

CandidTree: Visualizing Structural Uncertainty in Similar Hierarchies

Bongshin Lee¹, George G. Robertson¹, Mary Czerwinski¹, and Cynthia Sims Parr²

¹Microsoft Research
One Microsoft Way
Redmond, WA 98052, USA
{bongshin, ggr, marycz}@microsoft.com

²Human-Computer Interaction Lab
Institute for Advanced Computer Studies
University of Maryland
College Park, MD 20742, USA
csparr@umd.edu

Abstract. Most visualization systems fail to convey uncertainty within data. To provide a way to show uncertainty in similar hierarchies, we interpreted the differences between two tree structures as uncertainty. We developed a new interactive visualization system called CandidTree that merges two trees into one and visualizes two types of structural uncertainty: location and sub-tree structure uncertainty. We conducted a usability study to identify major usability issues and evaluate how our system works. Another qualitative user study was conducted to see if biologists, who regularly work with hierarchically organized names, are able to use CandidTree to complete tree-comparison tasks. We also assessed the “uncertainty” metric we used.

Keywords: Uncertainty visualization, Structural uncertainty, Tree comparison, Graphical user interfaces.

1 Introduction

Most current visualization systems generally suggest certainty. This means that when we show visualizations to users, they believe that what is currently displayed is ground truth. However, there are many cases where this is not true. For example, there exist several biological taxonomies and phylogenetic trees because not all biologists agree on one taxonomy or one phylogenetic tree and some analysis methods produce multiple possible trees. Current tree visualizations such as TaxonTree [10] and Hyperbolic Tree [9] typically show one taxonomy at a time without any certainty information, which often may not be easily computed. Hence, there is no way to see which parts of the tree are certain or uncertain. To address this problem, we interpreted the differences between two tree structures as uncertainty and developed a new interactive visualization system, called CandidTree (see Fig. 1), to visualize the differences.

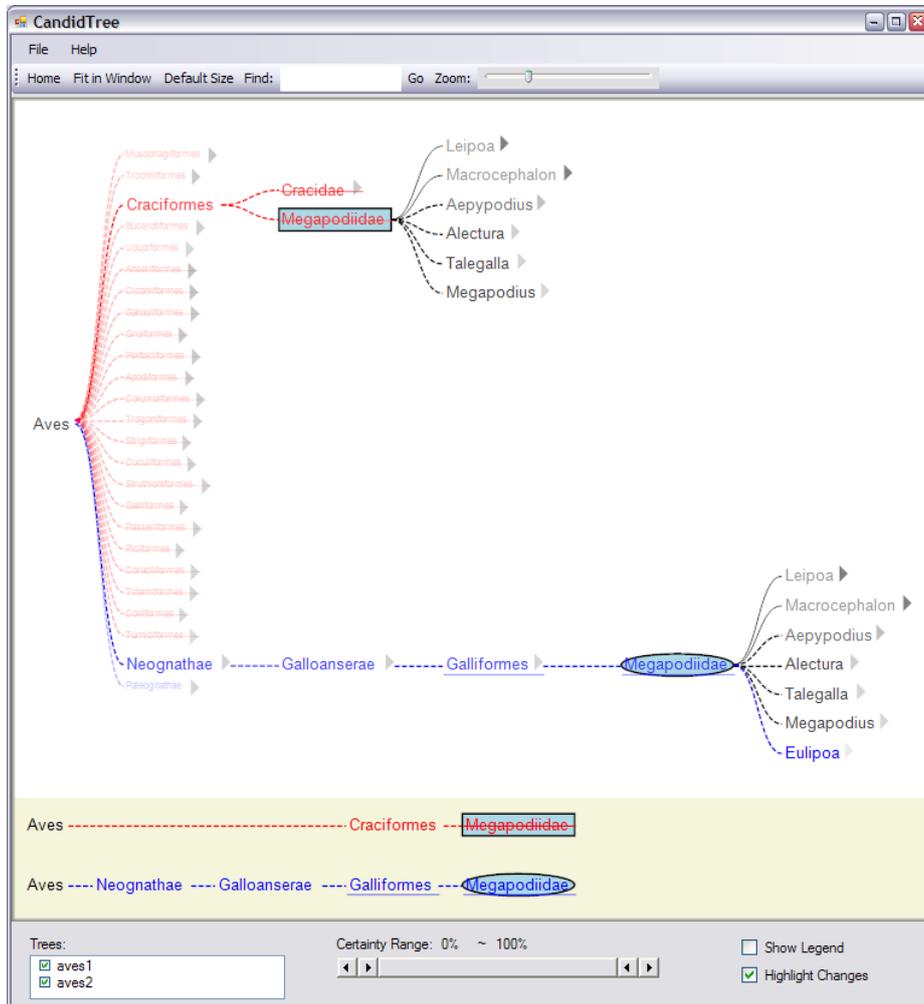


Fig. 1. CandidTree shows two types of structural uncertainty: location and sub-tree structure uncertainty, respectively, with color and transparency so that users can easily identify which parts are most certain or uncertain. CandidTree merges two trees into one and shows the merged tree in the tree browser at the top. The view of the paths at the bottom shows paths to the root in each tree from the currently selected node.

One of the most common approaches to comparing two tree structures is to use paired views side-by-side, using coupled interaction to allow users to compare and navigate two trees. This approach helps users identify where the differences are in two trees (usually by highlighting), but does not explicitly show the degree to which two parts are different.

CandidTree merges two tree structures into one and computes two types of structural uncertainty for each node: 1) location of a node relative to its parent and 2)

sub-tree structure of a node. CandidTree represents these uncertainties with color and with transparency, respectively, so that users can easily identify which parts are most certain or uncertain. It also enables users to interactively explore the merged tree to investigate those parts in more depth. Furthermore, when users select a node, CandidTree shows paths to the root in each tree so that they can see how its absolute location differs in the two trees.

While CandidTree was originally developed to show structural uncertainty, it can also be applied to visualize the differences between two tree structures. For example, when we have two (backup) directory structures (for two different time points) containing backups for the same folder, CandidTree can help users find added, deleted, or moved files, in addition to modified folders. Furthermore, it enables users to identify which folder has been changed the most or least.

After reviewing related work, we explain how these two types of the structural uncertainty are defined. We describe how CandidTree visualizes them using a set of two classifications of scientific names. We also report two user studies we conducted and then conclude with future work.

2 Related Work

To provide a complete and accurate visual representation of data, it is important to show uncertainty within the data. Uncertainty has been very broadly defined to include concepts such as error, inaccuracy/imprecision, minimum-maximum ranges, data quality, and missing data [21][23][26]. For more than a decade, much research has described approaches to handling these various aspects of uncertainties [2][3][20][25][26][28].

Geographic Visualization, Geographic Information Science, and Scientific Visualization communities have given particular attention to uncertainty visualization and many techniques have been developed [4][14][15][16][21][23][27]. The main techniques used to visualize uncertainty include adding glyphs [12][30], adding geometry, modifying geometry, modifying attributes, animation [6][12], sonification [13], and psycho-visual approach. While these techniques have been applied to a variety of applications such as fluid flow, surface interpolants, and volumetric rendering, only a few of them were actually evaluated. Furthermore, there has been very little research on visualizing uncertainty in tree structures. To our knowledge, only Griethe and Schumann proposed visual representations to represent uncertainty in parent-child relationship in structures [7]. For example, for node-link diagrams they used blurred or dotted links to indicate less certain relationships. However, they did not describe how to represent the degree of uncertainty. Moreover, while these authors brought up example applications for uncertainty visualization on structure information, their solutions were neither thoroughly investigated nor properly tested. In fact, it was beyond the scope of their paper to find effective metaphors in more challenging situations [7]. Since there is no formal definition of structural uncertainty, we propose two types of uncertainty for tree structures: location and sub-tree structure uncertainty, which will be explained in Section 3.

An error can be defined as a difference between a computed, estimated, or measured value and the true or correct value. There are many cases where we do not know correct values but can estimate those using different techniques or algorithms. It is common to use the differences between two results as an error (or uncertainty). For example, Pang and Freeman visualized differences between 3D surfaces generated by various rendering algorithms [22]. In fact, side-by-side comparison is one of the most commonly applied existing uncertainty visualization methods [23]. Therefore, theoretically we can use these kinds of visualization tools to show uncertainty in tree structures.

In the biology domain, there exist several biological taxonomies and phylogenetic trees because not all biologists agree on one taxonomy and one phylogenetic tree. In order to assess the quality of taxonomies and phylogenetic trees, it is important to understand which parts of two trees agree or disagree [19]. One of the commonly used approaches to comparing two trees is to use paired views side-by-side, using coupled interaction. In fact, many submissions to the InfoVis 2003 contest, Visualization and Pair Wise Comparison of Trees (<http://www.cs.umd.edu/hcil/iv04contest>), used side-by-side views. For example, TreeJuxtaposer automatically matches nodes in two trees based on the shared ancestors, and highlights where the differences are [18]. InfoZoom transforms a tree into a tabular representation, in which each leaf is represented as a column and the path from the root is stored in the attributes (rows) [24]. It displays both trees (in a tabular form) side-by-side and marks the cells of differences. Some visualizations provide a merged tree by combining two trees into one. For example, TaxoNote shows the merged tree on left, first tree at center, and second tree on right [17]. It uses multiple tables to provide taxonomic names that are common or different. Zoomology also provides a single overview of the merged tree with the indication of difference, and uses matched twin detail windows to show similarities and differences via a zoomable interface [8]. While these tools show where the differences are, they do not show the magnitude of the differences.

3 Structural Uncertainty

As for the cases where we do not know the correct tree structures, we interpret the differences between two tree structures as uncertainty. For each node, we measure two types of uncertainty; location and sub-tree structure.

3.1 Location Uncertainty

Within a tree structure, the location of a node can be represented in two different ways: 1) absolute path from the root and 2) relative path from its parent. The main drawback of the first representation is that a small difference close to the root would affect its whole sub-tree. Therefore, we decided to use the relative path to compute the uncertainty in node location. The location uncertainty is not scalar but categorical in value and three possible categories are as follows.

- 1) A node is in both trees at the same location (i.e., under same parent) – most certain
- 2) A node is in both trees but at different locations (i.e., under different parents)
- 3) A node is included in only one of the trees – most uncertain

3.2 Sub-tree Structure Uncertainty

Whether or not a node is in the same location in two trees, its sub-tree structures can be different. We compute the sub-tree structure uncertainty by measuring how many links overlap in two sub-trees. So, the sub-tree structure uncertainty function for a node v can be defined as follows.

$$\text{sub - tree structure uncertainty } y(v) = 1 - \frac{n(L_1(v) \cap L_2(v))}{n(L_1(v) \cup L_2(v))},$$

where $L_i(v)$ is the set of links in the sub - tree of v in the i th tree.

4 CandidTree

We developed a visualization system, named CandidTree, to show the structural uncertainty described in the previous section. CandidTree automatically merges two tree structures into one and computes uncertainty based on the differences between them. As shown in Fig. 1, it consists of two views: 1) a tree browser to show the merged tree and 2) a paths view to show paths to the root in each tree from the currently selected node.

To describe how CandidTree works in this chapter, we use a set of two classifications of scientific names of birds; one from the National Center for Biotechnology Information (NCBI, <http://www.ncbi.nlm.nih.gov>) and the other from the Animal Diversity Web (ADW, <http://animaldiversity.ummz.umich.edu>). The fact that these two authoritative sources disagree on these classifications illustrates the degree of uncertainty in biological classifications, which CandidTree helps expose.

4.1 Visualizing Structural Uncertainty

CandidTree shows the location uncertainty of a node relative to its parent by color (of the node label). The color black means that the node is included in both trees under the same node. Red and blue show that the node is included in the first and second tree, respectively. If the node is included in both trees but under different nodes, this means that the node moved from the first tree to the second one. To represent this move case, CandidTree shows the red node in the first tree with a strikethrough and the blue node in the second tree with an underline. For example, “Megapodiidae” was under “Craciformes” in the first tree and then moved under “Galliformes” in the second tree (Fig. 1).

CandidTree shows the sub-tree structure uncertainty of a node by transparency (of the node label). To make the node readable, even if uncertainty is very high, CandidTree uses 128 as a minimum alpha value (50% transparent). From the usability study, we learned that it is difficult to distinguish small differences (e.g., the difference between 1 and .9). To help users distinguish the 100% certain information from less certain data, CandidTree uses a solid link only when the sub-tree structure uncertainty is 0. For example, in Fig. 1, among the children of “Megapodiidae,” the sub-tree structure uncertainty of “Leipoa” and “Macrocephalon” is 0 and that of the others is non-0.

We also decided to use four discrete alpha values; 1) 255 (0% transparent) when u (uncertainty) = 0, 2) 214 ($\approx 16\%$) when $.5 \leq u < 1$, 3) 171 ($\approx 33\%$) when $0 < u < .5$, and 4) 128 (50%) when $u = 1$. To enable users to compare uncertainties in the same range, CandidTree provides exact values in a tooltip. The default set of alpha values means that the more certain the data, the more opaque (and readable) it is. However, users may want to focus on uncertain (different) information depending on data and tasks. CandidTree reverses the order of alpha values when users check the “Highlight Changes” check box (bottom right of Fig. 1) to make more uncertain information more readable.

Children of each node are grouped by location uncertainty; 1) nodes only in the first tree, 2) nodes moved from first tree, 3) nodes in both trees, 4) nodes moved to the second tree, and 5) nodes only in the second tree. This helps users capture only one tree from the merged tree. For example, if users want to focus on the first tree, they can ignore merged group 4) and 5). Within each group, children are sorted by sub-tree structure uncertainty.

4.2 Changes in Paths to the Root

When users click on a moved node (either red with strikethrough or blue with underline) in the tree browser, CandidTree finds the matching node in the other tree and opens them together. The selected node is indicated by a light blue background with a rectangular border and the matching node with an oval border. In the example of Fig. 1, once users click on “Megapodiidae” (in the first tree) under “Craciformes,” “Megapodiidae” (in the second tree) under “Galliformes” is automatically opened.

CandidTree’s tree browser and paths view are tightly coupled. So, when users select a node in the tree browser, paths to the root in both trees from the selected node are shown in the paths view. If the selected node is included in only one tree or its absolute location is the same in both trees, only one path is shown. When two paths are different, CandidTree vertically aligns the nodes with similar labels from two paths; the similarity of the labels is computed by Levenshtein distance [11]. This helps users see what changes are made between two trees. For example, two levels – “Neognathae” and “Neoaves” – are added to the second tree between “Aves” and “Strigiformes” (Fig. 2). This could also help users identify possible errors in node labels. For example, the parent of the “Musophagidae” is supposed to be “Musophagiformes” as in the first tree and “Musphagiformes” in the second tree is a typographical error (Fig. 3).

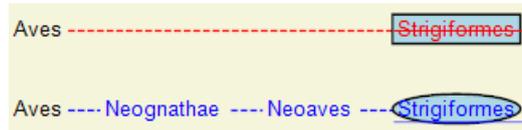


Fig. 2. CandidTree shows paths to the root in each tree from the selected node and its matching node (if exists) in the paths view.

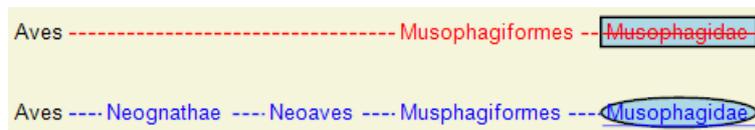


Fig. 3. CandidTree aligns nodes with similar labels from two paths. “Musophagiformes” is aligned with “Musphagiformes” because it is the most similar one among three nodes (“Neognathae,” “Neoaves,” and “Musphagiformes”) from the second tree.

When users click on a node in the paths view, CandidTree temporarily highlights the corresponding node in the tree browser with a thick purple rectangle surrounding the node so that users can recognize where the node is. In addition, if the corresponding node is off-screen, CandidTree pans the tree to view the node.

4.3 Emphasizing Nodes of Interest and Search

As users browse through the tree, especially for the trees with large fan outs, the selected node and its matching node could be very far away from each other and the matching node could be off-screen. Even when there is no matching node, the screen could be still cluttered. We assume that users are mainly interested in the selected node and its children, siblings, and ancestors. So, when users select a node, CandidTree uses a fisheye technique [5] and deemphasizes all other nodes by making them smaller and less opaque (Fig. 1). If it exists, the same rule is applied to the matching node of the selected node.

For the cases where users are only interested in the nodes with a specific certainty range, we provide a double-headed range bar (bottom center of Fig. 1) in the control panel to allow users to un-highlight nodes that do not meet the certainty requirement. Users can focus on one tree by manipulating the check boxes in the Trees list (bottom left of Fig. 1); the unchecked tree is un-highlighted. Since 128 (50% transparent) was used as an alpha value for the most uncertain nodes for readability, CandidTree uses 64 as an alpha value (75% transparent) for un-highlighted nodes.

CandidTree provides support for search, providing simple substring match with node labels. Typing a word and pressing the “Go” button displays the search results colored in orange and restricts the view to the nodes relevant to the search results (Fig. 4).

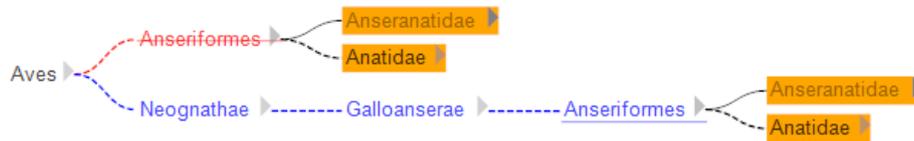


Fig. 4. A search for “Anatidae” shows four search results (containing the keyword) and nodes relevant to them.

4.4 Implementation Details

CandidTree is implemented in C# with Piccolo.NET, a shared source toolkit that supports scalable structured 2D graphics [1] (<http://www.cs.umd.edu/hcil/piccolo>). It uses a classical tree layout by Walker [29].

CandidTree reads data from files in an xml format. Each node is represented with a “node” element having two attributes; “id” and “name.” The id attribute, which should be unique, serves as an identifier to be used to match nodes in two trees. The name attribute is used as a label of the node. The current implementation can be extended to show other node attributes or handle other data formats.

CandidTree builds the merged tree in memory and computes uncertainty at startup. It loads the first tree and then merges the second tree with the first one. The location uncertainty is computed during this merge process. Once the merged tree is built, CandidTree computes the sub-tree structure uncertainty using the equation in Section 3.2 after recursively counting the number of links in the sub-trees of each node.

5 Evaluation

5.1 Study 1: Usability Study with Computer Scientists

To identify major usability issues, we conducted a preliminary usability study with six participants: two researchers, three developers, one research intern (all male computer scientists). We used two sets of backup directory structures, each with two trees representing the data at two different time points. Trees in these sets contained about 150 nodes and 500 nodes.

Participants were given a brief tutorial of the system for up to 15 minutes including the time to play with the system and ask questions. Next, they were asked to perform 10 tasks (1-5 for the small tree and 6-10 for the large tree) with the system, which were timed and scored for correctness. The tasks were meant to cover major tree-comparison tasks and to evaluate the usability of main features to represent structural uncertainty. All participants were asked to perform the tasks as quickly and accurately as possible. Once they completed all of the tasks, participants were asked to fill out a satisfaction questionnaire. Each session lasted about 30 minutes and participants were given a \$5 snack coupon for their participation. The list of tasks follows.

- 1) How many files were deleted from the “DynaVis” folder?
- 2) Among the sub-folders of “DynaVis,” which folder has been changed the most? (We asked users to ignore added or deleted folders.)
- 3) Describe the changes made to the “DynaVis\DynaTestWin\obj” folder.
- 4) Which file was moved from the “DynaVis\DynaVis\obj\Debug” folder?
- 5) Where did the file move to?

Tasks 6-10 were equivalent to tasks 1-5 but applied to a large tree.

There were only 6 incorrect answers provided out of 60 questions across participants. Three participants answered all questions right and the other three each incorrectly answered one, two, and three questions, respectively. Task 2 (and 7 for the large tree) got the most wrong answers because participants had difficulty distinguishing between the colors black and green (green was the color of the second tree at the time of experiment). Two participants each gave wrong answers to Task 4 and 9, respectively. Overall, average task times were fast. While Task 8 took longer than the others (46.5 seconds on average), it was because it takes time to describe all the changes not because participants had a hard time finding the changes.

Table 1 below shows the average satisfaction ratings on a 7-point Likert scale, with 1=Disagree and 7=Agree. There was clear user frustration around the use of transparency to represent the sub-tree structure change. This is related to the readability issue participants raised. The ratings are fairly consistent with the usability issues we identified. We here summarize the major iterations made to CandidTree based on the first user study.

- 1) Do not use the fourth color for the moved nodes; instead use strikethrough and underline.
- 2) Use solid links only when the uncertainty is 0; otherwise dotted links.
- 3) Use four discrete alpha values.
- 4) Provide an option to highlight changed (uncertain) information.
- 5) Provide a legend to show the color scheme.

Table 1. Average Likert scale ratings for CandidTree, using the scale of 1=Disagree, 7=Agree. Study 1 refers to the Usability Study with Computer Scientists and Study 2 refers to the Qualitative User Study with Biologists described in the next section.

	Study 1	Study 2
Overall, the system was easy to use	5.0	5.7
I felt comfortable using this system	5.3	5.7
It was easy to learn to use this system	5.2	5.7
It was easy to navigate through this system	6.0	5.7
It was easy to read the labels of the nodes	4.7	4.6
Colors representing the node location change was clear	4.2	4.5
Transparency representing the sub-tree structure change was clear	3.5	3.0

5.2 Study 2: Qualitative User Study with Biologists

We conducted a qualitative study with biologists who regularly work with hierarchically organized names, with two main goals. First, are these users able to use CandidTree to correctly and quickly complete tree-comparison tasks, and which tasks pose more difficulty for users? Second, does CandidTree support advanced information understanding and insight? Gaining insight from data does not lend itself easily to the metrics typically used in quantitative studies. Finally, we were interested in an assessment of the “uncertainty” metric we used.

Participants. We recruited 8 volunteers (3 females and 5 males, 28 to 59 years old) from the Smithsonian Institution and University of Maryland. They included two graduate students, two post-doctoral fellows, and four curators or research faculty. All were unfamiliar with the testing datasets, though all had previously used data from the same source before. Two were familiar with the tutorial data. Three mentioned regularly working with datasets of more than 150 terminal taxa (leaves); typical datasets include between 40 to 150 leaves (median 143). However, several of the biologists were associated with NSF projects dealing with trees of names of thousands of organisms and two participants mentioned that CandidTree might be useful for those projects. Each participant was given a \$20 Amazon.com gift certificate for his/her participation.

Datasets. For both demonstration and testing we used classifications of scientific names from the National Center for Biotechnology Information (<http://www.ncbi.nlm.nih.gov>) that were downloaded on different dates: December 14, 2005 and September 19, 2006. Tutorial data were from the Lepidoptera branch (moths and butterflies; tree 1=4103 nodes, tree 2=6262 nodes) while the test data were from the Aves branch (birds; tree 1=6912 nodes, tree 2=8140 nodes).

Procedure. Each participant filled out a background survey. They received a demonstration of CandidTree features and were told to freely explore and ask questions, for a total tutorial time of up to 20 minutes. The search feature was not described or tested, nor were default settings for transparency or filtering changed or described. Nodes that were darker were those with higher uncertainty scores. Participants who asked were told how uncertainty was calculated. Participants were then asked to perform eleven tasks, described below. Participants then completed the same preference survey as in the usability study (Table 1). We videotaped the computer screen throughout the tutorial and testing. Each session lasted 45 to 60 minutes.

Tasks. Biologically meaningful tasks (Table 2) were chosen based on 30-minute interviews with three biologists (one of whom was subsequently a participant). They were presented in order of increasing complexity. The first nine had single, correct answers. Task 10 was judged by the number of insights given by the participant, and Task 11 was an opinion.

Table 2. Biologist user study tasks, with results from eight participants.

Task	# Missed	Mean duration (s)
1. Which branch (child) of the Neognathae is the most uncertain, from tree 1 to tree 2?	0	73
2. In the children of Bucerotidae, are the changes additions, deletions, or changes in the placement of taxa?	1	28
3. How much growth (in number of new taxa) has Lampornis experienced from tree 1 to tree 2?	0	19
4. Of the children of Passeriformes, name all the taxa that are placed differently in the different trees.	3	31
5. For Estrildidae, describe the difference in the placement.	0	29
6. The genus Agelaius moved from one group to another. Did all of the other genera in its group move too?	1	51
7. Was the parent of Icterus promoted or demoted in rank?	2	104
8. What happened to the family Cracticidae?	1	57
9. Overall, how would you characterize the uncertainty or change in the tree: changes in deep relationships, mid-tree relationships, or among terminal taxa?	0	137
10. Summarize the changes to Furnariidae and to Sylviidae.	--	275
11. In your opinion, which of the two groups Furnariidae and Sylviidae is the most uncertain?	--	73

Results. Overall ease of use improved slightly over the usability study (Table 1). Still, transparency as an indicator of the sub-tree structure change scored particularly low. “Color representing the node location change was clear” and “It was easy to read the labels of the nodes” also received relatively low scores.

Of 72 possible answers (8 participants x 9 tasks), only 8 were incorrect (Table 2). The oldest participant had the most difficulty, answering incorrectly for 3 out of the 5 most complex tasks. Task 4, which required understanding the coding of location uncertainty, was missed most often. Task 7, which required understanding whether the location change was up or down a level, was missed by two participants but they both admitted a lack of concern about the utility of biological ranks.

Coding of location uncertainty continued to be problematic despite the iteration and the addition of a legend. Several participants answered question 4 incorrectly at first and then corrected themselves after glancing at the legend (they were counted as correct answers). Those who missed it gave all red names as answers instead of just the names in red with strikethrough. One participant remarked that she expected the two red codes to be additive: “It implies a hierarchy but in fact they signify different rather than nested ideas.”

Average times to complete biologically relevant tasks were somewhat longer than in the usability study (Table 2). Some participants took 1 to 2 minutes for Task 1, either to orient to the testing protocol (giving answers verbally) or to the sort order of uncertainty. Otherwise, the simplest tasks each typically took less than 30 seconds. For task 10, most participants explored for 3 to 6 minutes before being satisfied they had given a good summary.

For task 10, participants reported 3 insights (the participant with the least domain expertise who spent only 110 seconds exploring) to 9 insights (participant 4, who

systematically explored for 6 minutes). Overall, approximately 18 unique insights were given. Five of the 8 participants found a typographical error correction that required a correct interpretation of transparency, and 3 found a subtle location change that required understanding of location coding and the paths view. The example in Box 1 shows how Participant 4 built his insights in steps while exploring CandidTree. Participants also reported some insights that were unprompted during both the tutorial and test. For example, participants mentioned that the datasets were obviously large, had few changes, were different from a dataset they were familiar with, and had a sub-tree with lots of “problem children.”

Box 1. Participant 4 used CandidTree to make the following insights. Insight 9 builds on 6 and 7 which build on 5 which builds on 4.

Furnariidae sub-tree

1. Immediately identifies that there has been a spelling change between tree 1 and [requires opening the least certain sub-tree]
2. Counts 17 new genus-level nodes added
3. Finds in a sub-tree that there has been a new species added

Sylviidae sub-tree

4. Three nodes in tree 2 are not in tree 1
5. Two of these are entirely new subfamilies [judged by interpreting the label]
6. They contain both taxa that have been moved here from the first tree
7. and also include some taxa entirely new to the second tree
8. Elsewhere, some new subspecies have been added that were not in the first tree
9. Says, “Basically, taxa which had not previously been in subfamilies [a particular rank in the hierarchy] were moved into Acrocephalinae”

For task 11, six of 8 participants that sub-tree Sylviidae was more uncertain than the other. However, though the certainty scores were very close, half of the participants thought Sylviidae was much more uncertain because the changes involved numerous rearrangements and addition of internal nodes rather than simple addition of nodes at the leaves. Two participants thought it was not reasonable to compare them because the kinds of uncertainty were so different.

Consistent with the preference survey, the most common complaint was that transparency differences were too subtle to be usable. Also, two participants thought the more transparent names should be those with less sub-tree structure certainty. However, generally high task performance shows that these issues did not pose significant problems. Two participants thought that the colors should be reversed (red should represent the second, more recent or important tree). Several had trouble managing opening and then closing sub-trees and two suggested it would be useful to have a way to open or close one level of children all at once across the whole tree.

Participants offered many ideas for additional features or applications. Three participants wanted to use CandidTree to compare more than two trees, particularly with particular scientific datasets and websites such as NCBI or Tree of Life (<http://www.tolweb.org>). Some expressed interest in linking nodes to further information such as the GenBank sequence or host plants.

6 Conclusion and Future Work

We proposed two types of uncertainty for tree structures – location and sub-tree structure uncertainty – based on the differences between them. To visualize those structural uncertainties, a new interactive visualization system called CandidTree was developed. Since CandidTree computes uncertainty by comparing two trees, we were able to apply it to visualize the differences between them. For example, CandidTree helps users find added, deleted, or moved files as well as modified folders within two (backup) directory structures containing backups for the same folder.

We conducted two user studies to identify major usability issues and evaluate how our system works. Our qualitative study with biologists showed that while we have improved the uncertainty representation so that task performance and insight-building is high even with large trees, ways to improve satisfaction are needed. Also, while most users concur with relative uncertainty scores, there is not universal agreement on how to weigh the kinds of uncertainty. We are planning to conduct a controlled experiment to compare CandidTree with a traditional files and folders comparison tool to see whether users could perform better with CandidTree.

While the current implementation works only for two trees, we can handle more than two trees by providing the list of possible combinations of multiple trees and showing only one combination at a time. By ranking each combination based on the sub-tree structure uncertainty of the root node, we could enable users to easily identify most certain/uncertain (similar/dissimilar) tree combinations. As mentioned before, CandidTree loads each entire tree and builds the merged tree in memory, which is impractical for large trees. We can preprocess building of the merged tree and the computing structural uncertainty, and store them in a database. By accessing the data from a database when needed, CandidTree can be scaled up to support very large trees and with multiple attributes.

Acknowledgments. We would like to thank the participants of our two user studies for their participation and comments. Charlie Mitter and Ashleigh Smythe helped define biologist tasks as did Nathan Edwards who also provided NCBI data. Danyel Fisher reviewed our paper and gave thoughtful comments.

References

1. Bederson, B.B., Grosjean, J., and Meyer, J.: Toolkit Design for Interactive Structured Graphics, *IEEE Trans. on Software Engineering* 30(8), (2004) 535-546.
2. Cleveland, W.S.: *The Elements of Graphing Data*, (1985).
3. Eaton, C., Plaisant, C., and Drizd, T.: Visualizing Missing Data: Graph Interpretation User Study, *Proc. of Interact 2005, Lecture Notes in Computer Science*, Vol. 3585. Springer-Verlag, Berlin Heidelberg New York, (2005) 851-872.
4. Fegeas, R.G., Cascio, J.L., and Lazar, R.A.: An Overview of FIPS 173, *The Spatial Data Transfer Standard, Cartography and Geographic Information Systems* 19(5), (1992) 278-293.
5. Furnas, G.W.: Generalized Fisheye Views, *Proc. of CHI 1986*, (1986) 16-23.

6. Gershon, N.D.: Visualization of Fuzzy Data using Generalized Animation, Proc. of VIS 1992, (1992) 268-273.
7. Griethe, H. and Schumann, H.: The Visualization of Uncertain Data: Methods and Problems, Proc. of SimVis 2006, (2006) 143-156.
8. Hong, J.Y., D'Andries, J., Richman, M., and Westfall, M.: Zoomology: Comparing Two Large Hierarchical Trees, Posters Compendium of InfoVis 2003, (2003) 120-121.
9. Hyperbolic Tree for the Green Tree of Life,
<http://ucjeps.berkeley.edu/TreeofLife/hyperbolic.php>
10. Lee, B., Parr, C.S., Campbell, D., and Bederson, B.B.: How Users Interact with Biodiversity Information using TaxonTree, Proc. of AVI 2004, (2004) 320-327.
11. Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions, and reversals, Doklady Akademii Nauk SSSR 163(4), (1965) 845-848 (Russian). English translation in Soviet Physics Doklady 10(8), (1966) 707-710.
12. Lodha, S.K., Pang, A., Sheehan, R.E., and Wittenbrink, C.M.: UFLOW: Visualizing Uncertainty in Fluid Flow, Proc. of VIS 1996, (1996) 249-255.
13. Lodha, S.K., Wilson, C.M., and Sheehan, R.E.: LISTEN: Sounding Uncertainty Visualization, Proc. of VIS 1996, (1996) 189-196.
14. MacEachren, A.M., Robinson, A., Hopper, S., Gardner, S., Murray, R., Gahegan, M., and Hetzler, E.: Visualizing Geospatial Information Uncertainty: What We Know and What We Need to Know, Cartography and Geographic Information Science 32, (2005) 139-160.
15. Moellering, H.: Continuing Research Needs Resulting from the SDTS Development Effort, Cartography and Geographic Information Systems 21(3), (1994) 180-189.
16. Morrison, J.: The Proposed Standard for Digital Cartographic Data, American Cartographer 15(1), (1988) 9-140.
17. Morse, D.R., Ytow, N., Roberts, D.M., and Sato, A.: Comparison of Multiple Taxonomic Hierarchies Using TaxoNote, Posters Compendium of InfoVis 2003, (2003) 126-127.
18. Munzner, T., Guimbretiere, F., Tasiran, S., Zhang, L., and Zhou Y.: TreeJuxtaposer: Scalable Tree Comparison using Focus+Context with Guaranteed Visibility, Proc. of SIGGRAPH 2003, published as ACM Transactions on Graphics 22(3), (2003) 453-462.
19. Nye, T.M.W., Lio, P., and Gilks, W.R.: A Novel Algorithm and Web-based Tool for Comparing Two Alternative Phylogenetic Trees, Bioinformatics 22(1), (2006) 117-119.
20. Olston, C. and Mackinlay, J.D.: Visualizing Data with Bounded Uncertainty, Proc. of InfoVis 2002, (2002) 37-40.
21. Pang, A.: Visualizing Uncertainty in Geo-Spatial Data, Proc. of the Workshop on the Intersections between Geospatial Information and Information Technology.
22. Pang, A. and Freeman, A.: Methods for Comparing 3D surface Attributes, Proc. of SPIE-VDA 1996, (1996) 58-64.
23. Pang, A.T., Wittenbrink, C.M., and Lodha, S.K.: Approaches to Uncertainty Visualization, The Visual Computer 13(8), (1997) 370-390.
24. Spenke, M. and Beilken, C.: Visualization of Trees as Highly Compressed Tables with InfoZoom, Posters Compendium of InfoVis 2003, (2003) 122-123.
25. Sulo, R., Eick, S., and Grossman, R.: DaVis: A Tool for Visualizing Data Quality, Posters Compendium of InfoVis 2005, (2005).
26. Taylor, B.N. and Kuyatt C.E.: Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results, NIST Technical Note 1297, (1994).
27. Thomson, J., Hetzler, B., MacEachren, A., Gahegan, M., and Pavel, M.: A Typology for Visualizing Uncertainty, Proc. VDA 2005, (2005) 146-157.
28. Tukey, J.W.: Exploratory Data Analysis, (1977).
29. Walker II, Q.: A Node-Positioning Algorithm for General Trees, Software-Practice and Experience 20(7), (1990) 685-705.
30. Wittenbrink, C.M, Pang, A.T., and Lodha, S.K.: Glyphs for Visualizing Uncertainty in Vector Fields, IEEE Trans. on Visualization and Computer Graphics 2(3), (1996) 266-279.