

# Science and Technology for Augmenting Reading (STAR)

Tal August  
taugust@illinois.edu  
UIUC  
Champaign, IL, USA

Andrew Head  
head@seas.upenn.edu  
University of Pennsylvania  
Philadelphia, PA, USA

Alexa Siu  
asiu@adobe.com  
Adobe Research  
San Jose, CA, USA

Elena L. Glassman  
glassman@seas.harvard.edu  
Harvard University  
Boston, MA, USA

Jonathan K. Kummerfeld  
jonathan.kummerfeld@sydney.edu.au  
University of Sydney  
Sydney, New South Wales, Australia

Joseph Chee Chang  
josephc@allenai.org  
Allen Institute for AI (Ai2)  
Seattle, WA, USA

Lucy Lu Wang  
lucylw@uw.edu  
University of Washington  
Seattle, WA, USA

Marti A. Hearst  
hearst@berkeley.edu  
UC Berkeley  
Berkeley, CA, USA

## Abstract

The landscape of technology for consuming information is changing rapidly. One mode of information consumption, reading, stands to see profound changes due to its ubiquity and frequency as a cognitive task. Better reading technology could transform texts on demand so that they are easier to read, surface hard-to-find information, support synthesis, and better engage readers. The possibilities have been considerably expanded with the maturation of AI. The purpose of this workshop is to provide a platform for the growing cohort of HCI and AI researchers interested in augmented reading interfaces to define high-impact areas of development and standards of success. This platform will arise from two components of our workshop. The first is a brief and engaging format of introductions among community members through lightning talks. The second is a set of affinity group activities that will identify fresh opportunities for augmenting reading against the backdrop of reading theories, evaluation practices, and emerging technology.

## CCS Concepts

• **Human-centered computing** → *HCI design and evaluation methods*; **Interactive systems and tools**; *Natural language interfaces*.

## Keywords

reading, reading augmentations, comprehension, sensemaking, discovery, cognitive load

## ACM Reference Format:

Tal August, Andrew Head, Alexa Siu, Elena L. Glassman, Jonathan K. Kummerfeld, Joseph Chee Chang, Lucy Lu Wang, and Marti A. Hearst. 2026. Science and Technology for Augmenting Reading (STAR). In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3772363.3778695>



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI EA '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-XXXX-X/2018/06  
<https://doi.org/10.1145/3772363.3778695>

## 1 Motivation

Reading is one of the most ubiquitous modes for consuming information. You are learning about this workshop by reading this textual proposal. Researchers learn of the latest advances in their field by reading the scientific record. Professionals coordinate on work products consulting and reading reams of documents and text messages. Students look up help on homework problems by scanning textbooks. Recently, review and evaluation of AI systems has predominately relied on reading. We are reading all the time, and the information we consume in text is mediated by the interfaces that present that text whether that is a book, paper, or chat log.

Recent changes in the technological landscape are significantly changing the reading experience. AI has introduced many new possibilities for interfaces to augment or transform text to be more rapidly scanned, navigated, understood, and compared to other texts. In fact, many of these possibilities are becoming realities, as mainstream reading interfaces are augmented with new capabilities. Adobe, Apple, the Allen Institute for AI, and Google have all recently released AI features for interacting with one or multiple documents [1, 20, 33, 42]. The outputs of modern agents—including new deep research modes [2, 21, 38]—come populated with external links into other documents. And augmented reading interfaces have been deployed at scale for scientific reading, providing new facilities for finding key information and keeping track of citations [13, 33].

At the same time, the HCI and AI communities have picked up the pace in devising innovative kinds of augmentations that go beyond the the existing capabilities of deployed interfaces. These augmentations can enable more flexible skimming [23], simplify complex passages [7], synchronize complementary text and video [32] and create accompanying audio on-the-fly that serve as cognitive supports [12], among many other envisioned possibilities (e.g., [16, 28, 34, 35, 41, 46]). As AI becomes increasingly used, new needs for reading emerge, such as tooling that facilitates the auditing [22] (i.e., reading, verifying) of generative content [16, 29].

As the community around augmented reading broadens and as possibilities continue to unfold, it is the purpose of this workshop to set up our community to drive innovation in a productive, desirable,

and responsible way. It is a critical time for our community to share knowledge for two reasons:

**Synchronization.** While some researchers in augmented reading have focused on the area for a long time, many more are relative newcomers. Even among those experienced, many have recently been emboldened to shift focus to new problems. Our workshop will equip researchers with powerful ways of thinking about the problems that need to be solved in augmented reading by refreshing our knowledge of key theories of reading and assessing how these theories may, or may not, apply to new augmentations. Furthermore, many of the ideas in augmented reading may or may not transfer across different kinds of texts (e.g., scientific papers, news articles, novels, or chat logs). Our workshop will help participants become aware of what others have tried and what has worked.

**Ideation.** Against the backdrop of theories and recent AI capabilities, we will identify ripe directions for innovation in augmented reading. Through structured activities, we will strengthen these ideas so that they can serve as starting points for the next chapter of augmented reading research. Importantly, while AI augmentations will be a major theme we discuss, we will also make ample room to consider non-AI augmentation patterns which have been essential to prior HCI progress (e.g., [8, 10, 40, 45]), acknowledging that comprehensive reading support today will likely require advances in both AI and non-AI approaches.

## 2 Workshop length

The workshop will be a short (90-minute) workshop. Following the CHI 2026 change in workshop structure, we will be offering only in-person attendance for the workshop.

## 3 Organizers

This workshop's organizing team is comprised of leading researchers in SIGCHI and adjacent communities with substantial records designing and evaluating innovative augmented reading experiences. Organizers include heads of academic labs and researchers at industry labs, with many from HCI and some from NLP. Many organizers are responsible for augmented reading features and texts that have seen use by massive audiences. Organizers represent a number of different foci on reading, including sensemaking, information extraction, comprehension, navigation, and both AI and non-AI approaches. In turn, the organizers' work represents many of the theoretical frameworks that we think will be necessary for structuring the next steps forward in augmented reading (Section 4).

**Tal August** is an assistant professor at the University of Illinois Urbana-Champaign in the Siebel School of Computing and Data Science. His research seeks to improve the communication of knowledge-intensive text to a wider audience by adapting language to those audiences. His group studies and develops reading and writing innovations for medical [7], legal [26, 27], scientific [3, 4, 6, 31] and instructional [5, 44] text.

**Andrew Head** is an assistant professor at the University of Pennsylvania. His group develops novel interactions to support reading and reasoning. He is particularly interested in developing tooling that lowers the barriers to understanding complex technical content and which integrates complementary information that is fragmented across a longer text. Head was one of the earliest

junior research leads on AI2's *Semantic Reader*, developing the first published HCI paper [24] for the project as an intern and a postdoctoral young investigator at AI2. His group has recently developed interfaces for augmented reading of scholarly texts [14], math [25, 47], code [48], and medical documents [30].

**Alexa Siu** is a research scientist at Adobe Research. Her research focuses on designing and evaluating human-AI interactions that improve productivity and accessibility in complex information tasks. Her recent work has developed new experiences to support sensemaking over document collections [15] and NLP techniques to improve document Q&A [43] and editing [36]. Alexa's research has contributed to product innovations such as Acrobat AI Assistant, which uses generative AI to enhance the document reading experience [42].

**Elena L. Glassman** is an assistant professor at Harvard University's Paulson School Of Engineering and Applied Sciences. Her research focuses on building AI-resilient interfaces [17] that augment skimming [23], writing, and close reading of semi-structured code and text at scale, e.g., OVERCODE [18], EXAMPLORE [19], PARALIB [49], and Positional Diction Clustering [16]. Her group specializes in leveraging theories of human cognition about how humans form mental models from varying concrete examples.

**Jonathan K. Kummerfeld** is a Senior Lecturer (i.e., research tenure-track Assistant Professor) in the School of Computer Science at the University of Sydney. His research focuses on interactions between people and NLP systems, developing more effective algorithms, workflows, and systems for collaboration. In HCI venues, this has included systems for more rapidly reading text, without relying on automatic summarization, and methods for highlighting similarities and differences across sets of language model outputs. In NLP venues, he has worked on dialogue disentanglement, to ease reading of complex multi-party conversations, including creating the standard benchmark for the task.

**Joseph Chee Chang** is a senior research scientist at the Allen Institute for AI (AI2). His research focuses on human-AI systems and interactions for sensemaking and information overload. His recent work has centered around supporting literature reviews and other research support tools, for example, tools for searching, discovering, reading, and synthesizing research papers. His article "*The Semantic Reader Project: Augmenting Scholarly Documents through AI-powered Interactive Reading Interfaces*." appeared in the October 2024 issue of Communications of the ACM [33].

**Lucy Lu Wang** is an assistant professor at the University of Washington Information School. Her work aims to improve the accessibility of scientific text and its usefulness for supporting downstream decision-making in high-expertise domains such as healthcare. She is also an adjunct professor in the UW Paul G. Allen School of Computer Science & Engineering, Department of Biomedical Informatics & Medical Education, and Department of Human-Centered Design & Engineering, and is a visiting research scientist at the Allen Institute for AI. She was a co-organizer for the SciNLP workshop (2020-2021) and the Scholarly Document Processing workshop (2021, 2022, 2024).

**Marti A. Hearst** is a professor at the UC Berkeley School of Information. Her research encompasses user interfaces with a focus on search, information visualization with a focus on text, computational linguistics, and educational technology. She is the author of

*Search User Interfaces*, the first academic book on that topic. She is a former President of the Association for Computational Linguistics, a member of the CHI Academy and the SIGIR Academy, an ACM Fellow, and an ACL Fellow. Hearst was one of the leads on the *SemanticReader* project on augmented reading interfaces.

## 4 Workshop Activities

The two primary goals of this workshop are to bring together attendees and facilitate the construction of better ideas for developing and evaluating augmented reading experiences. The workshop begins with opportunities for connection (lightning talks), and then dives into hands-on experiences in co-constructing knowledge through affinity group activities. Our tentative schedule is outlined in Table 1 and we elaborate on each activity below.

### 4.1 Lightning talks

The workshop will begin with an engaging format of introductions among participants. Each participant will be asked to prepare a single slide introduction to describe, broadly, “what they are thinking about right now in the augmented reading space.” Visuals (screenshots, diagrams, sketches) will be highly encouraged. To keep the energy high, we will follow an enforced 1-minute auto-advance on the shared slide deck.

### 4.2 Affinity group activities

We believe that the best way for our community to work effectively towards better reading interfaces is to share and collaboratively construct new knowledge surrounding reading interfaces. To this end, after the lightning talks, workshop participants will co-construct knowledge within small groups of roughly 6 people. Each group will focus on establishing direction around a timely, unsolved question in augmented reading. We will have prepared example questions within three categories based on our experiences and participant contributions, though participants can propose additional questions on the day of the event. Convergence on questions will be facilitated through a Miro board maintained and updated during lightning talks. Below we discuss our initial plans for the activities. Our own proposed questions will belong to the categories of novel interface ideas, emerging problems in reading, and evaluation techniques:

*Novel interface techniques.* New technologies enable new ways for readers to interact with text. What are exciting new interfaces, tools, or features that the community should be exploring? Groups working on questions from this category will engage in a speculative design activity focused on collaborative prototyping. They will take example texts and attempt to augment them in a novel way that they have hypothesized will improve the reading experience via paper prototyping or digital tools. Participants will be asked to report back with a visual of their design, their intuition following the activity as to whether their design would actually have the intended effect, and a list of challenges that must be overcome for the design to be successful. Visuals will be added to a shared slide deck, and a scribe will take note of presenters’ names and commentary on the strength and feasibility of the augmentations.

*Emerging problems in reading.* The technology of reading is rapidly changing and with it, our relationship with text. How will

new reading affordances alter, for better or worse, the way we read? Groups working on questions from this category will do speculative design activities where they take a recent or anticipated reading technology and explore its possible consequences to reading. They will be encouraged to leverage reading theories (see “Theoretical primers” below) to identify detrimental effects to reading that are otherwise easy to overlook. Participants will be asked to report back on the emerging technology they chose and what effect it might have on how we read, including a list of challenges to achieving the best possible future with the technology.

*New evaluation techniques.* As more augmented reading interfaces are deployed, how should we evaluate their impact? Groups working on questions in this category will articulate a goal of evaluation, take inventory of known methods that address that goal, and ideate gaps where new methods are needed. Groups may focus on questions around how to more reliably assess comprehension, or how to better evaluate reading interactions in the wild. Participants will be asked share this review in a slide-deck at the end of the activity articulating an important goal in reading interface evaluations, the current challenges to measuring success reliably, and ideas they have come up with for doing it better.

*Supporting materials: theoretical primers.* To support deeper exploration in the affinity group activities, we on the organizing team will prepare short packets that briefly review major theories of the cognition of reading. Currently, we plan to prepare primers on comprehension [37], sensemaking [39], cognitive load [9], and discovery learning [11]. Primers will be brief, requiring only a minute or two to review. They will provide a definition of the reading task as viewed by that theory and an example of applying it to an augmented reading interface from the CHI community. If there is extra time after the lightning talks and before the activities we may include short, 3-minute introductions to the theories.

### 4.3 Closing discussion

In the last 15 minutes of the workshop, groups will report out discoveries from their activities. We anticipate contributions in the form of clarified ideas around exciting interface ideas, emerging problems in reading that need to be addressed, and improved ideas for evaluation. In the situation where our time runs short, groups will upload their activity materials and output (e.g., slides with possible designs, reviews or evaluation approaches) to a dedicated Discord channel to make their findings available to the group.

### 4.4 Post-Workshop Plans and Offline Materials

To support continued conversation and engage those of our community who are not able to attend in person, we will make the following outputs publicly available after the workshop:

- All accepted workshop submissions.
- A YouTube channel containing video recordings of the lightning talks (with permission of participants).
- Outputs of the small-group activities (with the permission of participants).
- A Discord channel for participants and organizers to continue their discussions.

Time	Activity
13:00	Welcome
13:05	Introductions (lightning talks)
13:35	Affinity group activities
14:15	Closing discussion

**Table 1: Schedule of workshop activities. Times are tentative given available workshop slots.**

Given the rising relevance of these topics, we will also recruit future organizers from workshop attendees to continue to grow the workshop and community in future HCI and/or NLP venues.

## 5 Logistics

### 5.1 Plans to Publish Workshop Proceedings

We will publish non-archival workshop proceedings composed of participants' submitted papers to arXiv using report numbers. We will also host these proceedings, lightning talk slides and other asynchronous materials like the primers and outputs of the design activities (Section 4.4) on the workshop's website.

### 5.2 Accessibility considerations

- We will provide all workshop materials to participants ahead of time.
- We will make live AI transcriptions of all activities.
- We will support multiple ways of contributing to discussions. During group activities, groups can choose what kinds of materials will best support their members.

### 5.3 Projected Attendance

We anticipate this workshop will draw about 30 participants. All attendees are expected to have had some stake in developing augmented reading experiences, with the largest cohort consisting of researchers who have developed or evaluated interfaces and/or technical pipelines. We anticipate and hope that those who attend have worked on a variety of augmentations, types of text, and languages. We also anticipate and hope that participants will represent proficiency in a variety of reading-related theories.

## 6 Call for Participation

Are you interested in what the future holds for how we read? Then join us at the first CHI workshop on the Science & Technology of Augmented Reading (STAR). Come together with researchers from HCI and AI to work on the most important questions for interfaces that alter the way we read. Topics include but are not limited to:

- Reading interfaces in different domains, such as as education or literature.
- Contexts where reading is an integral task, such as legal reasoning, scientific discovery, or writing.
- The future of reading with AI, or, does reading still matter?

To participate, submit a 2–4-page paper in 2-column SIGCHI format at <https://chi-star-workshop.github.io/>. Submissions are encouraged from junior and senior researchers alike, for ideas across the spectrum from hot takes to proofs of concepts and finished research results. The main requirement of these submissions is that it prepares you with an idea to bring to the workshop. Accepted

submissions will be posted or linked in non-archival form on the workshop website.

One author of each accepted submission must attend the workshop and one day of CHI '26. At the workshop itself, you will see and be a part of round-the-room 1-minute lightning talks, followed by speculative design and research activities in small interest groups formed at the event. Groups will be leveled up with brief primer documents on classic theories of reading. We can't wait to see you there!

## Acknowledgments

We would like to thank the reviewers for their valuable feedback. We would also like to thank Juho Kim, Advait Sarkar, Damien Masson, and Matthew Conlen for support and feedback on earlier versions of this proposal. This material is based upon work supported by the National Science Foundation under Grant No. IIS-2524724.

## References

- [1] Apple. 2025. Apple Intelligence - AI for the rest of us. *Apple Media* (2025).
- [2] Akari Asai, Jacqueline He, Rulin Shao, Weijia Shi, Amanpreet Singh, Joseph Chee Chang, Kyle Lo, Luca Soldaini, Sergey Feldman, Mike D'arcy, David Wadden, Matt Latzke, Minyang Tian, Pan Ji, Shengyan Liu, Hao Tong, Bohao Wu, Yanyu Xiong, Luke Zettlemoyer, Graham Neubig, Dan Weld, Doug Downey, Wen tau Yih, Pang Wei Koh, and Hannaneh Hajishirzi. 2024. OpenScholar: Synthesizing Scientific Literature with Retrieval-augmented LMs. *arXiv preprint arXiv:2411.14199* (2024).
- [3] Tal August, Lauren Kim, Katharina Reinecke, and Noah A Smith. 2020. Writing Strategies for Science Communication: Data and Computational Analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [4] Tal August, Kyle Lo, Noah A. Smith, and Katharina Reinecke. 2024. Know Your Audience: The Benefits and Pitfalls of Generating Plain Language Summaries Beyond the "General" Audience. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.
- [5] Tal August and Katharina Reinecke. 2019. Pay Attention, Please: Formal Language Improves Attention in Volunteer and Paid Online Experiments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, 1–11.
- [6] Tal August, Katharina Reinecke, and Noah A. Smith. 2022. Generating Scientific Definitions with Controllable Complexity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics, 8298–8317.
- [7] Tal August, Lucy Lu Wang, Jonathan Bragg, Marti A. Hearst, Andrew Head, and Kyle Lo. 2023. Paper Plain: Making Medical Research Papers Approachable to Healthcare Consumers with Natural Language Processing. *ACM Transactions on Computer-Human Interaction (TOCHI)* 30, 5, Article 74 (sep 2023), 38 pages.
- [8] Mark Bernstein. 2009. On hypertext narrative. In *ACM Conference on Hypertext & Social Media*.
- [9] Paul Chandler and John Sweller. 1991. Cognitive load theory and the format of instruction. *Cognition and instruction* 8, 4 (1991), 293–332.
- [10] Matthew Conlen, Megan Vo, Alan Tan, and Jeffrey Heer. 2021. Idyll studio: A structured editor for authoring interactive & data-driven articles. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 1–12.
- [11] Ton De Jong and Wouter R Van Joolingen. 1998. Scientific discovery learning with computer simulations of conceptual domains. *Review of educational research* 68, 2 (1998), 179–201.
- [12] Tiffany D. Do, Usama Bin Shafiq, Elsie Ling, and Nikhil Sarda. 2025. PAIGE: Examining Learning Outcomes and Experiences with Personalized AI-Generated Educational Podcasts. In *Proceedings of the 2025 CHI Conference on Human Factors*

- in *Computing Systems (CHI '25)*. Association for Computing Machinery, Article 896, 12 pages.
- [13] Elsevier. 2025. ScienceDirect AI. Accessed: 2025-09-10.
- [14] Raymond Fok, Hita Kambhmettu, Luca Soldaini, Jonathan Bragg, Kyle Lo, Marti Hearst, Andrew Head, and Daniel S Weld. 2023. Scim: Intelligent Skimming Support for Scientific Papers. In *Proceedings of the 28th International Conference on Intelligent User Interfaces (IUI '23)*. Association for Computing Machinery, 476–490.
- [15] Raymond Fok, Nedim Lipka, Tong Sun, and Alexa F Siu. 2024. Marco: Supporting Business Document Workflows via Collection-Centric Information Foraging with Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [16] Katy Ilonka Gero, Chelse Swoopes, Ziwei Gu, Jonathan K. Kummerfeld, and Elena L. Glassman. 2024. Supporting Sensemaking of Large Language Model Outputs at Scale. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery.
- [17] Elena L. Glassman, Ziwei Gu, and Jonathan K Kummerfeld. 2024. AI-Resilient Interfaces. *arXiv preprint arXiv:2405.08447* (2024).
- [18] Elena L Glassman, Jeremy Scott, Rishabh Singh, Philip J Guo, and Robert C Miller. 2015. OverCode: Visualizing variation in student solutions to programming problems at scale. *ACM Transactions on Computer-Human Interaction (TOCHI)* 22, 2 (2015), 1–35.
- [19] Elena L. Glassman, Tianyi Zhang, Björn Hartmann, and Miryung Kim. 2018. Visualizing API usage examples at scale. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 1–12.
- [20] Google. 2024. Summarize documents, text, and more with generative AI and LLMs. Data accessed: September 20, 2024.
- [21] Google LLC. 2025. Gemini Deep Research. Accessed: 2025-08-30.
- [22] Andrew D. Gordon, Carina Negreanu, José Cambronero, Rasika Chakravarthy, Ian Drosos, Hao Fang, Bhaskar Mitra, Hannah Richardson, Advait Sarkar, Stephanie Simmons, Jack Williams, and Ben Zorn. 2024. Co-audit: tools to help humans double-check AI-generated content. *Proceedings of the 14th annual workshop on the intersection of HCI and PL (PLATEAU 2024)* (5 2024).
- [23] Ziwei Gu, Ian Arawjo, Kenneth Li, Jonathan K Kummerfeld, and Elena L Glassman. 2024. An AI-Resilient Text Rendering Technique for Reading and Skimming Documents. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, 1–22.
- [24] Andrew Head, Kyle Lo, Dongyeop 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*. Association for Computing Machinery, 1–18.
- [25] Andrew Head, Amber Xie, and Marti A Hearst. 2022. Math augmentation: How authors enhance the readability of formulas using novel visual design practices. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [26] Ziheng Huang, Tal August, and Hari Sundaram. 2025. TermSight: Making Service Contracts Approachable. *arXiv preprint arXiv:2506.12332* (2025).
- [27] Hang Jiang, Xiajie Zhang, Robert Mahari, Daniel Kessler, Eric Ma, Tal August, Irene Li, Alex Pentland, Yoon Kim, Jad Kabbara, and Deb Roy. 2024. Leveraging Large Language Models for Learning Complex Legal Concepts through Storytelling. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- [28] Nikhita Joshi and Daniel Vogel. 2024. Constrained Highlighting in a Document Reader can Improve Reading Comprehension. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, 1–10.
- [29] Hita Kambhmettu, Jamie Flores, and Andrew Head. 2025. Traceable Texts and Their Effects: A Study of Summary-Source Links in AI-Generated Summaries. In *Proceedings of the Extended Abstracts of the 2025 CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, Article 538, 7 pages.
- [30] Hita Kambhmettu, Danaë Metaxa, Kevin Johnson, and Andrew Head. 2024. Explainable Notes: Examining How to Unlock Meaning in Medical Notes with Interactivity and Artificial Intelligence. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [31] Priyanka Kargupta, Ishika Agarwal, Tal August, and Jiawei Han. 2025. Tree-of-Debate: Multi-Persona Debate Trees Elicit Critical Thinking for Scientific Comparative Analysis. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics, 29378–29403.
- [32] Tae Soo Kim, Matt Latzke, Jonathan Bragg, Amy X. Zhang, and Joseph Chee Chang. 2023. Papeos: Augmenting Research Papers with Talk Videos. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (2023).
- [33] Kyle Lo, Joseph Chee Chang, Andrew Head, Jonathan Bragg, Amy X. Zhang, Cassidy Trier, Chloe Anastasiades, Tal August, Russell Authur, Danielle Bragg, Erin Bransom, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Yen-Sung Chen, Evie (Yu-Yen) Cheng, Yvonne Chou, Doug Downey, Rob Evans, Raymond Fok, F.Q. Hu, Regan Huff, Dongyeop Kang, Tae Soo Kim, Rodney Michael Kinney, Aniket Kittur, Hyeonsu B Kang, Egor Klevak, Bailey Kuehl, Michael Langan, Matt Latzke, Jaron Lochner, Kelsey MacMillan, Eric Marsh, Tyler Murray, Aakanksha Naik, Ngoc-Uyen Nguyen, Srishti Palani, Soya Park, Caroline Paulic, Napol Rachatasumrit, Smita Rao, Paul L Sayre, Zejiang Shen, Pao Siangliulue, Luca Soldaini, Huy Tran, Madeleine van Zuylen, Lucy Lu Wang, Christopher Wilhelm, Caroline M Wu, Jiangjiang Yang, Angele Zamarron, Marti A. Hearst, and Daniel S. Weld. 2024. The Semantic Reader Project: Augmenting Scholarly Documents through AI-Powered Interactive Reading Interfaces. *Communications of the ACM (CACM)* (2024).
- [34] Damien Masson, Sylvain Malacria, Géry Casiez, and Daniel Vogel. 2023. Charagraph: Interactive Generation of Charts for Realtime Annotation of Data-Rich Paragraphs. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI '23)*. Association for Computing Machinery, Article 146, 18 pages.
- [35] Damien Masson, Sylvain Malacria, Géry Casiez, and Daniel Vogel. 2023. Statslator: Interactive Translation of NHST and Estimation Statistics Reporting Styles in Scientific Documents. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, Article 91, 14 pages.
- [36] Puneet Mathur, Alexa Siu, Varun Manjunatha, and Tong Sun. 2024. DocPilot: Copilot for Automating PDF Edit Workflows in Documents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, System Demonstrations*. 232–246.
- [37] Danielle S McNamara and Joe Magliano. 2009. Toward a comprehensive model of comprehension. *Psychology of learning and motivation* 51 (2009), 297–384.
- [38] OpenAI. 2025. Introducing Deep Research. Accessed: 2025-08-30.
- [39] Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, Vol. 5. McLean, VA, USA, 2–4.
- [40] Morgan N. Price, Bill N. Schilit, and Gene Golovchinsky. 1998. XLibris: the active reading machine. *CHI 98 Conference Summary on Human Factors in Computing Systems* (1998).
- [41] Napol Rachatasumrit, Jonathan Bragg, Amy X Zhang, and Daniel S Weld. 2022. Citeread: Integrating localized citation contexts into scientific paper reading. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*. Association for Computing Machinery, 707–719.
- [42] Adobe Research. 2024. How Adobe Research is helping unlock the intelligence inside trillions of PDFs. *Adobe Research News* (2024).
- [43] Jon Saad-Falcon, Joe Barrow, Alexa F. Siu, Ani Nenkova, Ryan Rossi, and Franck Dernoncourt. 2023. PDFTriage: Question Answering over Long, Structured Documents. *arXiv preprint arXiv:2309.08872* (2023).
- [44] Jackson Stokes, Tal August, Robert A Marver, Alexei Czeskis, Franziska Roesner, Tadayoshi Kohno, and Katharina Reinecke. 2023. How Language Formality in Security and Privacy Interfaces Impacts Intended Compliance. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 12 pages.
- [45] Craig S. Tashman and W. Keith Edwards. 2011. LiquidText: a flexible, multitouch environment to support active reading. *Proceedings of the 2011 CHI Conference on Human Factors in Computing Systems* (2011).
- [46] Lucy Lu Wang, Isabel Cachola, Jonathan Bragg, Evie Yu-Yen Cheng, Chelse Haupt, Matt Latzke, Bailey Kuehl, Madeleine N van Zuylen, Linda Wagner, and Daniel Weld. 2021. SciA11y: Converting Scientific Papers to Accessible HTML. In *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '21)*. Association for Computing Machinery, Article 85, 4 pages.
- [47] Zhiyuan Wu, Jiening Li, Kevin Ma, Hita Kambhmettu, and Andrew Head. 2023. FFL: A Language and Live Runtime for Styling and Labeling Typeset Math Formulas. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–16.
- [48] Litao Yan, Alyssa Hwang, Zhiyuan Wu, and Andrew Head. 2024. Ivie: Lightweight anchored explanations of just-generated code. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 1–15.
- [49] Litao Yan, Miryung Kim, Bjoern Hartmann, Tianyi Zhang, and Elena L Glassman. 2022. Concept-annotated examples for library comparison. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–16.