KNOWLEDGE
MEDIA
DESIGN
INSTITUTE

DocuBurst: Radial Space-Filling Visualization of Document Content

Christopher Collins

Technical Report Series

KMDI-TR-2007-1

DocuBurst: Radial Space-Filling Visualization of Document Content

Christopher Collins
Department of Computer Science
University of Toronto
10 King's College Road, Toronto, Ontario, Canada M5S 3H5
ccollins@cs.utoronto.ca

Abstract

We present the first visualization of document content which takes advantage of the human-created structure in lexical databases. We use an accepted design paradigm to generate visualizations which improve the usability and utility of WordNet as the backbone for document content visualization. A radial, space-filling layout of hyponymy (IS-A relation) is presented with interactive techniques of zoom, filter, and details-on-demand for the task of document visualization. The techniques can be generalized to multiple documents.

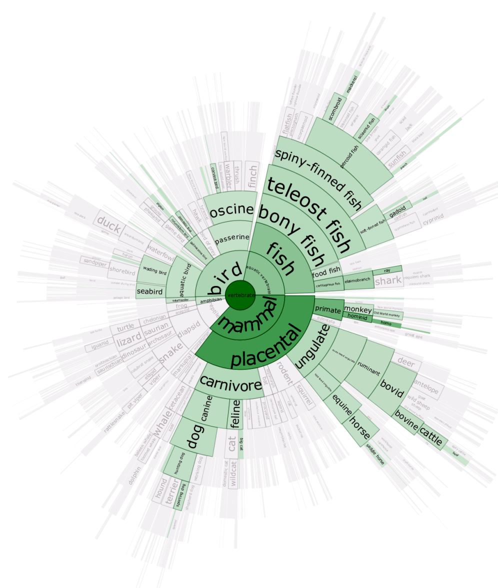

Figure 1: Hyponymy of single sense {*vertebrate*}, fully expanded. Subtrees with words occurring in the reference document are green with opacity directly related to strength of occurrence.

1 Introduction

'What is this document about?' is a common question when navigating large document databases. Overviews of document content have been an active area of research in information visualization for many years. Most reported works do not make use of human-annotated linguistic structure in the visualization, providing detail on topic content without a consistent view that can be compared across documents. We provide a visualization of document content based on the human-annotated IS-A noun hierarchy of WordNet (Fellbaum, 1998) (see Figure 1) and embedded in the multi-view visualization system WordNet Explorer (see Figure 2).

Despite the growing dependence on statistical methods, many NLP techniques still rely heavily on human-constructed lexical resources such as WordNet (Fellbaum, 1998). Senses are the most important data unit in WordNet and are defined as sets of synonymous words called *synsets* and an associated definition or *gloss*. WordNet is a lexical database of *words* and *synsets* and edges representing over 200,000 word MEMBER-OF synset relationships and many more synset-synset relationships. Throughout this paper, we will refer to *words* in single quotes (e.g 'thought'), and synsets using standard bracketed set notation (e.g. {*thought, idea*}). A word may be a member of multiple synsets, one for each sense of that word. Synsets in WordNet are related by many types of relationships, depending on the part of speech (noun, verb, etc.). WordNet contains 28 different types of relationships, but the most widely used part of WordNet is the hyponymy (`IS-A`) partial order. We remove all cycles (they are very rare) by taking a depth-first spanning tree at the user-selected root. In this work we focus on the noun hyponymy (`IS-A`) relationships in English WordNet (v2.1), rooted under the synset {*entity*} and having $73,736$ nodes (synsets) and $75,110$ edges. Verb hyponymy is also supported (but the tree is much smaller). The visualizations produced can be generalized to any partial order of a lexicon.

While development on WordNet continues, the interfaces for interacting with WordNet have not progressed to take advantage of advances in the field of information visualization. Currently available interfaces, both textual and graphical, focus on regions of local interest, for example by searching for the relationships for a single synset (ThinkMap, 2005; Bou, 2003; Alcock, 2004). In recent work, we created a working prototype of a visualization suite for WordNet which allows for an overview of the data, as well as the ability to focus on specific synsets of interest and obtain details. After developing these visualizations, we realized that the linguistic structure provided by WordNet could be useful not only for abstract visualization of the network itself, but that by applying other linguistic measures upon the nodes, we could better understand other aspects of language. Of particular interest to many in the information visualization and information retrieval communities is document structure and topic content.

In the following sections we will describe related work in document content visualization and present our interactive, animated, space-filling radial graph visualization of document content and WordNet hyponymy.

2 Related Work

2.1 Document Content

Several visualizations for document content have been reported. Most use a subset of the language (user-specified or algorithmically-selected key terms) to create a glyph based on word occurrence counts. Starstruck (Hetzler et al., 1998) creates glyphs by arranging lines in a circular pattern, where each line corresponds to a word and line length to word occurrence count. Gist Icons (DeCamp et al., 2005) builds on this idea using latent semantic indexing to group semantically-related search terms and reinforces the document as a glyph by drawing a smoothed iso-surface around the starstruck backbone. Blobby Text (Rohrer et al., 1999) is a 3D visualization which distorts a spherical surface to represent counts of up to 14 pre-selected terms of interest. Individual 3D glyphs are placed in the space according to their similarity. All three systems allow for inter-document comparison using arrays of glyphs, but do not investigate content within a single document.

Other visualizations of document content focus on the vocabulary and structure of a single document, such as TextArc (Paley, 2002), which arranges the sentences of a document in a circular layout with the individual words placed in the center. Self organizing (Kohonen) maps (Lin, 1992) have been used to reflect the relative strength of topics in a document, and the Document Lens (Robertson and Mackinlay, 1999) is an approach to focus-in-context distortion viewing of an entire document.

The document visualizations by Rembold and Späth (2006) compare and contrast themes and keywords within a collection of related documents while simultaneously revealing the thematic and typographic structure within an individual document. However, these visualizations are not interactive — they are printed graphics relating the essays in a collection. None of these approaches make use of formal linguistic structures such as WordNet.

TileBars (Hearst, 1995) is a document content visualization for information retrieval. It creates parallel small multiple visualizations, one for each user-specified search term set. Each small multiple consists of a bar divided into squares, each representing a 'text tile' (Hearst, 1997). The value of gray in the squares represents the occurrence of the corresponding search term in the associated text tile. This allows for overview of not only what terms appear in the documents, but also where they occur. Again, a small multiples view allows for inter-document comparison. ThemeRiver (Havre et al., 2002) is a 2D visualization of topic predominance in

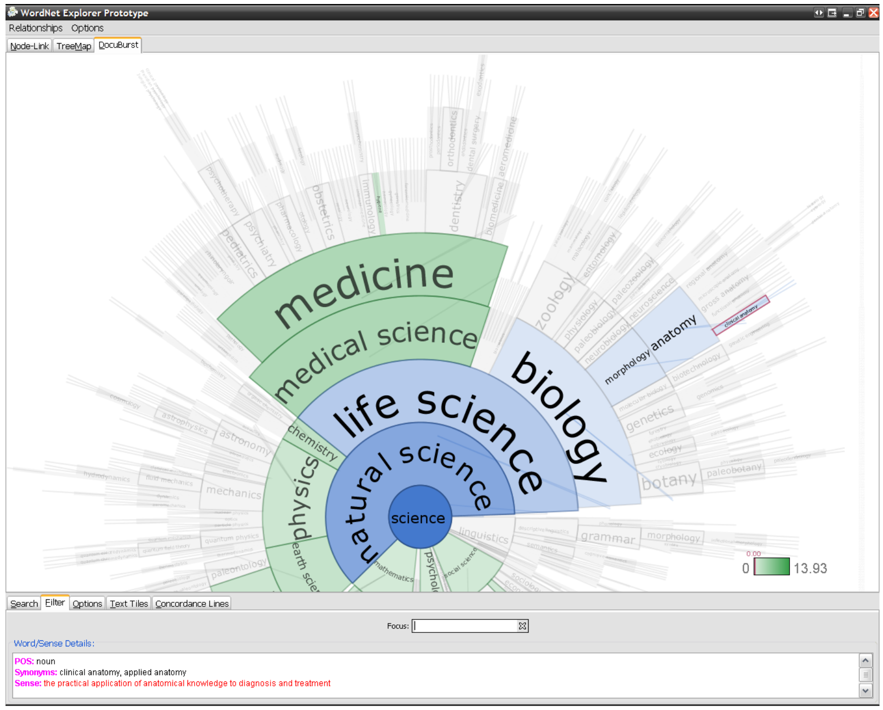

Figure 2: WordNet Explorer DocuBurst document content view rooted at {*science, scientific discipline*}. Fill opacity (green and blue) is proportional to number of occurrences of word in reference document. The mouse focus highlights {*clinical anatomy, applied anatomy*} in blue with pink outline and the path to the root is coloured blue to enhance understanding of the hyponymy structure of the focused subtree. Details of the synset under the mouse appear in the detail panel on the bottom.

a large document corpus. It depicts content changes through time, but not on an individual document's scale. Finally, the Galaxies visualization (Wise et al., 1999) visualizes document similarity through clustering and the related Themescape (Wise et al., 1999) overlays the Galaxies clusters with a surface of varying height, revealing the 'peaks' of strong keyword occurrence. Again, keywords are not related to one another nor clustered based on their semantic similarity.

2.2 WordNet

Many interfaces for WordNet exist, the most popular of which is the text-based WordNet Search[1] which is part of the publicly available WordNet package. With the exception of (Kamps and Marx, 2002), the existing interfaces for WordNet either provide for drill-down textual or graphical interaction with the data starting at a single synset of interest or provide path-tracing between two synsets. No reported visualization of WordNet uses the graph structure to enhance a visualization of other data (such as document content).

[1] http://wordnet.princeton.edu/perl/webwn

2.3 Graph Drawing

Radial graph-drawing techniques have been previously reported and serve as the basis of this work. Of particular interest are the semi-circular radial space filling (RSF) hierarchies of Information Slices (Andrews and Heidegger, 1998) and the focus + context interaction techniques of the fully circular Starburst visualization (Stasko and Zhang, 2000). The InterRing (Yang et al., 2002) visualization expands on the interaction techniques for RSF trees, supporting brushing and interactive radial distortion. TreeJuxtaposer (Munzner et al., 2003) illustrates methods for interacting with very large trees, where nodes may be assigned very few pixels. We adapt techniques such as tracing the path from a node of interest to the root and performing interactive accordion expansion from this work.

3 Design Paradigm

The most influential and succinct design framework is the information-seeking visualization paradigm of Shneiderman (1996): "Overview first, zoom and filter, then details-on-demand". This three-step description of visualization usage can be read as a design guideline summarizing many of the requirements of effective information visualization design. Most importantly, it captures the need for visualizations to be effective on both a macro and micro level. A visualization first provides an *overview* of the entire data set, displaying high-level features of the data to allow the user to then specify a region of interest. *Zoom and filter* functionality allows the user to target a region of interest using one of several methods: (1) remove the context from the display, (2) provide more detail on a focal region, abstract and display surrounding data, or (3) show detail in a new window, highlight region of enlargement on the overview display. We provide the first two forms. Finally, *details-on-demand* provides more detailed features of the data, for example by opening a list of synsets containing a word when it is selected with the mouse.

4 Linguistic Pre-processing

In order to populate the hyponymy tree with word counts, several pre-processing steps are necessary. Starting with raw text, for example, a book, we perform text tiling (Hearst, 1997). For each tile, we label parts of speech (`NOUN`, `VERB`, etc.) using the tagger by Brill (1993). Nouns (and verbs) are then extracted and stemmed (*e.g.*, books → book, going → go) using the morphological processor of Didion (2003). Punctuation is omitted.

Next we look up in which WordNet synsets the (*word, part-of-speech*) pairs occur. Because pairs usually occur in multiple synsets and we do not perform word sense disambiguation, we divide the word count evenly by the number of available synsets. This results in word counts that are not integers, but we think the overall results more accurately reflect document content, than, for example, assigning the full word count to all synsets in which it occurs. By dividing the counts, we dilute the contribution of highly ambiguous terms. The full text of tiles and their associated (*word, part-of-speech, count*) pairs are then read into the WordNet Explorer. This processing was carried out on a general science textbook (Clark, 1912) and this is the document from which counts are drawn for all examples in this paper.

5 DocuBurst Visualization

The combined structure of WordNet hyponymy and document lexical content is visualized in WordNet Explorer as the DocuBurst visualization, which uses a radial space filling layout technique. The root node is shown as a circle. All other nodes are assigned to a sector of an annulus with angular width which is part of the parent node's width. Angular width can be either (a) proportional to the number of leaves in the subtree rooted at that node (*leaf count*) or (b) proportional to the number of word occurrences counted for synsets in the subtree rooted at that node (*word count*). The leaf count view has the advantage of identical structure independent of which document is being viewed, supporting direct comparison of DocuBurst glyphs for different documents. The word count view maximizes screen space for synsets whose words actually occur in the document of interest, but the shape, as well as node colouring, will differ across documents. Depth in the hyponymy tree determines on which concentric ring a node appears (increasing depth corresponds to increasing radius). The width of each annulus is maximized to allow for all visible graph elements to fit within the display space (on initial display with neutral zoom factor).

A *root* search box is provided to select a root node for the visualization. Views can be rooted at any synset in the language, and are populated with all hyponyms of that synset. As the user types in the search box, the number of available synsets (senses) and parts-of-speech are constantly updated with each key press. If no synsets match the search query, the query is stemmed and tried again. A numbered listing of sense glosses can be used to assist selection of a synset of interest. After searching, the user can choose (a) to view all synsets containing that word (i.e. all senses of the word) (*e.g.*, Figure 4) or (b) a specific synset (sense) of interest (*e.g.*, Figure 1). For the all-senses view, the word is assigned to the initial central node and all its synsets and their

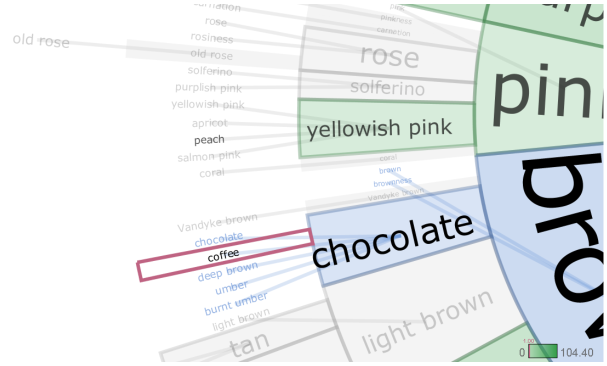

Figure 3: Semi-transparent edges distinguish the word `MEMBER-OF` synset relationship from the synset `IS-A` synset relation, which is implied by adjacency in the RSF structure. As the mouse hovers over 'coffee,' all its sibling words are highlighted in blue, as is the path to the root. 'Coffee' has an occurrence count of 1.00 as shown by the legend in the lower right. Note 'peach' is also in a darker font, indicating this specific word has a non-zero occurrence count, which is propagated to its synset, {*yellowish pink, apricot, peach, salmon pink*}.

Parameter	**Possible Values (*default*)**
Tree root	*synset*; word (all synsets)
Transparency mapping	*cumulative*; single node
Angular width	*leaf count*; word count
Word visibility	*on*; off
Zero-count visibility	*on*; off
Fish-eye filter depth	*3*; 0—20

Table 1: User-adjustable parameters for the DocuBurst visualization, and default values.

hyponyms are loaded into the visualization. For the single-sense view, the selected synset of interest is assigned as the initial central node and all its hyponym synsets and their word members are loaded into the visualization.

Nodes are rendered as labeled annulus sectors. Synset nodes have coloured backgrounds, word nodes are rendered as labels without fill. To further distinguish words from senses, semi-transparent edges connect words to senses while sense-sense relationships are implied by graph structure (see Figure 3). Font size is maximized and labels are rotated to allow the label to fit within the node and minimize label overlap. As senses are collections of words associated to a single gloss, senses are labeled with their first word member. To improve graph clarity and visibility of document content words, for nodes with angular width < 0.25 degrees, edges, labels, and sector outlines are omitted. For views rooted at a single word and containing all its senses, all nodes corresponding to the same sense of the root word are assigned the same hue to ensure unrelated subtrees are distinguished (see Figure 4). Trees rooted at a single synset appear as a single hue (green) (see Figure 1).

Document content is visualized through the transparency of the fill colour of the nodes. Gray hue is also used to distinguish nodes with zero occurrence counts. Highly coloured nodes have many occurrences; almost transparent nodes have few occurrences. This is a natural mapping as colour value and hue are ordered (Carpendale, 2003; Bertin, 1983) and pre-attentive (Ware, 2004) visual properties (with limits, of course, on how many variations are can be present at the same time). Transparency performs a blending of value and hue as in the limit of transparency both hue and value are reduced to the background value of these properties (Zuk and Carpendale, 2006). Words and senses that are more prominent in the document of interest stand out easily against a more transparent context.

Two ways to visualize word occurrence are provided: single-node and cumulative. In the *single-node* visualization, only synset nodes whose word members occur in the document are coloured. In the *cumulative* view, counts are propagated up to the root of the tree. In both views, transparency is normalized so maximum counts achieve full opacity. The single-node view allows for immediate viewing on precise concepts in the document and selection of synsets whose word members appear directly in the document being analyzed. The cumulative, or subtree, view uses 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. The cumulative view is useful for tracing where counts occur deep in regions of the tree that may be hidden by a fisheye filter. In addition, for a fully expanded graph, the single node view may highlight nodes that are too small to notice. Similar to the TreeJuxtaposer (Munzner et al., 2003) techniques for visualizing differences embedded deep in a large tree, by highlighting the entire subtree containing the node, salient small

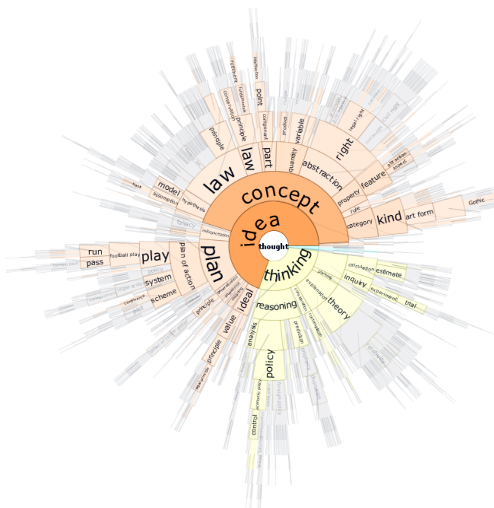

(a) Node fill transparency proportional to number of occurrences of words in subtree (cumulative view).

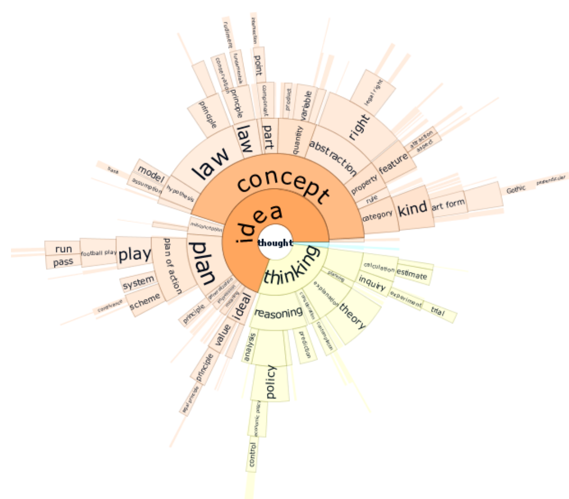

(b) Cumulative view with zero-count nodes hidden.

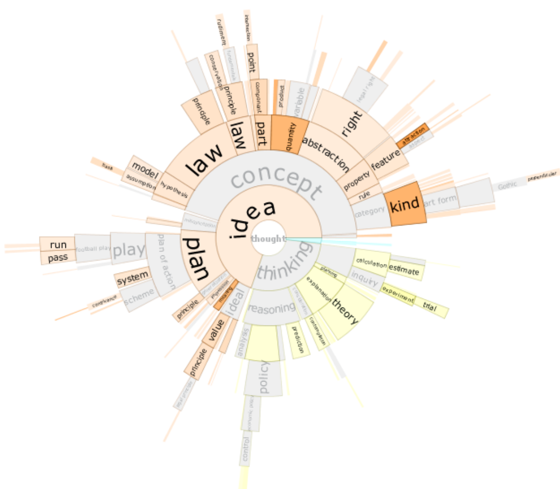

(c) Node fill transparency proportional to number of occurrences of words in a single synset (single node view).

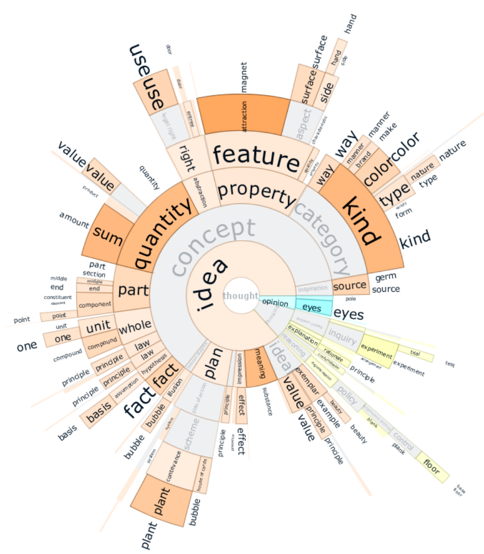

(d) Single node view, node size proportional to occurrence count total for subtree rooted at that node.

Figure 4: Four viewing modes allow for in-depth analysis of the data. View rooted at the word ‘thought’; children at depth 1 represent synsets containing ‘thought’ (*i.e.*, the different senses of ‘thought’). Node hue distinguishes the senses of ‘thought’. For (a–c) node angular width is proportional to the number of children in the node’s subtree. For (d) node angular width is proportional to the number of occurrences of words in the subtree rooted at that node.

nodes can be more easily located. The subtree and cumulative views are compared in Figure 4. Table 1 lists a summary of all view parameters and their default values.

This visualization is based on the structure of WordNet, which has been created by a team of linguists and cognitive psychologists over a period of 20 years. We hypothesize that this structure will provide more intuitive abstractions of document content than those developed through statistical techniques such as latent semantic analysis.

5.1 Zoom and Filter

The system can be used to explore the document content coverage of the entire language, but when viewing the more than $75,000$ English noun synsets, the animation slows and label illegibility becomes a serious problem. Such problems of scale are well-managed by following the visualization design paradigm: 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 gold background and border, and a darker font colour, drawing attention to even the smallest of search results. The transparency of the highlight (gold) background is attenuated to the node occurrence counts so as to not disrupt this data-carrying value and to provide for stronger "pop-out" of search results which high occurrence counts. Second, we implement a kind of generalized fisheye views (Furnas, 1986) 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 cumulative colouring, 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 (see Figure 5).

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 data navigation through the hyponymy relation.

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. This interaction can provide for increased detail on nodes of interest (see Figure 6).

The counts that are used to determine node transparency are based on text tiles, or automatically-determined subtopic regions of the document. The initial view is based on all text tiles in the document, but tile range selectors allow for limiting the tiles from which counts are drawn (see Figure 7).

Unrestricted visual pan and zoom of the display space are also supported, as well as a zoom-to-fit control (right click on background) to reset the pan and zoom to a best-fit for the currently visible tree. Semantic zoom is provided in that node borders are not rendered when the nodes are very small, and labels are not rendered when, through node size or visual zoom, they would not be legible. 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.

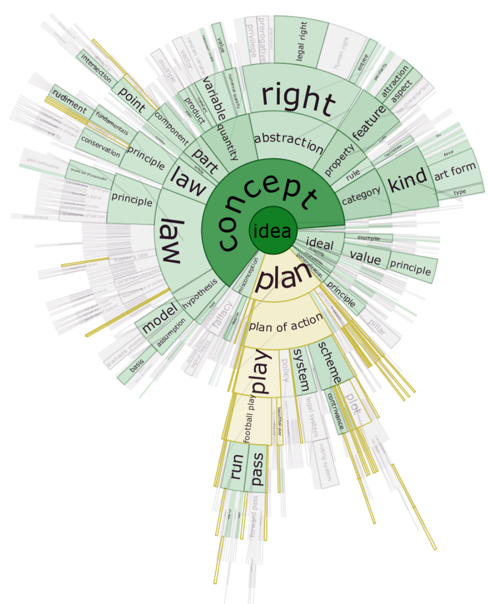

Figure 5: Hyponymy of {*idea*} with fish-eye filter at depth = 3. Nodes matching search query (starting with 'pl') are expanded by increasing their *a priori* importance. All nodes containing a search result are highlighted in gold.

Figure 6: Hyponymy of {*energy*}. Mouse wheel interaction used to roll up subtree rooted at {*radiation*} and expand subtree under {*heat*} exposing previously illegible node {*heat of fusion*}. Nodes more than 3 edges from the central focus are collapsed, for example, the hyponyms of {*beam*}.

Figure 7: Document content for hyponymy of {*energy*}. At left, counts are based on all 210 text tiles. At right, counts are restricted to only the first 10 tiles.

5.2 Details on Demand

Because of the potentially high data density, it may be difficult to discern parent-child relationships using DocuBurst. To facilitate understanding of the hyponymy relations in the visualization, when the mouse pointer rests over a node it is highlighted in a saturated green and all hyperonyms (ancestors under the `IS-A` relation) of that node are highlighted in blue. We think that this makes the task of reading the `IS-A` hierarchy for a synset of interest very easy. Details of the synset under the mouse pointer are provided in the synset details window at the bottom of the interface (see Figure 2). While transparency is an effective data-carrying visual variable for distinguishing large differences and trends, it is impossible to read exact values using it. To facilitate exact reading of synset occurrence counts for the selected text tiles, we provide a dynamic legend (see Figure 8).

A full text details tab is also provided at the bottom of the interface, containing the full text of the document whose content is being visualized (divided into text tiles). To access the text tiles, we use a linked visualization: the text tile browser. A numbered list of text tiles appears in a linear, vertical array to the right of the DocuBurst glyph. A fish-eye distortion (Bederson, 2000) facilitates navigation and selection of this list without scrolling. By clicking any tile number, that text tile is brought into view. Furthermore, this visualization can be used to see overall trends in the document. By clicking nodes in the DocuBurst visualization, synsets and words can be selected. Text tiles in which those synsets and words appear show as orange in the text tile browser. Occurrences of those synsets and words are also highlighted in orange in the full text window, and concordance lines[2] are extracted and shown in the concordance window. Patterns of orange in the tile browser can indicate how localized concepts are in the document. For example, in Figure 9, we see that {*electricity*} appears more frequently toward the end of the document. We can use the tile browser and full text window to quickly find occurrences of the terms of interest in context. By clicking the text tile numbers in the tile browser, we find, in the tile detail window, that there is a chapter on 'electricity' at the end of the book.

[2]A concordance line is a standard linguistic analysis tool in which all occurrences of a word of interest are extracted and displayed with their left and right n (usually 5) context words.

6 Conclusion and Future Work

By following commonly accepted design and interaction principles of information visualization we believe DocuBurst provides a useful overview visualization of document content which is based on a human-centered view of language where previous works were based on statistical analysis. Semantic and visual zooming, filtering, search, and details-on-demand provide a detailed view of what subset of language is covered by a document. This work introduces several interesting future work opportunities and challenges. Most importantly, the original goal of viewing what parts of an entire language are included in a document was not adequately met. As in other related works, it was necessary to view a subset of language due to limited display space and computational resources. Views rooted at {*entity*} appear cluttered and interaction is slow (see Figure 10). It is commonly held that WordNet sense-divisions are too fine-grained for many computational applications; investigation into other ways to abstract WordNet may help alleviate this problem. Alternatively, providing hints for which synsets may be of interest as tree roots for a particular document or set of documents may assist an analyst to find views of interest.

Currently, an occurrence of a word is divided among all senses (synsets) in which is appears. Thus 'river *bank*' will augment the count for {*bank, savings bank, depository financial institution*}. Word sense disambiguation is an area of active research in computational linguistics. Incorporating word sense disambiguation into the preprocessing step would greatly enhance the value of the visualization. For example, in the general science textbook used for the examples in this paper, 'man' occurs quite often in the sense of 'people' ("man invented the wheel"). However, these occurrences are also counted for the biological 'hominid' sense of 'man', resulting in the incorrectly strong appearance of the 'primate' subtree.

The use of transparency to indicate word occurrence is useful for the intuitive mapping between

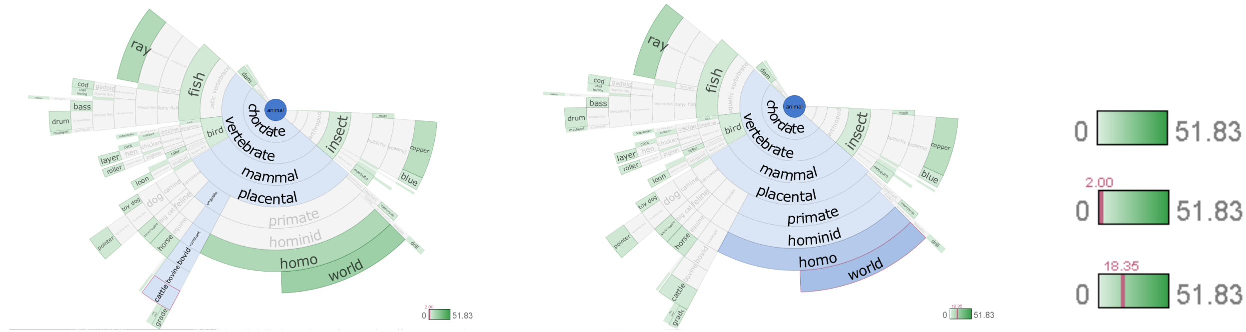

(a) **(b)** **(c)**

Figure 8: Occurrence count legend. **(a)** Synset {*cattle, cows, kine, oxen, Bos taurus*} highlighted. **(b)** Synset {*world, human race, humanity, humankind, human beings, humans, mankind, man*} highlighted. **(c)** Detail of legend for no selection (top), (a) middle, and (b) bottom.

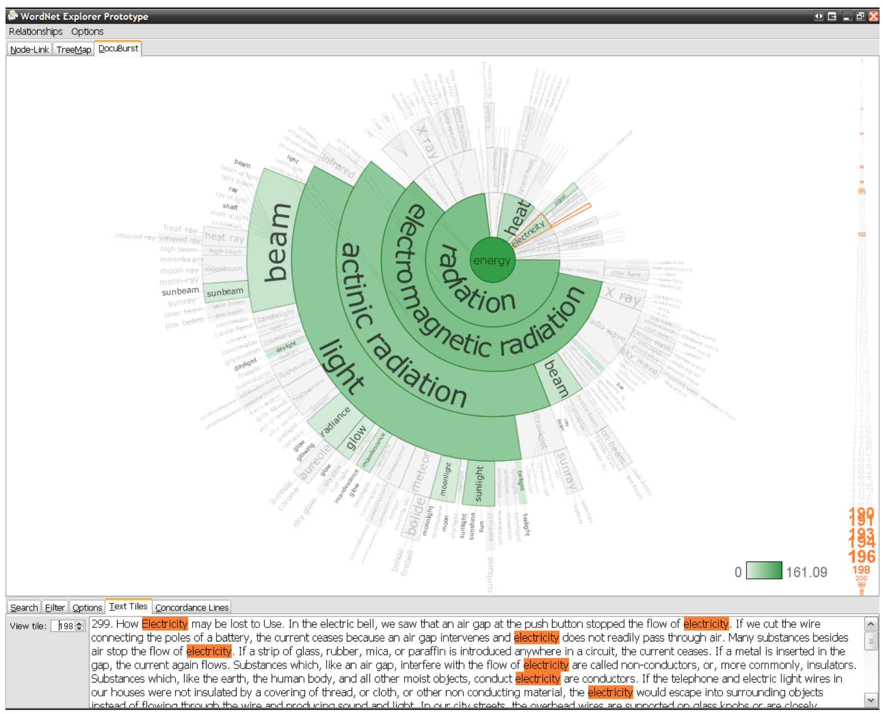

(a) The details window showing a selected text tile with occurrences of the selected synset {*electricity*} highlighted. Text in this window is shown one text tile at a time.

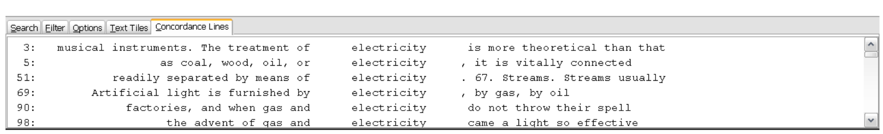

(b) The details window showing the concordance lines for the selected synset.

Figure 9: Linked visualizations for details-on-demand. The tile browser visualization on the right provides for navigation and selection of text tiles using a fish eye distortion. Text tiles containing any synset selected in the DocuBurst glyph appear in orange. A linked details window at the bottom of the interface reveals detailed information about the document being visualized. To maximize display space, the entire control panel acts as an auto-hide task bar, lowering when not in use.

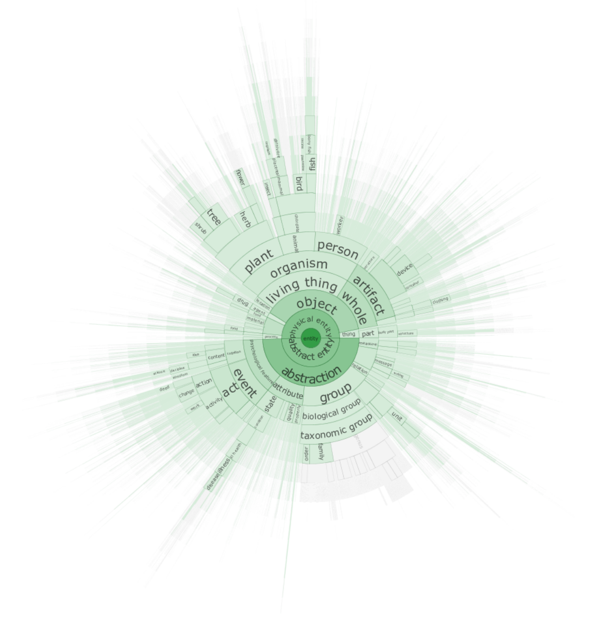

Figure 10: DocuBurst with subtree occurrence highlighting, rooted at {*entity*}, containing all English nouns. Nodes are too small and interaction too slow to be useful at this scale.

data and visual appearance. However, it also introduces the possibility of misleading illusions. Siblings in the DocuBurst are unordered. Furthermore, non-sibling nodes may be adjacent. By chance, unrelated nodes which both have high occurrence counts can appear as a large swath of strong colour. Gestalt perception may lead viewers to impart significance on this coincidence. Stronger node borders may distinguish these regions, but node borders become obstructive on small nodes.

This work leads well into an investigation of the DocuBurst technique to view the difference between two or more documents, which may be useful for plagiarism detection, document categorization, and authorship attribution. Arrays of DocuBurst icons could be compared against one another or a baseline reference corpus. Evaluation of the radial space-filling layout against and in conjunction with the radial node-link and tree map layouts also provided by WordNet Explorer on exploratory tasks would provide additional data on the usefulness of coordinated and different views of the same data, an area of information visualization needing further investigation.

References

Keith Alcock. 2004. WordNet relationship browser. Available from: `http://www.ultrasw.com/alcock/`.

Keith Andrews and Helmut Heidegger. 1998. Information slices: Visualising and exploring large hierarchies using cascading, semi-circular discs. In *Proc. of IEEE Symp. on Information Visualization (InfoVis), Late Breaking Hot Topics*, pages 9–12.

Ben Bederson. 2000. Fisheye menus. In *Proc. of the ACM Symposium on User Interface Software and Technology (UIST 2000)*, pages 217–226. ACM Press.

Jacques Bertin. 1983. *Semiology of Graphics: Diagrams, Networks, Maps*. University of Wisconsin Press.

Bernard Bou. 2003. WordNet treewalk. Available from: `http://wntw.sourceforge.net`.

Eric Brill. 1993. POS tagger. Software. Available from: `http://www.cs.jhu.edu/~brill/RBT1_14.tar.Z`.

M.S.T. Carpendale. 2003. Considering visual variables as a basis for information visualisation. Technical Report 2001-693-16, Department of Computer Science, University of Calgary, Calgary, Canada.

Bertha M. Clark. 1912. *General Science*. American Book Company, New York, project gutenberg edition. Available from: `http://www.gutenberg.org`.

Philip DeCamp, Amber Frid-Jimenez, Jethran Guiness, and Deb Roy. 2005. Gist icons: Seeing meaning in large bodies of literature. In *Proc. of IEEE Symp. on Information Visualization (InfoVis 2005), Poster Session*, Minneapolis, USA, October.

John Didion. 2003. Java WordNet Library. Available from: `http://jwordnet.sourceforge.net`.

Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, USA.

George W. Furnas. 1986. Generalized fisheye views. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 16–23. ACM Press, April.

Susan Havre, Elizabeth Hetzler, Paul Whitney, and Lucy Nowell. 2002. ThemeRiver: visualizing thematic changes in large document collections. In *IEEE Transactions on Visualization and Computer Graphics*, volume 8, January.

Marti A. Hearst. 1995. 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.

Marti Hearst. 1997. Texttiling: Segmenting text into multi-paragraph subtopic passages. *Computational Linguistics*, 23(1):33–64.

Beth Hetzler, Paul Whitney, Lou Martucci, and Jim Thomas. 1998. Multi-faceted insight through interoperable visual information analysis paradigms. In *Proc. of IEEE Symp. on Information Visualization (InfoVis)*, pages 137–144, North Carolina, USA, October.

Jaap Kamps and Maarten Marx. 2002. Visualizing WordNet structure. In *Proc. of the 1st International Conference on Global WordNet*, pages 182–186, Mysore, India.

Xia Lin. 1992. Visualization for the document space. In *Proc. of the 3rd Conference in Visualization*, pages 274–281, Boston, USA. IEEE.

Tamara Munzner, Francois Guimbretière, Serdar Tasiran, Li Zhang, and Yunhong Zhou. 2003. Treejuxtaposer: Scalable tree comparison using focus+context with guaranteed visibilit. *ACM Transactions on Graphics*, 22(3):453–462. SIGGRAPH 2003.

W. Bradford Paley. 2002. TextArc: Showing word frequency and distribution in text. In *Proc. of IEEE Symp. on Information Visualization (InfoVis)*, Poster, Boston, USA, October. IEEE Computer Society.

Magnus Rembold and Jürgen Späth. 2006. Graphical visualization of text similarities in essays in a book. Available from: `http://www.munterbund.de/visualisierung_textaehnlichkeiten/essay.html`.

George G. Robertson and Jock D. Mackinlay. 1999. The document lens. In Stuart M. Card, Jock D. Mackinlay, and Ben Shneiderman, editors, *Readings in information visualization: using vision to think*, pages 562–569. Morgan Kaufmann Publishers Inc.

Randall M. Rohrer, John L. Sibert, and David S. Ebert. 1999. A shape-based visual interface for text retrieval. *IEEE Computer Graphics and Applications*, 19(5):40–47, September.

Ben Shneiderman. 1996. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. of the IEEE Symp. on Visual Languages*, pages 336–343, Los Alamitos, USA. IEEE Press.

John Stasko and Eugene Zhang. 2000. Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *Proc. of IEEE Symp. on Information Visualization (InfoVis)*, pages 57–65.

ThinkMap. 2005. ThinkMap visual thesaurus. Available from: `http://www.visualthesaurus.com`.

Colin Ware. 2004. *Information Visualization: Perception for Design*. Morgan Kaufmann, 2nd edition.

James A. Wise, James J. Thomas, Kelly Pennock, David Lantrip, Marc Pottier, Anne Schur, and Vern Crow. 1999. Visualizing the non-visual: spatial analysis and interaction with information for text documents. In S. K. Card and J. D. Mackinlay, editors, *Readings in Information Visualization: Using Vision to Think*, pages 442–450. Morgan Kaufmann Publishers Inc., San Francisco, USA.

Jing Yang, Matthew O. Ward, and Elke A. Rundensteiner. 2002. InterRing: An interactive tool for visually navigating and manipulating hierarchical structures. In *Proc. of IEEE Symp. on Information Visualization (InfoVis)*, pages 77–84.

Torre Zuk and Sheelagh Carpendale. 2006. Theoretical analysis of uncertainty visualizations. In Robert F. Erbacher, Johnathan C. Roberts, Matti T. Gröhn, and Katy Börner, editors, *Proc. of SPIE-IS&T Electronic Imaging*, volume 6060, 606007.

The University of Toronto's

Knowledge Media Design Institute

The University of Toronto's Knowledge Media Design Institute (KMDI) – the first of its kind in Canada – is dedicated to research and graduate education in knowledge media design (KMD).

With more than 50 faculty from 25 faculties and departments participating, KMDI is a catalyst for collaborative endeavours and cross-disciplinary research. KMDI faculty and students design, build, use and critically evaluate novel knowledge media. They do this within the context of a research agenda that situates these developments in the broader social and cultural context. There is need to be aware of emerging trends in technology and to build malleable systems that take advantage of these developments to accommodate the practices of the communities with whom we work. The social and technical thus co-evolve in the context of people and their practices. Design – whether in terms of systems architecture, representation or process – informs the entire approach.

A Collaborative Master's and Doctoral Program in Knowledge Media Design began in 2002. The program is organized around four themes: *Human-centred design*, *Technologies for knowledge media*, *Knowledge media for learning*, and *Knowledge media, culture, and society*. Participating departments and faculties include: Information Studies; Mechanical and Industrial Engineering; Sociology; Computer Science; Architecture, Landscape, & Design; and the Institute of Medical Science.

Located in Canada's largest research university in Canada's largest city, KMDI is in the heart of an intellectually, culturally, economically and technologically dynamic community. Our research is helping to shape the products, processes and practices of the 21st Century.

Knowledge Media (KM)
The building blocks of a learning society

Knowledge Media Design (KMD)
The design, development and evaluation of media intended to support and enhance the ability of individuals and groups to think, communicate, learn and create knowledge.

For further information please contact:

Knowledge Media Design Institute
Bahen Centre for Information Technology
University of Toronto
40 St. George St.
TORONTO M5S 2E4

Phone/Fax: 416-978-KMDI
Email: info@kmdi.utoronto.ca
Web site: www.kmdi.utoronto.ca