
Beyond iTunes for Papers: Redefining the Unit of Interaction in Literature Review Tools

Xin Qian

Matt J. Erhart

xinq@umd.edu

mattjerhart@gmail.com

University of Maryland, College Park

College Park, Maryland

Wayne G. Lutters

Joel Chan

lutters@umd.edu

joelchan@umd.edu

University of Maryland, College Park

College Park, Maryland

Aniket Kittur

nkittur@cs.cmu.edu

Carnegie Mellon University

Pittsburgh, Pennsylvania

ABSTRACT

Conducting an effective literature review is an essential step in all scientific work. However, the process is difficult, particularly for interdisciplinary work. Here, we articulate a key avenue for improvement for literature review tools: supporting the appropriate unit of interaction, which we argue is a “grounded claim”, a concise statement linked to key contextual details such as evidence. However, there are

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CSCW'19 Companion, November 9–13, 2019, Austin, TX, USA

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6692-2/19/11.


<https://doi.org/10.1145/3311957.3359455>

Claim

Scientists primarily read specific fragments of articles.

Context

Evidence



Related claims

Online journal logs show that scientists view only 1 - 3 pages on average; life scientists self-report mostly reading the findings sections of articles.

Provenance

Reported in mid-2000's by European computer-scientists in *Int'l Journal of Digital Libraries*.

Figure 1: An example of a claim that is grounded with different kinds of context: a key figure as evidence, corroboration by other related claims, and information about the provenance of the claim.

significant cognitive and interaction costs in creating them. We share insights from our development of a prototype literature review tool, the KNOWLEDGE COMPRESSOR, that aims to lower these costs.

INTRODUCTION

To advance knowledge scientists must order what is currently known and unknown about research problems [10]. This synthesis typically manifests as the challenging task of *literature reviewing*. Conducting effective literature reviews is challenging: beginning researchers often struggle to distill and organize key claims from papers [5], and experienced researchers often avoid conducting, updating, and publishing literature reviews, due to the imbalance of returns on the amount of time and effort needed [10]. **How can we make literature reviewing less painful and more commonplace?**

PROBLEM ANALYSIS

Researchers need to work with grounded claims

Many commercial (e.g., Mendeley, Zotero) and research literature review tools [11, 21] explicitly or implicitly assume that the *paper* is the unit of interaction. Functioning like “iTunes for papers”, researchers import papers, reason about relationships between papers (e.g., with tags, stars, citation links), annotate papers, and export papers as citations in written literature reviews. But is the paper the right unit of interaction for literature reviewing tools?

We argue that systems should enable researchers to work directly with **grounded claims**. Grounded claims are concise statements grounded in context that helps a user understand, interpret, judge, and use a claim (see Fig. 1). For example, a claim might be contextualized by its *evidence*, such as a key figure or experiment details, or *related claims* that corroborate, oppose, or clarify a focal claim. The *provenance* of a claim, such as its source collaboration networks, institutional dynamics, and prestige, can be important context for understanding its validity and impact [7]. Our idea for grounded claims is rooted in models of scientific argumentation that also specify a scientific statement, linked to evidence, as a basic unit of scientific discourse [8, 9, 16].

To see why researchers need to be able to work with grounded claims, not just papers, it is helpful to understand that sensemaking [14] is a core subtask of literature reviewing [21]. In the conceptual language of sensemaking theory [14], researchers search for a new representation (i.e., conceptual model, theory, or argument) that is based on existing ideas in the literature to guide their subsequent research. Importantly, the process of searching for these representations is highly iterative, involving multiple shifts in mental representations supported by reinterpretations of data [14]. For example, a researcher attempting to prototype a system supporting scientists’ paper reading behavior might need to reconcile conflicting claims on how scientists read, with some claim that people frequently read through whole documents, and others claim that scientists primarily read specific fragments of articles. The researcher may need to dig into the contextual details of evidence to consider whether

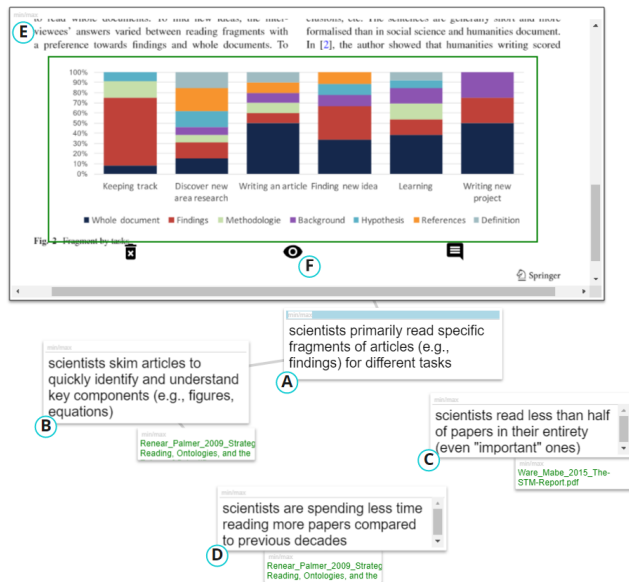


Figure 2: Four example grounded claims on KNOWLEDGE COMPRESSOR's claim canvas. The focal claim (A) is explicitly connected to a claim from another study that closely replicates it (B), and implicitly connected to a similar (but not exactly alike) claim (C). Another claim about the amount of reading scientists must do (D) suggests a potential explanation for scientists' reading behaviors. Clicking on a "min/max" button on the context segments toggles whether they are expanded or collapsed (E). If the user needs more information, there is a direct link (F) back to the location surrounding the claim.

¹An example usage of the KNOWLEDGE COMPRESSOR is in <https://vimeo.com/354474886>.

these contradictory findings might be explained by research methods, user populations, or user goals (i.e., what the scientists are reading for, such as discovering a new research area or writing an article).

The preceding example illustrates a key point from CSCW research on knowledge reuse: knowledge items must be identified and evaluated, but also be recontextualized in order to be reused effectively [1]. The lack of support for navigating de/recontextualization with grounded claims in existing systems may not be a critical barrier to scientific progress within disciplinary boundaries, where researchers are able to hold more information in memory as they gain expertise. However, there may still be room for substantial efficiency gains that can reduce the overhead of doing research. Research on sensemaking demonstrates that even small changes in the cost structure of individual access actions can have compounding effects on efficiency [13], which can impact downstream sensemaking, such as the degree to which researchers explore different ways of understanding existing findings and what research should be done next. Further, providing direct support for reuse is likely to be important for interdisciplinary, collaborative, and distributed research, which are increasingly becoming predominant forms of research [18, 19].

Efficiently creating grounded claims is challenging

As noted, literature review tools largely do not support working with grounded claims. Instead, researchers adapt other tools to approximate the ability to work with grounded claims, such as spreadsheets and text editors [20], and qualitative analysis software [3]. Some research systems have been proposed, such as "micropublications" [6] in bioinformatics, but uptake has been limited.

We assert that the central challenge to enabling grounded claims is the *cost structure* [14] of creating them. CSCW studies of knowledge reuse describe the **cognitive costs** of deciding in advance what details need to be retained as context for future reuse [2], with costly consequences for downstream task performance if a premature and inaccurate decision is made early on [15]. Also, systems that require high amounts of precision for specifying context during sensemaking may impose unnecessary **interaction costs** [4, 15].

THE KNOWLEDGE COMPRESSOR SYSTEM

To address these challenges, we are prototyping the KNOWLEDGE COMPRESSOR¹, a **claim-oriented** lit-reviewing system that enables users to create and work with grounded claims (not just papers) as a unit of interaction. In KNOWLEDGE COMPRESSOR, claims can be grounded by two kinds of context: 1) *evidence*: claims can be easily linked to one or more "segments" of a source PDF document (e.g., table, figure, methods details, quote) (see Fig. 3) and 2) *related claims*: claims can be connected (either with explicit links, or implicitly, by spatial proximity) to other claims on our claim canvas, which operates similarly to argument diagramming and modeling software (e.g., [17]; see Fig. 2).

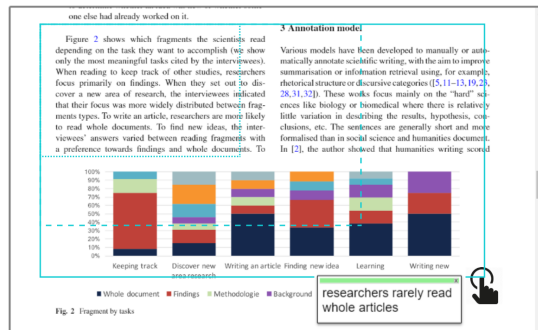


Figure 3: Users create claims by selecting a segment of the source document in the reading pane and writing a (re)description of the associated claim.



Figure 4: Users can adjust “flexibly compressed” context segments. For example, a user can scroll up in the segment to note that evidence for this claim comes from interviews (see the yellow highlight) and reason about reliability/diversity of evidence, or extract more details about which sections tend to be read for what kinds of tasks to inform design of a system that highlights specific sections for specific tasks.

At the heart of KNOWLEDGE COMPRESSOR is the “flexible compression” mechanism for lowering the cognitive and interaction costs of creating grounded claims. We deliberately designed the flexible compression to work similarly to conventional annotation, to reduce the cost of learning a new interface: users selecting a segment of a PDF document and typing a note in a text box to associate to that segment (see Fig. 3). Unlike conventional annotation, however, these segments are *flexible*, and can be adjusted/expanded by the reuser, because they are “live slices” of the source PDF (powered by a pdfjs React component) (see Fig. 4). Reusers can also link directly back to the source document in the reading pane from the segment, if even more context is needed. The flexibility of the context segment eases both the *cognitive costs* of deciding precisely which parts of the document count as context, and the *interaction costs* of precisely specifying contextual details, since the context can be adjusted more naturally by the reuser during the recontextualization step of knowledge reuse.

Example Usage

To illustrate the potential value of KNOWLEDGE COMPRESSOR, we return to our earlier example of the researcher attempting to reconcile conflicting findings about scientists’ reading behavior. In a typical setup, such as a combination of notes in MS Word and papers in Zotero, the researcher would need to track back the paper citations, open up each source PDF, and read again to locate the context - one claim at a time. Alternatively, they might need to create a new spreadsheet in MS Excel and record these contextual details for each claim, or modify their notes accordingly.

In contrast, if the claims were captured as grounded claims in KNOWLEDGE COMPRESSOR, they could explore different hypotheses by toggling the adjustable/expandable context segments for the claims, in a spatial layout, all at once. With a few scrolls from the context segments, they could quickly notice differences in contextual details that may explain the contradiction, such as whether the claim comes from observational studies or self-reported surveys, or variations in research expertise (number of years in conducting research). These contextual details drive them to the design decision of optimizing scientific reading through personalisation, such as highlighting different sections for different users.

CONCLUSION AND FUTURE WORK

We plan to conduct systematic evaluations of how creating grounded claims could change the cost structure, including controlled experiments against conventional annotation, measuring the cognitive load, response time, and number of behavioral signals (e.g. clicks, re-writes, re-selections). We also plan to explore other techniques to easing the cost structure, including automatic identification of specific candidate contexts (e.g., participant details, measures, parameters) to attach to claims, and intelligent text input [12] to further speed up and ease the process of redescribing key claims.

ACKNOWLEDGEMENTS

We thank Katrina Fenlon, Susan Winter, Joseph Chang, and Ge Gao for helpful discussions.

REFERENCES

- [1] Mark S. Ackerman, Juri Dachtera, Volkmar Pipek, and Volker Wulf. 2013. Sharing Knowledge and Expertise: The CSCW View of Knowledge Management. *Computer Supported Cooperative Work (CSCW)* 22, 4-6 (2013), 531–573.
- [2] Mark S Ackerman and Christine Halverson. 2004. Organizational memory as objects, processes, and trajectories: An examination of organizational memory in use. *Computer Supported Cooperative Work (CSCW)* 13, 2 (2004), 155–189.
- [3] Anuja Cabraal. 2012. Why Use NVivo for Your Literature Review? (Aug. 2012). <https://anujacabraal.wordpress.com/2012/08/01/why-use-nvivo-for-your-literature-review/>
- [4] Joseph Chee Chang, Nathan Hahn, and Aniket Kittur. 2016. Supporting Mobile Sensemaking Through Intentionally Uncertain Highlighting. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST'16)*. ACM, New York, NY, USA, 61–68.
- [5] Der-Thanq Victor Chen, Yu-Mei Wang, and Wei Ching Lee. 2016. Challenges confronting beginning researchers in conducting literature reviews. *Studies in Continuing Education* 38, 1 (2016), 47–60.
- [6] Tim Clark, Paolo N Ciccicarese, and Carole A Goble. 2014. Micropublications: a semantic model for claims, evidence, arguments and annotations in biomedical communications. *Journal of biomedical semantics* 5, 1 (2014), 28.
- [7] James A Evans and Jacob G Foster. 2011. Metaknowledge. *Science* 331, 6018 (2011), 721–725.
- [8] Alexander Hars. 2001. Designing scientific knowledge infrastructures: the contribution of epistemology. *Information Systems Frontiers* 3, 1 (2001), 63–73.
- [9] D. S. McCrickard. 2012. *Making Claims: The Claim as a Knowledge Design, Capture, and Sharing Tool in HCI*. Morgan Claypool.
- [10] Carole L. Palmer. 2001. *Work at the Boundaries of Science*. Springer Netherlands.
- [11] Antoine Ponsard, Francisco Escalona, and Tamara Munzner. 2016. PaperQuest: A Visualization Tool to Support Literature Review. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. ACM, New York, NY, USA, 2264–2271.
- [12] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models Are Unsupervised Multitask Learners. *OpenAI Blog* 1, 8 (2019).
- [13] Daniel M Russell, Malcolm Slaney, Yan Qu, and Mave Houston. 2006. Being Literate with Large Document Collections: Observational Studies and Cost Structure Tradeoffs. In *System Sciences, 2006. HICSS'06. Proceedings of the 39th Annual Hawaii International Conference On*, Vol. 3. IEEE, 55–55.
- [14] Daniel M. Russell, Mark J. Stefik, Peter Pirolli, and Stuart K. Card. 1993. The Cost Structure of Sensemaking. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems (CHI '93)*. ACM, New York, NY, USA, 269–276.
- [15] 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. *Computer Supported Cooperative Work (CSCW)* 8, 4 (1999), 333–352.
- [16] Stephen E Toulmin. 2003. *The uses of argument*. Cambridge university press.
- [17] Victoria Uren, Simon Buckingham Shum, Michelle Bachler, and Gangmin Li. 2006. Sensemaking Tools for Understanding Research Literatures: Design, Implementation and User Evaluation. *International Journal of Human-Computer Studies* 64, 5 (May 2006), 420–445.
- [18] Matt Willis, Sarika Sharma, Jaime Snyder, Michelle Brown, Carsten Østerlund, and Steve Sawyer. 2014. Documents and Distributed Scientific Collaboration. In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW Companion '14)*. ACM, New York, NY, USA, 257–260.
- [19] S. Wuchty, B. F. Jones, and B. Uzzi. 2007. The Increasing Dominance of Teams in Production of Knowledge. *Science* 316,

- 5827 (May 2007), 1036–1039.
- [20] Alyson L. Young and Wayne G. Lutters. 2017. Infrastructuring for Cross-Disciplinary Synthetic Science: Meta-Study Research in Land System Science. *Computer Supported Cooperative Work (CSCW)* 26, 1 (April 2017), 165–203.
 - [21] Xiaolong Zhang, Yan Qu, C. Lee Giles, and Piyou Song. 2008. CiteSense: Supporting Sensemaking of Research Literature. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 677–680.