

# Understanding and Supporting Academic Literature Review Workflows with *LitSense*

Nicole Sultanum  
University of Toronto  
Toronto, Ontario, Canada  
nicolebs@cs.toronto.edu

Christine Murad  
University of Toronto  
Toronto, Ontario, Canada  
cmurad@taglab.ca

Daniel Wigdor  
University of Toronto, Chatham Labs  
Toronto, Ontario, Canada  
daniel@dgp.toronto.edu

## ABSTRACT

It is increasingly difficult for researchers to navigate and reach an understanding of a growing body of literature in a field of research. While past works in HCI and data visualization sought to support such activities, few investigated how these workflows are conducted in practice and how practices change in view of support tools. This work contributes a more holistic understanding of this space via a user-centered approach encompassing (a) a formative study on literature review practices of 15 researchers which informed (b) the design of *LitSense*, a proof-of-concept tool to support literature review workflows, and (c) a week-long study with 12 researchers performing a literature review with *LitSense*.

## CCS CONCEPTS

• **Human-centered computing** → *User interface management systems; User studies*; • **Applied computing** → *Document management*.

## KEYWORDS

Bottom-up literature review, Sensemaking, Longitudinal study

### ACM Reference Format:

Nicole Sultanum, Christine Murad, and Daniel Wigdor. 2020. Understanding and Supporting Academic Literature Review Workflows with *LitSense*. In *International Conference on Advanced Visual Interfaces (AVI '20)*, September 28–October 2, 2020, Salerno, Italy. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3399715.3399830>

## 1 INTRODUCTION

Literature review workflows are complex activities consisting of “active and constructive exploration” [6] that lead to a gradual, “bottom-up” understanding of a domain [9]. This process entails searching for published research in a given field, synthesizing articles, organizing a subset of articles into a growing library of papers, and nesting new additions within existing thematic frameworks in the library [34]. Researchers face many challenges when conducting a literature review, from the growing scale and scattered nature of academic literature to the experience and expertise needed to forge meaningful connections between publications, tasks which are particularly challenging for novice researchers [10, 32].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

AVI '20, September 28–October 2, 2020, Salerno, Italy

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7535-1/20/09...\$15.00

<https://doi.org/10.1145/3399715.3399830>

Several commercial tools emerged to assist researchers in these tasks [3] (e.g., Mendeley [2] and Zotero [1]). Past research in data visualization proposed overview strategies for academic publications, such as citation connections [7, 15, 16, 23, 24, 26, 31], meta-data [5, 15, 26, 28], and publication venues [15]. These tools provide a top-down view of literature and facilitate paper discovery, but arguably support only part of the whole knowledge *workflow*.

We define the *literature review workflow* as the set of knowledge-building tasks that entail the gradual discovery of relevant papers, the bottom-up creation and refinement of semantic connections, and the emergence of new information [12, 34] leading to a broad and thorough understanding of a research question. Nowadays, support for segments of this workflow is distributed across separate tools, which is detrimental to the understanding and summarization of literature. While better tool integration may lead to improved workflows [6, 34], little has been done to investigate what and how tools should be integrated to support the whole workflow, and how they are affected by this integration.

This work fills this gap by revisiting literature review practices from a more holistic and grounded perspective to guide interface design efforts. Our contributions include (a) findings from a formative study about literature review practices; (b) the design and development of *LitSense*, a tool to support bottom-up literature review workflows; and (c) findings and lessons learned from a 1-week longitudinal evaluation of *LitSense*.

## 2 RELATED WORK

Commercial tools such as Mendeley [2] and Zotero [1] are widely used to help manage collections of papers but provide limited discovery and organizational support. A number of research works also focus on supporting paper discovery based on metadata [18], paper content [4], reading history [22], and citation networks [9, 17, 27] in isolation. On a more integrated approach, CiteSense [34] encompasses paper searching, organization structures, and note-taking to enable fluid transitions between tasks. These works support bottom-up expansion from a set of seed papers and assist in a gradual definition of a semantic framework to represent knowledge from literature, but provide limited support for knowledge synthesis and visualization. In addition, evaluations of these systems were limited to preliminary assessments [4, 27, 34] or controlled studies using predefined tasks [4, 9, 22], thus making it unclear how these systems would fare in real world scenarios.

A larger body of research has focused on visualizing aspects of academic document collections [13], with an emphasis on the complex relationships between papers [12, 14, 19, 26]. A citation network, *i.e.*, what papers have been cited, and by whom [10, 14–16, 23, 24, 26, 27, 31], is often presented as a node-link graph [9, 10, 26,

31]. Other works have focused on visualizing clusters of topics [26, 33], author-paper relationships [20, 21, 35], and other types of paper metadata [5, 11, 15, 26, 28]. A common limitation of these systems is that the underlying data is largely static; despite interactivity that enables one to explore different facets of the data, they do not offer users the ability to augment papers with user-defined concepts to allow for the gradual construction of meaning [12, 29], and lack sense-making support into one's own personal library.

### 3 FORMATIVE STUDY

Prior studies on literature review workflows [12, 34] were done before the advent of modern reference management tools [3]. We thus sought to reassess current workflows, including:

1. *How do researchers find new publications in a field?*
2. *How do researchers judge whether a paper is relevant to them?*
3. *How do researchers synthesize information from papers?*
4. *What tools do researchers rely on for their literature review?*

We interviewed 15 researchers (F1-F15), including 8 PhD students, 1 Master's student, 3 Post-Doctoral Fellows, 2 Industry researchers, and 1 Faculty, from various fields (Computer Science, Biology and Religious Studies). Interviews were audio recorded, transcribed, and analyzed via open coding.

#### 3.1 Findings and Discussion

Four themes emerged from our study: how researchers use *current management tools*, *organize* their library of papers, *discover papers* and *grow* their library, and *make sense of literature*.

**Modern tool use.** Many participants (8/15) used reference management tools, while some (5/15) tried but eventually abandoned them. While tools collectively covered a range of literature review tasks (e.g., organizing, searching, content overview), they suffered from a lack of integration, e.g.: *"...you try and find a new tool, and then you go through the whole process of trying to re-tag everything and then it just gets messy... you end up having ten different ways of keeping your own data and then you have to search five different places to find what you want"*(F13).

**Library management.** Publications were commonly organized by topic, project, publication date and relevance. Participants found that maintaining organizational structures was challenging. Some (6/15) stated that the way they organized papers evolved during a literature review (e.g., folder structures being renamed and reorganized), and the overhead to adjust organization structures was significant: *"I would spend too much time trying to refine the structure, and then coming back a year or two years later and saying, 'You know what, this structure doesn't work anymore'"* (F12). Intersectional research often led to duplication, e.g., *"some papers can potentially go in multiple categories...sometimes I'll put them in both because later it will be difficult to find it in one."* (F13). These points touch on well known issues around defining and updating early structures [29].

**Growing libraries.** Expansion strategies often began with *seed papers* provided by a supervisor or fellow researcher, or via search engines. Many (7/15) stated starting with an idea of what concepts to search for, and many (7/15) claimed new themes emerged as their knowledge evolved, with some losing their train of thought after continuous searches. References in a paper were another source of expansion (12/15): some (5/15) paused reading to look up citations,

which caused them to lose context of where and what they were reading; others (7/15) preferred finishing reading a paper before looking for references, which still required side notes on references to search later. While this echoes prior work on visualizing citation relationships [7, 15, 16, 23, 24, 26, 31]), there is still little support for keeping track of relevant citations to retrieve later.

**Synthesis of Literature.** Note-taking was used for summarization, questions, relevant points and quotes (11/15). While many liked having their notes directly on the article (be it on paper or digital form), most (10/15) wrote their notes in separate documents, for easy accessibility. Note-taking was seen as an important step to sensemaking, but reportedly with poor integrated support.

#### 3.2 Preliminary Design Goals

Based on our findings and prior work, we outline five design goals for a more integrated support to literature review workflows:

**(G1) Facilitate content discovery and inclusion**, via multiple expansion strategies (e.g., from keyword search and related work).

**(G2) Support content overview & details on demand**, with gradual access to metadata for quick relevance assessments and to reduce cognitive overload (e.g., title, authors, then abstract).

**(G3) Allow for flexible, user-driven organization structures**, to (re)organize papers based on user-defined relations or categories.

**(G4) Support multi-level note-taking**, both at a paper note level and at a library summary level.

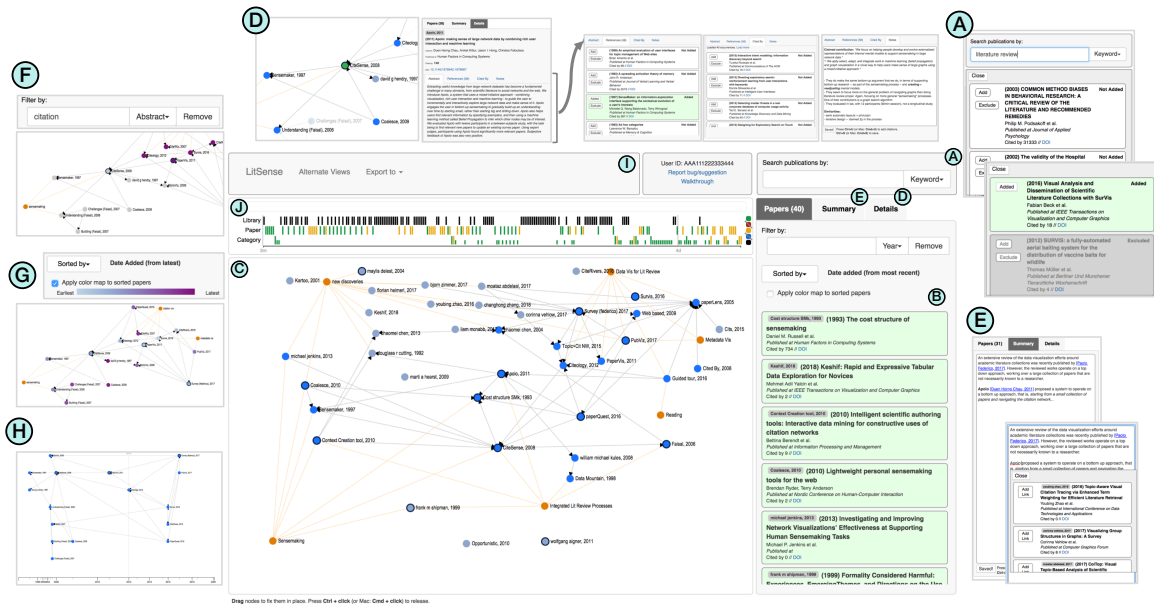
**(G5) Support situational awareness**, by tracking and conveying relevant changes to a paper collection, (e.g., relationships between papers, past paper searches, or if a papers has been read).

### 4 LITSENSE




*LitSense* is an end-to-end tool to support literature review workflows from paper discovery to synthesis. Unlike existing reference management tools that maintain "collections" of papers, *LitSense* is designed to focus on a particular *research question*, a distinction relevant to cases where a paper contributes different perspectives to various literature reviews and mental models (G3, G5). Following, we cover the various stages of a literature review with *LitSense*: *expanding*, *organizing*, *writing*, and *visualizing* the library.


**Discovery and Expansion.** From the *Search* panel (Fig. 1(A)), a researcher can look for seed papers based on keywords, paper titles or author names (G1). Search results use the *Microsoft Academic Knowledge API* [25, 30] and list metadata, including title, authors, source, number of citations, and a DOI link (G2). From search results, papers can be added to the library (G1) or "excluded" (i.e., flagged as not useful to the research question and shown grey in future searches, Fig. 1(A)) (G5). Added papers appear in the *Publications list* on the right-side panel (Fig. 1(B)), and in the *Topic Graph View* (Fig. 1(C)) as blue nodes in a node-link graph. A grey *citation connection* is automatically created when a paper cites another (G5). Clicking on a paper node in the *Topic Graph view* updates the right-side panel with details (Fig. 1(D)), including the abstract (G2), a space for notes (G4), a *References list* for backward citations, and a *Cited By list* for forward citations (G2). Papers in these lists can be easily added and excluded from the library (G1).

**Organizing the Library.** Papers in the library can be organized in several ways. The *Topic Graph* allows users to define *category*













**Figure 1: *LitSense* components.** (A) Search panel; (B) Publications list, (C) Topic Graph; (D) Paper Details tab, active when a paper node is selected; (E) Summary tab; (F) Filter by keyword; (G) Sort by criteria, with color map ; (H) Timeline view ; (I) the top menu, featuring export functions and an option to alternate the Topic Graph and Timeline views; and (J) the History view.

nodes  to represent semantic labels (e.g., keywords and topics) for a group of paper nodes  (G3). While nodes in the Topic Graph are automatically positioned via a force directed layout, they can also be manually re-positioned when dragged by a user, which "pins" them in place (with a darker hue ) while the rest of the graph readjusts the remaining "free" nodes. Paper node labels, follow a "<first author>, <year>" format by default but can also be customized to more memorable identifiers for a paper (G3, G4).

**Writing.** Notetaking is supported on a paper level (via the Details tab, Fig. 1(D)) and on a library level (via the Summary tab, Fig. 1(E))(G4). Users can cite library papers within notes via a keyboard shortcut (*Ctrl+U*, adding an interactive link to that paper for easier retrieval within *LitSense* (Fig. 1(E)). Papers featuring notes appear highlighted with a dark outline  (Fig. 1(C)) (G5).

**Viewing and Sensemaking.** *Litsense* offers several visual features to reorganize and filter content as the library grows. A paper node selection in the Topic Graph emphasizes papers and categories connected to that node (Fig. 1(D)) (G2). Nodes in the Topic Graph can be highlighted according to the keyword and source (i.e., year, title, author or abstract) provided in the Filter By bar (Fig. 1(B,F)) or colored according to sorting order defined in the Sorted by menu (e.g., date added to library, publication year, title, and citation count, Fig. 1(B,G)) (G2). Paper nodes can also be rearranged by publication year with the Timeline view (Fig. 1(H), accessible via the "Alternate Views" menu, Fig. 1(I)) (G2).

To ground users on their past activities (e.g., past search queries), the History View (Fig. 1(J)) provides an overview of user actions (G5). Actions are represented as coloured lines distributed across three tracks: library actions (top), paper actions (middle track), and category actions (bottom). Colour maps to action types:  for added items (papers /categories tracks) (G1),  for excluded items (G1),  for edited node labels (G3),  for notes saved (papers/library tracks)

(G4) and  for search actions (library track) (G1). On the category track, shorter action markers map to connection actions, i.e., when connections were created or removed. Temporal separators provide a perspective of how long ago these actions happened based on the present date, which is marked with a "T". Hovering over action markers displays a panel with more details on each action (G2). Coloured squares on the right     can be toggled (e.g., ) to filter out the corresponding action type from view.

## 5 LONGITUDINAL STUDY

Following the design and development of *LitSense*, we conducted a qualitative evaluation to assess how it supported literature review workflows in *ecologically valid* scenarios. We recruited 12 participants (P1-P12, 8 female) with at least 2 years of research experience to use *Litsense* for 1-week at their own pace.

The study included two interview sessions. In the first session, we assessed participant backgrounds, conducted a walkthrough of *LitSense* on the participant's laptop, and explained requirements for their literature review task: (a) adding about 20 papers to their library and (b) writing a short review on the topic (about 4-6 paragraphs, using the Summary View) that cited papers in their library. Participants returned after one week for the second session, which included three questionnaires (including a SUS [8] questionnaire) and a semi-structured interview about their experiences. We also collected *LitSense* usage data, including the action logs feeding the History View. In the next segment, we discuss emerging findings and outline «*lessons learned*».

### 5.1 Findings and Lessons Learned

**Impact on literature review workflows.** *LitSense* was found expressive enough to support a variety of literature review practices,

		Mean	Median	Mode
1	Understanding the relationships between papers	2.75	4.75	3.0
2	Avoiding redundant work	2.67	4.08	2.0
3	Keeping things organized as the library grows	3.00	4.25	3.0
4	Being confident that all relevant references are captured	2.75	4.00	3.0
5	Identifying groups of similar papers	3.50	4.58	3.5
6	Developing a higher-level understanding of the literature	3.08	3.83	3.0
7	Recalling content of a paper in library (that I've read)	3.17	3.75	3.0
8	Maintaining my own custom organization/structures	3.50	4.00	3.5
9	Locating papers in my library	3.58	4.00	4.0
10	Finding relevant papers	3.83	4.08	4.0
11	Investigating a question from scratch	3.67	3.75	4.0

Agreement	LitSense - Current process	LitSense - Current process
1 2 3 4 5	0 10.11 11.21 12.31 13.41	
Strongly Disagree	Strongly Agree	

Figure 2: Questionnaire findings on literature review tasks, comparing *LitSense* (right subcolumns) to participant's current processes (left subcolumns). Color encodes the extent of differences, with *LitSense* being higher rated overall.

with overall experience rated positively and a SUS score of 77.5. *LitSense* was also rated better than participants' current practices for various literature review tasks (Fig. 2). When commenting on their experience with *LitSense*, most appreciated avoiding redundancy and improving one's understanding of the literature (8/12). At the same time, many claimed their processes did not change significantly (4/12), indicating that *LitSense* fit well within their practices. We attribute part of this success to the «*end-to-end support and cohesive transitions between all stages of the workflow*». On the other hand, participants wished for tighter integration with PDF documents (6/12), which is the lesser supported part of the process. While a need to «*support more seamless integration with reading media*» follows from this (e.g., an integrated PDF reader and annotator), it also entails opportunities to support tasks beyond reading, e.g., trigger paper searches directly from documents.

**Expanding the library.** When determining the relevance of a paper, some argued that they practiced more scrutiny before adding something to their library compared to their usual approaches for handling PDFs (3/12). On one hand, *LitSense* may have contributed to a more selective inclusion habit (which helps reduce clutter) but may have also discouraged more exploratory searches. Many participants also wished to have been provided more information about papers within *LitSense* to help assess relevance, e.g., abstracts (2/12) and direct PDF links (P3). As such, «*sufficient context on a paper should be provided to assess usefulness*», including at least the abstract and potentially a quick look at the paper (without a need to download it) before adding it to the library. To minimize reluctance to add papers and encourage breadth, creating «*an intermediate assessment tier (e.g., a "reading list")*» could potentially lead to more liberal explorations while reducing cognitive overload within more exploratory stages of the workflow.

**Organizing the library.** Most participants (11/12) used categories in the *Topic Graph* to organize papers under themes, (averaging 4 categories and 18 connections per participant). Spatial node positioning in the *Topic Graph* was also important in assigning additional meaning to groups of papers. Half (6/12) commented on ways that spatial layout encoded meaning, e.g., positioning categories and papers to mirror the flow of the written summary (P10) or how the lack of connectivity between two paper clusters denoted knowledge gaps between them (P2). Also, several (6/12) felt that the visuals and the spatial layout helped them recall details about

papers in their library, with a few (2/12) pushing recall further by relabeling paper nodes with keywords to summarize their content. Overall, the «*aggregated visual representation of relationships and user defined semantics over a user's library*» emerged as key to supporting an understanding of relationships and coverage (Fig. 2(1,4)), but perhaps with some room to improve expressivity of organizational structures (Fig. 2(8)).

**Sensemaking and insight.** While the majority agreed that «*a visual representation of the papers and relationships*» helped their understanding of the literature, we found they were leveraged in different ways. Participants looked at citation density to find influential publications (3/12), assessed connections between paper clusters to identify gaps and guide exploration (4/12), and made use of the *Timeline* view (6/12). In particular, most (9/12) stated that «*citation connections help assess coverage*», i.e., inspiring more confidence that a review of the space was sufficiently comprehensive. This illustrates the expressive nature of *LitSense* organizational structures and suggests a need to support diverse practices via «*multiple complementary views of the library*».

**Sensemaking and writing.** The *Topic Graph* categories played an important sensemaking role by serving as labels to reflect one's current understanding of the literature (10/12). This task is key to bridge progressive abstraction of concepts leading to literature understanding of a higher order [29]. The *Topic Graph* categories also offered a natural transition to subsections and helped guide presentation flow, which supported writing tasks (6/12). Participants organized topics via relabeling, deletion, and updating of categories and their connections to papers, but missed the ability to specify relations between categories, e.g., define a super category or a custom semantic relationship (4/12). Future efforts should focus on more «*expressive means to take apart, combine, and hierarchically organize categories and other semantic structures*». *Paper Notes* (4/12) also helped distill relevant information, especially when paper metadata was not informative enough to learn or recall key details of papers. A few (2/12) missed the «*ability to aggregate all paper notes into a single view*», similar to an annotated bibliography, which was argued to facilitate comparisons and further guide summarization efforts.

**Tracking progress.** Lastly, although several found the *History View* useful (5/12), few stated using it (2/12), and mainly to look up past search keywords. This may be partially attributed to the short period of the study and how fresh information was in participants' memory. Nonetheless, users still found it contained relevant information, which we argue would become increasingly relevant as literature review efforts span longer time frames.

## 6 CONCLUSION

In this paper we conducted a more holistic and ecologically valid assessment of literature review workflows. Our investigations followed a user centered approach and examined practices of 27 researchers, assessing current practices and how workflows fare under a more integrated approach to literature review with *LitSense*. While assessing the value of *LitSense* as a tool, we compiled lessons as design guidance for a more integrated approach to literature review workflows. We hope our efforts encourage a stronger focus on integration in future work.

## REFERENCES

- [1] 2006. Zotero. <https://www.zotero.org/> [Retrieved 21-September-2018].
- [2] 2008. Mendeley. <https://www.mendeley.com> [Retrieved 21-September-2018].
- [3] 2015. 101 Innovations. 2015. Timeline of tools. <https://101innovations.wordpress.com/2015/12/25/timeline-of-tools/> [Retrieved 21-September-2018].
- [4] Michelle Q. Wang Baldonado and Terry Winograd. 1997. SenseMaker: an information-exploration interface supporting the contextual evolution of a user's interests. In *Human Factors in Computing Systems*. 1–7. <https://doi.org/10.1145/258549.258563>
- [5] Fabian Beck, Sebastian Koch, and Daniel Weiskopf. 2016. Visual Analysis and Dissemination of Scientific Literature Collections with SurVis. In *IEEE Transactions on Visualization and Computer Graphics*. 343–354. <https://doi.org/10.1109/TVCG.2015.2467757>
- [6] Bettina Berendt, Beate Krause, and Sebastian Kolbe-Nusser. 2010. Intelligent scientific authoring tools: Interactive data mining for constructive uses of citation networks. *Information processing & management* 46, 1 (2010), 1–10.
- [7] Peter Bergström and Darren C. Atkinson. 2009. Augmenting the exploration of digital libraries with web-based visualizations. (2009). <https://doi.org/10.1109/ICDIM.2009.5356798>
- [8] John Brooke et al. 1996. SUS-A quick and dirty usability scale. (1996), 189–194.
- [9] Duen Horng Chau, Aniket Kittur, Jason I. Hong, and Christos Faloutsos. 2011. Apollo: making sense of large network data by combining rich user interaction and machine learning. 21, 2 (2011), 321–335. <https://doi.org/10.1145/1978942.1978967>
- [10] Jia-Kai Chou and Chuan-Kai Yang. 2011. PaperVis: literature review made easy. In *IEEE VGTC Conference on Visualization*. 79–101. <https://doi.org/10.1111/j.1467-8659.2011.01921.x>
- [11] Ao Dong, Wei Zeng, Xi Chen, and Zhanglin Cheng. 2019. VISTory: Interactive Storyboard for Exploring Visual Information in Scientific Publications. In *Proceedings of the 12th International Symposium on Visual Information Communication and Interaction*. 1–8.
- [12] Sarah Faisal, Paul A. Cairns, and Ann Blandford. 2006. Developing User Requirements for Visualizations of Literature Knowledge Domains. In *Conference on Information Visualization*. 17–24. <https://doi.org/10.1109/IV.2006.42>
- [13] Paolo Federico, Florian Heimerl, Steffen Koch, and Silvia Miksch. 2017. A Survey on Visual Approaches for Analyzing Scientific Literature and Patents. 36, 6 (2017), 201–225. <https://doi.org/10.1109/TVCG.2016.2610422>
- [14] Jianguan He, Qing Ping, Wen Lou, and Chaomei Chen. 2019. PaperPoles: Facilitating adaptive visual exploration of scientific publications by citation links. *Journal of the Association for Information Science and Technology* 70, 8 (2019), 843–857.
- [15] Florian Heimerl, Qi Han, Steffen Koch, and Thomas Ertl. 2016. CiteRivers: Visual Analytics of Citation Patterns. 4, 2 (2016). <https://doi.org/10.1109/TVCG.2015.2467621>
- [16] Angelo Di Iorio, Raffaele Giannella, Francesco Poggi, Silvio Peroni, and Fabio Vitali. 2015. Exploring Scholarly Papers Through Citations. 30, 3 (2015), 721–730. <https://doi.org/10.1145/2682571.2797065>
- [17] Taraneh Khazaei and Orland Hoerber. 2017. Supporting academic search tasks through citation visualization and exploration. *International Journal on Digital Libraries* 18, 1 (2017), 59–72.
- [18] Christin Katharina Kreutz, Peter Boesten, Alex Witry, and Ralf Schenkel. 2018. FacetSearch: A Faceted Information Search and Exploration Prototype. In *LWDA*. 215–226.
- [19] William Kules, Max L. Wilson, m.c. schraefel, and Ben Shneiderman. 2008. From Keyword Search to Exploration: How Result Visualization Aids Discovery on the Web. 23, 9 (2008), 2179–2198.
- [20] Shahid Latif and Fabian Beck. 2018. VIS Author Profiles: Interactive descriptions of publication records combining text and visualization. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 152–161.
- [21] Bongshin Lee, Mary Czerwinski, George G. Robertson, and Benjamin B. Bederson. 2005. Understanding research trends in conferences using paperLens. In *Human Factors in Computing Systems*. <https://doi.org/10.1145/1056808.1057069>
- [22] Simon Lehmann, Ulrich Schwanecke, and Ralf Dörner. 2010. Interactive visualization for opportunistic exploration of large document collections. In *Information Systems*. 55–62. <https://doi.org/10.1016/j.is.2009.10.004>
- [23] Chun-Hung Lu, Chih-Chien Wang, Min-Yuh Day, Chong-Shyong Ong, and Wen-Lian Hsu. 2008. Using Cited by Information to Find the Context of Research Papers. (2008). [https://doi.org/10.1007/978-3-540-69304-8\\_34](https://doi.org/10.1007/978-3-540-69304-8_34)
- [24] Justin Matejka, Tovi Grossman, and George W. Fitzmaurice. 2012. Citeology: visualizing paper genealogy. In *Human Factors in Computing Systems*. 107–116. <https://doi.org/10.1145/2212776.2212796>
- [25] Microsoft. [n.d.]. Project Academic Knowledge. <https://labs.cognitive.microsoft.com/en-us/project-academic-knowledge> [Retrieved 21-September-2018].
- [26] Rina Nakazawa, Takayuki Itoh, and Takafumi Saito. 2015. A Visualization of Research Papers Based on the Topics and Citation Network. 35, 2 (2015), 260–269. <https://doi.org/10.1109/IV.2015.58>
- [27] Antoine Ponsard, Francisco Escalona, and Tamara Munzner. 2016. PaperQuest: A Visualization Tool to Support Literature Review. In *Human Factors in Computing Systems*. 264–269. <https://doi.org/10.1145/2851581.2892334>
- [28] Alexander Rind, Andrea Haberson, Kerstin Blumenstein, Christina Niederer, Markus Wagner, and Wolfgang Aigner. 2017. PubViz: Lightweight Visual Presentation of Publication Data. 8, 4 (2017), 333–352. <https://doi.org/10.2312/eurovisshort.20171152>
- [29] Frank M. Shipman and Catherine C. Marshall. 1999. Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems. In *Conference on Computer Supported Cooperative Work*. <https://doi.org/10.1023/A:1008716330212>
- [30] Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-june Paul Hsu, and Kuansan Wang. 2015. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th international conference on world wide web*. ACM, 243–246.
- [31] John Stasko, Jaegul Choo, Yi Han, Mengdie Hu, Hannah Pileggi, Ramik Sadana, and Charles D Stolper. 2013. CiteVis: Exploring conference paper citation data visually. *Posters of IEEE InfoVis* 2 (2013).
- [32] Yun Wang, Dongyu Liu, Huamin Qu, Qiong Luo, and Xiaojuan Ma. 2016. A Guided Tour of Literature Review: Facilitating Academic Paper Reading with Narrative Visualization. In *Visual Information Communication and Interaction*. 677–680. <https://doi.org/10.1145/2968220.2968242>
- [33] Shaopeng Wu, Youbing Zhao, Farzad Parvinzamin, Nikolaos Th Ersotelos, Hui Wei, and Feng Dong. 2019. Literature Explorer: effective retrieval of scientific documents through nonparametric thematic topic detection. *The Visual Computer* (2019), 1–18.
- [34] Xiaolong Zhang, Yan Qu, C. Lee Giles, and Piyong Song. 2008. CiteSense: supporting sensemaking of research literature. 22, 1 (2008), 190–199. <https://doi.org/10.1145/1357054.1357161>
- [35] Jian Zhao, Michael Glueck, Fanny Chevalier, Yanhong Wu, and Azam Khan. 2016. Egocentric analysis of dynamic networks with egolines. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 5003–5014.