

CiteSpace II: Detecting and Visualizing Emerging Trends and Transient Patterns in Scientific Literature

Chaomei Chen

College of Information Science and Technology, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104-2875. E-mail: chaomei.chen@cis.drexel.edu

This article describes the latest development of a generic approach to detecting and visualizing emerging trends and transient patterns in scientific literature. The work makes substantial theoretical and methodological contributions to progressive knowledge domain visualization. A specialty is conceptualized and visualized as a time-variant duality between two fundamental concepts in information science: research fronts and intellectual bases. A research front is defined as an emergent and transient grouping of concepts and underlying research issues. The intellectual base of a research front is its citation and co-citation footprint in scientific literature—an evolving network of scientific publications cited by research-front concepts. Kleinberg's (2002) burst-detection algorithm is adapted to identify emergent research-front concepts. Freeman's (1979) betweenness centrality metric is used to highlight potential pivotal points of paradigm shift over time. Two complementary visualization views are designed and implemented: cluster views and time-zone views. The contributions of the approach are that (a) the nature of an intellectual base is algorithmically and temporally identified by emergent research-front terms, (b) the value of a co-citation cluster is explicitly interpreted in terms of research-front concepts, and (c) visually prominent and algorithmically detected pivotal points substantially reduce the complexity of a visualized network. The modeling and visualization process is implemented in CiteSpace II, a Java application, and applied to the analysis of two research fields: mass extinction (1981–2004) and terrorism (1990–2003). Prominent trends and pivotal points in visualized networks were verified in collaboration with domain experts, who are the authors of pivotal-point articles. Practical implications of the work are discussed. A number of challenges and opportunities for future studies are identified.

Introduction

Scientific literature is characterized by two distinct citation half-lives of articles: classic articles with persistently

high citations and transient ones with their citations peaked within a short period of time (Price, 1965). Transient ones are much more common than classics (van Raan, 2000). The average length of time that a research article continues to be cited in the scientific literature is closely connected to the growth speed of the underlying research area (Abt, 1998). Understanding the dynamics of how transient articles transform the intellectual landscape of a scientific field has significant practical implications for scientists in a wide variety of disciplines.

Emergent trends and abrupt changes in the scientific literature can be associated with internal as well as external causes. Typical internal causes include new discoveries and scientific breakthroughs such as the discovery of an impact crater in mass-extinction research or the discovery of a supermassive black hole in astronomy. External ones may provoke scientists to study a subject matter from new perspectives. For example, the September 11, 2001, terrorist attacks have raised a variety of new issues to be addressed by researchers in national security, health care, posttraumatic stress disorder (PTSD) research, and many other areas. Detecting and understanding emerging trends and abrupt changes caused by such events in scientific disciplines can significantly improve the ability of scientists to deal with the changes in a timely manner. It is worth noting that large-scale changes in complex systems characterized by self-organized criticality may take place without apparent triggering events (Bak & Chen, 1991). There is limited evidence to suggest that the growth of scientific literature may be connected to self-organized criticality (van Raan, 2000). In this article, we concentrate on changes associated with significant events.

The concept of a *research front* was originally introduced by Price (1965) to characterize the transient nature of a research field. Price observed what he called the immediacy factor: There seems to be a tendency for scientists to cite the most recently published articles. In a given field, a research front refers to the body of articles that scientists actively cite. According to Price, a research front may consist of 40 to 50 recent articles. A research front has been studied in at least

Received July 26, 2004; revised August 11, 2004; accepted February 7, 2005

© 2005 Wiley Periodicals, Inc. • Published online 14 December 2005 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/asi.20317

three forms: (a) a cluster of co-cited articles (Chen & Morris, 2003; Small & Griffith, 1974), (b) a cluster of co-cited articles and all articles that cite the cluster (Garfield, 1994; Garfield & Torpie, 1964), and (c) a cluster of articles that cite a common group of articles (Morris, Yen, Wu, & Asnake, 2003).

The concept of an intellectual base is useful to further clarify the nature of a research front (Persson, 1994). If we define a research front as the state of the art of a specialty (i.e., a line of research), what is cited by the research front forms its intellectual base. A specialty can be conceptualized as a time-variant mapping $\Phi(t)$ from its research front $\Psi(t)$ to its intellectual base $\Omega(t)$.

$$\Phi(t) : \Psi(t) \rightarrow \Omega(t)$$

The goal of our research is to develop a generic approach that can be used to detect and visualize emerging trends and abrupt changes in $\Phi(t)$ over time. In particular, $\Psi(t)$ is a group of words and phrases (i.e., *terms*) associated with emerging trends and sudden changes at time t . These terms are called *research-front terms*. $\Omega(t)$ consists of groups of articles cited by articles in which research-front terms were found. The following notions summarize these relations, for example, S_{title} denotes a set of title terms; $IsHotTopic$ (term, t) denotes a Boolean function, and $article_0 \rightarrow article$ denotes that $article_0$ cites $article$. In particular, we illustrate the use of the new approach with two examples; one shows the impact of internal events, and the other shows the impact of external events.

$$\begin{aligned} \Phi(t) : \Psi(t) &\rightarrow \Omega(t) \\ \Psi(t) &= \{term \mid term \in S_{title} \cup S_{abstract} \cup S_{descriptor} \\ &\quad \cup S_{identifier} \wedge IsHotTopic(term, t)\} \\ \Omega(t) &= \{article \mid term \in \Psi(t) \wedge term \in article_0 \\ &\quad \wedge article_0 \rightarrow article\} \end{aligned}$$

The rest of the article is organized as follows. First, we summarize related work in detecting and visualizing emerging trends and abrupt changes in transient networks as well as quantitative studies of research fronts. Then we describe a new approach to detecting trends and changes in scientific literature, especially those due to triggering events. The goal is to enable domain analysts and scientists to identify and understand structural and temporal patterns clearly. We illustrate the new approach with two real-world examples: mass-extinction research (1981–2003) and terrorism research (1990–2003). The resultant research front maps are verified with domain experts. The results also are evaluated by examining the content of some of the relevant articles.

Detecting and Tracking the Evolution of a Specialty

Researchers have studied quantitative methods that can be used to identify and track the essence of a research front as it evolves over time. The transient nature of a research front poses challenges for scientists, research policy makers,

and many others to keep up with the rapid advances of the state of the art in science. Understanding the dynamics of a research front is essential for scientists, analysts, and decision makers to be able to identify emerging trends and sudden changes in the body of scientific knowledge.

Research Fronts and Intellectual Bases

A research front represents the state-of-the-art thinking of a research field. Price (1965) noticed the interesting tendency that the most frequently cited articles tend to be the most recent ones in the citation network of scientific articles. The *immediacy* effect gives a good explanation of the well-known phenomenon of articles being considered obsolete a few years after their publication. Price described an earlier conjecture made by Burton and Kebler (1960) that the periodical literature may be composed of two distinct types of literature with very different half-lives—the classic and the transient parts. The transient part essentially corresponds to what is called a *research front*. According to Price, “the research front builds on recent work, and the network becomes very tight” (p. 515). He estimated about 30 to 40 articles published before a citing article would constitute the research front relative to the citing article. Price arranged 200 articles on the N-rays subject chronologically and used a matrix of citations (column articles cite row articles) to depict the research front of the subject, showing that the boundary of the research front was about 50 articles published prior to the citing article. Clearly, it would be useful if the movement of such relatively small networks of articles can be tracked in the overwhelmingly large body of the literature.

Typical questions regarding a research front may include: How did it get started? What is the state of the art? What are the critical paths in its evolution? To address such questions, we need to detect and analyze emerging trends and abrupt changes associated with a research front over time. We also need to identify the focus of a research front at a particular time in the context of its intellectual base, to reveal significant intellectual turning points as a research front evolves, and to discover the interconnections between different research fronts.

Researchers studied the evolution of a field of study from a variety of perspectives. Small and Griffith (1974) represented currently activated scientific specialties as clusters of co-cited articles. The nature of a co-citation cluster was determined by word profiles. Suppose an article D_i in a cluster has been cited by K citing articles $D_{i1}, D_{i2}, \dots, D_{ik}$. The four most frequently found words in the titles of D_{ij} 's are selected to form the word profile of the article D_i . By aggregating all the word profiles within a cluster and selecting the top- N most frequent words, one can easily extend the method to generate an N -word cluster profile. The distinct advantage of using such word profiles is its simplicity; however, the first-order word profiles were limited to the titles of the K citing articles, thus they may not reveal the dynamics of underlying themes.

Braam, Moed, and Raan (1991) defined a specialty as “focused attention by a number of scientific researchers to a set of related research problems and concepts” (p. 252). They studied the continuity and stability of a specialty in terms of the similarity between co-citation clusters across consecutive years. The similarity between two co-citation clusters is determined by comparing aggregated word profiles of the clusters. They identified a series of similar co-citation clusters over years as different phases of the same specialty, although they did not present graph-theoretical or other visual representations to illustrate the transformation from one phase to another.

In a study of research fronts and intellectual bases based on articles published in the *Journal of the American Society for Information Science* between 1986 and 1990, Persson (1994) made a distinction between a research front and an intellectual base: “In bibliometric terms, the citing articles form a research front, and the cited articles constitute an intellectual base” (p. 31). The intellectual base was represented by co-citation clusters. By gradually lowering the co-citation threshold, he demonstrated what he called a step-wise enlargement of the intellectual base—the intellectual base grew larger as more articles with lower co-citations were added to the co-citation clusters. Persson found that an intellectual base was remarkably stable over a long period of time. He noted that the correspondence between research-front clusters and intellectual-base clusters is not self-evident since there is no restriction in terms of what cited articles were active when the research front was formed. In this article, we conceptualize a specialty as a time-variant mapping between its research front and its intellectual base.

Morris et al. (2003) described a timeline visualization of a research front, defined as groups of articles that consistently cited a fixed, time-invariant group of base articles. The research-front articles were clustered based on bibliographic coupling (Kessler, 1963).

The Special Topics series (<http://www.esi-topics.com>) is a Web-based service directed by Henry Small at the Institute for Scientific Information. The series reports scientific research areas that have experienced distinct advances in terms of citation profiles. The Special Topics series has drawn upon decades of work by Small (1977, 1999b, 2003; Small & Griffith, 1974). Clusters of co-cited articles are extracted from each year’s citation data. Clusters from adjacent years are compared to identify new members as a sign of an emerging trend. The extent to which co-citation clusters in adjacent years overlap with one another provides a means for detecting the shift of research focus. Table 1 summarizes related work regarding research fronts and intellectual bases.

In this article, we chose not to restrict the citations to a fixed, time-invariant intellectual base; instead, we expect that a specialty’s intellectual base will change over time along with the movement of its research front. In part, this is because we define a research front differently to emphasize emerging trends and abrupt changes as the defining features of a research front. A research front is the domain of a time-variant mapping, and its intellectual base is the co-domain of

the mapping. Our new definition means that now the intellectual base of a research front is simply the citation trails of the research front in the literature.

Our goal is to determine to what extent the movement of a research front is coupled with the stability of its intellectual base. It also is of particular interest to identify what impact external events may have on the dynamics of the interrelationship between a research front and an intellectual base. For example, how does the discovery of compelling evidence shift the course of a research front?

Detecting Emerging Trends

There is a growing interest in detecting emerging trends and tracking topics. Some researchers visualize various temporal patterns to aid trend detection (Erten, Harding, Kobourov, Wampler, & Yee, 2003), although it is still common to identify a trend based on statistical profiles (Popescul, Flake, Lawrence, Ungar, & Giles, 2000; Roy, Gevry, & Pottenger, 2002).

Research in Topic Detection and Tracking has primarily focused on five types of tasks: (a) story segmentation, (b) topic detection, (c) topic tracking, (d) first-story detection, and (e) story-link detection. Roy et al. (2002) surveyed existing methods for trend detection. Kontostathis, Galitsky, Pottenger, Roy, and Phelps (2003) gave a comprehensive survey of trend detection in text data mining.

Allan, Papka, and Lavrenko (1998) described a single pass algorithm for identifying new stories. Given a newly received story and a collection of existing stories, if the collection has no similar stories then the story is identified as a new story. Their algorithm worked well to distinguish stories on the Oklahoma City bombing in 1995 from the earlier ones on the World Trade Center attacks, but it did not single out the O.J. Simpson trial stories from other court cases.

Swan and Allan (1999) constructed a χ^2 distribution-based statistical model to determine the significance of received information. They tested their approach with CNN broadcast news and Reuters newswires, which is a subset of the Topic Detection and Tracking pilot study’s corpus.

Text analysis may help to identify the most popular words used in articles over a particular period of time. The movement of an underlying thematic flow can be captured by the changes of such words. ThemeRiver (Havre, Hetzler, Whitney, & Nowell, 2002) is a visualization system that uses the metaphor of a river to depict thematic flows over time in a collection of articles. The thematic changes are shown along a time line of corresponding external events. A thematic river consists of frequency streams of terms; the changing width of a stream over time indicates the changes of term occurrences. The occurrence of an external event may be followed by sudden changes of thematic strengths.

Erten et al. (2003) described a temporal graph visualization of the categorization of articles in ACM conference proceedings using their system TGRIP to find the “hottest” topics in computing and identify steadily declining areas as well as rapidly growing ones. Their category graph consisted

TABLE 1. Definitions of research fronts and intellectual bases.

Author	Year	Research front	Intellectual base	Cluster	Labeling
Price	1965	Transient clusters of actively cited recent (30–50) articles from the perspective of a given citing article	Not defined	Recentness of citations	N/A
Small & Griffith	1974	Co-citation clusters	Not defined	Co-citation	Cited articles labeled by word profiles derived from citing articles
Braam et al.	1991	Focused attention to a set of related problems and concepts		Co-citation	Word profiles
Garfield	1994	Co-citation clusters plus citing articles		Co-citation	
Persson	1994	Articles citing the same literature	Citation image of a research front	Co-citation	Articles labeled by title words
Morris et al.	2003	Groups of articles consistently citing a fixed, time-invariant group of base articles	Fixed, time-invariant groups of articles	Bibliographic coupling	Cluster labeled by manually examined title words
CiteSpace II (this article)		Emerging thematic trends and surges of new topics	Co-citation network	Hybrid networks of co-cited articles and terms citing these articles.	Terms from titles, abstracts, and descriptors of abrupt frequency increase

of multiple time slices. Changes between adjacent time slices were shown as edges weighted in proportion to the corresponding percent change. Trends in each year were identified by the top five title words that have the highest percentage of use. Words such as *design*, *system*, and *simulation* were persistently found on the list over time whereas *ADA*, *database*, and *parallel* were listed only for a few years before they disappeared. Their study did not use citation data.

A clustering technique was reported for identifying temporal trends in the CiteSeer document database (Popescul et al., 2000). First, influential articles were identified based on citations to seed clusters. Then, articles co-cited with the seeding article of a cluster were assigned to the cluster. The similarity between two clusters was determined by articles they have in common. Clusters were characterized by most frequently used title words in member articles.

Identifying and Visualizing Specialties

In traditional co-citation analysis, much of the focus is on individual clusters of co-cited articles, and few studies have particularly examined the interrelationships between specialty clusters except for the work of Griffith, Small, Stonehill, and Dey (1974) and Small (1999a) on disciplinary connections. Griffith et al. (1974) found that between-cluster co-citation links tend to be weaker than within-cluster co-citation links.

In a social network, the influence of an actor may depend on how he or she is connected to other members in the network. Granovetter (1973) emphasized the strength of weak ties. Burt (1992) introduced the concept of structural holes. Within-cluster links are often stronger than between-cluster links. The focus of traditional citation-network analysis is on

strong links (i.e., within cluster links). To understand how specialties and different thematic trends interact with each other, it is essential to study the nature of long-range, between-cluster links and understand why articles in different specialties were connected.

The centrality of a node is a graph-theoretical property that quantifies the importance of the node's position in a network. A commonly used centrality metric is the betweenness centrality (Freeman, 1979). It measures the percentage of the number of shortest paths in a network to which a given node belongs. Nodes with high-betweenness centrality tend to be found in paths connecting different clusters. This feature has been used in community-finding algorithms to identify and separate clusters (Girvan & Newman, 2002).

Centrality metrics provide a computational method for finding pivotal points between different specialties or tipping points in an evolving network. An advantage of such a graph-theoretical method is that it is applicable to a wide variety of subject domains because the method is independent of any domain knowledge. Working with a small number of pivotal points, instead of the entire network, has several desirable practical implications. First, if these pivotal points characterize the structural and dynamical nature of a domain, the user's cognitive load would be considerably reduced. Second, subsequent processing such as automatic text summarization and natural language processing algorithms can be effectively limited to the small number of more focused pivotal points. Third, a pivotal-point-based approach extends word-profile-based approaches because pivotal points are identified based on their global graph-theoretical properties whereas word profiles are based on first-order frequencies.

Citation networks and co-citation networks have been long studied in information science and other disciplines. Garfield's historiography, for instance, is a directed graph of articles depicting citation links. HistCite is the most recent example along this line of research (Garfield, Pudovkin, & Istomin, 2003). In contrast, co-citation networks are undirected graphs (Small, 1986; Small & Griffith, 1974). An earlier example of visualizing co-citation networks is the SCI-Map system developed at the Institute for Scientific Information. Rather than using the more commonly used multidimensional scaling, SCI-Map uses a two-dimensional layout algorithm based on a geometric triangulation process. The triangulation starts with an arbitrarily chosen article and maps it to the origin of the coordinate system. Next, it places the article that has the strongest co-citation link with the first article. The position of the third article is triangulated using distances from the first two positions. The method gives the priorities to the strongest links. As a result, the network grows from the center onward. According to Small (1999b), a drawback of the process is its dependence on the order in which objects are assigned. SCI-map scales the radius of an article to the cube root of its citation frequency to represent the number of citations as the volume of a sphere.

Link reduction, or network scaling, is a practical challenge in visualizing complex networks such as co-citation networks. Real-world co-citation networks often have a large number of links. Reducing the number of links shown at a time may improve the clarity of a visualized network.

Spatial configurations of multidimensional scaling do not show links explicitly. In contrast, minimal spanning trees and Pathfinder network scaling (Schvaneveldt, 1990) select a subset of the original links and form a simplified network representation. Pathfinder networks were initially adapted for author co-citation analysis and subsequently extended to general co-citation analysis (Chen, 1999; Chen & Paul, 2001). Pathfinder network scaling relies on a triangle inequality test to determine whether a particular link should be preserved or eliminated. The selection criterion is that a single-link path's weight should not exceed that of alternative paths of multiple links. A Pathfinder network is superior than a multidimensional scaling in terms of preserving the chronological growth patterns in co-citation networks (Chen & Morris, 2003). Progressive Pathfinder network scaling is the most recent development for visualizing the evolution of a knowledge domain over time (Chen, 2004).

Several studies explored techniques for visualizing changes of bibliographic networks, including a study of the evolution of the AIDS research (Small & Greenlee, 1980). VxInsight, developed at the Sandia National Laboratories, depicts a body of the scientific literature as a three-dimensional landscape (Boyack, Wylie, & Davidson, 2002). It allows users to spot trends over time by using a time-slider. Temporal patterns of the evolution of a three-dimensional co-citation landscape were animated (Chen & Kuljis, 2003).

CiteSpace is a Java application for analyzing and visualizing co-citation networks (Chen, 2004). Its primary goal is to facilitate the analysis of emerging trends in a knowledge

TABLE 2. CiteSpace I vs. CiteSpace II.

	CiteSpace I	CiteSpace II
Time slicing	√	√
Co-citation networks/intellectual bases	√	√
Time-slice-based thresholds	√	√
Pathfinder pruning	√	√
Progressive network patching	√	√
Time-coded network visualization	√	√
n-gram term extraction	–	√
Burst detection	–	√
Sharply increased topical terms/research fronts	–	√
Betweenness centrality	–	√
Citation tree-ring visualization	–	√
Heterogeneous networks	–	√
PubMed MeSH heading networks	–	√
Time-zone views	–	√
Multithread Pathfinder	–	√

domain. It allows the user to take a time series of snapshots of a domain and subsequently merge these snapshots. The initial version of CiteSpace was used to reveal turning points in superstring revolutions in physics; however, several issues remained unresolved when we implemented the first version of CiteSpace. To distinguish the substantial changes since our earlier report (Chen, 2004), we refer to the initial version as CiteSpace I and the new version as CiteSpace II. Table 2 lists the major improvements in CiteSpace II. The most distinctive new feature is the combination of computational metrics and visual attributes of pivotal points. The motivation is to substantially reduce the user's cognitive burden as he or she searches for pivotal points in a knowledge structure.

New features of CiteSpace II are related to three central concepts: burst detection, betweenness centrality, and heterogeneous networks. A specialty is conceptualized as a mapping function between a research front and its intellectual base. This mapping function provides the basis of our conceptual framework to accommodate the three concepts. These concepts have instrumental roles in addressing three practical issues: (a) identifying the nature of a research front, (b) labeling a specialty, and (c) detecting emerging trends and abrupt changes in a timely manner.

CiteSpace II

We next analyze the three practical issues and describe how they are addressed in CiteSpace II. The procedure of applying the new approach to the study of what impact triggering events may have on a research field is outlined. The conceptual model of CiteSpace II is shown in Figure 1.

Labeling Co-Citation Clusters

Labeling clusters is concerned with the clarity and interpretability of co-citation clusters. The standard approach relies on word profiles derived from articles citing a cluster of co-cited articles. A pragmatic assumption is that word

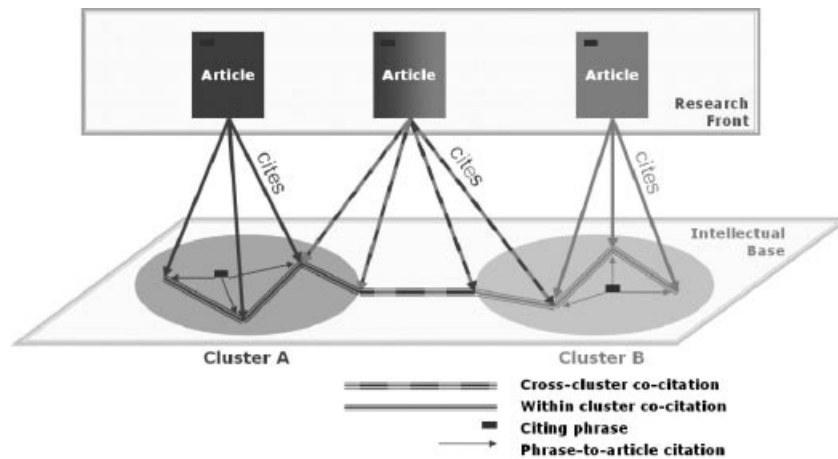


FIG. 1. The conceptual model of CiteSpace II. Time-sliced snapshots are devised to highlight changes of prominent specialties over time.

profiles characterize the nature of a co-citation cluster because the words are from articles that shape the cluster in the first place. Word-profile approaches have drawbacks. First, word profiles may not converge to a focused message. Analysts and users will make a substantial amount of sense-making efforts to synthesize a diverse range of word profiles. Second, cluster labels based on aggregating word profiles tend to be too broad to be useful. In practice, many users would be interested in not only the most commonly used terms but also terms that can lead to profound changes. Terms associated with an emerging trend could be overshadowed by a broader and more persistent theme. For instance, a surge of interest in *health care* with reference to the threats of biological and chemical weapons may be easily outnumbered by the more dominating terms such as *biological weapons*.

An ideal labeling method should distinguish emerging trends and rapid changes at the foreground from more persistent themes at the background. Small (1999a) dealt with a similar situation when he needed to explain the nature of cross-disciplinary co-citation links. He suggested that a useful explanation should distinguish how the articles at the two ends of such a link differ and what they have in common.

Kleinberg's (2002) burst-detection algorithm can be adapted for detecting sharp increases of interest in a specialty. Although Kleinberg's original algorithm was developed to detect the bursts of single words, the algorithm is generic enough to be applied to a time series of multiword terms or citations of articles. In CiteSpace II, a current research front is identified based on such burst terms extracted from titles, abstracts, descriptors, and identifiers of bibliographic records. These terms are subsequently used as labels of clusters in heterogeneous networks of terms and articles.

Improving the Timeliness of Visualized Co-Citation Networks

Using citing terms to characterize a research front may improve our understanding of its intellectual base. Co-citation analysis traditionally focuses on homogeneous networks of

co-cited articles. Newly published articles tend to be under-represented in such networks because they may not have enough citations. In addition, word profiles are not normally regarded as part of a co-citation network. Burst-detection algorithms can identify emergent terms regardless of how many times their host articles are cited. Therefore, a new research front can be featured in the "big picture" even before it attracts enough citations.

In addition to the standard force-directed view, CiteSpace II supports a time-zone view to highlight temporal patterns between a research front and its intellectual base. A time-zone view consists of an array of vertical strips as time zones. The time zones are arranged chronologically from left to right so that a research front points back to its intellectual base. The layout algorithm is a modified spring embedder algorithm such that the horizontal movement of an item is restricted to its own time zone, but its vertical movement is completely determined by its connections to items in other time zones. The goal is to make a specialty easily recognizable. The design of a time-zone view resembles the overall layout of time-line visualization (Morris et al., 2003). In terms of the major differences, time-zone views in CiteSpace II simultaneously show cited articles and citing terms to highlight the mapping between a research front and its intellectual base.

Computationally Identifying Pivotal Points

We assume that transitions between transient research fronts are traceable in terms of citations and co-citations. Assume Ψ_α are Ψ_β are prominent research fronts at time t and $t + \Delta t$, with corresponding landmark articles α and β in their intellectual bases $\Omega_\alpha = \Phi(\Psi_\alpha)$ and $\Omega_\beta = \Phi(\Psi_\beta)$. Two clusters of co-cited articles would emerge; α would be the center of one and β would be the center of the other. Articles on paths exclusively connecting the two clusters $\{p^{(i)}\}$ characterize the transition from Ψ_α to Ψ_β . We refer to such $\{p^{(i)}\}$ as pivotal points, turning points, or tipping points.

In CiteSpace I (Chen, 2004), users only have one way to identify pivotal points by visually scanning a visualized

network for nodes that bridge different clusters. An advantage of this approach is that no additional computing is required; however, a drawback is that users have no way to ensure that all the important pivotal points are found except by doing a brute-force visual search of all nodes in the network.

CiteSpace II makes it easier for users to identify pivotal points. In addition to inspecting salient visual attributes, the user easily can see nodes with high betweenness centrality (Freeman, 1979). Pivotal points are computationally identified and rendered so that they become preattentive, or pop out, in the visualized network. Pivotal points are highlighted in the display with a purple ring so that they stand out in a visualized network. Graph-theoretically identifiable pivotal points allow us to reduce networkwide operations to the subset of pivotal nodes only to improve the interpretability of the network.

Procedure

The procedure of using CiteSpace II is described in the following steps, followed by illustrative examples to demonstrate the new features.

1. Identify a knowledge domain using the broadest possible term. This is to ensure that subsequent analysis covers all major components of a knowledge domain (e.g., *mass extinction*).
2. Data collection: Currently, the primary source of data is the Web of Science. CiteSpace II also allows users to download bibliographic records directly from PubMed, a primary repository of medical literature. The topical terms identified in the first step are used to retrieve bibliographic records from the Web of Science, including titles, abstracts, and cited references. Each bibliographic record represents a citing article whereas a cited reference included in each record is called a cited article.
3. Extract research front terms: CiteSpace II first collects n-grams, or terms, from titles, abstracts, descriptors, and identifiers of citing articles in a dataset. The present study

used single words or phrases of up to four words. For example, *terrorism* and *posttraumatic stress disorder* are valid terms. Research-front terms are determined by the sharp growth rate of their frequencies.

4. Time slicing: In CiteSpace II, users specify the range of the entire time interval and the length of a single time slice.
5. Threshold selection: CiteSpace II allows users to specify three sets of threshold levels for citation counts, co-citation counts, and co-citation coefficients. The specified thresholds are applied to three time slices, namely the earliest slice, the middle one, and the last one. Linear interpolated thresholds are assigned to the rest of slices. In this study, bipartite networks of research-front terms and intellectual base articles were constructed. Such networks contain three types of links: (a) co-occurring research front terms, (b) co-cited intellectual base articles, and (c) a research-front term citing an intellectual base article.
6. Pruning and merging: Pathfinder network scaling is the default option in CiteSpace II for network pruning (Chen, 2004; Schvaneveldt, 1990). Users choose whether to apply the scaling operation to individual networks. Pathfinder network scaling is an asymptotically expensive algorithm. CiteSpace II implements a concurrent version of the algorithm to process multiple networks simultaneously, which substantially reduces the overall waiting time. Networks $G_i = (V_i, E_i)$ and $G_j = (V_j, E_j)$ from different time slices may have overlapping vertices and edges, i.e., $V_i \cap V_j \neq \emptyset$ and $E_i \cap E_j \neq \emptyset$. CiteSpace II merges individual networks by taking a set union of all the vertices and selecting links that do not violate a triangle inequality condition in overlapping areas between networks. Users can choose whether to prune the merged network as a whole.
7. Layout: CiteSpace II supports a standard graph view and a time-zone view.
8. Visual inspection: CiteSpace II enables users to interact with the visualization of a knowledge domain in several ways. The user may control the display of visual attributes and labels as well as a variety of parameters used by the underlying layout algorithms. Figure 2 illustrates the visual attributes of an article node.




FIG. 2. Citation tree rings represent the citation history of an article. The color of a citation ring denotes the time of corresponding citations. The thickness of a ring is proportional to the number of citations in a given time slice. The small number next to the center of a node is the citations throughout the entire time interval.

- Verify pivotal points. The significance of a marked pivotal point can be verified by asking domain experts (e.g., authors of pivotal-point articles) and/or examining the literature (e.g., passages containing citations of a pivotal-point article). A particularly interesting direction of research is the development of tools that can automatically summarize the value of a pivotal point. Digital libraries, automated text summarization, machine learning, and several other fields are among the most promising sources of input.

Case Studies

We demonstrate the new features of CiteSpace with case studies of two research fields: mass-extinction research (1981–2003) and terrorism research (1990–2003). Both fields have experienced profound changes as a result of triggering events.

Mass-extinction research (1981–2003). Five major mass extinctions have taken place on Earth in the past. The most widely known mass extinction is the K-T extinctions 65 million years ago, which led to the extinction of dinosaurs. Causes of mass extinctions have been debated by catastrophists and gradualists. Catastrophists stress the role of catastrophic events in triggering mass extinctions whereas gradualists believe that mass extinctions in general are a prolonged process and that no single event can determine the course of mass extinctions. The most famous catastrophic theory is known as the impact theory, proposed in the early 1980s to explain the K-T extinctions (Alvarez, Alvarez, Asaro, & Michel, 1980). According to the theory, a gigantic asteroid hit the earth and led to the extinctions of many species. The impact crater was identified in 1990. The goal of this case study was to see to what extent emerging trends and sudden changes can be detected with reference to specialties associated with the impact theory.

There are several key questions to be answered by the case study. What is the latest research front in mass-extinction research? How did the impact-theory research front evolve over the last 20 years? What are the “hottest” research-front terms? Which articles are associated with these terms?

The input data for CiteSpace II were retrieved from citation index databases via the Web of Science based on a topic search for articles published between 1981 and 2003 on *mass extinction*. The scope of the search included four topic fields in each bibliographic record: title, abstract, descriptors, and identifiers. The search was limited to articles in English only. The resultant dataset contains a total of 771 records. A total of 333 research-front terms were detected from the four topic fields of these records. The 23-year time span between 1981 and 2003 was divided into twelve 2-year time slices. Three sets of threshold levels, namely citation threshold (c), co-citation threshold (cc), and co-citation coefficient threshold (ccv), were set as follows: (2, 1, 10), (3, 1, 0), and (3, 2, 10).

Terrorism research (1990–2003). The 1995 Oklahoma City bombing and the September 11, 2001, terrorist attacks are among the most traumatic terrorism events. Unlike the mass-extinction case, each event may alter the course of research. How did the research community respond to such events and ramifications? What are the emerging research fronts in this field of research? How are they related to earlier research fronts?

The terrorism research (1990–2003) dataset consists of 1,776 records resulted from a topic search on *terrorism* in the Web of Science. A total of 1,108 research-front terms were found. The entire time interval of 1990 to 2003 was divided into seven 2-year slices. The seven corresponding networks were subsequently merged into a panoramic network to depict the changes of dominating research issues over the entire 14-year time span. CiteSpace II configurations are summarized in the next section.

Results

This section is organized in two parts: mass extinction (1981–2004) and terrorism (1990–2003). The following notations are used: D —a set of articles, T —a set of terms, $G(T^1, T^2, \dots, T^N)$ —a multipartite network of node types $\{T^{(k)}\}$, for example, $G(D)$ and $G(T, D)$, $\Pi\{G_i\}$ —a network merged from snapshot networks $\{G_i\}$, Π is the merge operation, and $PF(G)$ —a Pathfinder network of G .

Mass Extinction (1981–2004)

We describe four visualizations of the mass-extinction dataset produced by four configurations in CiteSpace II: (a) a merged network of co-cited articles $\Pi\{G_i(D)\}$, (b) a merged network of Pathfinder-pruned individual co-citation networks $\Pi\{PF(G_i(D))\}$, (c) a Pathfinder-pruned hybrid network $PF(G(T, D))$, and (d) the Pathfinder-pruned hybrid network $PF(G(T, D))$ shown in time-zone view.

Table 3 shows the first configuration for a 515-article co-citation network $\Pi\{G_i(D)\}$. Note that no articles were selected from the first two slices. It took less than 10 seconds to complete the process in CiteSpace II.

The most prominent article in the visualization is Alvarez-1980. It was this article that first introduced the impact theory (see Figure 3). The central area was surrounded by five densely connected clusters. In CiteSpace II, the user can select a cluster of items and find matching records in PubMed. Medical Subject Headings (MeSH) assigned to matched articles were ranked by their occurrences so that they serve as labels of the selected cluster. The number at the corner of each rectangle in Figure 3 is the number of articles in the selected region. For example, MeSH term *Geologic Sediments* was assigned to 8 of the 100 selected articles. Since PubMed primarily covers the medical literature, depending on the subject domain, articles in a visualized network may or may not have matching records in PubMed.

This $\Pi\{G_i(D)\}$ network is the result of merging several time-slice networks. Co-citation networks do not explicitly

TABLE 3. The construction configuration for a 515-article co-citation network.

2-year slices	c	cc	ccv	No. Articles	No. Nodes	No. Links
1981–1982	2	1	0.15	64	0	0
1983–1984	2	1	0.19	27	0	0
1985–1986	2	1	0.23	163	7	14
1987–1988	2	2	0.28	205	8	16
1989–1990	2	2	0.32	331	7	10
1991–1992	2	2	0.36	1,687	123	353
1993–1994	3	3	0.40	2,004	31	59
1995–1996	3	3	0.38	2,908	47	30
1997–1998	3	3	0.36	4,503	142	522
1999–2000	3	3	0.34	5,994	156	436
2001–2002	3	3	0.32	7,431	300	1,181
2003–2004	4	4	0.30	7,092	101	243
Total				32,409	922(515)	2,864

feature research fronts; rather, they represent the footprint of a research front. As a result, such networks do not readily lend themselves to answer questions such as: “Why did a particular cluster of co-cited articles emerge? Which research front is being involved?”

The second configuration generated a 541-article network $\Pi\{PF(G,(D))\}$, merging individual Pathfinder networks of co-cited articles. In addition, lower threshold levels were assigned to the earlier time slices. It took 24.6 seconds to complete the process. The results are summarized in Table 4.

The lower threshold levels increased the number of selected articles in the 1981–1982 and 1983–1984 slices, although these newly selected articles were cited only once in the dataset. These articles were not selected by the first configuration. These articles appeared to be widely co-cited. In 1981–1982, 1,555 co-citation links were found among 62 articles. Similarly, 2,964 co-citation links were found among 156 articles in 1985–1986. As shown in Figure 4, these articles are essentially one or two co-citation links away from Alvarez-1980 (Alvarez et al., 1980). The cluster located to the

right of the visualization was apparently formed recently because the prevalent red rings in this cluster indicate that they were cited in 2003–2004. Highly cited articles in this cluster include Bowring-1998, Renne-1995, Hallam-1997, and Wignall-1992. Raup-1982 is the second most highly cited article in the older cluster centered at Alvarez-1980.

The third configuration specifies a hybrid network of research-front terms and intellectual-base articles. The most recent cluster is positioned as the upper cluster in Figure 5. This cluster includes research-front terms such as *Late Frasnian*, *Atmospheric CO2*, and *carbon cycle*. The central area is associated with research front terms such as *Cretaceous-Tertiary Boundary* and *Late Cretaceous*, which are the key concepts of the impact theory.

In addition to the three visualizations, a time-zone view is shown in Figure 6, which emphasizes the temporal relationships. The time-zone view reveals three remarkable clusters stretched over time. Each cluster corresponds to a specialty in mass-extinction research. The lowest major cluster in the map, full of green lines if shown in color, is the K-T extinction

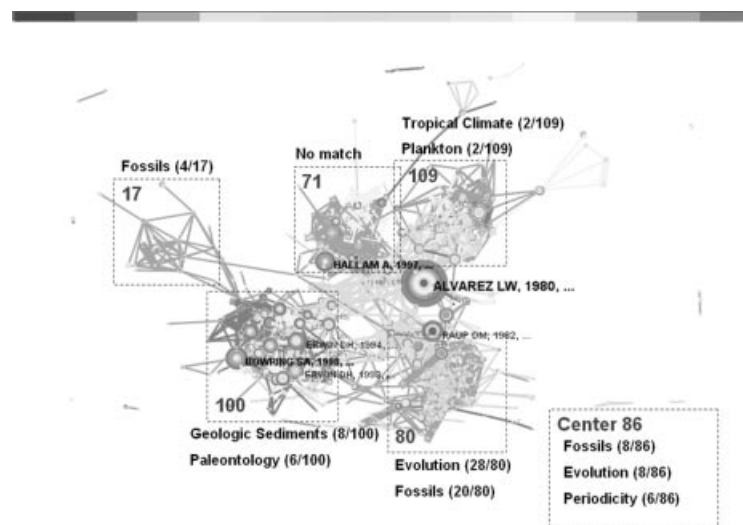


FIG. 3. A 515-node network of co-cited articles on mass extinction (1981–2004) based on twelve 2-year slices. Five clusters and the central area are marked by MeSH terms assigned to articles in individual clusters. MeSH terms were retrieved from PubMed on demand. The number at the corner of each rectangle is the number of articles found in the cluster. The youngest cluster is a 100-article cluster in the lower left of the visualization.

TABLE 4. The second configuration, including Pathfinder pruning and lower threshold levels.

2-year slices	c cc cvv	No. Articles	No. Nodes (unique)	No. Links
1981–1982	1 1 0.35	64	62	1,555
1983–1984	1 1 0.35	27	26	325
1985–1986	1 1 0.35	163	156	2,964
1987–1988	2 2 0.35	205	8	16
1989–1990	2 2 0.35	331	7	10
1991–1992	2 2 0.35	1,687	123	353
1993–1994	3 3 0.35	2,004	31	60
1995–1996	3 3 0.35	2,907	47	30
1997–1998	3 3 0.35	4,497	142	523
1999–2000	4 4 0.35	5,990	59	68
2001–2002	4 4 0.35	7,427	145	364
2003–2004	5 5 0.35	7,091	53	93
Total		32,393	859(541)	6,361




FIG. 4. The network of 541 co-cited articles based on 12 Pathfinder-pruned snapshot networks. The first three slices now present in the visualization. The most prominent cluster (i.e., the cluster to the right) also is the most recent one.

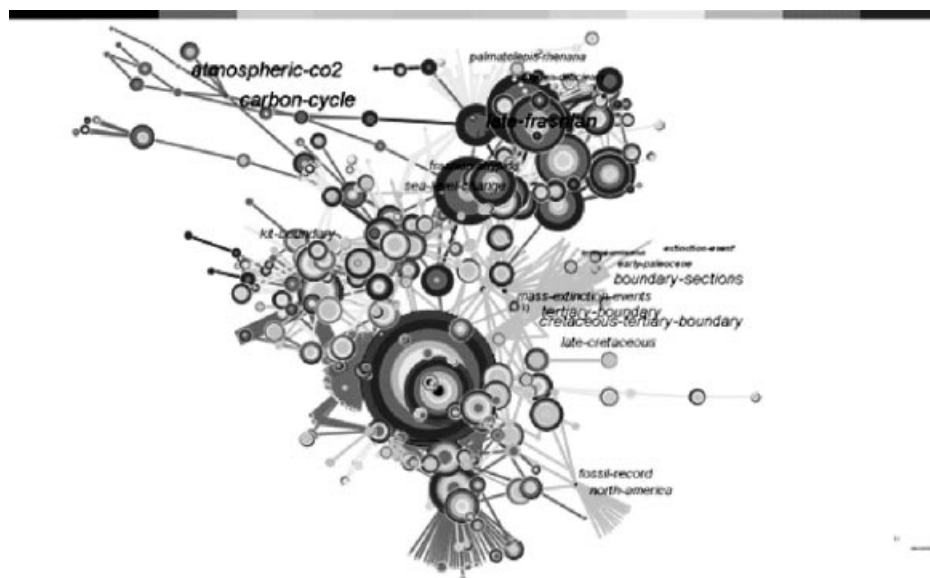


FIG. 5. A 569-node hybrid network of cited articles and citing terms (Time taken = 28 s).




FIG. 6. A time-zone view of mass-extinction research, showing three prominent lines of research: the K-T extinction thread (the second-lowest cluster), the Late Frasnian extinction thread (the middle cluster), and the Permian extinction thread (the upper cluster).

debate thread. Starting with Alvarez-1980, this thread maintained a strong presence for the next 10 years until it started to diminish in 1991–1992 and disappeared from the map after 1995–1996. Many articles in this cluster were still cited in 2003–2004, as indicated by the color of their outmost rings of citations. The most recent research-front terms include *North America, fossil record*, and *Late Cretaceous*.

The middle cluster, consisting of yellow links, started before the 1980s and ended in 1997. The thread contains articles Johnson-1985 and Copper-1986. The size of this thread is smaller than the K-T extinction thread. Because the thread has a research-front term *Late Frasnian*, we call it the Late-Frasnian thread. On a geological scale, the Frasnian Age is part of the Late Devonian Period between 386 and 375 million years ago. Late Devonian mass extinction is one of the five major mass extinctions at the boundary between the Frasnian and Famennian ages. Explanations for this mass extinction include an episode of global cooling, and associated lowering of sea level, and Meteorite impacts.

The cluster near the top of the map is predominated by orange links, representing an evolving thread. It started with Wignall-1992 (Wignall, 1992), reinforced by Bowring-1998, and recently added Becker-2001 (Becker, Poreda, Hunt, Bunch, & Rampino, 2001). *Permian-Triassic-boundary* and *South China* are two research-front terms detected in 2001. Becker-2001 is entitled “Impact Event at the Permian-Triassic Boundary.” A simple search on the Web for “Becker, Permian, and Triassic” confirmed that Becker-2001 presented evidence for a major impact associated with the Permian mass extinction. Becker et al. (2004) published a more recent article in *Science* entitled “Bedout: A Possible End-Permian Impact Crater Offshore Northwestern Australia.” Bedout is the name of a crater. Its significance to the Permian mass-extinction research is expected to be

comparable to that of the Chicxulub crater to the K-T impact theory. The discovery of the Chicxulub crater dramatically boosted the credibility of the K-T impact theory. Encouraged by the successful puzzle-solving experience, many scientists appear to have adapted the same approach to solve a different puzzle—by applying the impact theory to an earlier mass extinction. Finding the impact crater is the next logical step. Identifying a Permian-Triassic boundary impact crater has attracted the attention of many researchers. It was in this context that the current research front has emerged.

The time-zone visualization along with three questions (see Figure 7) were sent to Dr. Paul Wignall at Leeds University in the United Kingdom to verify the aforementioned observation because he authored two pivotal-point articles in the Permian extinction thread (Wignall & Hallam, 1992; Wignall & Hallam, 1993). He confirmed that the K-T extinction debate was slowing down since the mid 1990s because much of the debate was resolved as a result of the discovery of the Chicxulub impact crater: “Many people still work on this extinction event but many others have now moved to look at the end-Permian mass extinction, which is why there is continued interest in my work.” The research front term *Late Frasnian* was first found in 1997 in this thread. In answering the question of why the thread in the middle appears to come to an end, Wignall explained that the thread was related to Copper’s work on the Late Devonian mass extinction. It has not been an active line of research primarily due to the lack of supporting data.

In summary, the time-zone view conveys a picture of the decade-long trends consistent with the observation of a domain expert. The K-T extinction thread ended about the time the Permian extinction thread started to take off. Was it because scientists were encouraged by the outcome of the K-T extinction debate? The recent series of work by Becker et al.




FIG. 7. The time-zone visualization along with three questions sent to Dr. Wignall for comment. He is the author of two pivotal-point articles identified in the visualization.

(2004) indeed bears a considerable amount of resemblance to the K-T debate in the 1980s, notably the impact theory and the search for the impact crater. It would be very revealing to track how this thread will evolve in the near future.

Compared with earlier visualizations, the new methods in CiteSpace II have improved the clarity and interpretability of visualizations. In our 2002 case study of mass-extinction research (Chen, Cribbin, Macredie, & Morar, 2002), the goal was to reconstruct the history of a scientific field by visualizing the evolution of its co-citation network not only chronologically but also retrospectively. The most cited article can be identified by the tallest citation bar, and the mainstream specialty can be identified by the color groupings based on Principal Component Analysis. Walter Alvarez's (1997) book was used as the primary source to interpret the three-dimensional models—a two-dimensional co-citation network as the base map and a one-dimensional citation profile. Although three-dimensional models vividly reproduced

the growth of the citation landscape, identifying emerging trends in such models turned out to be a challenging task. The new visualizations in CiteSpace II enable us to identify not only groundbreaking articles such as Alvarez et al. (1980) but also critical points in the course of transition from one specialty to another. Identifying pivotal points purely based on the topological properties of a network has practical values. For example, one can quickly identify a small number of domain experts and seek their advice on specific issues. In the following example, several domain experts are identified in this way as part of an in-depth validation of the method.

Terrorism Research (1990–2003)

Merging seven snapshot networks led to a 335-node hybrid network of terms and articles. Threshold levels are summarized in Table 5. It took 8 seconds to build the merged

TABLE 5. The CiteSpace configuration used for visualizing the terrorism research dataset (1990–2003). Threshold levels are set for citations (c), co-citations (cc), and co-citation coefficients (ccv). The number of articles in each slice is the total number of articles that were cited within the slice. The number of nodes is the number of articles and terms that pass the thresholds. The number of links is the number of links within a slice.

2-year slices	c	cc	ccv	No. Articles	No. Nodes (unique)	No. Links
1990–1991	2	1	0.10	1,455	15	27
1992–1993	2	1	0.11	1,372	32	137
1994–1995	2	1	0.12	1,527	12	19
1996–1997	2	1	0.14	2,735	109	1,504
1998–1999	3	2	0.15	3,037	31	75
2000–2001	3	2	0.18	5,518	47	211
2002–2003	3	2	0.22	18,366	182	649
Total				34,010	428(335)	2,622

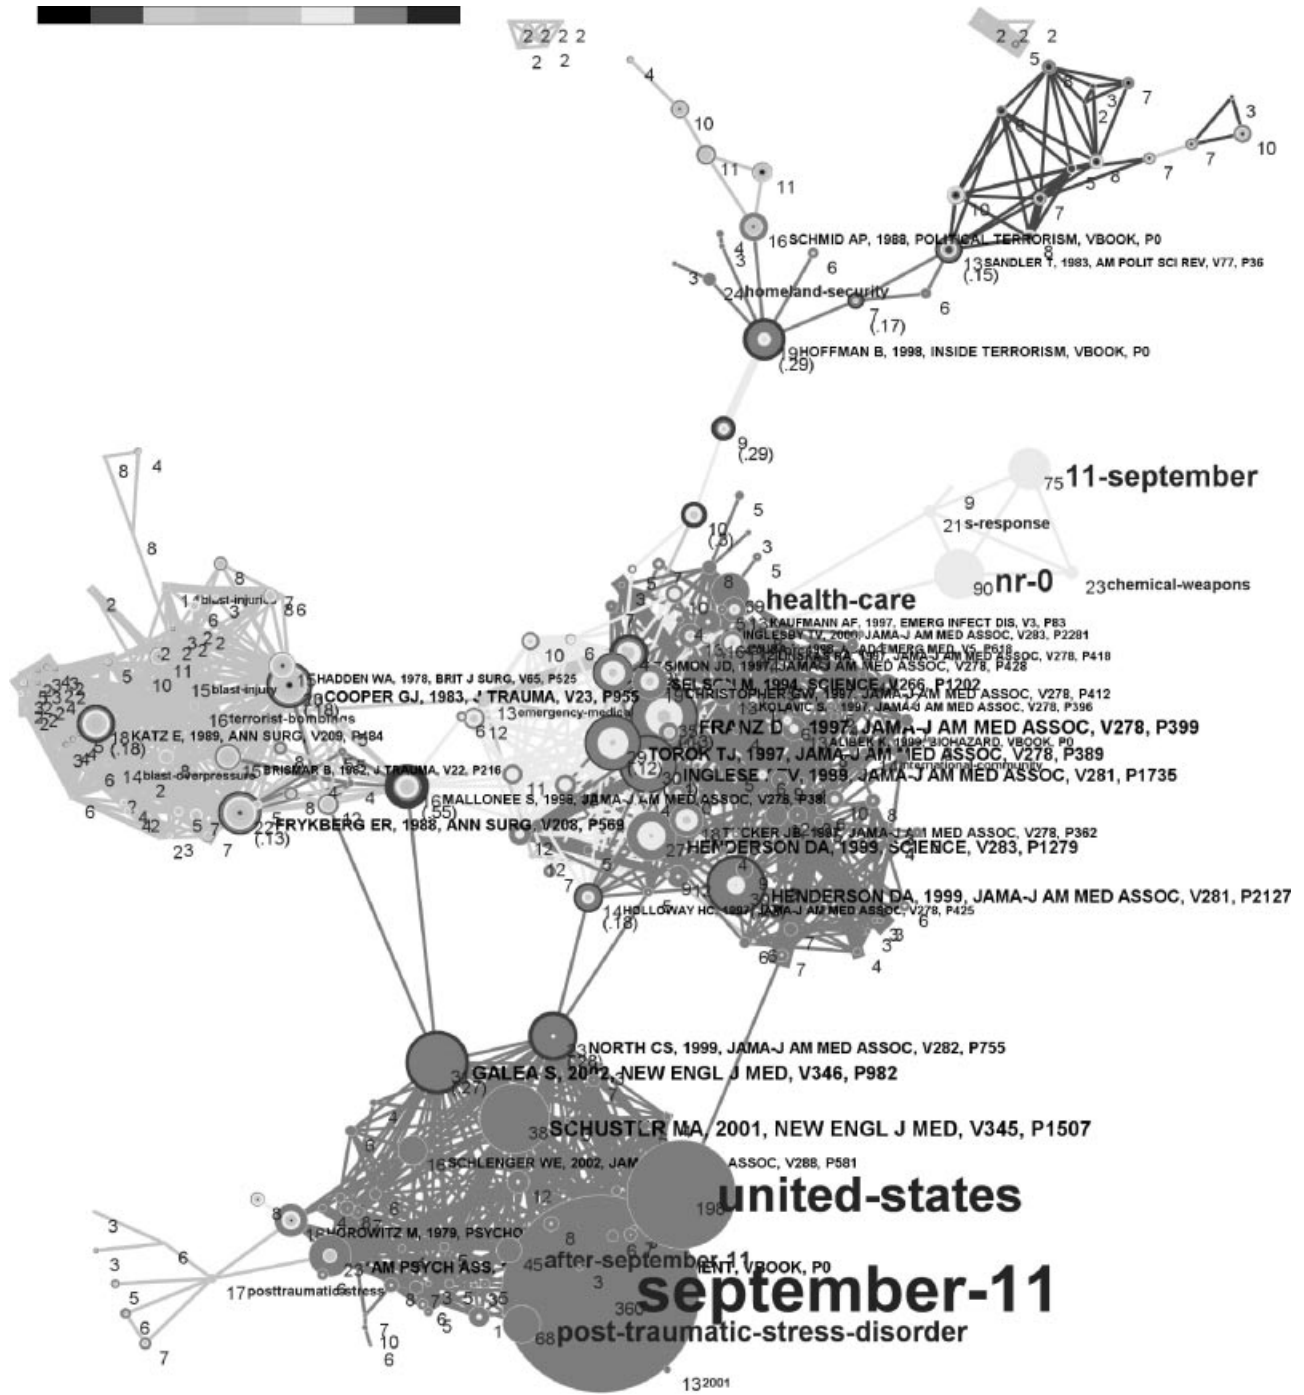


FIG. 8. Three prominent clusters emerge in the 335-node hybrid network of cited article (intellectual bases) and citing terms (research fronts): C_A: body injuries in terrorist bombing (upper left, green), C_B: health care systems in response to the threat of mass destructive biological and chemical weapons (upper right, yellow-orange), and C_C: psychological-psychiatric impacts of September 11, 2001, terrorist attacks (lower central, orange).

network in CiteSpace II on an IBM ThinkPad with a 1.6 GHz Pentium processor and 1 GB RAM.

CiteSpace II implements a multithread version of Pathfinder to reduce the overall waiting time. When CiteSpace II was configured to extract a Pathfinder network from each slice in this example, the process took 20 seconds to complete. Given the rich information it represents, this is expected to be a reasonable trade-off.

The terrorism research map has three prominent clusters (see Figure 8): (a) Cluster A (C_A) is body injuries in terrorist bombing (upper left, green), containing research-front terms *body injuries* and *terrorist bombing*; (b) Cluster B (C_B) is health care in response to the threat of biological and chemical weapons (upper right, yellow-orange), containing terms *health care* and *chemical weapons*, and (c) Cluster C (C_C) is psychological and psychiatric impacts of the September 11, 2001,

attacks (lower middle, orange), containing the terms *September 11, United States*, and *posttraumatic stress disorder*.

C_A represents the specialty on body injuries in terrorist bombing, especially from surgeons' perspectives. It includes intellectual-base articles such as Cooper-1983 (Cooper, Maynard, Cross, & Hill, 1983), Frykberg-1988 (Frykberg, Tepas, & Alexander, 1989), and Katz-1989 (Katz, Ofek, Adler, Abramowitz, & Krausz, 1989). C_B represents the health-care specialty, containing two substructures. Its left half is in bright yellow, and its right half is in orange. The left half includes the term *emergency medical*, and the right one is dominated by the term *health care*. Both parts are concerned with the vulnerability of the health care and emergency systems in response to the potentially massive damages caused by biological and chemical weapons. C_A and C_B are connected via a single pivotal point Mallonee-1996 (Mallonee et al., 1996) (centrality = 55%). The pivotal article reported a comprehensive study of the casualties caused by the 1995 Oklahoma City bombing. It is connected to C_A because it dealt with body injuries and casualties in a terrorist bombing. On the other hand, it also is connected to C_B because of the large number of injuries and casualties caused by the Oklahoma City bombing, which is a central concern of C_B .

C_C represents the specialty that is concerned with psychological and psychiatric disorders as a result of the September 11, 2001, attacks. The PTSD thread includes a 2001 PTSD survey Schuster-2001 (Schuster et al., 2001), a 2002 PTSD survey Galea-2002 (Galea et al., 2002), and a 1999 PTSD study of the Oklahoma City bombing North-1999 (North et al., 1999). Galea-2002 and North-1999 also are identified as pivotal points. Galea-2002 is the single bridge between C_A and C_C whereas North-1999 is the single

bridge between C_B and C_C . C_B and C_C are connected by a pivotal link between Holloway-1997 (Holloway, Norwood, Fullerton, Engel, & Ursano, 1997) and North-1999. Holloway-1997 resides in C_B , which is about the health care preparedness in response to possible mass destructive terrorism attacks; the psychological perspective of Holloway-1997, among other issues, echoes the PTSD aspect of terrorism research. The only other link between C_B and C_C is via a single term: *United States*.

The significance of North-1999 in the entire "big picture" became clearer as we examined the context of its citations. At a glimpse, one can easily and intuitively find articles that are worth a closer look. In CiteSpace II, users can examine the details at several levels. At the top level, the labels of the most prominent items are displayed by default. CiteSpace II provides a number of interactive controls for users to adjust the amount of information shown at this level based on various threshold values. At the second level, the user can bring up one-line reference details by a single click on a node in the network. The user can access further details by a double click on a node.

Figure 9 shows a local view of the overview visualization in Figure 8. This view shows a number of pivotal points between the three clusters (i.e., C_A , C_B , and C_C) and the titles of pivotal-point articles. Two of them are associated with the 1995 Oklahoma City bombing: Mallonee-1996 on body injuries and North-1999 on PTSD among the survivors. Post-September 11, 2001, studies of PTSD heavily drew upon the experiences and findings from the Oklahoma City bombing. On the other hand, the research focus shifted from the survivors of the direct blast in the Oklahoma City bombing, as studied by North-1999, to wider populations. PTSD studies




FIG. 9. Three prominent research fronts in the field of terrorism research.

of the September 11 attacks (e.g., Galea-2002; Schlenger-2001), focused on people who were indirectly exposed to the September 11 attacks in New York primarily from the mass media. CiteSpace II was able to draw our attention to the significant connection between North-1999 and Galea-2002 at a specialty level. The precise nature of the connection was discovered in a subsequent survey targeted to domain experts, namely the authors of the pivotal-point articles. The survey results are described in the following section.

We also generated a time-zone view of the terrorism research. Due to the limited space, we summarize the major findings only. Three major threads emerged. The Oklahoma City bombing thread came to an end in 1997. The most recent terms—*health care* and *after September 11*—were found in the research front of bioterrorism. The two terms were first detected in 2003. The PTSD thread was stretched across the entire time interval, suggesting that PTSD research has a long history despite the fact that it became prominent in the terrorism research only recently. PTSD research studies a broad range of stressors (e.g., triggers of stress such as natural disasters as well as terrorism). The events of September 11 added a new perspective and expanded the traditional scope of PTSD.

Table 6 lists research-front terms that emerged between 1990 and 2003. Terms associated with the September 11, 2001, attacks include *chemical-weapons* and *anthrax-vaccine* in 2001, *New-York* and *posttraumatic stress disorder* in 2002, and *health-care* and *after-September* in 2003. Terms associated with the Oklahoma City bombing, including *blast-injuries* and *federal-building*, surged in 1996. Note that PTSD was spelled out as *posttraumatic* in 1996, but

as *post-traumatic* in 2002 after the September 11, 2001, attacks. *Mental-health* is a closely related concept to PTSD; frequencies of both terms jumped up substantially in 1996 after the Oklahoma City bombing.

Evaluations

The nature of identified pivotal points was evaluated in two ways: (a) a questionnaire to authors of pivotal-point articles as domain experts and (b) content analysis of passages in which pivotal-point articles were cited and co-cited. Due to the limited space, we only present the questionnaire in this article.

A questionnaire containing four questions was e-mailed to authors of pivotal-point articles. The goal of the evaluation was to verify that these pivotal points indeed represent a crucial point in an intellectual transition, a shift of focus, or a change of perspective in the knowledge domain in question. In the questionnaire, we asked each domain expert about the significance of their pivotal-point articles and the nature of connections to neighboring pivotal-point articles. The questionnaire was e-mailed to eight domain experts along with hyperlinks to high-resolution images. Three experts were not reachable because their e-mail addresses were no longer valid. The remaining five experts replied and answered all questions. Due to the limited space, we discuss one expert's answers in detail and summarize other experts' answers. Further details are available from the author upon request.

Galea-2002 and North-1999 are the two pivotal-point articles in the PTSD cluster. Dr. Sandro Galea, the first

TABLE 6. Research-front terms surged between 1990 and 2003.

Terms	Frequency	Terms	Frequency	Terms	Frequency
	2003		1999		1995
health-care	47	biological-weapon(s)	16	peace-process	10
after-September	32	emergency-medical	11	couple-violence	7
care-workers	9	emergency-management	6	emergency-room	7
attacks	8		1998	common-couple	6
security	7	organophosphate-poisoning	9	sarin-vapor	6
	2002		1997	major-trauma	5
New-York	111	terrorist-bombings	15		1994
post-traumatic-stress-disorder	70	blast-overpressure	14	national-security	8
North-Korea	30	warfare-agents	10		1993
homeland-security	24	Tokyo-subway	9	united-states	21
international-community	14	mass-disaster	8	international-news	6
bin-laden	13	biological-terrorism	7	criminal-justice	5
health-protection	13	gastrointestinal-tract	6	nuclear-weapons	4
	2001	lung-injury	6		1992
11-september	75	traumatic-events	6	Northern-Ireland	8
chemical-weapons	23	Buenos-Aires	3	gulf-war	6
America's-response	20		1996	political-violence	6
anthrax-vaccine	11	blast-injuries (injury)	29	United-Kingdom	6
	2000	mental-health	19	blast-loading	5
law-enforcement	14	posttraumatic-stress	17	blast-lung	2
task-force	14	federal-building	6	Iowa-state	2
new-terrorism	7				

author of Galea-2002, provided the following answers regarding Galea-2002.

1. *What is the significance of the issue(s) addressed by your article?*

I think our article was the first to use a validated instrument to document psychological disorders that were linked to the September 11 terrorist attacks showing the psychological consequences of terrorism. There was intense interest in the September 11 attacks and our article provided some empiric rationale for the speculation about potential consequences of the attacks.

2. *What was the major impact or implication of your article on subsequent research?*

There have been a number of articles that have been published after the September 11 attacks, all discussing in some respect the psychological consequences of those attacks/potential implications of terrorism. I think our article (a) established the importance of terrorist events for population mental health (b) clearly laid out the fact that persons in the general population (not just victims) can have psychological disorders after a mass disaster.

3. *Your article has been frequently cited together with the following pivotal points (see details in the table). Could you explain the possible nature of such connections?*

- North, C.S., et al. (1999) Psychiatric disorders among survivors of the Oklahoma City bombing. *JAMA*, 282(8), 755–762.
- Mallonee, S., et al. (1996) Physical injuries and fatalities resulting from the Oklahoma City bombing. *JAMA*, 276(5), 382–387.

I think the “real link” is the first of these two articles. The September 11 attacks were perceived as the “next” big attack after the Oklahoma City bombing and as such early after the attacks there were a lot of references to Oklahoma City. The North paper was the “standard” paper after the Oklahoma City bombing and as such people inevitably made comparisons between the two (very different) papers.

4. *In your opinion, does your article represent a shift of focus, or a change of perspective?*

I think our article identified “everyone in the city” as potentially at risk for PTSD after terrorism, this is not an uncontroversial observation and broadens what people previously thought of as “exposures” to a disaster. In that regard yes, it represents a change in perspective. Subsequent articles have embraced that perspective, although others have contested it.

Other experts also offered their comments, including Dr. Eric Frykberg on pivotal-point article Frykberg-1988 in C_A and Dr. Todd Sandler on pivotal-point article Sandler 1983 (Sandler, Tschirhart, & Cauley, 1983). Dr. Carol North provided valuable insights and interpretations in a 30-minute telephone interview.

Frykberg explained the significance of Frykberg-1988 in C_A : “as a collective review of hundreds of bombing incidents in the literature, it represents the first time to my knowledge that the patterns and severity of injury of casualties of

terrorist bombings were analyzed to derive general principles of management, showing that once patterns were identified, preparation and planning is possible to improve outcome.” Frykberg-1988 introduced a new outcome measure of medical management, namely critical mortality, and concluded that triage is a critical factor in casualty outcome since accurate and rapid identification of that small minority needing immediate care is essential to saving lives in immediate jeopardy.

North-1999 is one of the two pivotal points in the PTSD cluster. North explained in a telephone interview that although the PTSD cluster appeared swiftly after the September 11 attacks, PTSD research has a long history. Earlier PTSD studies were associated with the Vietnam War, the Gulf War, and natural disasters such as earthquakes and hurricanes. Terrorism is a relatively new type of stressor in PTSD research. The sudden popularity of PTSD in terrorism research is largely due to a number of PTSD studies after the September 11, 2001, attacks. These studies, along with the earlier PTSD studies of the Oklahoma City bombing, have established PTSD as a prominent aspect of terrorism research.

North also noted the bioterrorism research front in C_B as an active area of research. Since bioterrorism is regarded as a special type of terrorism, lessons learned from terrorism research in general could be applicable to bioterrorism. As shown in a time-zone view, the research-front term *health-care* in C_B belongs to a current and active research front. North also suggested one way to distinguish the prominent PTSD studies in the cluster in terms of the populations they studied. For example, North-1999 studied the directly injured victims at the scene of the Oklahoma City bombing. Galea-2002 sampled from the entire population in Manhattan, a much broader population than the immediate victims at the World Trade Center scene. Schuster-2001 studied the prevalence of terrorism-triggered PTSD in an even broader population—the entire population across the United States. Schlenger-2002 studied populations increasingly further away from New York.

Discussions and Future Work

CiteSpace II is a system that could be potentially used by a wide range of users, notably scientists, science policy researchers, and research students, to detect and visualize emerging trends and radical changes in scientific disciplines over time. It can be used to explore the dynamics of a specialty in terms of a time-variant mapping from a research front to its intellectual base. The innovative combination of computational and visualization techniques has substantially sharpened the focus of progressive knowledge domain visualization techniques initially proposed and implemented in CiteSpace I. The results of the two case studies are encouraging, especially in light of the validations of domain experts.

The fully integrated representation of research fronts and intellectual bases in the same network visualization has three practical advantages. First, using surged topical terms rather

than the most frequently occurring title words is particularly suitable for detecting emerging trends and abrupt changes. In visualized networks, research-front terms are explicitly linked to intellectual-base articles. This design presents a compact representation of the duality between a research front and its intellectual base. Second, research-front terms naturally lend themselves to be used as labels of specialties. Third, it overcomes a common drawback of word-profile-based labeling approaches. Aggregated word profiles may not converge to an intrinsic focus. Terms selected based on sudden increased popularity measures are particularly suitable to characterize a current research front.

The divide-and-conquer strategy has captured the dynamics of evolving specialties efficiently. The ability to identify pivotal points algorithmically largely reduces users' cognitive load. In this study, we validated the pivotal points in terrorism research with domain experts. More comprehensive evaluative studies are recommended to investigate the effect of alternative centrality metrics.

The Pathfinder algorithm extracts the most salient patterns from a network, but it does not scale well. CiteSpace II implements a concurrent version of the algorithm. The concurrent Pathfinder algorithm has substantially optimized the network scaling module, although it still took 6,000 seconds to process 14 networks and merge them into a 1,704-node network. We also are investigating parallel solutions using cluster computing techniques. Currently, an optimal configuration on a single PC is to limit the visualization to a dozen snapshot networks of 200 to 300 nodes per network so that the process can be completed within tens of seconds.

A long-term challenge is to develop a seamless environment in which users can explore the literature of a scientific discipline at various levels of granularity. Pivotal points may serve as users' signposts and help them to track an evolving specialty. Another challenging issue to be addressed is an integration of citation index and co-citation index with text-processing techniques at the full-text level. Such integrations enable users to not only identify important patterns and relations based on citation indexing but actually access the firsthand information in original articles. In a sense, this line of research may fulfill what Vannevar Bush had envisaged in 1945.

In this article, we have described CiteSpace II as a system devised to facilitate the study of evolving specialties in the form of a time-variant duality between a research front and its intellectual base. There are several possible directions to extend the research. For example, the linkage between scientific literature and technological inventions, namely between articles and patents, has been a central topic in areas such as research policy, science and technology indicators, and knowledge diffusion (Carpenter, Cooper, & Narin, 1980; Tijssen & Vanraan, 1994); however, much work remains; users need tools that can turn a vast amount of data into crystal-clear and instructive messages.

An increasingly important link is between government funding agencies and scientists' outputs. Identifying trends in research-grant solicitations and awarded grants is vital for

scientists. The ability to identify the connections between a funding trend and an existing intellectual base can substantially improve the competitiveness of universities and research laboratories. In addition to the horizontal-expansion opportunities, a number of application domains are particularly in need of trend detection and topic-tracking tools, including bioinformatics and national security-related areas such as intelligence analysis.

A far-reaching research agenda is to develop deep knowledge discovery and data-mining techniques and synergize with knowledge domain visualization challenges. The work of Swanson (1986, 1989) on discovering logically related noninteractive literature and hypothesis formation (Smalheiser & Swanson, 1998) is particularly relevant.

Due to the limited access to the full-text version of articles in various sources, we were not able to locate as many instances of citation context as we had hoped. This experience raised a practical question concerning the extent to which analysts need to access the literature. In an ideal scenario, one would have access to the full text of all the citing and cited articles from a comprehensive digital library. In reality, the closest sources include the ACM Digital Library and the IEEE Explorer. In an even more realistic setting, one would get started with full-text access to just a few core journals in a field. Open-access repositories such as arXive and CiteSeer would be good testbeds. In future work, we plan to extend and tailor our tools to analyze trends in these sources.

Conclusions

In conclusion, the new features introduced to CiteSpace II for detecting and visualizing emerging trends and abrupt changes in a field of research have produced promising and encouraging results. The major findings are that

- the surge of interest is an informative indicator for a new research front;
- using heterogeneous networks of terms and articles provides a comprehensive representation of the dynamics of a specialty;
- research-front terms are informative cluster labels;
- citation tree-ring visualizations are visually appealing and semantically interpretable;
- betweenness centrality metrics identify semantically valid pivotal points.

Challenges and opportunities also are identified in this article for future work. In summary, the major challenges are to

- empirically evaluate the strengths and weaknesses of using burst terms as cluster summaries as opposed to using word profiles;
- empirically compare the pivotal points identified by betweenness centrality versus points identified by alternative centrality metrics;
- develop seamless environments by adapting and incorporating automated text summarization and natural language-processing techniques in addition to statistical approaches so that visual exploration and everyday scientific activities can be integrated closely;

- integrate the visualizations of research fronts and intellectual bases with full-text digital libraries;
- implement and compare trend detection and topic-tracking techniques;
- expand the work to heterogeneous data types such as patents and grant awards.

The provision of a tool-like CiteSpace II enables analysts to perform quantitative and qualitative studies of scientific-subject domains more easily. In terms of intellectual contributions, the study shows the potential of a practical tool for improving our understanding of a research field. The study also demonstrates how the experience and findings of the study have led to an understanding of more challenges and more opportunities ahead. By drawing the strengths from multiple disciplines, the research becomes a melting pot to turn otherwise isolated techniques into an integrative environment.

We have demonstrated the potential of an integrative approach to the visualization and analysis of the evolution of a field of research. In particular, we have focused on various practical issues concerning detecting emerging trends and abrupt changes in transient research fronts. The encouraging results indicate that this is a promising line of research with potentially wide-ranging benefits to users from different disciplines.

Acknowledgments

The author is grateful for generous help from Dr. Paul Wignall at Leeds University, England, on the mass-extinction study and Dr. Carol North, Professor of Psychiatry, Washington University School of Medicine; Dr. Sandro Galea, the Center for Urban Epidemiologic Studies, New York Academy of Medicine; Dr. Eric Frykberg, Professor of Surgery, University of Florida; and Dr. Todd Sandler, School of International Relations, University of Southern California, on the terrorism-research study. The author acknowledges a 2002 Citation Research Grant from the American Society for Information Science and Technology and the Institute for Scientific Information. CiteSpace is available at <http://cluster.cis.drexel.edu/~cchen/citespace>.

References

- Abt, H.A. (1998). Why some papers have long citation lifetimes. *Nature*, 395, 756–757.
- Allan, J., Papka, R., & Lavrenko, V. (1998). Online new event detection and tracking. *Proceedings of ACM SIGIR* (pp. 37–45).
- Alvarez, L.W., Alvarez, W., Asaro, F., & Michel, H.V. (1980). Extraterrestrial cause for the cretaceous-tertiary extinction. *Science*, 208(4448), 1095–1098.
- Alvarez, W. (1997). *T. Rex and the crater of doom*. New York: Vintage Books.
- Bak, P., & Chen, K. (1991). Self-organized criticality. *Scientific American*, 264(1), 46–53.
- Becker, L., Poreda, R.J., Basu, A.R., Pope, K.O., Harrison, T.M., Nicholson, C., & Iasky, R. (2004). Bedout: A possible end-Permian impact crater offshore of northwestern Australia. *Science*, 304(5676), 1469–1476.
- Becker, L., Poreda, R.J., Hunt, A.G., Bunch, T.E., & Rampino, M. (2001). Impact event at the Permian-Triassic boundary: Evidence from extraterrestrial noble gases in fullerenes. *Science*, 291(5508), 1530–1533.
- Boyack, K.W., Wylie, B.N., & Davidson, G.S. (2002). Domain visualization using VXinsight for science and technology management. *Journal of the American Society for Information Science and Technology*, 53(9), 764–774.
- Braam, R.R., Moed, H.F., & Raan, A.F.J.v. (1991). Mapping of science by combined co-citation and word analysis: II. Dynamical aspects. *Journal of the American Society for Information Science*, 42(4), 252–266.
- Burt, R.S. (1992). *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press.
- Burton, R.E., & Kebler, R.W. (1960). The “half-life” of some scientific and technical literatures. *American Documentation*, 11, 18–22.
- Bush, V. (1945). As we may think. *Atlantic Monthly*, 176(1), 101–108.
- Carpenter, M.P., Cooper, M., & Narin, F. (1980). Linkage between basic research literature and patents. *Research Management*, 23(2), 30–35.
- Chen, C. (1999). Visualising semantic spaces and author co-citation networks in digital libraries. *Information Processing and Management*, 35(2), 401–420.
- Chen, C. (2004). Searching for intellectual turning points: Progressive knowledge domain visualization. *Proceedings of the National Academy of Sciences, USA*, 101 (Suppl.), 5303–5310.
- Chen, C., Cribbin, T., Macredie, R., & Morar, S. (2002). Visualizing and tracking the growth of competing paradigms: Two case studies. *Journal of the American Society for Information Science and Technology*, 53(8), 678–689.
- Chen, C., & Kuljis, J. (2003). The rising landscape: A visual exploration of superstring revolutions in physics. *Journal of the American Society for Information Science and Technology*, 54(5), 435–446.
- Chen, C., & Morris, S. (2003). Visualizing evolving networks: Minimum spanning trees versus pathfinder networks. *Proceedings of IEEE Symposium on Information Visualization* (pp. 67–74), Seattle, WA: IEEE Computer Society Press.
- Chen, C., & Paul, R.J. (2001). Visualizing a knowledge domain’s intellectual structure. *Computer*, 34(3), 65–71.
- Cooper, G.J., Maynard, R.L., Cross, N.L., & Hill, J.F. (1983). Casualties from terrorist bombings. *Journal of Trauma*, 23(11), 955–967.
- Erten, C., Harding, P.J., Kobourov, S.G., Wampler, K., & Yee, G. (2003). Exploring the computing literature using temporal graph visualization (Tech. Rep. TR0304), University of Arizona.
- Freeman, L.C. (1979). Centrality in social networks: Conceptual clarification. *Social Networks*, 1, 215–239.
- Frykberg, E.R., Tepas, J.J., & Alexander, R.H. (1989). The 1983 Beirut Airport terrorist bombing: Injury patterns and implications for disaster management. *American Surgeon*, 55(3), 134–141.
- Galea, S., Ahern, J., Resnick, H., Kilpatrick, D., Bucuvalas, M., Gold, J., & Vlahov, D. (2002). Psychological sequelae of the September 11 terrorist attacks in New York City. *New England Journal of Medicine*, 346(13), 982–987.
- Garfield, E. (1994). Scientography: Mapping the tracks of science. *Current Contents: Social & Behavioural Sciences*, 7(45), 5–10.
- Garfield, E., Sher, I.H., & Torpie, R.J. (1964). *The use of citation data in writing the history of science*. Philadelphia: Institute for Scientific Information.
- Garfield, E., Pudovkin, A.I., & Istomin, V.S. (2003). Why do we need algorithmic historiography? *Journal of the American Society for Information Science and Technology*, 54(5), 400–412.
- Girvan, M., & Newman, M.E.J. (2002). Community structure in social and biological networks. *Proceedings of the National Academy of Sciences, USA*, 99, 8271–8276.
- Granovetter, M. (1973). Strength of weak ties. *American Journal of Sociology*, 8, 1360–1380.
- Griffith, B.C., Small, H., Stonehill, J.A., & Dey, S. (1974). The structure of scientific literatures: II. Towards a macro- and microstructure for science. *Science Studies*, 4, 339–365.
- Havre, S., Hertzler, E., Whitney, P., & Nowell, L. (2002). Themeriver: Visualizing thematic changes in large document collections. *IEEE Transactions on Visualization and Computer Graphics*, 8(1), 9–20.

- Holloway, H.C., Norwood, A.E., Fullerton, C.S., Engel, C.C.J., & Ursano, R.J. (1997). The threat of biological weapons: Prophylaxis and mitigation of psychological and social consequences. *Journal of the American Medical Association*, 278(5), 425–427.
- Katz, E., Ofek, B., Adler, J., Abramowitz, H.B., & Krausz, M.M. (1989). Primary blast injury after a bomb explosion in a civilian bus. *Annals of Surgery*, 209(4), 484–488.
- Kessler, M.M. (1963). Bibliographic coupling between scientific papers. *American Documentation*, 14, 10–25.
- Kleinberg, J. (2002). Bursty and hierarchical structure in streams. Proceedings of the 8th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 91–101), Edmonton, Alberta, Canada: ACM Press.
- Kontostathis, A., Galitsky, L., Pottenger, W.M., Roy, S., & Phelps, D.J. (2003). A survey of emerging trend detection in textual data mining. In M. Berry (Ed.), *A comprehensive survey of text mining* (pp. 185–224). Heidelberg, Germany: Springer-Verlag.
- Mallonee, S., Shariat, S., Stennies, G., Waxweiler, R., Hogan, D., & Jordan, F. (1996). Physical injuries and fatalities resulting from the Oklahoma City bombing. *Journal of the American Medical Association*, 276(5), 382–387.
- Morris, S.A., Yen, G., Wu, Z., & Asnake, B. (2003). Timeline visualization of research fronts. *Journal of the American Society for Information Science and Technology*, 55(5), 413–422.
- North, C.S., Nixon, S.J., Shariat, S., Mallonee, S., McMillen, J.C., Spitznagel, E.L., & Smith, E.M. (1999). Psychiatric disorders among survivors of the Oklahoma City bombing. *Journal of the American Medical Association*, 282(8), 755–762.
- Persson, O. (1994). The intellectual base and research fronts of JASIS 1986–1990. *Journal of the American Society for Information Science*, 45(1), 31–38.
- Popescul, A., Flake, G.W., Lawrence, S., Ungar, L.H., & Giles, C.L. (2000). Clustering and identifying temporal trends in document databases. Proceedings of IEEE Advances in Digital Libraries (ADL 2000) (pp. 173–182), Washington, DC.
- Price, D.D. (1965). Networks of scientific papers. *Science*, 149, 510–515.
- Roy, S., Gevry, D., & Pottenger, W.M. (2002). Methodologies for trend detection in textual data mining. Proceedings of the Textmine '02 Workshop at the 2nd SIAM Conference on Data Mining. Retrieved July 1, 2004, from <http://www.cse.lehigh.edu/~billp/pubs/ETDMethodologies.pdf>
- Sandler, T., Tschirhart, J.T., & Cauley, J. (1983). A theoretical analysis of transnational terrorism. *American Political Science Review*, 77(4), 36–54.
- Schlenger, W.E., Caddell, J.M., Ebert, L., Jordan, B.K., Rourke, K.M., Wilson, D., Thalji, L., Dennis, J.M., Fairbank, J.A., & Kulka, R.A. (2002). Psychological reactions to terrorist attacks. *Journal of the American Medical Association*, 288(5), 581–588.
- Schuster, M.A., Stein, B.D., Jaycox, L.H., Collins, R.L., Marshall, G.N., Elliott, M.N., Zhou, A.J., Kanouse, D.E., Morrison, J.L., & Berry, S.H. (2001). A national survey of stress reactions after the September 11, 2001, terrorist attacks. *New England Journal of Medicine*, 345(20), 1507–1512.
- Schvaneveldt, R.W. (Ed.). (1990). *Pathfinder associative networks: Studies in knowledge organization*. Norwood, NJ: Ablex.
- Smalheiser, N.R., & Swanson, D.R. (1998). Using Arrowsmith: A computer-assisted approach to formulating and assessing scientific hypotheses. *Computer Methods and Programs in Biomedicine*, 57(3), 149–153.
- Small, H. (1986). The synthesis of specialty narratives from co-citation clusters. *Journal of the American Society for Information Science*, 37(3), 97–110.
- Small, H. (1999a). A passage through science: Crossing disciplinary boundaries. *Library Trends*, 48(1), 72–108.
- Small, H. (1999b). Visualizing science by citation mapping. *Journal of the American Society for Information Science*, 50(9), 799–813.
- Small, H. (2003). Paradigms, citations, and maps of science: A personal history. *Journal of the American Society for Information Science and Technology*, 54(5), 394–399.
- Small, H., & Greenlee, E. (1980). Citation context analysis and the structure of paradigms. *Journal of Documentation*, 36(September), 183–196.
- Small, H.G. (1977). A co-citation model of a scientific specialty: A longitudinal study of collagen research. *Social Studies of Science*, 7, 139–166.
- Small, H.G., & Griffith, B.C. (1974). The structure of scientific literatures: I. Identifying and graphing specialties. *Science Studies*, 4, 17–40.
- Swan, R., & Allan, J. (1999). Extracting significant time varying features from text. Proceedings of 8th Conference on Information Knowledge Management (CIKM '99) (pp. 38–45), Kansas City, MO: ACM Press.
- Swanson, D.R. (1986). Undiscovered public knowledge. *Library Quarterly*, 56(2), 103–118.
- Swanson, D.R. (1989). Online search for logically related noninteractive medical literatures—A systematic trial-and-error strategy. *Journal of the American Society for Information Science*, 40(5), 356–358.
- Tijssen, R.J.W., & Vanraan, A.F.J. (1994). Mapping changes in science and technology—Bibliometric co-occurrence analysis of the r-and-d literature. *Evaluation Review*, 18(1), 98–115.
- van Raan, A. (2000). On growth, ageing, and fractal differentiation of science. *Scientometrics*, 47(2), 347–362.
- Wignall, P.B., & Hallam, A. (1992). Anoxia as a cause of the Permian/Triassic extinction: Facies evidence from northern Italy and the western United States. *Palaeogeography, Palaeoclimatology, Palaeoecology*, 93, 21–46.
- Wignall, P.B., & Hallam, A. (1993). Griesbachian (earliest Triassic) palaeoenvironmental changes in the salt range, Pakistan and Southeast China and their bearing on the Permo-Triassic mass extinction. *Palaeogeography, Palaeoclimatology, Palaeoecology*, 102, 215–237.