Demonstration of ReadFlect: Scaffolding Intent-driven, Multi-session, and Reflective Reading of Academic Papers

Yuwei Xiao
University of California, Los Angeles
Los Angeles, California, USA
yuweix@ucla.edu

Ollie Pai
Computation & Discovery Lab
University of California, Los Angeles
Los Angeles, California, USA
o.pai@ucla.edu

Brian Roysar
University of California, Los Angeles
Los Angeles, California, USA
brianroysar@ucla.edu

Michael Shi
University of California, Los Angeles
Los Angeles, California, USA
michaelshi@g.ucla.edu

Eunice Jun
UCLA
Los Angeles, California, USA
emjun@cs.ucla.edu

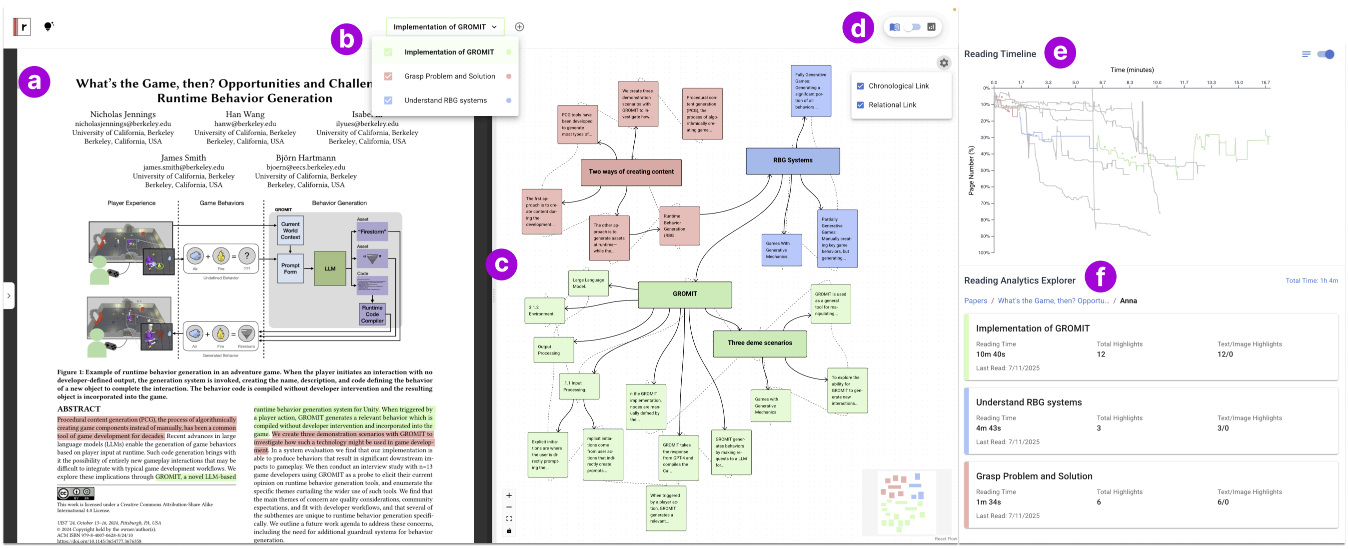

Figure 1: Interface and Key Features of ReadFlect. (a) Users highlight paper content with color-coded reading intentions. A drop-down menu (b) lets users select the active intention and toggle the visibility of other intentions to promote focus. On the interactive canvas (c), highlights appear as nodes linked by two edge types: chronological edges (dashed gray) reflect the order in which highlights were made, while relational edges (solid black) are created by users to connect conceptually related ideas. A toggle control (d) allows users to switch between reading and analysis modes. In analysis mode, a timeline visualization (e) presents users' reading progression over time, showing page positions and highlight events, while a dashboard (f) summarizes metrics such as reading time and highlight distribution.

Abstract

Reading academic papers is a routine yet cognitively demanding task for researchers and students. The reading process is often fragmented across sessions and shaped by evolving goals, yet most tools treat it as a single-pass, intention-agnostic activity. We present ReadFlect, an interactive system that scaffolds intent-driven, multi-session, and reflective reading of academic literature. ReadFlect enables users to define and manage reading intentions, annotate papers in context, and externalize insights through a diagramming canvas. It further supports reflection through visual analytics that capture reading behaviors over time. We demonstrate how ReadFlect supports more structured engagement with scientific literature and discuss its potential for enhancing comprehension, facilitating synthesis, and supporting collaborative and pedagogical use cases.

CCS Concepts

• **Human-centered computing** → **Interactive systems and tools**; *Information visualization*; Empirical studies in HCI.

UIST Adjunct '25, Busan, Republic of Korea

ACM ISBN 979-8-4007-2036-9/25/09
https://doi.org/10.1145/3746058.3759011

Keywords

Intelligent reading interfaces, scientific papers, information visualization, active reading, intent specification, highlights, sensemaking.

ACM Reference Format:
Yuwei Xiao, Ollie Pai, Brian Roysar, Michael Shi, and Eunice Jun. 2025. Demonstration of ReadFlect: Scaffolding Intent-driven, Multi-session, and Reflective Reading of Academic Papers. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST Adjunct '25), September 28–October 01, 2025, Busan, Republic of Korea.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3746058.3759011

1 Introduction

Reading academic papers is a common, often daily practice for students and researchers. Despite its frequency, it remains a cognitively demanding and complex task. Academic papers are typically dense with domain-specific terminology, unfamiliar methodologies, and complex data representations. As a result, the reading process is rarely linear. Instead, readers often pursue diverse and evolving intentions, ranging from skimming for general understanding, critically reviewing content, to deep diving for methodological details or practical application [9].

To mitigate the challenges of reading scientific literature, various strategies and tools have been developed. Approaches such as active reading [2] are widely adopted in academic settings, and experienced researchers have shared structured techniques for reading papers effectively, such as Keshav's three-pass method [7].

Complementing these strategies, a growing body of interactive systems has explored ways to support paper reading through interface design and automation. Systems like texSketch [11] and LiquidText [12] incorporate pen-and-touch interaction to support active reading and diagramming. SCIM [4] helps readers navigate key insights in papers by surfacing salient sentences. Qlarify [3] supports gradual disclosure of dense content through expandable abstracts. Tools like ScholarPhi [6] and Augmented Physics [5] provide just-in-time explanations and simulations to support comprehension of mathematical or visual material.

Despite these advancements, most tools treat reading as a single-session, intention-agnostic activity. In practice, however, reading is distributed across time, documents, and cognitive goals—and often involves significant re-visitation [1, 10]. Readers frequently return to previously material with new questions or intentions, yet lack sufficient support for tracking and reflecting on their evolving understanding. Without effective mechanisms to capture and organize the fragmented reading process, they risk losing track of reasoning, duplicating effort, or overlooking key insights.

To address these challenges, we present ReadFlect, an interactive system that tightly integrates reading and reflection to support intent-driven, multi-session engagement with academic papers. ReadFlect enables users to explicitly specify reading intentions, annotate in context, and externalize insights through active diagramming. At the same time, the system logs users' reading behaviors and progress, allowing them to retrospectively inspect their engagement and effectively re-enter previously read materials. We demonstrate ReadFlect through a detailed walkthrough of its features across various reading tasks, illustrating how it supports more structured, intentional, and reflective scholarly reading.

2 ReadFlect

ReadFlect is built on three concepts: **intentional reading**, **active diagramming** and **reflective analysis**. These concepts are integrated across three panels in our user interface: the reading panel, canvas panel, and analysis panel. Together, these components help users externalize their evolving goals, organize insights, and reflect on their engagement over time.

2.1 Reading Panel

The reading panel serves as the central interface for reading and annotating content. It supports traditional academic reading workflows while flexibly extending them through interactive affordances. Readers can load any number of PDFs into the system and switch between them seamlessly. Each document remains persistent across sessions, with highlights and annotations retained and synchronized in real-time.

2.1.1 Reading Intention Generation. Reading sessions in ReadFlect begin with the explicit articulation of a reading intention, such as *"understand methodology," "survey related work,"* or *"skim for the main contribution."* To support goal formulation, the system offers structured suggestions informed by existing pedagogical frameworks, including Keshav's three-pass method [7] and Bloom's taxonomy of learning objectives [8].

2.1.2 Reading & Paper Interaction. Once an intention is selected, ReadFlect tracks user interactions—including highlights, scroll behavior, and timestamps—and associates them with the active goal. Highlights can be created over text or images using a cursor-based selection. They are then color-coded, visually overlaid on the PDF, and backed by a structured representation on the canvas (Section 2.2). Readers may switch between intentions during a session and optionally filter out unrelated highlights to reduce distraction. These intention-specific interaction traces serve as a foundation for subsequent synthesis and reflection.

2.2 Canvas Panel

To support organization and synthesis, ReadFlect offers a visual workspace for active diagramming. Every time a user creates a highlight in the paper panel, a corresponding node is instantiated on the canvas. The user can interact with this canvas to explore relationships between ideas, build clusters, and surface deeper insights. This representation enables readers to externalize their evolving understanding, connect ideas across papers and sessions, and revisit prior insights with contextual structure, facilitating deeper engagement and knowledge integration.

2.2.1 Nodes. To support different stages of the reading and reflection process, ReadFlect introduces multiple node types.

- Highlight Node: A minimal representation of a highlight, shown by default. It displays only the raw content (text or image) and optional user-provided labels or tags. This view

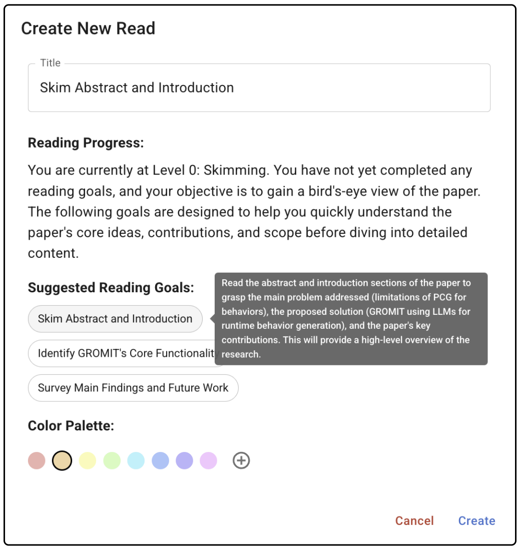

Figure 2: Reading intention creation in ReadFlect. Users can create a new reading intention by specifying one manually or selecting from system-generated suggestions. To provide actionable and context-specific recommendations, ReadFlect first analyzes the full paper, then evaluates the user's current reading progress (including prior intentions, highlights, and time spent) and generates a list of tailored reading goals with detailed descriptions.

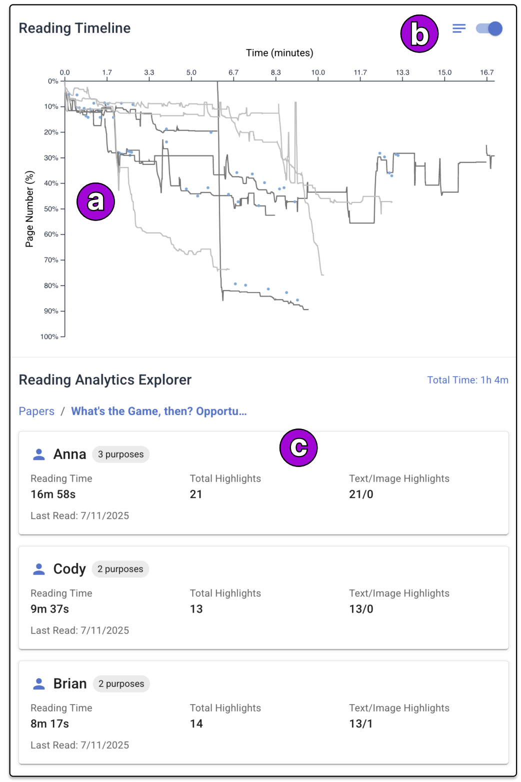

Figure 3: Multi-user analytics view in ReadFlect. In the reading timeline visualization (a), paths unrelated to the current selection are grayed out to help users focus on relevant activity. A toggle control (b) enables users to show or hide highlights on the timeline, represented as color-coded circular markers. The dashboard element (c) summarizes metrics based on the current analysis level, such as paper, user (as shown), or session-level, and clicking it navigates to a more detailed view (e.g., the session-level display in Fig. 1e–f).

minimizes clutter and supports quick scanning across the graph.

- Group Node: Users can select multiple nodes to create a group node. This group node represents a conceptual cluster and can be treated as a higher-level semantic unit.

2.2.2 Edges. To capture the structure and flow of a reader's thought process, ReadFlect includes multiple edge types—each designed to represent a different kind of relationship.

- Relational Edge: A user-defined edge that connects two nodes based on a semantic relationship. Users can create these edges manually by dragging from one node to another. Relation edges are meant to express connections such as cause-effect, question-answer, support-contradiction, or thematic similarity. These links may span across different pages of the same paper or different reading sessions.
- Chronological Edge: An automatically generated edge that links every consecutive pair of highlights made during a session. This encodes a sequential path that reflects the reader's chronological flow of attention.

2.3 Analysis Panel

ReadFlect includes a reflective analytics module to support retrospective engagement with the reading process. Users can toggle analysis mode via a control in the upper-right corner. When activated, the interface transitions from reading and graph interaction to an analytics dashboard. The analysis module supports not only single-paper analysis but also cross-document and cross-user comparisons. Users may select multiple papers to examine similarities and differences in their reading behavior, enabling reflection on patterns across literature.

2.3.1 Reading Timeline Visualization. A chronological visualization of user activity within a reading session, segmented by intention. Each highlight appears as a node plotted against time and location in the paper (as a percentage) and is color-coded by reading intention. This timeline helps users identify bursts of insight or attention, skimmed versus deeply-read sections, and interruptions or breaks in reading flow.

2.3.2 Reading Analytics Explorer. A multidimensional dashboard that surfaces metrics about a user's reading activity. Key metrics include time spent per paper or per reading session, total number of highlights, and ratio of image to text highlights. These metrics support reflection on reading depth, engagement, and knowledge structuring.

2.3.3 Multi-User Analytics. In collaborative settings, such as classrooms or reading groups, ReadFlect supports multi-user analytics. Through toggle controls, the analyst can select and compare reading progress and engagement patterns for different individuals or groups, as well as across various papers. This enables comparison of how different individuals approached the same paper, including which intentions were pursued, what sections were emphasized, and how interpretations varied. These comparisons support pedagogical reflection, peer learning, and scholarly discussion. Instructors may use this functionality to assess student engagement and guide critical reading instruction. Likewise, authors may use it to reflect on how their work is interpreted by different audiences.

3 Conclusion and Future Work

We presented ReadFlect, an interactive web-based system designed to support intentional, multi-session, and reflective reading of academic literature. By enabling users to articulate reading goals, externalize insights across sessions, and track their engagement, ReadFlecttransforms fragmented reading activities into a coherent and structured workflow. This approach offers a novel means of managing cognitive effort, fostering deeper comprehension, and sustaining long-term engagement with complex scientific texts.

In future work, we aim to enhance ReadFlect with a more comprehensive and intuitive workflow for specifying reading intentions. We also plan to refine the diagramming canvas and analytics visualizations to better support users in synthesizing insights and reflecting on their reading process. Beyond system refinement, we also plan to deploy ReadFlect in real-world settings, such as research labs or classrooms, to evaluate its functionality and usability in practice. This deployment will also enable the collection of authentic usage data, allowing us to investigate reading behaviors and uncover patterns in how users engage with academic literature over time.

References

[1] Eytan Adar, Jaime Teevan, and Susan T Dumais. 2008. Large scale analysis of web revisitation patterns. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 1197–1206.

[2] Mortimer J Adler and Charles Van Doren. 1972. *How to read a book*. Simon and Schuster.

[3] Raymond Fok, Joseph Chee Chang, Tal August, Amy X Zhang, and Daniel S Weld. 2023. Qlarify: Bridging scholarly abstracts and papers with recursively expandable summaries. *arXiv preprint arXiv:2310.07581* 6, 3 (2023).

[4] Raymond Fok, Hita Kambhamettu, 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*. 476–490.

[5] Aditya Gunturu, Yi Wen, Nandi Zhang, Jarin Thundathil, Rubaiat Habib Kazi, and Ryo Suzuki. 2024. Augmented Physics: Creating interactive and embedded physics simulations from static textbook diagrams. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–12.

[6] 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*. 1–18.

[7] Srinivasan Keshav. 2007. How to read a paper. *ACM SIGCOMM Computer Communication Review* 37, 3 (2007), 83–84.

[8] David R Krathwohl. 2002. A revision of Bloom's taxonomy: An overview. *Theory into practice* 41, 4 (2002), 212–218.

[9] Kenton O'Hara. 1996. *Towards a typology of reading goals*. Technical Report.

[10] Abigail J Sellen, Rachel Murphy, and Kate L Shaw. 2002. How knowledge workers use the web. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 227–234.

[11] Hariharan Subramonyam, Colleen Seifert, Priti Shah, and Eytan Adar. 2020. Texsketch: Active diagramming through pen-and-ink annotations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[12] 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*. 3285–3294.