Constellation: A Visualization Tool for Linguistic Queries from MindNet

Tamara Munzner
Stanford University *

François Guimbretière
Stanford University †

George Robertson
Microsoft Research ‡

1 Abstract

Constellation is a visualization system for the results of queries from the MindNet natural language semantic network. Constellation is targeted at helping MindNet's creators and users refine their algorithms, as opposed to understanding the structure of language. We designed a special-purpose graph layout algorithm which exploits higher-level structure in addition to the basic node and edge connectivity. Our layout prioritizes the creation of a semantic space to encode plausibility instead of traditional graph drawing metrics like minimizing edge crossings. We make careful use of several perceptual channels both to minimize the visual impact of edge crossings and to emphasize highlighted constellations of nodes and edges.

2 Introduction

Constellation is a tool for visualizing fragments of MindNet [5, 11], a system that constructs a large semantic network by parsing the text of machine-readable dictionaries and encyclopedias. MindNet was developed at Microsoft Research, and its possible applications include grammar checking, intelligent agent help systems, machine translation, and common-sense reasoning. MindNet's creators and users wanted a tool to help them check the plausibility of the results returned from a query about the connections between two words. The result of each query can be displayed as a medium-sized directed graph of less than one thousand nodes.

We found that relying on generic graph layout techniques to display this complex structure led to inadequate results. In most traditional graph drawing systems, spatial position bears most of the perceptual burden, and interaction is used simply for basic navigation. One of the major constraints on spatial positioning in traditional graph layout approaches is the need to minimize edge crossings, in order to avoid the visual impression of attachments that do not actually exist.

We instead incorporate a great deal of domain-specific information into our custom layout algorithm. Although spatial position is the strongest perceptual channel, we reduce false attachments by using several other perceptual channels in concert to create dynamically changeable foreground and background visual layers. The user can interactively explore highlighted subsets (*constellations*) of the graph while retaining the context of the entire dataset.

Our system incorporates ideas from graph drawing, human-computer interaction, computer graphics, cognitive psychology, and graphic design. While Constellation itself is quite narrowly targeted for the MindNet developers, the visualization and interaction techniques that we have developed have a wider scope and can be reused in other domains. Color Plates 1 and 2 show the visual effect that we have achieved.

3 MindNet

The MindNet parsing process turns a dictionary or encyclopedia entry sentence into a small *definition graph* of roughly one dozen nodes. The edges represent directed labelled relations between words, such as "is-a", "part-of", or "modifier". The nodes represent *word senses*: a natural language word may have several meanings depending on context, for instance "bank" as "financial institution" or "side of a river". MindNet distinguishes between these word senses by adding a numerical suffix and treats them as separate nodes. Two definition graphs that share a node can be combined into a larger graph, and this unification process results in a huge semantic network that can contain millions of nodes.

The semantic networks generated by MindNet are sufficiently large and interconnected that its developers find it impractical to study their global structure. They instead rely on a query engine to probe a small subsection of the network. Each of these snapshots is checked for potential problems and the system is modified to address them. The user provides a query consisting of two words and the number of *paths* to return. MindNet computes the best paths between the words, using (among other things) the edge weights in its unified network of definition graphs. Our target users requested that

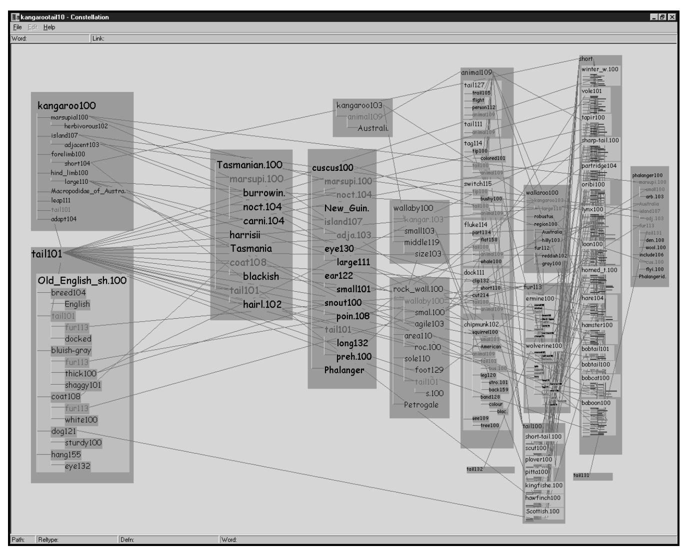

Figure 1: The Constellation visualization system showing the MindNet query result of the top ten paths between "kangaroo" and "tail", along with the definition graphs used to construct those paths. This black and white figure shows our spatial layout but since our visual layering scheme requires color, we have postprocessed the image for legibility purposes. Plate 2c shows a color version of this dataset.

*http://graphics.stanford.edu/~munzner, munzner@cs.stanford.edu, (650) 723-3154, 3B-360 Gates CS Bldg, Stanford, CA, 94305

†francois@cs.stanford.edu, (650) 723-0618, 3B-381 Gates CS Bldg, Stanford, CA, 94305

‡http://www.research.microsoft.com/~ggr, ggr@microsoft.com, One Microsoft Way, Redmond, WA, 98052

```
Kangaroo100
    Hyp>    marsupial100
                Mod>    herbivorous
    Locn>   island
                Mod>    adjacent
    Part>   forelimb100
                Mod>    short
            hind_limb100
                Mod>    large110
    <Part   Macropodidae_of_Australia
    Purp>   leap111
            tail101
    <Tobj   adapt104
```

(a)

```
Tasmanian_devil
    Hyp>    marsupial100
                Mod>    burrowing
                        nocturnal
                        carnivorous
            harrissi
    Locn>   Tasmania
    Part>   coat108
                Mod>    blackish
            tail101
                Mod>    hairless
```

(b)

1	kangaroo100 Purp> tail101	kangaroo100
2	kangaroo100 Hyp> marsupial100 <Hyp Tasmanian_devil100 Part> tail101	kangaroo100 Tasmanian_devil100
3	kangaroo100 Hyp> marsupial100 <Hyp cuscus100 Part> tail101	kangaroo100 cuscus100
4	kangaroo103 <Hyp wallaby100 <Hyp rock_wallaby100 <Tsub sole110 Tobj> tail101	wallaby100 rock_wallaby100
5	kangaroo103 Hyp> animal109 Part> tail101	kangaroo103 (tail127 tail111 tag114 switch115 fluke114 dock111 chipmunk102)
6	kangaroo103 <Hyp wallaroo100 Part> fur112 Join> fur113 Mod> tail132	wallaroo100 (phalanger100 ermine100)
7	kangaroo100 Part> forelimb Mod> short104 Join> short <Mod tail100	kangaroo100 (winter_wren100 vole101 tapir100 short-tailed_shrew100 sharp-tailed_grouse100 scut100 plover100 pitta100 partridge104 oribi100 lynx100 loon100 kingfisher100 horned_toad100 hawfinch100 hare104 hamster100 bobtail101 bobtail100 bobcat100 baboon100 Scottish_terrier100)
8	kangaroo103 <Hyp wallaroo100 Part> fur112 Join> fur113 Part> tail101	wallaroo100 Old_English_sheepdog100
9	kangaroo103 <Hyp wallaroo100 Part> fur112 Join> fur113 Part> tail100	wallaroo100 wolverine100
10	kangaroo100 Hyp> marsupial100 <Hyp phalanger100 Part> fur113 Mod> tail131	kangaroo100 phalanger100

(c)

Figure 2: The previously existing textual views in MindNet: the definition graphs for kangaroo100 and Tasmanian devil100 are shown in (a) and (b), and the ten best paths between kangaroo and tail are shown in (c). The first path is only one hop since tail101 is present in the definition of kangaroo100. The second shows that Tasmanian devils and kangaroos are both marsupials, and Tasmanian devils have tails. The words to the right of each path are the definition graphs which were used by MindNet to construct the path. This output has been simulated for legibility purposes. Figures 1 and 3 and Plate 2c show this full dataset laid out in Constellation.

we treat the internal MindNet path computation algorithms [5, 11] as a black box, so we simply use the ordering of the returned paths as a *plausibility* measurement. The system returns both the paths themselves and the definition graphs used to compute each path. Although each path usually contains less than ten words, the number of associated definition graphs can range from one to dozens. The users typically request the best ten or fifty paths.

A single word sense may appear in multiple places: it could lie on several different paths, appear as a *leafword* inside several definitions, and also appear as a *headword* with its own definition graph. These shared words are the reason it is difficult to understand how paths relate to each other and to the definition graphs used to create them. The existing MindNet software infrastructure provides dozens of specialized local textual views as in Figure 2, but none of them provides the MindNet developers with a global picture of the relationships between paths and definition graphs. One of the main design goals of Constellation is to provide such a spatial framework to help the users analyze these interrelationships.

4 Spatial layout

Color plates 1 and 2 show the results of our graph layout algorithm, which uses the domain-specific elements of paths and definition graphs. The broad layout parameters are based on the two orderings returned directly by MindNet: horizontal flow is derived from the plausibility ordering between the paths, and vertical flow is based on the internal ordering of words within a path, with the source on top and the sink on the bottom.

We start on the left with the most plausible path, and draw its *pathwords* in tan boxes one at a time at points along a curved vertical band from top to bottom. The plausibility ranking is used for the horizontal ordering and width of the path bands: the most important band on the left is straightest and widest, while lower ranking path bands farther to the right are more curved and narrow. Important words on the left are drawn larger than implausible ones on the right since wider bands allow larger font sizes. The left-to-right plausibility size gradient is usually much more visually salient than the underlying bands. After the vertical sweep for the pathwords of the leftmost path is finished, we move on to the next band on the right. A pathword that occurs in multiple paths is drawn only once, in the first (most plausible) path position encountered on the above sweep.

MindNet provides an association between a path and all the definition graphs used in its computation. In some cases a pathword is the *headword* of a definition graph, whose *leafwords* are drawn beneath it in a ladder-like structure with blue label boxes, as in the two tan boxes in Plate 1b. This local structure as shown in Figure 3 is deliberately similar to the previously existing definition graph view shown in Figures 2a and 2b.

In other cases a pathword is a leafword rather than the headword of a definition graph, so the entire definition graph is drawn in a green box nested within the tan pathword box, as in the leftmost box of Plate 1b. Some path computations involve the pooled influence of many definition graphs for a single pathword, so there may be many green boxes vertically stacked inside a single tan pathword box, as in the lower left corner of Plate 2a. Path 7 of the kangaroo-tail 10 path dataset is an extreme example, visible as text in Figure 2c and in our layout in Plate 2c.

Although pathwords and thus entire definition graphs are drawn only once, the individual words inside a definition graph may be drawn multiple times so that the entire definition can be easily read in the local view. The first time a word is drawn it is colored black, and subsequent instances are drawn in grey and connected to the master by long slanted lines.

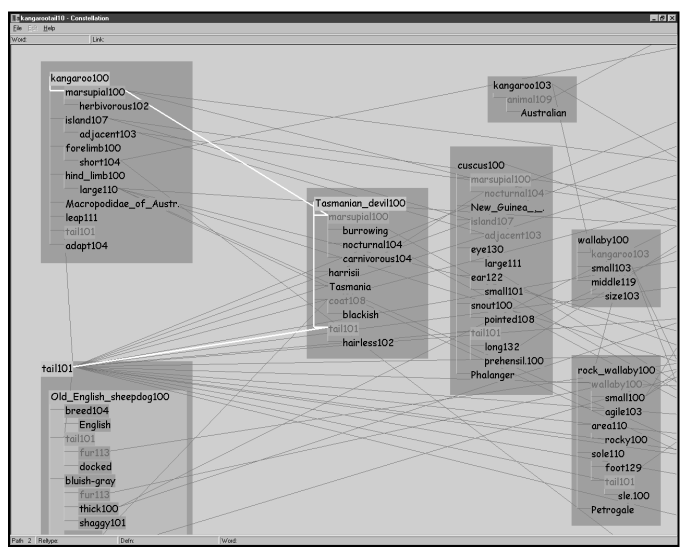

Figure 3: We have zoomed in on the dataset of Figure 1 to focus on the highlighted path 2 elements, for comparison with path 2 in Figure 2. We have postprocessed this image for black and white legibility purposes. Our visual layering scheme is visible in the color plates.

5 Navigation and legibility

Constellation is optimized for three viewing levels: a global view for inter-path relationships, an intermediate view for pathword associations, and a local view for reading individual definition graphs. In the global view the path-level words are emphasized by devoting a large part of the box space to the top word. Leafwords are smaller than headwords, and may be drawn as a line or omitted completely if there is no vertical room, as in the Plate 1a inset. All the blue leafwords in a tan pathword box are drawn the same size, to present equal visual salience, so some of the longer words are horizontally elided as in the large box on the left in Plate 1a. The word drawing strategy changes at higher zoom levels. At the local level, the task is reading a single definition graph, so the box space is more equally divided among words as in Plates 1b and 1e.

Although mouse dragging offers direct fine control over panning and zooming, the main navigation method is a mouse click inside any enclosure box, which triggers an animated transition [12]. The zoom level is computed so that the box is vertically framed within the window and there is enough horizontal space to draw every word string completely with no ellipsis.

6 Perceptual channels

The visual encoding of abstract information into perceptual channels is one of the central issues of information visualization [8, 9]. Section 4 describes the use of quantitative spatial position to encode plausibility and proximity to encode association. We also allow exploration of the dataset through the selective highlighting of *constellations* of boxes and edges. The four constellation categories are shown in the color plates: paths in Plates 1a, 1c, 2a, 2c; definition graphs in Plates 1d, 2a; word senses in Plates 1e and 2a, and relation types in Plate 1f.

Although no other perceptual channel alone is as salient as spatial position, combining several of them has proven to be very effective at creating visual popout to distinguish a foreground from a background visual layer [14]. The background layer with its many edge crossings is visible at all times for context, but is unobtrusive since the edges have low saturation and brightness. We emphasize the foreground layer by increasing both saturation and brightness. In the case of lines, we also increase the size, for the synergistic result that the hue differences in wide lines are much more discriminable than in the unhighlighted narrow ones.

The colored text background boxes use grouping and enclosure to encode the hierarchical relationship between pathwords and definition graphs. These boxes also provide a colored area large enough for effective hue discrimination and to maximize the legibility of the black label text. The long slanted lines between master and proxy instances of the same word sense encode association with a connection cue. The short axis-aligned lines between words in a single definition graph are visually distinguishable from the long proxy lines by the orientation channel.

Finally, we use hue as a nominal variable, to distinguish between the types of enclosure boxes and the types of relation lines. Each of the eight relation types is color coded with with hues 45 degrees apart on the HSB color wheel, while the three hues for the enclosure boxes (tan, green, and blue) were empirically chosen to complement them. Our color scheme draws heavily on ideas from Reynolds [10], who presented a set of color palettes to improve the legibility of air traffic control displays.

7 Interactive visual emphasis

The previous section discusses the use of multiple perceptual channels to bring a particular subset of the data to the emphasized foreground visual layer. This interactive visual emphasis is similar in spirit to the dynamic queries of previous information visualization systems such as FilmFinder [1]. Our lightweight hover mode allows quick visual inspection with no need to navigate. In hover mode, moving the mouse over a link marks it visually and shows full details about its origin and destination in an upper status bar. Hovering over a word will temporarily draw it at maximum size so that it is legible even from the overview position, and visually mark all shared instances of that word, as in Plates 1a and 2c.

Our new *pie flipper* interaction technique is a translucent, popup, minimal screen footprint visual interface that exploits the scrolling mouse. Holding a mouse button down and dragging the cursor into a slice of the radial display picks a category type, and then scrolling the wheel with the button still held down selects instances in that category. The popup display is shown in Plate 2b. Visual feedback is provided by both the selective highlighting visible in the main window through the translucent popup, and auxiliary information in a lower status bar. The sensory feedback of actively holding down the mouse button during scrolling minimizes mode errors [13], and the popup radial display was inspired by pie menus [2].

8 Discussion

The target users of our system were three computational linguists involved in the development of MindNet. This extremely small user population allowed us to count on using the latest hardware and gave us the opportunity to follow a user centered design approach as much as possible given their time constraints. The current version of Constellation has evolved in response to detailed user feedback on several previous paper and software prototypes. We relied on informal task analyses of the various iterations instead of formal user studies, since we had such a small target audience.

Our final layout algorithm is the result of many iterations as we explored the tradeoff between legibility and the semantic use of space on a finite resolution display. At the former extreme, we could tile the window with a rectangular grid containing 300 words, but there would be no spatial encoding whatsoever. At the latter extreme, a very strict spatial encoding would allow an exact encoding of the desired attributes, but vast amounts of navigation would be required to actually read anything because of low information density. Our horizontal plausibility gradient is a middle ground where more important words are allocated more room in the overview position. To compensate for our finite resolution, we offer easy navigation with animated transitions and intelligent zooming, where the relative amount of space devoted to words changes based on the zoom level. Rapid visual emphasis through hovering is useful in situations where navigation would be a cognitive burden.

The layout provides a great deal of structural information about the paths and definitions which were returned by a MindNet query, at the expense of many edge crossings. Our visual layering approach of using many perceptual channels in concert proved to be quite effective at both avoiding false edge attachments and visual emphasis. The psychophysical literature on color coding is extensive, and we benefited from it by following recommendations of Reynolds [10].

Our current implementation is a strong foundation, but further polishing could make it a more productive tool for our target users. Possibilities include adding incremental visual search capability, increased support for finding high-connectivity “hotspot” word constellations, and better integration with the main MindNet text views. The Constellation visualization was tuned to show the relationships between computed paths and their constituent definition graphs, so it only addresses a subset of the potential visualization needs of its target users. Additional visualization tools that support other tasks could be built.

9 Related work

Although there are several information visualization systems for visualizing relationships between or within documents, there are fewer aimed at showing the relationships between words in a semantic network.

The SemNet system [6] for visualizing a large knowledge database used in linguistics tackles a problem similar to our own. Fairchild et al advocate a 3D representation to avoid edge crossings, and present several approaches for placement including mapping functions, proximity placement and heuristics. Despite an effort to rely solely on syntactic information, they acknowledge the utility of exploiting semantic information to achieve a more effective visualization.

The Visual Thesaurus [3] is a more recent attempt to show connections between words. Its simple spatial layout algorithm does not scale past a few dozen words, and the constant motion of the scene makes it ill-suited for extended analysis.

There are many graph drawing systems of varying degrees of generality [4]. One of the more flexible and popular two dimensional layout systems is *dot* [7], which we found useful for prototyping some of our early ideas. However, its main focus is to provide a visually pleasing graph using only graph topology, with a strong emphasis on hierarchy and only minimal user steering possibilities. Our goal in Constellation was to explicitly use domain-specific information in the spatial layout.

10 Conclusion

Constellation is an interactive visualization system for linguistic queries. The key visual encoding choices are the spatial position of nodes and edges, and the creation of visual layers using several other auxiliary perceptual channels. We present a special-purpose graph layout algorithm that constructs a semantically structured space with a left to right plausibility gradient and a top to bottom path flow. The selective highlighting of boxes and edges creates a visually salient foreground layer distinct from the unobtrusive background layer. We exploit the scrolling mouse with a new *pie flipper* interaction technique for selecting instances of a constellation category. Intelligent zooming helps us achieve high word readability at global, intermediate, and local scales.

Although Constellation was designed for a small target audience, our design principles are relevant for many information visualization systems. The interactive visual emphasis capability is as fundamental to the dataset exploration as the interactive navigation, and these interactions are as important as the base spatial layout for understanding the full dataset.

11 Acknowledgments

We gratefully acknowledge the time and ideas of our target users Lucy Vanderwende, Bill Dolan, and Mo Corston-Oliver. Thanks to Mike Barnett for providing MindNet support. Mary Czerwinski contributed significantly to the user centered design and evaluation process. We also thank Pat Hanrahan, Terry Winograd, Maneesh Agrawala, and Chris Stolte for useful discussions. Diane Tang and Greg Humphreys provided comments on previous versions of this paper. This work was partially supported by a Microsoft Graduate Research Fellowship.

References

[1] Christopher Ahlberg and Ben Shneiderman. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Human Factors in Computing Systems CHI '94 Conference Proceedings*, pages 313–317, 1994.

[2] Jack Callahan, Don Hopkins, Mark Weiser, and Ben Shneiderman. An empirical comparison of pie vs. linear menus. In *Proceedings of the CHI'88 Conference on Human Factors in Computing Systems*, pages 95–100, May 15–19 1988. Washington, D.C.

[3] Plumb Design. Thinkmap Visual Thesaurus. http://www.plumbdesign.com/thesaurus

[4] Giuseppe Di Battista. *Graph drawing: algorithms for the visualization of graphs*. Prentice Hall, 1999.

[5] William B. Dolan, Lucy Vanderwende, and Stephen Richardson. Automatically deriving structured knowledge base from on-line dictionaries. In *Proceedings of the Pacific Association for Computational Linguistics*, April 21–24 1993. Vancouver, British Columbia.

[6] K. Fairchild, S. Poltrok, and G. Furnas. SemNet: Three-dimensional graphic representations of large knowledge bases. In R. Guindon, editor, *Cognitive Science and its Applications for Human-Computer Interaction*, pages 201–233. Lawrence Erlbaum, 1988.

[7] Emden R. Gansner, Eleftherios Koutsofois, Stephen C. North, and Kiem-Phong Vo. A technique for drawing directed graphs. *IEEE Transactions on Software Engineering*, 19(3):214–229, March 1993.

[8] Marc Green. Toward a perceptual science of multidimensional data visualization: Bertin and beyond. http://www.ergogero.com/dataviz/dviz0.html, 1998.

[9] Jock Mackinlay, Stuart Card, and Ben Shneiderman. Readings in information visualization: Using vision to think. Morgan Kaufmann, 1999. Chapter 1.

[10] Linda Reynolds. Colour for air traffic control displays. *Displays*, 15(4):215–225, 1994.

[11] Stephen D. Richardson, William B. Dolan, and Lucy Vanderwende. MindNet: acquiring and structuring semantic information from text. In *Proceedings of COLING '98*, 1998. To appear.

[12] George G. Robertson, Stuart K. Card, and Jock D. Mackinlay. Information visualization using 3D interactive animation. *Communications of the ACM*, 36(4):57–71, April 1993.

[13] A. Sellen, G. Kurtenbach, and W. Buxton. The prevention of mode error through sensory feedback. *Human Computer Interactions*, 7(2):141–164, 1992.

[14] Edward Tufte. *Envisioning Information*. Graphics Press, 1991.