The Challenge of Visualising Multiple Overlapping Classification Hierarchies

Martin Graham, Jessie B Kennedy and Chris Hand
School of Computing, Napier University, 42 Craiglockhart Ave, Edinburgh EH14 1LT
{marting, jessie, cph} @dcs.napier.ac.uk

Abstract

Techniques for visualising hierarchies have concentrated on displaying static structures or, in the case of dynamic hierarchies, adding or deleting nodes from the hierarchy. However none have adequately dealt with the situation of visualising change occurring in the structure between existing nodes. We present initial work that outlines the difficulties of visualising such an event, with specific regard to the multiple classifications of information that are produced in botanical taxonomy.

1 Introduction

The field of Information Visualisation (IV) concerns itself with taking complex information structures and making them easily understandable and useable to the user by choosing appropriate ways to address perception, representation and interaction with visually displayed structures.

One area to which visualisation is commonly applied is hierarchical information structures, such as file systems or organisational structure charts. However, much of this research has focused on single instances of such structures, whereas a general problem occurs in understanding how they change over time, often a process of organic, *ad hoc* growing and shrinking.

This paper addresses the problem of supporting users in working with *multiple* information hierarchies that represent different classifications of a number of sets of objects. In particular, this problem occurs in the field of taxonomy, in which the repeated classification and re-classification of objects such as plants and animals results in large, complex databases of historical information which need to be understood and manipulated by biologists.

The paper is broken down into four main sections, firstly describing present IV techniques that are categorised using an existing classification. This is followed by a description of the general visualisation problem we are concerned with plus a description of the specific area where such a visualisation could be of benefit. We then discuss the visualisation techniques that address situations that have the greatest similarity to our own problem, and explain why they still lack suitability for our purposes. Finally, some possible characteristics of a suitable visualisation for multiple classification hierarchies are suggested, and the paper finishes with a short conclusion.

2 Information Visualisation

Gershon, Card et al [1] classify information visualisation (IV) research into sub-topics such as human visual perception, display techniques, and interaction techniques. A review of visualisation research is presented using this classification framework. Although useful for reviewing IV, the boundaries between the sub-groupings are fuzzy and some work has relevance to more than one of these categories.

2.1 Human Visual Perception

A number of pieces of research have addressed the strength of human perception of visual information with experimental evidence. Generally, they deal with issues such as motion as an aid for recognising structure [2; 3] and motion used for encoding further dimensions of the information set [4]. Other work shows the importance of colour in differentiating groups of objects or information [5; 6], and the conflicting effects of several perceptual visual cues operating simultaneously, such as colour, shape, and motion [7; 8]. The experimental findings provide a proven basis for developing IV systems that co-operate with our visual capabilities. At the very least, they can be used to ensure the wrong visualisations aren't communicated to the user.

2.2 Display Techniques

Information that is to be visualised may be abstract, but it often has a structure that categorises it, such as a network or a hierarchy. Research in this area has concentrated on tackling the conflicting issues of size, layout, and legibility on limited screen area. Other information sets have less obvious structures, if any, and require display techniques that accent the important dimensions whilst still retaining an overview of the others.

2.2.1 Hierarchies.

The original IV hierarchy visualiser was Cone Trees [9]. Hierarchical information, more generally known as a tree structure, is displayed in three dimensions in an attempt to increase the number of nodes that could be presented on-screen. Selecting any node would bring that node to the front of the view of the Cone Tree in a smooth animated sequence. This use of animation preserves the users' model of the visualisation as a change is taking place. It's superiority to the alternative; an abrupt move to the final position without any intermediate views; has been shown [3].

Another hierarchical visualisation is *"Information Pyramids"* [10]. A tree, in this specific case a file directory, is viewed firstly as a flat plane. Each sub-directory is then viewed as a raised block on the plane, it's area proportional to the combined size of the files in the sub-tree. This process continues recursively down the file tree, with further sub-directories placed on the blocks their parent directories formed, until eventually the leaves are placed on top. Its advantages are that the important parts, the files, are always on top and visible, as they form the leaves of the tree. Also enough of the underlying directories are still visible to obtain a view of the whole structure.

Other IV hierarchy visualisations include Johnson and Shneiderman's Treemaps [11] for efficient use of screen space, Rapley and Kennedy's WINONA [12] for object-oriented database visualisation, and Herman's visual tree path navigation [13].

2.2.2 Node and Link structures (Graphs & Networks)

The second structure that accommodates a large variety of abstract information sets is the graph, or network. Visual representations of graphs and networks suit high-dimensional, discrete information, such as document collections, where nodes can represent individual items of information, and links represent the relations or dimensional correlation between them. Hyperspace [14] visualises the hypermedia structure of the World Wide Web as a graph, with individual pages forming nodes, and links between pages forming links in the visual representation. Hyperspace visualises the graph by letting the user select an area or keyword of interest, whereupon related pages concerning this topic move closer together, and dissimilar pages repel. In the end a graph with clusters of related pages is formed and displayed.

This type of self-organising structure, based on the spring-mass metaphor, occurs in several systems such as Hyperspace's successor Narcissus [15], and others [16; 17; 18; 19]. The effect of this clustering is analogous to the concept of chunking in drop-down menus. Similar items are grouped together, and the user recognises them as sharing common attributes due to their visual proximity.

2.2.3 Other Display Techniques

One of the main difficulties with IV is the mapping of many abstract dimensions to the few spatial dimensions displayable on a computer monitor. 3D projections can help, and VR techniques can increase the depth perception of 3D visualisations, but even this only gives us one extra spatial dimension. It also introduces its own set of problems such as occlusion and effective depth cueing [2; 20; 21]. The general answer appears to be mapping of dimensions to non-spatial cues such as colour, brightness, transparency, and shape [22]. If these are not enough, another approach consists of nesting co-ordinate systems within the points of other co-ordinate systems, and thus viewing only a subset of the actual dimensions present in the information [23; 24; 25].

A number of IV approaches use Benedikt's ideas of extrinsic dimensions (position and orientation of object) and intrinsic dimensions (properties of object), to model large-dimensional information sets [26; 27]. However, an initial problem is deciding which dimensions should take the seemingly more important extrinsic representations. This can only be decided by analysing the user's task and deciding what aspects of the information they are most likely to be searching for. Even the intrinsic qualities have a certain precedence, established by the experiments on human visual perception mentioned earlier. A simple ordering of these, in descending order of perceptual effectiveness, are motion, colour intensity, colour hue, and lastly shape [7].

A further problem with abstract dimensions is that they are not easily given to fitting on the numeric scales that Benediktine dimensions use, extrinsic and intrinsic. The fallback position is to use one of the arbitrary scales that Benedikt proposes, which are alphabetical, geographical, and chronological. Dimensions should be able to map to at least one of these orderings, though doing so might appear to make no difference to comprehending it.

2.3 Interaction Techniques

There are a range of techniques that have been developed to allow IV users to search information sets, and discover specific pieces of information. These are general techniques that can be applied to the display techniques previously described, and are grouped by Gershon into categories such as focusing, filtering, and linking interaction techniques.

2.3.1 Focusing Techniques

Focusing techniques are mainly concerned with the distortion of graphical displays, to give greater prominence to a certain area of the visualisation space, and hence greater prominence to whatever is displayed there. They are used to allow navigation of large structures, but at the same time allow close detailed inspection of specific pieces of information. This distortion could be in terms of an actual focal point, utilising the numerous lens-style viewers, or a uniform increase in magnification, using zooming, or perhaps other qualities such as colour brightness, or a combination of many.

Lens viewers [28], increase the size of the information at the focal point of a screen, to the detriment of information which is visualised further away, and so reduced in scale, but still visible. Sarkar and Brown's paper is a specific example of a general lens distortion technique [29]. Munzner [30], and Lamping and Rao [31] apply hyperbolic lenses to networks and trees respectively, whilst Carpendale [32] pushes the lens metaphor to three dimensions.

Closely linked to lenses are zoom methods; the difference being that in a zoom, the entire screen is always at the same magnification. The effect is still to focus on a particular piece of information, but to lose a lot of the periphery information. Schaffer [33] describes a way of combining zoom and lens distortions.

2.3.2 Filtering Techniques

Filtering techniques are used when the user wishes to home in on information that has common attributes or values. The query mechanisms underlying the filtering are outside the scope of this paper; the filter referred to is a visual filter on the screen, where a set of conditions initiated by the user affects the visualisation in some way. The usual effect is to highlight information that matches the desired conditions, or removal from the visualisation of information that doesn't correspond.

Fishkin [34], Eick [35], and Ahlberg [36] describe the user setting filter conditions on unstructured information (i.e. not networks or trees.) The use of the filters results in the removal of the visual clutter i.e. the unwanted information, and hence the visual promotion of the information the user is interested in. Colby and Scholl [37] showed that a similar effect of visual promotion could be achieved by using transparency and blur effects on information that did not match user requirements.

Kumar's [38] filtering of visualisations operates on a structured information set, namely a tree. The user sets filter conditions and the effect is to filter out sub-trees instead of unstructured groups of information. What remains are the parts of the tree, and thus the leaves and paths to the leaves, that the user wishes to find.

Therefore these issues should be considered when devising useful visualisations. One of the challenges of visualisation is to find the best way of representing the data. Each type of data has its own particular characteristics, so that although one technique may work well for one type of data it may not be applicable to another. A good visualisation will address the perceptual issues of IV, display techniques, and interaction techniques as described above.

2.4 The problem of visualising structural change in hierarchies

To summarise so far, the visualisation of hierarchical structures has been a focus of research within the field of information visualisation since Xerox PARC's work on Cone Trees. As discussed, research has since continued along the lines of increasing the information density of visualised hierarchies, focusing and filtering techniques, and improved 2D and 3D layout algorithms. However, to our knowledge, there are currently no visualisation methods that allow for the tracking of *structural change* within hierarchies, a phenomenon that frequently occurs with the reclassification of an existing hierarchy. Alternatively, the situation could be described as tracking structural differences across a set of hierarchies holding the same node information.

One discipline where such information sets occur and the ability to track information from one classification to another is required is taxonomy, the study of scientific classification. A more detailed explanation of the field and its particular problems follow.

3. Multiple Classification Hierarchies in Taxonomy

Taxonomists study and then classify organisms to generate a classification hierarchy depicting their presumed natural relationships. These classifications are hierarchical structures where specimens are grouped into *taxa* (singular: taxon) which are then placed in higher level taxa according to some criteria (e.g. DNA relationships, morphological similarities). Taxa are assigned to ranks that specify the level of a taxon in a classification hierarchy. The levels (or ranks) used in generating the classification hierarchies vary for different groups of specimens and between taxonomists.

Over time some specimens may end up classified in different ways. These classifications are all valid, even though more recently revised versions exist. Taxonomists do not have the concept of ‘correct classification’: they regard all published classifications as valid viewpoints.

When taxonomists choose a group to be studied, they collect preserved and living specimens on which to base their work. At the same time, they compile information about past classifications of this group from the literature. The taxonomist examines the specimens and decides on the criteria to differentiate them. Using these criteria the specimens are classified into different groups and attributed to taxa. Finally a name is assigned to each taxon defined in the classification. This involves application of the nomenclatural code, which in simple cases involves finding the name of the oldest published specimen of its type in each group. If no type specimen exists in a newly defined taxon, a new name is created.

The classification is published for other taxonomists to use and is now considered a valid classification. If other taxonomists disagree with this classification then they must undertake a revision of the group and publish their conflicting viewpoint.

A challenge generated by the way taxonomists work is the management of the accumulation of old historical classifications. Indeed, when a classification is revised, it stays valid (e.g. because of references to it in the literature) even if it is not the classification that is recognised by the majority of taxonomists.

A second challenge is that the choice of criteria and the way a classification is created (e.g. a revision of previous work or a new study) is largely free. Even the nomenclatural code has varied over time and hence will affect the naming of taxa. Thus it is likely that two taxonomists working on the same set of data will not produce the same classification. The same specimens may be seen differently by different taxonomists and may be classified under many different taxonomic groups.

Prometheus (EPSRC/BBSRC ref. BIO10516), a collaborative project between Napier University and the Royal Botanic Garden Edinburgh, is developing a database to support taxonomic working practice. Full details of taxonomic working practice can be found in Pullan et al [39], and a description of the Prometheus database to support multiple classifications in Raguenaud et al [40]. The visualisation issues described here have arisen out of the work with the taxonomists at RBGE.

3.1 An Example of How Multiple Classifications Evolve

Figure 1 depicts a simplification of the kind of scenario found in taxonomy. The information available grows over time, the criteria used for classification vary and the number of levels (ranks) used in the classification process varies. The grey shapes at the leaf nodes represent individual specimens to be classified.

The top figure is the earliest classification based on a smallish set of specimens. This classification was based on the shape and resulted in a two-level hierarchy. Square specimens are *typified* by the mid-grey square, triangular specimens by the dark equilateral triangle and circular specimens by the light-grey circle (i.e. these specimens are chosen as representative types of the taxons). Shapes in general are typified by squares and hence are represented by the mid-grey square.

Subsequently a second taxonomist decides that an intermediate level in the classification would make things clearer and introduces the general type square, triangle and circle and 2 sub-types of triangle, equilateral and right angle and two sub-type of round shape, circles and ovals. Due to the naming conventions, squares are still typified by the same mid-grey square, triangles by the dark equilateral triangle, and circular shapes by the light-grey circle. However new types are required for right-angled triangles (the black one) and ovals.

A third taxonomist then decides that shape is not an important characteristic and reclassifies the previous specimens along with some newly found ones, according to their brightness. This creates a two level classification with five groups (he ignores one particular shade as there is only one instance of it). Co-incidentally each group contains an existing type specimen and therefore no new types are required to be defined from the classification. In practice often several types will end up in one group, requiring the oldest type specimen to be chosen.

Finally a fourth taxonomist comes along, and reclassifies the specimens by shape again.

The reality in taxonomy is much more complicated and involves many more specimens. However, the general principle and reason for the existence of multiple classifications should be clear.

The lack of tools that handle multiple contradictory classifications limits the compilation and comparison of useful global data. In essence taxonomists have a need to represent overlapping multiple classifications to allow them to compare and contrast the classifications produced by different taxonomists or to try what-if scenarios on a classification. A visualisation that supports the work of taxonomists must allow them to explore the similarities and differences between the classifications.

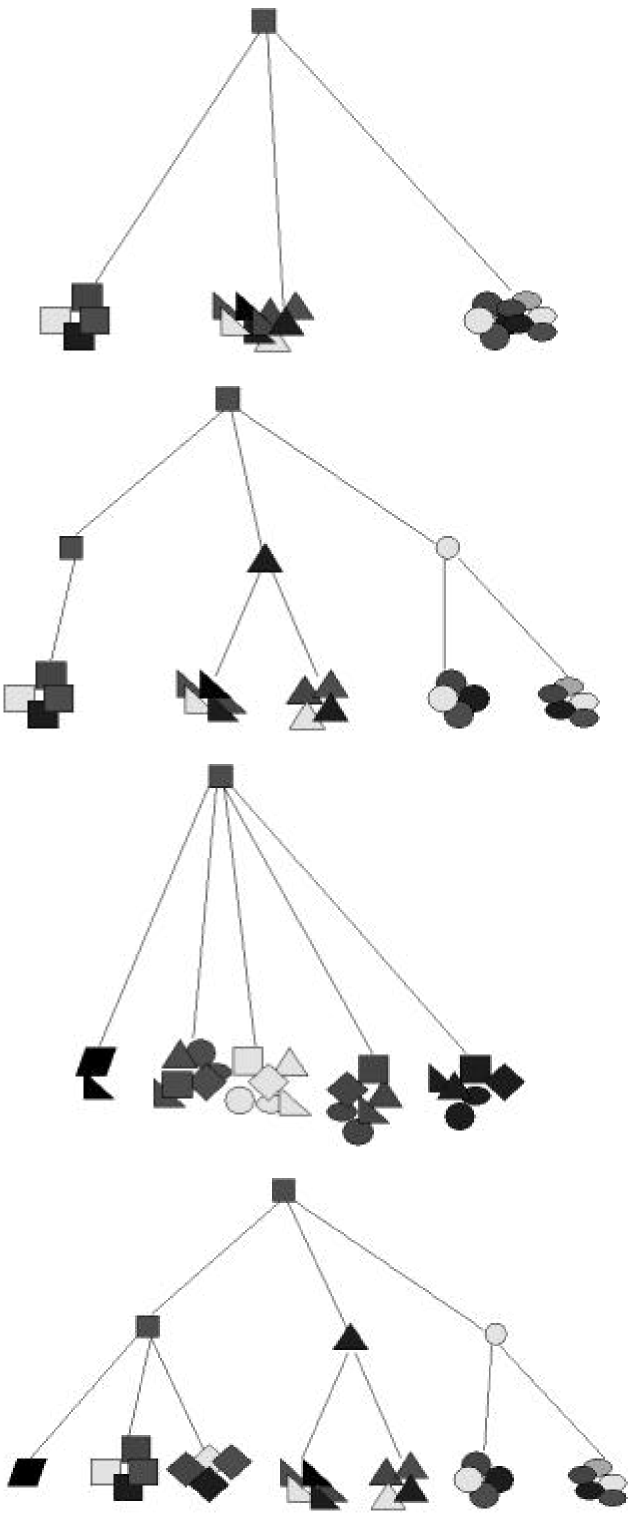

Figure 1 Four classifications with overlapping specimens and concepts

4 Limitations of Existing Visualisation Techniques

From the taxonomic problems described it is important to provide a visualisation that allows taxonomists to compare and contrast classification hierarchies. Previous visualisations that have aimed at showing change in structures have settled on one of two approaches. Firstly, visualisations have animated the information regarding the structure, relying on the animation's inherent perceptual qualities to enable users to comprehend the change, as is pointed out by Bartram [4]. Secondly, some visualisations have laid out a series of snapshots of the structure at critical phases, allowing the user to visually compare all or most of the relevant information at the same time. This type of visualisation is best known as Tufte's 'small multiples' [41, Ch. 2]. These visualisations are evaluated below as possible solutions to the multiple classification problem.

A number of visualisation techniques, including Huang and Eades' visualisation of huge graphs [42], and Wittenburg and Sigman's Treeviewer [43], use animation as the cue to show change in the structure of a hierarchy. In Huang's technique, a large graph is visualised as a hierarchy by omitting certain links in the display. Shrinking and expanding of sub-trees within this hierarchy cause the animated change in the visualisation, rather than the display of a succession of differing structures as we require. Wittenburg's Treeviewer is a mainly textual visualisation of web-search queries, and again animates addition/deletion of nodes, rather than the reclassification of the hierarchy's existing nodes.

Animation itself has two intrinsic drawbacks for the type of information we would like to visualise. Firstly, animation allows only direct visual comparison between two states, the last and the next stage of the animation. Other comparisons between states resulting from the animation must be recalled from memory. Secondly, whilst animation works for visualising gradual changes, such as Chi's web ecology evolution, it would become overly complex for drastic structural changes caused by reclassification of existing nodes. Wittenburg acknowledges this point by stating that their system could employ a fade-in/fade-out approach between two states or structures where there is a poor degree of correlation, instead of utilising the 'tweening' style of animation. Animation, however, may have advantages in its perceptual qualities for attracting the user's attention.

Two visualisations allow comparison using Tufte's idea of small multiples, namely Chi's web ecology visualisation [44], and Turo and Johnson's application of Treemaps [45]. The closest problem we have found to our own is discussed by Chi et al and involves displaying the evolution of a website over a number of months. In this visualisation, certain points during the period are visualised by displaying the sites' hierarchy in the form of a compact 'disk tree', with successive disk trees displayed next to each other, enabling visual comparison. The set of multiple disk trees is termed a 'time tube'.

The drawbacks of this technique, when considering our requirements, are that the visualisation is designed to highlight evolution of a hierarchy resulting from addition or deletion of nodes, rather than a *restructuring of the hierarchy* resulting from a reclassification of the nodes from which it is formed. Consequently, the main visual prominence for the change of structure is given to the addition or deletion of nodes. Nodes that already exist but have changed their links, and hence altered their position within the hierarchy, are not differentiated from nodes that have remained static within the hierarchies' organisation. There is a provision for highlighting a particular node's progress through a time tube, but due to the manner in which the disk trees are drawn, all nodes are displayed at the same position within each tree. Therefore seeing 'movement' due to restructuring is not aided.

Another visualisation prompt that is not present here, nor in the other techniques, is the ability to see the *context* in which a particular node has been restructured. By context, we are referring to the other nodes with which it has relations in the hierarchies, namely its parent, sibling, and child nodes as appropriate. This is needed, as nodes do not actually move when conceptual structures are re-organised (reclassified). Rather, they are *grouped* differently, possibly with different nodes. This principle also applies to sub-trees of various sizes, as well as individual nodes. Knowledge of the other nodes or sub-trees with which it shares relations in each version of a hierarchy will hint at the methodology behind that particular classification.

Turo and Johnson's visualisation technique, based on Johnson and Shneiderman's Treemaps [11], also includes an option to visualise change in trees or sub-trees over time, again using the small multiple approach. However, the changes they are concerned with are to do with information attached to individual nodes, rather than changes in the structure of the hierarchy.

The main disadvantage of the small multiples approach is a simple lack of space on-screen due to the resolution and size of the average monitor. To visualise a larger set of hierarchies will require smaller, more compact visualisations that are still intelligible to the viewer. Also, it lacks animation's pre-attentive visual cues, hence placing more cognitive load on the user when using the visualisation.

To summarise, these visualisations lack the ability to track one or more nodes and their contexts across multiple versions of a dissimilarly structured but similarly populated hierarchy. Our problem is directly related to this need, and hence our visualisation will require the abilities to show these relationships.

5. Visualisation of Multiple Overlapping Hierarchies

From our reflections on the visualisations mentioned above and our problem requirements, we have sketched a number of approaches for evaluation as a suitable visualisation. Essentially it is a choice between showing change over time or the change over space. (Chi et al [44] express this more formally.)

As mentioned in Section 2, IV can be approached from a number of angles, namely human visual perception, display techniques for the information in question, and interaction techniques.

Display techniques for hierarchies concentrate on increasing the number of nodes on screen and general layout. These will also be considerations for our visualisation. However, as Chi states, a number of trees will require an extra dimension either in space or time to be visualised. If we approach the notion of using an extra dimension in space to show the structural correlation of a number of trees, we are effectively visualising a 3D network or a graph, and some of the issues of graph visualisation such as occlusion and line crossing enter our problem domain. The problem of how to visualise the structure of a general graph with regard to placement of nodes is not such an issue, as the overall network is formed from hierarchies, giving a coherent conceptual structure enforced by the individual hierarchies.

Interaction techniques at our stage of development are difficult to visualise, and will need interactive prototypes for their demonstration. The principle approaches of interaction in IV, filtering and focusing, are methods of increasing the visual impact of information the user has expressed an interest in. We mentioned that existing visualisations did not provide a basis for distinguishing changing structure from static structure. One possibility would be to incorporate some focusing and filtering principles in the visualisation under user and/or automatic control to help in this undertaking.

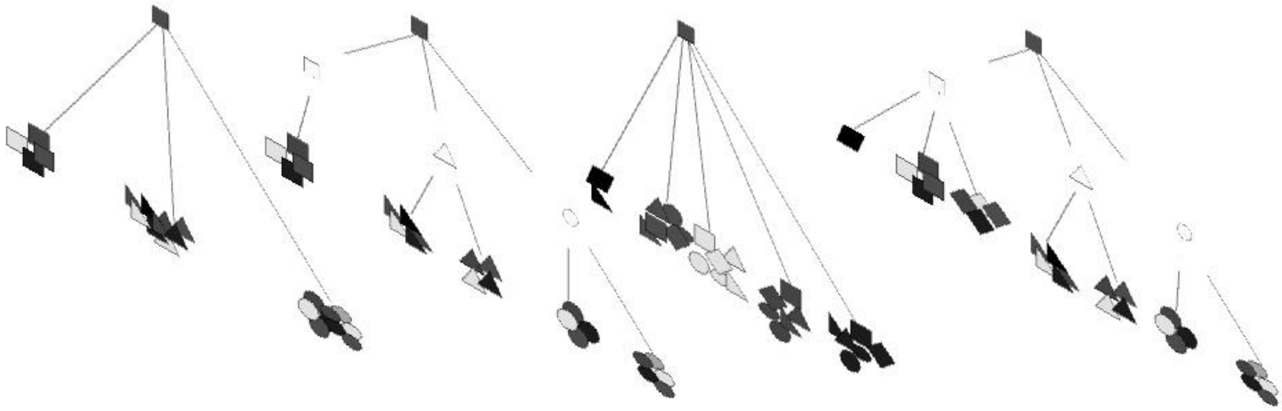

Figure 2 Filtering of intermediate levels in hierarchy

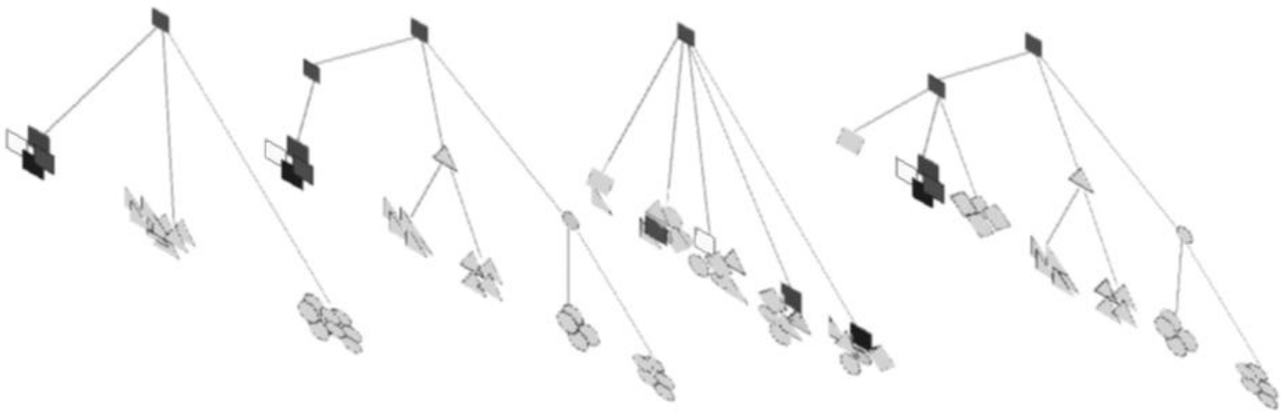

Figure 3 Tracking of a sub-tree through hierarchies

For example, nodes and sub-trees that did not change across multiple versions of the hierarchies could be lessened in visual prominence in a similar manner to Kumar's [38] pruning of hierarchies, or via a clustering method. By comparison, the information that does differ across the hierarchies would achieve greater prominence, resulting in a basic form of focusing. It would also enable saving of screen space as the amount of duplication visualised across the hierarchies would be reduced. Alternatively the user could choose which portions of which hierarchies they wished to observe. A form of filtering such as observing a slice of the entire structure could enable a user to see just one hierarchy or the progress of one particular sub-tree and its component nodes through all hierarchies.

Human visual perception aspects would relate to visual cues that could be used to differentiate certain information, either to avoid clutter or as a form of non-spatial focusing. An example is shown in Figure 2, where the intermediate level used in two of the classifications is faded out to allow easier visual comparison of the levels common to all hierarchies. In Figure 3, all nodes apart from the squares are set to the same level of contrast, an example of filtering, which allows the user to focus more easily on a sub-group (the squares) of particular interest.

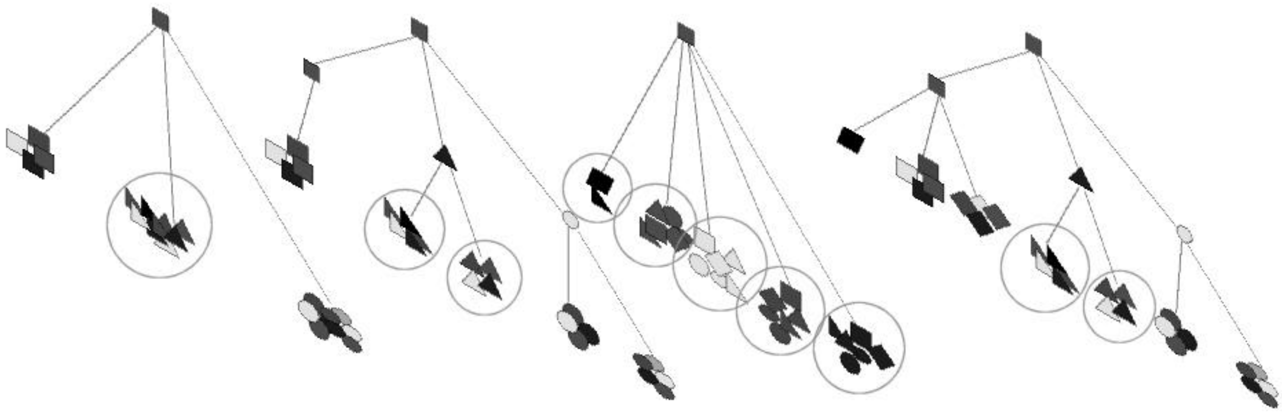

Figure 4 Highlighting of all sub trees that contain triangles

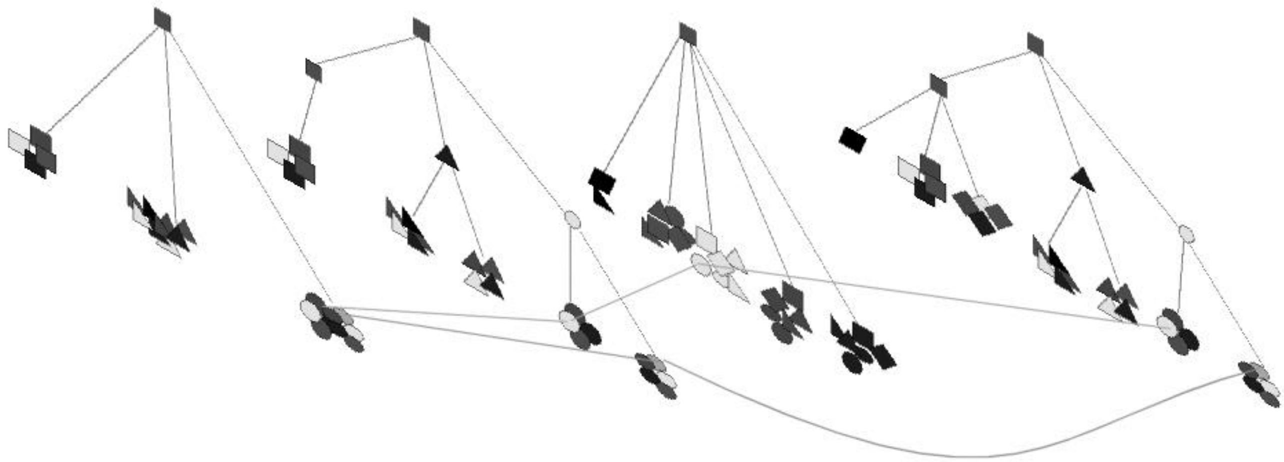

Figure 5 Tracking of individual shapes across hierarchies

Further, in Figure 4, a simple highlighting technique shows all the sub-trees that triangles appear in. This has a similar effect to Figure 3 as it draws the user's attention to the location of the specimens of interest. For taxonomists these visualisations would help determine the method used to formulate the hierarchies. Different methodologies would result in differing patterns of distribution, as can be seen from figure 4's indication that triangles appear in all of the third hierarchy's sub-trees.

Another situation would be in order to visualise a node's progress through a number of hierarchies, its path could be highlighted. Other visual cues of lesser prominence for the sibling, child and parent nodes would also show their journeys through the multiple classifications. This would enable a viewer to see the context change of the various hierarchies, a point we outlined the need for in the previous section. A small example of this is shown in Figure 5, tracking one node and a sibling. It also highlights missing information as one of the nodes is not represented in the third classification, and so the nodes' path skips this hierarchy.

6 CONCLUSION

We have presented a review of current information visualisation techniques and described a visualisation problem for which no adequate visualisation technique has been developed. Issues facing taxonomists in comparing multiple overlapping classifications exemplify the generic problem of visualising structural change in hierarchies. Techniques that might have approached the visualisation of comparative multiple hierarchies have been evaluated and their shortcomings highlighted.

The requirements from a visualisation for this problem were presented and sketches of visualisations and techniques to address the issue were described.

Further work will include applying to our sketches Brath's Metrics for Information Visualization [46], which judge the appropriateness of a non-interactive screenshot of a visualisation. Following this, a prototype demonstration in Java 2.0 for evaluation with the taxonomists will be developed. A full implementation of the visualisation will then proceed, which will undergo full user testing with the taxonomists at the Royal Botanic Garden Edinburgh.

Acknowledgement

Martin Graham is supported by an EPSRC PhD studentship. We would like to thank the taxonomists Mark Watson and Mark Newman at the RBGE for their time spent discussing the classification and visualisation issues in taxonomy.

7 REFERENCES

[1] Gershon, N., Card, S. K. and Eick, S. G., Information Visualization Tutorial, in *Proc. ACM CHI '98*, pp. 109-110, Los Angeles, USA, April 18-23, 1998. ACM Press.

[2] Hubona, G. S., Shirah, G. W. and Fout, D. G., "The effects of motion and stereopsis on three-dimensional visualization." *International Journal of Human-Computer Studies*, 1997, **47** (5): 609-627.

[3] Ware, C. and Franck, G., "Evaluating Stereo and Motion Cues for Visualizing Information Nets in Three Dimensions." *ACM Transactions on Graphics*, 1996, **15** (2): 121-140.

[4] Bartram, L., Perceptual and interpretative properties of motion for information visualization, in *Proc. NPIV '97. New Paradigms in Information Visualization and Manipulation*, pp. 3-7, Las Vegas, November 13-14, 1997. ACM Press.

[5] Drevermann, H. and Travis, D., "Visualization using Colour: visual presentation of events in particle physics." *Behaviour & Information Technology*, 1998, **17** (1): 18-26.

[6] Healey, C. G., Choosing Effective Colours for Data Visualization, in *Proc. IEEE Visualization 1996*, pp. 263-270, Oct. 27 - Nov. 1, 1996. IEEE Computer Society Press.

[7] Healey, C. G., Booth, K. S. and Enns, J. T., "High-speed Visual Estimation using Pre-attentive processing." *ACM Transactions on Human-Computer Interaction*, 1996, **3** (2): 107-135.

[8] Healey, C. G., Booth, K. S. and Enns, J. T., "Visualizing real-time multivariate data using preattentive processing." *ACM Transactions on Modelling and Computer Simulation*, 1995, **5** (3): 190-221.

[9] Robertson, G. G., Mackinlay, J. D. and Card, S. K., Cone Trees: Animated 3D Visualizations of Hierarchical Information, in *Proc. CHI '91 : Human Factors in Computing Systems*, pp. 189-194, ACM Press.

[10] Andrews, K., Pichler, M. and Wolte, J., Information Pyramids: A New Approach to Visualising Large Hierarchies, in *Proc. Late Breaking Hot Topics, IEEE Visualization '97*, pp. 49-52, Phoenix, October 19-24, 1997. IEEE Computer Society Press.

[11] Johnson, B. and Shneiderman, B., Treemaps: A Space-Filling approach to the visualization of hierarchical information structures, in *Proc. IEEE Visualization '91*, pp. 284-291, San Diego, IEEE Computer Society Press.

[12] Rapley, M. H. and Kennedy, J. B., Three Dimensional Interface for an Object Oriented Database, in *Proc. Second International Workshop on User Interfaces to Databases*, pp. 133-158, Ambleside, UK, July 13th-15th, 1994.

[13] Herman, I., Delest, M. and Melancon, G., Tree Visualisation and Navigation Clues for Information Visualisation, Technical Report INS-R9806, CWI (Centruum voor Wiskunde en Informatica), 1998.

[14] Wood, A., Drew, N., Beale, R. and Hendley, B., HyperSpace: Web Browing with Visualization, in *Proc. Third International World-Wide Web Conference*, pp. 21-25, Darmstadt, Germany,

[15] Hendley, R. J., Drew, N. S., Wood, A. M. and Beale, R., Narcissus: Visualising Information, in *Proc. IEEE Information Visualisation Symposium 95*, pp. 90-97, October 30-31, 1995. Computer Society Press.

[16] Donath, J. S., Visual Who: Animating the affinities and activities of an electronic community, in *Proc. ACM MultiMedia '95*, pp. 99-108, San Francisco, USA, November 5-9, 1995. ACM Press.

[17] Bentley, C. L. and Ward, M. O., Animating Multidimensional Scaling to Visualize N-Dimensional Data Sets, in *Proc. IEEE Information Visualization 1996*, pp. 72-73, San Francisco, October 28-29, 1995. IEEE Computer Society Press.

[18] Sprenger, T. C., Gross, M. H., Eggenberger, A. and Kaufman, M., A Framework for Physically-Based Information Visualization, in *Proc. Visualization in Scientific Computing '97*, pp. 71-83, Boulogne-sur-Mer, France, April 28-30, 1997. SpringerWienNewYork.

[19] Xiong, R., Smith, M. and Drucker, S., Visualizations of Collaborative Information for End-Users, Technical Report MSR-TR-98-52, Microsoft Research, 1998.

[20] Rokita, P., "Generating Depth-of-Field Effects in Virtual Reality Applications." *IEEE Computer Graphics & Applications*, 1996, **16** (2): 18-21.

[21] Zhai, S., Buxton, W. and Milgram, P., "The Partial-Occlusion effect: Utilizing Semi-Transparency in 3D Human-Computer Interaction." *ACM Transactions on Human-Computer Interaction*, 1996, **3** (3): 254-84.

[22] Benedikt, M., Cyberspace: some proposals. *Cyberspace: First Steps*. M. Benedikt, MIT Press, 1991, pp. 119-224.

[23] Lucas, P., Roth, S. F. and Gomberg, C. C., Visage: Dynamic Information Exploration, in *Proc. CHI '96 conference companion on Human factors in computing systems:*

Common Ground, pp. 19-20, April 13-18, 1996. ACM Press.

[24] Parker, G., Franck, G. and Ware, C., “Visualization of Large Nested Graphs in 3D: Navigation and Interaction.” *Journal of Visual Languages and Computing*, 1997, **9** (5): 299-317.

[25] Wegenkittl, R., Loffelmann, H. and Groller, E., Visualising the behavior of higher dimensional dynamic systems, in *Proc. IEEE Visualization 1997*, pp. 119-125, Phoenix, October 19-24, 1997. IEEE Computer Society Press.

[26] Benford, S. and Mariani, J., Virtual Environments for Data Sharing and Visualisation - Populated Information Terrains, Research Report CSCW/5/94, SECAMS, Lancaster University, Lancaster, 1994.

[27] Graves, A., Hand, C. and Hugill, A., Interactive Visualisation of Musical Form using VRML, in *Proc. Fourth UK VR-SIG Conference*, pp. 98-109, Brunel University, England, 1st Nov 1997.

[28] Leung, Y. K. and Apperley, M. D., “A review and taxonomy of distortion-oriented presentation techniques.” *ACM Transactions on Human-Computer Interaction*, 1994, **1** (2 (June 1994)): 126-60.

[29] Sarkar, M. and Brown, M. H., “Graphical fisheye views.” *Communications of the ACM*, 1994, **37** (12): 73-84.

[30] Munzner, T., “Exploring Large Graphs in 3D Hyperbolic Space.” *IEEE Computer Graphics & Applications*, 1998, **18** (4): 18-23.

[31] Lamping, J. and Rao, R., Visualizing Large Trees Using the Hyperbolic Browser, in *Proc. ACM CHI '96*, pp. 388-389, Vancouver, April 13-18, 1996. ACM Press.

[32] Carpendale, M. S. T., Cowperthwaite, D. J. and Fracchia, F. D., “Extending Distortion Viewing from 2D to 3D.” *IEEE Computer Graphics & Applications*, 1997, **17** (4): 42-51.

[33] Schaffer, D., Zao, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S. and Roseman, M., “Navigating hierarchically clustered networks through fisheye and full-zoom methods.” *ACM Transactions on Human-Computer Interaction*, 1996, **3** (2): 162-88.

[34] Fishkin, K. and Stone, M. C., Enhanced Dynamic Queries via Movable Filters, in *Proc. ACM CHI '95*, pp. 415-420, Denver, May 7-11 1995. ACM Press.

[35] Eick, S. G., Data Visualization Sliders, in *Proc. UIST '94*, pp. 119-120, November 2-4, 1994. ACM Press.

[36] Ahlberg, C. and Shneiderman, B., Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays, in *Proc. ACM Conference on Human Factors in Software, CHI '94*, pp. 313-317, Boston, ACM Press.

[37] Colby, G. and Scholl, L., Transparency and blur as selective cues for complex visual information, in *Proc. SPIE Image Handling and Reproduction Systems Integration*, pp. 114-125, August 1991. SPIE.

[38] Kumar, H. P., Plaisant, C. and Shneiderman, B., “Browsing hierarchical data with multi-level dynamic queries and pruning.” *International Journal of Human-Computer Studies*, 1996, **46** (1): 103-24.

[39] Pullan, M., Watson, M., Kennedy, J., Raguenaud, C. and Hyam, R., The Prometheus Taxonomic Model, Napier University / RBGE, Edinburgh, 1999.

[40] Raguenaud, C., Kennedy, J. and Barclay, P. J., A Database System for Supporting Taxonomic Work, Technical Report School of Computing, Napier University, Edinburgh, 1999.

[41] Tufte, E. R., *Envisioning Information*, Graphics Press, Cheshire, Connecticut, 1991.

[42] Huang, M. L., Eades, P. and Wang, J., “On-line Animated Visualization of Huge Graphs using a Modified Spring Algorithm.” *Journal of Visual languages and Computing*, 1998, **9** (6): 623-645.

[43] Wittenburg, K. and Sigman, E., Visual Focusing and Transition Techniques in a Treeviewer for Web Information Access, in *Proc. Visual Languages '97*, pp. 20-27, Capri, Italy, Sept 23-26, 1997. IEEE Computer Society Press.

[44] Chi, E. H., Pitkow, J., Mackinlay, J., Pirolli, P., Gossweiler, R. and Card, S. K., Visualizing the Evolution of Web Ecologies, in *Proc. ACM CHI '98*, pp. 400-407, Los Angeles, April 18-23, 1998. ACM Press.

[45] Turo, D. and Johnson, B., Improving the Visualization of Hierarchies with Treemaps: Design Issues and Experimentation, in *Proc. IEEE Visualization '92*, pp. 124-131, May 1992. Computer Society Press.

[46] Brath, R., Concept demonstration metrics for effective information visualization, in *Proc. IEEE Symposium on Information Visualization 1997 (InfoVis '97)*, pp. 108-111, Phoenix, Arizona, USA, October 20-21, 1997. IEEE Computer Society Press.