

A Comparison of Set-Based and Graph-Based Visualisations of Overlapping Classification Hierarchies

Martin Graham, Jessie B Kennedy, Chris Hand

School of Computing, Napier University,
219 Colinton Road, Edinburgh, EH14 1DJ, UK
+44 (131)

{marting, jessie, cph}@dcs.napier.ac.uk

ABSTRACT

The visualisation of hierarchical information sets has been a staple of Information Visualisation since the field came into being in the early 1990's. However, at present, support for visualising the correlations between multiple, overlapping sets of hierarchical information has been lacking. This is despite the realisation that for certain tasks this information is as important as the information that forms the individual hierarchies. In response to this, we have produced two early visualisation prototypes, one based on a graph visualisation, and the other on a set-based metaphor, that endeavour to display such information in a readily perceived form to potential users. The science of botanical taxonomy is used as an example of a field where such a visualisation would be useful, and also as a resource for example information sets that the prototypes can act upon. Technical and perceptual issues involved in the design and implementation of both prototypes are discussed. Following this, informal user testing on both prototypes is described, which utilised user observation techniques to elicit qualitative feedback from the taxonomists. These findings are then used to emphasise the shortcomings and advantages of each prototype, and from these probable issues for future prototyping and development are drawn.

Keywords

Guides, instructions, authors kit, conference publications.

1. INTRODUCTION

Information visualisation techniques for displaying and examining single instances of hierarchical information sets have a history dating back to Robertson et al's Cone Trees [25]. However, there is a need for visualisation techniques that have the capacity to show multiple hierarchies that result from the re-classifying of information and allow exploration of the relationships between them [10].

Such multiple overlapping hierarchies can occur in a number of

areas, such as document classifications, computer file systems, and biological taxonomies. The particular information that we are concerned with visualising is a product of the practice of botanical taxonomy, the study and classification of plants.

Some current visualisation systems can display multiple hierarchies, but lack either the ability to interactively explore the relationships between them, or can only show simple changes in structure such as addition and deletion, not reclassification of existing information.

This paper describes the development and early user evaluation of two prototypes designed to rectify this shortcoming and visualise multiple overlapping hierarchies. The early user evaluations will be reported, and from these findings, conclusions on the prototypes' suitabilities and the directions for future work will be put forward. Firstly, we will describe the practice of taxonomy using a simple example to clarify the type of information we are aiming to visualise.

2. 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 or morphological similarities. Taxa are assigned to ranks that specify the level of a taxon in a classification hierarchy, though the levels (or ranks) used in generating the classification hierarchies vary for different groups of specimens and also between taxonomists.

The classification is then 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.

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

Thus a challenge generated by the way taxonomists work is the management of the accumulation of old historical classifications. As stated, 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.

Permission to make digital or hand copie of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AVI 2000, Palermo, Italy.

© 2000 ACM 1-58113-252-2/00/0005..\$5.00

A second challenge is that the choice of criteria and the way a classification is created, such as 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 (RBGE), is developing a database to support taxonomic working practice. Full details of taxonomic working practice can be found in Pullan et al [23], and a description of the Prometheus database to support multiple classifications in Raguenaud et al [24]. The visualisation issues described here have arisen out of the work with the taxonomists at RBGE.

2.1 An Example of How Multiple Classifications Evolve

This section describes a small example of how different classifications arise. 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 left figure (a) 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 taxa.

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 types square, triangle and circle. They also introduce two 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 and for ovals. This is classification (b).

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, ignoring one particular shade as there is only one instance of it, shown as classification c. Each group contains an existing type specimen and therefore no new types need 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, including some new examples, shown in Figure 1 as classification (d).

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

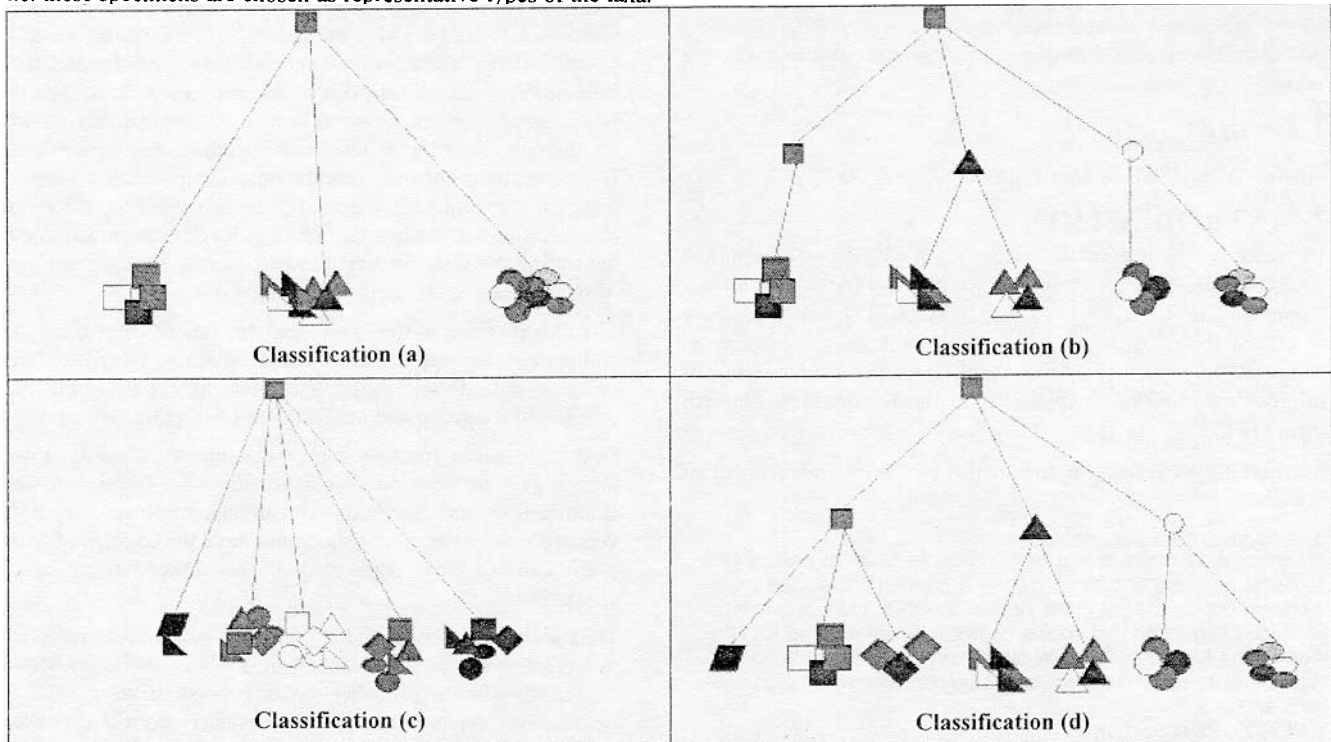


Figure 1 Four classifications with overlapping specimens and concepts

scenarios on a classification. Hence, a visualisation that supports the work of taxonomists must allow them to explore the similarities and differences between the classifications.

2.2 Taxonomist requirements

From our preliminary meetings with the taxonomists we gleaned a number of tasks that a proposed visualisation should be able to carry out or support. These tasks at the moment are quite basic, as the paper-based nature of their current system prevents them from carrying out such work at length or in much detail. However, these are the tasks that the taxonomists expect to be able to perform using a visualisation:

1. To track a particular genus's siblings and parents across re-organised taxonomic structures, if present.
2. To track a particular higher-level node's children across re-organised taxonomic structures, if present.
3. To compare the number of distinct levels within and across a set of taxonomic hierarchies.
4. To compare the structure of whole classifications against each other, though this was stated to be an infrequent and secondary task.

Surprisingly, comparison of different sub-trees across hierarchies was stated to be a less useful concept by the taxonomists, due to the nature of the way in which different taxonomies are arrived at. Therefore, support of this task would not be an overriding concern.

3. PREVIOUS TECHNIQUES FOR VISUALISING MULTIPLE HIERARCHICAL STRUCTURES

Previous techniques for the display of multiple hierarchies have generally consisted of two approaches, namely animation [12; 11], and the technique described by Tufte [27, Ch. 4] as 'small multiples', where numerous snapshots of information sets are displayed together [28; 5]. A detailed breakdown of the pros and cons of visualisations that use these techniques is given in Graham et al [10], though essentially the conclusions for both approaches can be summarised as follows.

Animation has strong pre-attentive cues that attract a user's visual focus and is useful for showing gradual change in a structure. However, the types of re-classifications we are concerned with are sometimes abrupt and therefore not easily shown in a gradual fashion. Furthermore, animation at any one point in time shows the change between one state and the next, in our case the change between only two classifications. Other previous states in the animation would have to be recalled from memory. We regard this as a serious problem as one aim of visualisation techniques is to reduce the cognitive load on the user.

The second technique, small multiples, has the advantage of theoretically displaying all the different structures at the same time. However, the disadvantages are that for 'N' small multiples to be displayed, each will only have 1/Nth amount of the screen space at most, otherwise some cannot be displayed at their full resolution or shown at all. Also, changes across the structure will have to be discovered by comparing each small multiple against each other in turn, as the pre-attentive cue of motion in animation is not present here.

This is not to say that either approach is wrong or entirely unsuitable, but it is important to be aware of the limitations associated with each technique, so compensations and solutions for particular situations can be explored.

3.1 Other Approaches

Another approach of interest to this field is Furnas and Zacks' Multitrees [9], which describes a graph structure that enables the re-use of hierarchical structure, along with a visualisation of such a structure. Furnas and Zacks' work diverges from our specific problem as their system re-organises entire sub-sections of an existing hierarchy to give different viewpoints on the hierarchy. The example they give is professors giving different reading lists composed of fragments of a hierarchy of volumes, books, chapters and sections. Conversely, taxonomy re-organises and re-uses just the bottom-level nodes of a particular group of ranks. Pieces of previous taxonomies may be incorporated into another taxonomy, but this is not always the case.

Multitrees' specification also forbids two nodes from having more than one distinct path between them. Multiple hierarchy taxonomies often have this characteristic, whereby one particular node would have two distinct parents in different taxonomies, and both of these would have the same parent node themselves in their hierarchies. Therefore between the two nodes (grandchild and grandparent nodes) would lie more than one distinct path, so Multitrees could not fully describe the inter-relations between the multiple hierarchies we are concerned with. Furnas and Zacks acknowledge this, and describe it as a systematic problem for the Multitrees' structure and visualisation. They suggest that such structures could be broken down into separate Multitrees, and the classifications could then be visualised separately. This means that correlations between different classifications cannot be shown directly, and it is also indicated that the approach was not carried out but remains a putative solution to the problem.

Wittenburg et al's research on "Group Asynchronous Browsing on the WWW" [29] is probably the closest work in concept to the problem we are tackling. It combines the graph structure of Multitrees, which is used to amalgamate a number of bookmark structures, with a small-multiple Treemap-based visualisation [14] of the separate bookmark hierarchies. Treemaps is itself a space-efficient method for displaying single hierarchies, and works by recursively dividing a given area according to the hierarchy's structure. The whole visualisation runs under the Pad++ zooming user interface environment. [1]

It can be argued that as the individual bookmark hierarchies re-use only the individual bookmarked pages, and not any higher structures; it is not a true example of Multitrees as Furnas and Zacks envisaged the concept, and closer in methodology to the node-only re-classifications of taxonomy. However, this is a moot point as it is the visualisation that interests us most.

Each bookmark hierarchy, displayed as a Treemap, is given an associated colour that is used to mark its particular nodes. Additional cross-reference nodes are placed in a Treemap if a particular bookmark is shared between bookmark hierarchies, and these cross-reference nodes are coloured according to the colour of the hierarchy into which they link.

This system would be enough, when applied to our multiple taxonomies, to show the shared nodes between each classification. However, it can give no indication of how one particular sub-tree

in a particular hierarchy, equivalent to a family or genus, is distributed by itself or with respect to its peers. This information is unavailable as the colour-coding of the nodes is static and set for each hierarchy. Information regarding the distribution of parts of a particular tree is therefore indistinguishable among the other parts of the hierarchy.

This is an important distinction as, unlike the bookmark hierarchies, there is a high degree of overlap between taxonomic hierarchies. Applying Wittenburg et al's visualisation would simply show a large number of cross-reference nodes with no detailed indication of how they group with each other across the other hierarchies. However, Wittenburg et al's system is the nearest technique we have found with regard to the type of information we need to visualise.

4. PROTOTYPES

After consulting with the taxonomists to realise their requirements, work began on prototyping visualisations that could overcome the drawbacks described in the previous visualisation section. What follows is a description of the prototypes and the ideas behind their application along with the problems encountered.

Both prototypes were programmed in Java 1.2 (now known as Java 2.0) using the Kawa Java environment, and run on a 400-MHz Pentium-II-based computer. Java was chosen as the language to build our prototypes for two main reasons. Firstly, any prototypes built could be shown on the WWW as an applet to

a wide, critical population to gain general feedback. Huang's visualisations of very large graphs [12] and Inxight's hyperbolic tree visualisation [13] have on-line demonstrations of this nature. It is our aim to shortly have our visualisations accessible over the WWW. Secondly, novel visualisations' graphical interfaces tend not to be composed of standard UI objects, for example spring-mass models for graphs, fisheye lenses etc. As a result, these objects are not available as standard objects in rapid prototyping environments, and trying to implement or mimic such features is often harder than coding them in a programming language such as Java.

4.1 Design of a Graph-based prototype

The first prototype attempted to combine the individual taxonomy hierarchies into what would structurally be a Directed Acyclic Graph (DAG), shown in Figure 2. This could be done due to the fact that the information in the different hierarchies has a high degree of correlation. As previously stated, the nodes remain relatively invariant, it is mainly the links and hence the organisation of the nodes that constitute the major change between hierarchies. The prototype visualises this DAG structure using the spring-mass metaphor commonly used in network visualisations.

The approach appears to go against common wisdom in that developing good visualisations of networks/graphs is much harder than producing good visualisations of individual hierarchies. Indeed, Mukherjea et al [19] take the approach of deconstructing a graph structure, in their case hypermedia structures such as the

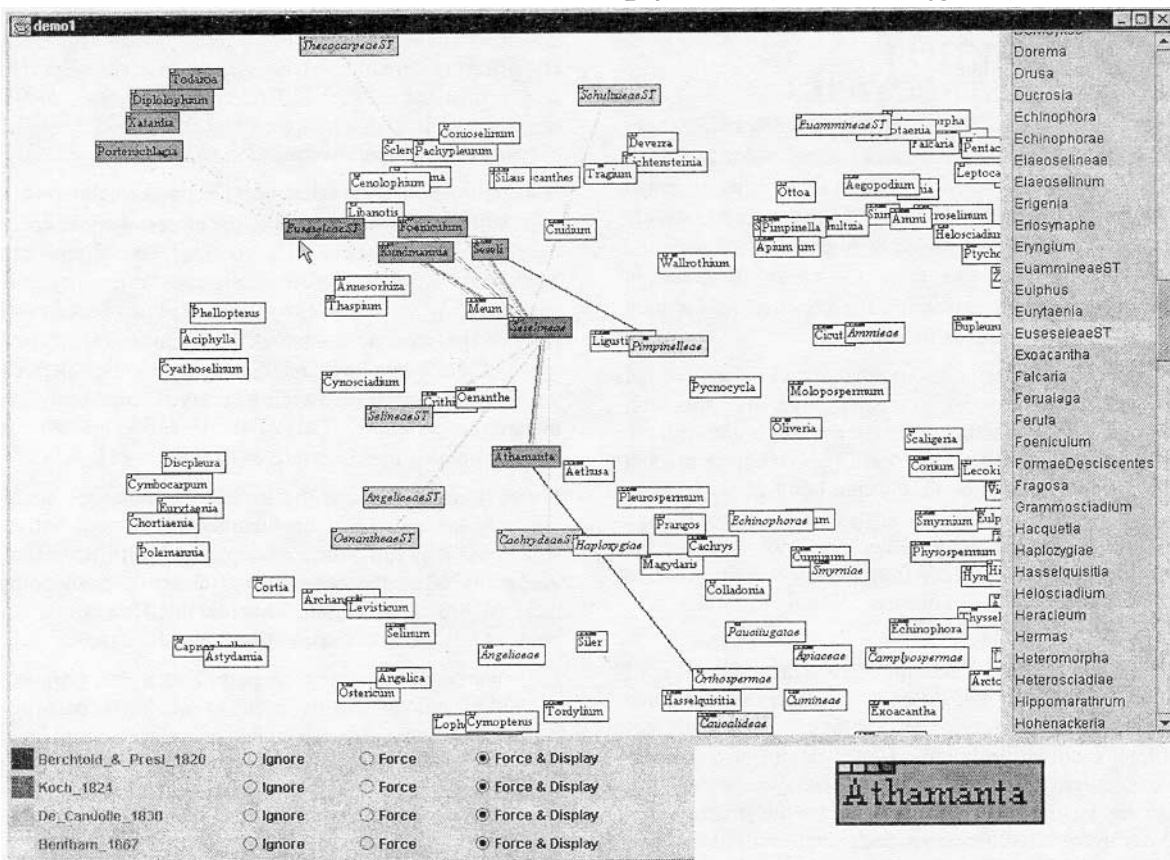


Figure 2 Graph-based Prototype

WWW, into a number of different hierarchical structures, each of which can be viewed separately. In defence of the design's methodology, Mukherjea et al were focusing on a task that involved seeing hierarchical organisation within a graph. Conversely, the taxonomists wish to see how multiple hierarchies correlate with each other. In effect, they want to know how the simpler, hierarchical structures interact to produce more abstract information, the reciprocal of the task Mukherjea et al were concerned with. User testing will show the outcome of the effect of visualising the users' task more directly versus the increased perceptive and cognitive load imposed by a graph-based visualisation.

The main advantage of this prototype is the integration of the many hierarchies spatially and temporally, avoiding the drawbacks and problems previously described with the animated and small multiple approaches. All the hierarchies can be seen at once in a single visualisation.

The visualisation of the DAG is displayed using a spring-mass model to move nodes to their positions. Unlike most spring-mass model based systems, these positions are not final and are constantly re-calculated by the prototype, though Donath's Visual Who system [7] works on the same approach. This is because the prototype allows the user to switch on and off the display of individual hierarchies, and the visualisation can be allowed to re-adjust itself to suit only the hierarchies that are to be displayed.

It must be stated that this approach has disadvantages, both in human factors and algorithmic complexity. Primarily the spring-mass model is not guaranteed to generate the same layout for the same data. Also, Misue et al [17] claim that layout creation methods such as the spring-mass model are not always the most suitable methodologies for layout update. The spring-mass model performs both layout creation and update operations in this prototype. Another problem with this approach is the amount of time needed to recalculate and redisplay the nodes, especially in Java, designed as it is for portability and reusability, not for speed.

The visualisation displays the links between the nodes using a colour coding designed to differentiate the links belonging to each separate hierarchy that forms the graph. The colours are from a linear colour scale (see Levkowitz [15]), in which the perceived difference between colours is proportional to the distance they are separated by on the scale. The colours used are evenly spread along the scale so they appear visually to be evenly spread along the colour range, making differentiation easier.

The nodes are displayed individually as labelled rectangles, a magnified example of which is shown inset at the bottom right-hand corner of Figure 2. Each node is coloured along a grey-scale according to its depth from the taxonomic root. We can do this, despite the notion that nodes might be at differing depths in different hierarchies, due to the fact that botanical taxonomy enforces a strict organisation on the nodes it contains. A node's depth is judged as its rank in the taxonomic structure of families, species, genera etc, rather than its depth in any particular hierarchy. To complement the links, a small set of coloured glyphs is shown running along the top of each node, indicating that node's membership of the individual trees in the graph. This enables membership of particular classifications to be noted for nodes when the links are not displayed.

The reasoning behind this particular allocation of grey and colour indicators is that colour scales and grey scales are perceived differently. Grey scales are perceived as being ordinal, having some quantitative meaning, and colour scales are generally seen as being nominal, only indicating membership or some other qualitative meaning [6]. Therefore, in this case it makes sense to use colour to indicate membership of a particular tree, a qualitative attribute, and the grey scale to show distance from the root, a metric that can be quantified.

At the time of writing, the user's interaction entails adding or removing particular hierarchies from the graph, along with panning and zooming controls to gain an overview or a close-up of a particular area of the graph. The user also has the ability to view only one node and its relations (siblings, children etc), therefore filtering out the other nodes and links.

The prototype approached a limit of showing and updating roughly 250 node positions at a rate of 4/5 refreshes per second. The algorithm for calculating updates on the spring-mass model is of $O(N^2)$ complexity, so to reduce this it was decided to attempt to incorporate Chalmers' linear time layout algorithm [4] for the spring-mass model, which performs in the order of $O(N)$. Unfortunately the approach foundered for two reasons, the first being that the graphical update on-screen took a considerable amount of time also, and the algorithm could not reduce this. Secondly, Chalmers' algorithm was designed to move objects rapidly to their final positions, upon which the model would then be visualised. Our prototype continuously displays the positions of nodes during updates, and as Chalmers' algorithm uses a degree of stochastic sampling, this introduced an unacceptable amount of visual jittering whenever changes were made in the spring-mass model. This is an example of the layout creation versus update issue that Misue et al [17] remarked upon. Therefore, the use of this method had to be ruled out in the visualisation prototype.

4.2 Design of a Set-based Prototype

The second prototype, shown in Figure 3, was influenced by a number of factors that emerged from background reading and the initial implementation of the graph-based prototype, though it was mostly developed in parallel with this prototype. These factors included difficulties such as the speed problems of Java and the visual clutter caused by merging all the hierarchies into one visual structure.

In visualising an organisational hierarchy, we are able to draw on a strong delineation between leaf nodes and non-leaf (internal) nodes. Essentially in such a hierarchy, leaf nodes are objects of some type, and internal nodes are categories of varying abstraction used to impose an organisation on the leaf nodes. For example, a file directory consists of files as leaf nodes and directories as internal nodes.

On the other hand, trees that show navigation routes such as Huang's web-browsing visualisation [12] have no such distinction. Leaf nodes indicate the same type of object as internal nodes, the difference being that no further navigation can be or is made from the leaf nodes.

It can be argued from a taxonomic standpoint that the objects being re-categorised in taxonomic hierarchies are families or genera, and hence abstract categorisations themselves, the physical objects being the instances of plant specimens. However,

it is these categories that are being re-organised at a higher level, remaining constant across the different hierarchies. As such, for our purposes, they can be regarded as objects or indivisible sets of objects to be categorised, and not as categories themselves.

Hence, this visualisation moves away from the node-and-link metaphor and towards a set-based visualisation. This could be more productive as Parunak [21; 22] states that taxonomic reasoning, the categorisation of objects, is essentially set-based. Thus the user's mental model of such a process is inclined towards a set-based model, rather than a node-link system. A visualisation for aiding such a task should therefore benefit from mirroring a set-based metaphor on-screen.

As a result of this distinction, the second prototype was developed as a set of small multiples, accepting the restrictions of this approach that have been previously stated in Section 3, with the leaf nodes given a different representation to the internal nodes. To reduce screen space problems, the leaf nodes of each bottom level category are arranged in a grid formation, as opposed to the normal style of a linear layout for each level of a hierarchy. This reduces the chances of individual hierarchies' displays spreading horizontally across the screen width and exceeding the overall display dimensions.

Interaction is performed using a linking and brushing technique. A selection of a particular node, particular category or set of categories in one hierarchy is shown dynamically in the context of

given a separate colour, and its distribution throughout the other hierarchies can then be seen, as the correlating nodes in the other hierarchies are similarly highlighted. This enables comparisons of correlations to be made across the set of hierarchies. We believe this to be a novel use of linking techniques, as it is applied across a set of visualised hierarchies, not scatterplots or parallel coordinates as in Brodbeck et al's work [3], or two differing visualisations of the same hierarchy as performed by Fua et al [8]. It is also this functionality that we believe is an advantage over Wittenberg et al's visualisation, in which distributions are also colour-coded but only statically at the top level.

Previous selections are shown in a history bar along the bottom of the screen. It was thought at first to represent past selections as miniature versions of the visualisation, but the space taken up by such representations would have been too great. Instead, small boxes containing the names of the node selected for past interactions were displayed. Clicking on one of these boxes would take the visualisation to the state that clicking on the node itself would invoke.

The second prototype also avoids the visual clutter of the first prototype at the cost of losing a level of detail. The leaf nodes are not individually labelled, and only one set of distinct relationships between the hierarchies can be seen at any one time. However, it avoids the speed restrictions of the first prototype, as the visualisation and internal model are not updated continuously,

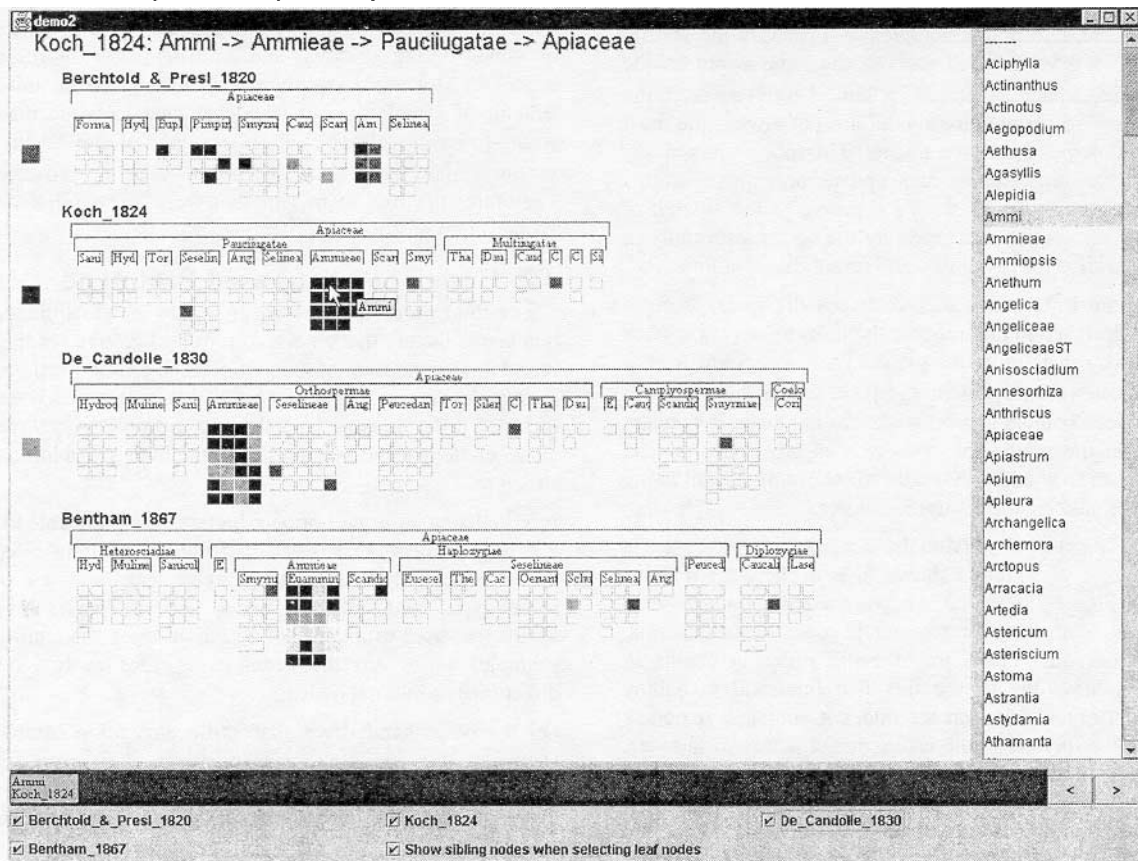


Figure 3 Set-based prototype

the other hierarchies, by a simple use of colour on the selected nodes. The set of nodes in each particular highlighted category is

only when the user initiates interaction.

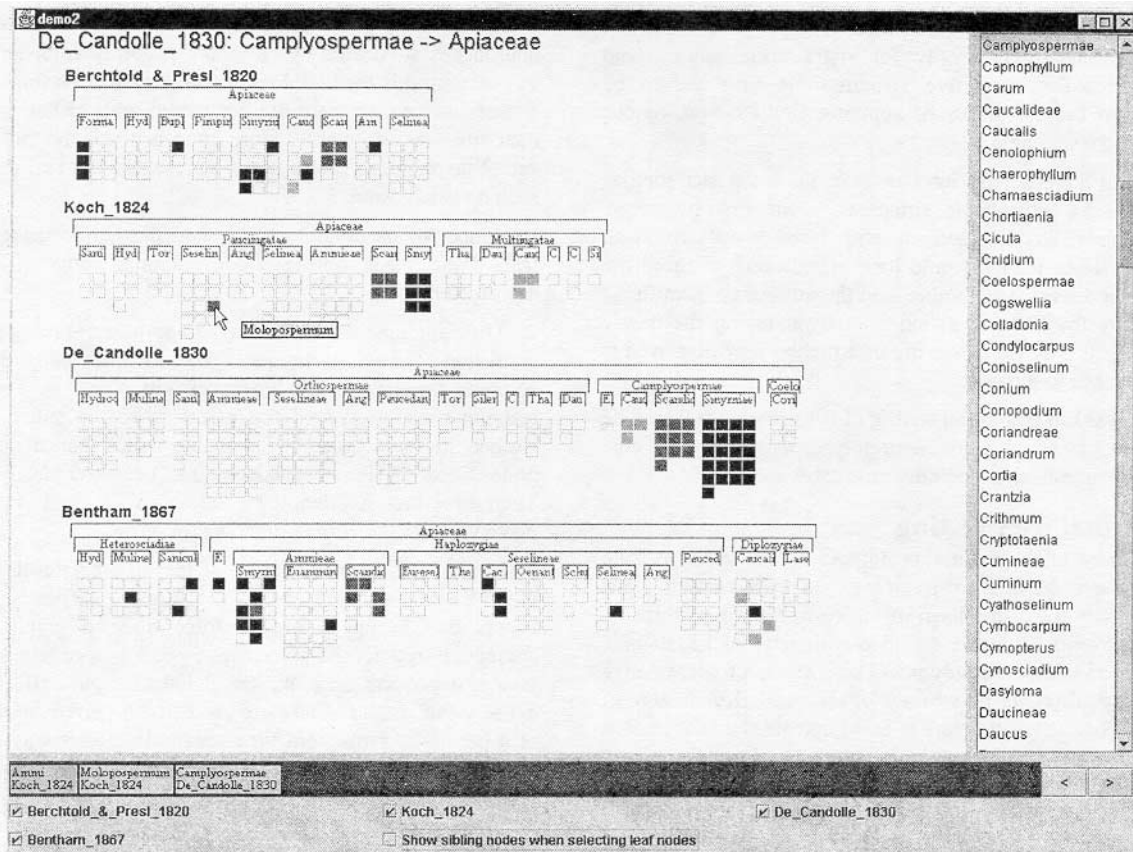


Figure 4 Groupings of De Candolle's Camplyospermae genus across other classifications

4.3 Example of use

An example of how the second prototype can be used is now described. In the screenshot of Figure 4, the sub-family known as Camplyospermae has been selected in the third hierarchy down (the hierarchy under 'De Candolle 1830' in the figure), to view the distribution of its component genera throughout the other hierarchies. This represents the task of tracking a higher-level node's genera from one classification across a number of other classifications, equivalent to the second task in the taxonomists' original requirements. This has shown a pattern of distribution that indicates, via the colouring of groups of specimens, that the correlation between the second and third taxonomies ('Koch' and 'De Candolle') is quite strong, with only one node from De Candolle's taxonomy being split from its original grouping. This is the mid-grey node that is being investigated by the mouse pointer, a genus called Molopospermum, which is a member of Scandiceae in De Candolle's classification.

In reality as Koch's is the earlier classification (from 1824), the visualisation shows that this particular node has been taken from one part of Koch's hierarchy and grouped with the other nodes in De Candolle's classification. Looking at the distribution of the mid-grey Scandiceae nodes in the other hierarchies, it shows again a strong grouping across the hierarchies, with only one mid-grey node in Bentham's classification being grouped dissimilarly.

Further investigation of the rogue node Molopospermum in Koch is warranted. The option to show all sibling nodes of this node is activated and the node clicked upon, which performs the task of displaying the distribution of a particular genus's sibling genera

across all classifications (task 1 in our group of requirements in section 2.2). The resulting visualisation is shown in the next screenshot, Figure 5.

This shows that the dark-grey nodes that are present in Koch's classifications, and that form Molopospermum's siblings in this classification, are also classed together in De Candolle's classification, all with the exception of Molopospermum which is ringed by a border. De Candolle has chosen to group it with a new family, whose distribution across the hierarchies is indicated as the mid-grey nodes. In Bentham's classification, Molopospermum has been taken out from the mid-grey nodes and grouped with a new family called Smyrniae, indicated by the set of light-grey nodes, and in effect this is the third different group it has been associated with in as many classifications. This shows that Molopospermum was also the differently grouped mid-grey node in Bentham's classification in Figure 4. This gives an indication that Molopospermum could be a difficult node to classify, perhaps exhibiting many features in common with many of the other genera.

5. MEASURING PROTOTYPE EFFECTIVENESS

5.1 Choosing evaluation methods

Visualisation prototypes can be subjected to metrics such as those proposed by Brath [2], which attempt to give a quantitative measure of characteristics such as on-screen data density, percentage of occlusion in 3D systems, and percentage of

identifiable data points. However Miller et al [16] point out that these metrics are suitable only for static visualisations, and useable metrics for interactive visualisations have yet to be established, so the usefulness of applying Brath's work to our prototypes is questionable.

An example of the problems interaction could cause such metrics can be given by a probable situation in our first prototype. Suppose a user has zoomed in and filtered out irrelevant information. These actions would have significantly reduced the on-screen data density by zooming, and the number of identifiable data points by the filtering action, thereby reducing the metric scores, though it is obvious that the information now displayed is of more relevance to the user.

Therefore it was felt that initial testing of our prototypes would be best served by having users try the prototypes using representative tasks, rather than rely on potentially unsuitable metrics.

5.2 Informal user testing

To gain feedback on these initial prototypes, the taxonomists from the RBGE were invited to participate in user testing. This informal testing took the form of a co-operative evaluation, wherein the taxonomists were supplied with a list of 12 example tasks to perform with the prototypes. These tasks, a representative example being "discover all siblings of the genus Kundmannia", were designed to cover the earlier requirements the taxonomists had supplied us with. The taxonomists were asked to describe out loud their actions, observations, and difficulties as they progressed through the tasks. This style of evaluation, known as a think-aloud protocol [26], is suitable for bringing out qualitative

issues in early prototype testing, and is one of a suite of approaches advocated by "discount usability engineering" [20]. In line with this methodology, the testing was performed on only 3 users, but as we were not concerned with gathering statistical data, this was acceptable. More important was the fact that being actual taxonomists, these users represented the real end-users of such a visualisation.

A number of observations related to the performance of certain tasks were made during the evaluation, some of which are described below:

1. Overlapping nodes in the graph-based prototype caused confusion. Though techniques exist for alleviating this, it is a common problem in graph/network visualisations. The zooming technique also worsened this effect when the entire graph was shrunk. In some tasks, it led to the belief that all the relevant nodes were visible, whereas some had been completely obscured. Therefore the overlapping nodes could lead to incorrect perceptions of the correct outcome for a task.
2. The lack of spatial ordering of the different levels within the graph-based prototype caused problems, a difficulty that again occurs due to the layout of force directed graphs. Previous attempts to restrict the placement of nodes to certain areas of the layout, dependent on rank, had failed as it gave rise to areas of dense visual clutter. Therefore tasks that involved finding a name at a particular rank were time-consuming, as it was difficult to gain visual cues to tell if the name was indeed at the required rank. In short, users found it impossible to derive rank information from spatial layout alone.

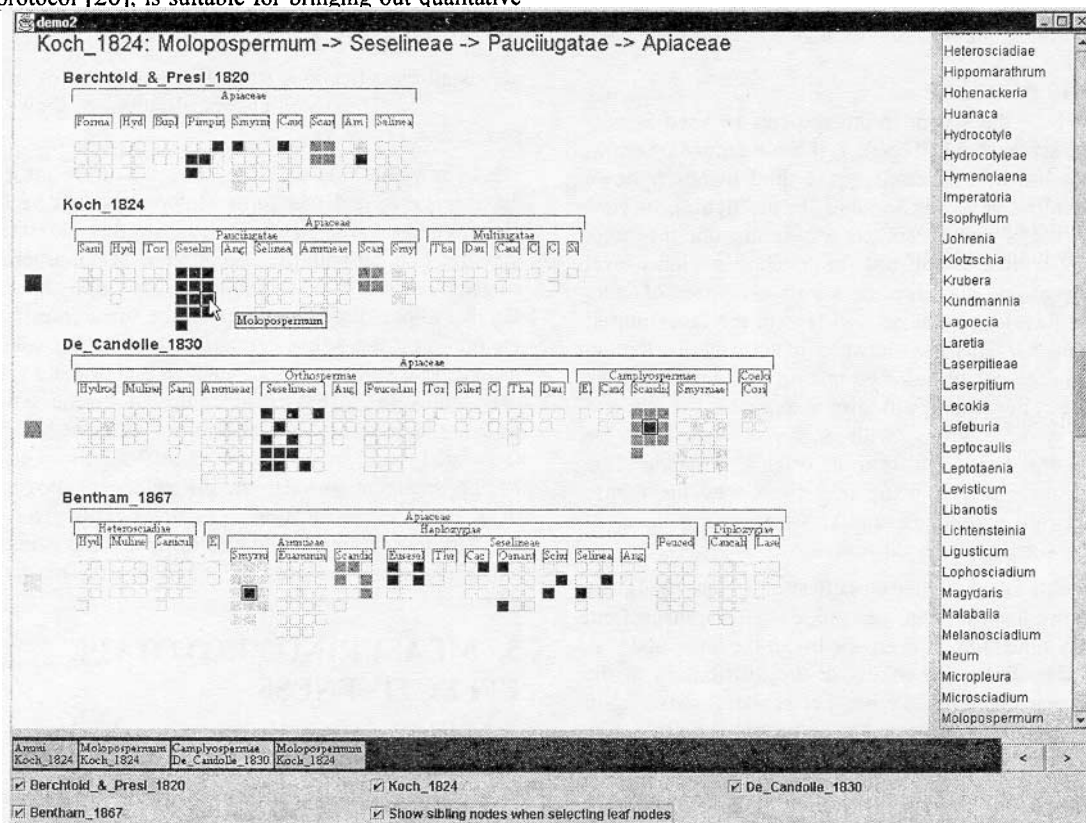


Figure 5 Display of siblings of Molopospermum across all classifications

3. The set-based prototype abbreviated the names of the higher rank nodes. This was noted to be a problem as rank within the taxonomic structure can also be distinguished by the end of a name, and these were cut off by the abbreviation. Tasks that required names at a specific rank to be discovered were therefore hindered. This indicated that purely spatial cues were not enough to indicate rank in this prototype.

4. In using the set-based prototype when the task involved finding the existence of a specific non-genus node in a hierarchy, the visualisation showed, as in all other tasks, the distribution of the genera classified under that non-genus node. This had two effects. Firstly, if these genera were distributed across more hierarchies than the non-genus node itself was present in, it appeared to indicate that the non-genus node was also present in more classifications than was actually the case. Secondly, as it was the genera that were highlighted, it sometimes caused the perception that one of the highlighted genera was actually the node the user was supposed to find.

Following Monk and Wright's [18] idea of using task-based observations to produce corresponding inventions, each of the above observations can lead to a possible solution (the invention), that can either be applied directly or used to generate further possible solutions. For example, the observations above lead to the following possible solutions:

1. All directly or indirectly selected nodes, should be displayed on top of all other nodes, ensuring they are not obscured by information that is currently not of interest.
2. As enforcing spatial ordering had already failed, a non-spatial cue for indicating rank appears to be a solution. In addition to the glyphs that indicate individual classification membership, each node could also carry a prominent glyph that indicates rank.
3. Additional rank information could be indicated by textual labels aligned to the respective ranks at the side of each classification.
4. Non-genus as well as genera nodes could be highlighted when selected. This in conjunction with the previous invention should give more clarity to the non-genus node when it is selected.

These inventions, and others, may be implemented on the next iteration of the prototypes, tackling the problems encountered by the taxonomists who used the prototypes.

While both prototypes could show the information necessary for answering the questions posed by basic tasks, there was a definite preference for the set-based prototype. The taxonomists stated that the multiple tree effect and the grouping of sets was closer to how they viewed classifications when working with taxonomic data, rather than a node-link diagram. This supported our original assumption that a set-based presentation style would match their mental model more closely than the graph-based visualisation. Also, it was felt that even with filtering mechanisms the graph-based prototype showed too much visually, leading to clutter and confusion. It was also said that the zooming technique felt unnatural, as the taxonomists were used to the Microsoft Windows style of scrolling to extend a windows' range of view. The initialisation period necessary for the spring-mass diagram to settle into a preliminary layout was also detrimental.

Further tasks that the users envisaged after seeing the prototypes included the ability to see if certain nodes only appeared in one

classification, a task neither visualisation prototype can adequately perform at the moment. Another wish was to see the extension of the visualisation of a single genus's sibling distribution to a larger set of genera.

6. CONCLUSIONS

In conclusion, we have presented two separate prototypes that indicate different approaches for tackling the problem of visualising multiple overlapping hierarchies. The technical and initial usability problems of both approaches have been noted and their effects evaluated. Indeed, given the comments and observations resulting from the user evaluation, and other technical considerations, it is felt that development should continue primarily on the set-based visualisation. The graph-based visualisation could still be explored but its limitations and drawbacks seem to indicate it is a poorer metaphor for displaying our information than the set-based prototype.

The users stated that both prototypes gave them the ability to see information that was not easily understandable or available in their current paper-based system. As reported, viewing the prototypes also encouraged the users to consider extending the visualisations to further tasks, such as the ability to see if certain nodes only appeared in one classification.

Though further development is necessary, we believe that the set-based prototype in particular indicates a technically feasible and user-centred approach for visualising multiple overlapping hierarchies. Therefore, immediate future work will consist of refining this prototype in line with the implications of the observations gained from the informal user testing. The focus will also move to tackling information sets that contain more hierarchies, and also hierarchies with larger structures, to see how the prototype visualisations handles them. In the long term, the aim is to discover whether the techniques employed in the visualisation are suitable for viewing other similar information sets, such as document classifications or changing organisational structures.

7. ACKNOWLEDGEMENTS

We would like to thank Mark Watson, Mark Newman and Martin Pullan at the RBGE for their input and valuable feedback on the prototypes. Martin Graham's PhD is funded by an EPSRC studentship.

8. REFERENCES

- [1] Bederson, B. B. and Hollan, J. D. Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. Proc. of UIST '94 (Marina Del Ray, California, USA, November 2-4, 1994), ACM Press, 17-26.
- [2] Brath, R. Concept demonstration metrics for effective information visualization. Proc. of IEEE Symposium on Information Visualization 1997 (InfoVis '97) (Phoenix, Arizona, USA, October 20-21, 1997), IEEE Computer Society Press, 108-111.
- [3] Brodbeck, D., Chalmers, M., Lunzer, A. and Cotture, P. Domesticating Bead: Adapting an Information Visualization System to a Financial Institution. Proc. of IEEE InfoVis '97 (Phoenix, Arizona, USA, October 20-21, 1997), IEEE Computer Society Press, 73-80.

- [4] Chalmers, M. A Linear Iteration Time Layout Algorithm for Visualizing High-Dimensional Data. Proc. of IEEE Information Visualization '96 (San Francisco, California, USA, October 28-29, 1996), Computer Society Press, 127-132.
- [5] Chi, E. H., Pitkow, J., Mackinlay, J., Pirolli, P., Gossweiler, R. and Card, S. K. Visualizing the Evolution of Web Ecologies. Proc. of ACM CHI '98 (Los Angeles, California, USA, April 18-23, 1998), ACM Press, 400-407.
- [6] Csinger, A. The Psychology of Visualization, University of British Columbia.
- [7] Donath, J. S. Visual Who: Animating the affinities and activities of an electronic community. Proc. of ACM MultiMedia '95 (San Francisco, USA, November 5-9, 1995), ACM Press, 99-108.
- [8] Fua, Y.-H., Ward, M. O. and Rundensteiner, E. A. Navigating Hierarchies with Structure-Based Brushes. Proc. of InfoVis '99 (San Francisco, California, USA, October 25-26, 1999), IEEE Computer Society Press
- [9] Furnas, G. W. and Zacks, J. Multitrees: Enriching and Reusing Hierarchical Structure. Proc. of ACM CHI '94 (Boston, Massachusetts, USA, April 24-28, 1994), ACM Press, 330-6.
- [10] Graham, M., Kennedy, J. B. and Hand, C. The Challenge of Visualising Multiple Overlapping Classification Hierarchies. Proc. of User Interfaces to Data Intensive Systems (UIDIS '99) (Edinburgh, UK, September 5-6, 1999), IEEE Computer Society Press, 42-51.
- [11] Herman, I., Melancon, G., Ruiters, M. M. d. and Delest, M. Latour - a tree visualisation system, Centrum voor Wiskunde en Informatica (CWI), Netherlands.
- [12] 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, 9 (6), 623-645.
- [13] Inxight. Inxight Software Inc. Site Lens Demos. http://www.inxight.com/Demos/SLS_Demos/Site_Lens_Studio_Demos.html
- [14] Johnson, B. and Shneiderman, B. Treemaps: A Space-Filling approach to the visualization of hierarchical information structures. Proc. of IEEE Visualization '91 (San Diego, California, USA, Oct 22-25, 1991), IEEE Computer Society Press, 284-291.
- [15] Levkowitz, H. and Herman, G. T. Color Scales for Image Data. IEEE Computer Graphics & Applications, 12 (1), 72-80.
- [16] Miller, N., Hetzler, B., Nakamura, G. and Whitney, P. The Need For Metrics In Visual Information Analysis. Proc. of New Paradigms in Information Visualization and Manipulation 1997 (NPIV '97) (Las Vegas, USA, November 13-14, 1997), ACM Press
- [17] Misue, K., Eades, P., Lai, W. and Sugiyama, K. Layout Adjustment and the Mental Map. Journal of Visual Languages and Computing, 6 (2), 183-210.
- [18] Monk, A. F. and Wright, P. C. Observations and inventions: new approaches to the study of human-computer interaction. Interacting with Computers, 3 (2), 204-216.
- [19] Mukherjea, S., Foley, J. D. and Hudson, S. Visualizing Complex Hypermedia Networks through Multiple Hierarchical Views. Proc. of ACM CHI '95 (Denver, Colorado, USA, May 7-11, 1995), ACM Press, 331-7.
- [20] Nielsen, J. Guerrilla HCI: Using Discount Usability Engineering to Penetrate the Intimidation Barrier. In R. G. Bias and D. J. Mayhew (eds). Cost-Justifying Usability, Chapter 11, Academic Press Professional.
- [21] Parunak, H. V. D. Don't Link Me In: Set Based Hypermedia for Taxonomic Reasoning. Proc. of ACM Hypertext '91 (San Antonio, Texas, USA, December 15-18, 1991), ACM Press, 233-242.
- [22] Parunak, H. V. D. Hypercubes Grow on Trees (and other Observations from the Land of Hypersets). Proc. of ACM Hypertext '93 (Seattle, Washington, USA, November 14-18, 1993), ACM Press, 73-81.
- [23] Pullan, M., Watson, M., Kennedy, J., Raguenaud, C. and Hyam, R. The Prometheus Taxonomic Model. To appear in Taxon (February).
- [24] Raguenaud, C., Kennedy, J. and Barclay, P. J. A Database System for Supporting Taxonomic Work. Edinburgh, School of Computing, Napier University.
- [25] Robertson, G. G., Mackinlay, J. D. and Card, S. K. Cone Trees: Animated 3D Visualizations of Hierarchical Information. Proc. of CHI '91 : Human Factors in Computing Systems (New Orleans, Louisiana, USA, April 27 - May 2, 1991), ACM Press, 189-194.
- [26] Tognazzini, B. User testing on the cheap. TOG on Interface, 79-89, Chapter 14, Addison-Wesley.
- [27] Tufte, E. R. Envisioning Information. Graphics Press, Cheshire, Connecticut, 1991.
- [28] Turo, D. and Johnson, B. Improving the Visualization of Hierarchies with Treemaps: Design Issues and Experimentation. Proc. of IEEE Visualization '92 (Boston, Massachusetts, USA, October 19-23, 1992), Computer Society Press, 124-131.
- [29] Wittenburg, K., Das, D., Hill, W. and Stead, L. Group Asynchronous Browsing on the World Wide Web. Proc. of Fourth International World Wide Web Conference (Boston, Massachusetts, USA, Dec. 11-14, 1995), 51-62.