Improving the Legibility of Virtual Environments

Rob Ingram and Steve Benford

Department of Computer Science
The University of Nottingham
Nottingham, NG7 2RD, UK

Tel.: +44 602 514203
Email: s.benford, r.ingram @cs.nott.ac.uk
Fax: +44 602 514254
http://www.crg.cs.nott.ac.uk/

1. Introduction

Years of research into hyper-media systems have shown that finding one's way through large electronic information systems can be a difficult task. Our experiences with virtual reality suggest that users will also suffer from the commonly experienced "lost in hyper-space" problem when trying to navigate virtual environments.

The goal of this paper is to propose and demonstrate a technique which is currently under development with the aim of overcoming this problem. Our approach is based upon the concept of *legibility*, adapted from the discipline of city planning. The legibility of an urban environment refers to the ease with which its inhabitants can develop a cognitive map over a period of time and so orientate themselves within it and navigate through it [Lynch60]. Research into this topic since the 1960s has argued that, by carefully designing key features of urban environments planners can significantly influence their legibility.

We propose that these legibility features might be adapted and applied to the design of a wide variety of virtual environments and that, when combined with other navigational aids such as the trails, tours and signposts of the hyper-media world, might greatly enhance people's ability to navigate them. In particular, the primary role of legibility would be to help users to navigate more easily as a result of experiencing a world for some time (hence the idea of building a cognitive map). Thus, we would see our technique being of most benefit when applied to long term, persistent and slowly evolving virtual environments. Furthermore, we are particularly interested in the automatic application of legibility techniques to information visualisations as opposed to their relatively straight forward application to simulations of the real-word. Thus, a typical future application of our work might be in enhancing visualisations of large information systems such the World Wide Web.

Section 2 of this paper summarises the concept of legibility as used in the domain of city planning and introduces some of the key features that have been adapted and applied in our work. Section 3 then describes in detail the set of algorithms and techniques which are being developed for the automatic creation or enhancement of these features within virtual data spaces. Next, section 4 presents two example applications based on two different kinds of virtual data space. Finally, section 5 presents some initial reflections on this work and discusses the next steps in its evolution.

2. What is legibility?

Legibility, in the context of navigation and wayfinding, is a term which has been used for many years in the discipline of City Planning. Work on legibility in this area has been concerned with the way in which people are able to 'read' an environment and hence perform wayfinding tasks. In his book "The Image of the City" [Lynch60] Kevin Lynch defines the legibility of a city as: "…the ease with which its parts may be recognised and can be organised into a coherent pattern…" Here, Lynch is referring to the formation of a *cognitive map* within the persons mind [Passini92], a structure which is an internal representation of an environment which its inhabitants use as a reference when navigating to a destination. The Image of the City describes experiments carried out in a number of major US cities which show how the cognitive map is built up over time through experience of the city. The experiments involved obtaining information from long term inhabitants of the cities in the form of, for example, interviews, written descriptions of journeys through the city and drawn maps. By examining this data Lynch identified five major elements of urban landscapes which are identified by the inhabitants and used as the building blocks of the cognitive maps. These features are:

- **Landmarks**. Static and recognisable objects which can be used to give a sense of location and bearing
- **Districts**. Sections of the environment which have a distinct character which provides coherence, allowing the whole to be viewed as a single entity
- **Paths**. Major avenues of travel through the environment such as major roads or footpaths
- **Nodes**. Important points of interest along paths, e.g. road junctions or town squares
- **Edges**. Structures or features providing borders to districts or linear obstacles

3. Legibility techniques for virtual environments

The aim of ongoing research at Nottingham University is to apply the work described above not to real environments but to the artificial spaces of virtual reality systems. More specifically we are developing techniques to automatically construct the five legibility features in the abstract spaces produced by data visualisation systems such as database or document store visualisers. One of our main aims is to accomplish this with without requiring the users of the system to perform the placement of the features manually. Essentially the system should, wherever possible, identify and place the features using information available from the database and visualisation systems alone.

To do this we have constructed a prototype system called LEADS (LEgibility for Abstract Data Spaces) which is designed to provide a layer on top of existing visualisation systems and which performs the addition of legibility information to the space. LEADS acts on the position of data items provided by this underlying system, as well as accessing the raw data where necessary, to add and emphasise the objects which are used to improve the legibility of the environment.

LEADS is designed to be applied to spaces which satisfy three main criteria. They should: be persistent over relatively long periods of time; be relatively stable so that they evolve over their lifetime and are rarely disturbed by major upheavals in the database; be accessed repeatedly by a number of independent users. An example of a large space to which application of LEADS techniques might be appropriate is the WWW space, which is constantly evolving but which rarely undergoes global scale restructuring.

To place the legibility features LEADS uses districts as a starting point as this allows for a number of relatively simple techniques to be used to form the other features. Districts are areas of cities which are identified as a single unit because of some common theme which runs through the buildings contained within them. In a data space this idea maps onto clusters of items which have some sort of similarity to each other which they do not share with nearby items in the rest of the space. In order to identify districts within an arbitrary data space the LEADS system uses the technique of cluster analysis. In this area a number of algorithms have been developed to carry out just this task. For the initial prototype of the system we needed to identify one of these algorithms which would have a relative simple basis, making it easy and quick to implement, but which carried out the task effectively, and generally, enough to be able to identify natural clusters in a wide range of data without being too computationally expensive. The particular algorithm which was eventually chosen is Zhan's Minimum Spanning Tree Algorithm [Zhan71].

This clustering technique is basically a two stage process. In the first stage a minimum spanning tree of all the items in the space is formed using any well-known algorithm. This assumes that the space contains some sort of graph structure from which the spanning tree may be produced. Whether this occurs naturally depends on the type of space being visualised, for example a visualisation of a large computer network will necessarily contain a graph structure already but this is may not be the case with a representation of a document store. In such cases we would assume the space to be fully connected, with each item having a link to all others. This has an impact for the amount of computation required to form the tree but has not proved a major problem with the size of spaces examined so far. To form the minimum spanning tree all of the links must have a weight associated with them which is their value. In a space which already contains a network these values may already exist but where they do not it is necessary to compute some measure of the 'distance' between the data items. How this is carried out must depend on the contents of the database itself and the structure of the data items, and so cannot be determined in a general sense for all spaces. Comparison methods for different types of data have been developed in the area of cluster analysis and seem appropriate to be applied to this problem [Jain88].

The second stage of the clustering process involves walking through the spanning tree and comparing the length of the links with those nearby. If the link is found to be significantly larger than those nearby it is marked for elimination. This removal of links leaves a number of disconnected subtrees whose items will form the clusters. It is necessary to provide the algorithm with criteria with which to decide if a link should be eliminated. The system at the moment uses as its default measure criteria similar to those used by Zhan in the referenced paper, that a link should be eliminated if it is greater than 2.5 times the standard deviation of the average length of the neighbouring links *longer* than that average and also the overall length of the link is twice that of the average. Links used to form the average are those within two steps of the candidate edge. Every effort has been made to allow these values to be changed easily to adapt to different data spaces.

The requirements of landmarks are that they form stable reference points in areas of the space where they will be most useful. A number of methods of placing landmarks based on the district data were considered. The most simple solution would be to position the feature at the centre point of each district. Two possible values exist for this point, the geometric centre of the area and the centroid of the cluster based on the values of the data. The latter is a useful value to discover, being the mean point of all the data in the cluster and therefore in some respects being the value that most accurately represents the *character* of the cluster. It was felt however that this positioning may not be the most useful as it takes no real account of any local conditions or variations of the data. A second suggestion places landmarks at points where more than two districts intersect. This method was thought more useful as the landmarks will provide useful external references for observers within the clusters but again it does not really take account of the actual data nearby. The solution that was finally chosen was to find groups of three clusters which were mutually adjacent and compute their centroids. The landmark would be placed in the centre of the triangle formed by these centroids. This has the advantage of providing useful external references but only where clusters are measurably adjacent, and therefore relatively dense.

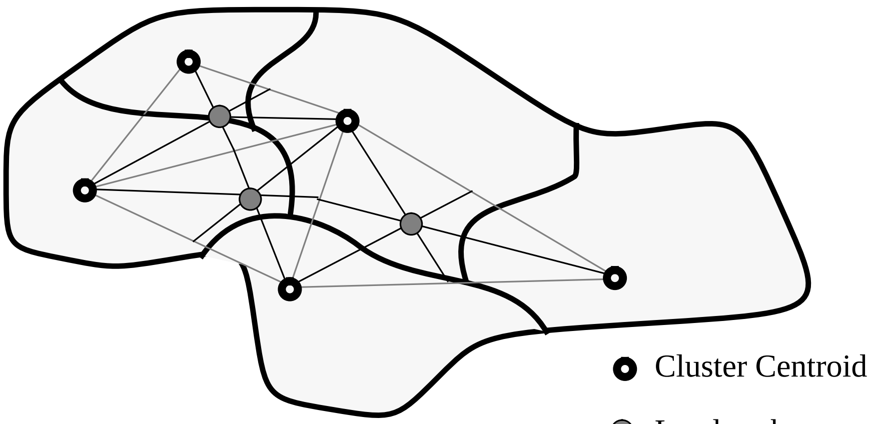

fig.1: The Centroid Triangulation Method of positioning Landmarks

Edges in LEADS are structures which help to define the borders of districts. They are placed in the space between those districts which are 'significantly' large. The value of significance is currently defined in a simple manner as 'having more than a single entry' although this measure could easily be changed to some more complex criterion. Ideally, edges would be relatively complex surfaces which follow the shape of the edges of the clusters they divide. For the purposes of the initial prototype concept demonstrator of the system a simpler method was chosen where edges are placed between the nearest neighbours in adjacent clusters and aligned along the line that joins them. In most cases, especially where the clusters are essentially spherical, this results in an edge placement which effectively separates the space but does not cut into the individual districts.

Finally the system must place the path and node elements. These features are very strongly linked in the real world and so in the LEADS system they are co-dependant. The eventual aim is to develop method whereby the paths will evolve from the data space on the basis of the way in which the users interact with the database and move through the environment. This means that all access and modification of the database will be logged and this usage information used to form the paths, therefore placing these thoroughfares along the most accessed routes though the data. Nodes will be formed at the intersection of paths and at very well used items along them. From this it can be noted therefore that while paths form additional objects in the space the formation of nodes will be accomplished by enhancing or highlighting data items which already exist in the virtual environment.

This method of developing paths is not viable at the current stage of work and with the data sets now available therefore the initial prototype currently uses an interim method to produce paths so that their impact in the space may be, at least visually, considered. This placement method identifies nearest neighbour data items between districts and designates them as gateway nodes, placing a path between them. Within districts a spanning tree of all the nodes is produced and these edges are also added as path links.

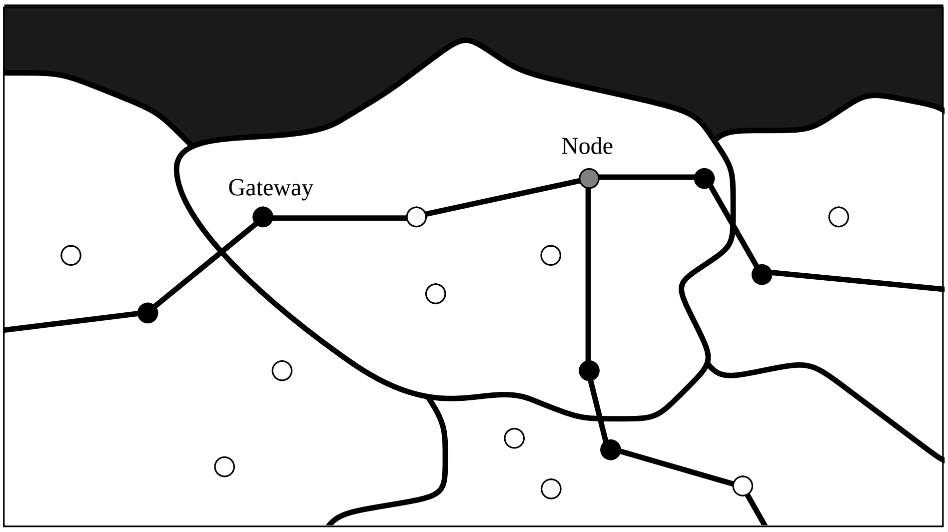

fig2 Paths joining four Districts

3.1 Other Navigation Devices

In addition to the legibility features LEADS also contains other devices to aid navigation which we will describe here briefly. In the field of city planning signposts are often considered to be something of an admission of failure [Canter84], that the planner was unable to produce a sufficiently legible environment, but we concur with the view of Passini that they are in fact a useful, even invaluable tool during wayfinding tasks. For this reason we are endeavouring to include signposts in our environments where possible. Currently these appear attached to the paths near to the node objects and identify the item to which the path leads but the use of signs will continue to be expanded as the system develops.

Another tool in the space is an optional axis object which floats near the occupant of the space and remains constantly aligned to the three major axes. The purpose of this object is to allow users to maintain bearing in spaces where there may be none of the references we are familiar with such as a horizon line. It is also planned to extend this object to become something like a ‘portable sign’ by annotating it with text to indicate interesting nearby objects in the direction in which the axes are pointing.

Finally, we wish to allow users to backtrack and move to interesting places by providing history mechanisms and the ability to store ‘bookmarks’ such as those widely used in hypertext systems.

4. Two example applications

We have so far applied LEADS to the visualisation produced by two existing information visualisation tools. The first of these, Q-PIT, is a system which works on databases where the items have a number of well defined fields [Benford94]. Three of these fields may be chosen so that the values they contain are mapped onto the three major axes of the space to give the position of the data items in a scatter-graph style layout. The remaining fields may be used to define aspects of the representation of the items such as their shape, colour or speed of rotation. This system could be used to visualise a variety of databases having the ability to map both numerical and string values to the axes and could provide a number of different visualisations of the same data by changing which fields are chosen to give the object positions. This gives the users of the system the ability to choose the visualisation which best represents their interest in the data. Q-PIT also allows multiple users to simultaneously access the data space and interact with the data. Any changes thus made to the underlying database are animated in real time in the visualisation.

The second system, FDP (Force Directed Positioning), is a 3D graph drawing tool. This takes a representation of an arbitrary network and produces a 3D visualisation by representing the nodes as masses and the links as springs. Initially the items will be placed randomly and the system will then go through a cycle of repositioning the nodes based on the tension in the springs until a relatively stable formation is found. Each link in the graph can be given a weight which will alter the way in which the tension value is calculated. This system also allows multiple users to enter the data space and query the values given to the nodes. Each user can therefore see which other people are interested in the data and on which items they are focusing their attention, which could potentially lead to interaction around a common subject.

LEADS attempts to enhance the legibility of both systems through the above techniques. Figure 3 shows before and after shots from Q-PIT. The data shown is around 120 items which represent people, having fields such as name, location, age and occupation. LEADS identifies a number of significant districts in the data as indicated by the different shapes (and colours) of objects and the presence of the edges and paths. One landmark is placed between a cluster of three districts. Figure 4 shows before and after shots from FDP[1] where the space is a randomly generated test file of more than 200 simple nodes and unit weight links. The second image shows a close up of the intersection between two of the clusters found in this space and the edge that is placed between them.

5. Reflections

Our initial experiences of developing and testing LEADS have been largely positive and we suspect that this approach does have the potential to significantly improve the legibility of virtual environments. However, we have encountered a number of difficulties which suggest some immediate improvements to the system:

First, legibility features such as paths seem to sit uncomfortably with 6 degrees of freedom navigation. Paths in a city are actually travelled along and the environment is therefore experienced from the perspective of the path. We suspect that users of our legible virtual environments should also directly experience paths as part of navigation. Thus, we need to develop simple interface techniques to support navigation via paths (e.g. “take me

1. Because of the extensive use of colour in LEADS these plates can not fully demonstrate the effectiveness of the enhancement. We hope however that the do give some idea of the effect of the addition of legibility features.

along the nearest path in this direction").

Second, automatically determining an appropriate scale and appearance for legibility features has not always been easy. In particular, features such as landmarks and edges must be visible without being intrusive. Creating useful edges has proved to be a particularly difficult task as an edge should ideally follow the contours of a district. In a 3-D space, determining a sensible size for edges has been difficult. At one extreme an edge might be a hull completely surrounding a district. At the other it might be a thin flat surface dividing two districts. The former is likely to be visually intrusive; the latter is likely to provide an insufficient sense of separation between districts.

Third, other features and tools are clearly needed to help people navigate. For example, the use of textual information in the form of signposts is an important part of navigating conventional urban environments. We can imagine adding signposts to virtual environments and placing them at nodes pointing along paths. Furthermore, we can imagine that signposts might refer to districts and landmarks that lie along a path. However, this gives rise to the problem of how to name districts and landmarks. More specifically, given that districts and landmarks are automatically created by LEADS, we are left with the problem of automatic name generation or alternatively the use of non-textual symbolic identifiers on signposts.

Finally, we need to be careful that by adding additional objects to an information visualisation, we do not increase the rendering overhead thereby degrading system performance. However, we suspect some extensions to LEADs might enable it to actually improve system performance. Specifically, LEADS might enable a general distancing effect for data spaces by allowing the individual objects in a district to be replaced by an overall representation of the district when viewed from a distance.

6. Summary

We have described a number of general techniques for improving the legibility of virtual environments so that their users might more easily construct cognitive maps to help them navigate. Our work has adapted techniques from the discipline of city planning where decades of experience have identified key features of urban landscapes which are critical to their legibility. The primary goal of our work has been to develop a set of algorithms for *automatically* creating or enhancing these features within information visualisations. These include:

- the use of clustering algorithms to create districts;
- the creation of edge objects separating districts;
- the placement of landmark objects at a central point between districts;
- emphasising nearest neighbour node objects within districts and creating paths between them.

We have implemented these techniques in a system called LEADS which is intended to provide an additional legibility layer sitting on top of current information visualisations. So far, we have applied LEADS to two existing and contrasting information visualisation tools. Our paper included some before and after screen-shots to show the overall effect of LEADS on these systems.

Our early experiences have been positive and suggest that this approach is promising. However, we have encountered several difficulties including the need to experience paths when navigating; the difficulty of getting an appropriate scale for features such as edges and landmarks; and problems with automatically naming features so that they may be referred to by signposts.

Longer term work will include combining LEADS with larger scale, more widely used information sources (e.g. visualising the World Wide Web) and carrying out a more formal evaluation of our work. For the latter, we intend to revisit and adapt the initial experiments carried out by Lynch within real cities, but this time within virtual environments.

7. References

[Benford94] Steve Benford, John Bowers, Lennart Fahlén, Chris Greenhalgh, John Mariani and Tom Rodden, *Networked Virtual Reality and Co-operative Work*, Presence, MIT Press, (in press).

[Canter84] David Canter, Way-finding and Signposting: Pennance or Prosthesis, Nato Conference on Visual Presentation of Information, published as Information Design (Easterby and Zwaga eds.), Wiley, 1984

[Jain88] Anil K. Jain and Richard C. Dubes, *Algorithms For Clustering Data*, Prentice-Hall, 1988

[Lynch60] Kevin Lynch, *The Image of the City*, M.I.T. Press, 1960

[Passini92] Remedi Passini, *Wayfinding in Architecture*, Van Nostrand Reinhold, 1992

[Zahn71] Zahn, C.T., *Graph-theoretical Methods for Detecting and Describing Gestalt Clusters*, IEEE Transactions on Computers, C 20, 68-86