

CROSSWEAVE: Bridging Perspectives Across Online Communities with a Dual-Pane Design

Fei Fang*	Reva Hirave [†]	William Jurayj*
Yuqi Li*	Brian Lu*	Tarik Metin*
Tsugunobu Miyake [‡]	Kateryna Morhun [§]	Yash Permalla*
Kenan Rustamov*	Allen Shen*	Haojun Shi*
Prabhav Singh*	Xiheng Tom Wang*	Kevin Xu*
Qingcheng Zeng [¶]	Jiayi Zhang*	Daniel Khashabi*
Andrew Perrin*	Tiziano Piccardi*	Ziang Xiao*
	Jason Eisner*	
	jason@cs.jhu.edu	

Abstract

Social media systems typically display conversations among already familiar contributors, which can be predictable and one-sided. In civic discourse, this design narrows discussion, reinforces divides, and distorts the perception of public opinion. To encourage cross-community engagement, we present CROSSWEAVE, an AI-powered bridging system that augments the standard social media feed. As the user reads a post, CROSSWEAVE surfaces diverse relevant posts from other threads in a side pane and highlights the connections. Users are invited to venture out of their echo chamber, explore a broader range of views and arguments, and “click across” to engage with their authors. When they do, CROSSWEAVE facilitates constructive posting, not only by showcasing relevant past content but also by simulating possible reactions as the user drafts a post.

CCS Concepts

• **Human-centered computing** → Collaborative and social computing systems and tools; *Hypertext / hypermedia*; • **Information systems** → Specialized information retrieval; • **Computing methodologies** → Discourse, dialogue and pragmatics.

*Johns Hopkins University, Baltimore, Maryland, USA

[†]Princeton University, Princeton, New Jersey, USA

[‡]Georgia Institute of Technology, Atlanta, Georgia, USA

[§]Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

[¶]Northwestern University, Evanston, Illinois, USA



This work is licensed under a Creative Commons Attribution 4.0 International License. *CSCW Companion '26, Salt Lake City, UT, USA*

© 2026 Copyright held by the owner/author(s). This is the authors' recommended version of the work, slightly modified to correct authorship, improve formatting, add the demo URL, and slightly update the interface description and images. It is posted here for your personal use. Not for redistribution.

ACM ISBN 979-8-4007-2378-0/2026/10

<https://doi.org/10.1145/3785651.3831421>

Keywords

bridging systems; information exploration; applications of LLMs

ACM Reference Format:

Fei Fang, Reva Hirave, William Jurayj, Yuqi Li, Brian Lu, Tarik Metin, Tsugunobu Miyake, Kateryna Morhun, Yash Permalla, Kenan Rustamov, Allen Shen, Haojun Shi, Prabhav Singh, Xiheng Tom Wang, Kevin Xu, Qingcheng Zeng, Jiayi Zhang, Daniel Khashabi, Andrew Perrin, Tiziano Piccardi, Ziang Xiao, and Jason Eisner. 2026. CROSSWEAVE: Bridging Perspectives Across Online Communities with a Dual-Pane Design. In *Companion of the Computer-Supported Cooperative Work and Social Computing (CSCW Companion '26)*, October 10–14, 2026, Salt Lake City, UT, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3785651.3831421>

1 Introduction

Social media has been blamed for destabilizing democratic societies. Large majorities in several countries, and especially in the U.S., believe that social media has made us more divided in our political opinions, less civil in our political discourse, and more vulnerable to misinformation [7]. Yet social media is profitable, popular, and likely here to stay. So we must meet users and companies where they are. The challenge is to *mitigate* problems such as echo chambers [8, 24], unchallenged misinformation [13, 27], selective exposure [3, 23], and perceived polarization [14]—and in an appealing way.

We introduce CROSSWEAVE, a design that aims to broaden the information diet of social media users, connecting them to a wider range of high-quality ideas that bear on whatever they are currently reading or writing. It does not intervene in the main social media feed, but rather *augments* it with a curated side pane of social media posts from across various platforms. Construction of such a side pane is only now becoming possible, using large language models (LLMs) to automatically identify diverse relevant content and highlight connections. We outline the AI challenges in section 5, but this paper focuses on the UI concept and its prosocial goals.

A demo video is at <https://youtu.be/QFW9i7KoBhs>. At time of publication, a live demo is at <https://x-weave.org>.




Figure 1: CROSSWEAVE’s side (right) pane provides scrollable content related to the focused post in the main (left) pane. (a) Hovering the mouse over a highlighted span in the side pane reveals the corresponding span in the focused post. A hover tooltip (“2 more Industrial Policy And Investment”) indicates how many other side posts address the span’s topic and invites the user to click to see them. (b) Without hovering, highlights are visible on all posts in the side pane, color-coded by contribution type (Evidence (Public), Proposals, etc.). In this screenshot, the user is drafting a reply: CROSSWEAVE displays real-time AI-generated feedback and simulates several community reactions, encouraging the user to consider how their post might resonate before submitting.

2 Design Principles

Since the early days of Usenet and Facebook, social media and news comment sites have organized discussions into disjoint threads with different readerships. Arguably, this is the *core* reason for online echo chambers [8, 24] and for the limited range and quality of many online discussions [26].

CROSSWEAVE is designed to **augment** rather than replace this popular traditional social media design. Sustained conversations

within existing social circles are clearly engaging and presumably serve real needs (see [30]). Users come to social media for many reasons—to maintain offline social connections [29], to follow and react to news [21], to monitor others’ reactions [29], to share their own views and seek validation [18], to attract attention and gain social status [25], to cultivate followers and demolish opponents [15], and so on. They may not be actively seeking information, deep understanding, or new conversation partners [18, 25]. However, by

augmenting their display with a side pane (fig. 1), we hope to make them at least peripherally aware of other views and communities and draw them into exploration and engagement [16, 17].¹

CROSSWEAVE is intended as a **conduit** rather than a referee. All of the content is still written by actual human users: there is no AI voice summarizing comments [28] or “telling me what to think.” The side pane showcases a diverse set of views and supports easy exploration, so we hope users will regard it as interesting and empowering rather than biased or coercive. The side pane is not personalized to a particular user, but only customized to the current context.

CROSSWEAVE aims at both informational **broadening** and social **bridging**. First, it can broaden users’ understanding of the current topic by showing them additional perspectives from all angles. Second, it offers them bridges to new social communities [19]: a user can click on any post in the side pane to jump into its original thread and reply if they choose.

Finally, our design strengthens the **incentive** to *write* useful posts—the sort that will be promoted via the side pane. A user may put in the effort to rebut a friend’s misinformation, knowing that their effort will not be trapped in its original thread, but may be shown again next to similar misinformation when it appears in other threads across the platform (cf. X Community Notes [31]).²

3 Design Elements

A CROSSWEAVE user begins as in other social media interfaces, by scrolling through a list of top-level posts in the **main pane**.³ They may click on any post to open the conversation that it is part of. The main pane then changes to show the post with its thread ancestors above it and its replies below it (as in X, Mastodon, or Bluesky). The user may then post their own reply or continue navigating.

A permanent **side pane** (fig. 1) displays a ranked list of posts whose contributions relate to the currently **focused post**—that is, the main-pane post that was most recently clicked on or scrolled to. The user can click on any side post to open its thread in the main pane (thus bridging to that related thread).⁴

3.1 Side Pane Curation

Recall that the entire side pane provides commentary on a single focused post. Overall, we regard a side pane as good if browsing it would help the reader *reply thoughtfully* to the focused post.⁵ We select and rank side posts by balancing three criteria: **relevance**, **added value**, and **diversity** (as assessed by AI: see section 5).

A side post is **relevant** to the focused post to the extent that it engages with a *specific* claim, stance, or concern raised in the focused post. That is, the two posts share a *topic*—preferably a fine-grained one. (Either post may also make additional points.)

¹A side pane that surfaces relevant content should have further applications beyond social media. It might also be useful when reading or writing news stories [32], academic papers, patents, or legal briefs, and in other discussion forums designed to serve education or democratic deliberation [1, 20, 22].

²Ideally, the user would be notified and commended when their work is shown elsewhere and draws attention. This feature is on our to-do list.

³The main pane may be chronological, an algorithmic feed, or search results.

⁴Side pane posts have been plucked from their original threads and are sometimes hard to understand out of context (what is “he,” “us,” or “this”?). In these cases, we lightly rewrite the post with AI. By hovering over the Modified label, the user can see what changes were made (as edit markup) and verify that the meaning was preserved.

⁵A possible task for a future user study to evaluate our interface.

A side post **adds value** when it contributes something *beyond* what the focused post already provides. For example, it may introduce new evidence, offer a fresh argument or counterargument, propose a solution, or suggest a reframing. In linguistic terms, it supplies a novel *comment* on the existing shared *topic*—much as an ordinary social media reply might [2, 5, 11], though the connection may be less apparent in a side post that was not a direct reply.

The side pane as a whole should be **diverse** [6]. As the user scrolls down, successive posts should incrementally add non-redundant value, providing *new* topics, comments, and contribution types. This directly serves our broadening goals. A diverse side pane should also be more engaging to users than a list of redundant posts that repeat the same obvious or popular talking points [4], as we have noticed in the comments sections of many news articles.⁶

3.2 Enhanced Reading Interface

How can a user interact with the side pane? Most simply, they can skim the related posts that are shown, optionally scrolling downward to see more, or clicking to open interesting ones. Beyond this, CROSSWEAVE includes a suite of affordances to invite substantive engagement with the side pane. In our design, visually salient elements are always invitations to explore. Accepting an invitation ought to be cheap, rewarding, and non-disorienting.

To draw the user’s attention, we reveal *how* each side post *B* is related to the focused post *A*. After all, both posts may also make unrelated points, and *B* was probably not originally written as a reply to *A*,⁷ so the *A–B* point of contact is often not obvious. To help the user locate it, CROSSWEAVE highlights one or more text spans in each side post *B* that justify why *B* was displayed alongside *A*. These spans were identified by the backend AI.

The highlighted spans are color-coded based on their *contribution types*, which can be one or more of the following: public evidence, personal experiences, connections to other subjects, proposals, predictions, value statements, pointers to further resources, questions, issue framing, humor, as well as agreement and disagreement.

When the user hovers over a highlighted span in *B*,⁸ CROSSWEAVE temporarily highlights the corresponding span in *A* (and temporarily dims other highlights in the side pane), as shown in fig. 1a. These two corresponding highlighted spans generally contribute different *comments* on the same topic, which CROSSWEAVE calls out by temporarily bolding phrases within them.

Conversely, the user can also hover over some interesting or questionable part of *A*. Although *A* does not show any highlights by default (to minimize reading distraction), hovering over *A* will reveal its bold phrases. Hovering over a *specific* bold phrase will show its topic as a tooltip, and will temporarily dim all off-topic highlights in the various *B* posts. Clicking on the bold phrase filters the side pane to just the *B* posts with on-topic highlights.⁹

⁶By raising the ranking of unusual views, our diversity criterion does risk misleading the user about the prevalence of these views on the platform. We can mitigate this by ranking posts with popular views higher, other things equal, and also by using other means to indicate prevalence, such as the tooltips in section 3.2 below.

⁷When *B* is unclear outside its original thread, we use AI to lightly rewrite the side pane’s version (e.g., changing “they” to “auto manufacturers”). We disclose rewriting.

⁸Or taps it, on a mobile/touchscreen interface (wide phone or landscape mode).

⁹If a bold phrase has multiple topics, then clicking rotates through them and also the tooltip becomes a clickable selection menu.

The posts in the side pane can also be filtered by contribution type, via clickable icons on each post and at the top of the pane. The side pane normally shows a diverse mix of contribution types.

Finally, we let the user override the AI’s ranking in the side pane. While the initial ranking favors diversity (breadth-first), sometimes the user may be curious to dive (depth-first) into a particular topic raised in a side post *B*: “Are others making this point and how did they phrase it?” Thus, when the user hovers over a highlighted span in *B*, a tooltip like “8 more **AI Hype**” signals the prevalence of posts making this point and hints that the user can click to see more of them. Clicking will raise the ranking of the next 3 “AI Hype” side posts so that they fall immediately below *B*. The user can scroll through this topical block of posts, then click a link at the bottom to expand the block further with 3 more, or else just scroll past the current end of the block to resume reading other posts. All posts within the block highlight their “AI hype” spans, with the other highlights dimmed.¹⁰

While these features are primarily designed to facilitate navigation of the side pane, we also support them for the focused post’s direct replies in the main pane. This should help when a post attracts a large number of replies (e.g., a breaking news article from a trusted source). Rather than reading replies in chronological or popularity order, the CROSSWEAVE user can efficiently broaden their view by surveying the diverse perspectives scattered across hundreds of responses.

3.3 Enhanced Posting Interface

CROSSWEAVE should also improve average content quality on the platform. As the user drafts a reply to the focused post, the side pane of related posts provides food for thought, as section 3.1 noted. Some related posts may provide counterarguments that the user may care to address; others may offer supporting evidence or interesting angles that the user could incorporate. The user can quote a related post by dragging it into their own draft.

We also surface real-time feedback as the user writes (fig. 1b). It is generated by an AI persona [12]—a “wise and patient teacher and discussion moderator”—who both acknowledges the contributions of the current draft and provides suggestions to further strengthen its articulation. CROSSWEAVE also previews how the current draft might resonate with various communities, by displaying three AI-simulated replies to the draft—each adopting a distinct persona—and disabling submission for several seconds. With this preview, the user is encouraged to assess whether they would be satisfied with how the post might be perceived, and revise accordingly if not.

4 Data and Goals for the Live Demo

Our demonstration focuses on conversations about U.S. news (rather than about books or movies, say). Our demo system is currently populated with over 300 reader comments harvested from Reddit, Fox News, and The New York Times.

We invite users of the demo to consider whether CROSSWEAVE could improve civic discourse on social media:

- (1) Does the system feel trustworthy and usable?
- (2) When you are reading a post, does the side pane add value by genuinely broadening your perspective?
- (3) Would the side pane and the posting feedback help you take more informed and defensible positions on contested political issues, given your *own* values?
- (4) Would the system help you understand your political *opponents’* positions, arguments, and values well enough that you can have more constructive dialogue with them? Would it draw you into such dialogue?

5 AI Challenges

Our demo uses precomputed side pane data for each possible focused post. We first use an LLM (currently GPT-5.4) to describe and classify the contributions of each post, to identify a highlightable span for each contribution, and to assess the relevance and added value. Then, from these judgments, a cheaper secondary pipeline (GPT-5.4 and GPT-4.1-mini) induces a small set of subtopic labels to enable the depth-first drilldown. To finalize the ranking, we first use a deterministic, interpretable function to combine relevance and added value into a single score, and then rank greedily via maximal marginal relevance [6], penalizing redundancy in subtopic and contribution type.

While this implementation illustrates the UI vision and confirms that AI can indeed be used to identify and highlight relevant content, it is quadratically expensive and runs offline. Handling live traffic would require CROSSWEAVE to construct each side pane efficiently on demand (at least when no recent version is cached). We have been working on this problem for some time, in parallel to developing the UI design. Challenges include scalability and latency. A reasonable workflow would identify high-quality reusable posts as they arrive, preprocess them by extracting their substantive key points (see the contribution types in section 3.2), semantically index these key points for rapid retrieval from a precomputed topic graph or vector database, and later rerank retrieved posts on the fly using a fast scoring function that considers diversity [6].

Of course, rapidly assembling a diverse list of relevant documents is the central task of information retrieval. IR and NLP researchers have recently been incorporating LLMs into the indexing process [e.g., 10]. Still, our criteria in section 3.1 may be particularly hard to index for.¹¹ They seem to require closer reading, deeper reasoning, and more structured human evaluation than traditional IR settings. We find that it is challenging even for a human to assess numerically how relevant a post is, how valuable its contributions are, and how distinct they are from the contributions of other posts. Furthermore, such assessments are value-laden, subjective, and likely gameable. Yet we must take care to make them fairly, so as not to inadvertently suppress some viewpoints or lose user trust on suspicion of doing so. In short, we regard the retrieval of broadening content as a challenging new task for NLP.

¹⁰The side pane allows at most one topic block at a time. Canceling the current topic block or switching to a new one removes the current block formatting, but leaves the posts in their position so that the user does not have to encounter them again when scrolling downward.

¹¹For example, a real reply might challenge a **virtual standpoint** that was only quietly assumed by the focused post [9]. Finding non-replies that can play this role seems especially difficult using only vector-based retrieval and reranking, even if a powerful LLM could read them in the context of the focused post and recognize their relevance.

6 Conclusion

Social media is a consequential user interface. CROSSWEAVE is our effort to reimagine it with AI, not by adding a chat pane but by designing new helpful and prosocial features that are quietly enabled by LLMs. CROSSWEAVE attempts to connect isolated social media threads to the broader landscape of viewpoints, arguments, and users, in the hope of giving users a more mature understanding of this landscape and enabling them to participate more constructively through their own posts.

Acknowledgments

We thank Sanjeev Khudanpur and the 10th Frederick Jelinek Memorial Summer Workshop on Speech and Language Technology (JSALT 2024) for hosting and funding the 6-week residential hackathon in which we developed a previous version of this interface. Our project has subsequently been funded by a 2025 Johns Hopkins University Discovery Award, “Using AI to Make Social Media More Prosocial.” Our gratitude also goes to Evan Prodromou and Andy Piper from Mastodon and to the authors of many open-source tools.

As section 4 mentions, our demo makes use of harvested public comments. We thank the authors of those comments and reproduce them in our demo under fair use, with their username shown. The official ACM version of this paper has supplementary material that credits the authors of example comments from the official paper and video.

References

- [1] [n. d.]. *vTaiwan: rethinking democracy*. <https://info.vtaiwan.tw/>
- [2] Kelsey Allen, Giuseppe Carenini, and Raymond Ng. 2014. Detecting Disagreement in Conversations using Pseudo-Monologic Rhetorical Structure. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1169–1180. doi:10.3115/v1/D14-1124
- [3] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to Ideologically Diverse News and Opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132. doi:10.1126/science.aaa1160
- [4] Roy Bar-Haim, Yoav Kantor, Lilach Eden, Roni Friedman, Dan Lahav, and Noam Slonim. 2020. Quantitative Argument Summarization and Beyond: Cross-Domain Key Point Analysis. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 39–49. doi:10.18653/v1/2020.emnlp-main.3
- [5] Farah Benamara, Diana Inkpen, and Maite Taboada. 2018. Introduction to the Special Issue on Language in Social Media: Exploiting Discourse and Other Contextual Information. *Computational Linguistics* 44, 4 (12 2018), 663–681. doi:10.1162/coli_a_00333
- [6] Jaime G. Carbonell and Jade Goldstein. 1998. The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries. In *Proceedings of SIGIR*. 335–336. https://www.cs.cmu.edu/~jgc/publication/The_Use_MMR_Diversity_Based_LTMIR_1998.pdf
- [7] Pew Research Center. 2022. *Social Media Seen as Mostly Good for Democracy Across Many Nations, but U.S. Is a Major Outlier*. <https://www.pewresearch.org/global/2022/12/06/social-media-seen-as-mostly-good-for-democracy-across-many-nations-but-u-s-is-a-major-outlier/>
- [8] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The Echo Chamber Effect on Social Media. *Proceedings of the National Academy of Sciences* 118, 9 (2021), e2023301118. doi:10.1073/pnas.2023301118
- [9] Frans Van Eemeren, Rob Grootendorst, Curtis Scott Jacobs, and Sally Jackson. 2002. The Pragmatic Organization of Conversational Argument. In *Reconstructing Argumentative Discourse*. University of Alabama Press, Chapter 5, 91–116. <https://www.uapress.ua.edu/9780817312299/reconstructing-argumentative-discourse/>
- [10] Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2023. Precise Zero-Shot Dense Retrieval without Relevance Labels. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1762–1777. <https://aclanthology.org/2023.acl-long.99/>
- [11] Debanjan Ghosh, Smaranda Muresan, Nina Wacholder, Mark Aakhus, and Matthew Mitsui. 2014. Analyzing Argumentative Discourse Units in Online Interactions. In *Proceedings of the First Workshop on Argumentation Mining*. Association for Computational Linguistics, 39–48. doi:10.3115/v1/W14-2106
- [12] Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, Xin Zhou, Enzhi Wang, and Xiaohang Dong. 2024. Better Zero-Shot Reasoning with Role-Play Prompting. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. 4099–4113. doi:10.18653/v1/2024.naacl-long.228
- [13] David M. J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The Science of Fake News. *Science* 359, 6380 (2018), 1094–1096. doi:10.1126/science.aao2998
- [14] Matthew S. Levendusky and Neil Malhotra. 2016. (Mis)perceptions of Partisan Polarization in the American Public. *Public Opinion Quarterly* 80, S1 (01 2016), 378–391. doi:10.1093/poq/nfv045
- [15] Anuja Majumdar, Jon-Patrick Allem, Tess Boley Cruz, and Jennifer Beth Unger. 2018. The Why We Retweet Scale. *PLOS ONE* 13, 10 (10 2018), 1–12. doi:10.1371/journal.pone.0206076
- [16] Randall Munroe. 2008. xkcd: Duty Calls. Web comic. https://www.explainxkcd.com/wiki/index.php/386:_Duty_Calls
- [17] Randall Munroe. 2018. xkcd: Bad Opinions. Web comic. https://www.explainxkcd.com/wiki/index.php/2051:_Bad_Opinions
- [18] Mor Naaman, Jeffrey Boase, and Chih-Hui Lai. 2010. Is It Really About Me? Message Content in Social Awareness Streams. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (Savannah, Georgia, USA) (CSCW '10)*. Association for Computing Machinery, New York, NY, USA, 189–192. doi:10.1145/1718918.1718953
- [19] Aviv Ovadya and Luke Thorburn. 2023. *Bridging Systems: Open Problems for Countering Destructive Divisiveness across Ranking, Recommenders, and Governance*. Technical Report. Knight First Amendment Institute. <https://knightcolumbia.org/content/bridging-systems>
- [20] Jiaxin Pei, José Ramón Enriquez, Umar Patel, Alia Braley, Nuole Chen, Lily Tsai, and Alex Pentland. 2025. Deliberation.io: Facilitating Democratic and Civil Engagement at Scale with Open-Source and Open-Science. <https://deliberation.io/Workingpaper>
- [21] Pew Research Center. 2025. *Social Media and News Fact Sheet*. <https://www.pewresearch.org/journalism/fact-sheet/social-media-and-news-fact-sheet/>
- [22] Collective Intelligence Project and Contributors. 2025. Polis: Open Source AI for Large-Scale Open-Ended Feedback. <https://github.com/compdemocracy/polis>. Accessed: 2025-11-24.
- [23] Natalie Jomini Stroud. 2010. Polarization and Partisan Selective Exposure. *Journal of Communication* 60, 3 (09 2010), 556–576. doi:10.1111/j.1460-2466.2010.01497.x
- [24] Cass R. Sunstein. 2018. *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press. <http://www.jstor.org/stable/j.ctv8xnhtd>
- [25] Olivier Toubia and Andrew T. Stephen. 2013. Intrinsic vs. Image-Related Utility in Social Media: Why Do People Contribute Content to Twitter? *Marketing Science* 32 (2013), 368–392. <https://doi.org/10.1287/mksc.2013.0773>
- [26] Marshall Van Alstyne and Erik Brynjolfsson. 2005. Global Village or Cyber-Balkans? Modeling and Measuring the Integration of Electronic Communities. *Management Science* 51, 6 (June 2005), 851–868. doi:10.1287/mnsc.1050.0363
- [27] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151. doi:10.1126/science.aap9559
- [28] WashPostPR. 2024. The Washington Post Launches a New Commenting Experience Exclusively for Subscribers. *Washington Post* (Dec. 2024). <https://www.washingtonpost.com/pr/2024/12/16/washington-post-launches-new-commenting-experience-exclusively-subscribers/>
- [29] Anita Whiting and David Williams. 2013. Why People Use Social Media: A Uses and Gratifications Approach. *Qualitative Market Research: An International Journal* 16, 4 (08 2013), 362–369. doi:10.1108/QMR-06-2013-0041
- [30] Wikipedia contributors. 2026. Chesterton’s Fence—Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=Chesterton's_Fence&oldid=1151687374. [Online; accessed 4-August-2026].
- [31] X Corp. [n. d.]. *Community Notes Guide: Introduction*. <https://communitynotes.x.com/guide/en/about/introduction>
- [32] Kangyu Yuan, Guanzheng Chen, Sizhe Liang, Hehai Lin, Qingyu Guo, Dingdong Liu, Xiaojuan Ma, and Zhenhui Peng. 2026. CoNewsReader: Supporting Comprehensive Understanding and Raising Critical Thoughts on Social Media News Through Comments. <https://arxiv.org/abs/2604.27905>