First publ. in: Proceedings of the IEEE Symposium on Information Visualization (InfoVis ´02), 2002, pp. 145-151

Visual Unrolling of Network Evolution and the Analysis of Dynamic Discourse*

Ulrik Brandes
Dept. Computer & Information Science
University of Konstanz
Ulrik.Brandes@uni-konstanz.de

Steven R. Corman
Hugh Downs School of Human Communication
Arizona State University
steve.corman@asu.edu

Abstract

A new method for visualizing the class of incrementally evolving networks is presented. In addition to the intermediate states of the network it conveys the nature of the change between them by unrolling the dynamics of the network. Each modification is shown in a separate layer of a three-dimensional representation, where the stack of layers corresponds to a time line of the evolution. We focus on discourse networks as the driving application, but our method extends to any type of network evolving in similar ways.

1. Introduction

Evolving networks are networks that change over time, thus presenting a unique challenge to network visualization. They occur in diverse areas such as Web graph analysis, social network analysis, software engineering, and so on, and are typically visualized using methods from dynamic graph drawing [13, 5]. These methods usually produce a sequence of interdependent visualizations that represent intermediate states of the network and try to preserve the user's mental map [15]. Graph animation (see, e.g., [11]) is often used to lower the cognitive effort required to follow the transition from one visualization to the next.

If the actual changes occurring between consecutive states of an evolving network are an integral part of the data that should be open to analysis, it is certainly not sufficient to compute interdependent layouts and animate between them. To facilitate simultaneous analysis of state and change, we therefore propose a method in which the evolution of the network is unrolled and each step represented as a layer in a three-dimensional network visualization.

Our method was developed with a particular application in mind, the analysis of dynamic discourse, but it is straightforwardly applied to any type of network that evolves in similar ways. Discourse networks are introduced briefly in Sect. 2. Our approach to visualizing them, and some illustrative examples, are given in Sect. 3. In Sect. 4 we demonstrate the usage of our method on real-world data. Some conclusions are offered in Sect. 5.

*Research supported in part by a gift from Intel Corporation and by Deutsche Forschungsgemeinschaft (DFG) under grant BR 2158/1-1.

2. Evolving Discourse Networks

In this section, we extend a recently introduced method for static text analysis to dynamic discourse such as human conversations or newsgroup archives. The method is described in some detail to emphasize the need for a visualization method that displays the nature of change that leads to new states of an evolving network.

2.1. CRA of Static Text

Centering Resonance Analysis (CRA) [9] is a method of network text analysis [7] that is designed for the study of complex discourse systems. This encompasses a wide range of phenomena, including interpersonal conversations, group discussions, interaction in large organizations, the Internet and other mass media, as well as even larger social groups.

Though other text analysis methods could be used for studying discourse, CRA offers distinct advantages in that it is a representational method that does not rely on context-specific semantic rules, training sets, or predefined document collections [9]. It produces stand-alone, abstract representations of texts that can be analyzed alone or arbitrarily combined or compared with other CRA representations.

Let T be a grammatically correct text. CRA extracts a graph $G(T) = (V, E)$ from T in the following way. First, noun phrases that make up individual sentences are identified using linguistic analysis. Nouns[1] and adjectives that

[1] We here disregard pronouns because the identity of particular individuals is not especially relevant to the analysis. See [9] for more information on this choice and its implications.

"One of the **most essential branch**es of **English liberty** is the **freedom** of one's **house**. A **man**'s **house** is his **castle**; and whilst he is **quiet**, he is as **well guard**ed as a **prince** in his **castle**."

Figure 1. Sample text with highlighted noun phrases (stemmed words in boldface).

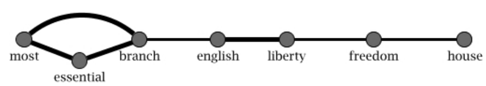

words linked by first sentence

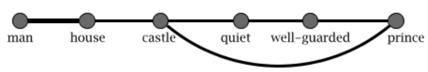

words linked by second sentence

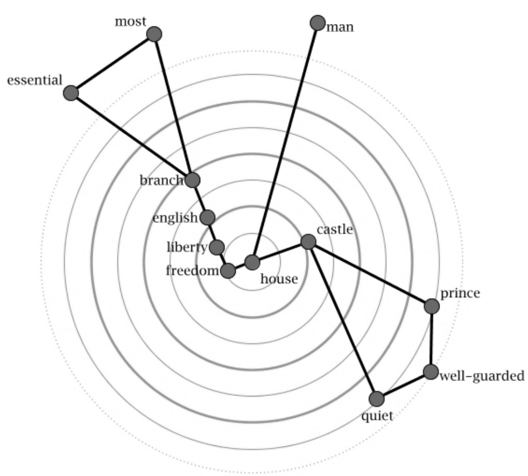

CRA network $G(T)$
according to which 'house', 'freedom', and 'castle'
are most influential (highest betweenness centralities)

Figure 2. CRA network extracted from text T in Figure 1. For each sentence, edges inside noun phrases are shown bold. When visualizing the merged (i.e., a single, static) network, concentric circles can be used to indicate levels of centrality [3].

make up these noun phrases form the set V, so that each vertex corresponds to a word that occurs in T. Words are considered linked if they co-occur inside noun phrases or occur on adjacent ends of consecutive noun phrases within a sentence. Accumulating these over all the sentences in a text yields a network of words comprising the subjects and objects of the text and how these are related to one another. The graph $G(T)$ is called the *CRA network* of text T. An example is shown in Figure 2.

2.2. Importance of Words

Competent writers and speakers deploy words strategically to create a sensible, coherent message. Therefore the structural positions of words in a CRA network $G = (V, E)$ reflect their importance or influence in structuring the meanings of the text. While many measures of structural position are possible, CRA uses a particular measure that reflects the extent to which a word is involved in chains of association between other words in the network. Define the *betweenness centrality* [1, 10], c_v, of a vertex $v \in V$ as

$$\sum_{s \neq v \neq t \in V} \frac{\sigma_G(s,t|v)}{\sigma_G(s,t)}$$

where $\sigma_G(s,t)$ and $\sigma_G(s,t|v)$ are the number of shortest paths between s and t and those that pass through v, respectively. Thus, a word with high betweenness is influential because it is greatly involved in channeling flows of meaning in the network. In the example in Figure 2 nodes are placed at a distance from the center that reflects their betweenness centrality, with nodes outside the outer circle having centrality equal to zero. On undirected graphs with n vertices and m edges, betweenness centrality is computed in time $\mathcal{O}(nm)$ [2].

2.3. Dynamic Discourse

CRA has so far been applied to the analysis of static texts. However, discourse is best thought of as dynamic, involving a sequence of speaking turns or a set of time-ordered written texts (e.g., e-mails, press releases). Advances in computer networking in the last decade have provided unprecedented access to sequences of texts, and advanced voice recognition promises to do the same for spoken communication [6].

The basic CRA framework can be straightforwardly applied to the analysis of dynamic discourse [9]. However, change is a primary feature of interest, so analyzing a CRA network of the entire discourse neglects what happens within turns or texts and the relationship of words between them and across sets of them. The following is an extension of CRA to a dynamic setting and referred to as Dynamic Centering Resonance Analysis (DCRA).

A *discourse* episode is defined as a sequence $T = (S_1, \ldots, S_k)$ composed of *slices* S_t which represent, e.g., a speaking turn or a serial item of written text. Simply deriving a CRA network $G(S_t)$ for each slice would ignore another crucial feature of discourse, namely the fact that humans have *conversational memory* [16] for more than one turn. A parameter d is therefore used to represent the *depth of conversational memory* appropriate in a given context. At any given point in time t, $1 \leq t \leq k$, we define the *state* of the discourse to be $T_t = (S_{t-d'+1}, \ldots, S_t)$ where $d' = \min\{d, t\}$, i.e. the text obtained from concatenating slices in the conversational memory. The importance of words is relative to the conversational memory, and hence calculated in each $G(T_t)$. Thus importance is dynamic and relative to the changing circumstances of the conversation. The sequence $G(T) = (G(T_1), \ldots, G(T_k))$ constitutes the evolving *discourse network*.

3. Unrolling Network Evolution

Evolving networks are characterized by the (in general gradual) change of the state of a network between consecutive points in time. Standard methods for dynamic graph visualization generate a drawing for each such state (taking into account the drawings of previous states) and animate between consecutive states, thus maintaining the user's mental map and smoothing the transition from one state to another. An example of such an approach is [18].

However, the nature of the change that led to a state cannot be recognized from that state's visualization. We propose a visualization method that allows us to simultaneously examine the state of a network and the changes that led to it by unrolling the sequence of events.

3.1. Evolutionary Cross-Sections

Networks evolving in discrete steps are composed of subnetworks that exist at particular points in time. Since the abstract information inherent to these networks consists of both the structure of subnetworks and changes between them, we propose a layered visualization design in which each newly introduced part of the network is shown in a layer of its own, and the composition of consecutive layers represents the corresponding state of the network. In other words, we unroll the dynamics of networks into its constituent layers.

In dynamic discourse, each slice induces a modification of the current state through the addition of edges, the potential creation of vertices, and the deletion of edges and vertices that drop out of conversational memory. The layer introduced into our visualization at time t therefore shows only the new elements introduced by the corresponding slice, i.e. $G(S_t)$. By displaying only layers for slices in

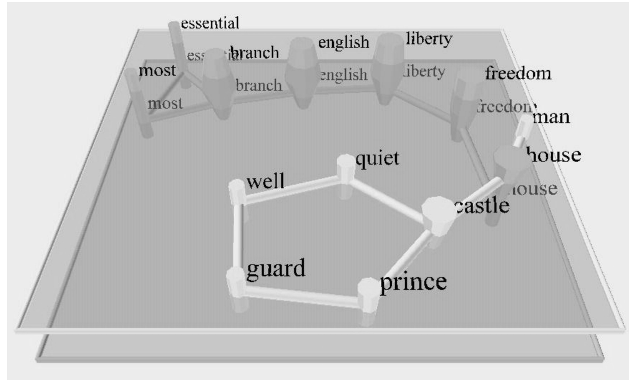

(a) layers, colors: dynamic discourse slice-by-slice

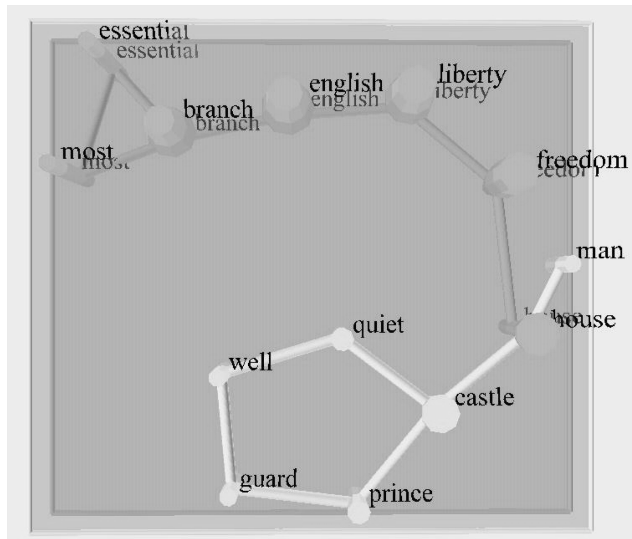

(b) transparency: fading from conversational memory layout: separation of topics

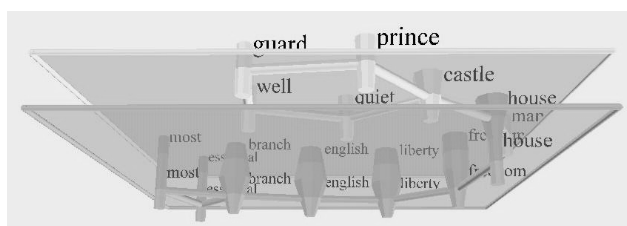

(c) column thicknesses: changing importance of words

Figure 3. Visualization design illustrated with the quote from Figure 2, which is treated as if sentences were from different speakers. The sentences have almost disjoint subjects, but are linked through the topic 'house'. Hence, the addition of the second sentence causes 'house' to become more influential while decreasing the importance of most other words with respect to the first layer.

the current conversational memory, and by making them semi-transparent, we obtain a visual representation that can be viewed from the top to infer the entire state $G(T_i)$ of the conversation, but also shows the contribution of slices $S_{t-d+1}, \ldots, S_t$ with past slices slowly fading away. In Figure 3, each sentence of the quote in Figure 2 is interpreted as a slice.

3.2. Colors and Sizes

Several aspects of a layered visualization can be varied to enrich its information content. In the context of discourse networks, the most interesting additional information are the speakers in each turn, and the changing importance of words.

We therefore assign a distinct color to each speaker and border each layer according to whose turn it represents. Edges are colored according to the speaker inducing them. Finally, we use the color of the speaker that introduces a word into the conversational memory for the corresponding vertex. This is illustrated in Figure 3(a).

As outlined in Sect. 2.2, the current importance of a word is captured by its betweenness centrality in the current state of the conversation. The change in centrality is conveyed by varying the thickness of columns that represent vertices. Let c_- and c_+ be the minimum and maximum centrality score over the entire discourse, then the radius of a vertex v with centrality $c_v(t)$ at time t is defined as

$$r_- + \frac{c_v(t) - c_-}{c_+ - c_-} \cdot (r_+ - r_-)$$

where r_- and r_+ denote the minimum and maximum desired thickness of vertices, so that the entire interval is actually used. In a top-down view, this results in central words being larger than peripheral words. Figure 3(c) shows how, when viewed from the side, the changing thickness of a column depicts the centrality profile of a word.

To reduce the complexity of large data sets, an importance threshold can be specified to exclude less important words from the visualization and emphasize the main structure of the discourse.

3.3. Layout

Any graph layout algorithm can be used to determine the layout of each layer, and in fact the algorithm should be chosen in close accordance with the substantive background of the application generating the graph. In principle, there is no look-ahead in a conversation and the layout of each layer should be contingent only on the layout of preceding layers [4]. However, our column-like representation of vertices prevents subsequent adjustments of positions, so that only newly introduced vertices can be placed freely, while all others are fixed at their position in previous layers.

Since this highly constrained, serial layout results in poor readability of later layers, we take a different approach. At least today, dynamic discourse analysis is not carried out in real time, so that all slices of a discourse are available from the beginning. We can therefore compute a layout of the CRA network $G(T)$ of the entire discourse $T = (S_1, \ldots, S_k)$, and use its vertex positions for the layout of all states $G(T_t)$, $1 \leq t \leq k$.

As our global layout algorithm we choose a spring embedder variant that corresponds to metric multidimensional scaling of pairwise distances [14], because the results graphically support the intuition of words being close or distant according to the structure of the discourse. See Figure 3(b) for an example. The objective is to minimize the squared differences between distances in the layout and distances in the graph, weighted by the square of that distance. For $G(T) = (V, E)$ and vertex positions $(x_v)_{v \in V}$ we want to minimize

$$\sum_{u,v \in V} \frac{\left(\|x_u - x_v\| - d_{G(T)}(u,v)\right)^2}{d_{G(T)}(u,v)^2}$$

where $d_{G(T)}(u,v)$ is the length of a shortest path between u and v in $G(T)$. This is an $\mathcal{NP}$-hard task, so that gradient methods are typically used to obtain a local minimum. A straightforward implementation runs in time $\mathcal{O}(n^2)$ per iteration, but typically few iterations are required.

An extension in which vertices are added incrementally is introduced in [8] and shown to have empirically better performance. In our setting this method is particularly appropriate, because we can introduce vertices one slice at a time and thus combine the advantages of global and incremental layout according to the evolution of the graph.

4. Application Examples

To evaluate our approach for exploratory analysis of evolving networks, we are implementing a visualization system using Java and the Data Structures Library in Java (JDSL) [17]. Our prototype produces interactive visualizations through the Extended Application Interface (EAI) for VRML, so that users can move forward and backward in the discourse and rotate, zoom, and pan the scene to view it from any point and orientation.

The following are examples in which the system is used to support Dynamic Centering Resonance Analysis (DCRA) on recorded conversational data. Any other type of incrementally evolving networks, in which states are defined by cross-sections, would be appropriate as well. While the first example is crafted, the other two are extracts of real discourses.

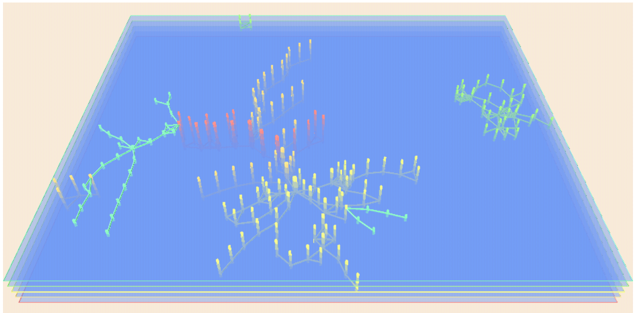

Figure 4. Incoherent discourse.

4.1. Incoherence

An important feature of discourse is coherence. If we have a sequence of random statements that make sense in themselves but have nothing to do with one another, these cannot form a coherent body of discourse. Visual DCRA should be able to identify these cases, and not create a lot of artifacts that make them appear integrated or sensible. To test for this we created a set of five arbitrary clips from unrelated sources and treated these as slices of a discourse. The results are shown in Figure 4. These clips clearly occupy distinct areas in the layout and colors representing them are mostly confined to their own slices. There are only a few cases where the slices connect due to chance common occurrence of the same word.

4.2. Two-Party Conversation

The first recorded example demonstrates the ability of our visualizations to indicate significant interventions by participants in a conversation. We analyzed a transcript of a video-taped conversation between Osama bin Laden and a man identified as "Shaykh", which was released in December 2001.[2] The conversation begins with pleasantries and is structured almost exclusively by "Shaykh". After several turns, bin Laden makes an abstract reference to Koranic stories, thus marking a topic shift to reminiscing about the World Trade Center attacks. The DCRA visualization in Figure 5 shows this transition clearly. The red nodes introduced by "Shaykh" can still be seen on recent layers, but the top layer has words introduced by bin Laden, shown in yellow, that connect to only two of "Shaykh"'s words ('allah' and 'people'). Otherwise, new words are introduced and connections are mostly between them.

[2] The transcript has been published by CNN at `http://www.cnn.com/2001/US/12/13/tape.transcript`. The first segment of the transcript was subjected to DCRA, with small disconnected components of the network (paths of length at most four) have been removed and evaluated separately.

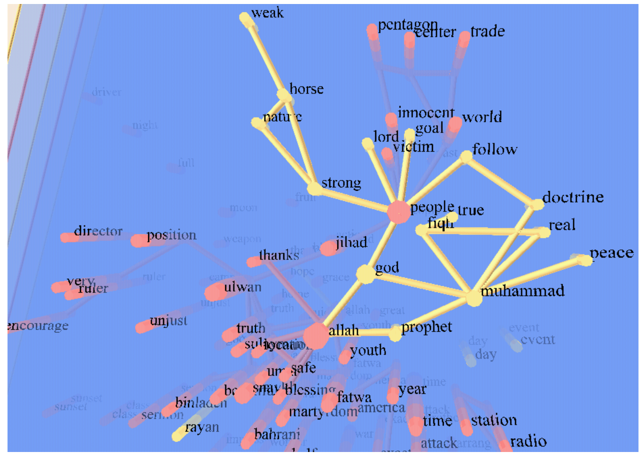

Figure 5. Bin Laden initiating a topic shift.

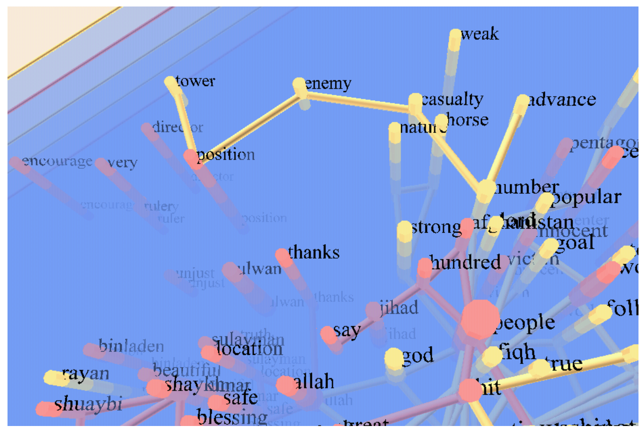

Figure 6. "...we calculated in advance the number of casualties from the enemy..."

The examples in Figures 6 and 7 display other crucial excerpts from the conversation. While the first is probably the most frequently cited statement in the conversation, the second is an example of how a theme is picked up and reflected upon from a different angle.

4.3. Group Discussion

A different kind of change in discourse can be described more as gradual movement than wholesale change. Here we are interested in processes where an idea expressed at one point in a conversation is linked to an idea at a subsequent point in the conversation, and so on in a chaining-out fashion. In a DCRA visualization, such a structure appears as a staircase-like configuration. A brief real-life example of this form comes from a transcript of a graduate seminar discussion on group decisions support systems (GDSS). Figure 8 shows one part of the discussion where the issue

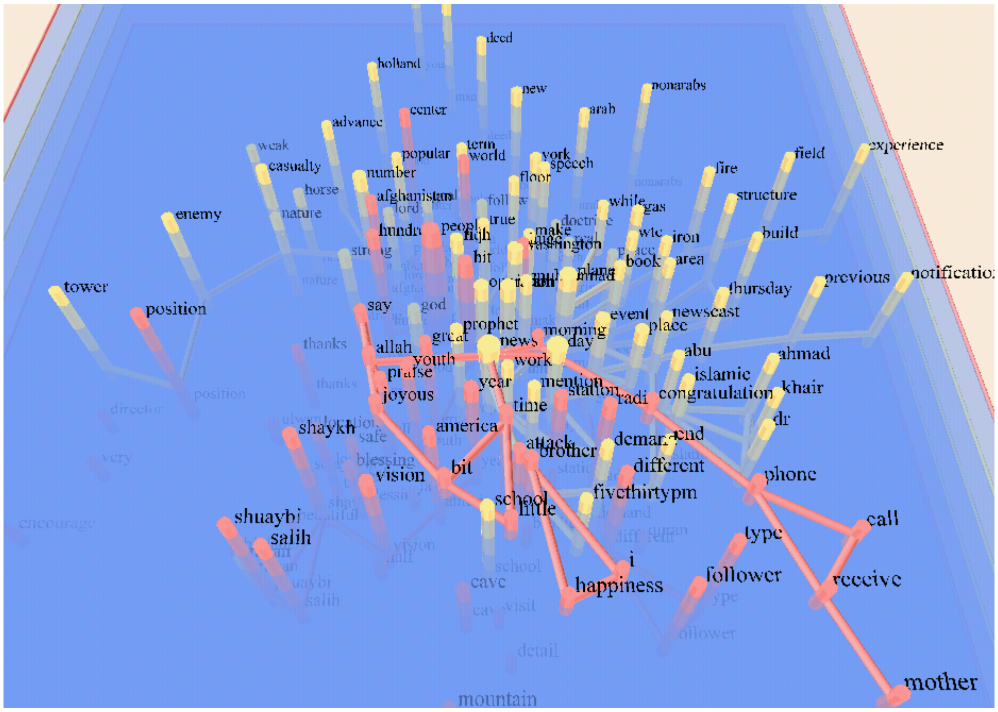

Figure 7. After bin Laden has completed his account of the 'event', "Shaykh" reflects on coverage in the 'news' and other reactions.

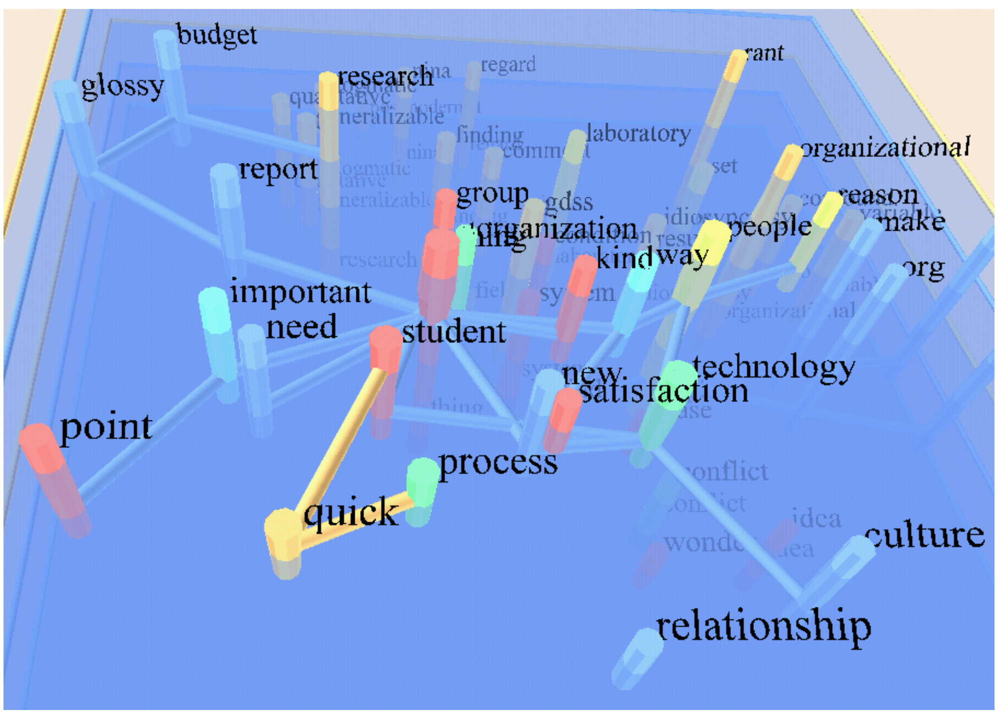

Figure 8. Topical movement in group discussion illustrated by a "staircase" structure: 'technology'–'student'–'quick'–'process'.

was whether results from GDSS studies using students are valid for real organizational settings.

In the slice below the surface, a student had made a rather complicated point about how students were unlikely to feel the same pressure to make quality decisions using the technology as in real organizations. On the top slice the instructor draws out the point that students are more willing to engage in quick processes to get their decisions made. This forms a small staircase, 'technology'–'student'–'quick'–'process', representing a movement in the conversation from students' relationship with the technology to their relationship with a particular kind of process.

5. Conclusions

We have presented a visualization approach for evolving networks and applied it to visualize networks created by Dynamic Centering Resonance Analysis (DCRA).

Based on the above examples, we conclude that visualizations of DCRA are capable of indicating prominent features of discourse. They can show major interventions where a participant in a conversation changes the direction of a conversation. They can also show the more subtle movement involved in chaining ideas across turns. Finally, they can indicate cases where discourse is incoherent. Assessment of its potential for detecting more complex

features of discourse awaits further research. For instance, it is planned to apply visual DCRA to court-room negotiations to identify or highlight strategies used by lawyers in particular cases. Potential other application include newsgroup postings, organizational communication, and press releases. The sequence of importance scores for topics generated by DCRA can also be used as input for visualization approaches that do not depict network structure (see, e.g., [12]).

Visual unrolling exploits a characteristic property of dynamic discourse and similar types of evolving networks, namely that change occurs only through the addition of new edges and vertices, and the deletion of the least recently introduced vertices and edges in the current state. It will be interesting to extend our method to types of networks that evolve with arbitrary deletions.

Another direction for future research is the extension to real-time situations, e.g. by using dynamic graph layouts methods that produce suitable layouts of layers subject to many fixed positions inherited from earlier layers, or by relaxing the latter constraint and allowing vertex-representing columns to bend.

Acknowledgments. We thank Odilo Oehmichen for implementing large parts of the DCRA visualization prototype.

References

[1] J. M. Anthonisse. The rush in a directed graph. Technical Report BN 9/71, Stichting Mathematisch Centrum, Amsterdam, Oct. 1971.

[2] U. Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.

[3] U. Brandes, P. Kenis, and D. Wagner. Communicating centrality in policy network drawings. *IEEE Transactions on Visualization and Computer Graphics*, 2002. To appear.

[4] U. Brandes and D. Wagner. A Bayesian paradigm for dynamic graph layout. In *Proc. 5th Intl. Symp. Graph Drawing (GD '97)*, LNCS 1353:236–247. Springer, 1997.

[5] J. Branke. Dynamic graph drawing. In M. Kaufmann and D. Wagner, editors, *Drawing Graphs: Methods and Models*, LNCS Tutorial 2025:228–246. Springer, 2001.

[6] E. W. Brown, S. Srinivasan, A. Coden, D. Ponceleon, and J. W. Cooper. Toward speech as a knowledge resource. *Knowledge Management*, 40(4), 2001.

[7] K. M. Carley. Network text analysis: The network position of concepts. In C. W. Roberts, editor, *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences from Texts and Transcripts*, pages 79–100. Lawrence Erlbaum Associates, 1997.

[8] J. D. Cohen. Drawing graphs to convey proximity: An incremental arrangement method. *ACM Transactions on Computer-Human Interaction*, 4(3):197–229, 1997.

[9] S. R. Corman, T. Kuhn, R. D. McPhee, and K. J. Dooley. Studying complex discursive systems: Centering resonance analysis of communication. *Human Communication Research*, 28(2):157–206, 2002.

[10] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40:35–41, 1977.

[11] C. Friedrich and M. E. Houle. Graph drawing in motion II. *Proc. 9th Intl. Symp. Graph Drawing (GD '01)*, LNCS 2265:220–231. Springer, 2002.

[12] S. Havre, E. Hetzler, P. Whitney, and L. Nowell. ThemeRiver: Visualizing thematic changes in large document collections. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):9–20, 2002.

[13] I. Herman, G. Melançon, and M. S. Marshall. Graph visualization and navigation in information visualization: a survey. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):24–43, 2000.

[14] T. Kamada and S. Kawai. An algorithm for drawing general undirected graphs. *Information Processing Letters*, 31:7–15, 1989.

[15] K. Misue, P. Eades, W. Lai, and K. Sugiyama. Layout adjustment and the mental map. *Journal on Visual Languages and Computing*, 6(2):183–210, 1995.

[16] L. Stafford, V. R. Waldron, and L. L. Infield. Actor-observer differences in conversational memory. *Human Communication Research*, 15:590–611, 1989.

[17] R. Tamassia, M. T. Goodrich, L. Vismara, M. Handy, G. Shubina, R. F. Cohen, B. Hudson, R. S. Baker, N. Gelfand, and U. Brandes. JDSL: The data structures library in Java. *Dr. Dobb's Journal*, 323:21–31, April 2001.

[18] K.-P. Yee, D. Fisher, R. Dhamija, and M. Hearst. Animated exploration of dynamic graphs with radial layout. *Proc. IEEE Symp. Information Visualization 2001 (INFOVIS '01)*, pages 32–50, 2001.