

Innovation Trajectories for Information Visualizations: Comparing Treemaps, ConeTrees, and Hyperbolic Trees

Ben Shneiderman, *Senior Member, IEEE*, Cody Dunne, Puneet Sharma, and Ping Wang

Abstract— This paper reviews the trajectory of three information visualization innovations: treemaps, conetrees, and hyperbolic trees. These three ideas were first published in the early 1990s, so we are able to track academic publications, patents, trade press articles over almost two decades. We describe the early history of each approach, problems with data collection from differing sources, appropriate metrics, and strategies for visualizing these longitudinal data sets. This paper makes two contributions: (1) it offers the information visualization community a history of how certain ideas evolved, influenced others, and were adopted for widespread use and (2) it provides example of how such trajectories of historical trends can be gathered and visualized. Guidance for innovators and future analysts is offered.

Index Terms—information visualization, innovation trajectories, graphical user interfaces, treemap, cone tree, hyperbolic tree.



1 INTRODUCTION

THE healthy outpouring of innovations from the information visualization community has raised important questions about how to measure the efficacy, adoption, and durability of these innovations. These measures could help retrospective analyses that seek to compare several technologies, but the greatest interest is predictive models that forecast eventual impact of novel technologies. This goal may be difficult to attain, but simpler descriptive and explanatory theories can be helpful to guide future entrepreneurs [21], corporate research managers [1], government funding agency staff [20], and historians of science as they seek to understand the evolution of technology [4][5]. Another goal is to develop prescriptive theories based on these metrics, which suggest guidelines and policies for promoting technological innovations.

Initial questions of *efficacy* can be partially settled by empirical testing with traditional controlled experiments with a few dozen subjects for a few hours that compare an existing visualization against the innovation [32]. Critics of this approach suggest that short training periods with standard tasks are insufficient to test innovations that may require more substantial training, a wider variety of tasks, and even changes to familiar problem-solving strategies [33] and work processes and practices. These critics advocate extended case studies with users who work for weeks and months using the innovative visualization on their own tasks. Longer term measures of efficacy account for other changes such as improved interfaces, better training, simpler integration with other tools, and a community of like-minded users who are capable of discussing advanced uses of an innovative visualiza-

tion. Evaluations of efficacy often lead to refinements, clarification of which tasks are more helped, and provide a basis for promoting an innovation.

Measuring *adoption* (also called acceptance or diffusion) beyond the originators is also a challenge as innovations may spread in different ways [30]. Simple measurements in the 2-5 years following the initial presentation of an innovation, include the references in later papers, implementations by multiple open source or commercial organizations, and use by a growing set of users. Some impacts are more noteworthy than others, such as solving important problems, sales of commercial products, or inspiring further innovations. However, impacts may take years or decades to emerge and innovations may need much transformation or integration with other ideas to have impact. The longer-term (5-15 years) diffusion through organizations, industries, and countries is likely to be segmented, as in reaching novice or experienced users, professionals or consumers, and old or young users. The Technology Acceptance Model [12] focused on perceived usefulness and perceived ease of use, but later versions incorporated other parameters [37].

Measuring *durability* (also called sustainability) over decades is important in understanding why innovations lose their enthusiastic advocates or satisfied users, how later innovations replace earlier innovations, and how innovations become so accepted as to become invisible, unmentioned, and taken-for-granted.

These measures may be difficult to capture, difficult to compare, imprecise, and unstable over time. For example, the number of websites mentioning an innovation only began to be relevant by the later 1990s, so comparisons with earlier innovations is difficult. Download counts for some software is useful, but sometimes a few hundred downloads by valued users can be a success while in other situations millions of downloads may be needed to

- Ben Shneiderman, Cody Dunne, and Puneet Sharma are with the Department of Computer Science and Human-Computer Interaction Lab, University of Maryland, College Park, 20742. E-mail: {ben, cdunne, puneet}@cs.umd.edu.
- Ping Wang is with the College of Information Studies, University of Maryland, College Park, 20742. E-mail: pwang@umd.edu.

claim success.

Another class of difficulties relate to names used for search engines, digital libraries, or databases. Some innovations have distinct names that are easy to track, while others have generic names that are harder to search on. An innovation may be given different names by those who make refinements and credit may not always be given to the originators. Often an innovation becomes included in a larger product so it is difficult to track.

This paper uses counts of academic papers, patents, and trade press articles as measures to describe the innovation trajectories of three tree visualization methods: treemaps, cone trees, and hyperbolic trees. A secondary measure is the number of citations to key academic papers and patents. We use these measures and other sources to try to explain their relative successes and give guidance to promoters of information visualization innovations.

A still more difficult goal is to use early measures to predict which future information visualization innovations will become successful. This latter goal is aligned with current U.S. National Science Foundation efforts to develop a “Science of Science”, which would help program managers allocate funding more effectively to high-payoff research proposals.

This paper begins with a personal historical review of tree browsing and visualization methods that focuses on treemaps, cone trees and hyperbolic trees. Then it discusses how to gather data for tracking the trajectory of these methods from academic publications to commercial applications, followed by data visualization to show these trajectories. The conclusion discusses the difficulty of such reviews, lessons learned, and future work.

2 TREE BROWSING AND VISUALIZATION

Even early uses of computers, such as Doug Engelbart’s famous 1968 demonstration of his Augment system presented strategies for browsing tree-structured information, often described as hierarchies. Some tree browsers simply used indented textual representations, which became standard in many directory browsers such as the Microsoft Windows Explorer. Indented textual representations allow rapid scanning down the list, alphabetical ordering, and comprehensible expand/contract strategies to support exploration. The disadvantages include the need for frequent scrolling as the number of nodes grows and the difficulty in discovering large subtrees that might be several levels down.

The appeal of graphical user interfaces encouraged many developers to create node-link diagrams which became widely used during the 1980s. These visualizations were readily understandable, whether drawn with the root node at the top, bottom, left, or right side. With small trees of 10-50 nodes this strategy was effective, but with larger trees that had large depth (many levels) or high branching factor (high fan-out), drawing a complete tree was impossible, even with megapixel displays. While panning and zooming facilitated exploration, and overviews helped even more, the node-link diagram had its

limitations.

2.1 Early History

By the early 1990s several research groups developed innovative methods of tree browsing that offered fresh overview and browsing strategies. In each case they developed a space-limiting approach which ensured that the entire tree would always be seen on standard displays, thereby avoiding the need to pan and zoom, but requiring some new user actions. This paper focuses on three of these tree browsing methods:

Treemaps used nested rectangles to show tree structure, producing not only a space-limiting but also a screen-filling algorithm. The area of each leaf rectangle was determined by one of its attributes, and interior rectangles were sized by the sum of the attribute values of its subtrees and colored by any other attribute value. The original recursive algorithm, slice-and-dice, organized subtrees in a meaningful order (alphabetical, chronological, etc.), but often resulted in long thin rectangles (high aspect ratios) [31][17]. Improved layouts produced more square-like aspect ratios (i.e. closer to 1) [41][7], but sacrificed the lexicographic ordering. Bruls et al.’s “squarified” algorithm caught on widely because it was visually appealing (typically placing large squares in the upper left and small squares in the lower right), eased selection, and often permitted better labeling. Still further refinements offered ordering, low aspect ratios, and some new strategies [6]. The University of Maryland HCIL treemap history page offers numerous links¹.

Cone trees used a 3-dimensional layout with the root node at the top and the first level of nodes were hung down in a circular layout connected by links to the root node, thereby looking like a cone [29][28]. Lower levels of nodes were also laid out in cones all the way down to the leaf nodes at the bottom of the screen. The appeal of 3-dimensional tree structures was strong and the animated rotation of cones to bring occluded nodes to the front made for eye-catching demonstrations. Perspective effects, lighting models, and shadows added to the impressive visual appeal. When the root node was placed at the left side and lower levels to the right, the design was called “cam tree” since the appearance was similar to an automobile cam shaft. Tree-structure traversal and node name searching were the main emphasis, so node color and size coding was not discussed. Robertson, inspired by early scientific visualization researchers, applied the power of high-end graphics workstations to create abstract data animations based on physical processes. Their published video produced a vigorous response and great interest.

Hyperbolic tree browser retained the appealing node-link visual presentation but placed the root node in the center with first-level nodes arranged around it in a circle or oval. Further levels were placed in larger concentric

¹ www.cs.umd.edu/hcil/treemap-history/

circles or ovals [19], thus preserving a 2-dimensional planar approach. To ensure that the entire tree would be visible, outer levels were shrunk according to a hyperbolic formula. The compelling aspect of the hyperbolic tree browser was that users could drag any node to the center, thereby redrawing the tree in a smoothly animated way that was innovative, playful, and eye-catching. The developers were inspired by Escher's art work and drew on the popular notion of fisheye layouts that magnified the central area while diminishing the peripheral areas. They adopted the then current terminology of "focus+context", suggesting that the main item of interest was entirely viewable, while the context remained visible, but reduced in size. Node labels were readable for nodes in the focus area while truncated labels were used at the periphery. Using node size and color to represent node attribute values was introduced in later versions.

Treemaps were invented by the first author of this paper (Shneiderman), while working at the University of Maryland. He was seeking to solve an immediate problem of understanding disk space usage on a machine shared by 14 users. Inclined to visual solutions, he tried many variations until he had the inspirational moment that led to the compact recursive algorithm. He then worked with graduate student Brian Johnson to implement the algorithm and refine the design. The original paper dragged through the journal review and production process, so it was published in 1992, months after the implementation description appeared at a 1991 conference which had a more rapid turnaround. Cone trees and the hyperbolic tree browser originated at the Xerox Palo Alto Research Center (PARC), where a well-funded team of researchers was developing advanced visual interfaces from as early as 1988. While cone tree authors Robertson, Card, and Mackinlay were distinct from the hyperbolic tree browser authors, Lamping, Rao, and Piroli, they collaborated on other projects and co-authored related papers. The University of Maryland did not file any patents for treemaps, but Xerox submitted multiple U.S. patents on aspects of cone trees (5295243, 5689628, 6088032, and others) and hyperbolic trees (5590250, 5619532, 6300957, and others).

All these researchers were enthusiastic about their innovations and had many opportunities to present their ideas in journals, at conferences, and in professional talks. They were all respected and productive researchers, whose work would draw attention from colleagues. As a university researcher Shneiderman had the advantage of graduate students who would implement and evaluate variations on treemaps, while the Xerox PARC researchers were a capable well-funded group of professionals who were embedded in a community skilled at producing innovations.

2.2 Proving Efficacy

Proving efficacy of new ideas is often difficult, as the innovators struggle to identify relevant tasks, decide on the degree of training, choose appropriate subjects, and assess long-term performance. First versions of new ideas typically require many refinements to reach maturity plus integra-

tion with existing technologies. While the two initial tree-map papers did not offer evaluations of efficacy, work was already underway on two evaluations [35]. The first, using 12 participants, showed performance benefits for treemaps over UNIX commands for specific directory browsing tasks. The second, using 18 participants, showed performance benefits for treemaps over paper reports containing financial data. These modest studies were especially helpful in understanding the strengths and weakness of treemaps, which led to many improvements.

Cone tree developers did not conduct traditional evaluations:

"Evaluation of this work [on cone trees] was a topic of frequent conversations. We had a firm belief that we needed a more holistic approach to evaluation of the whole user experience rather than the usability techniques which were in use at the time (which tended to focus on specific design features, and we believed might be misleading). Lacking that holistic evaluation approach, we did no formal user studies. So, the only informal evaluation that was done was based on our own experience as information workers, and the observations we made when using the new techniques compared to the existing ways of looking at the same information structures." (Robertson, personal communication, 3/21/2010).

The initial hyperbolic tree browser paper described an evaluation of four users with a World Wide Web hierarchy of URLs. The counterbalanced within-subjects design produced subjective preferences that favored the hyperbolic tree browser compared with a "conventional 2-d scrolling browser with a horizontal tree layout." However, there were no significant performance differences.

Subsequent evaluations of tree searching were conducted by many other researchers. A comparison of several space-filling tools for tree visualization found problems with treemaps but performance improved with experience [34]. However, their tasks emphasized traversal and nesting levels for which treemaps are known to be weak, while tasks involving node attribute values were not included.

2.3 Adoption and Durability

Many sources describe the evolution of novel ideas into mature technologies and possibly commercially successful innovations [4][5]. Measuring the efficacy, adoption, and durability of an innovation remains difficult, but clearly some ideas such as the World Wide Web have enormous impact and widespread adoption, while other ideas like pie menus have a narrower impact and limited adoption. For the innovations discussed in this paper, the remainder of this section covers critical incidents, influence on other research groups, commercialization by companies, and visibility in popular media. Section 3 covers citation analysis for academic papers, patents, and trade publication articles.

As Moore [21] points out innovators face a modest challenge in gaining early visionary adopters and the much greater challenges in crossing the chasm to reach an early majority of the more pragmatic business users. His arguments are especially relevant to disruptive innova-

tions that require changes to existing practices and ways of thinking. Novel visualizations still face this challenge and have the additional burden that many people are resistant to visualizations compared to textual or tabular numerical presentations. Educational efforts to promote visual literacy are spreading as more interactive visual strategies are developed and as the benefits in corporate utilization grow. Innovators are typically enthusiastic about the benefit of their work and they successfully engage with other innovators who are attracted to novelty. However, reaching the wider circles of early adopters, the early majority, and the later majority that Moore describes can be difficult.

Proponents of novel information visualizations frequently report that an effective method is to show potential users their own data or at least familiar data sets using the novel method. This is in harmony with Moore's advice to focus on narrow vertical market segments where measurable advantages can be shown.

Treemaps were developed for disk directory browsing which had natural and meaningful variables for size (file size) and color coding (file type or age). A memorable incident was a presentation by Shneiderman to the University of Washington Dept. of Computer Science on January 13, 1993. Faculty and student attendees were skeptical about the treemap as a directory browser and even hostile to the idea of replacing UNIX, DOS, or current graphical user interfaces. After the talk, Shneiderman went to their lab and inserted an early demonstration disk into one of their workstations to show them their disk directory. The gasps from viewers were immediate - they saw that much of their space was wasted by three copies of a large compiler that had been installed by separate users. The University of Maryland did not file patents but did succeed in securing at least 17 licenses of the software, bringing enough income to support several student Research Assistants.

While several companies began offering a treemap addition to their directory browser adoption was slow. Jarke van Wijk and his team developed a free program, SequoiaView, that applied his refinements of squarified and cushion treemaps. SequoiaView produced hundreds of thousands of downloads, a popular following, and a spin-off company called MagnaView.

A big increase in adoption and an expansion to a very different application domain came in 1999 when Martin Wattenberg working for Smartmoney.com developed Map-of-the-Market² to show more than 500 stocks in 11 industry sectors, which were divided into subsectors, all as leaf nodes at the third level. Wattenberg's clustered layout, developed independently of the squarified layout, used size to indicate market capitalization and color to indicate degree of rise or fall since the previous day.

This free web application which had good labels and helpful tooltips with links to extensive background data generated a huge following. However, Smartmoney.com was only moderately successful with charging a monthly fee for more sophisticated treemaps or in licensing their

software. Shneiderman became a consultant to Smartmoney.com in 1999 and developed a working relationship with Wattenberg that led to further research results [6].

Another free public website that increased impact was Marcos Weskamp's Newsmap³ which used a treemap to visualize the data on active news stories as reported by the Google news aggregator. In April 2004, just after Weskamp released Newsmap, it generated 114 comments on Slashdot.org⁴.

By 2007, Wattenberg, working at IBM Research with others, produced ManyEyes⁵, a popular collaborative visualization web site with 20 visualizations including treemaps. This open tool expanded adoption as thousands of users were able to upload their data and view the results in treemaps or other visualizations. Also in 2007, The New York Times began using their own interactive treemap software. They continue to use treemaps to help tell stories such as auto industry sales, financial trends, budget allocation, and inflation factors.

These and other treemap applications as well as several open source implementations increased the impact and broadened the adoption, but specific numbers of users are hard to find. The movement to treemap commercialization was advanced by companies such as The Hive Group (for which Shneiderman is a Technical Advisor), Panopticon, Lab Escape, Macrofocus, Magnaview, ILOG, and others listed on the Wikipedia page⁶. The Hive Group's CEO, Jim Bartoo reports: "Hundreds of organizations have licensed the company's Honeycomb visualization software including some of the largest banking, oil & gas, pharmaceutical, and consumer products companies in the world... With the exception of the economic crisis associated with 2009, The Hive Group's sales increased between 40% and 50% each year since 2005" (personal communication, 3/28/2010). In addition, the Wikipedia page lists 16 versions of treemap software in development libraries in these languages: Adobe Flex, Java, JavaScript, Silverlight, Perl, Python, Ruby, and SVG. Section 3 reports on the citation and patent history. It may be too early to discuss durability, but numbers of companies, websites, and users continues to grow.

Cone trees produced a strong immediate response with great interest in the academic community. In conjunction with the first published paper [29] at the 1991 ACM CHI Conference they also published a video "Information Visualization Using 3D Interactive Animation," which was distributed by the *SIGGRAPH Video Review* as a videotape to academic researchers, course instructors, and industrial laboratory developers. A group at the University of Waterloo created their own cone tree implementation with fewer 3D effects and what they believed to be refinements such as a better node layout [8]. They conducted a small user test with five users and three tasks with mixed results but some guidance for refinements. A variation of the cone tree for MySQL databases was shown in a 2006 YouTube video by

³ newsmap.jp

⁴ news.slashdot.org/article.pl?sid=04/03/31/2332203

⁵ manyeyes.alphaworks.ibm.com

⁶ en.wikipedia.org/wiki/List_of_Treemapping_Software

² www.smartmoney.com/marketmap

Daniel Bierwirth⁷. A commercial application of cone trees was created by Xerox spinoff XSoft under the name Visual Recall but it did not reach many users. A demonstration of the S3 graphics chip used cone trees in their 1997 Visidrive 3D product. No Wikipedia article was found for cone trees.

Hyperbolic trees was first published in a two-page description in a 1994 conference [18], but the fuller description appeared in the 1995 ACM CHI Conference, which generated excitement and widespread interest [19]. The original video, “Visualizing Large Trees Using the Hyperbolic Browser,” published by the CHI Conference by way of the *SIGGRAPH Video Review*, shows examples of organization charts and web browsing⁸. Hyperbolic trees (HT) generated variations from university groups and was spun off from Xerox PARC as a commercial product from Inxight. There were commercial and public uses such as web-based hyperbolic trees for popular web sites like NASA. Inxight CTO Ramana Rao reports that they “licensed toolkits to implement HT in vendor’s or solution provider’s packages (~100 deals) and several end user products including an enterprise server product (~500 servers) that allowed for incorporating HT into internal web apps and an end user product for publishing site maps on websites (~5000 sites) and maybe 500k users of a Windows navigation freebie called Magnifind. The revenue over 10 years was on the order of \$30M” (personal communication, 3/27/2010).

Hyperbolic trees were renamed Startrees by Inxight, which was acquired by Business Objects in 2007, which in turn was acquired by SAP, which continues to offer licensing at its website⁹. Inquiries to SAP about sales history were not responded to. Rao laments the absence of a champion at SAP and the declining number of public web-based uses.

The Wikipedia article on hyperbolic trees¹⁰ gives seven external links including examples that remain on the web such as the Green Tree of Life from Univ. of California-Berkeley¹¹. An unidentified developer provided an open source version called “Hypergraph” on the Sourceforge website¹². However, strong patent protection by Xerox PARC raises questions about legality of usage of this code or the Inxight source code, which was also posted at Sourceforge with clear warnings.

One enthusiastic writer for an information professionals’ magazine wrote about “The Hype Over Hyperbolic Browsers” [2]: “For libraries, this approach could revolutionize subject searching” but she cautions that “Unfortunately, I predict that it is unlikely we will see hyperbolic browsers or visually orientated relational databases replacing traditional search engines any time soon. Current

⁷ www.youtube.com/watch?v=1eO1pgTVu-g

⁸ www.youtube.com/watch?v=pwpze3RF55o

⁹

www.sap.com/solutions/sapbusinessobjects/large/business-intelligence/dashboard-visualization/advanced-visualization/index.epx

¹⁰ en.wikipedia.org/wiki/Hyperbolic_tree

¹¹ ucjeeps.berkeley.edu/map2.html

¹² hypergraph.sourceforge.net/examples-general.html

TABLE 1
CITATION AND DOWNLOAD COUNTS FOR EACH INNOVATION

	Google Scholar Cites	ACM DL Cites	ACM DL Downloads last 6 weeks	ACM DL Downloads last 12 months	CiteSeer Cites
Treemap					
ACM TOG 1992 [28]	764	145	61	373	237
IEEE Vis 1991 [16]	754	155	52	321	164
Cone tree					
ACM CHI 1991 [29]	1132	242	48	450	382
CACM 1993 [28]	576	126	75	557	227
Hyperbolic tree					
ACM CHI 1995 [19]	833	154	18	246	263
UIST 1994 [18]	207	48	16	53	76

Citation and download counts as of March 20, 2010 for the key academic papers for each innovation. All three designs show strong and continuing patterns of interest, with cone trees generating the strongest showing.

conventions for searching and information retrieval are so ingrained that will take an enormous cultural shift among information professionals to pave the way for graphically oriented search tools.” This comment, painful to all information visualization researchers, may have an important message about the difficulty in gaining acceptance to many of our innovations.

Competition over browsing strategies for hierarchies produced a lively event at the ACM CHI’97 conference in which six teams competed to browse a large hierarchy [24]. Results were complex but Ramana Rao using the hyperbolic tree browser clearly stood out. However, the tasks chosen involved only navigation over named nodes with no attribute values so the benefits of treemaps’ size and color coding were not applicable.

Other evaluations of the hyperbolic tree browser helped reveal its strengths and weaknesses for specific tasks [27][26][42]. A survey article covering overview+detail, zooming, and focus+context interfaces discussed tree browsing plus other applications, reviewed the empirical studies, and provided a principled analysis based on perceptual psychological issues [10].

3 CITATION COUNTS, PATENT HISTORY, AND TRADE PRESS ARTICLES

Tracking academic papers would seem to be relatively easy since these data are extremely important to many people and there is a long history of analyzing these data for research projects, tenure decisions, scientometric analyses, and history of science studies. Academic authors have an obligation to cite relevant previous work which

TABLE 2
SEARCH RESULTS FOR EACH INNOVATION

	Google	Google Scholar	ACM Digital Library	HCI BIB	Web of Science	IEEE Xplore	Cite-seer	USPTO	LexisNexis
Treemap	318,000	7090	132	16	24	57 + 27	602 + 1199	38	112
Cone tree	25,700	1150	19	3	0	4	59 + 396	44	4
Hyperbolic tree	87,700	1020	20	5	5	6	197	38	45

Frequency counts from various services as of March 20, 2010. The term *treemap* has a strong showing, but the Google numbers especially are inflated since some references are to a phylogenetic term, a Java program, and national park maps. The ideas and papers related to cone trees and hyperbolic trees are often cited, but the term is not always used, giving it an unfairly lower count. The Web of Science results were checked to verify relevance. IEEE Xplore and Cite-seer results are for (*treemap*) + (*tree map*) or for (*conetree*) + (*cone tree*).

makes tracking influence possible, especially when facilitated by some search services. Similarly, there is strong interest in patents and patent authors are required to point to prior art that influenced their innovations. The U.S. Patent and Trademark Office provides free public access to the patents, but extracting links to previous patents or academic papers requires diligent effort, since these links are not well supported. Trade press articles are available through commercial search services such as Lexis/Nexis and free services such as Google's or Bing's search engine. However, trade press articles rarely link to or mention academic papers, patents, or related trade press articles. Instead they name corporations or products, mention key terms or concepts, and quote from key individuals, but extracting this information is difficult.

An initial study of academic citation patterns for the six primary papers shows the strongest and most continuing interest in the original cone tree papers (Table 1). An initial study of use of the terms treemaps, cone trees, and hyperbolic trees shows the strongest pattern of mentions for treemaps possibly indicating broader interest beyond the research community (Table 2). These searches give an overall impression, but detailed analysis is necessary as they may include extraneous citations and mentions while omitting relevant ones. The terms are fairly specific, but treemaps are also used in phylogenetic research and as a Java program, while cone trees generate some results from botanical studies.

To find more detailed, relevant results as well as analyze the trajectory of citations over time we decided to pursue a more conservative approach. We compiled datasets for the three innovations from three sources. For academic papers we obtained from ACM an XML database dump of the ACM Digital Library¹³. This provided us with metadata for the vast majority of the papers published dealing with these visualization innovations and should provide a representative sample of the visualization literature. Moreover, it also includes cleaned citations between papers with approximately 30% of all references within our post-1990 dataset linked to papers within the

ACM Digital Library. We obtained patent records through automated scraping of the U.S. Patent and Trademark Office patent database¹⁴, which includes abstracts, full texts, linked citations to other patents in the database (100% linked), and other references to academic papers (unlinked).¹⁵ The USPTO database is less usable than Google Patents for reading individual papers but it was easier for us to extract references to other patents and papers from it. Finally, we obtained trade journal and newspaper articles from the Lexis/Nexis Academic Database using an automated scraping tool. It provides a wide selection of news, political, legal, business, and reference information including trade journals, newspaper articles, and press releases.

Within the ACM Digital Library database dump we searched for papers where the title, abstract, or keywords matched one of the three innovations. We called each of the matching papers *flagged* papers. This search methodology limited the results to papers that we believed made contributions to one of the three tree browsing strategies. The search on title, abstract, and keywords was performed using these regular expressions:

```
\b(tree[- ]?map)[s]?\b
\b(cone[- ]?tree)[s]?\b
\b(hyperbolic[- ]?tree)[s]?\b
```

Each regular expression is composed of a starting word boundary "\b", the first part of the name, e.g. "tree", an optional hyphen or space "[-]?", the second part of the name, e.g. "map", an optional plural form "[s]?", and finally another word boundary. Using treemap as an example, this regular expression will match all 6 of the following variations: "treemap", "treemaps", "tree-map", "tree-maps", "tree map", "tree maps".

This accounted for all but one of the innovation name variations we discovered within the dump. The last, a lone instance of "hypertree" used for the visualization among a huge number of parallel processing usages, was added to the flagged papers manually. There were, how-

¹³ portal.acm.org (subscription required)

¹⁴ patft.uspto.gov

¹⁵ www.lexisnexis.com/us/lnacademic (subscription req.)

ever, a handful of critically important papers that did not match these restrictions. These included [18][19][28], the pivotal work on hyperbolic trees and cone trees. These three omissions were quickly discovered due to our prior knowledge of the field, though we present a technique later to help identify omitted potentially relevant publications. We added them to the flagged papers manually after their discovery and checked for additional omitted but relevant papers in the dataset.

Additionally, we added any publications that cited a flagged paper, as well as any publications cited by flagged papers as well. We call these papers *alters* to the flagged papers, or just *alters*. We added to our citation or edge dataset all citations (going either direction) between flagged papers, flagged papers and alters, and between just alters. The benefits of including the alters and the edges between them are threefold. First, it gives us a way to count the number of citations each paper receives, including binning citations into time periods. Second, it allows us to find commonly cited papers within an innovation, providing leads to potentially relevant publications that didn't match our search terms. Finally, it gives us a rough idea of the publication and citation trends in the field to compare against. Adding any papers to the dataset manually necessitates including all citations between it and all flagged papers and alters (in either direction), as well as adding the requisite new alters and all citations between them and all flagged papers and alters. This is required to ensure proper citation counts and to find missed papers that would have been highlighted by the new addition.

Patents were added similarly, though instead of searching a database dump for the matching keywords the automated scraping tool we developed downloaded only the flagged patents and their alters. For matching it used the same regular expressions as we used for academic papers, searching across the abstract and full text of the patents. This wonderfully clean dataset of patent-to-patent citations was easy to automatically extract as well, including all citations in either direction between flagged patents, flagged patents and alters, and between just alters.

The links between patents and papers were more difficult to extract as only the plain text references section was available. We wrote a simple title and year matching tool to compare the references to the cleaned academic papers dataset which was able to extract the vast majority of the citations from patents to papers. The reverse citations from papers to patents were more difficult as there is a huge range of formats for citing patents in the papers in our dataset¹⁶, along with OCR failures in the ACM references. However, none of the flagged papers cited any relevant patents, so these edges are excluded from the analysis.

Lexis/Nexis articles were automatically scraped by

¹⁶ Here is an example lightly tested regular expression to match most patent citations within the ACM DL dump (all one line):
`((United States|US|U\..S\..|U\.. S\..|U\..S))
 ?patent(</i>)?,? ?(US|No[\..]?|Number)?,? ?
 (US|#)?[\d, -\..]{4,9}[\d\b`

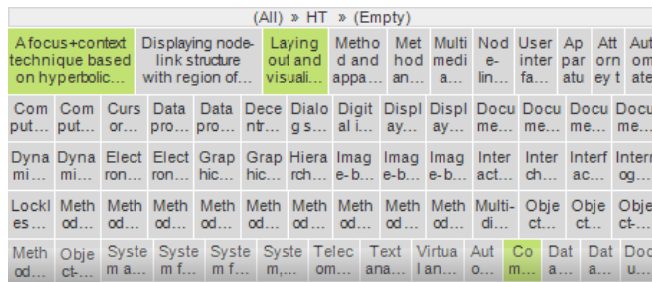


Fig. 1: Treemap visualization for finding highly cited hyperbolic tree papers (green) and patents (gray) by drilling down and filtering on an edge list that has source and target publication attributes available. Each cell is a target publication, and the size of each cell represents the number of source publications that cite it. The most highly cited (largest) float to the top left to be analyzed. Here the top three publications weren't properly flagged as HT papers and patents. We created a hierarchy based on the matching innovations or *flags* of the source and target publications: SourceFlags -> TargetFlags -> TargetNode. Then we drill down to the portion of the treemap with a particular SourceFlag like HT (for hyperbolic tree) and an *empty* Target-Flag. Those publications are cited by HT papers but are not flagged for any innovation.

querying keywords for different tree structures and exporting the results in bibliographic format. It was parsed by automatic scripts to make it suitable for visual analysis. We also manually read each article and collected the data about the major organizations and products discussed. There were no easily identifiable citations to these articles as even the URLs weren't available. Even with proper URLs, web citation networks are much harder to accurately extract and have much less metadata attached to them. As such, we didn't collect any citation data for them.

Our conservative approach in these searches provided us with a low estimate for the total impact of these innovations. More liberal studies have larger numbers, but for the purposes of this paper we wanted to identify direct contributors and the impact of their contributions. Moreover we can do a much more careful and accurate analysis uncluttered by the chaff. As an example, compared to the straightforward basic search provided by the ACM Digital Library, which includes references and reviews in the search, our technique matched 132 treemap papers instead of 357, many of which were not appropriate. Similarly, we found 19 cone tree and 20 hyperbolic tree papers compared to 75 and 100 found online. We were able to verify each rejected match from the online results manually, though many of them ended up being added as alters due to the references that the basic search matched against. The majority of the alters weren't particularly relevant to the innovation, only citing the work in a single line or paragraph. This was even more the case with the patents, which have a tendency to cite even tangentially related patents and papers in an effort to be comprehensive.

In order to ensure that our search terms weren't omitting any relevant publications we analyzed the resulting network of publications and citations between them to




Fig. 2: “Tire tracks” diagrams showing academic research, industrial research, and the emergence of \$1B commercial markets.

find any publications that were highly cited by publications with a particular flag but did not have that flag themselves. A treemap is a straightforward visualization for finding omitted publications by using a hierarchy based on the citations themselves. We created a hierarchy with the innovations or flags of the source paper at the first level, the flags of the target paper at the second, and finally the target papers themselves at the third and last level. From that we built the treemap visualization of papers and patents seen in Fig. 1 in Spotfire. It represents the potentially omitted papers as large rectangles floated to the top left of the screen, sized by the number of citations matching the source and target flags specified (in this case, a source matching *HT* for hyperbolic trees and a target that doesn’t match any of our three search terms (*empty*). Here the papers are colored light green and the patents gray, and the publications with low citation numbers matching the hierarchy have been cut off at the bottom.

Here the top three publications weren’t properly flagged as HT papers and patents because they didn’t mention “hyperbolic tree” anywhere within their title, abstract, or keywords (for papers) or their full text (for patents). As it turns out, these three are critical hyperbolic tree publications that would have substantially impacted the validity of our results, though we had noted their absence before due to our prior knowledge. We easily found a critical omitted cone tree paper and patent through the same technique.

For more details on the complexities of our data cleaning and analysis, see the follow-up paper.

4 VISUALIZING INNOVATION TRAJECTORIES

While there are many tools for showing timelines of numerical data, categorical event data, and historical time lines, there is relatively little work on showing the multiple complex longitudinal relationships among different types of data such as academic papers, patents, and trade press articles. Even more complex is showing the historical pattern of linkages among commercial suppliers, products, press releases, and adopting organizations.

Visualizations for historical academic citation patterns is usually traced to Garfield’s arguments for “historiographs” to show the impact of a single paper [14][15]. A more general tool, DIVA shows multiple horizontal time lines for each topic with nodes showing papers and links to show citations within and across the time lines [22][23]. The impact of key authors at four human-computer interaction conferences is thoroughly examined and visualized with several forms of networks but less emphasis on temporal relationships [16]. Citation patterns across sub-

topics are shown in CiteSpace, which recently added the capacity to locate nodes on timelines [9]. The progress of scientific arguments is visualized to show opposing opinions and supportive research studies [36]. Attribute based network visualization which shows temporal patterns for academic papers as well as links across topics provides an inspiration for our current research [3].

The business literature discusses the process of innovation adoption, especially as measured by trade press articles [39][40] and income from sales [38]. Wang and Swanson [38] show a 5-year sales histogram, plus growth rate, and textual annotations of significant events. The famed Gartner’s Hype Cycle is a chart with five stages in the evolving expectations for and visibility of an innovation: Technology Trigger, Peak of Inflated Expectations, Trough of Disillusionment, Slope of Enlightenment, and Plateau of Productivity [13]. However, critics complain that this widely discussed chart is not really a cycle, does not reflect reality, and fails to provide a basis for action.

A compact visualization that shows relationships among academic research, industrial research, and the emergence of billion dollar industries was described as “tire tracks diagrams” because of their visual similarity to the tracks a wet tire might make on pavement [11][20]. These diagrams are appealingly simple and suggest the strong impact of academic research on commercial success, but they are based on the knowledge of and discussions among selected researchers (Fig. 2). Attempts to replicate them from public data have not proven successful. Historical time lines of many kinds are a popular method, but there are few strategies that can take data from multiple and diverse public sources and automatically create complex timelines that show innovation trajectories.

4.1 Publications

Perhaps the most straightforward way of visualizing the adoption or impact of each innovation is by using line charts to show the number of publications for each. Fig 3., prepared with Spotfire, shows separate line charts for the trade press *articles* (top), academic *papers* (middle), and *patents* (bottom). Each line shows the trajectory for one innovation, with the publication year on the x-axis and the number of publications that year on the y-axis. Green represents treemaps (*TM*), blue hyperbolic trees (*HT*), and red cone trees (*CT*). A few key features are immediately apparent. There is a large spike of 17 trade press articles published for hyperbolic trees in 2000. Drilling down into the data it became apparent that Inxight was heavily promoting their hyperbolic tree products with press releases, which accounted for almost all the publications. It’s possible their promotion helped swell hyperbolic trees in academic papers slightly in the subsequent years.

For both trade press articles and academic papers treemaps became steadily more mainstream after 2002 and were promoted heavily by the Hive Group, Panopticon, SequoiaView, and others. This increase was matched by two years of somewhat increased treemap patents, predominantly assigned to the Hive Group, IBM, and Microsoft; though the patents tapered off while trade ar-

ticles and academic papers continued to increase. These patent figures seem to fall off sharply during 2006-2008 (shaded gray). This change was traced to the time it took the USPTO to process patents increased dramatically to an average of 32 months during 2005-2008 (the publication year we use for patents is the filing date)¹⁷. Our dataset of matching patents and their alters contains only 1 patent from 2009, compared with 200, 152, 66, and 19 in 2005, 2006, 2007, and 2008, respectively. We omitted all 2009 data from the charts due to incompleteness from all our data sources.

Treemaps were a little late to the patent party, with cone trees and hyperbolic trees already heavily patented from 1997 to 2003. Cone trees were patented by 12 unique assignees led by Xerox with 18, 41% of the total, followed by Sun with 7 and Microsoft with 5. Hyperbolic trees were less concentrated by assignee with 20 unique assignees, though again Xerox led with almost a quarter of the total (9). The other top contenders were IBM (4), Rosetta Inpharmatics (3), and HP (3). Treemaps were concentrated similarly with 19 unique assignees. IBM was the leader in treemaps with 12 (31% of the total), with runner ups Microsoft and the Hive Group tied with 4.

4.2 Citations

Instead of simply looking at the number of publications each year for an innovation we can analyze the number of citations the related publications receive over time. Fig. 4 provides an example, where the y-axis represents how many times matching publications have been cited per year instead of showing the number of publications per year. Unfortunately there are no easily measurable citation data for trade journal articles so they are omitted in further analyses. Separate line charts show academic papers (top) and patents (bottom), with the publication year on the x-axis and line color green for treemaps (TM), blue for hyperbolic trees (HT), and red for cone trees (CT). The y-axis is the count of all citations from both academic papers and patents published that year (src) where the cited publication matches the innovation. We include patent-to-paper citations in the citation count for academic papers, though no paper-to-patent citations existed for term-matching patents in our dataset.

This visualization paints a somewhat different picture for cone trees and hyperbolic trees, where instead of shorter spurts of academic papers and patents published there is a more prolonged and sizable measured impact. Even though there were only a few cone tree academic papers published each year, they were cited 44 times in 1998 and 55 in 2002. Citations for cone tree patents peaked in 2000 (194) and for hyperbolic trees later in 2005 (144), also a local maxima for cone trees (153). Treemap patents were never comparatively highly cited but peaked 2003-2004 with 45. All have been following a steady descent in the period of 2006-2008 (shaded gray), though this matches closely the reduction in patents in our dataset during that time period, possibly due to the backlog at the USPTO.




Fig. 3: These line charts show the number of trade press *articles* (top), academic *papers* (middle), and *patents* (bottom) published each year matching each innovation. Each line represents a different innovation: treemaps (TM/green), hyperbolic trees (HT/blue), and cone trees (CT/red). Note: the sharp fall in patent figures in the faded area may be to average 32-month USPTO processing time in 2005-2008.

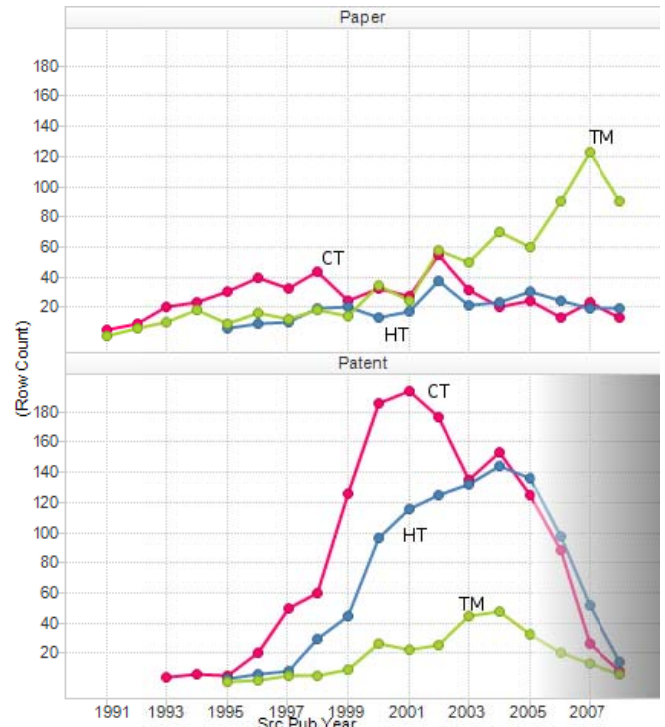


Fig. 4: These line charts show the number of citations academic *papers* (top) and *patents* (bottom) matching each innovation received each year. Each line represents a different innovation: treemaps (TM/green), hyperbolic trees (HT/blue), and cone trees (CT/red). Note: the sharp fall in patent figures in the faded area may be to average 32-month USPTO processing time in 2005-2008.

¹⁷ www.ipwatchdog.com/2009/04/22/uspto-backlog-patent-pendency-out-of-control/id=2848

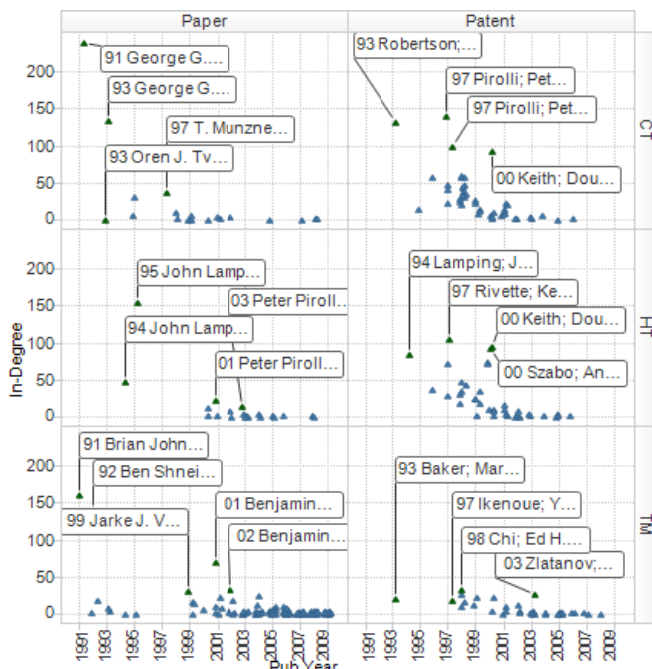


Fig. 5: The total number of citations each academic paper and patent receives is plotted here, providing another measure of its influence. The number of citations, or in-degree, a publication receives is plotted on the vertical axis and the publication year is plotted on the horizontal. The rows divide the innovations: cone trees (CT/top), hyperbolic trees (HT/middle), and treemaps (TM/bottom) and the columns publication type: academic papers (left) and patents (right).

Another approach to measure adoption or impact is to dig down to the level of individual papers and patents to see their influence. The scatterplots in Fig. 5 show a triangle for each publication with the y-axis representing the total number of citations it received from both academic papers and patents and the horizontal publication year. The left column shows academic papers and the right patents, with a row for each innovation: cone trees (CT/top), hyperbolic trees (HT/middle), and treemaps (TM/bottom). Here the key publications for each innovation are apparent, with labels where there is space. The 1991 ACM CHI paper on cone trees [29] has been cited an amazing 242 times, with the 1993 CACM paper [28] cited a respectable 132. The first patent in 1993 by Robertson et al., “Display of hierarchical three-dimensional structures with rotating substructures”, was cited 130 times and the two most cited followups came in 1997: Pirolli et al., “System for categorizing documents in a linked collection of documents” (138) and Pirolli et al., “System for predicting documents relevant to focus documents by spreading activation through network representations of a linked collection of documents” (103). Those only touch on cone trees as an application of their patent, though.

For hyperbolic trees, the 2-page 1994 UIST description [18] garnered 47 citations, though dwarfed by the 156 for the 1995 ACM CHI paper [19]. The enhanced version by the University of Waterloo [8] received only 30 citations. There is an interesting mix of highly-cited patents during 1994-2000. There were not any highly cited treemap pa-

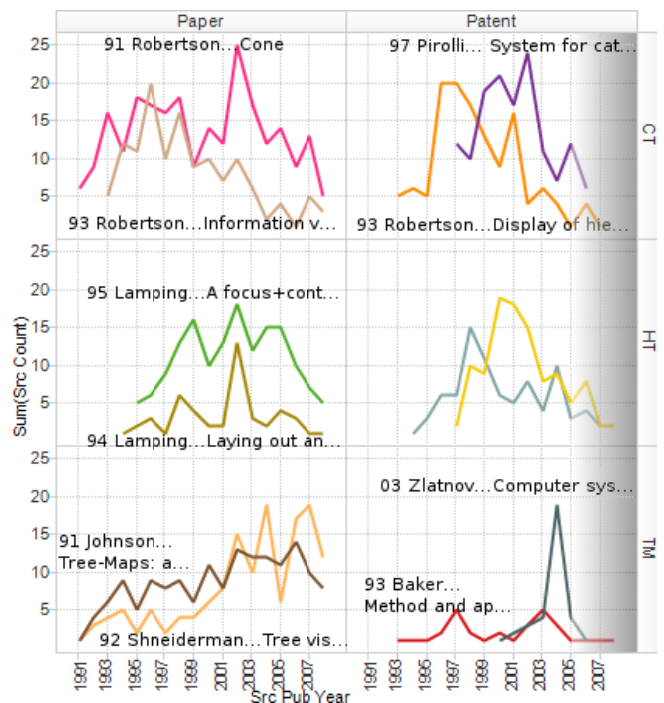


Fig. 6: The citation trajectory of the two top cited academic papers and the two patents with the most consistent or interesting citation patterns is shown here. Each colored line represents a publication and indicates how many citations it received each year. This provides another view of their long-term impact. The rows divide the innovations: cone trees (CT/top), hyperbolic trees (HT/middle), and treemaps (TM/bottom) and the columns publication type: academic papers (left) and patents (right). Note: the sharp fall in patent figures in the faded area may be to average 32-month USPTO processing time in 2005-2008.

tents, though there was an early 1993 patent by Baker with 27 and the original papers did quite well. The 1991 IEEE Vis paper [16] was cited 158 times and the 1992 ACM TOG paper [28] got 143. A notable departure from the other innovations is the highly cited 2001 PhotoMesa paper and two new design papers in 1999 and 2002 that may have helped usher in a renewed interest in treemaps.

For an even finer-grained look at the citation history of individual publications we can plot the number of citations each receives per year as a line chart (Fig. 6). There each line represents an individual publication, and its height at each year indicates how many citations it received that year. Again, columns are by publication type (academic papers on the left and patents on the right) with innovations separated by row: cone trees (CT/top), hyperbolic trees (HT/middle), and treemaps (TM/bottom). The cone tree papers and the hyperbolic tree papers stand out at the top with quick uptake and relatively horizontal trajectories, while the two treemap papers start from humbler beginnings but consistently climb, ending between hyperbolic trees and cone trees for total citations. Both the cone tree and hyperbolic tree patents are highly cited over time, while the main treemap patent was consistently but not frequently cited. The 2003 treemap patent by Zlatanov and Furlong is an aberration, highly cited in 2004 but back to previous norms in 2005. It will be interesting to see its trajectory

over the next decade.

5 DISCUSSION

Our sample of these three information visualizations for tree structures is a narrow one, so generalizations are difficult, but some topics invite discussion.

Some new ideas catch on rapidly and spread widely, but most fail, while another group takes decades of steady refinement and tuning to fit appropriate application domains. The entrenched command line interface slowly gave way to the graphical user interface, which has opened the doors for many users. Similarly, the entrenched textual, numeric, and tabular interfaces for data are giving way slowly to new information visualization ideas. Change is possible, although new application domains, platforms, tasks, and users may offer fertile possibilities, much as Moore [21] and Ogle [25] suggest. Rao comments on these issues in his Information Flow blog in 2003, but even his prediction for 2007 has not yet been realized¹⁸.

These three cases suggest that an important promoter of adoption is having a readily understandable demonstration using a familiar application domain, while carrying out commonly used tasks. Treemaps started with directory browsing, but became a broader success when other appropriate applications appeared, such as Smartmoney's MarketMap, supply chain management, insurance fraud, stock portfolio analysis, and production management. Cone trees began with directory browsing and spread to web browsing, but few other applications. Hyperbolic trees were used for information-rich applications such as organization charts, web browsing, and document libraries.

Both sides of the long-standing controversial issue of 2-dimensional vs. 3-dimensional information visualization can make points based on tree browsing. While several 3-dimensional treemaps were implemented, they quickly faded and empirical evidence showed few benefits over 2-dimensional versions. The 3-dimensional cone tree generated strong interest and may have inspired many developers, but the lack of commercial diffusion raises questions about the complexity of using 3-dimensional representations for these tasks.

Enthusiasts for user-controlled animated visualizations included the designers of cone trees and hyperbolic trees, which were sufficiently appealing to generate strong interest but the modest commercial diffusion raises questions of efficacy. Does the frequent change of position of nodes undermine recall and discovery? By contrast the treemaps had limited animation, mostly related to zooming in on a subtree, but until recently, this was usually done with a simple jump zoom rather than a smooth animation. Researchers on treemaps focused on spatial stability, especially as node sizes changed, to prevent the disturbing movement and reshaping of rectangles. Later refinements to treemaps addressed these issues but have

yet to appear in commercial versions.

While Shneiderman has written and spoken in favor of intellectual property protection, these cases add support to those who argue that patents inhibit adoption and diffusion. The unpatented treemaps led to many implementations with refinements and a lively competition among commercial and open source developers. By contrast, the patented cone trees and hyperbolic trees limited the number of derivative implementations. Ironically, the number of treemap refinement patents approximately matches the number of patents for cone trees and hyperbolic trees. Card points out the further irony that the academically-oriented University of Maryland produced a more broadly adopted and commercialized technology, while the industrially-oriented Xerox PARC produced strong and more broadly cited papers (personal communication 3/30/2010).

4 CONCLUSION

Understanding the innovation trajectory for novel research ideas from invention to successful innovation is a difficult task. Even when there are evaluations of efficacy and precise records of early adoption it seems difficult to predict durability. Academic citation counts are not always tied to patent counts or trade press articles, and certainly many factors govern commercial success. A few of the determinants of success are the entrenched alternatives, resistance to change, interactions with existing technologies, intellectual property rights, perceived usefulness and ease-of-use, availability of free versions, influence of entrepreneurial individuals, and responsiveness of trade press editors. This complex mix of personalities, ideas, institutions, and economic constraints makes it difficult to predict outcomes even 2-3 years in the future.

On the more positive side, readers might take away an appreciation of how many components there are to commercial success and the sobering realization of how much effort is required in many cases. Early evaluations of efficacy may help lead to refinements, as well as a clearer understanding of which tasks and application domains are a good match for an information visualization invention. Presentations at conferences, through videos, and at academic seminars can provide further feedback and a chance to win over visionaries and early adopters. The shift to a commercial effort is still more demanding, probably requiring enthusiastic champions (as Ramana Rao certainly was for Inxight's hyperbolic trees), a devoted team of developers, and a determined sales force. Small or large companies can help advance an invention to a broadly used innovation.

There are also many dimensions to success. Academic success is typically measured by citation counts which are presumed to be an indicator of influence on other researchers. Commercial success could be measured in dollars, number of customers, or numbers of users. For inventors, there is great satisfaction and pride in seeing their ideas influencing others and being put to use for important applications. The information visualization community faces substantial challenges in bringing their

¹⁸ www.ramanarao.com/informationflow/archive/2003-02.html

ideas to broad adoption, but each success creates a greater visual literacy, which sharpens users' critical thinking while making them still more sympathetic to the next innovation.

ACKNOWLEDGMENT

The authors wish to thank Jim Bartoo, Stu Card, Jock Mackinlay, Ramana Rao, and George Robertson for their comments on the history, Catherine Plaisant for comments on drafts, and the ACM for supplying us with a dump of their digital library. This work was partially supported by National Science Foundation grant SBE-0915645.

REFERENCES

- [1] Allen, T. J., *Managing the Flow of Information Technology*, MIT Press, Cambridge, MA (1977).
- [2] Allen, M. M., The hype over hyperbolic browsers, *ONLINE* 26, 3 (May/June 1992), Available at <http://www.onlinemag.net/may02/allen.htm>.
- [3] Aris, A., Shneiderman, B., Qazvinian, V., and Radev, D., Visual overviews for discovering key papers and influences across research fronts, *Journal of the American Society for Information Systems and Technology* 60, 11 (November 2009), 2219-2228.
- [4] Arthur, B., *The Nature of Technology: What It Is and How It Evolves*, Free Press, New York, (2009).
- [5] Basalla, G., *The Evolution of Technology*, Cambridge University Press, New York (1988).
- [6] Bederson, B. B., Shneiderman, B., and Wattenberg, M., Ordered and quantum treemaps: Making effective use of 2D space to display hierarchies, *ACM Transactions on Graphics (TOG)* 21, 4 (October 2002), 833-854.
- [7] Bruls, M., Huizing, K., & van Wijk, J. J., Squarified treemaps. *Proc. Joint Eurographics and IEEE TCVG Symposium on Visualization (TCVG 2000)* IEEE Press, Piscataway, NJ (2000), 33-42.
- [8] Carriere, J. and Kazman, R., Research report: Interacting with huge hierarchies: beyond cone trees, *IEEE Symposium on Information Visualization (InfoVis '95)*, IEEE Press (1995), 74-81.
- [9] Chen, C., CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. *Journal of the American Society for Information Science and Technology*, 57, 3 (2006), 359-377. ***ADD*** Chen, C., Ibekwe-SanJuan, F., & Hou, J., The structure and dynamics of co-citation clusters: A multiple-perspective co-citation analysis. *Journal of the American Society for Information Science and Technology*, 61, 7 (2010), 1386-1409.
- [10] Cockburn, A., Karlson, A., and Bederson, B. B., A review of overview+detail, zooming, and focus+context interfaces, *ACM Computing Surveys* 41, 1 (December 2008), Article No. 2.
- [11] Computer Science and Telecommunications Board, *Innovation in Information Technology*. National Academies Press, Washington, DC (2003). http://www7.nationalacademies.org/cstb/pub_itinnovation.html
- [12] Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13, 3 (1989), 319-340.
- [13] Fenn, J. and Raskino, M., *Mastering the Hype Cycle: How to Choose the Right Innovation at the Right Time*, Harvard Business Press, Boston, MA (2008).
- [14] Garfield, E., Historiographic Mapping of Knowledge Domains Literature, *Journal of Information Science* 30, 2 (2004), 119-145.
- [15] Garfield, E., Pudovkin, A. I., and Istomin, V. I., Why do we need algorithmic historiography? *Journal of the American Society for Information Science and Technology (JASIST)* 54, 5 (2003), 400-412.
- [16] Henry, N., Goodell, H., Elmqvist, N., and Fekete, J.-D. 20 Years of Four HCI Conferences: A Visual Exploration, *Int'l Journal of Human-Computer Interaction*, 23, 3(2007), 239-285, Available at: <http://dx.doi.org/10.1080/10447310701702402>
- [17] Johnson, B. and Shneiderman, B., Tree-maps: A space filling approach to the visualization of hierarchical information structures, *Proc. IEEE Visualization '91*, IEEE Press, Piscataway, NJ (October 1991), 284-291.
- [18] Lamping, J. and Rao, R., Laying out and visualizing large trees using a hyperbolic space, *Proc. 7th Annual ACM symposium on User Interface Software and Technology*, ACM Press, New York (1994), 13-14.
- [19] Lamping, J., Rao, R., and Pirolli, P., A focus+context technique based on hyperbolic geometry for visualizing large hierarchies, *Proc. SIGCHI Conference on Human Factors in Computing Systems*, ACM Press, New York (1995), 401-408.
- [20] Lazowska, E., Envisioning the future of computing research, *Communications of the ACM* 51, 8 (August 2008), 28-30.
- [21] Moore, G. A., *Crossing the Chasm*, HarperCollins, New York, 1991.
- [22] Morris, S. and Martens, B., Mapping Research Specialties, *Annual Review of Information Science* 42 (2007), Chapter 6.
- [23] Morris, S. A., Yen, G., Zheng, W., and Asnake, B. (2003). Time Line Visualization of Research Fronts. *Journal of the American Society for Information Science and Technology (JASIST)* 54, 5 (2003), 413-422.
- [24] Mullet, K., Fry, C., and Schiano, D., On your marks, get set, browse! (The Great CHI'97 Browse Off), *ACM CHI'97 Conference Extended Abstracts*, ACM Press, New York (1997), 113-114. Available at: <http://www.sigchi.org/chi97/proceedings/panel/kem.htm>
- [25] Ogle, R., *Smart World: Breakthrough creativity and the New Science of Ideas*, Harvard Business School Press, Boston, MA (2007).
- [26] Pirolli, P., Card, S., and Van Der Wege, M., The effects of information scent on visual search in the hyperbolic tree browser. *ACM Trans. Computer-Human Interaction* 10, 1 (2003), 20-53.
- [27] Plaisant, C., Grosjean, J., and Bederson, B., SpaceTree: Supporting exploration in large node link tree, design evolution and empirical evaluation. *Proc. IEEE Symposium on Information Visualization (InfoVis)*, IEEE Press, Piscataway, NJ (2002), 57-64.
- [28] Robertson, G. G., Card, S. K., and Mackinlay, J. D., Information visualization using 3D interactive animation, *Communications of the ACM*, 36, 4 (1993), 51-71.
- [29] Robertson, G. G., Mackinlay, J. D., and Card, S. K., Cone trees: Animated 3D visualizations of hierarchical information, *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems*, ACM Press, New York, (April 1991), 189-194.
- [30] Rogers, E. M., *Diffusion of innovations* (5th ed.). New York: Free Press, New York (2003).
- [31] Shneiderman, B., Tree visualization with tree-maps: A 2-dimensional space filling approach, *ACM Transactions on Graphics* 11, 1 (January 1992), 92-99.
- [32] Shneiderman, B. and Plaisant, C., Strategies for evaluating information visualization tools: Multi-dimensional In-depth Long-term Case Studies, In *Proc. Beyond time and errors: novel evaluation methods for Information Visualization*, Workshop of the Advanced Visual Interfaces Conference, Available in ACM Digital Library (2006), 1-7.
- [33] Shneiderman, B. and Plaisant, C., *Designing the User Interface: Strategies for Effective Human-Computer Interaction: Fifth Edition*, Addison-Wesley Publ. Co., Reading, MA (2010).
- [34] Stasko J, Catrambone R, Guzdial M, and McDonald, K., An evaluation of space-filling information visualizations for depicting hierarchical structures, *Int'l J. of Human-Computer Studies* 53, 5 (Nov. 2000), 663-694.
- [35] Turo, D. and Johnson, B., Improving the visualization of hierarchies with treemaps: design issues and experimentation, *Proc. IEEE Conference on Visualization 1992*, IEEE Press, Piscataway, NJ (1992), 124-131.
- [36] Uren, V., Buckingham Shum, S., Bachler, M., and Li, G. (2006). Sensemaking tools for understanding research literatures: Design, implementation and user evaluation, *Int'l Journal of Human-Computer Studies* 64, 5 (2006), 420-445.

- [37] Venkatesh, V., Morris, M. G., Davis, G. B. and Davis, F. D., User acceptance of information technology: Toward a unified view, *MIS Quarterly* 27 (2003), 425-478.
- [38] Wang, P. and Swanson, E. B., Launching professional services automation: Institutional entrepreneurship for information technology innovations, *Information and Organization*, 17, 2 (2007), 59-88.
- [39] Wang, P. and Ramiller, N.C., Community learning in information technology innovations, *MIS Quarterly*, 33, 4 (2009), 709-734.
- [40] Wang, P., Popular concepts beyond organizations: Exploring new dimensions of information technology innovations, *Journal of the Association for Information Systems*, 10, 1 (2009), 1-30.
- [41] Wattenberg, M., Visualizing the Stock Market. Proc. Extended Abstracts of Human Factors in Computing Systems (CHI 99) ACM Press, New York (1999), 188-189.
- [42] Xiang, Y., Chau, M., Atabakhish, H., and Chen, H., Visualizing criminal relationships: Comparison of a hyperbolic tree and a hierarchical list, *Decision Support Systems* 41, 1 (Nov. 2005), 69-83.