A Colour-Filling Approach For Visualising Trait Evolution With Phylogenies

Savrina F. Carrizo
School of Information Technologies
University of Sydney
New South Wales, Australia, 2006
`scarrizo@it.usyd.edu.au`

Abstract

The development of powerful visualisation tools is a major challenge in bioinformatics. Phylogenetics, a field with a growing impact on a variety of life science areas, is experiencing an increasing but poorly met requirement for software supporting the advanced visualisation of phylogenies. This paper describes a new method to visualise trait evolution across a phylogenetic tree by using colour to represent ancestral trait values. This space-efficient method supports trait visualisation on very large phylogenetic trees, and adds a new dimension to the conventional phylogenetic tree representation in order to convey extra information. Initial usability tests by phylogeneticists show that this technique indeed enhances their data analysis processes.

Keywords: phylogenetic tree, trait evolution, colour

1 Introduction

Visualisation is an integral part of the field of bioinformatics. It has roles not only in analysis, but also in building more use-friendly interfaces, implementing methods to navigate large information spaces intuitively, and powerful techniques to browse and query data interactively the via visualisation (Robinson and Flores, 1997). The value of visualisation in bioinformatics lies in the potential of new visual metaphors to bring insight to unfamiliar biological data sets, and in the capability to stimulate collaborative and dynamic enquiry (Rhyne, 2000).

Phylogenetics is a field with a growing impact on a variety of life science areas. The central aim is to discover the evolutionary relationships i.e. *phylogenies* of the tree of life. The discovery and analysis of the patterns and processes that underpin present-day biodiversity has many practical applications including drug discovery, conservation and forensics, to name a few.

Phylogenetic trees, like the one in Figure 1, are the visual representation of inferred phylogenies. The leaf nodes represent the species or molecules under study, the internal nodes represent inferred hypothetical ancestors, and the edges, when weighted, represent the amount of evolutionary change that has occurred between two evolutionary events (i.e. those that are represented by the internal and leaf nodes of the tree).

Phylogenetic analyses rely heavily on visual inspection, structural comparison, manipulation and exploration of phylogenetic trees and thus present a number of visualisation challenges. The available tools (Felsenstein, 2002; Swofford 1999; Huelsenbeck and Ronquist, 2001), however, are characterised by a lack of effective visualisation techniques. It is not uncommon for biologists to "(fall) back on paper, tape and highlighter pens" due to current deficiencies in phylogenetic visualisation programs (Munzner et al., 2003).

Phylogenetic trees convey much biological information relative to their simple two-dimensional (2D) form, but phylogenetic data often contain more information than can currently be visualized by phylogenetic tools; one such case is that of how characteristics or *traits* of the species in question have changed over evolutionary time.

This paper presents a new method of visualising the evolution of traits across a phylogenetic tree based on the use of colour for data representation. The evolution of traits has not been the subject of information visualisation research before. The new method adds dimensionality to the traditional 2D phylogenetic tree and increases the information richness of the display. The result of an initial usability analysis with phylogeneticists has confirmed the utility of the method in phylogenetic research.

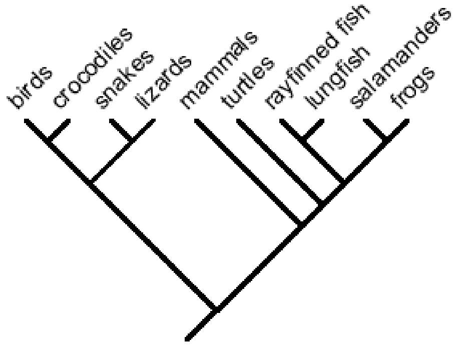

Figure 1: A phylogenetic tree (Maddison and Maddison, 1992)

In section 2 an overview of trait evolution and current methods of analyses are presented, followed by colour

---

This paper appeared at the Australasian Symposium on Information Visualisation, Christchurch, New Zealand, 2004. Conferences in Research and Practice in Information Technology, Vol. 35. Neville Churcher and Clare Churcher Eds.

research in section 3. Sections 4 and 5 present the design and implementation of the new colour-filling method. Experimental results are presented in section 6. Use cases for the new method are presented in section 7. Section 8 presents an evaluation of the method and possibilities for future research, followed by the conclusion in section 9.

2 Trait Evolution and Previous Work

Discovering how traits may have evolved for a given set of species is the first step to understanding why they evolved in such a manner. Trait evolution is the subject of much phylogenetic research, for example see (Ho et al., 2004; Pagel, 1999; Pagel, 1997).

Trait types are commonly referred to as ecological, anatomical, physiological or genomic. They may be continuous/quantitative characters, such as metabolic rate estimates, or discrete/qualitative, such as the presence or absence of a character or various character states.

Existing methods of analysis for trait evolution are predominantly statistical and to our best knowledge only one other program (Maddison and Maddison, 2002) provides a visualisation of the evolution of traits. Applications such as Discrete (Pagel, 2000), Continuous (Pagel, 2000) and Compare (Martins, 1999) are used to estimate the ancestral trait values. States are inferred for all ancestral nodes based on present-day values and the given phylogeny.

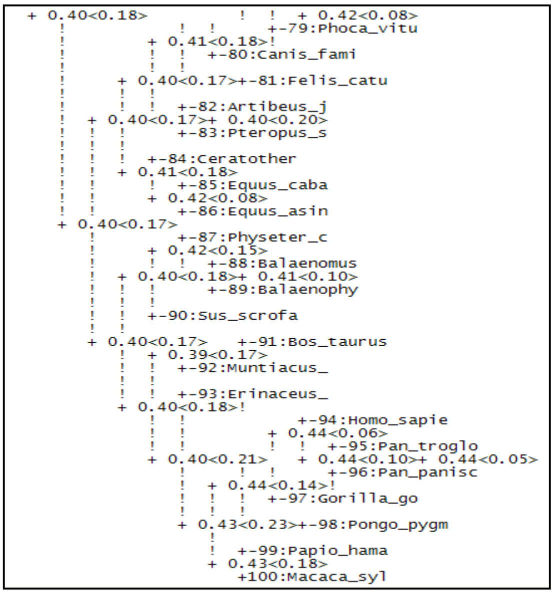

```
+ 0.40<0.18>         !  !  + 0.42<0.08>
   !           !  !     +-79:Phoca_vitu
   !        + 0.41<0.18>!
   !        !  !  +-80:Canis_fami
   !        !  !
   !     + 0.40<0.17>+-81:Felis_catu
   !     !  !
   !     !  !  +-82:Artibeus_j
   !  + 0.40<0.17>+ 0.40<0.20>
   !  !  !     +-83:Pteropus_s
   !  !  !
   !  !  !  +-84:Ceratother
   !  !  + 0.41<0.18>
   !  !     !  +-85:Equus_caba
   !  !     + 0.42<0.08>
   !  !        +-86:Equus_asin
   + 0.40<0.17>
      !        +-87:Physeter_c
      !     + 0.42<0.15>
      !     !  !  +-88:Balaenomus
      !  + 0.40<0.18>+ 0.41<0.10>
      !  !  !     +-89:Balaenophy
      !  !  !
      !  !  +-90:Sus_scrofa
      !  !
      + 0.40<0.17>   +-91:Bos_taurus
         !  + 0.39<0.17>
         !  !  +-92:Muntiacus_
         !  !
         !  !  +-93:Erinaceus_
         + 0.40<0.18>!
            !  !           +-94:Homo_sapie
            !  !        + 0.44<0.06>
            !  !        !  !  +-95:Pan_troglo
            + 0.40<0.21>   + 0.44<0.10>+ 0.44<0.05>
               !     !  !     +-96:Pan_panisc
               !  + 0.44<0.14>!
               !  !  !  +-97:Gorilla_go
               !  !  !
               + 0.43<0.23>+-98:Pongo_pygm
                  !
                  !  +-99:Papio_hama
                  + 0.43<0.18>
                     +100:Macaca_syl
```

Figure 2: Output format of the 'Compare' application

Figures 2 and 3 show attempts by biologists to visualise trait evolution. These representations are clearly inadequate by information visualisation standards. Figure 2 is part of a tree 'diagram' output by Compare. The ancestral traits in this case are quantitative and are represented by numerical values placed at each internal node. The tree structure itself is difficult to see, making the interpretation of the trait values equally difficult. Figure 3 shows a *manually* constructed visualisation of the same data using pie charts. This representation is an improvement on Figure 2; however, close visual inspection is required to determine where the trait changes occur.

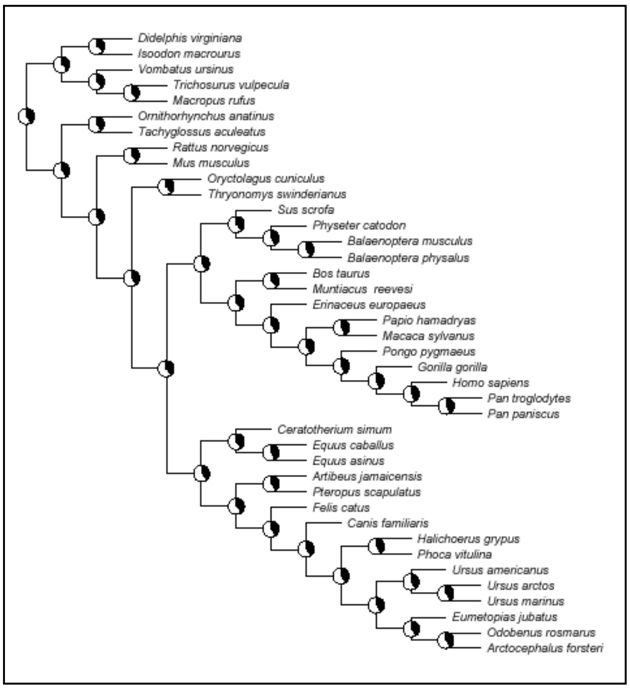

Figure 3: A visualisation constructed manually, i.e. using vector graphics software to draw and add pie charts. Pie charts represent relative proportions of nucleotide content in genes. The black represents the GC content and the white represents the AT content (Ho et al. 2004).

MacClade (Maddison and Maddison, 2002), a commercial application, visualises traits by applying "pattern, shade or colour" to branches, or to lines or glyphs placed along the branches. Brewer (1999) reports that small coloured objects are more difficult to identify than larger coloured areas (Brewer, 1999). Also, fewer colours are able to be distinguished with small point symbols (glyphs) or thin lines (Brewer, 1999). Indeed, the authors of MacClade have found this to be a problem and in the user manual state, "*In general you can look at the character legend to see which colours correspond to which states but the shades may be difficult to distinguish*".

The above representations are not based on sound visualisation principals, i.e. those that take advantage of the perceptual capabilities of the human visual system, thus providing the motivation for this study. This study applies previous research in the use of colour for data representation to address the above issues.

3 Colour use for Data Visualisation

The application of colour to visualisations plays an important role in data representation and analysis. Much research has been conducted on how colour is perceived by the human visual system and thus how it can best be applied to data visualisation. Most notably, studies in the use of colour for cartographic purposes and census data (Brewer, 1994; Olson and Brewer, 1997; Brewer, 1999), which are also perceivable by people with colour vision impairments, have pioneered the use of colour for data visualisation.

3.1 Perceptual Dimensions of Colour

Colour is a three dimensional phenomenon consisting of *hue, lightness* and *saturation* (Brewer, 1999). *Hue* is the perceptual dimension of colour associated with colour names, like blue and green (Brewer, 1999). *Lightness* is the most important dimension for data representation and is "a relative measure of how much light appears to reflect from an object compared to what appears white in the scene" (Brewer, 1999). *Saturation* is a measure of the vividness of a colour (Brewer, 1999). While careless use of colour will obscure patterns in data, matching the organization of the perceptual dimensions of colour to the organization of the data being represented will allow interrelationships and patterns within the data to be easily observed (Brewer, 1994).

3.2 Types of Colour Schemes

There are three basic colour schemes (Brewer, 1999) each of which can be used in conjunction with the new colour-filling approach to cater for all types of traits.

Sequential schemes: Lightness is best used to represent quantitative data that are monotonically increasing, such as ranked data or interval scales of numerical data. Colours that progress from light to dark through a series of adjacent hues are also often used to represent low to high data values (Brewer, 1999). Thus a variety of hues may be used but the scheme should first and foremost be ordered by lightness (Brewer, 1999).

Diverging schemes: A light colour is recommended to emphasize a critical value such as the mean, median, zero or threshold value in a data range, which then diverges towards different hues for high and low data extremes (Brewer, 1999).

Qualitative and binary schemes: Differences in hues are commonly used to represent categorical data. Binary schemes are a special case, where lightness and/or hue are appropriate for the two categories (Brewer, 1999).

4 The Colour-Filling Method

The goal of this research was to develop a new visualisation to convey rapidly and effectively the potentially complex nature of trait evolution. Information visualisation should help scientists understand the nature of the underlying data and enhance their abilities to generate hypotheses. The new method was developed with reference to well-established colour research and was combined with the familiar 2D phylogenetic tree to display the fluctuations of trait values over evolutionary time. This method complements and enhances existing statistical methods of analysis.

4.1 Design

The *colour-filling* method uses the perceptual dimensions of colour to represent ancestral traits on an *information background*, i.e. an added data dimension onto which the phylogenetic tree is superimposed. The background consists of colour-filled rectangles. Parts of the tree contained within a given rectangle take the trait value represented by the colour of the rectangle. Figure 4 illustrates this concept.

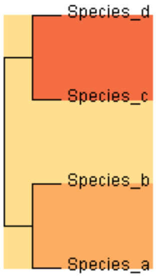

Figure 4: A simple tree superimposed on the tessellated background

While the colour-filling method may resemble *space-filling* tree-maps (Johnson and Schneiderman, 1991), it differs in that the placement and size of rectangles is dependent on the spatial layout of the phylogenetic tree, while the placement and size of rectangles in tree-maps is based on the hierarchical relationships of the tree-like data itself. See for example Figure 5.

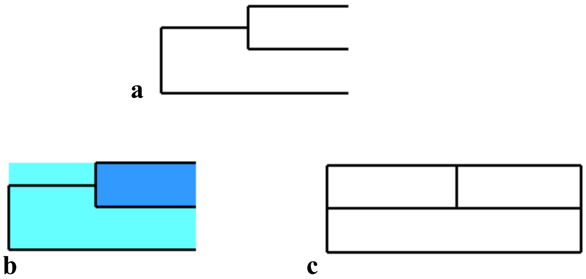

Figure 5: a) A tree structure. b) The corresponding spatial tessellation of the background in the colour-filling approach. c) A tree map showing the corresponding hierarchical tessellation.

One of the main advantages of the colour-filling method is *space-efficiency,* as it supports the visualisation of trait evolution on very large trees. Typical sizes of phylogenetic trees can range from ten to more than ten thousand leaf nodes. Visualisation techniques developed for the field must therefore ensure scalability of the methods developed.

A further advantage of the colour-filling method is the application of colour to the background rather than on glyphs or branches. This provides larger coloured areas that can be more easily perceived and thus colour (trait) trends across the screen are easily noted.

Not only does the colour-filling approach better support the visual-pattern processing abilities of researchers, it also allows for the visualisation of larger trees than is possible with the approach used in MacClade where the use of glyphs limits a tree to approximately 35 leaf nodes before scrolling becomes necessary. In contrast, the colour-filling approach has been shown to clearly display 111 leaf nodes in a single view.

4.2 Selection of Colour Schemes

The sequential and diverging colour schemes were selected with the use of the Color Brewer application (Brewer, 2002). Figure 6 shows the eleven-class diverging, nine-class sequential and five-class spectral diverging-colour schemes used in this study. The first two colour schemes were chosen based on their suitability for people with red-green colour blindness. The spectral scheme was chosen due to its popularity in scientific visualisations.

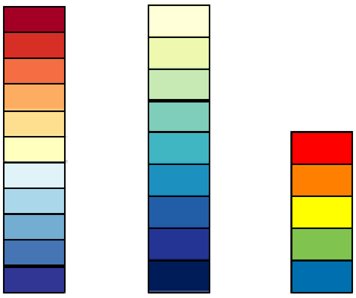

Figure 6: 11-class diverging colour scheme, 9-class sequential colour scheme, 5-class spectral diverging colour scheme.

The final colour scheme selection was decided upon after having inspected the colours on the intended visualisation. Brewer (1999) reports that comparing colours only in the key where they appear in one order on a uniform background is not adequate to ensuring an effective visualisation. Perceptual processes, such as 'assimilation' and 'simultaneous contrast', can change perceptions of individual colours substantially so that they may not match corresponding colours in the key (Brewer, 1999). Little can be done to prevent these effects as the data distribution dictates the size and position of coloured features in a visualisation (Brewer, 1999).

5 System Implementation

The system was developed using the Java (version 1.4.1) Abstract Window Toolkit (AWT) package on a Windows 2000 platform.

The method to visualise trait evolution associates the trait value to its corresponding colour in a vector containing the colour scheme. For a given internal node, the colour is applied to the bounding box of the subtree rooted at that node. Bounding boxes are recursively computed starting from the root. Coloured boxes are progressively rendered on top of each other as the algorithm proceeds.

5.1 Features

The final appearance of the visualisation is obviously dependant on the characteristics of the underlying data set. Trees with skewed[1] branch lengths are not uncommon in phylogenetics and can present some visualisation problems. Features allowing the user to interactively manipulate the visualisation not only aid understanding of the data, but also provide a means of customising the visualisation to match the features of the dataset.

Several features to address the above issues and support trait evolution analyses were implemented, including:

- Phylogram and topological tree layouts
- User-controlled colour scheme selection
- Tree manipulation operations

The details of these features are now explained.

Two layouts of the phylogenetic tree were implemented to handle cases where the data are skewed. The phylogram in Figure 7a presents a phylogeny in which the branch lengths are skewed; this makes it difficult to view the tree topology within the lower part of the tree, as the branch lengths are very short in relation to the topmost subtree. The ability to view colour changes would be impaired if the traits were rapidly changing in this lower area. To address such cases, an alternative to the phylogram is available, which displays only the topology of the tree and therefore trait changes clearly. Figure 7b shows the topological view that can be toggled on or off by the user.

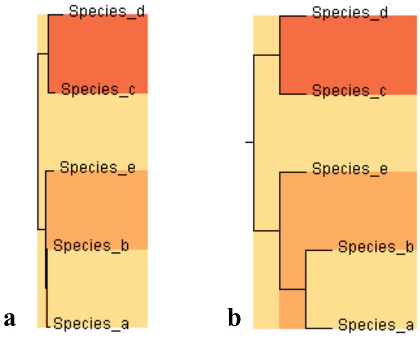

Figure 7: a) Relatively short branch lengths may obscure colours from view b) The topological layout shows colour changes clearly in this situation.

The user can also control the colour scheme selection and the number of classes within the scheme. The option of sequential, diverging and qualitative schemes is necessary to cater for all trait types. The suitability of a sequential or diverging scheme depends on the data being analysed. If there were no logical separation of extremes in the data set, then a sequential scheme would be adequate; however, if such a distinction can be drawn (e.g., *AT-rich* vs *GC-rich* genes), then a diverging colour scheme may be more appropriate. Currently, users have the option of selecting up to nine classes for a sequential scheme and up to eleven classes for the diverging scheme. An option for a greater number of classes can be easily implemented and can certainly make the display richer in information. However, many colours may also increase the cognitive load of the user and thus may make the display difficult to interpret. Further tests are required in this area.

---

[1] If a phylogenetic tree has very long branches as well as relatively short ones, the branching structure of the short ones may not be easily visible.

The ability to manipulate a phylogenetic tree, whereby a user rotates a subtree or swaps the positions of two subtrees of a given ancestral node, are necessary functions for phylogenetic visualisations, and are primarily used for illustrative and explanatory purposes as well as a visual aid in hypotheses testing. Implemented with smooth animation, the user can do these operations intuitively, without the loss of their mental map. When integrated with the colour-filling visualisation scheme, rotation and swapping operations result in the corresponding movement of background colours. Smooth animation for tree manipulation was implemented using the direct linear interpolation technique.

6 Experimental Results

For illustrative purposes the same data set used in Figures 2 and 3 has been used to demonstrate the effectiveness of the colour-filling approach for visualising trait evolution. See Figures 8 - 10. A brief description of the data is presented here. Guanine (G), Cytosine (C), Adenine (A) and Thymidine (T) are the nucleotides from which our DNA is composed and thus they are the building blocks of our genetic traits. The data under consideration are the total GC content of all mitochondrial protein-coding genes for the 111 metazoan taxa in the tree. The tree was inferred using the nuclear 18S rRNA gene and was downloaded from the Ribosomal Database Project II website (http://rdp.cme.mse.edu). The program Compare (Martins, 1999) was used to estimate the ancestral GC content values of the mitochondrial data.

The results in these visualisations effectively show the fluctuations of GC content. At a glance a user can clearly distinguish between areas of relatively high and low values.

Figure 8 provides an overview of the gradual change in GC content across the tree. It is interesting to note the fluctuations of the lower two subtrees; from the root value we observe an increase, followed by a decrease and then an increase again. The top-most subtree gradually decreases in GC content while the two larger subtrees below it display an increase in GC content.

Figure 9 reflects a number of extra changes, when compared with Figure 8, by virtue of the greater number of colour classes used. It also emphasizes the areas of the tree with the lower and higher values of GC content.

The 5-class spectral colour scheme shown in Figure 10 is also effective; it highlights the most distinctive groups of trait values; high, low and mid range values. Figure 10 demonstrates that even with a skewed data set, the colour-filling approach can clearly convey the trait values associated with the various subtrees; however should it be difficult to see the colour changes along the short branches, then the topological view may be used as an alternative.

The colour-filling visualisations take advantage of the pattern processing abilities of the human visual system. For a large tree such as the one applied here, there may be a large range of trait values; a large number of colour classes provide a richer visualisation that would benefit researchers. As demonstrated by Figure 9, the method is suitable for use with a large number of classes, the level of detail represented may provide insights into the suitability of analyses. For example, the level of GC content may affect the methods that should be used to infer the tree. Thus, the more information conveyed, the clearer is the ‘picture’ obtained by the phylogeneticists on the evolution of a trait.

7 Use Cases

Using the program developed here, many biological analyses are supported by the colour-filling visualisation and many more are possible with further development. Researchers may use the visualisation in the following way:

1. Observe the groups demarcated by trait values (colours) and see whether they are consistent with the current classification schemes.
2. Observe the rates of trait (colour) change throughout the tree, or along individual root-to-leaf paths.
3. Observe whether traits (colours) vary randomly or whether there is a distinct phylogenetic pattern in the evolution of the trait.
4. Compare sequential and diverging scheme views for quantitative characters to determine if traits exhibit skewness or have a more even distribution both numerically and in terms of the spatial distribution on the phylogenetic tree.
5. Rotate and swap operations can be used to split contiguous areas of colour to highlight convergent evolution, i.e. subtrees that have converged to the same trait value but are distantly related.
6. Similarly, rotate and swap operations may be used to make areas with the same or similar traits contiguous to highlight coevolution i.e. subtrees that have evolved similar traits due to a close phylogenetic relationship.

8 Usability Evaluation

The usability analysis was conducted with phylogenetics students and researchers to gather subjective feedback regarding the clarity and utility of the new trait evolution visualisation. The results show that the techniques developed support phylogeneticists undertaking trait evolution analyses by offering an information-rich representation. For the purposes of this paper, the responses of the senior researchers are focused on.

Three researchers were shown the examples of the visualisation as seen in Figures 8, 9 and 10. Researchers were told that the underlying trait data represented by the colours were overall GC content of mitochondrial protein-coding genes and were provided with a question sheet for assessing the visualisation.

The following questions regarding the clarity and utility of the visualisation were asked:

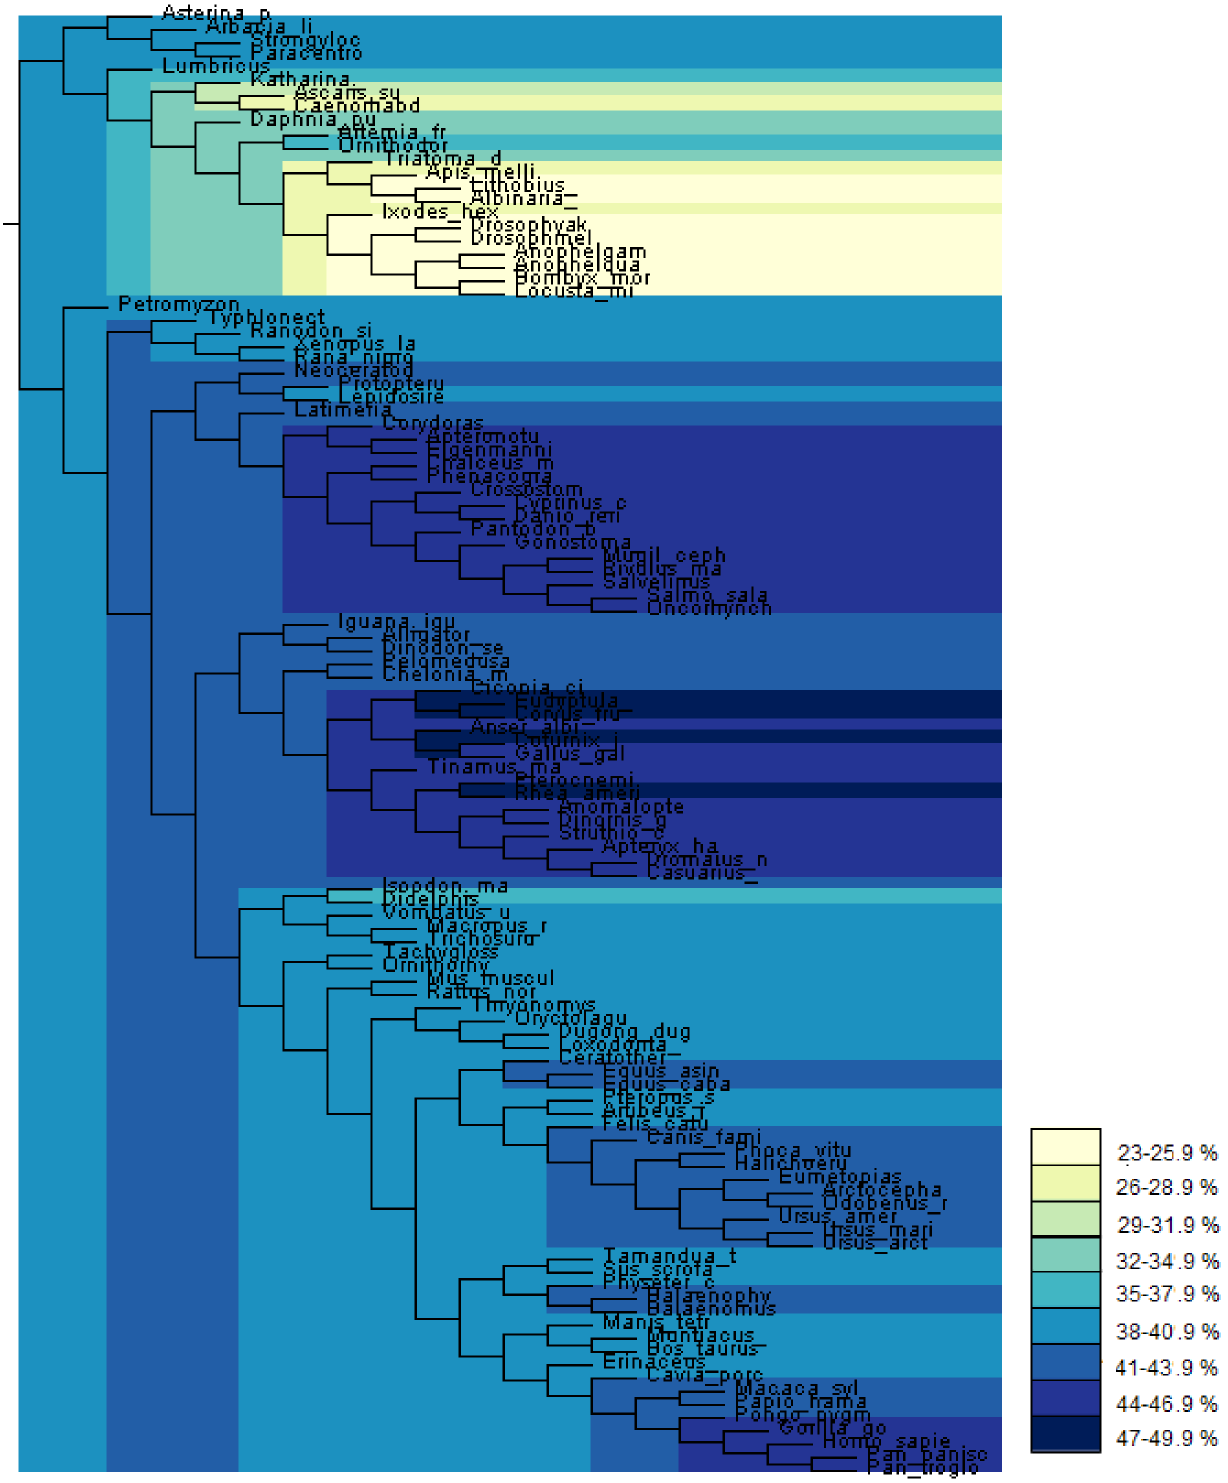

Figure 8: The colour-filling method visualising GC content from mitochondrial protein-coding genes. A nine-class sequential colour scheme is used (data from Ho et al. (2004)).

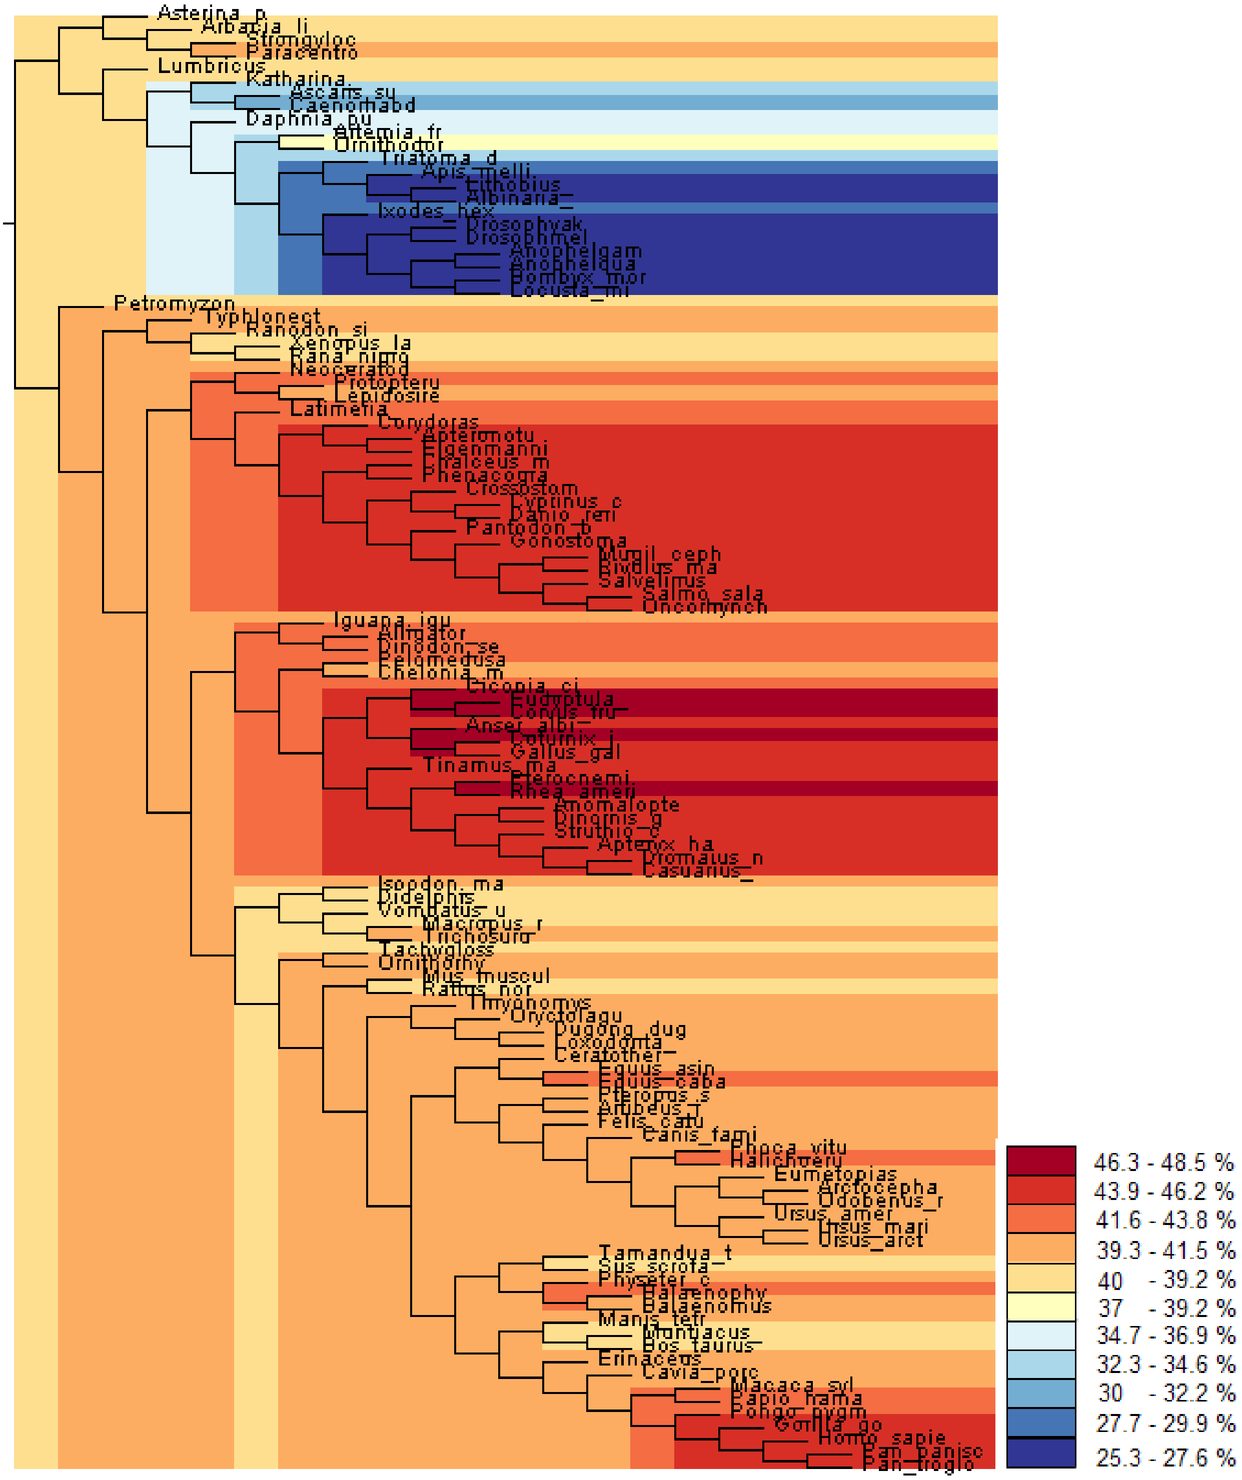

Figure 9: The colour-filling method visualising GC content from mitochondrial protein-coding genes. An eleven-class diverging colour scheme is used (data from Ho et al. (2004)).

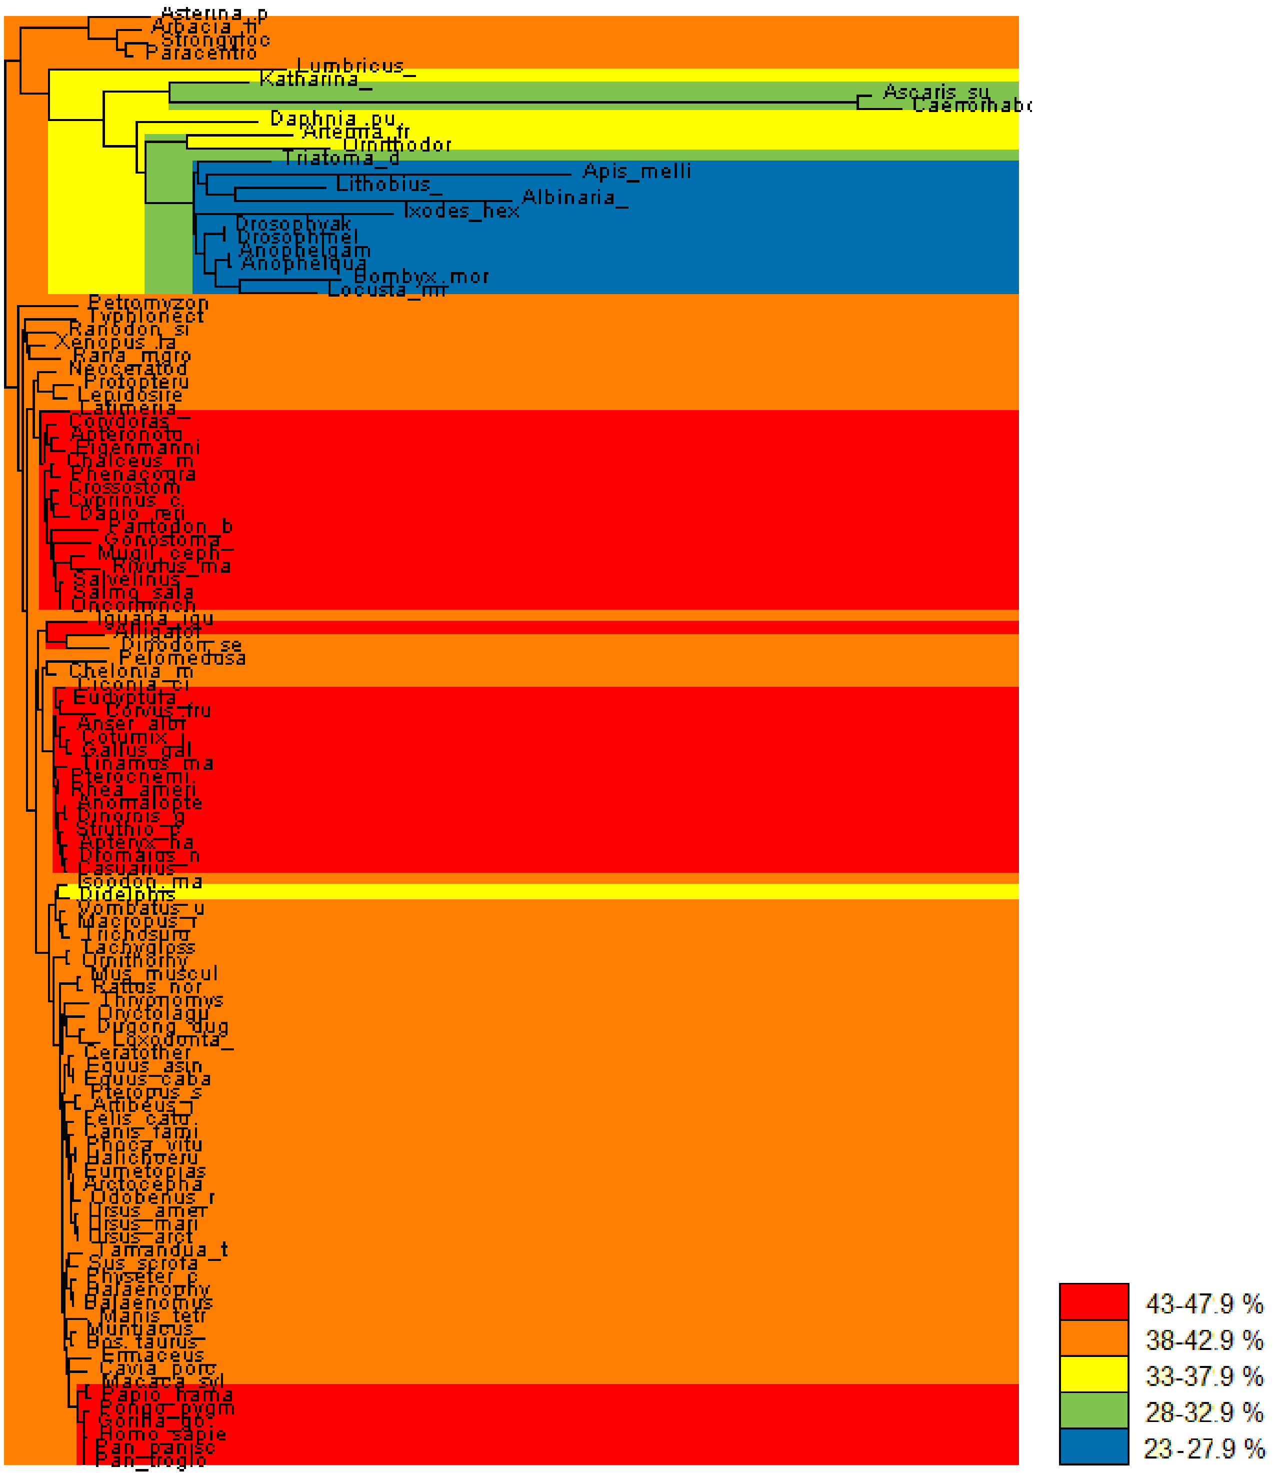

Figure 10: The colour-filling method visualising GC content from mitochondrial protein-coding genes. A five-class diverging spectral colour scheme is used. This view highlights the deviations from the mid-range values 33-37% (data from Ho et al. (2004)).

1. **Do you think that the visualisations are easy to understand?** Users indicated that the visualisation is straightforward and easy to understand.
2. **What aspects, if any, are unclear or confusing?** Users indicated that no aspects of the visualisation are unclear. All agree that the potential problem of a rapidly changing trait in situations where the internal edges are short is adequately resolved by the option for viewing tree topology.
3. **Which scheme (sequential or diverging), do you think would support analysis of quantitative trait evolution better?** All users agreed that the option for both is definitely necessary due to data dependency. It was also suggested that a diverging scheme might be easier to understand if there were many data values.
4. **Did the visualisations lead you to form questions or hypotheses about the data under consideration?** The visualisation led users to "question the processes underpinning the trait fluctuations at a biological level". In particular, "I am amazed and fascinated by the temporal fluctuation of GC content and want to know whether this occurs at different codon sites in the mitochondrial DNA and whether tRNA and rRNA genes are equally variable".
5. **Is the visualisation a good analysis tool?** All users agreed that it is a useful tool and a particular response was "it is very interesting to be able to *see* trait change across a tree". Other users indicated the visualisation "sends a powerful image about variation of the character in question" and is "a useful way to trace and describe evolutionary history".
6. **Do you think the inclusion of present day trait values (i.e. colours for leaf labels) would improve the visualisation?** Opinions indicated that it would at least be a good option to toggle present day values on or off. Another indicated that the visualisation of present day values would definitely enhance hypothesis generation and the utility of the method.

Overall, the feedback demonstrates that the colour-filling approach for visualising trait evolution with phylogenies is effective and was well received by the phylogeneticists involved.

8.1 Suggestions For Future Research

Further discussions with the users have highlighted a number of open research questions.

1. Ancestral state inferences are dependent on tree topology. If uncertain about which tree is correct, comparison of ancestral traits (colours) with those inferred from a different tree is possible.
2. Similarly, it would be useful if researchers could map the variation of one trait onto one or more others to observe correlation or coevolution of traits. This would also be useful in determining whether traits can simply be analysed from the data at leaf nodes, or whether ancestral patterns must be accounted for as well.
3. A final suggestion regarded the variance and uncertainty associated with the traits. A way of determining how robust inferences are is to generate a number of alternative topologies, or vary the edge lengths and observe any changes in inferred states.

These suggestions involve structural comparisons of phylogenetic trees, which can be complicated, and the development of more specialised techniques is required to support comparisons between many trees, large trees or many traits on a single tree.

9 Conclusion

The new colour-filling method provides an effective and intuitive means of analysing trait evolution across a phylogenetic tree. It full-fills two important roles of scientific visualisation;

1. It provides a new way of looking at complex biological data.
2. It stimulates dynamic enquiry.

With further research and refinement, the colour-filling method has the potential to further augment existing statistical methods of analysing trait evolution.

10 Acknowledgements

I gratefully acknowledge the support and feedback of my supervisors, Dr Seokhee Hong and Dr Lars Jermiin. Thank you also to Dr Masahiro Takatsuka for his valuable insight.

11 References

Brewer, C. A. (2002) Color Brewer. Available at: colorbrewer.org

Brewer, C. A. (1999) Colour Use Guidelines For Data Representation. *Proceedings of the Section on Statistical Graphics, American Statistical Association,* Alexandria VA. pp. 55-60.

Brewer, C. A. (1994) Colour Use Guidelines For Mapping and Visualisation. Chapter 7 pp.123-147 In: Visualisation in Modern Cartography. MacEachren A. M. and Talor, D. (eds) Elsevier Science, N.Y.

Cole, J.R. et al. The Ribosome Database Project (RDP-II) : previewing a new autoaligner that allows regular updates and the new prokaryote taxonomy. *Nucleic Acids Research* 31, 442-443 (2003). Available at: http://rdp.cme.mse.edu

Felsenstein, J. (2002) PHYLIP. A computer program available at: http://cmgm.stanford.edu/phylip/

Ho, S.Y.W., Childs, L.H., and Jermiin, L.S. (2004) A survey of compositional variation in metazoan mitochondrial genomes. In preparation.

Huelsenbeck, J. P. and Ronquist, F. (2001) MRBAYES: Bayesian inference of phylogenetic trees. *Bioinformatics* 17: 754-755.

Johnson, B and Shneiderman, B (1991) Tree-maps: A Space-filling approach to the visualization of hierarchical information structures. In *Proceedings of IEEE Visualization '91*, pp. 284-291.

Maddison, D. R., and W. P. Maddison. 2002. MacClade: Analysis of Phylogeny and Character Evolution. Sinauer Associates, Sunderland, Massachusetts.

Martins, E (1999) Compare. A computer program distributed by the author. Available at: http//compare.bio.Indiana.edu/

McConnell, P., Johnson, K., Lin. S. (2002) Applications of Tree_Maps to Hierarchical biological data. *Bioinformatics*18: 1278-1279.

Munzner, T., Guimbretiere, F., Tasiran, S., Zhang, L. and Zhou, Y. (2003): TreeJuxtaposer: Scalable Tree Comparison using Focus+Context with Guaranteed Visibility. In *Proc. SIGGRAPH 2003.*

Olson, J. M. and Brewer, C. A. (1997) An Evaluation of Colour Selections to Accommodate Map Users with Colour Vision Impairments. *Annals of the Association of American Geographers* 87: 103-134.

Pagel, M. 2000. Discrete. A computer program distributed by the author. Available at: http://sapc34.rdg.ac.uk/meade/Mark/

Pagel, M. 2000. Continuous. A computer program distributed by the author. Available at: http://sapc34.rdg.ac.uk/meade/Mark/

Pagel, M. (1999) Inferring the historical patterns of biological evolution. *Nature* 401: 877-884.

Pagel, M. (1997) Inferring evolutionary processes from phylogenies. *Zoologica Scripta* 26: 331-348.

Rhyne T. M. (2002) Evolving Visual Metaphors and Dynamic Tools for Bioinformatics Visualisation. *IEEE Visualization Conference Proceedings*: 579 - 582

Robinson, A. J., and Flores, T. P., (1997): Novel Techniques for Visualising Biological Information. *Intelligent Systems In Molecular Biology* 5: 241-9.

Swafford, D. (1999) PAUP, A computer program distributed by the author. Available at: http://paup.csit.fsu.edu/index.html