CitePeek: Contextual Citation Exploration Within Research Papers

Yasaman A.Basti
Science Department
Ontario tech University
Oshawa, Ontario, Canada
Yasaman.akhoundzadehbasti@ontariotechu.net

Heidar Davoudi
Ontario Tech University
Oshawa, Ontario, Canada
heidar.davoudi@ontariotechu.ca

Ali Neshati
Ontario Tech University
Oshawa, Ontario, Canada
ali.neshati@ontariotechu.ca

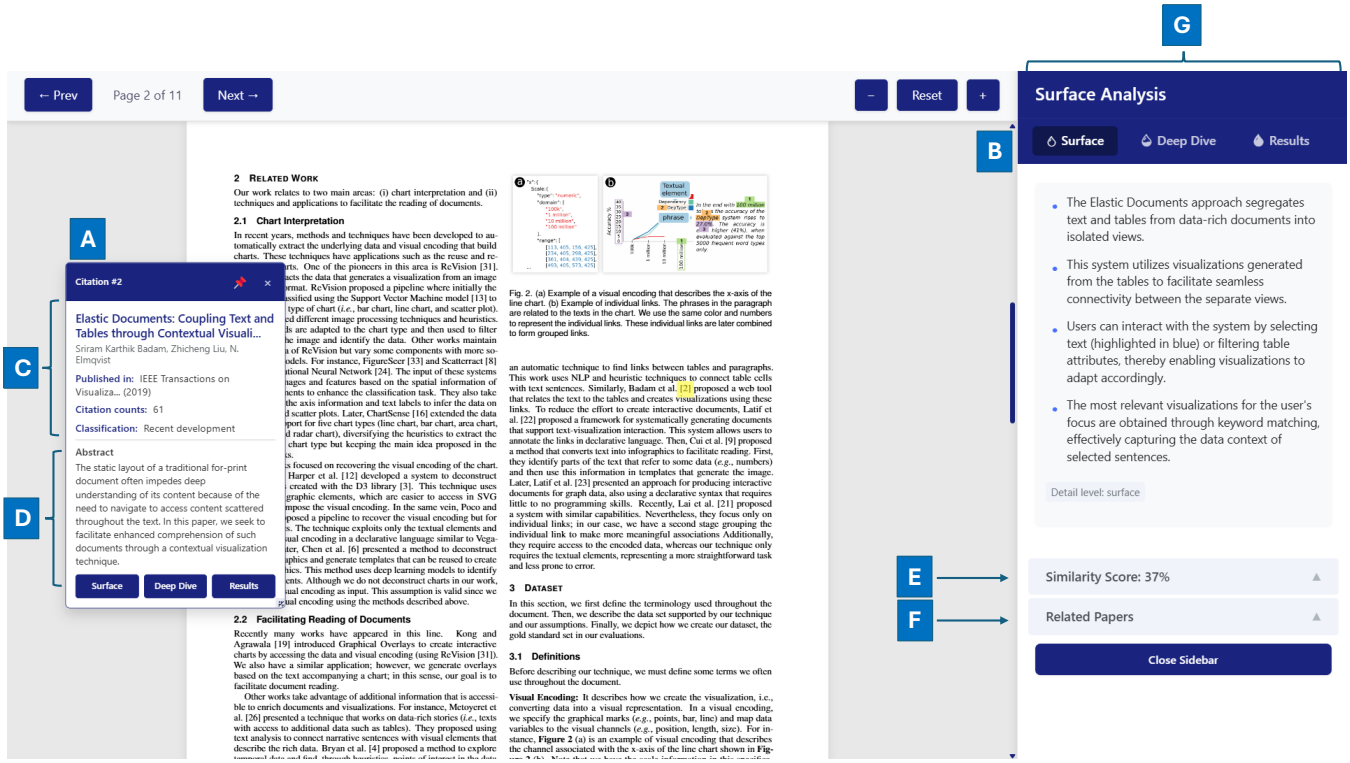


Figure 1: The main system. (A) The pop-up window presents the metadata of the selected cited paper. (B) The sidebar includes multiple features. (C) The metadata section displays the details of the publications and authors. (D) Abstract summary providing a quick overview of the cited paper's content. (E) Similarity score indicator (37%) showing the contextual relevance between the current reading context and the cited document. The system uses this score to highlight the most relevant sentences. (F) The Related Papers section shows connections within the citation network. (G) Navigation controls for accessing different information levels: Surface, Deep Dive, and Results.

Abstract

Academic reading requires frequent citation engagement, but traditional PDF readers do not provide contextual support to understand cited works. Current tools require readers to navigate away from the main text to access citation details. This disrupts the reading flow, highlighting a significant gap in seamless citation exploration. We present CitePeek, an interactive PDF assistant that transforms citations in a research paper into interactive elements. Clicking a citation reveals a pop-up with metadata and an abstract for quick critical insights. A comprehensive sidebar offers Surface and Deep Dive information levels and cited paper results. Surface provides key points, Deep Dive gives narrative explanations with visuals, and Results highlight significant experimental outcomes with data visualizations. CitePeek calculates contextual relevance through similarity scoring, presenting relevant paragraphs and relevance-based highlighting in original PDFs. CitePeek identifies related

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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CUI '25, Waterloo, ON, Canada
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1527-3/25/07
https://doi.org/10.1145/3719160.3737639

papers within the reference network, delivering essential citation details directly within the reading environment.

CCS Concepts

• **Human-centered computing** → **Interactive systems and tools**; **User interface management systems**.

Keywords

Interactive Reading, PDF Assistant Tool, Information Retrieval, Reference Navigation, Metadata Retrieval

ACM Reference Format:

Yasaman A.Basti, Heidar Davoudi, and Ali Neshati. 2025. CitePeek: Contextual Citation Exploration Within Research Papers. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces (CUI '25), July 08–10, 2025, Waterloo, ON, Canada.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3719160.3737639

1 Introduction

Understanding why citations are included and how they support the surrounding text is a critical yet unsupported aspect of academic reading. Our research is motivated by the need to reduce the cognitive burden associated with citation exploration and to provide readers with immediate, contextual insights into cited works. When engaging with academic literature, researchers often struggle to comprehend citations without interrupting their reading flow. Switching between a paper and its references imposes a cognitive load, breaks contextual flow, and slows comprehension. Chang et al. [1] note that current tools fail to track previously seen citations, making it hard to identify overlapping references. Park et al. [11] emphasize how these disruptions can be particularly challenging for junior researchers lacking deep domain knowledge.

Our system, CitePeek, addresses this gap by using Natural Language Processing (NLP) and Application Programming Interface (API) integration to deliver context-aware information from cited papers directly within the PDF interface (Figure 1). It presents two levels of detail: a surface view offering 3–5 bullet-point takeaways for quick reference (Figure 1.B), and a deep dive view that includes extended content such as methods, frameworks, and figures when available (Figure 3). A separate results panel highlights findings and visualizations to support faster comprehension.

CitePeek also constructs a scholarly network map based on reference relationships (Figure 6), helping users visualize clusters and key influences. Also, similarity scoring highlights the most relevant paragraphs from cited papers, with a red-to-green gradient for quick skimming (Figure 5). Users can access the original PDF with guided highlighting. All citation content appears in a sidebar, accompanied by metadata and summaries in Pop-ups that can be pinned or resized based on user interest. This design helps users stay focused without switching contexts.

Our main contribution is an interactive PDF system that retrieves and presents multilevel, contextually relevant citation content. By combining citation analysis, visual guidance, and integrated summaries, CitePeek improves the efficiency and depth of academic reading.

2 Background and Literature Review

Prior research on academic reading primarily focuses on two areas: interactive reading environments and citation context comprehension. Despite progress, most systems provide isolated features and do not support seamless navigation between papers and references without breaking the reading flow.

Researchers have developed various tools to support more effective academic reading. These tools typically focus on helping readers interact with one paper at a time and understand its citations. However, most of these tools do not make it easy to move between a paper and the papers it cites. In the following, we describe prior efforts to address these challenges and identify remaining gaps.

2.1 Interactive and Contextual Reading Environments

Research in interactive reading environments has explored various strategies to enhance digital reading comprehension. Early efforts focused on improving interaction with individual digital documents. Systems such as XLibris [12] and Beyond Paper [14] introduced annotation tools that allowed users to highlight, organize, and manage digital text effectively. These tools marked an essential shift in academic reading, from static to interactive documents, but primarily focused on engagement with a single document without addressing the challenges of navigating through references. Later systems like LiquidText [15] and Passages [4] built on this foundation by enabling multi-touch interaction and flexible manipulation of snippets, allowing users to compare and organize ideas more fluidly. Despite these advances, the focus remained on improving interaction with a single document rather than integrating citation navigation. To address cross-document exploration, systems like Threddy [6] enable readers to collect and organize "threads" of related literature while reading. Although this helps users follow multiple references on a theme without losing their place, it requires manual curation. It does not automatically provide additional context beyond what is mentioned in the current paper. Similarly, Synergi [8] supports synthesis in multiple papers using AI and graph algorithms, but its multi-paper focus makes it heavyweight for readers seeking lightweight, inline citation assistance during normal reading. While these tools improved document-level interaction, they lack automated, context-sensitive features that help readers interpret citations as they appear. Their reliance on user-driven curation or global document views often misses the local, sentence-level function and purpose of citations. Efforts like Scim [3] provide "intelligent skimming" through faceted highlights that guide the reader's attention to salient passages. However, these highlights are derived solely from the paper's own text and offer no information about citations or referenced works. Papeos [9] links segments of conference talk videos to corresponding passages in a paper's PDF, aiding comprehension of the main paper but not citation handling. Moreover, the system depends on the availability of author talk videos, which are not always provided. A fundamental limitation across these systems is their focus on interaction with individual documents without addressing the need to seamlessly navigate between interconnected scholarly works. They also tend to treat citations as static or secondary without offering deeper semantic context or explanation. In contrast, CitePeek offers

a more integrated reading experience by dynamically identifying and presenting citation details that are semantically aligned with the reader's current focus.

2.2 Citation Context and Comprehension

Several systems aim to improve citation comprehension by offering contextual insights, but important limitations remain. These tools reflect the growing interest in helping readers understand how citations contribute to scholarly arguments, yet they often stop short of delivering the depth and relevance needed during real-time reading.

CiteRead [13] enhances documents with commentary from subsequent papers to show how prior work has been cited, helping readers assess scholarly impact. However, it focuses only on incoming citations and does not support understanding of the references within the document, an essential part of academic reading. Its delivery is also non-personalized, offering the same experience regardless of user background.

CiteSee [1] introduces a personalized touch by providing visual citation cues based on prior reading behaviour. While useful for returning users, it delivers only brief metadata or snippets, with limited value for newcomers lacking a rich reading history. ScholarPhi [5] enhances reading flow through on-demand definitions but does not extend support to citation comprehension, leaving references as non-interactive elements.

Other tools like Action Science Explorer [2] and Relatedly [10] enable exploration of citation networks across document collections, but their macro-level focus often pulls readers away from the main text and disrupts attention. Similarly, ComLittee [7] supports author-based discovery but lacks tools for understanding citations in context within a paper. Commercial platforms, including Google Scholar's PDF reader[1] and Elsevier's ScienceDirect[2], offer limited previews of cited content. Though Google Scholar shows parts of cited papers, these systems generally lack deep contextual insights, visual support, or recursive navigation features, making it difficult for users to evaluate citation relevance without interrupting their reading flow.

Compared to these approaches, CitePeek applies a more context-sensitive and document-aware method for citation understanding. Across these systems, a common problem is that they mostly rely on surface-level information like citation titles, counts, or short summaries. They do not help the reader understand why the citation is there or what part of the cited paper actually supports the sentence. In contrast, CitePeek goes deeper. It not only shows metadata; it also processes the full text of the cited paper to identify the most relevant parts based on the reader's immediate context. Using methods like semantic similarity, keyword matching, term frequency, and section importance, it extracts the most relevant segments and presents them inline, allowing readers to understand the citation without leaving the page or opening another document. In summary, current academic reading tools disrupt comprehension by forcing readers to switch contexts when exploring citations. They often lack depth, visual support, smooth navigation, and multi-level citation tracing. We propose an integrated platform that embeds contextual insights and visual cues into the reading interface, enabling seamless engagement with cited content and a better understanding of scholarly relationships.

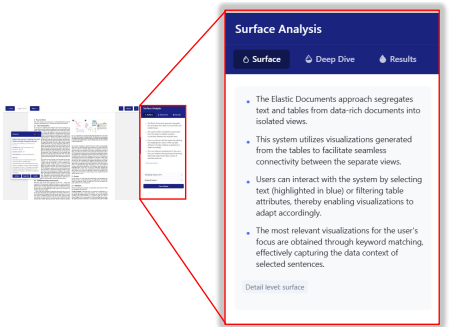


Figure 2: Clicking on the surface level will present bullet points of relevant information at a quick glance without requiring navigation away from the current document.

3 System Overview and Implementation

In this section, we present the design and implementation of an interactive PDF system that enhances academic reading through contextual citation analysis and progressive information disclosure. Our AI-powered assistant is a web-based application that improves navigation and comprehension by embedding citation-related features directly into the PDF interface. By minimizing the need to leave the document to access external references, the system offers a more seamless and focused reading experience.

CitePeek uses a client-server architecture combining a React-based front-end with a Django-powered Python backend to support these functionalities. The system presents a streamlined interface for academic PDF uploads, rendered through an enhanced PDF.js [3] implementation that supports interactive features and seamless integration with our analysis tools. When a user uploads a paper, the system automatically downloads all cited references. If a reference cannot be accessed online through Semantic Scholar or DOI links, our analyzer retrieves the PDF from a pre-populated local storage directory, which contains previously downloaded or user-uploaded cited papers. This ensures uninterrupted access and smooth interaction even when APIs are unavailable or network access is limited. The local storage functions as a fallback repository of reference PDFs that the system indexes and analyzes, just like online sources.

The backend uses PyMuPDF [4] to extract text and analyze the structural elements of the document. It specifically identifies sentences that immediately precede inline citations to extract contextual cues, which are then used to search and rank relevant information within the corresponding referenced papers. The backend system's core functionality includes a sophisticated multi-stage

[1]Google Scholar is an academic search engine that provides limited previews of cited documents. https://scholar.google.com
[2]ScienceDirect is a scientific article database and PDF viewer provided by Elsevier. https://www.sciencedirect.com
[3]PDF.js is an open-source JavaScript library developed by Mozilla that renders PDF documents in the browser using HTML5 and JavaScript. More information is available at: https://mozilla.github.io/pdf.js/
[4]PyMuPDF is a Python binding for the MuPDF library, allowing access to PDF files for text extraction, rendering, and document structure analysis. More information is available at: https://pymupdf.readthedocs.io/

content-ranking mechanism designed to identify the most relevant information within cited documents. After extracting sentences that immediately precede inline citations, the backend uses this context to initiate a sentence-ranking algorithm that integrates four weighted metrics:

- **Semantic Similarity**: This measures how closely the content in the cited paper matches the context around the citation in the main document. We use SentenceTransformer[5] models to generate vector representations of sentences and compute their cosine similarity. Higher scores indicate more semantically aligned content, allowing us to prioritize relevant passages that align with the reader's current focus.
- **Term Overlap**: This metric evaluates the number of shared terms between the citation context and candidate sentences in the cited paper. While simpler than semantic similarity, it captures direct textual matches that often indicate relevance, especially for domain-specific keywords.
- **TF-IDF Similarity**: TF-IDF (Term Frequency-Inverse Document Frequency) weighs the importance of terms based on how unique they are within and across documents. This metric emphasizes distinctive terms in the citation context and helps highlight unique contributions or rare terminology from the cited paper.
- **Section-Based Weighting**: Not all parts of a paper are equally informative. This metric assigns greater weight to sentences found in critical sections, such as Methodology and Results, under the assumption that these sections are more likely to contain substantive content that justifies the citation.

High-scoring sentences are then contextually expanded by extracting surrounding paragraphs using intelligent boundary detection techniques that preserve the coherence of the presented information. This enriched content is passed through a two-stage summarization pipeline: first distilled using BART [6] summarization tailored to the selected analysis level, and then refined using GPT-based[7] prompting to generate structured, informative, and reader-friendly outputs.

The application's key components and features include:

3.1 Interactive Inline Citations

When readers encounter citations, our system transforms static references into interactive elements that provide immediate context (Figure 1). Upon selection, a concise popup appears adjacent to the citation (Figure 1.A), displaying metadata of the cited paper, including title, authors, publication year, citation count, and venue (Figure 1.C), and a summary of the cited work's abstract, and options for further exploration (Figure 1.D). This contextual approach, fetched from Semantic Scholar's API [8], allowing readers to gain essential information without losing their place in the main document. The popup's positioning adapts to screen space and reading position to minimize disruption to the reading flow. At the bottom of the popup, there are three options: surface, deep dive, and result. Clicking on them will open the sidebar for further information, which is explained in the following section.

In comparison to ChatGPT, while large language models like ChatGPT can summarize and answer questions about the text, CitePeek is specifically designed for academic reading and citation navigation. Unlike ChatGPT, which operates as a general-purpose assistant without document structure awareness, CitePeek integrates directly into the PDF interface and uses citation context to retrieve relevant segments from the cited papers themselves. Its features are not prompted manually but are automatically triggered by reader interaction with inline citations. Thus, rather than replacing ChatGPT, CitePeek complements it by offering a task-specific, citation-aware experience embedded within the reading workflow.

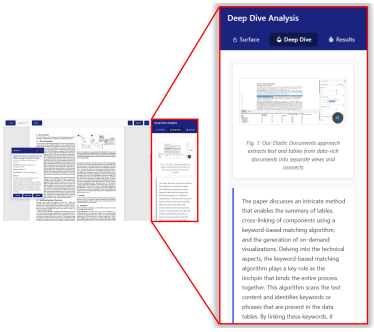


Figure 3: Clicking on the Deep Dive level will present a thorough analysis of the paper, representing relevant information and the key contributions of the cited paper alongside the main figure or graph of the paper.

3.2 Contextual Information Sidebar

Upon selecting a specific citation, a detailed sidebar appears, offering in-depth analyses at three distinct and progressively detailed levels:

Surface Level extracts and displays 3-5 bullet points containing the most relevant sentences from the cited paper that directly support the specific context where the citation appears (Figure 2). Relevant sentences are identified using hybrid ranking algorithms and then summarized using the BART model, and then refined using the GPT-4 API, which transforms the summarized content into clearer, more coherent outputs aligned with academic reading needs. These brief insights facilitate rapid understanding without extensive reading.

Deep Dive Level Extracts relevant information from the cited paper and provides an in-depth exploration by analyzing content that is directly related to the citation's context in the main document (Figure 3). This mode emphasizes methodological details and

[5]SentenceTransformer is a Python library for generating semantically meaningful sentence embeddings using pre-trained models like BERT. See: https://www.sbert.net/
[6]BART (Bidirectional and Auto-Regressive Transformers) is a sequence-to-sequence transformer model developed by Facebook AI, used for tasks such as summarization and text generation.
[7]GPT-based prompting refers to the use of OpenAI's Generative Pre-trained Transformer models for generating human-like text based on provided instructions.
[8]Semantic Scholar is a free, AI-powered research tool for scientific literature maintained by the Allen Institute for AI. It provides access to over 200 million academic papers across various disciplines with enhanced search capabilities and citation information.

processes larger content segments to construct a comprehensive understanding. High-ranked sentences are grouped into coherent paragraphs and summarized using BART with extended output parameters. These summaries are then refined using GPT-4 to produce detailed, readable explanations enriched with examples and clarifications. Rather than randomly selecting paper sections, the system prioritizes information directly connected to how and why the paper was cited in the main text. This targeted approach ensures that readers receive explanations that directly address the citation's purpose and relevance to the current discussion. The system transforms these relevant elements into coherent, contextually appropriate explanations that preserve technical accuracy while enhancing accessibility. To further enhance comprehension, the system leverages PyMuPDF to extract relevant and important figures, tables, and diagrams from the cited paper. It prioritizes the main figure, the most cited figure in the text, that aligns with the citation context, integrating them directly into the explanation to provide visual reinforcement of key concepts. This focused presentation enables readers to thoroughly understand why the paper was cited and how its contributions specifically relate to the current document, all without disrupting their reading flow or requiring navigation to external sources (Figure 3).

Results Level highlights precise experimental outcomes, critical metrics, and statistically significant findings. Essential result visuals, such as charts and graphs, are included to clearly present the quantitative data reported in the original paper. These visuals are organized logically, styled to match the data type, and designed for clarity by removing unnecessary elements. When available, statistical details like confidence intervals and p-values are included. All visuals follow consistent formatting and strictly reflect the reported findings without additional interpretation. This presentation is achieved through a custom GPT-4-powered implementation that organizes findings according to scientific reporting standards (Figure 4).

To implement this, the system explicitly targets experimental outcomes by employing dedicated extraction techniques to isolate results sections. It uses both pattern recognition, identifying headings like "Results" or "Findings", and statistical marker analysis, detecting terms such as "p-value," "significant," or "participants." Once relevant segments are identified, they undergo statistical highlight identification to capture key metrics and indicators of significance. The system then extracts visuals, such as graphs and charts, associated with these outcomes. These are refined using GPT-4 to emphasize quantitative data, statistical relevance, and comparative insights while deliberately omitting methodological or contextual information not directly tied to the results. This layered approach ensures that the Results Level presents a focused, scientifically clear, and trustworthy summary aligned with academic reporting standards (Figure 4).

3.3 Semantic Similarity and Contextual Highlighting

Another feature in the sidebar is the Relevance Highlighter, which identifies and displays the most relevant paragraph from the cited paper, with the most related sentence visually emphasized (Figure 5). To achieve this, the system calculates the semantic similarity

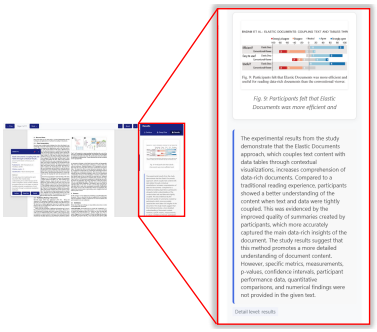


Figure 4: Clicking on the result level will present a summary of experimental outcomes, critical metrics, and statistically significant findings with the proper chart or graph representing the outcome of the cited paper.

between the user's current reading context and the content of the cited paper. Based on this analysis, it ranks and highlights the most contextually aligned text. The interface then shows the top-ranked paragraph with key sentences highlighted for quick comprehension (Figure 5). By clicking "View Paper with Relevance Highlights," users are directed to the original PDF, where the content is color-coded from green (most relevant) to red (least relevant). Sentences scoring above 85% similarity are marked green, indicating strong contextual relevance. Those scoring between 75% and 85% are highlighted in orange, and sentences with scores between 60% and 75% appear in red. This allows for focused reading of the most important information in the cited paper while preserving the original formatting of the paper (Figure 5.B). This color-coded system provides flexible navigation and consistently guides the reader's attention to the most relevant parts of the reference material.

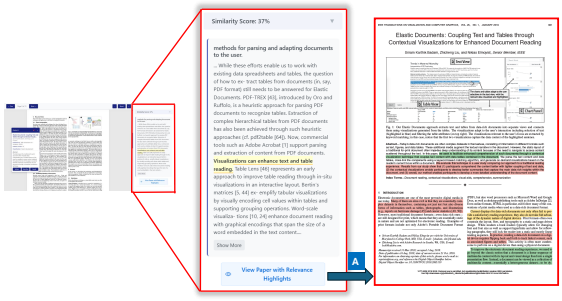


Figure 5: similarity score retrieving the most related paragraph containing the most relevant sentences highlighted with a link to the original PDF with relevant information highlighted from red to green.

3.4 Related Papers Network

CitePeek examines the content and metadata of all cited papers to reveal meaningful relationships, such as shared techniques, overlapping research areas, or similar contributions (Figure 6). To identify

these connections, the system automatically analyzes all the references cited in the paper to find those that are related to each other. This comparison is based on shared contributions, research topics, or other semantic similarities derived from metadata such as titles, abstracts, and summaries. This data is sent to a GPT-4 model using a prompt designed to extract shared research topics and methodological similarities. The model returns a brief relationship summary and a classification of the connection. These results are parsed and displayed in the sidebar, where users can see related reference numbers and relationship descriptions. When users click on any of the cited paper numbers (e.g., 3, 4, or 22), the system displays the paper's title along with a brief metadata summary, including the authors and abstract, to support a quick evaluation of its relevance. This feature helps readers quickly understand how cited works are connected and which clusters of references represent a coherent research thread. By surfacing these related papers, the system enables users to grasp the state-of-the-art around a specific topic without manually reviewing every citation, streamlining the process of identifying which references are most aligned with their interests (Figure 6).

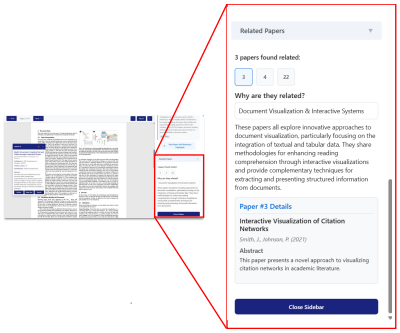


Figure 6: The related papers section presents the cited papers that are related to each other; clicking on the mentioned papers will present a brief metadata of the paper.

4 Limitations and Future Work

While our system offers meaningful improvements for academic reading and citation exploration, several limitations remain. The current implementation depends on internet connectivity for metadata retrieval via external APIs, which can be problematic in low-connectivity environments. When references are unavailable online, users must rely on local downloads, limiting real-time updates and overall convenience.

Citation detection accuracy is another challenge, especially with various citation formats across disciplines. The system performs best with well-structured citations but may struggle with older works, inconsistent formatting, or unusual characters, leading to missed or misparsed citations that affect citation network accuracy. One other limitation of CitePeek is that certain features, such as the related papers network, require analyzing the full reference list and paper content, making the process computationally intensive and slower for large documents. Our prototype demonstrates promising functionality, but empirical validation is still needed. As the system remains under development, our efforts have focused on core features rather than extensive testing. To address this, we plan to conduct a structured user study comparing CitePeek with traditional PDF readers and basic citation tools. This study will evaluate the system's effectiveness in enhancing citation comprehension and reading efficiency through quantitative metrics (e.g., task completion time, comprehension accuracy, citation tracing speed) and qualitative feedback (e.g., perceived usefulness, trust, and usability). These evaluations will validate the system's core claims and inform future refinements based on real user needs. In addition, to enhance citation detection, we plan to combine multiple approaches to support a broader range of citation styles. We will explore rule-based parsing techniques and pattern recognition using regular expressions tailored to known citation structures. By integrating heuristic rules and configurable templates that reflect domain-specific citation conventions, we aim to improve the system's adaptability and robustness across academic disciplines.

5 Conclusion

The Interactive PDF Assistant illustrates the potential of conversational interface principles to enhance academic reading by enabling in-context access to relevant information, supporting better comprehension, and reducing cognitive load. These early insights suggest this approach can enhance academic reading tools. The progressive disclosure model helps balance content richness and usability, enabling readers to decide how much detail they wish to explore. As digital reading evolves, our work highlights the value of contextual citation analysis in supporting more effective knowledge acquisition.

References

[1] Marti Chang, Kyle Lo, Daniel S. Weld, and Marti A. Hearst. 2023. CiteSee: Augmenting Citations in Scientific Papers with Persistent and Personalized Historical Context. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 543, 15 pages. doi:10.1145/3544548.3581168

[2] Cody Dunne, Ben Shneiderman, Robert Gove, Judith Klavans, and Bonnie Dorr. 2012. Rapid Understanding of Scientific Paper Collections: Integrating Statistics, Text Analytics, and Visualization. In *Journal of the American Society for Information Science and Technology*, Vol. 63. American Society for Information Science and Technology, 2351–2369. doi:10.1002/asi.22652

[3] Raymond Fok, Andrew Head, Ruijia Huang, Kyle Lo, Tris Worsley, Se Hwang Oh, Daniel S. Weld, and Marti A. Hearst. 2023. Scim: Intelligent Skimming Support for Scientific Papers. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. ACM, 268–286. doi:10.1145/3581641.3584063

[4] Donghan Han, Yasha Yang, Tianyi Hu, Tsukasa Kono, Takanori Fujiwara, and Kwan-Liu Ma. 2022. Passages: A Document Analysis Workflow Enabling Text-Visualization Integration and Interaction. In *2022 IEEE Visualization and Visual Analytics Conference (VIS)*. IEEE, 131–135. doi:10.1109/VIS54862.2022.00031

[5] Andrew Head, Kyle Lo, Daniel Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. 2021. Augmenting Scientific Papers with Just-in-Time, Position-Sensitive Definitions of Terms and Symbols. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, 1–18. doi:10.1145/3411764.3445648

[6] Daniel Kang, Andrew Head, Marti A. Hearst, and Jane Kim. 2022. Threddy: An Interactive System for Personalized Thread-Based Exploration and Organization of Scientific Literature. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. ACM, 1–14. doi:10.1145/3526113.3545663

[7] Daniel Kang, Taesung Lee, Kyle Lo, Daniel S. Weld, and Marti A. Hearst. 2023. ComLittee: Literature Discovery with Personal Elected Author Committees. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, 1–16. doi:10.1145/3544548.3581307

[8] Daniel Kang, Eddie Yuan, Andrew Head, and Marti A. Hearst. 2023. Synergi: A Mixed-Initiative System for Thread-Based Synthesis of Scientific Literature. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, 1–14. doi:10.1145/3586183.3606764

[9] Jane Kim, K. Arvind Manasvini, Hwiyeon Cho, Hyun Cheol Kum, Woosuk Park, Hee-Tae Suk, and Bongwon Suh. 2023. Papeos: Augmenting Research Papers with Talk Videos. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, 1–16. doi:10.1145/3586183.3606743

[10] Saraschandra Palani, Joseph Robinson, Taesung Lee, and Marti A. Hearst. 2023. Relatedly: Scaffolding Literature Reviews with Existing Related Work Sections. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, 1–18. doi:10.1145/3544548.3581060

[11] Jeongwook Park, Hwiyeon Cho, and Bongwon Suh. 2023. QuickRef: Should I Read Cited Papers for Understanding This Paper?. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 240, 7 pages. doi:10.1145/3544549.3585733

[12] Morgan N. Price, Gene Golovchinsky, and Bill N. Schilit. 1998. Linking by Inking: Trailblazing in a Paper-Like Hypertext. In *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia*. ACM, 30–39. doi:10.1145/276627.276632

[13] Napol Rachatasumrit, Yufang He, Katherine Tate, Adam Perer, Haotian Li, and Michael L. Littman. 2022. CiteRead: Integrating Localized Citation Contexts into Scientific Paper Reading. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*. ACM, 17–29. doi:10.1145/3490099.3511106

[14] Bill N. Schilit, Gene Golovchinsky, and Morgan N. Price. 1998. Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 249–256. doi:10.1145/274644.274680

[15] Craig S. Tashman and W. Keith Edwards. 2011. LiquidText: A Flexible, Multitouch Environment to Support Active Reading. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3285–3294. doi:10.1145/1978942.1979430