Augmented Reading through Word-scale Graphics (WSGs)

Ruishi Zou
University of California, San Diego
La Jolla, CA, USA
ruzou@ucsd.edu

Abstract

This position paper explores how Word-scale Graphics (WSGs)–graphics displayed in situ within paragraphs at typographic size–can be used as a medium for augmented reading. Grounded in a narrative review of research on WSGs, I argue that WSGs' properties make them an ideal information scent to uncover disclosure gaps–information present in raw data/fact but lost in the disclosure process of composition (by both human and AI authors). I then show two use-case-specific designs that employ WSGs as contextual cues to display disclosure gaps and call for research on tools to support the creation and reading of WSG-augmented documents.

1 Introduction

One way to scaffold the process of reading (especially technical materials such as academic papers) is that it's the opposite of writing–readers untangle the written narratives to uncover the data, facts, and evidence embedded behind them. However, being a "reconstruction" process, the content readers consume is inherently lossy because of the writing strategies authors take, creating a ***disclosure gap*** between facts and narratives. For instance, researchers typically report only the "success" path of their study analysis [3], omitting the failure paths and alternative explanations to readers in favor of a more coherent research narrative. With the advancements of AI, the definition of authors also goes beyond human–disclosure gaps also occur when AIs are generating content, such as in the case of reading an automated summarization where AI selectively picks content to highlight "salient" information [2].

In face of disclosure gaps, readers might apply *reflective reading*, a technique that involves deeply engaging with and analytically thinking about the content. However, reflective reading is time-consuming and cognitively demanding–and oftentimes, readers might not find every critical point where a disclosure gap occurs. Compounded by the democratization of AI technologies (especially large language models), huge amounts of paraphrased text are being authored every day, which might be risky to careless readers (overtrusting generated text that causes misunderstanding), knowledge workers [14] (the convenience impedes critical thinking), and personnel working in critical societal domains (*e.g.*, legal, medical, with direct risk of malpractice due to wrong understanding).

Given the broad need and importance of identifying disclosure gaps, I argue that ***we need reading augmentations that provide information scents from the reading content to help users identify disclosure gaps.*** I argue that Word-scale Graphics (WSGs), or graphics displayed in situ within paragraphs at typographic size, are a potential vehicle to bridge such disclosure gaps. I connect my argument to the vision of AI-resilient interfaces [4, 10], extending it to include reading augmentations that not only enable readers to identify AI errors but also reveal disclosure biases in narratives in general (both human-authored and AI-authored).

In the following, I first present a review of what is known about WSGs and how they influence reading. Then, I envision two use cases for WSGs, showing 1) how we can repurpose them as information scents to reveal disclosure gaps, and 2) how future work can conceptualize designing WSGs for disclosure. I finally wrap up with a conclusion and call for future research.

2 What is known about WSGs for Reading Augmentation?

Word-scale graphics (WSG) have been discussed in the literature under several expressions. One of the more widely known terms is *sparkline* [17], defined as data-dense (commonly without visual marks other than data), word-sized graphics positioned within text.

More recent research [8] expanded sparkline with *word-scale visualization.* The key difference between word-scale visualization and sparkline is a more flexible integration between the graphic and text: word-scale visualizations entail everything that is a sparkline, but allow graphics 1) larger than a single line's height and 2) a floating position. Goffin et al. [6] further expanded the concept of word-scale visualizations to *word-scale graphics* (WSGs) to include both data-driven (*i.e.*, word-scale visualizations) and non-data-driven graphics. In this position paper, I use the term WSGs to refer to the broad set of graphics at typographic size.

Research has discussed several attributes of WSGs, including their positioning [8], design, function [6], and interaction [5, 9]. While not in abundance, a few applications designed for augmented reading also applied WSGs, most notably for augmenting code comprehension [12] and data-rich documents [18] (*e.g.*, Fig. 1). Besides, other research has investigated the in-the-wild use of WSGs, specifically their application in scientific communication [1, 15].

Meanwhile, evidence on the efficacy of WSGs as a means to augment reading is relatively scarce–in this position paper, I summarize themes that appear to be consistent across those studies. Because WSGs situate "naturally" within the textual space, they do not seem to impact readers' typical reading flow [18], agnostic to the positioning [7], and do not present significant signals in improving reading comprehension [7, 12, 18]. Instead, what appears where WSGs reap benefit is largely perceptual and cognitive–readers perceived to have better performance [12], induces lower cognitive load [18], and affords readers in using cues presented in WSGs to foster understanding [7, 18]. Those findings are particularly formative because they display two unique properties of WSGs: *1) they are a welcoming type of lightweight additions to text, and 2) they can be significant and informative enough to draw readers' attention to convey critical information.*

In this position paper, I argue that research has yet to fully tap the potential of WSGs. I argue that, instead of proposing WSGs with the primary goal of fostering reading comprehension (which has been found less effective), they should be repurposed as information

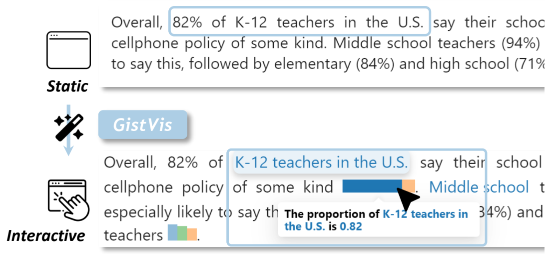

Figure 1: GistVis automatically transforms data-rich documents into interactive WSGs to facilitate insight discovery.

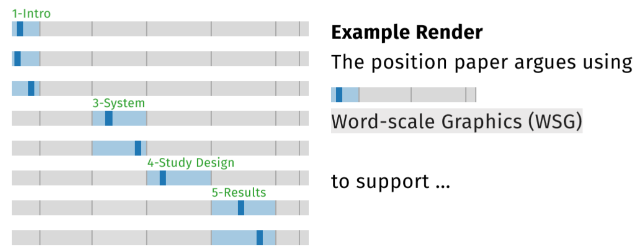

Figure 2: *Left* - A stack of WSGs showing the position of sources used in the automated summary. *Right* - Example render of the augmented document.

scents–intermediate representations between facts and narrative to highlight disclosure gaps. WSGs' unassuming presence might just be enough to hint to readers certain disclosure gaps.

3 WSGs as Information Scents

In this section, I substantiate my vision through two use cases that demonstrate how WSGs can reveal disclosure gaps in different contexts. Building on the two examples, I propose a conceptual model to inform the future design of reading augmentations to reveal disclosure gaps through WSGs.

3.1 Example Use Cases

3.1.1 Augmenting Data Claims. Data claims readers consume (*e.g.*,, academic papers, data-rich news) can be misleading. For example, consider the following statistical report–"*Mean*=5, *SD*=2.03"–a result computed by aggregating individual responses in a 10-point Likert scale questionnaire. While such a report seems robust at a glance, the summative nature of the disclosure strategy creates a gap between the real data distribution and its summative figures:

- **Case 1**: *Mean*=5, *SD*=2.03 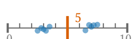
- **Case 2**: *Mean*=5, *SD*=2.03 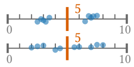

Specifically, Case 1 shows a "bimodal" response (centered around 3 and 7), indicating that the study's participants are divergent on the measure. Case 2's distribution is more "normal/uniform," suggesting a broad range of samples included in the study. To readers, such information is lost in the disclosure process from raw data facts to paper narratives, and could influence how readers interpret the validity of the findings. For example, in some study types (*e.g.*, phenomenography), Case 1 might not be acceptable–the authors failed to include participants with diverse experiences.

As shown above, WSGs can bridge the disclosure gap between raw data facts and their summary statistics. Such designs explicitly display the original data points, prompting readers to critically reflect on the implications of the data's shape. Two subsequent design challenges arise using WSGs to augment data claims: 1) in data-dense paragraphs, integrating WSGs can create a messy visual clutter and make reading difficult, and 2) the data context might not be readily available.

Future instantiations of intelligent, interactive reading tools beyond PDFs could perhaps address these challenges. For instance, with academia's push towards open science, more and more publications are making their experimental data publicly accessible–AI could access the raw data underlying the claims and automatically generate WSG augmentations. For instance, in a related but distinct use case, Zou et al. [18] demonstrated the feasibility of augmenting data-rich documents with WSGs using a large-language-model-based computation pipeline. Interactions, such as transformations between text layouts and WSGs, can also help reduce visual clutter and provide affordances that promise further insights–for example, aligning multiple WSGs' axes can enable comparison across multiple data reports.

3.1.2 Augmenting Automated Document Summaries. Having discussed how WSGs can help reveal disclosure gaps in data claims, I showcase another potential application: helping users identify disclosure gaps in automated summaries. Summaries can misrepresent original documents, for instance, by omitting content from certain sections. Consider summaries for academic papers–they might exclude the related work or discussion sections because they are only peripheral to "what" the paper did and do not contain the direct actions from the paper's author. Less experienced readers might not realize that the automated summary omits the gist of several sections–sections that could inspire subsequent research.

Fig. 2 shows one potential approach for how WSGs can provide additional information to help readers identify the disclosure gap. In this early design concept, sources referenced in the generated summary are highlighted in blue and rendered in situ within the summary text (see Fig. 2-*Right*). By reviewing the eight WSGs that show where sources are cited in the original document (see Fig. 2-*Left*), readers can easily identify that the information in Sections 2, 6, and 7 is missing, prompting them to probe those sections. In this use case, the WSGs do not represent facts in the document itself; instead, they provide "metadata" for the inner workings of AI algorithms, an approach familiar in explainable AI.

3.2 Conceptual Model

The two use cases presented above showed that WSGs can act as a medium to bridge the disclosure gap between facts and narrative. In the first use case, WSGs invite readers to reflect on how summative data reports omit information about data distribution in the raw data; in the second use case, WSGs provide cues indicating the omission of information between original documents and summaries. In both cases, ***WSGs do not disrupt the normal narrative flow***. They can be ignored–readers are not expected to retrieve dense information from them. Instead, they provide information scents–cues that, at a glance, invite reflective reading.

Given the above-mentioned design goal of WSGs for augmented reading, I envision the following three types of disclosure gaps WSGs could provide scent for. Inspired by recent research about designing visualizations for disclosure [16], I adapt their taxonomy of disclosure tactics to the context of reading augmentation, sketching how WSGs might be instantiated for each gap type:

- **Summarization Gaps**: Gaps caused by reducing the complexity of raw data/information, with the intention to keep as much original information as possible. Both use cases introduced above fall under this category.
- **Sampling Gaps**: Gaps caused by the raw data/information being consciously selected. For example, medical notes might frequently exhibit sampling gaps, such as heavy jargon usage (the omission of definitions of terms) [13]. A potential instantiation could be typographic graphics showing synonyms (*e.g.*, layman's term & expert term) to help readers understand clinical jargon.
- **Narration Gaps**: Gap caused by narrative writing techniques that sidelined alternative paths of a narrative. One potential instantiation to reveal this gap is to present a network WSG that shows alternative narrative paths.

I present the above conceptual model as a scaffold for future research to investigate: 1) what other disclosure gaps WSGs might address, 2) what visual forms they might take, and 3) how such augmentations could be generated and delivered to readers (AI-based technical solutions and/or better authoring tools, *e.g.*, [11]).

4 Conclusion

In this position paper, I reflect on the use of WSGs for reading augmentation. Grounded in prior literature, I argue for repurposing WSG as an information scent to reveal disclosure gaps rather than a primary location for data retrieval. Through two example use cases, I show how WSGs can be repurposed to reveal disclosure gaps and contribute a conceptual model for future WSG design. I hope this position paper can provide alternative viewpoints on WSG and augmented reading science and technology, and foster subsequent research that makes theoretical and empirical contributions, empowering readers to better employ reflective reading techniques while reading.

References

[1] Fabian Beck and Daniel Weiskopf. 2017. Word-Sized Graphics for Scientific Texts. *IEEE Transactions on Visualization and Computer Graphics* 23, 6 (June 2017), 1576–1587. doi:10.1109/TVCG.2017.2674958

[2] Isabel Cachola, Kyle Lo, Arman Cohan, and Daniel S. Weld. 2020. TLDR: Extreme Summarization of Scientific Documents. doi:10.48550/arXiv.2004.15011 arXiv:2004.15011 [cs]

[3] Pierre Dragicevic, Yvonne Jansen, Abhraneel Sarma, Matthew Kay, and Fanny Chevalier. 2019. Increasing the Transparency of Research Papers with Explorable Multiverse Analyses. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3290605.3300295

[4] Elena L. Glassman, Ziwei Gu, and Jonathan K. Kummerfeld. 2024. AI-Resilient Interfaces. doi:10.48550/arXiv.2405.08447 arXiv:2405.08447 [cs]

[5] Pascal Goffin, Tanja Blascheck, Petra Isenberg, and Wesley Willett. 2020. Interaction Techniques for Visual Exploration Using Embedded Word-Scale Visualizations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376842

[6] Pascal Goffin, Jeremy Boy, Wesley Willett, and Petra Isenberg. 2017. An Exploratory Study of Word-Scale Graphics in Data-Rich Text Documents. *IEEE Transactions on Visualization and Computer Graphics* 23, 10 (Oct. 2017), 2275–2287. doi:10.1109/TVCG.2016.2618797

[7] Pascal Goffin, Wesley Willett, Anastasia Bezerianos, and Petra Isenberg. 2015. Exploring the Effect of Word-Scale Visualizations on Reading Behavior. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '15)*. Association for Computing Machinery, New York, NY, USA, 1827–1832. doi:10.1145/2702613.2732778

[8] Pascal Goffin, Wesley Willett, Jean-Daniel Fekete, and Petra Isenberg. 2014. Exploring the Placement and Design of Word-Scale Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 2291–2300. doi:10.1109/TVCG.2014.2346435

[9] Pascal Goffin, Wesley Willett, Jean-Daniel Fekete, and Petra Isenberg. 2015. Design Considerations for Enhancing Word-Scale Visualizations with Interaction. In *Posters of the Conference on Information Visualization (InfoVis)*. IEEE.

[10] 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, New York, NY, USA, 1–22. doi:10.1145/3613904.3642699

[11] Jeffrey Heer, Matthew Conlen, Vishal Devireddy, Tu Nguyen, and Joshua Horowitz. 2023. Living Papers: A Language Toolkit for Augmented Scholarly Communication. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3586183.3606791

[12] Jane Hoffswell, Arvind Satyanarayan, and Jeffrey Heer. 2018. Augmenting Code with In Situ Visualizations to Aid Program Understanding. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–12. doi:10.1145/3173574.3174106

[13] Hita Kambhamettu, 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 (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3613904.3642573

[14] Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–22. doi:10.1145/3706598.3713778

[15] Siyu Lu, Yanhan Liu, Shiyu Xu, Ruishi Zou, and Chen Ye. 2026. Graphing Inline: Understanding Word-scale Graphics Use in Scientific Papers. doi:10.48550/arXiv.2603.10533

[16] Krisha Mehta, Gordon Kindlmann, and Alex Kale. 2026. Designing for Disclosure in Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 32, 1 (Jan. 2026), 1317–1327. doi:10.1109/TVCG.2025.3634781

[17] Edward Rolf Tufte. 2006. *Beautiful evidence*. Vol. 1. Graphics Press Cheshire, CT.

[18] Ruishi Zou, Yinqi Tang, Jingzhu Chen, Siyu Lu, Yan Lu, Yingfan Yang, and Chen Ye. 2025. GistVis: Automatic Generation of Word-scale Visualizations from Data-rich Documents. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713881