Semantic Wiki Visualization

Matt Gedigian

Abstract—
Semantic wikis are wikis which incorporate machine-readable semantic structure. The data they expose can be visualized for many different purposes using various tools. One particularly interesting use of visualization is to provide users with an overview of the wiki contents. In this work we develop a variant of radial space filling trees which employ a novel type of interaction to help users manage constrained spaces. An application to an existing wiki is shown.

Index Terms—Ontology, wiki, semantic similarity, spacing filling.

1 Introduction

There has been much interest recently in attempts to unite Semantic Web technologies with the design patterns collectively known as Web 2.0. One product of these efforts is a new class of applications known as ”semantic wikis”. Semantic wikis extend the simple, community oriented editing process of traditional wikis with specialized handling for formally representable semantics [2].

While current semantic wikis have exploited their formal structure to provide powerful semantic search capabilities, they have not changed the way that readers move through the wiki. Specifically, there is one shortcoming of wikis that they have not rectified, namely that it is easy for users to get lost [3]. Users view a single article at a time with no context other than a link to the enclosing category for the current article. In Wikipedia, this problem is partially mitigated by having articles specifically written to provide overviews (for example pages which contain a “List of ...”). However this is a manual process which is limited, time-consuming, and error prone.

To provide maximum benefit, the overview should not be presented as a separate page that users must navigate to and from. Instead, it should use the ”Overview Plus Detail” user interface design pattern [6]. The simultaneous display of overview and details provides ongoing feedback that enables a user to understand where they are in the wiki. The overview also enables a type of rapid navigation that isn’t possible using normal wiki-links. In particular, via an overview, a user can jump directly to non-adjacent pages.

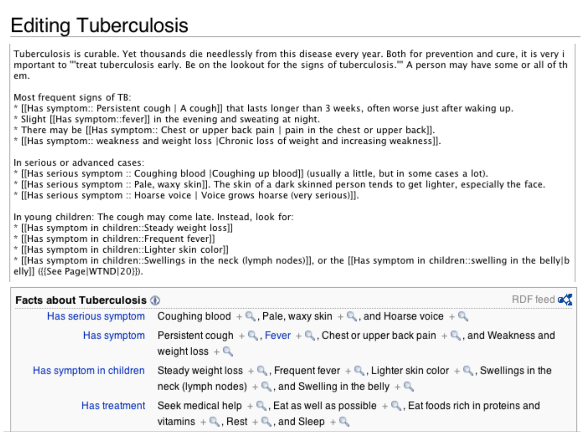

Fig. 1. Markup used on a semantic wiki and the formal properties it produces

2 Related Work

A number of visualization techniques have been developed to depict hierarchies. Below we survey several of these and discuss their applicability before focusing on the most closely related work.

2.1 Text Based Trees

One simple method of depicting hierarchical relationships is a textual tree view. Each item is displayed on its own row, with children indented relative to their parent. This method requires large amounts of vertical space to display nodes and may require vertical scrolling to see different parts of larger lists. It is not good at conveying the *overall* structure of large collections. Each sibling is separated from the preceding sibling by all of that nodes’ children. As a result nodes that have many descendents are displayed far from their siblings. As the vertical spacing between nodes grows, it becomes increasingly difficult to determine their relationship.

A textual tree layout of a small animal ontology is shown in Figure 2. In this space, only 16 rows are visible. Even with this small number of items, it becomes difficult to compare the relative depths of “Koala” and “Rainforest”. Also notice how the descendants of “Animal” increase the spread between “Animal” and its sibling node “Degree”.

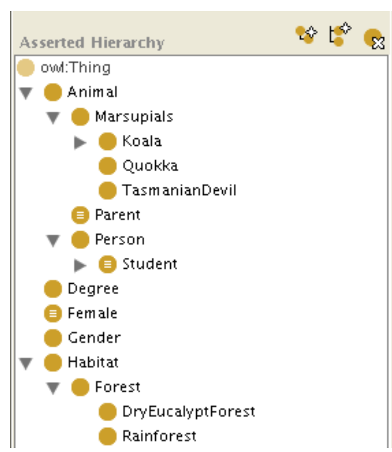

Fig. 2. Text Based Tree view of koala.owl ontology in Protegé

2.2 Node-Edge Layouts

Node-edge layouts are extremely general and can used to represent all sorts of graphs in addition to trees. These use edges to represent the parent-child relationship between nodes. While these are effective at representing schemas, there is not a clear distinction between classes and instances. Additionally, they require empty space separating the nodes.

• *Matt Gedigian, U.C. Berkeley School of Information, E-mail: matt_g@ischool.berkeley.edu.*

Submitted 15 December 2008

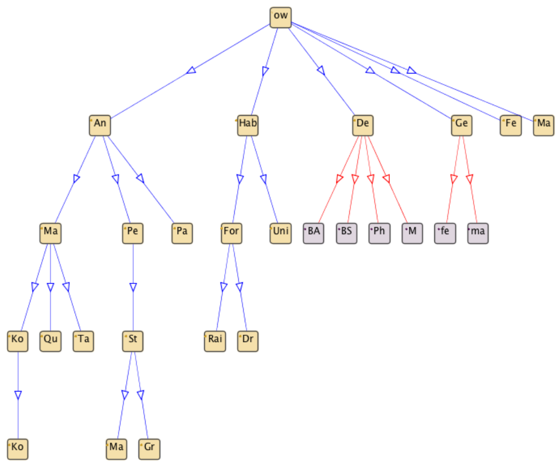

Fig. 3. Node Edge Layout of koala.owl in Jambalaya

2.3 Treemaps

It is not necessary to use edges to represent the relationships. Instead of using edges, the hierarchy can be depicted using containment, geometry, and positioning. This gives rise to space-filling layouts. Treemaps consist of rectangles which alternatively divide horizontal and vertical space as they descend the hierarchy [4]. They are extremely space efficient and can represent a quantitative variable associated with each instance using the area of that node. However, while the relationship between a parent and its descendents is clear, other relationships can be difficult to read.

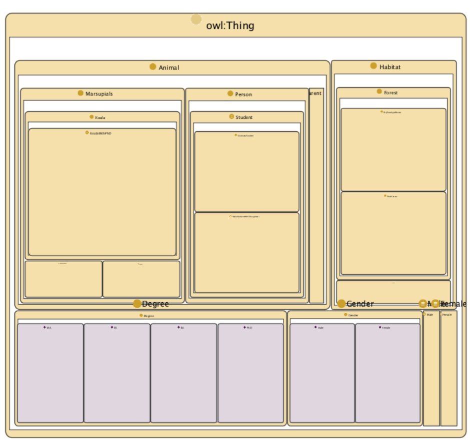

Fig. 4. Treemap Layout of koala.owl ontology in Jambalaya

2.4 Radial Space Filling Trees

Radial space filling trees were developed to present hierarchies more clearly than treemaps [1]. One problem with sunbursts is that as the hierarchy grows in size, deeper items become small and can be difficult to discern. Various approaches have been used to address this by allowing users to selectively focus on a subtree – magnifying it to show detail at a particular level [5].

It is usually acknowledged that radial space filling trees are not as space efficient as treemaps. For instance in a square viewing area of side s (and area s^2), the largest inscribed circle has diameter s, (and area: $s^2\frac{\pi}{4}$), resulting in more than 20% empty space. However, it is important to note that in a rectangular space (such as one would use for a navigation bar on the periphery of a web page) the effect is even more pronounced. For a rectangle s by $4s$, the largest inscribed circle still has diameter s, resulting in more than 80% empty space.

3 Methods

The visualization we use is a variant of a radial space filling tree. Categories and individuals are drawn as wedges. Distance from the center of the sunburst correponds to depth in the hierarchy. Children are arranged to span the angle of their parent node's wedge. Categories are given a uniform gray color while instances are colored according to the class to which they belong. When the mouse hovers over a node, it is highlighted and a tooltip displays its name and parent category.

To enable users to focus on a portion of the hierarchy, we enable a novel interaction. Users can drag individuals subtrees from a sunburst to break them off. Upon release they expand into independent sunbursts. Users can pan around in the visualization and reposition the sunbursts to accomodate the constrained space.

The visualization is presented in a thin html iframe above the wiki. Shift-clicking on a node loads the corresponding article in the wiki. The node corresponding to the article currently being viewed on the wiki is highlighted and is updated as the user navigates among pages using standard wiki-links.

Each node is scaled in proportion to the amount of information about it that is contained in the ontology. That is, the angle of each instance node is proportional to the number of statements about it in the ontology. The full 2π distance of the central sunburst is divided evenly among the ontology statements about the instances within it.

Although it is conventional to use the angle of a sunburst to encode a quantity, this is somewhat problematic. The angle and depth together determine a two-dimensional area. Since wedges are enclosed and shaded, it is natural for viewers to judge them by area rather than angle. Area comparisons are valid only for nodes at the same depth of the hierarchy. For example the area of an annulus with $r = 1$ and $R = 2$ is 3π while that of an annulus with $r = 2$ and $R = 3$ is 5π.

4 Results

To demonstrate these techniques, they were applied to the Hesperian Digital Library, a semantic wiki containing information on rural healthcare. The "Disease" section of the ontology contains 13 categories and 116 instances. The resulting visualization is shown in Figure 5.

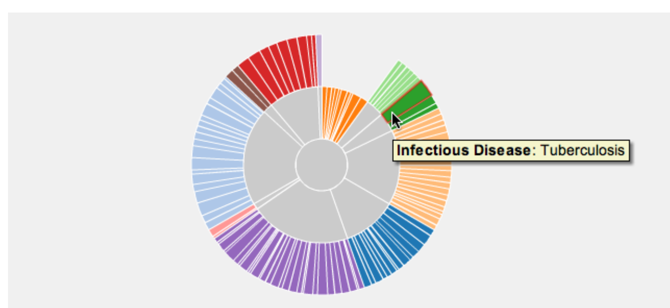

Fig. 5. The hierarchy of Diseases from the Hesperian Digital Library

At the center of the sunburst is the "Disease" node itself. Between the 12 o'clock and 1 o'clock positions, at depth 1 and colored in orange, are nodes which are direct instances of the disease class itself. All other instances are at depth 2. The mouse is hovering over "Tuberculosis", an instance of the "Infectious Disease" class.

Notice the large amount of empty space on each side of the sunburst shown in Figure 5. If the user wishes to focus on the "Common Sickness" subtree at the left, they can grab it, drag it, and release where ever they want. This interaction is shown in Figure 6. Upon release, the wedge expands to become a new sunburst, thereby increasing the size of its child nodes. Notice that in part 3 of the figure, it is easier to discern that "Cough" is one of the larger "Common Sickness" nodes.

This indicates that there are a large number of statements in the ontology about "Cough", relative to some of its thinner siblings. Also notice that the amount of blank space in the layout has been reduced.

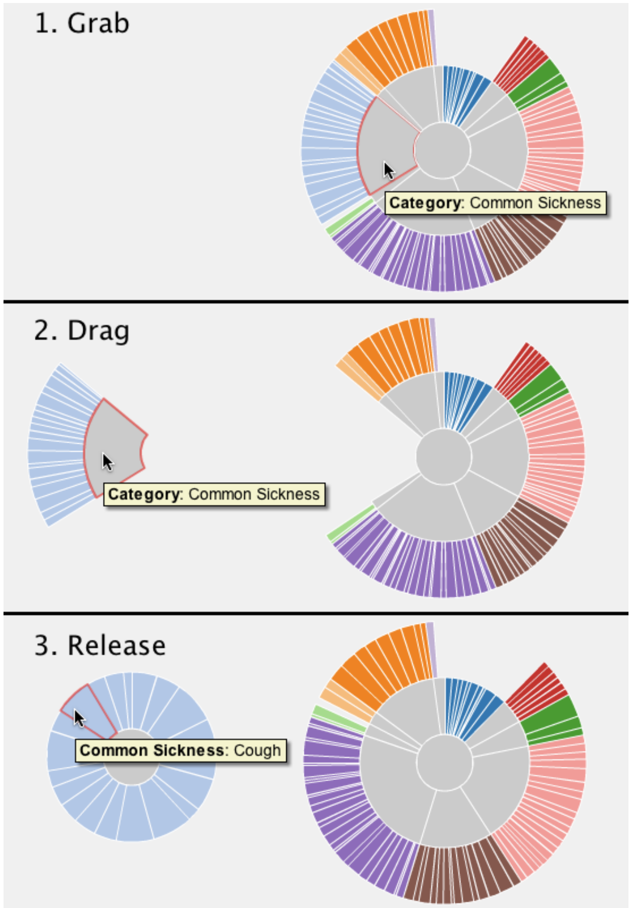

Fig. 6. Dragging a subtree

The preceding figures have shown the visualization in isolation, but in practice the link to the wiki is essential. Figure 7 shows the visualization in context. After following a wiki-link to "Children's Diseases", the visualization is updated to highlight (in red) the currently displayed article.

5 Discussion

This work demonstrates that visualization can be used to provide a structural overview of the contents of a semantic wiki. It shows how structured information can be conveyed using the size and arrangement of nodes. Enabling users to create multiple sunbursts affords them greater control over managing the constrained space. Multiple sunbursts are effective at reducing the amount of wasted space in an elongated rectangular area.

While other sunburst variants have allowed users to focus on a subtree, these have been limited to a single subtree at a time. When the users focuses on one area, the previous selection is lost. In contrast, this approach allows the user to focus on multiple areas by spawning multiple sunbursts simultaneously. This flexible approach is well-suited to handle non-square bounds while single sunburst layouts can be extremely space inefficient.

6 Future Work

One direction of future work would be to improve the communication between the wiki and visualization. In this work, the entire ontology is downloaded and processed to create a GraphML file. While the process can be scheduled to run at regular intervals, it would be better

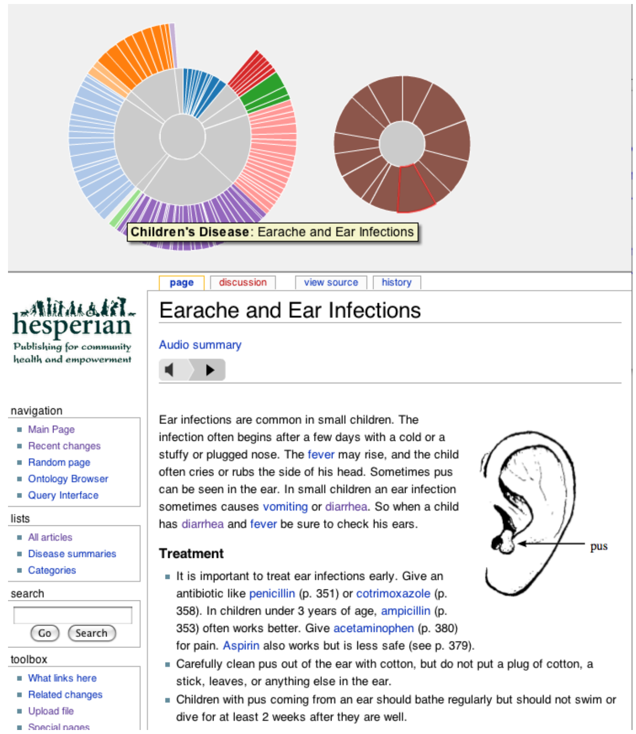

Fig. 7. Navigating the wiki

to be able to update the visualization in realtime as edits are made to the wiki.

Another style of visualization that would be well-suited to this application would be a Degree of Interest (DOI) Tree [**?**]. In this application, there is a natural point-of-interest at the node corresponding to the currently viewed article. The degree of interest is modelled as a function that assigns a score to other nodes in the tree. Several such functions could be developed which use properties from the ontology. The user could choose between such functions. Since changing interest scores only affects whether nodes are expanded, they produce trigger dramatic changes in node position which might be confusing.

Another direction is to integrate visualizations with the semantic query system of the wiki. Presently there is support for displaying query results rendered as HTML lists and tables. This could be extended and combined with a data driven visualization system. This would allow members of the community to design and share visualizations of the site content (rather than just administrators).

Acknowledgements

The author wishes to thank the Hesperian Foundation for the access to the material from the books and the contributors to the Hesperian Digital Library for their work in building the ontology used in this project.

References

[1] K. Andrews and H. Heidegger. Information slices: Visualising and exploring large hierarchy using cascading, semi-circular discs. In *CIMS'07: Proceedings of IEEE Information Visualization (Late Breaking Hot Topics Procs)*, pages 9–12. IEEE Computer Society, Oct. 1998.

[2] A. Ankolekar, M. Krötzsch, T. Tran, and D. Vrandecic. The two cultures: mashing up web 2.0 and the semantic web. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 825–834, New York, NY, USA, 2007. ACM Press.

[3] D. Aumüller. Shawn: Structure helps a wiki navigate. In *Proceedings of the BTW-Workshop "WebDB Meets IR"*, Mar. 2005.

[4] D. Aumüller and S. Auer. Towards a semantic wiki experience - desktop integration and interactivity in wiksar, 2005.

[5] J. Heer and S. K. Card. Doitrees revisited: scalable, space-constrained visualization of hierarchical data. In *AVI '04: Proceedings of the working conference on Advanced visual interfaces*, pages 421–424, New York, NY, USA, 2004. ACM.

[6] A. P. S. Leonidas Deligiannidis, Krys J. Kochut. Rdf data exploration and visualization. In *CIMS'07: Proceedings of the ACM first Workshop on CyberInfrastructure: Information Management in eScience*, pages 39–46, Nov. 2007.

[7] C. Sayers and F. Graph. Rdf, visualization, resource description, 2004.

[8] B. Shneiderman. Tree visualization with tree-maps: 2-d space-filling approach. *ACM Trans. Graph.*, 11(1):92–99, January 1992.

[9] J. Stasko and E. Zhang. Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *Proceedings of IEEE Information Visualization*, pages 57–65. IEEE Computer Society, 2000.

[10] J. Tidwell. *Designing Interfaces : Patterns for Effective Interaction Design*. O'Reilly Media, Inc., November 2005.