

Interaction Design for Literature-Based Discovery

Meredith M Skeels¹, Kiera Henning², Meliha Yetisgen Yildiz³, and Wanda Pratt^{1,3}

¹Biomedical & Health Informatics
University of Washington
Seattle, WA, USA
{mskeels, wpratt}@u.washington.edu

²Computer Science & Eng.
University of Washington
Seattle, WA, USA
kierah@cs.washington.edu

³Information School
University of Washington
Seattle, WA, USA
melihay@u.washington.edu

ABSTRACT

Rapid growth in the scientific literature makes it increasingly difficult for scientists to keep abreast of findings outside their own narrowing fields of expertise. To help biomedical researchers address this problem, LitLinker uses literature-based discovery to find new connections between biomedical terms that could lead to new directions in research. In this paper, we discuss the design of an interface that supports researchers' interactive exploration of the identified connections. Because the interface suggests many possible new connections, researchers must be able to understand how connections are established and to evaluate those connections based on their own expertise. Based on the results of our user study, we have further tailored the interface to support the work processes of biomedical researchers. LitLinker's interaction design promotes user-comprehension of the complex relationships among connected terms and allows for dialogue with researchers on the use of literature in scientific discovery.

Author Keywords

Interface, Interaction Design, Text Mining, Visualization

ACM Classification Keywords

User Interface Design, Interaction Design, Usability Testing and Evaluation.

INFORMATION OVERLOAD

Information overload has become a significant problem for biomedical researchers. Scientific literature is readily available, but the sheer volume and growth rate of the literature makes it impossible for researchers to keep up with new findings outside their own narrowing fields of expertise. This isolation of scientists inhibits innovation across multiple fields. To combat this problem, we developed a system called LitLinker that identifies possible links across disparate sections of the literature.

Unlike the tools currently available to researchers, such as search engines, LitLinker generates results about possible new connections between terms. Given a term entered by a researcher, LitLinker identifies a large number of potential new connections between previously unconnected terms. For this information to be useful to researchers, the interface must facilitate comprehension, investigation, and evaluation of the connections proposed. This paper

discusses the evolution of the interactive visualization interface we designed to support this task.

FINDING NOVEL CONNECTIONS

Researchers have been working in the general area of biomedical literature-based discovery for nearly fifteen years. Swanson, a pioneer in the area, used a combination of citation analysis and manual review to identify plausible new connections across disjoint biomedical literatures [5]. In an early example, Swanson identified a hidden connection between the disjoint literatures on *migraine* and *magnesium*[5]. He noticed this hidden connection by identifying several linking medical terms, such as *epilepsy* and *calcium channel blockers*, that occurred frequently in the titles of both the *magnesium* literature and *migraine* literature. Swanson's work introduced the seminal idea of using an intermediate linking literature, which inspired the development of various literature-based discovery systems, such as Arrowsmith[3], Manjal[4], LitLinker[2], and Bitola[1].

Our system, LitLinker, was designed with what Swanson calls an open discovery approach. Our literature-based discovery begins with a **starting term** (i.e. *migraine*), the term the researcher is interested in investigating. Next, LitLinker uses a text-mining process to find a set of terms that are directly correlated with the starting term. We refer to this first set of correlated terms as the **linking terms** (e.g., *epilepsy*, *calcium channel blockers*). For each of the linking terms, LitLinker then uses the same text-mining process to identify a set of terms that are correlated with each linking term. We call these final terms **target terms** (e.g., *magnesium*). The first version of LitLinker operated only in batch mode and was used only to test its ability to replicate previous literature-based discoveries.

The amount of data produced by LitLinker for a typical starting term is immense. When LitLinker replicated Swanson's Migraine-Magnesium discovery, it processed over 4 million documents. It generated 349 linking terms and 545 target terms with 57,622 possible starting term-linking term and linking term-target term combinations. Thus, for biomedical researchers to benefit from the system they must be able to explore and evaluate the potential new discoveries.

Arrowsmith [3] was the first system based on Swanson's approach that incorporated an interactive, web-based user

interface. Unlike LitLinker, it functions in a closed-discovery fashion; the user must enter both a starting term and a term or category for a candidate target term. Arrowsmith's task is to generate and display the list of linking terms between the starting term and the target term.

Recently, other systems, Bitola [1] and Manjal [4], have been created using an open-discovery approach, like LitLinker's. Their interfaces enable researchers to enter a starting term and various settings to control aspects of their text mining algorithms. To display results, Bitola provides lists of linking and target terms that can be sorted by different parameters, such as connection strength and name. Manjal provides an additional tree-structured graph visualization, which displays the linking terms as children of the starting term and the target terms as children of the linking terms.

Arrowsmith, Bitola, and Manjal provide interfaces that enable user navigation through the generated results and corresponding literatures. However, they are limited in the ways researchers can navigate and evaluate their results. For example, with these interfaces, it is not possible to retrieve all the target terms connected to a selected linking term or all the linking terms that connect a selected target term to the starting term. Such multi-dimensional navigation capabilities could prove essential for researchers who want to evaluate the quality and validity of the potential connections that the system generates.

INITIAL DESIGN

LitLinker's success in facilitating new discoveries depends on the interface's ability to inform and engage its users as they attempt to interpret and evaluate proposed connections. Accordingly, the interface must achieve three goals. Goal 1: The interface must promote user comprehension of the complex relationships among the terms involved in each proposed connection. Goal 2: The interface must provide flexible navigation and a level of detail appropriate to the scope of each view without obscuring data necessary for the evaluation of connections. Goal 3: The interface should help researchers incorporate LitLinker's results into their own research discovery process.

In the initial design phase, we focused on the first and second goals. The biggest constraint on the display of LitLinker's results is the large number of connections identified for a typical starting term. A tree-structured graph, for example, is quite appealing, but the sheer number of terms used to establish potential new connections makes graph visualization impractical.

To promote comprehension of the results (Goal 1), the interface must convey clearly the relationships among the three different types of terms. We used both consistent color cues and consistent horizontal positioning on the screen to differentiate the three types of terms. The starting term appears in the left-most column against a green background, the linking terms appear in the center column

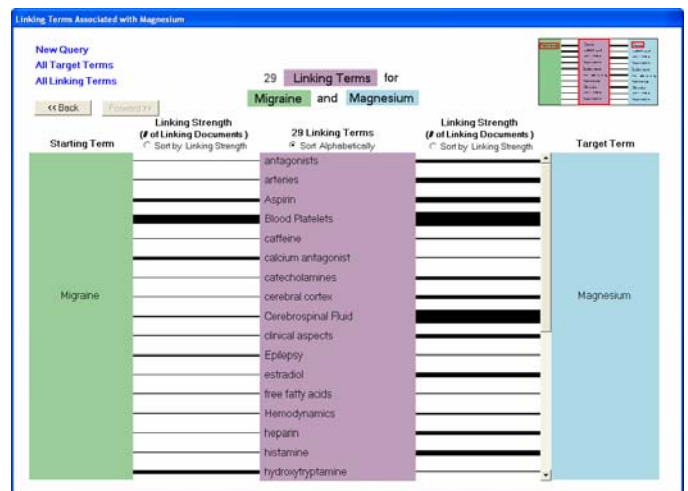


Figure 1. From this view users can see all the linking terms connecting *Migraine* and *Magnesium*.

against a purple background, and the target terms appear in the right-most column against a blue background (Figure 1). The relative positions of each type of term reminds users that target terms are never directly connected to starting terms; they are always connected through linking terms. We maintain this physical orientation and color scheme throughout the interface as visual reminders about which part of a connection is currently being viewed. The connections between terms are represented in the interface as connecting lines of varied thickness. The thickness of each line represents the strength of the connection between two terms (Figure 1). This visualization allows quick comparison of connection strengths among various terms.

The first view a researcher sees of the interface is a top-level overview of results. The overview contains the starting term in the left column and all the target terms in the right column with lines bridged over the linking terms column to connect the two terms. The thickness of these lines is based on the number of linking terms connecting the target term to the starting term.

When the researcher sees an interesting target term, she can investigate the connection between the target term and her starting term (Goal 2) by clicking on the line linking the two terms. This action reveals all the linking terms that were used to generate the potential connection (see Figure 1). In this view, the linking terms are connected to both the starting term and to the target term by documents, and correspondingly, the line thickness represents the number of documents that connect the terms. To evaluate these connections, the researcher can narrow her viewing scope further by clicking on a linking line and then examining citations and abstracts for the documents that comprise connections between terms.

Further navigational flexibility is furnished by supporting a second method of navigation. A researcher might prefer to begin by examining all linking terms and then investigating target terms generated by interesting linking terms. The LitLinker interface supports this method of exploration by

providing a view of all the linking terms. Throughout the interface, clicking on a linking term reveals all the target terms generated from that linking term and clicking on a target term reveals all the linking terms connecting that target term to the starting term.

USABILITY EVALUATION

After implementing a prototype of our initial design, we conducted a usability evaluation with ten participants, including nine graduate students and one faculty member. Participants were recruited from the Information School and Division of Biomedical and Health Informatics at the University of Washington. All had research experience. Because LitLinker is unlike any system currently available, prior experience with other tools was not a concern, and none of our participants were familiar with the LitLinker project. The evaluation consisted of three parts: a general introduction, a task-based questionnaire, and an interview. The participants completed the evaluation in 25-40 minutes, with most variation during the interview section.

The introduction included a scripted overview describing LitLinker's function and purpose. The introduction did not include any discussion of the interface or any views of the interface. After the introduction, participants were asked three questions to identify any initial misunderstandings about LitLinker's discovery process and were given opportunities to ask questions.

Participants used LitLinker with *Migraine* as the starting term, to complete a task-based questionnaire with nine questions, one of which had two parts. The tasks were designed to evaluate each participant's ability to find specific data, to navigate the interface, and to compare the strengths of connections. Participants were asked to speak aloud as they completed the tasks. The interviewer observed without answering questions and noted any difficulties participants experienced.

After participants completed the questionnaire, we interviewed them to discover any aspects of the interface that were confusing or were particularly helpful. This dialogue also helped explain the reasons behind incorrectly answered questions. For example, one participant incorrectly identified the strongest connection. Through dialogue, we discovered that she understood the meaning of connection line width, but was unaware of the sorting feature. Thus, her error was the result of inaccurate judgment of line width and LitLinker's inability to make the sorting feature visible. Probing questions were used to elicit feedback about improvements that might be made.

RESULTS

The results of our usability evaluation reveal that participants understood the relationships among terms and were able to navigate the results. As is evident from the comparison tasks (Figure 2), participants had no trouble understanding that the thickness of the lines connecting two terms indicated connection strength. The line thickness

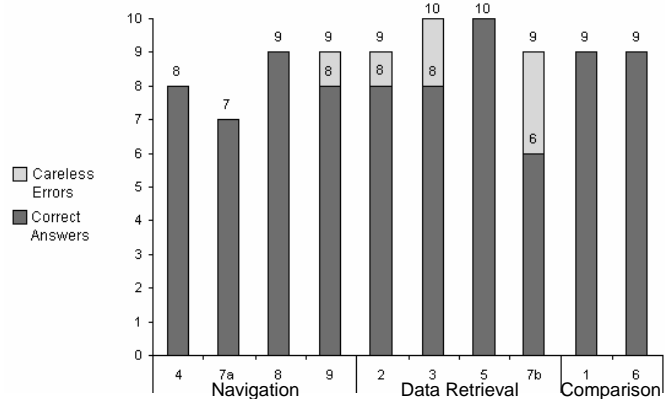


Figure 2. Questionnaire tasks answered correctly, or incorrectly due to careless errors, grouped by type of task.

metaphor was intuitive and useful, and through discussion with participants, it became clear that the consistent coloring and spatial location of the three types of terms promoted understanding of both the results and the literature-based discovery process. The evaluation also uncovered opportunities to improve the interface.

Five participants completed all the tasks correctly, and five answered at least one question incorrectly. We have grouped the incorrect answers into mistakes resulting from misunderstandings and mistakes that reflect careless errors that were independent from the interface (Figure 2). In an example careless error, one participant successfully located the oldest document connecting two terms, but recorded the date instead of the title. The task required him to record the title, and thus was not successfully completed, but this error did not reflect a misunderstanding about the interface.

The most straightforward problem identified in the evaluation was confusion about which components of the interface were interactive. Specifically, some participants did not understand that they could click on lines connecting terms. This oversight caused problems for the navigation tasks and caused one data retrieval mistake. Users who did recognize that lines were "clickable" accomplished tasks in the interface with much greater success.

The evaluation also uncovered a more challenging problem. Three participants either could not find the link to view all the linking terms, or misunderstood the difference in scope between the views. These participants missed question 7a, which required them to navigate to the list of all linking terms. Thus, a subsequent version of the interface must clarify the difference between the view of all linking terms and the view of a subset of linking terms.

From the results of our usability evaluation, we concluded that the fundamental layout of the interface promoted understanding of the complex discovery process in LitLinker. One participant commented that she was initially confused about the relationships among terms, but that using the interface improved her understanding of those relationships. We also concluded that future improvements

must include greater clarity about which parts of the interface are interactive, more visible sorting features, and better cues about the scope of linking terms being viewed.

REDESIGN AND ENHANCEMENT

The next phase in the interface evolution was to develop a web-accessible version (litlinker.ischool.washington.edu) that incorporated the conclusions drawn from participant feedback. To visually distinguish the interface components that are interactive from those that are not, all “clickable” components now highlight red when a user positions the cursor over them. The sorting features have been moved to a conspicuous location near the top of each view. To remind users of the difference between viewing all the linking terms and viewing just those used to generate a particular target term, the linking term overview is now listed along side the target term overview in the navigation menu as a reminder of its top-level scope.

In addition to these improvements, we added support for our third interface goal: integration of LitLinker into researchers’ normal discovery process. To support such integration, LitLinker needed to facilitate collaboration and evaluation of results.

To support a researcher’s evaluation process, we noted that the interface would need to support:

1. Sorting into personally meaningful categories
2. Marking terms for intended action
3. Marking terms for attention from a colleague
4. Identifying notably relevant or irrelevant terms
5. Eliminating irrelevant or uninteresting subsets from the result set

The first four activities in the list pertain loosely to organization and planning. LitLinker assists these tasks by providing two ways to annotate terms. The first type of annotation involves assignment of priority to terms. Priority is a mutually-exclusive categorization, allowing researchers to record their own interest level or assessment of a term’s potential before delving in to the details of each connection. All terms begin with neutral priority. A researcher can assign initial priority levels to push the most promising connections to the top. As the researcher explores the terms and the connections that produced them, she can change a term’s priority to reflect any changes in opinion that result from further investigation. The second type of annotation allows researchers to tag any linking term or target term with one or more meaningful icons. These icon tags are visual representations of user-determined organizational categories or intended actions. The icons provided as an example indicate “Save for Later”, “Attention!”, and “Not Useful.” Because they can be personalized, these annotations are designed to allow flexibility. Researchers can assign any number of icon tags to a term, allowing them to overlap categories or markers of intended action. To annotate a term, users toggle one or both types of annotation to change the interaction mode from navigation to annotation. Once the desired priority or icon tag is

selected, users annotate linking or target terms simply by clicking on them. By using toggle-able interaction modes for one unified view of the results, the annotation tools are seamlessly incorporated into the navigation environment.

Finally, if the researcher finds that a large number of related terms are irrelevant, she can perform the fifth activity by selecting subsets of linking terms or target terms to remove. In summary, the enhancements made to the LitLinker interface in this most recent phase of development should assist researchers in an iterative process of evaluation and refinement of identified connections.

CONCLUSIONS AND FUTURE WORK

The interface that we have developed and evaluated for LitLinker is designed to allow researchers to carefully assess the potential connections generated by LitLinker. After migrating to an interface with improved visual cues and evaluation features, we are now ready to gather direct feedback from biomedical researchers about the interface design and necessary improvements to the evaluation toolset. In the near future, we are planning longitudinal case studies with biomedical researchers using LitLinker. Researchers will use LitLinker to find possible new connections to their current research that would guide future research. Through this interaction, we will gain a better understanding of how biomedical researchers might incorporate LitLinker into their work and gain further insight into the role the literature can play in discovery. By pairing an innovative literature-based discovery process with an interface that supports the researchers’ need to understand and evaluate the results, LitLinker will be a powerful tool for scientific researchers.

ACKNOWLEDGMENTS

We would also like to thank Jinmei Xu and Brian Stoll for their work on earlier versions of LitLinker. The National Science Foundation Award #IIS-0133973 funded this work.

REFERENCES

1. Hristovski, D., Peterlin, B., Mitchell, J.A., and Humphrey, S.M. Improving literature based discovery support by genetic knowledge integration. In *Stud Health Technol Inform* 95, (1993), 68-73.
2. Pratt, W., and Yetisgen-Yildiz, M. LitLinker: Capturing Connections across the Biomedical Literature. In *Proc. K-Cap’03*, Florida (2003), 105-112.
3. Smalheiser, N.R., and Swanson, D.R. Using ARROWSMITH: A Computer Assisted Approach to Formulating and Assessing Scientific Hypothesis. In *Computing Methods Programs in Biomedicine* 57, 3 (1998), 149-153.
4. Srinivasan, P. Generating Hypotheses from MEDLINE. In *JASIST* 55, 5 (2004), 396-413.
5. Swanson, D.R. Migraine and Magnesium: Eleven Neglected Connections. In *Perspect Biol. Med.* 31, (1988), 526-557.