Mindless Visualisations

Claire Knight and Malcolm Munro

Visualisation Research Group,
Research Institute in Software Evolution.
Department of Computer Science,
University of Durham,
Durham, DH1 3LE, UK.
{C.R.Knight, Malcolm.Munro}@durham.ac.uk

Abstract. The wonder and, unfortunately, to the detriment of visualisation for the representation and comprehension of complex data sets is that to be most successful requires that they are tailored to suit the task and underlying data. Such a restriction enables visualisations to be well designed for the tasks to which they are known to be applied to, and also to accommodate the style and range of data to be expected as normal. The problem with this repeated redesign of visualisations is that the interface is often neglected, and can even be solely dependent on the implementing technology used for the visualisation. It is important to add such issues as the interface to visualisation considerations, and to provide reusable concepts that will integrate with a range of metaphors and displays. This position paper examines the issues surrounding such visualisation interfaces and presents a discussion of those issues.

1. INTRODUCTION AND AIMS

Visualisations are usually tailored to the data set that provides the underlying rationale for that visualisation to exist and have been developed. This is not so much a problem in its own right, as a position where the visualisation research community has not gone far enough to enable tailoring of the visual display. It is now common for there to be some form of control panel to allow which of the data items are displayed or which relationships that are part of that data are emphasised in the visualisation. What is much less common is the provision of many different representations and metaphors for the same data, and all possible from the same interface.

There is a requirement for such visualisation interfaces, simply because of the variability in users, their preferences, and their varied ways of working when performing assorted tasks. It has taken some time for the view that visualisations are task and data dependent – at least to some extent – to be accepted. Early attempts at information visualisation produced many useful displays, but trying to apply these too widely without change was misguided. Data transformation to fit the visualisation is not necessarily a good step when that visualisation is trying to provide avenues for human insight into that data. Any translation of that data may well affect any patterns or relationships hidden within the mass of raw information.

A related issue of task and data dependence is that of interface and interaction suitability. Issues of the number of dimensions used for the actual visualisation aside, it is often the case with current technologies and implementations that the interaction mechanisms are unadaptable. This is not good from the obvious view of any users who suffer from a disability that affects their interaction with computers and their various technologies. It is also not acceptable for the general user. Being able to specify the actions of the mouse or to use personal key combinations is an important, but relatively small step, often left out of "proper" user interfaces. But one that has long been embraced by the gaming community, where different games can be controlled by exactly the same key combinations should the user wish. These variations in interface from both interaction and task viewpoints need to be considered in order to produce the most accepted tools of the future.

Visualisations that allow for visual exploration of the data sets provide these opportunities for human insight; they act as intelligence amplification tools. In effect what is being aimed for is the inverse of intelligent user interfaces. What is required are interfaces that provide the flexibility required by users for carrying out a variety of tasks in the context of the visualisations with appropriate configuration options and thus allowing the user to exploit the tool by using it to augment their abilities. Such interfaces do not benefit from so called intelligent interfaces where much is made of data mining, learning algorithms, and suchlike. These techniques have numerous other users but in such a situation their second guessing and information presentation lend themselves only to possibly obscuring patterns within the data being visualised which may be of such a subtle nature that they require human intuition for recognition.

This position paper presents some visualisation issues that are then related to the need to be able to create configurable interfaces for such tools to enable better acceptance and flexibility. The view of the user in control is the preferred one, with intelligence amplification the driving force rather than trying to provide intelligence in interfaces. This is evident in the later sections of the paper where such systems are discussed.

2. BACKGROUND

This section provides a brief overview of the main areas of literature that influence a visualisation and thus its interface. They do not centre on particular applications or domains of visualisation, but on the general theory that is more widely applicable to a range of visualisations. These areas set the scene for the ideas presented later in the paper by providing a context for why tool interfaces need more research.

2.1 Metaphors

A metaphor is where a word or phrase (or in terms of visualisation, a graphical representation of that word or phrase) is used in place of another. This tends to suggest some form of analogy between the two concepts, although this may be at a higher level of abstraction than individual words or phrases. Blackwell [Blac96] poses the question of whether these abstractions should be seen as metaphor or analogy, although a discussion of the distinctions and use of these terms is beyond the scope of this thesis. From a VR perspective the metaphors act as a mapping from the concepts required in the virtual world to their graphical representation. This need was identified by Levialdi et al. [Levi95] in the construction of their database visualisation system.

> *"Using VR visualization techniques to represent the results of queries implies the definition of a mapping, or* metaphor*, among the objects of the database and the objects of some virtual world."*

According to Benford et al. [Benf96] the use of natural metaphors can aid the usability of virtual environments.

> *"… an attempt to exploit people's natural understanding of the physical world, including spatial factors in perception and navigation, as well as general familiarity with common spatial environments…"*

Fitzpatrick et al. [Fitz96] also apply the spatial metaphor to the level of social interaction possible within the virtual world representation of the metaphor.

> *"Even though space is an intuitive, familiar metaphor to work with, there can be a more encompassing meaning of space in the virtual world, independent of graphical and VR depictions, that is driven by social world needs and the needs of individuals participating in multiple social worlds."*

This view of using real world interface metaphors is also supported by Väänänen and Schmidt [Vään94, Vään93]. The authors are of the view that these types of metaphor solve many navigation issues because they impose familiar structures and interaction possibilities on the system and these are visually recognisable by the user.

Pettifer and West [Pett97a] suggest that the potential power of VR comes from the strength of its metaphor, and the fact that it is closer to natural interaction than many other forms of computer system. They also identify the benefits of natural metaphors, and making use of perceptual and spatial skills learnt and used in the real world in the virtual environment.

> *"A three-dimensional world metaphor has much more scope for direct human/computer interaction than the two-dimensional desktop because it engages in us those perceptual and spatial faculties that allow us to comprehend our surroundings and to process effortlessly the vast amounts of information that are presented to our senses second by second. It is the potential to directly engage these faculties that is the defining characteristic of virtual reality. As the immersive environment is far richer than the desktop, the metaphors for interaction assume a far greater significance. … The role and management of metaphors for the virtual environment therefore assumes key significance."*

It is obvious from the above that the design of the metaphor used in the virtual environment can play a large part in the usability of that system, both in terms of human-computer interaction, and in terms of enabling the user to carry out the required tasks. What is also of benefit is that in using three-dimensional environments some of the cognitive processing needed for navigation and visual interpretation can be shifted to the sub-conscious as these are activities that are carried out daily with no real thought.

Metaphors are often criticised for hiding the original data or causing the user to have false expectations of what an object does or is capable of. Monin and Monin [Moni94] have this view. There is the distinction between metaphors that work because of some direct resemblance between two things and others that work through some common attitude to both things. This common attitude often a direct result of accidental and extraneous reasons and that a disparity between the two facets can hold potential dangers relating to comprehension and expectation.

This problem of the use of metaphors is summed up eloquently by Wiss and Carr [Wiss98]:

> *"As always, metaphors are difficult to find and easy to abuse."*

They also say that proving a metaphor works is a difficult task, and certainly this has to be true in a wider sense because of the variability of the users and tasks. What is more contentious is that such systems require some formal way of assuring users or sponsors that metaphors work, simply because of this known variability. Ideally each system would be able to support a range of metaphors to cater for all tastes, but as with every interface the more it is used the easier it is to use it, and once a system has been "learnt" the metaphor will become more acceptable thus invalidating the need for such figures or proofs.

An experimental study carried out by Dutton et al. [Dutt99] support the use of metaphors as their results led to the conclusion that after the initial use of the metaphor based system, responses and performance were enhanced in systems with metaphors. Their studies also showed that (in this particular study) the two different systems founded on metaphors outperformed the system without, but that the two metaphors were as good as each other.

2.2 Spatial Orientation and Navigation

If the VR environment is a representation of the spatial world that we already know then there is a need to model orientation and navigation features found in the real world. In any spatial setting some form of base orientation needs to be found which can then be used for navigation and re-orientation as movement occurs. Hemmje et al. [Hemm94] relate this to their database visualisation work although what they write is readily extendible to all spatial visualisations.

> *"It is necessary to* move*, i.e. change position in the context space and* explore *information visible from each point of view. It is important to achieve an* orientation*, i.e. to determine the relation between a* current point of view *(e.g. from an information item) and the whole of an information space."*

Many authors document the problems of getting lost in "cyberspace" when dealing with spatial virtual environments. Ingram and Benford [Ingr95] write

> *"More recent experiences with virtual reality suggest that users will also suffer from the commonly experienced "lost in hyperspace" problem when trying to navigate virtual environments."*

They relate the orientation and navigation processes to the cognitive map the user has of the environment. Cognitive maps can be one of two sorts. Linear maps are based on movement through the space and the observations made during that movement. Spatial maps do not require movement through the space. Generally, linear maps are the first created of an environment, and over time the map may evolve to being a spatial map. Exploration rather than guidance through an environment encourages the development of a spatial map. Their research has focused on providing ways to ease the navigation (and orientation) problems that occur in VR.

Pettifer and West [Pett97a] also relate the problem to the systems and metaphors in use today.

> *"Losing a cursor on the desktop is one thing, losing yourself in cyberspace is quite another."*

Three-dimensional worlds are potentially infinite whereas desktops are of generally finite space even if current implementations are able to cover several screens.

Hubbold et al. [Hubb93] discuss design issues that are important to consider for VR systems and cover orientation when discussing perceptual consistency.

> *"More important is the creation of an environment in which the user remains comfortable and well oriented."*

Pettifer and West [Pett97b] also comment on the construction of virtual environments, and that the aim must be to construct these environments so that they correspond with human perceptual requirements. Backing up these comments made by the above authors, Pesce [Pesc93] asserts:

> *"The first perogative in the engineering of a holosthetic environment is: design to avoid disorientation. Disorientation represents a step towards the amputation of the self, and necessarily precedes the dislocation of self that concludes in holosthetic psychosis."*

Another aspect of perceptual orientation, often missed, is that of causality. It provides a continuity of experience in "reality" so by providing such continuity in virtual realities allows natural comprehension, interaction and orientation. This is not implying that the causalities need to model exactly the laws of time and motion, but that the "laws" used in the environment need to be continuous throughout that environment, allowing things to be comprehended, and to an extent, explainable. A ball floating in mid air is considered strange, but provide a context of outer space and the ball's behaviour is perfectly acceptable. Attention is given to the issue of causality by Pettifer and West in [Pett97b] and Pettifer in [Pett96].

Feiner and Beshers [Fein90] cover the concept of *n*-Dimensional virtual worlds, but restrict their work in this paper to abstract visualisations. This has much use, but is not the only sort of visualisation that may need to be viewed and navigated. One counter example is when virtual visualisation spaces are inhabited by more than one user at any one time. Nevertheless, for specific data and tasks this framework is useful.

Crossley et al. [Cros97] give reasons why VR interactive interfaces can allow an intuitive and natural way to explore and comprehend complex information:

> *"A well-designed user interface with good spatial representation of information can be effective in assisting the user in the following tasks:*
> - *browsing and navigation,*
> - *searching,*
> - *comparing,*
> - *grouping,*
> - *analysis,*
> - *creating new information."*

The authors also recognise the importance of metaphors and navigation when using such interfaces.

It is easy to cause navigation and orientation problems if attention is not given to the design of the virtual environment. This would obviously make the system worse than two dimensional graphics or plain text because the cognitive overload gets so large. Conversely, if suitable attention is paid to the design of the virtual environment, the metaphors used, the interface between the environment and the user, and the use of suitable "laws" (relating to the metaphor if the metaphor allows) then there is a great potential for the use of VR and virtual environments.

2.3 IA not AI

Intelligence amplification (IA) is the use of computers to aid and enhance human intelligence rather than the artificial intelligence (AI) aim of trying to substitute humans with computers. Intelligence amplification builds on the skills that humans already have, and tries to augment the areas that are lacking in some way. Frederick Brooks (documented in by Rheingold [Rhei92]) describes his beliefs about intelligence amplification in the following way

> *"I believe the use of computer systems for intelligence amplification is much more powerful today, and will be at any given point in the future, than the use of computers for artificial intelligence (AI). In the AI community, the objective is to replace the human mind by the machine and its program and its data base. In the IA community, the objective is to build systems that amplify the human mind by providing it with computer-based auxiliaries that do the things that the mind has trouble doing."*

Brooks identifies three areas in which humans are more skilled than computers. The first is *pattern recognition* (aural or visual). The second is in performing *evaluations*, and the third is the *overall sense of context* that allows previously unrelated pieces of information to become related and useful in a new situation.

Walker [Walk95] also touches on the subject of intelligence amplification in his discussion on the challenges of visualisation.

> *"A natural and intuitive visual interface can retain the critical contribution from human perceptual skills, ensuring that opportunities for lateral thinking or perhaps an unexpected leap of imagination are not lost. Programming a computer to "look for something interesting" in a database is a major undertaking, but given appropriate tools, it is a task for which humans are well equipped."*

The first sentence can be seen to be similar to the third skill identified by Brooks, that of a sense of context. The second sentence by Walker is essentially talking about the pattern recognition skill specified by Brooks (in [Rhei92]).

Intelligence amplification is of importance to software visualisation (and any other form of visualisation) because in representing large and complex data sets graphically the aim is to help the user to get a better understanding of content of the data sets. By aiding the user in this way visualisation tools are acting also as intelligence amplification tools. Reading through many thousands of pieces of information and then summarising them in a finite graphical space would be an immense, complex and possibly tedious task. For a computer with the right "instructions", it is a simple data processing exercise.

Hubbold et al. [Hubb93] make a similar connection with the field of VR (and therefore visualisations that make use of VR as an enabling technology). They also identify the pattern recognition and contextual abilities of humans.

> *"In our everyday existence we cope with, and filter out, tremendous amounts of information almost effortlessly and with very little conscious thought. Indeed, if the same information, in all its detail, were to be presented in a form that we had to think about consciously, then we would be overwhelmed quite easily. Spatial awareness, pattern recognition, information filtering, coordination of multiple information streams are things we take for granted. Rather than look for a solution in AI, part of the VR thesis is that information presented in a suitable way can be processed far more effectively and directly by people."*

The role of a visualisation system as an intelligence amplification tool rather than as a system that tries to second-guess the information the user requires is emphasised by Crossley et al. [Cros97]:

> *"…the role of the system is not to select documents similar to a user-supplied query but to organise and display information about many documents in such a way as to assist users to select useful documents on their own."*

This shows that the important challenges and research issues for visualisations are to be able to handle such tasks well and provide the necessary support as transparently as possible. Changing the query mechanism in order to improve performance (for example in the situation above) is not going to help in another situation or be widely applicable to other visualisations.

2.4 Task Dependence

As with facilitating the transfer of tacit knowledge, the task to which the visualisation will be put has a role to play in the design of the metaphor (thus representation) and the environment in which the visualisation is located. The importance the task places on the visualisation design is elucidated by Kennedy et al. [Kenn96] in their framework information visualisations. Eick [Eick97] also acknowledges that this is important:

> *"Since the analysis needs of each dataset are often unique, some of the best visualizations are task-oriented. These visualizations help frame interesting questions as well as answer them."*

It is only through the use of appropriate visualisations that the use of such systems will become accepted. In this case, "appropriate" considers not only the data set but also the analysis task. The visualisation has to lead to insight and understanding in some way to have any validity.

2.5 Tangible from Intangible

One of the main problems for software visualisation (and other forms of information visualisation) is of trying to create a tangible representation of something that has no inherent form. Therefore the aim is to visualise the intangible in an effective and useful way. Effective and useful here refers to the visualisation being able to increase the understanding of the user whilst reducing the perceived complexity.

Ball and Eick [Ball96] recognise this problem when they write

> *"Software is intangible, having no physical shape or size. After it is written, code "disappears" into files kept on disks."*

and

> *"The invisible nature of software hides system complexity, … ."*

Walker [Walk95] comments on the software being the intangible part of information systems when he writes

> *"Some aspects of an information system are tangible, but a major component is the software which is an abstract and invisible collation of computer instructions."*

Chapin and Lau [Chap96] also recognise the intangible nature of software

> *"Furthermore, software is intangible, and it is only the representation of the software which can be communicated between people and between people and computers."*

An important point to be drawn from this is the communication aspect. Since software is intangible and each programmer has his own mental representation then an effective visualisation can also act as a common frame of reference. In discussing pieces of the software either informally between colleagues or formally in meetings, if the participants do not concur over the code being discussed the discussion may as well not take place. Visualisation of the software can provide not only a graphical representation of the piece of code under discussion (for clarity over the section being discussed), but also allow the discussion to take place in the realm of that visualisation. This means that the discussion can be based around the visualisation and the code it represents rather than the piece of code. In doing this, the visualisation has provided a starting point for common understanding.

3. VISUALISATIONS

The above literature showed that visualisation can be a good way of providing intelligence amplifying tools for the analysis and understanding of complex data sets. This is essentially what all visualisations are trying to do, although variations along the teaching/learning and statistics avenues do exist. It can also be seen in the above sections. In the work done to date on software visualisation [Knig99a, Knig99b, Knig00a, Knig00b, Knig00c] the use of 3D has also been prominent. It can be seen that 3D visualisation are inherently spatial and therefore require navigation around the virtual space as well as the interface. Metaphors are used to create tangible representations from intangible data sources. The end result – the visualisation – then acts as an intelligence-amplifying tool for the purposes of comprehension and analysis. As an example, Figures 1 and 2 show sample images from the software visualisation papers cited in this paragraph. These show different views of the results of a static analysis of over 17,000 lines of Java source code.

Figure 1 - District within *Software World*

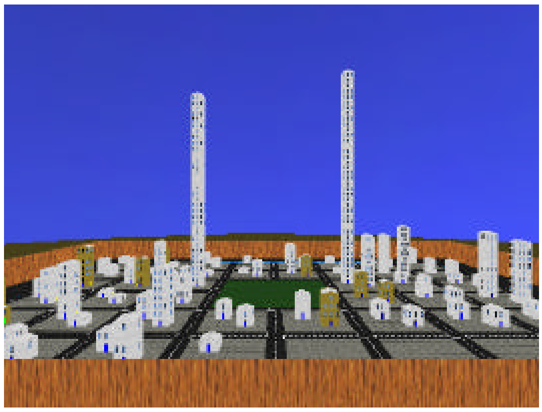

Figure 2 - Overview of entire district

In the past many of the issues surrounding software visualisation (more so than for other visualisations) have been based on the fact that all images are essentially nodes and arcs. The following issues also apply to these sorts of visualisation, but there is less scope for dealing with them. Graph layout is known to be a hard problem and layout algorithms have long been the focus of computer scientists; unfortunately focusing on mathematical properties of such rather than trying to address aesthetics and readability. To illustrate the background of some of this work Figure 3 has been included.

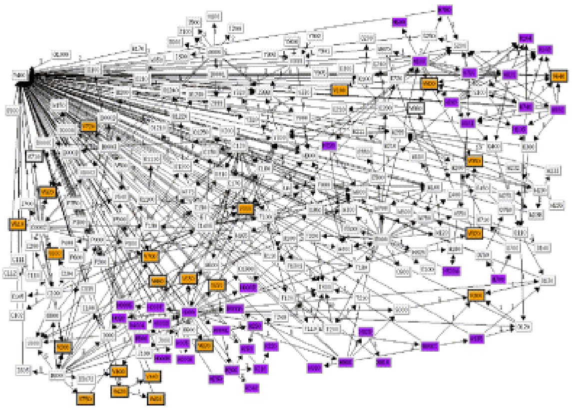

Figure 3 - Traditional software visualisation image

There are many unsolved visualisation issues [Knig00a] and these have a direct impact on are affected by interface issues. Those that are most pertinent here are those of:

- Evolution
- Scalability
- Navigation, Interaction, and Orientation
- Automation

In order to provide more detail each of these will be examined in the following four sections. The relation to interfaces will be highlighted within these sections.

3.1 Evolution

Evolution is an important issue with visualisations of program code, since software systems are known to change in a variety of ways and for a variety of reasons. It is also an overlooked issue for many more forms of visualisation. If a data set is large and complex, the sort visualisations are best suited to helping, then there is a high probability that some aspect of that data will change.

Once a visualisation has been generated it is useful for that visualisation to evolve as the underlying data evolves and to reflect the changes visually. Implementation issues of this aside, it is important for the visualisation representations and metaphor to be able to support this; if it cannot happen logically within the constrained framework that the metaphor provides, it might as well not happen at all. This is because the changes involved visually would cause too great a cognitive effort on the part of the user to be beneficial. In these cases (a) it would be better to generate the visualisation from scratch and (b) question whether the visualisation tool is of use for the work tasks of the user. It may be that the answer to (b) is that yes it is useful, even with a complete regeneration, since there is adequate relearning time but this cannot be assumed.

If the visualisation can evolve successfully, then does this affect the interaction or interface in any way? A first glance would suggest that it would have little bearing due to the metaphor and representation having to provide consistency for that change. A more detailed look, however, shows that control interfaces (for example) may be affected with choices as to which data is available for certain views. If that part of the data set ceases to exist, such parts interface is rendered obsolete. It may be the case that an alternative metaphor no longer works as the change was consistent for one metaphor and representations but a different one cannot generate suitable views. This would then affect the personalisation of the visualisation.

3.2 Scalability

Scalability of visualisations is related to the ability of a visualisation to evolve. Again, the only way to answer the question of how scalable a visualisation is requires it to be tested with varying amounts of source data. Scaling could be considered to be an evolution of the visualisation, but since it depends each time on the base code of the system, it is more of an issue with whether an initial development algorithm can handle a wide range of data sizes. A hard problem for designers of visualisations is that, on the whole, visualisations must be created to accommodate a very wide range of data. Essentially the visualisation has to be able to deal with one to an infinite number of items. Keeping this in mind during the development of the visualisations should enable them to scale better. It may be that some smaller visualisations developed for a very specific need, where the data is known to be limited, do not have to consider such scaling issues and this is quite acceptable. Just as long as when designing visualisations that can be applied to data that is known to vary in size and content this fact is borne in mind.

This would not have a direct impact on the interface and interaction mechanisms of a visualisation but does have side effects when considered in tandem with the issues of evolution (preceding section). Scaling would perhaps require interaction mechanisms to adapt to reflect the visual changes; moving around a virtual space that had expanded to ten times its original size would perhaps benefit from faster movement, but the control would otherwise be unaffected.

3.3 Navigation and Interaction

Navigation is important because it affects the usability of the visualisation. The visualisation should be designed and structured with navigation in mind. If navigational features are added as an afterthought it will then be hard to add the necessary paths and beacons. As Young and Munro [Youn98] write

> *"Well structured data terrain should also result in a more understandable layout and easier navigation".*

There are also guidelines for navigation and orientation that can be taken from city planning textbooks which indicate ways in which humans orient themselves in three-dimensional space.

Tied into navigation issues is the way in which any user of the visualisation is able to interact with it; to move around the landscape and to find the information they require must be as intuitive as possible to make people view visualisations as useful tools. Unfortunately for designers of visualisations all users have different wants and needs where interaction with computers is concerned. For this reason the more flexibility the system offers, the better. The ability for the user to have a degree of configuration is also likely to lead to the acceptance and use of the visualisation system.

3.4 Automation

The visualisation should be able to be generated from the data with minimal intervention. A configuration file of preferences is acceptable because the graphics are still created in a fully automatic manner. User generation of visualisations may allow tweaking for that user but the resulting visualisation is then only really suitable for that person. The cost of time needed to produce such visual displays is also high. The visualisation is then not really applicable to any visualisation aiming for consistent appearances between versions (changes in the data set). It also prevents a visualisation system being used as a common frame of reference for discussion. In the creation of multi-user visualisation environments the freedom for users to create their own landscapes would also completely destroy the notion of having a shared workspace – all users would have their own environments and each one (apart from their own) would be unfamiliar to everyone else. A much better solution to the problem of user preference is to provide various metaphors for a data set and allow personalisation through the choice of metaphor and which of the data set values to be included. The actual visualisation can then still be automatically generated.

4. USER VARIABILITY AND INTERFACE COMPONENTS

The previous section, in highlighting areas of concern for visualisation and interfaces, also touched on ways of handling those problems. As with any interface, familiarity breeds both contempt and an ease of working than generally only comes after many hours of exposure. The same can be said of visualisations; if the data set and metaphor are viewed often, anomalies or interesting areas will be visible much more quickly to the experienced user.

In order to minimise the impact of a new or changed interface, then a user should be able to customise it so that it is, at least in part, starting to become more like applications that they are already familiar with. The same applies to visualisations; metaphors and representations for data sets can be seen as contentious. By providing several alternatives, and by allowing data selections, or visual highlighting (within the visualisation), then the visualisation can be opened up to more users who would otherwise cite the display as a reason for not using it.

An approach currently being worked on for visualisations is the use of a *pluggable visualisation interface* that relies on the use of components containing visualisation definitions and graphics along with the relevant customisation options available for that display and data set. This allows for a choice of visualisations (metaphors and representations) for a given data set. It also allows for highlighting, isolated views, and a choice of which parts of the data set are visualised for a given view. The possibilities can become endless! This is not as easy a task as presented here. The choice of metaphors to provide for specific problems is not as defined as some might consider, and the generation of automatic 3D visualisations, whilst achievable, needs to be constructed to fulfil the scaling and navigation issues identified above. There is then the component interface, which needs to provide enough data for connectivity to the main application, whilst handling the visualisation and configuration options itself. This work is ongoing, but a screenshot of the current state can be seen in Figure 4. This illustrates multiple visualisations in the same application framework.

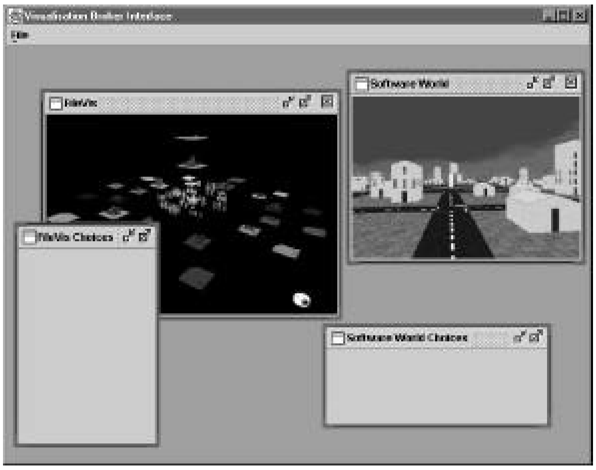

Figure 4 - Multiple visualisation interface

This is not so much about designed a visual (or visualisation interface) per se, but providing an application "container" to allow the user control of what they want to see and do. The component creators do impose some restrictions in what they provide for the user, but by having enough different components available should overcome this problem. There is obviously the side issue of having to define specifications of the interfaces at the component level, but this is hidden from a user of the visualisation application!

Such an approach is not without problems, but it does provide a step forward in handling user preferences in visualisations; more so than providing a small set of options on a panel or menu. A visualisation is normally used for some purpose; to achieve a (sub) goal. The visualisation application acts as an intelligence-amplifying tool by providing various facilities but not controlling what the user does. Its power comes from not being prescriptive in any way.

5. SUMMARY

The aim of these visualisation interfaces is to create the inverse of intelligent user interfaces. It is believed that the augmentation of user abilities is much more useful in this situation than trying to replace their intuition, experience, and pattern recognition skills. It is important that such interfaces are configurable, but not with some form of automated help. Such approaches are considered to be a hindrance when visualising for comprehension and analysis and wrong assumptions could abstract away from the most important areas of the data. These sorts of interfaces are also not necessarily aware of the ultimate goals to which the user's current activities may be contributing. Intelligence amplifying visualisations and their corresponding interfaces for configuration and use have great potential when used for analysis and comprehension of large and complex data sets. Anything that tries to overcome the inherent overload experienced with such data has to be worth more investigation. Starting to address the interface and usability of such diverse displays can only be a step in the right direction.

REFERENCES

[Ball96] **T. Ball** and **S. G. Eick**, *Software Visualization in the Large*, IEEE Computer, pp33-43, April 1996.

[Benf96] **S. Benford, C. Brown, G. Reynard,** and **C. Greenhalgh**, *Shared Spaces: Transportation, Artificiality and Spatiality*, Proceedings of the 1996 ACM Conference on CSCW, pp77-85, November 16-20, 1996.

[Blac96] **A. F. Blackwell**, *Metaphor or Analogy: How Should We See Programming Abstractions?,* Proceedings of the 8th Annual Workshop of the Psychology of Programming Interest Group, pp105-113, 1996.

[Chap96] **N. Chapin** and **T. S. Lau**, Effective Size: An Example of Use from Legacy Systems, Journal of Software Maintenance: Research and Practice, Vol. 8, pp101-116, 1996.

[Cros97] **M. Crossley, N. J. Davies, R. J. Taylor-Hendry,** and **A. J. McGrath**, *Three-dimensional Internet Developments*, BT Technology Journal, Vol. 15, No. 2, pp179-193, April 1997.

[Dutt99] **R. T. Dutton, J. C. Foster,** and **M. A. Jack**, *Please Mind the Doors - Do interface metaphors improve the usability of voice response services?*, BT Technology Journal, Vol. 17, No. 1, pp172-177, January 1999.

[Eick97] **S. G. Eick**, *Engineering Perceptually Effective Visualizations for Abstract Data*, In Scientific Visualization Overviews, Methodologies and Techniques, IEEE Computer Science Press, pp191-210, February 1997.

[Fein90] **S. Feiner** and **C. Beshers**, *Worlds within Worlds; Metaphors for Exploring* n-*Dimensional Virtual Worlds*, Proceedings of the ACM Symposium on User Interface Software and Technology, pp 76-83, October 3-5, 1990.

[Fitz96] **G. Fitzpatrick, S. Kaplan,** and **T. Mansfield**, *Physical Spaces, Virtual Places and Social Worlds: A study of work in the virtual*, Proceedings of the 1996 ACM Conference on CSCW, pp 334-343, November 16-20, 1996.

[Hemm94] **M. Hemmje, C. Kunkel,** and **A. Willett**, *LyberWorld - A Visualization User Interface Supporting Fulltext Retrieval*, Proc. of SIGIR '94, 1994.

[Hubb93] **R. Hubbold, A. Murta, A. West,** and **T. Howard**, *Design Issues for Virtual Reality Systems* Presented at the First Eurographics Workshop on Virtual Environments, 7 September 1993.

[Ingr95] **R. Ingram** and **S. Benford**, *Legibility Enhancement for Information Visualisation*, Proceedings of IEEE Visualization '95, October 30 - November 3 1995.

[Kenn96] **J. B. Kennedy, K. J. Mitchell,** and **P. J. Barclay**, *A Framework for Information Visualisation*, ACM SIGMOD Record, special issue on Information Visualisation, Vol. 24, No. 4, 1996.

[Knig99a] **C. Knight** and **M.Munro**, *Comprehension with[in] Virtual Environment Visualisations*, Proceedings of the IEEE 7th International Workshop on Program Comprehension, pp4-11, May 5-7, 1999.

[Knig99b] **C. Knight** and **M. Munro**, *Visualising Software – A Key Research Area*, Proceedings of the IEEE International Conference on Software Maintenance, Oxford, England, August 30 – September 3, 1999. (Short paper.)

[Knig00a] **C. Knight**, *Virtual Software in Reality*, PhD Thesis, Department of Computer Science, University of Durham, June 2000.

[Knig00b] **C. Knight** and **M. Munro**, *Should Users Inhabit Visualisations?*, Proceedings of IEEE WETICE 2000; Knowledge Media Networking Workshop, June 2000.

[Knig00c] **C. Knight** and **M. Munro**, *Virtual but Visible Software*, Proceedings of the IEEE International Conference on Information Visualisation 2000, July 2000.

[Levi95] **S. Levialdi, A. Massari,** and **L. Saladini**, *Visual Metaphors for Database Exploration*, A position paper relating the Virgilio system, available on-line from http://www.darmstadt.gmd.de/~hemmje/Activities/Virgilio/Publications/virgilio.ps.

[Moni94] **N. Monin** and **D. J. Monin**, *Personification of the Computer: A Pathological Metaphor in IS*, Proceedings of SIGCPR '94, ACM Press.

[Pesc93] **M. D. Pesce**, *Final Amputation: Pathogenic Ontology in Cyberspace*, Presented at the Third International Conference on Cyberspace, University of Texas, May 1993.

[Pett96] **S. R. Pettifer**, *Beyond the Metaphysics of Virtual Reality*, M. Sc. Thesis, Department of Computer Science, University of Manchester, UK. January 1996.

[Pett97a] **S. Pettifer** and **A. West**, *Deva: A coherent operating environment for large scale VR applications*, Presented at the first Virtual Reality Universe conference, April 1997.

[Pett97b] **S. Pettifer** and **A. West**, *Worlds within Worlds: on the topology of virtual space*, Presented at Cyberconf, Oslo (http://www.cyberconf.org), June 1997.

[Rhei92] **H. Rheingold**, *Virtual Reality*, Mandarin Science, 1992.

[Vään93] **K. Väänänen**, *Multimedia Environments: Supporting Authors and Users with Real-World Metaphors*, Proceedings of the Conference on Human Factors in Computing Systems (INTERACT '93), pp99-100, 1993.

[Vään94] **K. Väänänen** and **J. Schmidt**, *User Interfaces for Hypermedia: How to Find Good Metaphors?*, Proceedings of CHI '94 Short Papers (CHI Companion '94), pp263-264, April 24-28, 1994.

[Walk95] **G. Walker**, *Challenges in Information Visualisation*, British Telecommunications Engineering Journal, Vol. 14, pp17-25, April 1995.

[Wiss98] **U. Wiss** and **D. Carr**, *A Cognitive Classification Framework for 3-Dimensional Information Visualization*, Research report LTU-TR--1998/4--SE, Luleå University of Technology, 1998.

[Youn98] **P. Young** and **M. Munro**, *Visualising Software in Virtual Reality*, Proceedings of the IEEE 6th International Workshop on Program Comprehension, pp19-26, June 24-26, 1998.