Visual Overviews for Discovering Key Papers and Influences Across Research Fronts

Aleks Aris and Ben Shneiderman

Department of Computer Science and Human–Computer Interaction Lab, University of Maryland, College Park, MD 20742. E-mail: {aris, ben}@cs.umd.edu

Vahed Qazvinian and Dragomir Radev

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109. E-mail: {vahed, radev}@umich.edu

Gaining a rapid overview of an emerging scientific topic, sometimes called *research fronts*, is an increasingly common task due to the growing amount of interdisciplinary collaboration. Visual overviews that show temporal patterns of paper publication and citation links among papers can help researchers and analysts to see the rate of growth of topics, identify key papers, and understand influences across subdisciplines. This article applies a novel network-visualization tool based on meaningful layouts of nodes to present research fronts and show citation links that indicate influences across research fronts. To demonstrate the value of two-dimensional layouts with multiple regions and user control of link visibility, we conducted a design-oriented, preliminary case study with 6 domain experts over a 4-month period. The main benefits were being able (a) to easily identify key papers and see the increasing number of papers within a research front, and (b) to quickly see the strength and direction of influence across related research fronts.

Introduction

Researchers and analysts are often called on to learn about *research fronts*, which are defined as "clusters of documents that tend to cite a fixed, time invariant set of base documents" (Morris, Yen, Zheng, & Asnake, 2003; Thomson Corporation, 2008). These researchers and analysts need to know about emerging or hot topics to contribute to decisions about research support from government agencies or choices over resource allocation in companies and universities. They need to quickly learn about the new ideas, key personnel, influential labs, current controversies, and future directions.

The current pressure for interdisciplinary (sometimes called "cross-cutting") work increases the need to rapidly learn about research fronts. Furthermore, the large volume of scientific papers presented at many conferences and journals makes it difficult to track research fronts. Finally, the problem

is exacerbated by volatile terminology, which may not fit neatly into existing taxonomies and subject-heading lists.

Potential sources of reliable information on connections across ideas are the citations in key papers. These have author-chosen reliability, thereby providing credible clues such key papers. As the researchers and analysts begin to make insights about an emerging topic, they can strengthen their understanding by pooling citations across a set of papers. Their insights might include identification of a few key papers that most authors cite, recognition of productive research teams, and identification of high activity levels by publication frequency.

This article reports on a preliminary case study that applied a novel information-visualization tool to help researchers explore citations between scholarly publications. Our emphasis was on assessing the efficacy of our novel layout strategy and identifying potential improvements to the user interface. The next section describes previous work while the third section briefly discusses evaluation methods. The fourth section reports on the case study, which is followed by a final section that provides limitations and future work.

Previous Work

Algorithmic citation analysis and natural language summarization have a long history (Newman, 2001; Teufel & Moens, 2002; Teufel, Siddharthan, & Tidhar, 2006). Bradshaw (2002, 2003) demonstrated benefits from citations to determine the content of articles and introduce "Reference Directed Indexing" to improve the results of a search engine.

Going deeper into textual analysis, Nanba and colleagues analyzed citation sentences to build a tool for survey generation (Nanba, Abekawa, Okumura, & Saito, 2004a; Nanba, Kando, & Okumura, 2004b). They also discussed the same citation categorization to support a system for writing a survey (Nanba & Okumura, 1999). Elkiss et al. (2008) showed the importance of the citation sentences in understanding the contributions of an article. This key observation was then put

to work by extracting the set of citation sentences to produce a summary of an article's contributions (Qazvinian & Radev, 2008).

To show changes in citation networks over time, one project used the temporal evolution of node degree, and then applied the expectation-maximization (EM) clustering algorithm to identify communities of related papers (Leicht, Clarkson, Shedden, & Newman, 2007). Another approach to studying change and diversity is to use topic entropy as a measure of the rising or falling of topics and similarity across conferences (Hall, Jurafsky, & Manning, 2008).

While statistical methods are helpful, their value can be increased by proper information visualization that shows unexpected trends, clusters, gaps, outliers, or other features as well as contextual clues that may be valuable to domain experts (Card, Mackinlay, & Shneiderman 1999; Chen, 2004a,b; Thomas & Cook, 2005). Visualization of citation patterns for a research front is a growing theme, most notably with Garfield's (2004; Garfield, Pudovkin, & Istomin, 2003a, 2003b) historiographs that are based on Sugiyama-style layouts (Sugiyama, 1987; Sugiyama, Tagawa, & Toda, 1981). These show old papers as nodes on the top and recent papers on the bottom, with arrows going upward to show citations. Nodes are arranged in horizontal rows by years. The validation of utility was only by personal case studies in which Garfield and his colleagues (2003a,b) show examples of topics, such as Watson-Crick's unraveling the structure of DNA in 1953. While the examples are appealing and understandable, they focused on tracing the antecedent papers for an important publication. This is fine for retrospective analysis, but less effective since it is difficult to identify the important publication in an emerging research front. Furthermore, the largest case study published showed only 34 papers.

Other attempts at visual overviews of paper collections have focused on showing temporal patterns using horizontal time lines with nodes for older papers on the left (Morris, Yen, Zheng, & Asnake, 2003). The DIVA software shows 10 to 15 horizontal lines organized by a dendrogram to group related research fronts. Each time line shows the papers related to a research front as circular markers, jittered in the *y* direction to reduce occlusion, with size coding to show numbers of citations. Links can be selectively shown, detailed information on papers can be obtained, and rich bibliometric data can be extracted.

The contributions of key authors to four human–computer interaction conferences has been traced with multiple visualizations that show the evolution of topics, identify key papers, and reveal coauthorship patterns (Henry, Goodell, Elmqvist, & Fekete, 2007). An analysis of citation patterns within three subtopics in terrorism research (335 papers on body injuries in terrorist bombings, healthcare in response to biological and chemical weapons, and psychological and psychiatric impacts of September 11, 2001) showed tightly related clusters while the visualization revealed five crossover citations (Chen, 2006). Some visual overviews have gone beyond citation diagrams and have shown the progress of scientific arguments within a research front (Buckingham

Shum, Uren, Li, Domingue, & Motta, 2003; Carr 2003; Uren, Buckingham Shum, Bachler, & Li, 2006). Since current natural language processing technology is not capable of automatically recognizing claims, they must be extracted by hand, but the resulting diagrams are richly informative about controversies and progress.

More ambitious efforts to map all scientific literature, sometimes called *knowledge domain visualization*, have gained increased attention in recent years (Chen, 2003; Small, 2005). The motivations are varied: understanding the size and relationship of scientific disciplines such as physics, chemistry, or biology; discovering bridging papers that connect two subdisciplines; or identifying emerging interdisciplinary subdisciplines that might grow rapidly (Börner, Chen, & Boyack 2003; Chen, 2004a,b; Ginsparg, Houle, Joachims, & Sul, 2004; Morris & Martens, 2007; Shiffrin & Börner, 2004).

Evaluation Methods

One of the challenges of research in information visualization is to assess the value of a given tool for a given task (Plaisant, 2004). While controlled experimental methods have been used in numerous human–computer interaction projects, these methods are typically 30- to 120-min studies with novice users that focus on a specific input device, interactive widget, or display strategy. From the 132 papers of five recent information-visualization conferences, only 39 papers had any user evaluation, with the longest test being for 2 hr of usage on a single day (Perer & Shneiderman, 2008). Similarly, in the cited works in bibliometrics and knowledge-domain visualization, we did not find any reports from users other than the authors.

In recent years, researchers have become interested in studying how experienced professionals change their working strategies as they gain familiarity with a tool over weeks and months (Shneiderman & Plaisant, 2006). Instead of controlled experiments that measure time to completion for benchmark tasks, interest has grown in insight-based evaluation (Saraiya, North, & Duca, 2005) and longitudinal case studies with domain experts doing their own work over days and weeks. These case studies build on ethnographic methods that depend on close participant observation to capture the struggles users have and their successes. It requires that the experimenters have close knowledge of the work of the domain experts, and engage with them closely, possibly making suggestions and asking leading questions. This violates traditional experimental rules that keep a greater distance between experimenter and participant; however, the benefit is potentially deeper insights to the strengths and weaknesses of a proposed tool with rapid, specific guidance for improvement. These formative evaluation strategies can be complemented by more focused controlled studies of specific design features.

Preliminary Case Study With 6 Domain Experts

This article reports on a preliminary case study that applied the novel information-visualization tool Network

Visualization by Semantic Substrates (NVSS) to aid computational linguists working on summarization of emerging topics (Aris & Shneiderman, 2007; Shneiderman & Aris, 2006). NVSS was developed in close collaboration with political scientists studying citation patterns among the Supreme Court and circuit and district courts for the specific legal issue of regulatory takings over various periods of time. The NVSS project was taken on to test the value of showing network data using meaningful and stable node-layout strategies that differ from the commonly used force-directed approach. Force-directed approaches place nodes based on their linkage, which can be helpful in showing clusters. Unfortunately, such layouts undermine the utility of spatial cognition that assumes that left–right or up–down has meaning. Furthermore, force-directed approaches are unstable, in that the addition/deletion of a single node can produce a dramatically different display. Two network-visualization tools that apply force-directed approaches were tried with the dataset presented in this article, but they produced a tangle of nodes and links that made it difficult to discern meaningful relationships among research fronts.

With NVSS, users usually can define two to five regions to attain effective layouts and assign nodes to a region based on the node attributes. For example, in a social network, women might be placed in one region and men in another. Then, other node attributes are used to locate the nodes in a region (e.g., age might be used for the x axis and height for the y axis). In short, NVSS allows users to place nodes in scattergramlike regions with meaningful axes, and then show the links within and across regions. These semantic substrates might be described as attribute-based layouts in which nearby nodes have similar attribute values, which helps to easily spot outliers, and layouts are stable after additions/deletions. Of course, other problems such as node occlusion and sparse layouts, which waste screen space, are possible.

NVSS gives users rich control over link visibility with filters on x- or y-axis attributes to reduce clutter and allow users to follow links from source to destination. Node size can be determined by yet another attribute such as a person's weight.

The NVSS experts included 2 of the authors (Aris and Shneiderman) of this study, who had earlier developed NVSS in conjunction with political scientists. This case study was intended to study how well the ideas developed with one community of domain experts would benefit other domain experts who were not involved in the development process. The domain experts who were participants in this preliminary case study included 2 of the other authors (Qazvinian and Radev) of this study, whose role was to facilitate data gathering, suggest useful layouts, and develop meaningful questions. Other participants included 4 domain experts (Bonnie Dorr, Jimmy Lin, Judith Klavans, and Saif Mohammad) conducting research on topic summarization in a project supported by the National Science Foundation. The transfer of user interface design ideas from one community of domain experts to another is a topic that deserves further study.

First Session

The case study participants were highly experienced with data analysis either in information retrieval or in natural language processing. Some of our participants also were experienced in using visualization tools. We began the familiarization process by showing previous work on Supreme Court as well as circuit and district court cases, and then moved to a small dataset that contained publications from the Computer Science literature about treemaps (Shneiderman, 1992). The case-study participants had little knowledge of the courts or treemap literature. This 30-min demonstration generated interest in visualization methods as well as immediate suggestions for refinements such as larger nodes and jittering strategies to reduce clutter.

Second Session

A few weeks later, the 6 participants were shown another small dataset on a topic of strong interest to this group. The dataset consisted of 27 papers and 39 citations among them on Phrase-Based Machine Translation (PBMT) (Figure 1). These more familiar data were extracted from the Association for Computational Linguistics (ACL) Anthology, which is a collection of 11,000 papers from computational linguistics journals and proceedings from ACL conferences and workshops. To produce the ACL Anthology Network (AAN), Joseph and Radev (2007) preprocessed the data by parsing references and then built the citation and author collaboration networks. Further effort was required to produce a properly formatted dataset, with citations only to other papers in this dataset.

The 90-min session included a presentation and an open discussion. Participants' opinions and questions as well as their actions on the NVSS visualization tool were noted. They wanted to explore this topic to understand the citation patterns, identify the key papers, and follow the changing terminology in these topics. Within minutes, they confirmed their knowledge of two key papers in 2002 and 2003, and then explored the more recent papers to see whose work was included. Once the basic principles were understood, their enthusiasm was apparent.

Figure 1 shows the 27 PBMT papers in the NVSS. All papers are of type PBMT and are grouped into a region with the label "PBMT Papers" using a simple, one-region semantic substrate. The *Year* attribute is used to place the nodes from left to right along the x axis. The y axis is not used to represent an attribute, so the nodes are simply spaced evenly along the y axis. The control panel on the right-hand side allows users to control the visibility of links. In Figure 1, the Ranges & Filters tab is visible with its four sections. The first section, REGIONS, lists all the regions and how many nodes it contains. The second section, LINKS, lists by source and destination regions all the links. The third section, MODE, provides node-aggregation modes, which are not used in this example. The fourth section, RANGES, lists all the filters by placement attribute. Since *Year* is used as a placement attribute on the PBMT Papers region, a filter is provided for
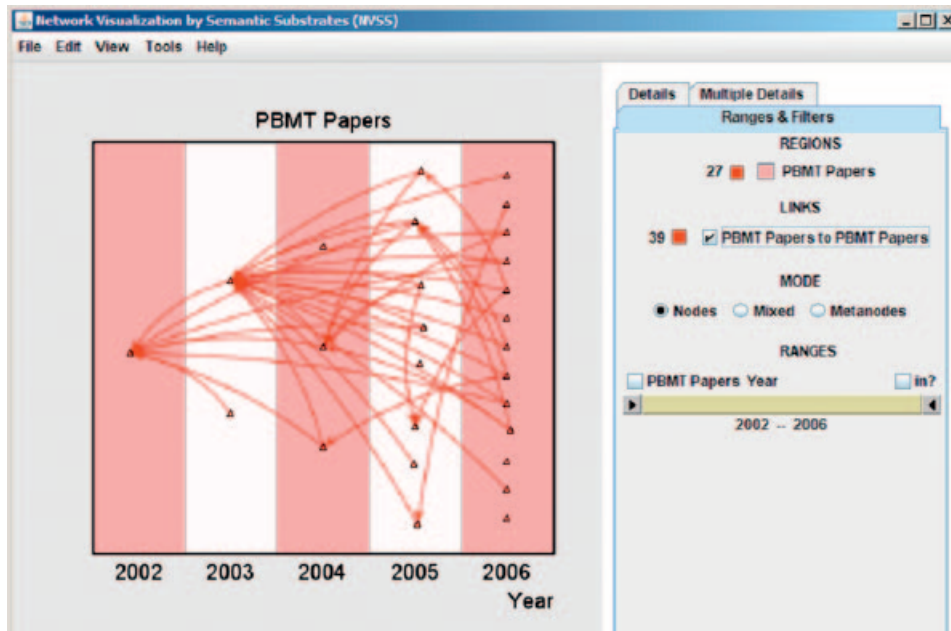
FIG. 1.   PBMT papers in NVSS with Year attribute on the x-axis. From this users can see that two seminal papers in 2002 and 2003 are widely cited in the community.
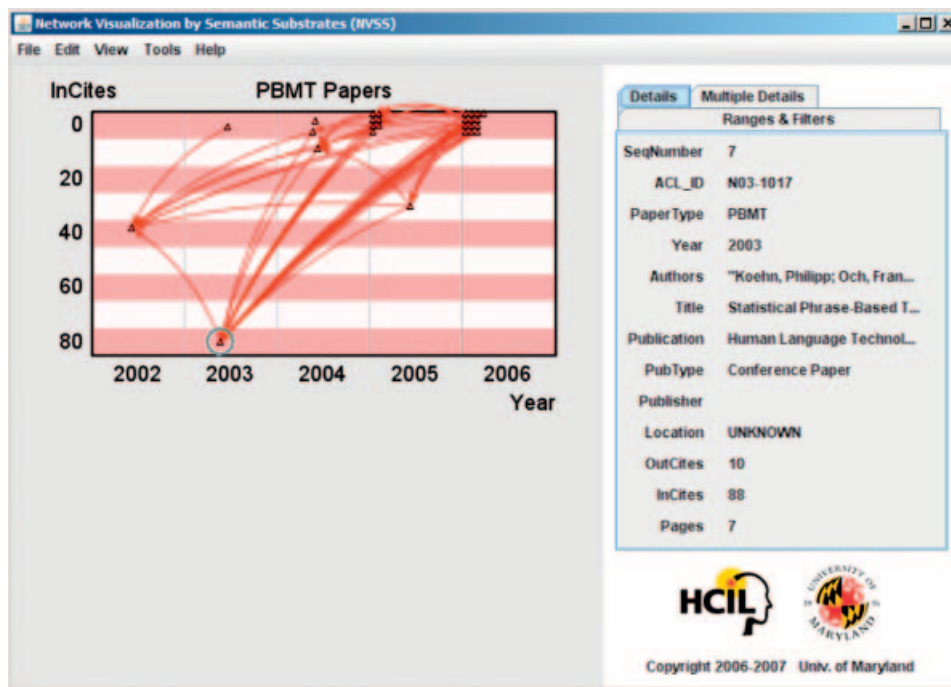

FIG. 2.   PBMT papers with highlight on key paper by Koehn et al. (2003). Using the InCites attribute, the number of citations a paper receives in the larger dataset is used to place nodes along the y axis, which helps easily spot the key paper by Koehn et al.

it under the RANGES section to filter the links according to *Year* values in that region.

While the simple version in Figure 1 spreads the PBMT papers on the y axis for link visibility, Figure 2 uses the number of incites, which is the number of citations a paper has received from other papers in the larger dataset, on the y axis to separate highly cited papers (lower on the figure) from the others. The key paper by Koehn, Och, and Marcu in 2003[1] (the one enclosed in a circle) and citations to it are salient.

---

[1]The full citation of the paper is: Koehn, P., Och, F.J., & Marcu, D. (2003). Statistical phrase-based translation. In Proceeding of the Human Language Technology Conference and Meeting of the North American Association for Computational Linguistics (pp. 48–54). Stroudsburg, PA: NAACL.
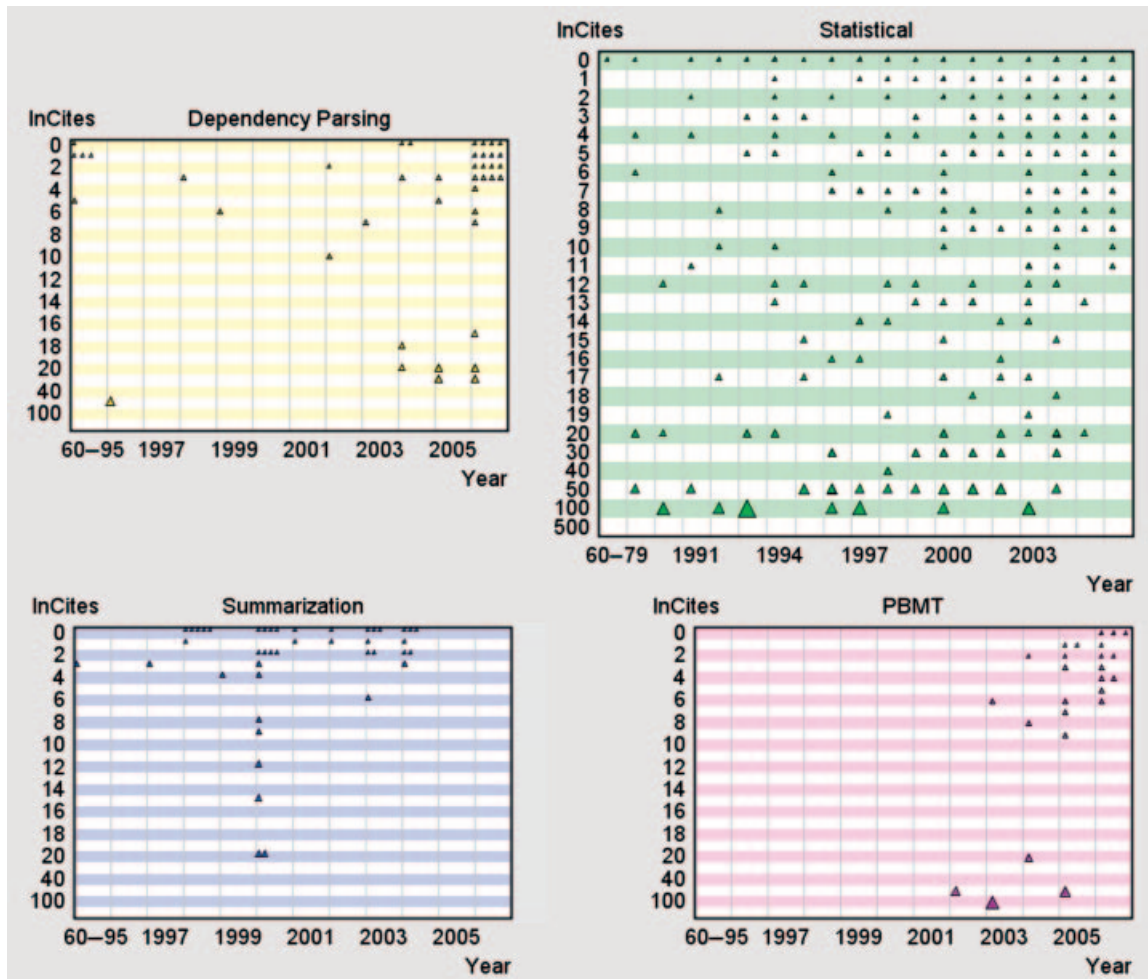
FIG. 3. The paper citation dataset with four topics without the links (citations) between them. This provides a general overview of the dataset in terms of topic (divided into four regions), year of publication (x axis) using the Year attribute, and the number of incoming citations from other papers in the larger dataset (y axis) using the InCites attribute.

Figure 2 shows the "Details" tab on the control panel on the right-hand side. The paper by Koehn et al. in 2003 is clicked, and the node-attribute values for that node appear in the "Details" tab on the control panel.

At the end of this second session, the 6 case-study participants chose four computational linguistics topics for exploration. The authors extracted the necessary papers from the AAN: PBMT (27 papers), statistical or machine learning methods (371), dependency parsing (46), and summarization (43). Each topic includes papers with the specific phrase in the title or content.

Third Session

Three months later at the third session, the participants were familiar with the visual layout. This session lasted approximately for 1½ hr. The session format was similar to the previous one; however, this time more structure was added, with several questions to ask to the participants. The iOpener group eagerly explored the dataset to discover relationships among the four topics (Figures 3–6). To guide their exploration and gain experience with the user interface

controls, we provided a set of progressively more complex questions. These questions are not a complete set of use cases or user needs, but they begin to clarify the goals for visualizations of research fronts (described as "topics" in the questions):

1. View the distribution of documents (nodes only) (no citation links).
   - Which topic has the longest duration?
   - Which topic has peaked and has fewer documents?
2. View citations/links only within the four topics.
   - What are the early key documents in each topic?
3. View citations only from the 3 small topics to the big topic.
   - Which topics depend most heavily on statistics?
4. View citations only among the 3 small topics.
   - What relationships are there across these 3 small topics?
5. View citations only from the big to the 3 small topics.
   - Which topics provide input to the big statistics topic?
6. View the many citations within the big statistics topic (You'll need to filter to reduce complexity.)
   - What are the key documents?
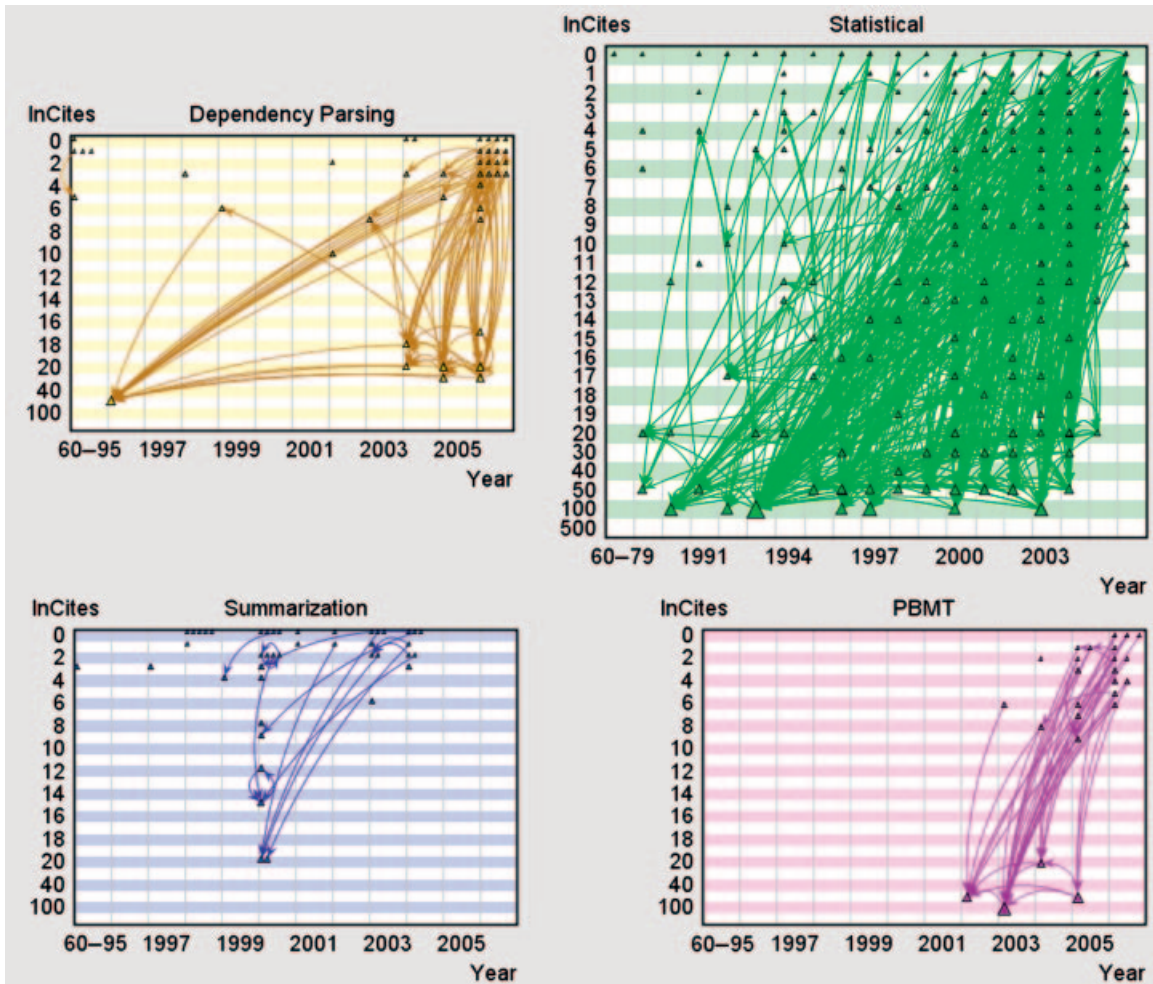   - What is the temporal change in citation patterns?

FIG. 4. Links only within each research front (region) are enabled. The large number of links in the Statistical Machine Translation (Statistical) region shows a mature and well-defined research front. The remaining three research fronts are still developing and show many fewer links.

After exploring the visualization, the 6 participants were unanimous in reporting that they better understood the history of the four topics (Question 1). They reported that visualizing the data on a single display, compared to talking about it in abstract terms or submitting multiple database-search queries, was a substantial advantage. They believed that the visual representation clarified the size of the research fronts and left a memorable visual pattern.

Their responses were encouraging, as the design of the semantic substrates (Figure 3) was intended to emphasize the greater size of the Statistical topic and allow users to discover the strength of relationships among the topics. As in the previous PBMT example, the x axis represents years while the y axis represents the number of citations a paper has received (inCites). In this substrate, there is a difference on the x axis from the previous ones in that there is custom binning enabling the first column of the small region to represent the years 1960 to 1995, and the first two columns in the Statistical region represent the year ranges 1960 to 1979 and 1980 to 1989, respectively. All other columns represent the following single years.

The next step for the participants (Question 2) was to select the check boxes to reveal the citation links within each region (Figure 4). The distinct colors of the regions for each topic were repeated with distinct colors for the links within each region. The region colors were kept light whereas darker, more saturated colors were used for links to make them stand out. As datasets become complex, choosing appropriate color palettes to highlight features is a significant challenge.

Our participants easily found the highly cited paper in the Dependency Parsing region and recognized it from their domain knowledge. They also found a few other highly cited papers, and recognized them as well. Then, they looked at the Summarization region to find one of the participant's (Radev) papers and found it among the highly cited ones, as they expected. They inspected the remaining regions in a similarly systematic way. When looking into the PBMT region, they found a highly cited paper but were surprised by its popularity. One team member (Lin) then recalled that it had introduced a methodology that many other papers had used and therefore was frequently cited.
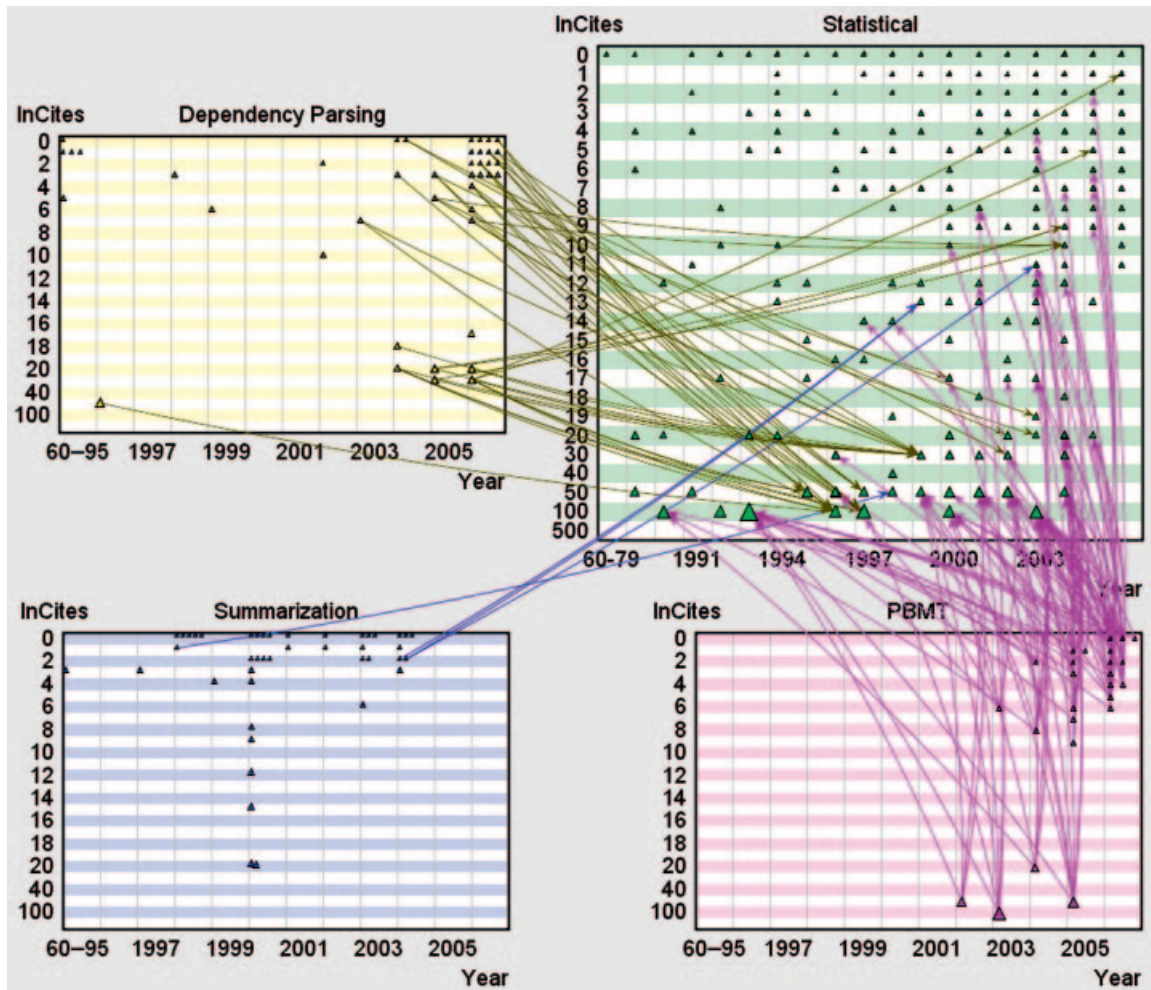
FIG. 5. Links from other regions to the Statistical region are enabled. This shows the citation patterns from other research fronts to the Statistical region. The Summarization papers are the least likely to cite the Statistical papers while DP and then PBMT have frequent citations to the Statistical papers.

With the Statistical Machine Translation region, filters became necessary because this region contains more nodes and more links than do the others, exceeding the threshold of comprehensibility when all links are viewed. After a single demonstration of the filters, the participants started to use them extensively. One of the researchers (Klavans) expressed enthusiasm to continue exploring this Statistical region because it was in her domain of expertise. She wanted to see where the patterns matched her understanding and whether there were surprises. She also provided an explanation for the contrast of citation patterns with respect to time between the Dependency Parsing and the Statistical topics. She commented that in the Statistical topic, the papers cite the recent papers, and that was good and expected, while this behavior would be considered undesirable in the Dependency Parsing topic as they would cite the early papers more often. Klavans attributed this to differences in the disciplinary culture in the ways previous research is cited between the two groups of researchers. She also commented that this visualization strategy could be widely used to compare citation cultures.

Next, the participants looked at the distribution of papers from other topics that cite papers in the Statistical region (Question 3) (Figure 5).

In this view (Figure 5), their expectations were confirmed that Dependency Parsing and PBMT cited many Statistical papers while the Summarization papers cited only four Statistical papers. In addition, the large number of links from the PBMT to the Statistical region confirmed that PBMT is more tied to Statistics than it is to Dependency Parsing. Another observation was that early papers in Dependency Parsing showed only one link to the Statistical topic.

Next, participants were asked what the relationships were among the small regions (i.e., Dependency Parsing, Summarization, and PBMT; Question 4). At first, they did not understand and did not have an answer. Perhaps this was not a point that they had thought of, but that when brought to their attention, they began to ponder it. When they realized that there was no relation due to the lack of links between regions in the NVSS, they found this fact interesting; however, this did not seem useful for them.
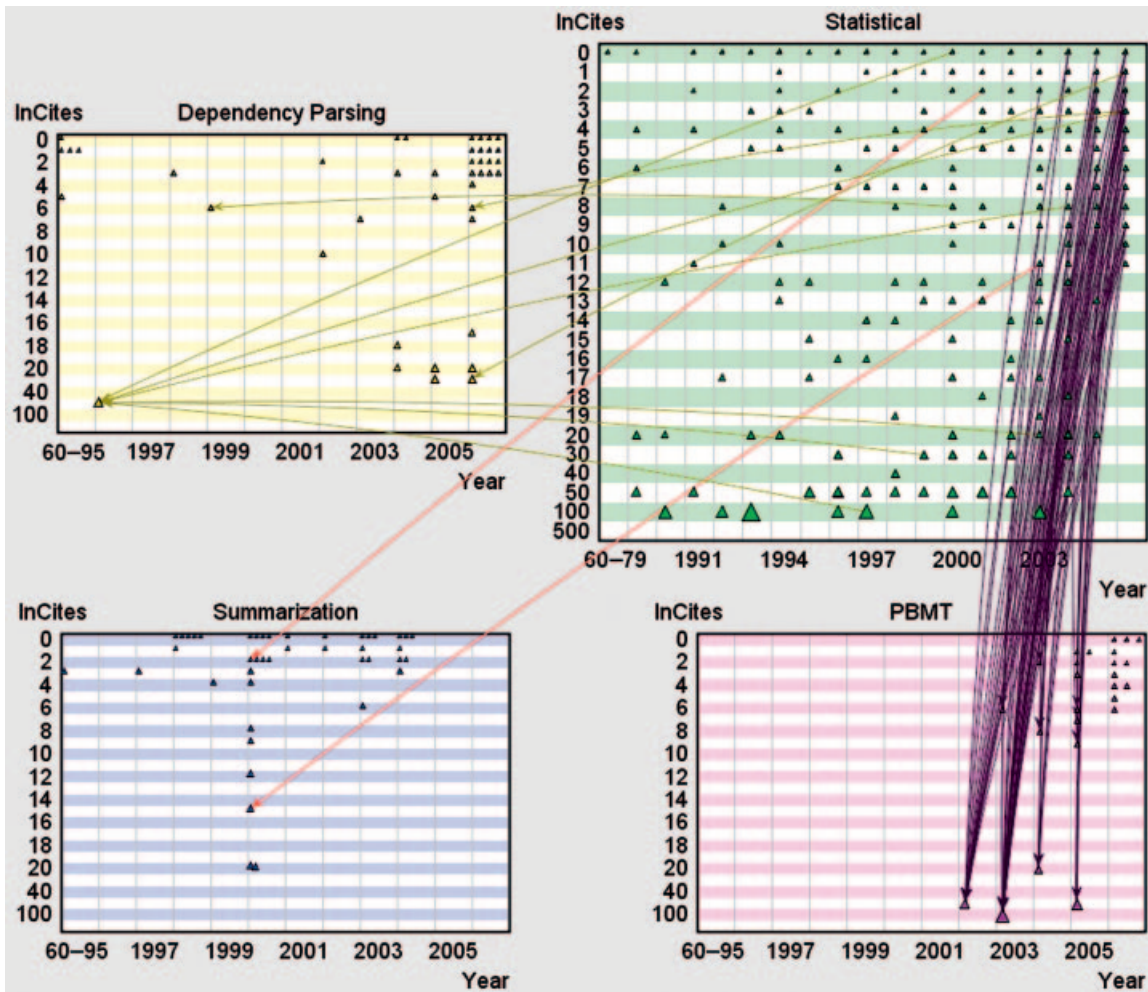
FIG. 6. Displaying links from the Statistical region to other topic areas. The summarization papers are cited least (as it was in the other direction, see Figure 5). One paper in Dependency Parsing is cited six times while PBMT papers receive the majority of citations.

Next, they looked at what papers were cited from the papers in the Statistical region (Question 5) (Figure 6). Klavans was surprised by the fact that the papers in the Statistical region cite only two papers in the Summarization topic. She said that the summarization topic generally is based less on statistical modeling, so fewer references would be expected; however, she expected more than two citations and was surprised by this small number.

Klavans and Dorr noted that the funding in Summarization as a field dropped in approximately 2004, so fewer citations would be the result of a lack of funding for that research area. In contrast, statistical approaches were heavily funded, and thus the number of articles (and citations) continued to rise. This clearly shows the impact of funding, both positive and negative, on scientific advancements.

Finally, the participants explored links within the statistical region by applying filters (Question 6) (Figure 7). Seeing the link patterns generated excitement. Klavans stated her appreciation for the NVSS, and the team expressed interest in exploring further datasets.

Even this modest case study (6 domain experts in three sessions over a 4-month period) gave us strong encouragement about the benefits of semantic substrate layouts while also generating suggested improvements to the user interface design: Node placement inside cells could be improved by considering the node sizes in each cell rather than considering the node sizes in the entire region (see Figure 7; each cell in the top row suffers from node occlusion; however, there is space to minimize the occlusion). When nodes are occluded, they resemble a single node (see Figure 7; each cell in the top row looks as if it contains a single node; however, in fact, each cell contains several occluded nodes). This could be ameliorated by jittering. The NVSS already includes a metanode feature that allows users to aggregate all the nodes in a cell to a larger metanode, thereby reducing clutter of nodes and links.

Limitations and Future Work

Our preliminary case study with 6 participants produced evidence that this visualization method for research fronts is
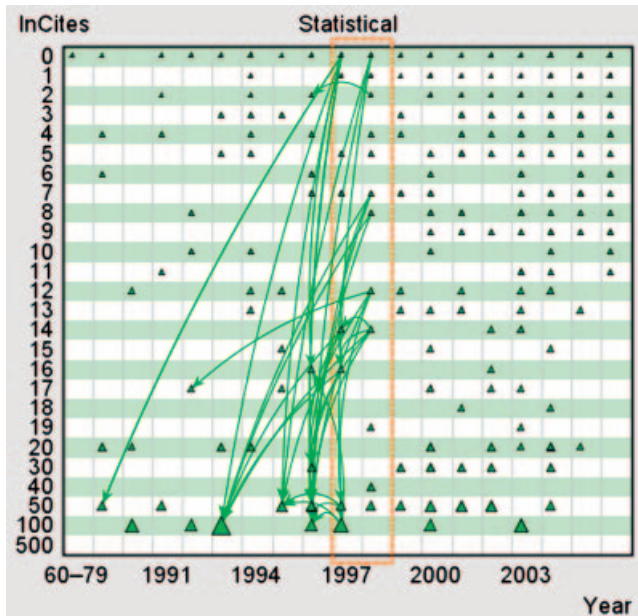
FIG. 7. Filtering links according to year in the Statistical region. This helps users to minimize the link clutter and focus on part of the links by an attribute, such as year (outgoing links from 1997–1998) in this figure. Most citations are to recent papers, so the one citation to an early paper is clearly visible.

interesting to researchers highly experienced in data analysis. These participants may be biased to support this effort, but in previous encounters, they had demonstrated healthy skepticism about novel strategies. Additional case studies with other participants and datasets would add evidence of benefits and lead to further design improvements. We hope that this form of case study will be adopted by other researchers, as it enables developers and designers to validate their designs with domain experts on problems that are meaningful to them.

A comparative case study with tools such as DIVA would be helpful to highlight the benefits and limitations of each tool for the set of tasks. A more refined set of user needs would be helpful to evaluate information visualization and other tools.

As a follow-up to this proof-of-concept study, controlled experimentation with specific tasks and a larger group of participants would help verify hypotheses and measure the strength of effect for individual features. Human performance times can be helpful to adjust specific design features such as screen management, filtering, history keeping, set saving, and aggregation to reduce clutter. Further exploratory data analysis with newcomers to these topics would validate the value of these interactive visualizations and lead to suggestions for other features.

In summary, we are greatly encouraged to see the value of multiple topic visualizations with user control over link visibility. Having the nodes in distinct regions, each temporally organized, provides clear understanding about the emergence of topics. The most dramatic effect is the capacity to quickly see the strength and direction of influence among research fronts.

Acknowledgments

References

Aris, A., & Shneiderman, B. (2007). Designing semantic substrates for visual network exploration. Information Visualization, 6(4), 281–300.

Börner, K., Chen, C., & Boyack, K. (2003). Visualizing knowledge domains. In B. Cronin (Ed.), Annual review of information science & technology (Vol. 37, pp. 179–255). Medford, NJ: Information Today, Inc./American Society for Information Science and Technology.

Bradshaw, S. (2002). Reference directed indexing: Indexing scientific literature in the context of its use. Unpublished doctoral dissertation, Northwestern University.

Bradshaw, S. (2003). Reference directed indexing: Redeeming relevance for subject search in citation indexes. In Proceedings of the 7th European Conference on Research and Advanced Technology for Digital Libraries (pp. 499–510).

Buckingham Shum, S., Uren, V., Li, G., Domingue, J., & Motta, E. (2003). Visualizing internetworked argumentation. In P.A. Kirschner, S. Buckingham Shum, & C.S. Carr (Eds.), Visualizing argumentation: Software tools for collaborative and educational sense-making (pp. 185–204). London: Springer-Verlag.

Card, S.K., Mackinlay, J.D., & Shneiderman, B. (1999). Readings in information visualization: Using vision to think. San Francisco: Kaufmann.

Carr, C.S. (2003). Using computer supported argument visualization to teach legal argumentation. In P.A. Kirschner, S. Buckingham Shum, & C.S. Carr (Eds.), Visualizing argumentation: Software tools for collaborative and educational sense-making (pp. 75–96). London: Springer-Verlag.

Chen, C. (2003). Mapping scientific frontiers: The quest for knowledge visualization. Berlin: Springer.

Chen, C. (2004a). Information visualization: Beyond the horizon (2nd ed.). London: Springer.

Chen, C. (2004b). Searching for intellectual turning points: Progressive knowledge domain visualization. Proceedings of the National Academy of Sciences, USA, 101(Suppl.), 5303–5310.

Chen, C. (2006). CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. Journal of the American Society for Information Science and Technology, 57(3), 359–377.

Elkiss, A., Shen, S., Fader, A., Güneş, E., States, D., & Radev, D.R. (2008). Blind men and elephants: What do citation summaries tell us about a research article? Journal of the American Society for Information Science and Technology, 59(1), 51–62.

Garfield, E. (2004). Historiographic mapping of knowledge domains literature. Journal of Information Science, 30(2), 119–145.

Garfield, E., Pudovkin, A.I., & Istomin, V.I. (2003a). Mapping the output of topical searches in the Web of Knowledge and the case of Watson-Crick. Information Technology and Libraries, 22(4), 183–187.

Garfield, E., Pudovkin, A.I., & Istomin, V.I. (2003b). Why do we need algorithmic historiography? Journal of the American Society for Information Science and Technology, 54(5), 400–412. Retrieved June 22, 2009, from http://garfield.library.upenn.edu/papers/jasist54(5)400y2003.pdf

Ginsparg, P., Houle, P., Joachims, T., & Sul, J.-H. (2004). Mapping subsets of scholarly information. Proceedings of the National Academy of Sciences, USA, 101, 5236–5340.

Hall, D., Jurafsky, D., & Manning, C.D. (2008). Studying the history of ideas using topic models. Proceedings of the Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics (pp. 363–371).

Henry, N., Goodell, H., Elmqvist, N., & Fekete, J.-D. (2007). 20 years of four HCI Conferences: A visual exploration. International Journal of Human–Computer Interaction, 23(3), 239–285. Retrieved June 22, 2009, from http://dx.doi.org/10.1080/10447310701702402

Joseph, M.T., & Radev, D.R. (2007). Citation analysis, centrality, and the ACL Anthology. Tech. Report No. CSE-TR-535-07, University of Michigan, Department of Electrical Engineering and Computer Science.

Leicht, E.A., Clarkson, G., Shedden, K., & Newman, M.E.J. (2007). Large-scale structure of time evolving citation networks. European Physics Journal B, 59, 75–83.

Morris, S., & Martens, B. (2007). Mapping research specialties. Annual Review of Information Science, 42.

Morris, S.A., Yen, G., Zheng, W., & Asnake, B. (2003). Time line visualization of research fronts. Journal of the American Society for Information Science and Technology, 54(5), 413–422.

Nanba, H., Abekawa, T., Okumura, M., & Saito, S. (2004a). Bilingual presri: Integration of multiple research paper databases. In Proceedings of the RIAO (pp. 195–211), Avignon, France.

Nanba, H., Kando, N., & Okumura, M. (2004b). Classification of research papers using citation links and citation types: Towards automatic review article generation. In Proceedings of the 11th SIG Classification Research Workshop (pp. 117–134), Chicago.

Nanba, H., & Okumura, M. (1999). Towards multidocument summarization using reference information. Proceedings of the International Joint Conference on Artificial Intelligence (pp. 926–931), Stockholm.

Newman, M.E.J. (2001). The structure of scientific collaboration networks. Proceedings of the National Academy of Sciences, USA, 98(2), 404–409.

Perer, A., & Shneiderman, B. (2008). Integrating statistics and visualization: Case studies of gaining clarity during exploratory data analysis. In Proceedings of the ACM CHI2008 Conference (pp. 265–274). New York: ACM Press.

Plaisant, C. (2004). The challenge of information visualization evaluation. In Proceedings of the Conference on Advanced Visual Interfaces (pp. 109–116). New York: ACM Press.

Qazvinian, V., & Radev, D.R. (2008). Scientific paper summarization using citation summary networks. In Proceedings of the International Conference on Computational Linguistics (pp. 689–696). Stroudsburg, PA: Association for Computational Linguistics.

Saraiya, P., North, C., & Duca, K. (2005). An insight-based methodology for evaluating bioinformatics visualization. IEEE Transactions on Visualization and Computer Graphics, 11(4), 443–456.

Shiffrin, R.M., & Börner, K. (2004). Mapping knowledge domains. Proceedings of the National Academy of Sciences, USA, 101, 5183–5185. DOI:10.1073/pnas.0307852100

Shneiderman, B. (1992). Tree visualization with tree-maps: 2-d space-filling approach. ACM Transactions on Graphics, 11(1), 92–99.

Shneiderman, B., & Aris, A. (2006). Network visualization by semantic substrates. Proceedings of IEEE Visualization/Information Visualization, IEEE Transactions on Visualization and Computer Graphics, 12(5), 733–740.

Shneiderman, B., & Plaisant, C. (2006). Strategies for evaluating information visualization tools: Multi-dimensional in-depth long-term case studies. In Proceedings of the BELIV'06 Workshop, Advanced Visual Interfaces Conference (pp. 1–7), Venice, Italy. New York: ACM Press.

Small, H. (2005). Paradigms, citations, and maps of science: A personal history. Journal of the American Society for Information Science and Technology, 54(5), 394–399.

Sugiyama, K. (1987). A cognitive approach for graph drawing. Cybernetics and Systems, 18(6), 447–488.

Sugiyama, K., Tagawa, S., & Toda, M. (1981). Methods for visual understanding of hierarchical system structures. IEEE Transactions on Systems, Man and Cybernetics SMC 11(2), 109–125.

Teufel, S., & Moens, M. (2002). Summarizing scientific articles: Experiments with relevance and rhetorical status. Computational Linguistics, 28(4), 409–445.

Teufel, S., Siddharthan, A., & Tidhar, D. (2006). Automatic classification of citation function. In Proceedings of the Empirical Methods in Natural Language Processing, (pp. 103–110). Burwood, Australia: BPA Digital.

Thomas, J.J., & Cook, K.A. (Eds.). (2005). Illuminating the path: Research and development agenda for visual analytics. Los Alamitos, CA: IEEE Press.

Thomson Corporation. (2008). Essential science indicators: Emerging research fronts. Retrieved June 22, 2009, from http://esi-topics.com/erf/

Uren, V., Buckingham Shum, S., Bachler, M., & Li, G. (2006). Sensemaking tools for understanding research literatures: Design, implementation and user evaluation. International Journal of Human–Computer Studies, 64(5), 420–445.