

single synset, all nodes are colored blue.

Document content is visualized through the alpha transparency of the fill color of the nodes. Highly opaque nodes have many occurrences; almost white (transparent) nodes have none. Words and senses that are more prominent in the document of interest stand out easily against a desaturated context. Two views are provided: In the *single-node view* (Fig. 1), only nodes whose word members occur in the document are colored, revealing the precise concepts in the document. In the *subtree view* (Fig. 2), counts are propagated up to the root of the tree, using the association of synonyms into synsets and synsets into a hyponymy tree to aggregate counts for related concepts and provide a higher level view of document content. In both views, alpha is normalized so maximum counts achieve full opacity. The subtree view reveals where concepts occur when a fisheye filter has been applied to the graph and some nodes are hidden. In addition, for a dense graph, the subtree view helps us find single nodes that are too small to distinguish.

4 Interaction

We provide several techniques to visually abstract the data. First, we provide a highlight search function which visually highlights nodes whose label matches any of the given search terms. *Highlight nodes* have a pink background and a larger font size, drawing attention to even the smallest of search results. Second, we implement generalized fisheye views [4] to collapse all subtrees which are more than a user-specified distance from the central root node. The presence of non-zero word occurrence counts within collapsed subtrees is indicated by using the subtree coloring, in which counts are propagated to the root. Optionally, all highlight nodes can be exempted from the distance filter (by increasing their *a priori* importance in the DOI function), effectively abstracting the graph to all synsets within a given distance from the root or highlight nodes.

Double clicking on a node of interest restricts the visualization to the hyponyms of the corresponding synset; double right-clicking reloads the graph at the parent of the clicked node, thus providing bi-directional semantic zoom. Node angular width can be manually adjusted using the mouse wheel to increase (up) and decrease (down) the width of the node. Changes to a node's angular width affect its children equally and its siblings in an inverse manner [7].

The counts that are used to determine node alpha values are based on text segments, or automatically-determined subtopic regions of the document. The initial view is based on all segments in the document, but range selectors allow for limiting the segments from which counts are drawn (for example, an analyst may be interested in only the first half of a document).

Visual pan and zoom of the display space are also supported. Word nodes can be shown or hidden to increase detail or decrease clutter as desired. Highlighting, roll-up, fisheye filtering, pan, and zoom are provided in real time.

Because edges are not shown, it may be difficult to discern parent-child relationships using DocuBurst, especially if the graph is particularly dense. To facilitate understanding of the hyponym relations in the visualization, when the mouse pointer rests over a node it is highlighted in a saturated green and all hyperonyms of that node are highlighted in a saturated blue, tracing a path to the root. Details of the synset under the mouse pointer are provided in a linked view: the synset details window at the bottom of the interface lists the synset members and the associated sense definition. A full text details tab is also provided, in which occurrences of the synset under the mouse pointer are highlighted in the text of the document being visualized. (see Figure 2).

5 Future Work

There are several areas for further research and improvements of this prototype tool. We are interested in investigating dynamically-linked visualizations, for example, a scatter-plot of concept occur-

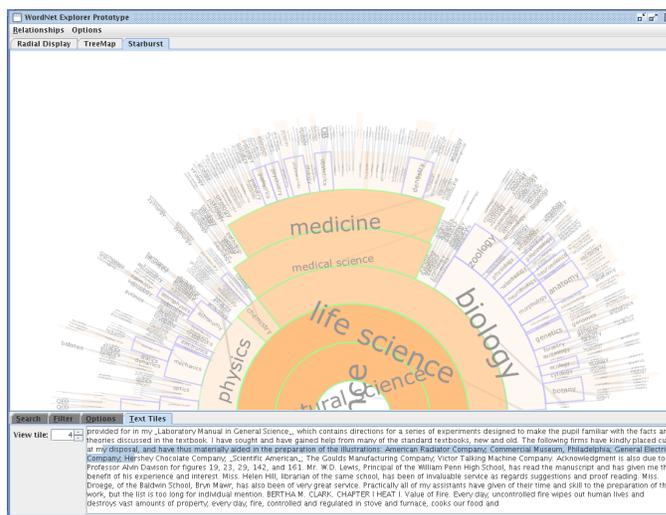


Figure 2: Subtree view rooted at {science}. A details window at the bottom of the interface shows the full text of the document being visualized. Controls provide filter, search, and graph-loading functions. To maximize display space, the entire control panel acts as an auto-hide task bar, lowering when not in use.

rence by text segments to quickly reveal where highly occurring words appear in the document. This work leads well into an investigation of the DocuBurst technique to view the difference between two documents, which may be useful for plagiarism detection, document categorization, and authorship attribution. Spreadsheets of many DocuBursts could also be used for comparison across document collections. Future evaluation of DocuBurst against content analysis using other tools and full-text documents will reveal the usability and power of combining radial space-filling layout with human-created linguistic structure in a document analysis tool.

References

- [1] Freddy Y. Y. Choi. Advances in domain independent linear text segmentation. In *Proceedings of NAACL*, pages 26–33, 2000.
- [2] Philip DeCamp, Amber Frid-Jimenez, Jethran Guinness, and Deb Roy. Gist icons: Seeing meaning in large bodies of literature. In *Poster Session, IEEE Symposium on Information Visualization*, Minneapolis, USA, October 2005.
- [3] Christiane Fellbaum, editor. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, USA, 1998.
- [4] George W. Furnas. Generalized fisheye views. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 16–23. ACM Press, April 1986.
- [5] Marti A. Hearst. Tilebars: visualization of term distribution information in full text information access. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 59–66. ACM Press, 1995.
- [6] John Stasko and Eugene Zhang. Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *Proceedings of the IEEE Symposium on Information Visualization (INFOVIS)*, pages 57–65, 2000.
- [7] Jing Yang, Matthew O. Ward, and Elke A. Rundensteiner. Interring: An interactive tool for visually navigating and manipulating hierarchical structures. In *Proceedings of the IEEE Symposium on Information Visualization (INFOVIS)*, pages 77–84, 2002.