The Knowledge Engineering Review, Vol. 15:4, 2000, 405–410
Printed in the United Kingdom

Advances in information visualisation: recent outcomes*

LAURA TARANTINO

Dip. di Ingegneria Elettrica, Università dell'Aquila, L'Aquila, Italy I-67040. laura@ing.univaq.it

Abstract

To an ever-growing degree, information technology is being based on interactive visual media, and visualisation techniques are emerging as the primary support to decision-making tasks. In this paper I briefly survey the most recent results in the field of advanced visual interfaces, focusing on visualisation techniques along with their design and their applications.

1 Introduction

> **Vi-su-al** *adj.* **1.** Of, serving, or pertaining to the sense of sight **2.** Capable of being seen by the eye, visible.
> **Vi-su-al-ize** *v.* To form a mental image or vision of.
>
> The American Heritage Dictionary

One of the distinctive features of our era is the dominance of the image, in almost every aspect of our life. Human existence is more visual and visualised than ever, to the extent that the term *visualisation* has changed its meaning: from being an internal construct of the mind, a visualisation has become an external artefact (Ware, 1999). The growing tendency to visualise things that are not in themselves visual is also favoured by the technological capacity to make visible those things that our eyes could not see unaided (e.g. satellite pictures, or medical images of the interior of the human body, which are in fact transpositions of frequencies our eyes cannot detect) (Mirzoeff, 1999).

The human ability to absorb and interpret visual information is the basis of industrial society and is becoming crucial in the information age: to an ever-growing degree, information technology (IT) is being based on interactive visual media, and visualisation techniques are emerging as the primary support to decision-making tasks. Allied to this trend is the growing computer capacity to process and render heavy graphics in real time. Visual user interfaces are no longer fancy, additional layers added on top of existing systems. They rather constitute a crucial aspect of modern applications, capable of determining the success on the market of software products.

In this paper I briefly survey some recent contributions coming from the visual interface field, presented at the fifth edition of the International Conference on Advanced Visual Interfaces, AVI 2000 (Di Gesu' *et al.* 2000). Information visualisation will be the starting point of our discussion, and certainly the topic does deserve a special interest, if only for the fact that over one-third of the papers presented at AVI 2000 focused on visualisation techniques. I will then survey the related design principles and guidelines, and we will see how all this translates into systems offering efficient solutions to application-specific requirements.

*Partially supported by the Italian Ministry of University and Scientific and Technological Research under the MURST 60% project on "Representation and Interaction Techniques for Spatial Data".

2 Visualisation to amplify cognition

I have noted how the notion of visualisation strayed from its coded meaning. A possible, modern, definition for *information visualisation* is "*the use of computer-supported, interactive, visual representation of (abstract) data to amplify cognition*" (Card *et al.*, 1999) (where cognition refers to acquisition and use of knowledge). Cognition enhancement is achieved when visual externalisations are designed so as to take advantage of the human perceptual processing system.

Explicit reference to this general guideline is found in the design of *3D geon diagrams* proposed by Irani and Ware (2000), based on structural object perception theory (Marr & Nishihara, 1978). According to this theory objects are perceived as being decomposed into 3D primitives called geons (short for geometrical ions: cylinders, cones, boxes), together with the skeleton structure connecting them (Biederman, 1987). In order to explore the effectiveness of geons in diagrams, the authors considered geon-based visualisations of the kind of information usually represented by the Universal Modeling Language (UML) (Booch *et al.*, 1998), a rich diagrammatic notation including heterogeneous nodes and arcs, with attributes attached to both nodes and arcs. The results of three experiments show that geon diagrams can be visually analysed more rapidly, with fewer errors, and can be remembered better in comparison with equivalent UML diagrams.

The third dimension in UML diagrams (along with interactivity and animation) is introduced also by Radfelder and Gogolla (2000). Here the aim is to use the foreground–background metaphor to arrange items in the space according to their current importance, or to place two different aspects of a system in the same view. It would be interesting to have experimental comparisons between 3D geon diagrams and 3D UML diagrams to get more insights about the influence on cognition enhancement of the diverse variables (e.g. shapes, dimensions, surface properties, interactivity).

Though 3D information visualisation is favoured by availability of inexpensive hardware and software toolkits, 2D visual representations still constitute a great part of the techniques used. Many visual languages are based on diagrams and graphs, and research on graph drawing can be, for example, of use for the visualisation of networks, provided that the techniques keep pace with the data volume growth. In this direction, a multi-scale layout algorithm for the aesthetic drawing of undirected graphs is proposed by Harel and Koren (2000). The algorithm is extremely fast, and capable of handling graphs of substantially larger size than other methods (examples with over 6000 nodes are shown in the paper). The basic idea is to consider small abstractions of the graph, in order to approximate coarse structure and achieve a global arrangement, while accomplishing local beautification by optimising small neighbourhoods. A different approach for coping with large graphs is given by Niggeman and Stein (2000), who define a meta-heuristic which decomposes a graph into clusters, and then applies to each subgraph optimal drawing methods. It has to be observed that the authors do not provide information about the size of manageable graphs.

Node-link diagrams can also be used to visualise hierarchically structured data (actually a special case of network), even if appropriate techniques (such as interaction, distortion and use of the third dimension) have to be adopted to overcome the problems arising in even moderately large trees (see, e.g., the *ConeTree* and the *Hyperbolic Browser*, in Card *et al.* (1999)). Another way to visualise trees is to use *enclosure*, which, unlike connection, fills the available space. A number of AVI 2000 papers discussing hierarchy visualisations are based on enclosure.

In Bjork (2000) the author presents the *Hierarchical Flip Zooming* focus + context technique, a generalisation of the basic flip zooming. Flip zooming requires the information to be split into pieces suitable for presentation in (sequentially ordered) tiles. The central area of the display is assigned to the focus tile, while the remaining tiles are arranged around the focus in smaller areas while preserving the sequential order; changing the focus modifies the overall arrangement of the tiles. In hierarchical flip zooming the technique is recursively applied within flip zooming visualisations to obtain an n-layer visualisation (in a box-in-a-box fashion (D'Atri *et al.*, 1991)), where individual visualisations can be interacted with independently.

Information visualisations based on enclosure are frequently used for information structures that have little or no information in non-leaf nodes. A technique straying from this general rule is discussed by Dix *et al.* (2000): unlike other hierarchy-oriented techniques aimed at visualising the hierarchies themselves, *pieTrees* is primarily aimed at visualising the numeric data associated with the hierarchy. The technique is inherently interactive, and allows the user to drill into data shown on a pie chart. Selecting an area makes it split into smaller segments while retaining a small pert-circle segment towards the center, representing the value associated with the interior node being exploded.

A more complex problem is faced by Graham *et al.* (2000), namely the support for visualising the correlations between multiple, overlapping sets of hierarchical information. A direct application is in biological taxonomies, where, over time, specimens may end up classified in different groupings in dissimilar classifications (made by different taxonomists), all regarded as valid viewpoints. The proposed technique stems from the observation that taxonomic reasoning (i.e. the categorisation of objects) is essentially set-based (Parunak, 1991), and that visual aid for such a task would benefit from mirroring a set-based metaphor on the screen. The display simultaneously shows several hierarchies while making others easily accessible. To reduce screen space problems, the leaf nodes of each bottom-level hierarchy are arranged in a grid formation. Interaction is based on a linking-and-brushing technique. The prototype was positively evaluated by taxonomists against a graph-based prototype spatially and temporally integrating many hierarchies.

It should be clear at this point that the effectiveness of a visualisation always depends at least as much on the dynamic aspects (the interaction) as on the static aspects (the visual structures). As argued by Dix *et al.* (2000), even simple interaction can enhance traditional paper-based representations. An example is offered by *dancing histograms*, the interactive version of paper-based stack diagrams, in which columns can be moved so as to change the baseline, thus allowing the user to compare categories different from the bottom one.

Dynamic aspects are fundamental also in *MultiNav*, proposed by Lanning *et al.* (2000) for multidimensional data exploration, with direct application in shopping on the Internet. MultiNav is designed to be a visualisation tool that can be linked with several data sources (XML documents, CSV files, relational databases). It comes into play when a set of items of appropriate number (not larger than 500), having a common attribute schema, has been specified. Each dimension is mapped on a row of bars, and each item has a "physical" position in each attribute row. Clicking on one of the bars corresponds to selecting an item (the focus) and highlights all its associated bars. The user can get a quick sense of where the item falls in the space as a whole (the context) by seeing its position in each attribute sort. At the same time, item details are displayed in the lower portion of the screen. Simple restriction queries are possible by acting on labels representing attribute values. Despite its intrinsic limitation (the number of bars in each row is bounded by the screen size), the technique seems very effective. A different view on the information space allows rows to slide back and forth. *Sliding rods* are aligned such that the focused item is always in the centre position and directions of sliding favour attribute-correlation analysis.

MultiNav and other focus + context techniques share the assumption that, among other things, cognition amplification can be achieved by increasing the amount of information placed for the user's attention. Analysis of this assumption and a study on focus + context impact on cognition and attention was carried out by Pirolli *et al.* (2000). The experiment aimed at understanding how use of the Hyperbolic Browser differs from the use of a more conventional browser (Microsoft Explorer). The two browsers were compared with respect to simple and complex retrieval tasks (aimed at finding leaf nodes in the tree), and local and complex relational tasks (involving examination of several nodes close together or in disparate parts of the tree). The Hyperbolic Tree showed many superior aspects, such as better performance times and better learning of the tree. However, results show that it is greatly affected by information scent: strong information scent expands the spotlight of attention, while crowding of targets in the compressed areas narrows it. The results suggest design improvements.

3 Steering the visual interface designer

The increasing complexity of the adopted visualisation and interaction techniques makes the job of the interface designer increasingly difficult. Codification, principles and guidelines are necessary to steer design decisions and evaluation tasks.

Visual design attempts to solve communication problems in a way that is at once functionally effective and aesthetically pleasing (Mullet & Sano 1995). And both *aesthetic* and *functionality* are actually stressed by papers presented in AVI 2000.

The designer of hierarchy visualisations based on enclosure may profit from Harel and Yashchin's work (2000). Their objective is to arrange an inclusion hierarchy of non-overlapping *blobs* (rectangles representing nodes) in a planar layout constituting an aesthetic representation that achieves the "best" positioning of the involved blobs. The paper provides general aesthetic local criteria (e.g. blob dimension uniformity between siblings, and blob proportions) and global criteria (e.g. blob dimension uniformity between leaves), along with a (customisable) layout algorithm. Blobs can be moved, resized and rearranged in any possible way, provided that the hierarchy structure is preserved. We have to observe that specific applications may pose more strict constraints on blob arrangement, thus forcing the designer to search for compromises among the proposed aesthetic criteria.

The emphasis is instead placed on functionality in Beaudouin-Lafon and Mackay's paper (2000). Starting from the observation that today's user interfaces often suffer from overabundance of functionality (making users sometimes reject software updates), they propose three design principles based on a trade-off between power and ease of use: *reification* turns concepts into objects (e.g. commands are reified into interaction instruments); *polymorphism* permits commands to be applied to objects of different types (essential when one wants to keep the interface simple while increasing its power); *reuse* makes both user input and system output accessible for later reuse. The power of these principles lies in their combination: reification creates new objects that can be acted upon by a small set of polymorphic commands, creating more opportunities for reuse. The authors then report on how these principles have provided a framework for redesigning a complex tool for editing and simulating coloured petri nets. The design has been influenced by two key decisions: first, to use a combination of new interaction techniques (marking menus, floating palettes, and toolglasses) rather than using a standard WIMP (Windows, Icons, Menus, Pointing) interface and, second, to support two-handed input with a mouse for the dominant hand, and a track-ball for the non-dominant hand (as in previous work by the authors, design is inspired by the way people interact with the physical world). The discussion is extremely interesting and anyone working in visual interface design should take a careful look at it.

No less important is the discussion conducted by Dix *et al.* (2000) with regard to the role of software architecture in effective visualisation construction and design. Architectural design improves the performance of the software, favours the splitting of the software into segments possibly running on different parts of a distributed system and, maybe more importantly, makes the software easier to construct and maintain thus encouraging designers to use more iterative design techniques and to try multiple design alternatives. As a benchmark, the authors discuss the functionality and low-level architecture of *onCue*, a context-sensitive desktop agent integrating lightweight visualisations with desktop and Internet data sources (pieTrees and dancing histograms, previously discussed, are parts of onCue).

Architectural design is a key aspect also in *Snap-Together Visualization*, a system proposed by North and Shneiderman (2000) to enable data users to rapidly and dynamically mix and match visualisations and coordinations to construct custom exploration interfaces without programming. Snap's conceptual model is based on the relational model of data: users load relations into visualisations and coordinate them based on relational joins; brushing, drill down, overview, detail view and synchronised scrolling are the supported coordination types.

Related to the Snap-Together goals are the problems faced by Wang Baldonado *et al.* (2000), who provide eight guidelines for the design of multiple view systems (two or more distinct views support the investigation of a single conceptual entity). A first group of rules (diversity, complementarity, parsimony and decomposition) are aimed at helping designers and users assess whether or not multiple views are appropriate for their applications. When multiple views are appropriate, a second group of rules (space/time resource optimisation, self-evidence, consistency and attention management) help designers make decisions about presentation and interaction, as well as help usability experts and system usability evaluators pinpoint trouble spots in an existing system.

Actually, the importance of adopting usability methods during the development cycle to increase the acceptability of new products on the market is now recognised. Usability inspection methods are emerging as preferred evaluation methods: they provide results in a limited amount of time and are less costly than traditional lab-based usability evaluation (no special equipment or lab-facilities are required) (Madsen, 1999). However, usability inspections are still highly subjective, and strongly dependent upon the skill, experience and ability of the inspectors.

To overcome these disadvantages, De Angeli *et al.* (2000) have proposed the *SUE* (systematic usability evaluation) inspection technique, oriented towards hypermedia applications. SUE is based on an evaluation pattern called Abstract Tasks, which precisely describes the activities to be performed by evaluators during the inspection. The paper present an empirical validation of the technique carried on as a comparison with a traditional heuristic technique for the evaluation of a commercial CD-ROM. Results show a clear advantage of the SUE technique with respect to effectiveness, efficiency and satisfaction.

4 Putting visualisations into use

In IT applications, visualisation may be used to increase communication and/or processing effectiveness in conventional domains, or to make the technology itself more easily utilisable for the final users. As examples of possible applications, we discuss here a few AVI 2000 papers related to the travel industry and computer-mediated conversation.

Visualisations of travel itineraries are discussed by Apperley *et al.* (2000), as components of a larger system supporting collaboration between a travel consultant and a client in the planning and implementation of an itinerary. The chosen interactive graphic representation, resulting from extensive experiments on paper prototypes, is really very informative, and allows one to quickly gather information about involved locations, time and time zones (through effective use of vertical and horizontal axes, respectively), daytime and night-time (through shadings) and events.

Kruger *et al.* (2000) discuss graphical presentations of way directions to pedestrians (e.g. airport passengers or tourists). The graphics process adapts to the available resources (user's time, computing power, media, output devices) to generate a wide variety of presentations from a single 3D object and knowledge base. Presentations (for handheld computers and information booths) range from simple 2D sketches (possibly enriched with 3D thumbnails when appropriate) to animated 3D walkthroughs of the surroundings presented from an egocentric perspective.

In the field of asynchronous text-based communication, Yabe *et al.* (2000) propose a technique for the animated visualisation of USENET discussions: a scenario from articles in a thread is generated using the quote relationship and translated into an animation casting as 3D avatars the authors of the articles. A 2D approach is instead proposed by Popolov *et al.* (2000) in *ConverSpace*, a spatial representation of computer-mediated conversation, which strays from the usual linear presentation in the attempt to combine the spontaneity of oral communication and the persistence of textual discourse.

5 Conclusions

In this paper I have presented some recent achievements in the visual interface field presented at the AVI 2000 conference. The conference programme included three invited papers, 24 research papers, 11 system papers and 14 posters. Space limitations forced me to limit the discussion to only part of the AVI 2000 collection of papers (covering only some of the topics of interest). The discussed techniques share the underlying assumption that computer-mediated communication should rely as much as possible on sensory languages using the perceptual processing power of the brain without learning. Sensory and perceptive interfaces (not limited to visuality) are actually the future in this field, and will be the main focus of the 2002 edition of the conference. In the meantime, information about purchase of the AVI 2000 Proceedings can be found at the AVI website (`http://www.ing.univaq.it/avi2000/`).

References

The American Heritage Dictionary, 1989, Second College Edition, Dell Publishing.

Apperley, M, Fletcher, D, Rogers, B and Thomson, K, 2000 "Interactive visualization of a travel itinerary", in Di Gesu' *et al.*, 2000, 221–226.

Baus, J, Butz, A and Kruger, A, 2000, "Smart graphics in adaptive way descriptions", in Di Gesu' *et al.*, 2000, 92–97.

Beaudouin-Lafon, M and Mackay, WE, 2000, "Reification, polymorphism and reuse: three principles for designing visual interfaces", in Di Gesu' *et al.*, 2000, 102–109.

Biederman, I, 1987, "Recognition-by-components: a theory of human image understanding" *Psychological Review* 94(2) 115–147.

Bjork, S, 2000, "Hierarchical flip zooming: enabling parallel exploration of hierarchical visualizations", in Di Gesu' *et al.*, 2000, 232–237.

Booch, G, Rumbaugh, J and Jacobson, I, 1998, *The Unified Modeling Language User Guide* Addison-Wesley.

Card, SK, Mackinlay, JD and Shneiderman, B (eds.), 1999, *Readings in Information Visualization: Using Vision to Think* Morgan Kaufmann.

D'Atri, A, Laenens, E, Paoluzzi, A, Snijders, J Tarantino, L, 1991, "A graphical browser to object-oriented knowledge bases" *Database Technology* **4**(1) 45–55.

De Angeli, A, Matera, M, Costabile, MF, Garzotto, F and Paolini, P, 2000, "Validating the SUE inspection technique", in Di Gesu' *et al.*, 2000, 128–135.

Di Gesu', V, Levialdi, S and Tarantino, L (eds.), 2000, *Proceedings of the International Conference on Advanced User Interfaces AVI 2000*, May 24 – 26 ACM Press.

Dix, A, Beale, R and Wood, A, 2000, "Architectures to make simple visualizations using simple systems", in Di Gesu' *et al.*, 2000, 51–60.

Graham, M, Kennedy, JB and Hand, C, 2000, "A comparison of set-based and graph-based visualizations of overlapping classification hierarchies", in Di Gesu' *et al.*, 2000, 41–50.

Harel, D and Koren, Y, 2000, "A fast multi-scale method for drawing graphs", in Di Gesu' *et al.*, 2000, 282–285.

Harel, D and Yashchin, G, 2000, "An algorithm for blob hierarchy layout", in Di Gesu' *et al.*, 2000, 29–40.

Irani, P and Ware, C, 2000, "Diagrams based on structural object perception" in Di Gesu' *et al.*, 2000, 61–67.

Lanning, T, Wittenburg, K, Heinrichs, M, Fyock, C and Li, G, 2000, "Multidimensional information visualization through sliding rods", in Di Gesu' *et al.*, 2000, 173–180.

Madsen, KH (ed.), 1999, special issue on "The diversity of usability practices", *Communication of the ACM* **42**(5).

Marr, D and Nishihara, HK, 1978, "Representation and recognition of the spatial organization of three-dimensional shapes" *Proceedings of the Royal Society of London*, B, 200 269–294.

Mirzoeff, N, 1999 *An Introduction to Visual Culture* Routledge.

Mullet, K and Sano, D, 1995, *Designing Visual Interfaces* Prentice Hall.

Niggeman, O and Stein, B, 2000, "A meta-heuristic for graph drawing", in Di Gesu' *et al.*, 2000, 286–289.

North, C and Shneiderman, B, 2000, "Snap-together visualitazion: a user interface for coordinating visualizations via relational schemata", in Di Gesu' *et al.*, 2000, 128–135.

Parunak, H. Van Dyke, 1991, "Don't link me in: set-based hypermedia for taxonomic reasoning" *Proceedings of ACM Hypertext'91* 233–242

Pirolli, P, Card, SK and Van Der Wege, MM, 2000, "The effect of information scent on searching information visualizations of large tree structures", in Di Gesu' *et al.*, 2000, 161–172.

Popolov, D, Callaghan, M and Luker, P, 2000, "Conversation space: visualising multi-threaded conversations", in Di Gesu' *et al.*, 2000, 246–249.

Radfelder, O and Gogolla, 2000, "On better underastanding of UML diagrams through interactive three-dimensional visualization and animation", in Di Gesu' *et al.*, 2000, 292–295.

Wang Baldonado, MQ, Woodruff, A and Kuchinsky, A, 2000, "Guidelines for using multiple views in information visualization", in Di Gesu' *et al.*, 2000, 110–119.

Ware, C, 1999, *Information Visualization: Perception for Design* Morgan Kaufmann.

Yabe, J, Shibayama, E and Takahashi, S, 2000, "Automatic animation of discussions in USENET", in Di Gesu' *et al.*, 2000, 84–91.