR'91*, published as a special issue of *SIGIR Forum*, ACM, October 1991, pp. 262–269.
16. Lynch, K. *The Image of the City,* MIT Press, 1960.
17. MacKinlay, J.D., Robertson G.G. & Card S.K. The Perspective Wall: Detail and Context Smoothly Integrated, *Proc. CHI'91* (New Orleans, April 1991), ACM, 1991, pp. 173–180.
18. MacKinlay, J.D., Rao, R. & Card, S.K. An Organic User Interface for Searching Citation Links, *Proc. CHI95*, Denver, May 1995, pp. 67–73.
19. Perlin, K. & Fox, D. Pad: An Alternative Approach to the Computer Interface, *Proc. SIGGRAPH 93*, ACM, 1993, pp. 57–64.
20. Perlman, G. The HCI Bibliography Project, *SIGCHI Bulletin* 23(3), 1991, pp. 15-20.
21. Robertson, G.G., MacKinlay, J.D. & Card S.K. Cone Trees: Animated 3D Visualizations of Hierarchical Information, *Proc. CHI'91* (New Orleans, April 1991), ACM, 1991, pp. 189–194.
22. Tufte, E. *Envisioning Information*, Graphics Press, 1990.
23. Wise, J. et al., Visualizing the Non-Visual: Spatial Analysis and Interaction with Information from Text Documents, *Proc. IEEE Information Visualization*, October 1995, pp. 51–58.