Trailblazing the Literature of Hypertext: Author Co-Citation Analysis (1989¾ 1998)

Chaomei Chen
Department of Information Systems & Computing
Brunel University
Uxbridge UB8 3PH, UK
Tel: +44 1895 203080
E-mail: chaomei.chen@brunel.ac.uk

Les Carr
Department of Electronics & Computer Science
Southampton University
Southampton SO17 1BJ, UK
Tel: +44 1703 594479
E-mail: L.Carr@ecs.soton.ac.uk

ABSTRACT

This paper presents the analysis and modelling of two data sets associated with the literature of hypertext as represented by the ACM Hypertext conference series. This work explores new ways of organising and accessing the vast amount of interrelated information. The first data set, including all the full papers published in this series (1987—1998), is structured and visualised as a semantic space. This semantic space provides an access point for each paper in this collection. The second data set, containing author co-citation counts based on nine conferences in the series (1989—1998), is analysed and mapped in its entirety and in three evenly distributed sub-periods. Specialties — major research fronts in the field of hypertext — are identified based on the results of a factor analysis and corresponding author co-citation maps. The names of authors in these maps are linked to the bibliographical and citation summaries of these authors on the WWW.

KEYWORDS: Author co-citation analysis, domain maps, information visualisation, virtual structures

INTRODUCTION

The ability to create ones own threads of association is central to Bush's visionary Memex [1]. The Memex was designed to organise information in the most intuitive way — the way we think. The World-Wide Web (WWW) has realised much of Bush's vision, but the core insight, the trail-building ability, remains unfulfilled on the Internet [11].

Trailblazing in an information space, also known as threading or trail-building, has been addressed in the literature of hypertext over the last decade from various perspectives. For example, Halasz included virtual structures as one of his famous 7 issues concerning the design of hypertext systems [9]. The ability to create virtual structures over hard-wired structures is a type of trailblazing. Open hypermedia systems such as Microcosm are designed to support dynamic link-node binding [2]: the same set of information may be accessed in different ways based on different configurations of links. This approach may draw trailblazers' attention to how they should construct various link structures in order to adapt existing information resources for a given task. A good example of how one may trailblaze the Web is Walden's Paths [7; 14]. With Walden's Pathds, teachers can effectively assembly a course of learning materials for young children based on existing documents on the Internet.

The ultimate goal of our work is to realise the vision of making the best use of an interrelated information space and building ones own threads of association. As one step in this direction, we explore in new ways of structuring and visualising a domain-specific information space. In this study, we choose the field of hypertext as the subject domain. The ACM Hypertext conference series is used as the core of the literature of hypertext.

In previous studies, we developed the Generalised Similarity Analysis (GSA) framework [4; 5] to unify a variety of analytical and modelling techniques for structuring and visualising complex information spaces. For example, we included content-similarity analysis as a means of structuring and visualising a collection of documents. In this paper, we extend the GSA framework to incorporate author co-citation analysis.

The greatest advantage of using author co-citations as the proximity measure is that not only can authors who are frequently cited together be mapped near to each other in author co-citation maps, but also a specialty — an active sub-field of research — is easily identifiable in association with a cluster or a sub-network of authors. In this study, author co-citation maps are generated for the field of hypertext and we want to find out whether these maps can be used to structure the literature more intelligibly.

The rest of the paper is organised as follows. First, we review related work in the area of author co-citation analysis. Next, we introduce the methodology of author co-citation analysis to be used in the work. Then, we present major results of the factor analysis and author co-citation mapping. Finally, we discuss issues remained and conclusions that we can draw from this work.

RELATED WORK

The idea of mapping the tracks of science is explained by Garfield in [8]. The aim of such work is to identify research front specialties in a field of study. A specialty is characterised by its influence on the development of a given field. One can tell a specialty by the number of citations that it receives. In this study, our aim is to identify some of the most predominant specialties in the field of hypertext and to use author co-citation maps as a means of trailblazing the literature of hypertext.

Specialties

In 1981, Institute for Science Information (ISI) published ISI Atlas of Science in biochemistry and molecular biology [10]. The Atlas was constructed based on co-citation index associated with publications in the field over a limited period of one year. 102 distinct clusters of articles were identified, which were called research front specialties, in order to give researchers a snapshot of significant research activities in biochemistry and molecular biology.

More than 100 people were involved in the project for several months. In the Atlas, journal articles were clustered according to associated co-citation index. The atlas provides a clear, distinct snapshot of the scientific network and how it was structured. More recently, ISI developed Sci-Map software for users to navigate the citation network [15; 16].

Unfortunately, the work at ISI remains largely unknown to the hypertext and digital library communities, for whom a major concern is to make the complex structure of knowledge more accessible.

The ACM Hypertext conference series is the primer forum for hypertext and hypermedia as a discipline and a community. The first conference was held in 1987. The conference has been held annually since then (except 1988 and 1995). If users can navigate through this literature as a whole, we would be one step closer to Bush's visionary Memex.

Another example of navigating through the citation network is the Butterfly system [12]. Butterfly was designed as a user interface for accessing Science Citation databases. It was an application of information visualisation techniques. The design of Butterfly highlighted the role of an organic user interface, in which a virtual landscape grows under user control as information is accessed automatically. In the Butterfly system, the head of a butterfly represents a currently searched article. The wings of the butterfly represent the citing and cited articles associated with the current one. Butterfly provides a new way of navigating through the literature via citation links.

Author Co-Citation Analysis

Author co-citation analysis (ACA) uses authors as data points in the literature. The aim is to reveal clusters of authors, i.e. specialties. The focus of ACA is on authors instead of articles or journals. In our study, ACA is used as a means of visualising a field through a representative slice of its literature.

White and McCain [17] used author co-citation analysis to map the field of information science. In their analysis, the literature was operationalised by the selection of 12 key journals in information science from 1972 through 1995. Their analysis included the top 120 authors ranked by citation counts. They generated maps of the top 100 authors in the field. Their study also included a factor analysis, in which major specialties were identified. One of the most remarkable findings is that the field of information science consists of two major specialties with litter overlap between their memberships: experimental retrieval and citation analysis.

Multidimensional scaling techniques are typically used in author co-citation analysis as a means of depicting the underlying patterns between authors. However, author co-citation studies have been limited by the number of authors that one can map with multidimensional scaling facilities in the SPSS statistical package. For example, they had to limit the maximum number of authors within the capacity of the multidimensional scaling routines. In 1980s, the limit was 40 authors and in 1990s, this number is raised to 100 authors. In this study, this limitation is considerably reduced.

Most co-citation analyses have taken the advantage of ISI's Science Citation Index and Social Science Citation Index. Unfortunately, these citation databases do not include conference proceedings for citation indexing purposes. Therefore, a dedicated citation database for the ACM Hypertext conference series is needed. Over the past few years a group of people at Southampton University have been building a bibliographical database for the conference series. Author citation and co-citation counts used in this study are based on this database.

Using authors as access points may also reduce the complexity of the literature, at least to certain extent. One may use prominent authors' names as a filter to select and partition the literature. In this sense, the extended GSA framework provides a new means of organising and accessing the literature of hypertext.

Why is This Relevant to Hypertext?

Author co-citation analysis is an analytical method that has been traditionally used to identify interrelationships between pairs of authors. It is believed that author co-citation can provide invaluable information about how authors, as domain experts, perceive the interconnectivity between published works. A group of authors may have developed special areas of focus, i.e., the specialties.

In a series of studies, we have been investigating the role of Pathfinder network scaling techniques in a virtual structure-generation methodology that can be used to reduce the excessive number of links and extract the most salient structures from a range of proximity data [3]. One problem we repeatedly encountered is an interpretation problem: users found hard to make sense the nature of links selected by Pathfinder. Partly this is due to the examples we used were based on vector space models and its extended versions, such as Latent Semantic Indexing (LSI) [6], which typically involve the reduction of high dimensional data into lower dimensional representations. In a recent experiment, we found that some of the links may not be intuitive to human beings.

As explained in [17], author co-citation is a more rigorous grouping principle than that of typical subject indexing, because it depends on repeated statements of connectedness by citers with subject expertise. Although author co-citation analysis does not change the high dimensional nature — different aspects of an author's work can be cited in different context, we now have more sources of information available to help us to understand the extracted structure. For example, the representative works of an author and the research group that the author belongs to all provide clues to why they are structured in a particular way. Therefore, we expect that author co-citation analysis can play a significant part in helping people to make sense of computationally generated structures.

METHODS

Two sets of data were used in our study in order to delineate both the structure of the hypertext literature and the structure of the hypertext field. The structure of the hypertext literature is derived from all the full papers published in the ACM Hypertext conference series (1987—1998). Most electronic copies of these papers were gathered from the ACM Digital Library and the rest of papers were retrieved from our own copy of the ACM Hypertext Compendium, which covers 1987, 1989, and 1991 conference proceedings. The title, the names of authors, and the abstract of each paper were included in the final collection for content-similarity analysis.

This collection was automatically analysed and modelled with our Generalised Similarity Analysis tools [3-5]. The resultant model was rendered in Virtual Reality Modelling Lanuage (VRML 2.0).

The second data set was needed for author co-citation analysis, including author co-citation counts from all the papers in nine conferences over the last ten years (1989—1998). In addition to full papers, the second data set also included data derived from short papers. Author co-citation counts were computed for all the authors who were cited five times or more during the whole period. This selection criterion resulted in a pool of 367 authors for the entire period. In order to discover significant advances and trends in the history of the field, the series of nine conferences were grouped into 3 sub-periods. Each sub-period consisted of three consecutive conferences (Table 1).

Period I	Period II	Period III
Hypertext '89	ECHT '92	Hypertext '96
ECHT '90	Hypertext '93	Hypertext '97
Hypertext '91	ECHT '94	Hypertext '98

Table 1. Three sub-periods of the series.

Similarly, author co-citation matrices were generated for three sub-periods using the same criterion. In particular, 196 authors were short-listed for Period I, 195 authors for Period II, and 195 authors for Period III as well. The second set of data was provided by Southampton University.

Factor analysis was conducted on SPSS for Unix Release 6.1 because of the size of the largest author co-citation matrix and the computation-intensive nature of the analysis. Following [17], the raw co-citation counts were transformed into Pearson's correlation coefficients using the factor analysis. These correlation coefficients were used to measure the proximity between authors' co-citation profiles. Self-citation counts were replaced with the mean co-citation counts for the same author. In the factor analysis, principal component analysis with varimax rotation was used to extract factors. The default criterion, eigenvalues greater than one, was specified to determine the number of factors extracted. Missing data were replaced by mean co-citation counts for corresponding authors.

Pearson's r was used as a measure of similarity between author pairs, because, according to [17], it registers the likeness in shape of their co-citation count profiles over all other authors in the set. Pearson correlation matrices were submitted to the GSA environment for processing, especially including Pathfinder network scaling and VRML-scene modelling. An author co-citation map was then generated automatically. Hypertext reference links are provided with these maps between the name of an author in the co-citation map at Brunel University's Web server and the corresponding bibliographical entry in a citation table maintained at Southampton University's Web server. In fact, much of the verification of author co-citation maps was based on the provision of these links. Three sub-period data sets were analysed with the same methodology.

RESULTS

The results are presented in four parts. The first part is a content-based similarity analysis. The second is the factor analysis. The third and fourth parts are an overall author co-citation map and three periodical author co-citation maps.

Content Similarity Analysis

Figure 1 is a screenshot of the visualisation of the semantic space derived from 269 full papers published in the ACM conference proceedings throughout the last decade (1987—1998). The latent semantic space was rendered in VRML 2.0. In this VRML model, each sphere represents a paper published in one of the nine ACM Hypertext conference proceedings. The colours of these spheres indicate the "age" of corresponding papers: papers in the earlier years are darker and papers in recent years are lighter. By looking at the patterns of colours, one may identify emerging research interests in a light-coloured cluster. One may also find out the origin of a recent topic by tracing back to its more central, darker ancestors.

In addition to the overall semantic space, the example in Figure 1 incorporates the results of a search to the same collection of documents indexed by Latent Semantics Indexing ([6]). The search query included terms such as *visualisation*, *spatial*, and *map*. These results are represented by red spikes vertical to the global structure. The higher a spike, the more closely the corresponding paper matched the search query. The distribution of these spikes suggests two major clusters of papers relevant to the search. The cluster at the far end has many recently published papers (in lighter colours), whereas the cluster at the near end includes papers some years ago.

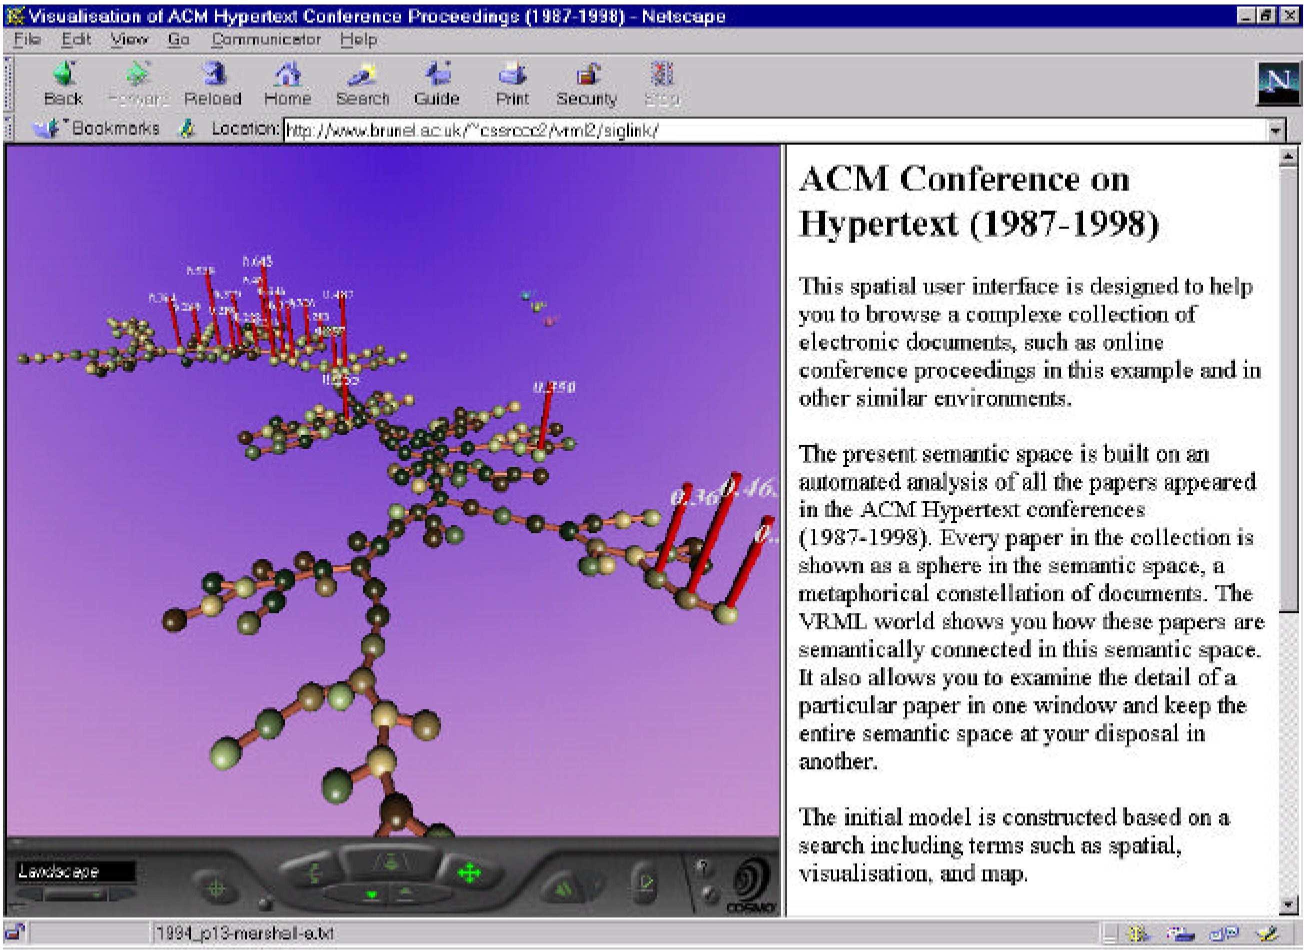

Figure 1. A content-similarity map of the ACM Hypertext conference proceedings (1987-1998).

Note that the centre of the semantic structure is occupied by a large number of dark spheres, which indicate papers published in the early years of ACM Hypertext conference series. Many of these papers have become classics. They are frequently and widely cited. This observation leads to a new perspective to interpret and understand the interrelationships between topics, subjects, and papers in our visualisation model. The closer a paper to the centre of the structure, the more likely that it belongs to a path connecting papers of different types. In the subsequent author co-citation analysis and associated co-citation maps, this tendency becomes even stronger and more intuitive.

Factor Analysis

Thirty-nine factors were extracted from the 367 × 367 author co-citation data set. These factors explain 87.8% of the variance. In particular, the top four factors alone explain 52.1% of the variance (Table 2). In terms of specialties, these four factors imply some substantial sub-fields of study in hypertext. These specialties will be discussed further with reference to author co-citation maps.

In order to understand the nature of each specialty associated with an identified factor, we ranked the authors according to their loadings on each factor and categorised the corresponding specialty based on the profiles of the top 20 authors. Of course, we could describe a specialty more accurately, if we have more detailed information about the nature of a particular co-citation pattern, for example, which particular piece of an author's work was referred to by citing authors.

Factor	Eigenvalue	Pct of Var	Cum Pct
1	110.72	30.2	30.2
2	39.53	10.8	40.9
3	23.39	6.4	47.3
4	17.63	4.8	52.1

Table 2. Eigenvalues of the top four factors.

Table 3 lists the top 20 authors on each factor. Note that these authors may also have significant loading on other factors. The names in Factor 1's column are well-known to the hypertext community and they represent the impact of classic hypertext systems. Therefore, we name Factor 1 as Classics. The nature of Factor 2 is also apparent — we can easily recognise the names of authors whose work was focused on Information Retrieval, for example, Salton, Crouch, and Sematon. The third factor includes names such as Rao, Pitkow, and Jones. We suggest that this factor is about graphical user interfaces and probably an information visualisation-related specialty. The membership of the fourth factor is clearly about links and linking mechanisms, especially the Southampton group. In the following discussion on author co-citation structures, we will identify more specialties based on topological patterns suggested by the co-citation maps.

Authors	F1	Authors	F2	Authors	F3	Authors	F4
Trigg_R	.921	Laurel_B	.592	Duda_A	.701	O'Hara_K	.523
Smith_J	.906	Chignell_M	.588	Sheldon_M	.686	Schilit_B	.523
Moran_T	.901	Gloor_P	.575	Weiss_R	.686	Levy_D	.519
Halasz_F	.885	Cousins_S	.570	Szilagyi_P	.681	Sellen_A	.505
Schwartz_M	.883	Schneiderman_B	.564	Velez_B	.681	Sawhney_N	.498
Conklin_J	.883	Crouch_C	.564	Gifford_D	.654	Golovchinsky_G	.456
Streitz_N	.875	Salton_G	.559	Jones_S	.620	DeRoure_D	.445
Coombs_J	.870	Crouch_D	.559	Danzig_P	.611	Maurer_H	.429
Rogers_R	.870	Andreas_G	.554	Rao_R	.587	Fountain_A	.426
Marshall_C	.868	Hearst_M	.547	Lucarella_D	.562	Carr_L	.414
Thuring_M	.848	McGill_M	.543	Pitkow_J	.544	Jones_R	.406
Engelbart_D	.848	Mylonas_E	.537	Goble_C	.541	Hill_G	.395
McCracken_D	.848	Egan_D	.523	Scholl_M	.531	Andrews_K	.392
Akscyn_R	.844	Allan_J	.521	Chen_C	.522	Bouvin_N	.381
Begeman_M	.843	Buckley_C	.521	Zanzi_A	.518	Davis_H	.378
Brown_P	.840	Smeaton_A	.519	Agosti_M	.482	Heath_I	.377
Meyrowitz_N	.836	Glushko_R	.513	Rada_R	.472	McKnight_C	.375
Yankelovich_N	.836	Landauer_T	.512	Mendelzon_A	.469	Hall_W	.375
Hannemann_J	.834	Guinan_C	.510	DeRoure_D	.455	Haas_C	.370
Parunak_H	.834	Waterworth_J	.508	Bruza_P	.454	Dillon_A	.367

Table 3. The top 20 authors with the highest loadings on each of the first four factors[1].

Overall Author Co-Citation Map

In this section, we present an all-in-one version of the author co-citation structure. This map includes 367 authors over the period from 1989 through 1998. The aim is to establish the overall structure of the hypertext literature, research fronts, and the evolving community of researchers. Then we compare author co-citation maps across three consecutive periods in order to reveal significant trends in the evolution of hypertext as a field.

The all-in-one map is shown in Figure 2, which is essentially a connected graph. Our analysis and interpretation have been based on several heuristics derived from the various computational algorithms used and our experiences with this approach. For example, a node on the main spine of the map is likely to belong to a much larger number of shortest paths connecting two nodes in the graph than a leaf node. In other words, a spinal node plays a more significant role in establishing associations between different nodes than a leaf node. A highly cited article in the field of hypertext is thus likely to be found near to the centre of the overall author co-citation map. Trailblazers may find it useful to include these nodes in their threads of association.

Indeed, we found names such as **Engelbart**, **Nelson**, **Halasz**, **Trigg**, and **Streitz** in the centre of the map. The significance of a central position is that many major branches are only connected because of them. For example, if we remove the links around **Streitz** in the map, several branches of nodes would fall apart from the main spine of the map. This example also illustrates the role of the triangular inequality condition used in Pathfinder network scaling.

We have explained that nodes near to the centre of the map are more essential to the field of hypertext in terms of co-citation patterns. Branching nodes are also significant in identifying the sub-fields of hypertext. We particularly verified branching nodes and used our knowledge about these nodes to suggest the nature of a specialty, which typically includes all the authors in the branch. Eight major specialties are identified from the map:

1. Classics
2. Design Models
3. Hypertext Writing
4. Information Retrieval
5. Open Hypermedia
6. Information Visualisation
7. Structural Analysis
8. User Interface

For example, in Figure 1 **Salton** is a branching node, which connects three branches to the central spine of hypertext. Salton's work in information retrieval and automated hypertext generation has been well-known in both Information Science and Hypertext. We also checked leaf nodes of these branches, because leaf nodes represent some unique characteristics. We found that the leaf node of the right-most branch is **van Rijsbergen**, whose name is prominent in Information Retrieval, it becomes apparent that this branch represents a specialty with strong connections with information retrieval. Therefore, this specialty is called the information retrieval specialty.

We incorporate this map into an interactive user interface on the WWW. Users can click on the name of an author in the map and follow links to the corresponding bibliographical entry of the same author. We actually used this interface to explore the map and verify various citation details. More specialties may be identified and discussed in this way.

Periodical Author Co-Citation Maps

The following analysis focuses on the evolution of the hypertext field over three 3-year sub-periods since 1989.

[1] The complete table is available on request from the first author.

We identify the predominant authors in each period based on author co-citation maps and compare them across these periods.

Figure 3 includes three author co-citation maps: the left-hand side map (L) for (1989—1991), the middle map (M) for (1992-1994), and the right-hand side map (R) for (1996—1998). The map L includes 196 authors who have five or more citations during the first period. We followed the links from the map to detailed citation information and found that authors were clearly grouped with reference to papers describing hypertext systems, which are all well-known today, including NoteCards, Intermedia, KMS, and Microcosm. A specialty of information retrieval was already in place in the first period.

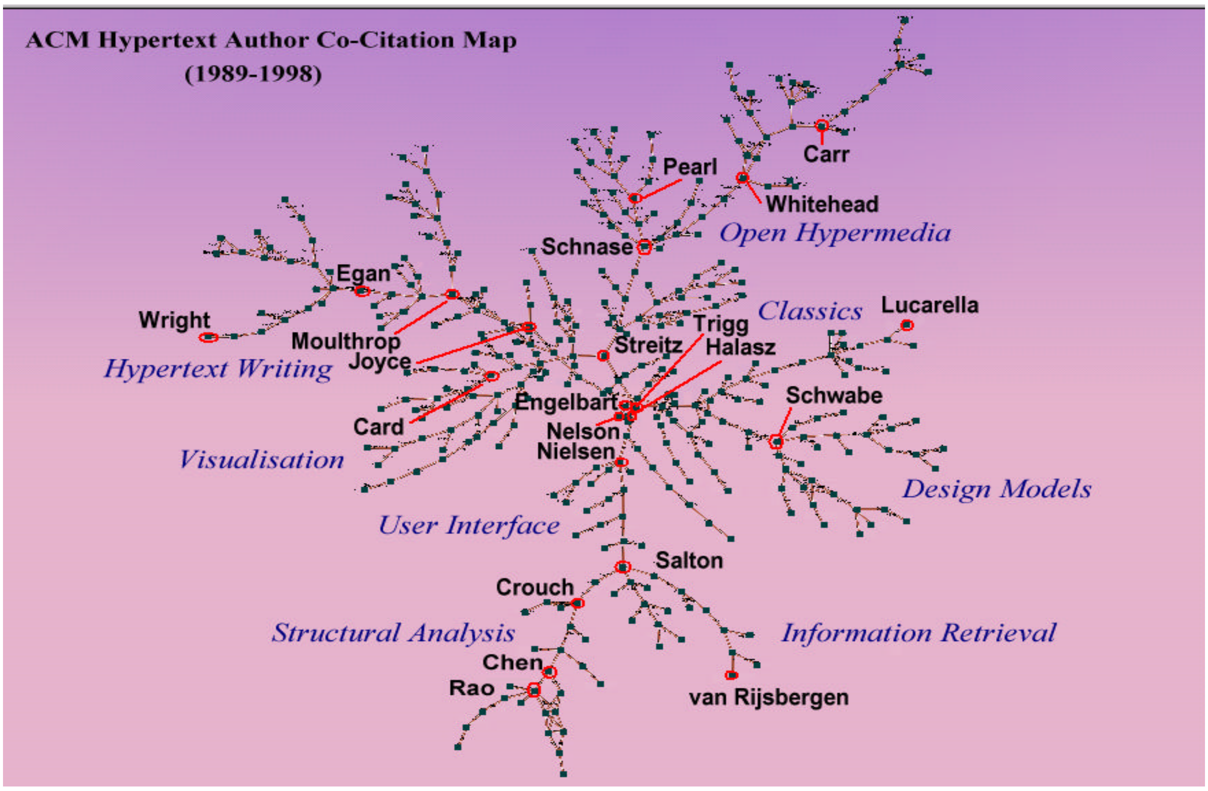

Figure 2. The ACM Hypertext author co-citation map and specialties (1989¾ 1998).

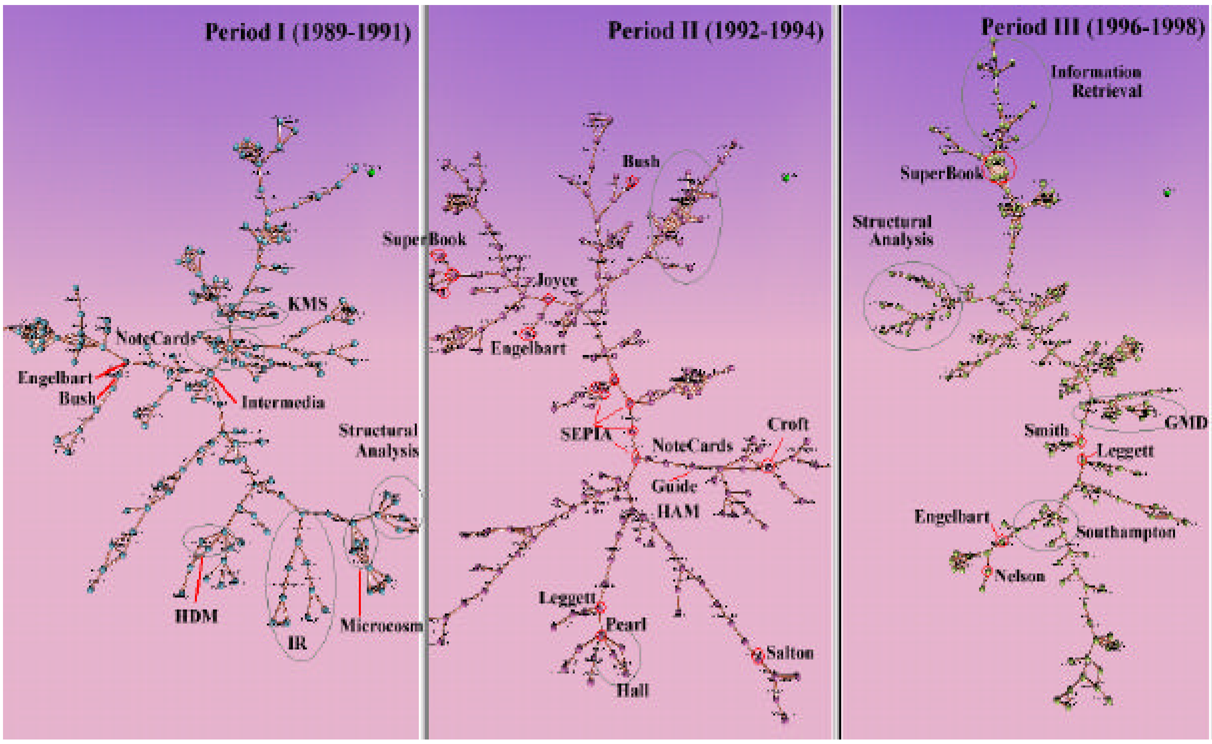

Figure 3. Three snapshots of the evolution of the hypertext field (1989-1991, 1992-1994, 1996-1998).

The second period, 1992—1994, was predominated by SEPIA, a neo-classic system. Six members of GMD, where SEPIA was developed, occupied the central area of the map M. In this map, Pearl became the branching node for the Microcosm group, and Leggett and the Pearl's branch appeared on the same major branch. We have identified an open hypermedia specialty in the overall author co-citation map. This movement indicated the emergent open hypermedia specialty. Remarkably, Salton and Croft are not in the same major branch in this map. Although we need to check more detailed citation information regarding both Salton's and Croft's work in this period, this example shows that author co-citation maps can be used to highlight the impact of particular aspects of one's work.

The latest period ranges from 1996 through 1998. We expected to identify a specialty of the World-Wide Web in map L, given the apparent influence of the rapid advances of the Web and the growth of the WWW as a research field. However, this is not clear from the map. We need to examine the citation details more thoroughly. The information retrieval specialty is located towards the North of the map. The structural analysis specialty is located to the West of the map.

CONCLUSIONS

This is our first attempt to combine a content-similarity analysis, a factor analysis, author co-citation analyses, and structural visualisation together on the field of hypertext as a whole. The factor analysis extracted 39 factors, which may be the roots of many new specialties to be identified in our future work. The nature of a specialty has been identified in the context of author co-citation maps. We have used Pathfinder networks to layout our maps. These maps are different from the multidimensional scaling-based maps typically used in author co-citation analyses, such as [17]. Pathfinder networks can provide more accurate information about local structures than multidimensional scaling maps [13]. We found that the provision of explicit links in our maps made it easier to interpret interrelationships among different data points.

The value of this work is not only in its clustering papers in the ACM Hypertext conference proceedings together based on latent semantics, but also in its ability to thread through the literature and extract the most salient associations among authors who have made significant contributions to the field. Furthermore, author co-citation maps provide a means of identifying research fronts, i.e. specialties in the field, and a visual aid of interpreting the results of factor analysis.

Technically, this is a walk-through of a potentially useful trailblazing methodology. With these tools, users can choose their own grouping principles and criteria in order to build a literature that is intelligible to them. Furthermore, unlike the author co-citation analysis in [17], in which 100 authors were chosen because of the limit of multidimensional routine in the SPSS, the 367 authors in our study is not limited by the resources. This set of authors allows us to take advantage of the power of self-organised modelling and visualisation techniques and depict a natural and comprehensive picture of the field. Citation analysts therefore are likely to benefit from the increased capacity and clarity.

We expect that this approach will be valuable to researchers who want to synthesize a field of study, to people who want to optimize the accessibility of the literature to their learning and communication needs, and to people who want to build intuitively organised digital libraries.

ACKNOWLEDGMENTS

This work is in part supported by the EPSRC (Research Grant GB/L61088). Special thanks to Janet Cole for her assistance in data preparation.

REFERENCES

1. Bush, V. As we may think. *The Atlantic Monthly*, 176, 1 (1945), 101-108.
2. Carr, L., DeRoure, D., Hall, W., and Hill, G. The distributed link service: A tool for publishers, authors and readers. *World Wide Web Journal*, 1, 1 (1995), 647-656.
3. Chen, C. Bridging the gap: The use of Pathfinder networks in visual navigation. *Journal of Visual Lanuages and Computing*, 9, 3 (1998), 267-286.
4. Chen, C. Generalised Similarity Analysis and Pathfinder Network Scaling. *Interacting with Computers*, 10, 2 (1998), 107-128.
5. Chen, C., and Czerwinski, M. From latent semantics to spatial hypertext: An integrated approach, in *Proceedings of the 9th ACM Conference on Hypertext and Hypermedia (Hypertext '98)* (Pittsburgh, PA), ACM Press, 77-86.
6. Deerwester, S., Dumais, S. T., Landauer, T. K., Furnas, G. W., and Harshman, R. A. Indexing by Latent Semantic Analysis. *Journal of the American Society for Information Science*, 41, 6 (1990), 391-407.
7. Furuta, R., Shipman, F., Marshall, C., Brenner, D., and Hsieh, H. Hypertext paths and the WWW: Experiences with Walden's Paths, in *Proceedings of Hypertext '97* (Southampton, UK), ACM Press, 167-176.
8. Garfield, E. Scientography: Mapping the tracks of science. *Current Contents: Social & Behavioural Sciences*, 7, 45 (1994), 5-10.
9. Halasz, F. Reflections on NoteCards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM*, 31, 7 (1988), 836-852.
10. ISI. *ISI atlas of science: Biochemistry and molecular biology, 1978/80*, Institute for Scientific Information, Philadelphia, PA, 1981.
11. Johnson, S. *Interface Culture: How New Technology transforms the way we create and communicate*, HarperEdge, 1997.
12. Mackinlay, J. D., Rao, R., and Card, S. K. An organic user interface for searching citation links, in

Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems, 67-73.

13. Schvaneveldt, R. W., Durso, F. T., and Dearholt, D. W. Network structures in proximity data. The Psychology of Learning and Motivation, 24, G. Bower, ed., Academic Press, 249-284, 1989.

14. Shipman, F. M., Furuta, R., Brenner, D., Chung, C.-C., and Hsieh, H.-W. Using paths in the classroom: Experiences and adaptations, in *Proceedings of Hypertext '98* (Pittsburgh, PA), ACM Press, 267-276.

15. Small, H. A SCI-MAP case study: Building a map of AIDS research. *Scientometrics*, 30, 1 (1994), 229-241.

16. Small, H. Update on science mapping: Creating large document spaces. *Scientometrics*, 38, 2 (1997), 295-293.

17. White, H. D., and McCain, K. W. Visualizing a discipline: An author co-citation analysis of information science, 1972---1995. *Journal of the American Society for Information Science*, 49, 4 (1998), 327-356.