MONADIC VISUALIZATION OF METADATA NETWORKS TO

SUPPORT EXPLORATORY BROWSING

An Undergraduate Research Scholars Thesis

by

KADE KEITH

Submitted to Honors and Undergraduate Research
Texas A&M University
in partial fulfillment of the requirements for the designation as an

UNDERGRADUATE RESEARCH SCHOLAR

Approved by
Research Advisor: Dr. Andruid Kerne

May 2015

Major: Computer Science

TABLE OF CONTENTS

ABSTRACT

Monadic Visualization of Metadata Networks to Support Exploratory Browsing. (May 2015)

Kade Keith
Department of Computer Science
Texas A&M University

Research Advisor: Dr. Andruid Kerne
Department of Computer Science

Interactive data visualizations tend to either show an entire collection or the individual element. For example, a page of search results shows the whole, while clicking on an individual link shifts focus to the individual. With these kinds of interfaces, it can be difficult to understand how individual pieces relates to each other and maintain a sense of place amongst the network being perused. To address this, I develop a focus+context browsing tool that uses a monadic approach to visualizing networks of semantic information in the form of metadata summaries. Focus+context visualizations present both the full detail of an object and the contextual information needed to relate that object to the entirety. The monadic approach, first suggested by sociologist Gabriel Tarde, is the idea that a network is best comprehended from the perspective of individual elements. This stands in opposition to the standard approach of viewing the entirety of the network at once. Semantic information, such as product specs and bibliography, provides essential and valuable contexts for people to make sense of information, assess values, and make decisions. I present a study, in which participants explored networks of academic publications, showed that a monadic metadata visualization helps users explore networked information, understand relationships, and maintain focus, when compared with traditional methods.

Allowing the user to traverse the network laterally while using metadata to provide a summary of the focus node serves to maximize the amount of focus and context available to the user.

ACKNOWLEDGMENTS

CHAPTER I

INTRODUCTION

What is the best way of not just viewing, but making sense of a network? The standard approach is to look at things from two levels, that of the individual and that of network, micro and macro. This approach is dominant, but there are alternatives. Sociologist Gabriel Tarde asserted the value of a *monadic* view, which focused on neither the individual nor the network [15]. It instead views the network from the perspective of one of its elements. This means the network is understood differently from different perspectives. For example, if you are looking at a recent research paper, its references and citations are the most relevant part of the network of all papers in the field. Over a hundred years later, this approach was examined in the digital space and it was found that, thanks to the growing availability of vast digital databases and advances in visualization technology, navigating networks via monads is an increasingly viable and valuable way to explore information [12]. Building upon this, Marian Dörk built a monadic exploration of the book, Beautiful Trouble [6]. In his study, the monadic approach to exploration excited and engaged users as they waded through articles.

While valuable for exploring connected information and understanding relationships, Dörk's monadic approach falls short in that it required a separate "article view" to show content or details in addition to the monadic view. A detailed view is necessary for users, for example, to comprehend individual scholarly articles after exploring citation chains. This is typically done by viewing the source article page in a separate context, such as a new browser tab or window, which results in disorientation, or being uncertain in users' perception of location in a network of

information, a common problem with hyperlink browsing [16]. Prior work on exploratory browsing interfaces, such as the Metadata In Context Expander (MICE) [14], reduces disorientation by presenting metadata summaries for individual objects in one context and making relationships visible. The effectiveness of metadata as document summaries is also demonstrated in prior work, such as Web Summaries [5].

This research integrates monadic exploration with metadata summaries. The integrated view shows a metadata summary of the object under focus within the exploration context as well as contextual linked information, the monad. This view reduces the frequency with which users need to change from one context to another and thus the disorientation that results from switching contexts. It enables users to comprehend individual objects and maintain focus while they explore the network and understand relationships. This research develops the emerging monadic exploration method towards a practical paradigm for exploring many kinds of interconnected information, including scholarly publications and social networks. This also significantly improves the usability of interlinked metadata summaries, and marks a new milestone for exploratory browsing.

CHAPTER II

METHODS

Design of a Monadic Interface

I set out to design a monadic interface with the following goals in mind:

- *Fluid navigation.* Users should be able to smoothly switch between elements, which can be thought of as graph nodes, and clearly see those elements' relationships.

- *Emergent connections.* By visualizing secondary connections as well as primary ones, the interface should allow users to see previously unknown connections. Recognizing and reacting to these connections is an inherently creative experience [3].

- *Sufficient summary.* The representation of the in-focus node should provide a sufficient summary of the web page it represents.

- *Backtracking.* Users need to see what nodes they visited to arrive at their current monad, and should be allowed to revisit them, as prior works identifies revisiting elements, re-finding, as an important aspect of browsing [13]. The interface features bookmarking to support backtracking as well.

To create a monadic interface for exploring networks, I divide the available visual space into three portions: the history visualization, the focus node, and the collection graph
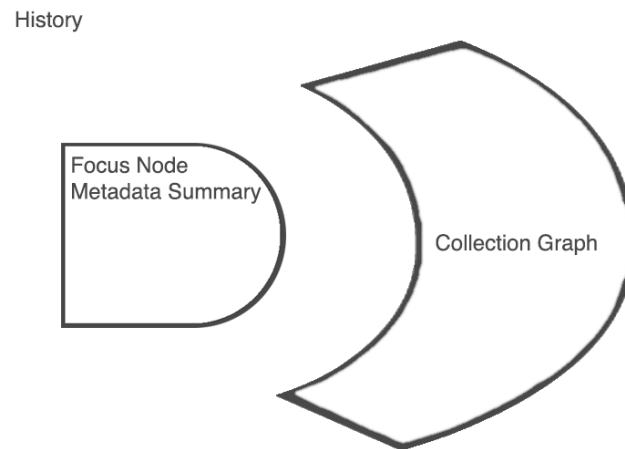
History

Figure 1. The layout of my monadic interface

History Visualization

The history visualization serves as a breadcrumb trail for users during exploration. It shows previously visited nodes in the order they were visited. Clicking on a node in the history will re-focus on that node. Previous monadic interfaces have lacked such functionality. There is a back button, which takes users to the node they were last focused on. I also allow users to bookmark nodes that they find significant. These are displayed next to the history and can be navigated to. The page titles of bookmarked nodes can be exported.

Focus Node

The most prominent element of the interface is a metadata summary of the focus node. I use a variant of the Metadata In-Context Expander (MICE) to display this summary. MICE relies on the BigSemantics architecture, which uses the Meta-Metadata [11] language to extracts salient

7

metadata from web pages. That metadata is then presented in the form of an expandable summary. I explicitly designed with the exploration of Association of Computing Machinery digital library [1] articles in mind, but the BigSemantics architecture supports over 200 web pages, and thus the interface could easily be extended to support a number of domains.

Figure 2. Shown is a partial metadata summary of the ACM page on a paper. Hovering over the abstract section expands it to show the full abstract.

Collection Graph

I call the most relevant portion of the network based on the focus node the collection graph. It contains direct (nodes linked directly to the focus node), and indirect (nodes linked by a direct node) connections. I call the set of nodes linked to by a particular node its relatives. The total number of nodes displayed is bounded. I choose which nodes to prioritize the visualization of by ordering them by number of relatives on the current graph. For example, in the domain of scholarly articles, if every direct connection has a reference in common, showing that reference will be prioritized over a showing a paper only referenced by only half of them. Nodes are rendered on a timer, as metadata for each node is loaded asynchronously. Unrendered nodes are kept in a max heap structure, and roughly every second the unrendered node with the most connections is rendered until the bound is reached. To render these nodes, I use a force-directed layout, which represents nodes and edges as through they were masses and springs [7].

Connected nodes pull towards each other as a function of the distance between them and how many relatives they have in common. Nodes that are not connected push each other away. This allows the perceived physical structure of the graph to convey information about the data in an aesthetically pleasing way.

I employ a number of visual cues to convey information about each element in the graph. Color denotes the relationship a node has to its parent. The icon of a node corresponds to that node's type. Icon size and text opacity differentiate local and global significance. The size of the icon is determined by how many relatives that node has on the current graph for example. Opacity is a function of global significance, which I base on overall citation count. So the papers and authors with the most overall citations appear bolder in the graph.
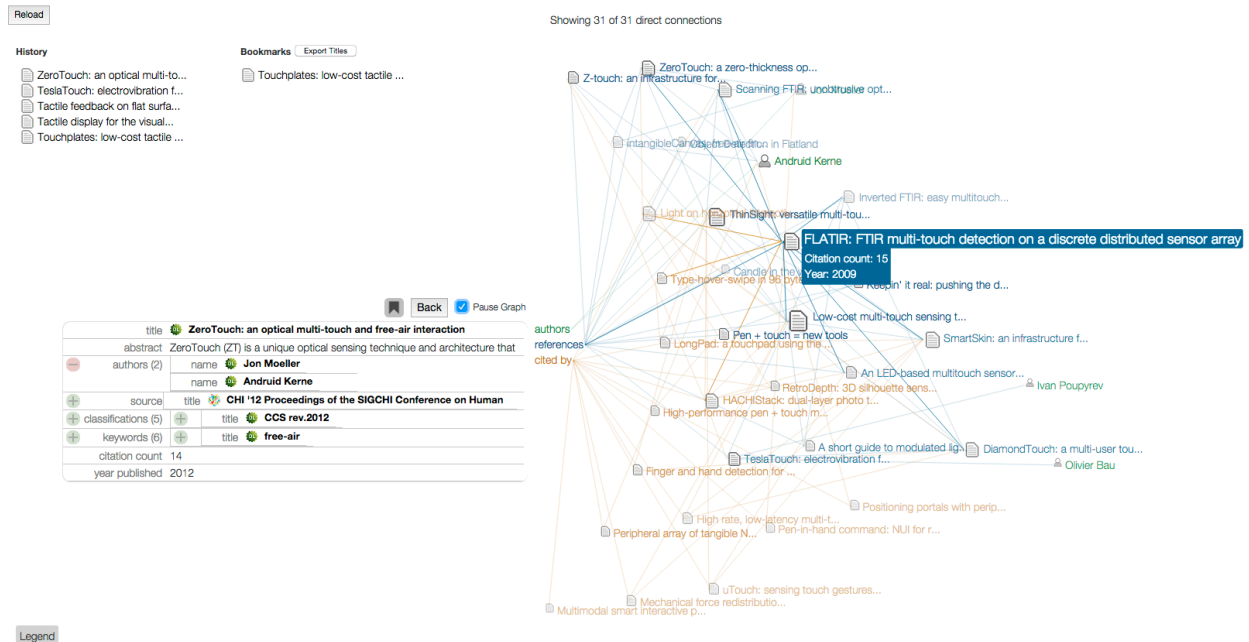

Figure 3. A screenshot of the application in use. In it, a user has focused on a scholarly article about ZeroTouch, for which the authors (green), citations (blue), and references (yellow) have been visualized as graph nodes. The user has hovered over one of those references.

Hovering over a node will display that nodes full title and increase the visual prominence of that node and it's connections. If that node is a paper whose metadata we have loaded then the citation count and year will be shown below the title. Clicking on a node refocus the graph around that node. Hovering over one of the node categories, which are listed as a part of the focus node, will emphasize the nodes in that category. Clicking on a category removes/adds the nodes of that category from the graph.

Evaluation

Laboratory Study

To evaluate the interface, I conducted an iterative user study of 5 undergraduate and graduate student researchers. Observations from the study informed both the system's functionality and the evaluation methods. Each participant came in for approximately an hour and a half. After being given the non-disclosure agreement, they were given a pre-questionnaire that asked about their research habits, how much experience they had reading scholarly articles, what tools they currently used, and what they would like to see in software for exploring citation networks. The participants were then shown how to use the software. Participants were asked to use the software for approximately 30 minutes. Some used it for upwards of an hour. Their task was to explore scholarly articles and authors relevant to their own research interests. If they found a paper or author of particular relevance they could bookmark it. After the study the list of bookmarks was emailed to them. Each participant provided a paper of his or her choice as a starting point for exploration. After doing so, students were given a post-questionnaire and a post interview was conducted, in which they were asked how using the monadic interface impacted

the way they explored citation chains. Questions about specific aspects of the interface were asked as well.

Metrics

As students used the software, their interactions with it were logged. I focus in particular on the number and nature of connections explored and the use of the history visualization and bookmarking functionality to backtrack to previously visited nodes.

To evaluate the effectiveness of my visualization of the collection graph, I focus on the ideation metric of Fluency [10]. Fluency, which is determined by the number of ideas considered, has been identified by prior research as a viable metric for judging creative processes, since the more ideas a person considers, the greater the likelihood that one of them will be creative. Since the task given involves the exploration of scholarly articles so that participants can formulate their own research interests and ideas, I apply this metric to their task. The more scholarly articles a participant is exposed to, the higher their Fluency. By visualizing indirect connections and exploration history in addition to direct connections, I theorize that my interface offers higher Fluency to users. If a user chooses to switch focus from one paper to a new paper that is a reference of a reference, an action not allowed by some interfaces, then they are taking advantage of the higher Fluency afforded by this interface.

To evaluate the sufficiency of the focus node, I measure how frequently users open up a view of the node's original web page in a separate tab. If this happens infrequently, then I can assume that the metadata summary I provide is sufficient for the task at hand. I also directly ask

participants in the post interview if they thought the metadata summary of the focus node was sufficient.

To evaluate the usefulness of the history visualization, I measure how frequently participants used it to navigate to previously visited nodes. The post questionnaire also asks participants about their experience with that part of the software.

Scenario

To better illustrate how the interface is used, I present a hypothetical use case that is similar to how study participants used it. A Ph.D. student has just finished reading a recently published article sent to her by her advisor. She found the paper quite interesting and thinks that some of the techniques mentioned in it could be relevant to her own research. She wants to learn more about the research this paper was based upon and decides to use my monadic metadata explorer to do so. With the paper she just read as her starting point, she views the portion of the citation network shown to her by the interface. As more nodes are rendered on the graph, one paper becomes particularly prominent visually since it, in addition to being referenced by the original paper, is referenced by 7 of the 14 other references. She hovers over that paper and sees that it has relatively low global citation count of only 15. Nonetheless she clicks on it, and it becomes the focus node. After reading the abstract provided by the metadata summary, the researcher decides that this paper is worth reading in full later, and bookmarks it. She then continues to explore the graph, bookmarking papers of interest and using the history visualization to navigate back to her starting point periodically. After half an hour of doing this, she has assembled a reading list of 6 papers that she believes to be the most promising for her own research so she outputs their titles with the intention of later reading them in full.

CHAPTER III

RESULTS

A Survey of Research Habits

Prior to using the software, participants were surveyed about their research habits in the form of a questionnaire. From their responses I observed the following:

- *Current Software.* The websites Google Scholar [8] and ACM digital library were listed as the primary places participants went to explore citation chains. Mentioned less frequently were Bibtex [4] and IdeaMÂCHÉ [9].

- *Shortcomings of Existing Software.* Participants pointed out that it is too easy to "get lost" using current software. When exploring many papers, participants said they just ended up with too many tabs open, unsure of where they were or how they got there. Another issue identified was that collecting and managing bibliographic information is difficult.

- *Desired Functionality.* Almost universally, participants expressed a desire to view their navigation history in a meaningful way.

The Value of Monadic Network Visualization

In the post questionnaire, participants were asked if they found themselves exploring citation chains differently with the monadic interface than they do with current software. 4 of the 5 participants answered yes to this. When asked if they learned anything new about the field they were exploring while using the interface, participants noted that they discovered papers that they likely would not have discovered using traditional interfaces.
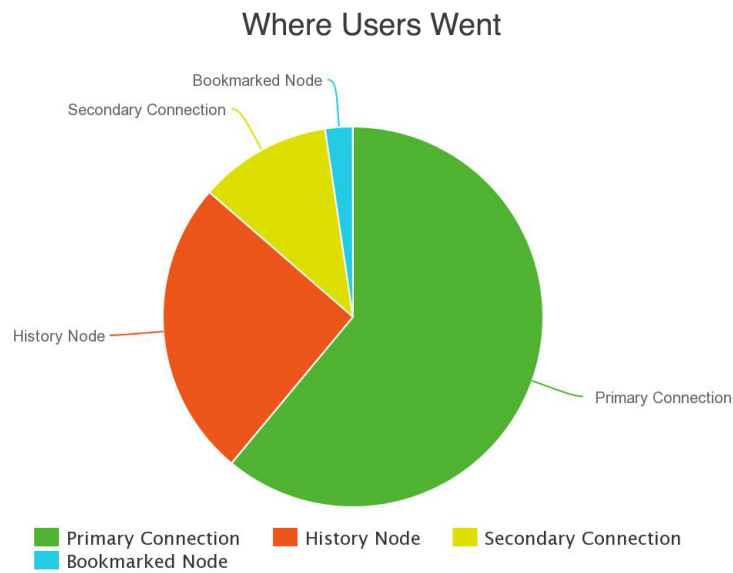
Where Users Went

Figure 4. The breakdown of where users chose to navigate.

The data on which nodes users chose to explore supports the hypothesis that the monadic interface affords new and varied types of exploration when compared to traditional interfaces. Of nodes navigated to, almost 40% were something other than a direct connection to the focus node. These non-linear types of explorations, which are not supported by traditional interfaces, confirm that users took advantage of the higher Fluency of ideas provided by the monadic interface.

In the post interview, multiple participants asserted that being able to see and navigate to secondary connections as well as primary was valuable. Additionally they called the interface "more fun" and "less cluttered" than traditional interfaces.

Focus Node Summary

When asked directly if the focus node metadata provided an adequate summary of the web page the first participant in the study noted that it would be helpful to also show the citation count and

year of papers. Those fields were then added. It was observed that, of the nodes that were in

focus, only 19.7% had their associated web pages opened in a new tab. Nodes were bookmarked

at a rate of 17.9%. From this, I conclude that the metadata of the focus node provided an

adequate summary of the web page it represented in the context of monadic exploration.

Supporting Backtracking

In the post questionnaire, each participant answered yes to the question of whether the history

visualization was useful for backtracking to previously visited nodes. This is supported by how

frequently (25.3% of navigations) users navigated to nodes in the history. In interviews,

participants identified the history visualization as useful, but also noted that it could be greatly

improved. Multiple participants mentioned that they would like to have a less linear and more

graphical visualization of their traversal of the citation network.

Discussion

Implications for Design

This study provided a number of valuable insights into the needs of monadic interfaces, and

established the interface I created as more viable than existing software for exploring networks

of scholarly articles and their authors. The first of these insights is that robust bookmarking

functionality is conducive to exploratory browsing. Participants requested the ability to

categorize, annotate, and organize bookmarks as a needed feature for the monadic interface.

Another feature requested multiple times was the ability to remove nodes from the collection

graph. Participants wanted to remove less relevant nodes so that other nodes would have room to

be shown. This desire is consistent with the "pruning" model of search, which prior research has

identified as valuable for exploring networked information [2,17].

Participants also had substantial insight on what metadata was needed for the focus node summary. What was requested often varied from participant to participant based on the domain they were exploring. For example, two participants exploring papers about Human Computer Interaction techniques said that figures from the papers would be the most useful addition. From this I hypothesize that presenting users with a list of possible metadata fields and letting them choose which ones are shown, both in the focus node and on hover of a graph node, would optimize what data is presented.

Future Research

The primary research question that arose from the study was how to best visualize the history of monadic exploration. The current implementation simply displays a linear list of nodes visited. It does not take advantage of the networked nature of the data underlying it. Participants hypothesized that visualizing the navigation history as a network graph would help them better understand how they arrived at the current focus node. How best to do this is a nontrivial question that I plan to investigate further.

REFERENCES

[1]   ACM Digital Library, 2015. Retrieved April 8, 2015, from Association for Computing Machinery: http://dl.acm.org

[2]   Bates, M., (1989),"The design of browsing and berrypicking techniques for the online search interface", Online Review, Vol. 13 Iss 5 pp. 407 – 424

[3]   Beghetto, R.A. and Kaufman, J.C. 2007. Toward a broader conception of creativity: A case for "mini-c" creativity. *Psychology of Aesthetics, Creativity, and the Arts*, Vol 1(2) (May 2007), 73-79.http://dx.doi.org/10.1037/1931-3896.1.2.73

[4]   BibTex, 2006. Retrieved April 8, 2015, from http://www.bibtex.org/

[5]   Dontcheva, M., Drucker, S.M., Wade, G., Salesin, D., and Cohen, M.F. 2006. Summarizing personal web browsing sessions. *Proc. 19th Annu. ACM Symp. User interface Softw. Technol. - UIST '06* (2006), 115.

[6]   Dörk, M., Comber, R., and Dade-Robertson, M. 2014. Monadic exploration: seeing the whole through its parts. *Proc. 32nd Annu. …* (2014).

[7]   Fruchterman, T. and Reingold, E.. 1991. Graph drawing by force directed placement. *Softw. Pract. Exp.* 21, NOVEMBER (1991), 1129–1164. DOI:http://dx.doi.org/10.1002/spe.4380211102

[8]   Google Scholar, 2015. Retrieved April 8, 2015, from Google: http://scholar.google.com

[9]   IdeaMÂCHÉ, 2015. Retrieved April 8, 2015, from Interface Ecology Lab: http://ideamache.ecologylab.net/

[10]   Kerne, A. et al. 2014. Using Metrics of Curation to Evaluate Information-Based Ideation. *ACM Trans. Comput. Interact.* 21, 3 (June 2014), 1–48. DOI:http://dx.doi.org/10.1145/2591677

[11]   Kerne, A., Qu, Y., Webb, A. M., Damaraju, S., Lupfer, N., and Mathur, A. 2010. Meta-metadata. *Proc. 19th ACM Int. Conf. Inf. Knowl. Manag. - CIKM '10* (2010), 1129. DOI:http://dx.doi.org/10.1145/1871437.1871580

[12]   Latour, B., Jensen, P., Venturini, T., Grauwin, S., and Boullier, D. 2012. "The whole is always smaller than its parts" - a digital test of Gabriel Tardes' monads. *Br. J. Sociol.* 63, 4 (2012), 590–615.

[13]   Morris, D., Morris, M., and Venolia, G. 2008. SearchBar. In *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*. New York, New York, USA: ACM Press, 1207.

[14]   Qu, Y., Kerne, A., Lupfer, N., Linder, R., and Jain, A.. 2014. Metadata Type System : Integrate Presentation , Data Models and Extraction to Enable Exploratory Browsing Interfaces. (2014).

[15]   Tarde G. and Lorenc, T. Monadology and sociology. Prahran, Vic: re.press, 2011. Retrieved April 8, 2015, from re.press: http://re-press.org/book-files/9780980819724-Monadology_and_Sociology.pdf

[16]   Utting, K. and Yankelovich, N. 1989. Context and orientation in hypermedia networks. *ACM Trans. Inf. Syst.* 7, 1 (January 1989), 58–84. DOI:http://dx.doi.org/10.1145/64789.64992

[17]   Wilkins, J. Järvi, J., Jain, A., Kerne, A., Kejriwal, G. Kumar, V. EvolutionWorks: Towards Improved Visualization of Citation Networks, *Proc. ACM INTERACT 2015*.