
Scaling Human Judgment in Community Notes with LLMs

Haiwen Li, Soham De, Manon Revel, Andreas Haupt, Brad Miller, Keith Coleman, Jay Baxter, Martin Saveski, and Michiel A. Bakker

Abstract. This paper argues for a new paradigm for Community Notes in the LLM era: an open ecosystem where both humans and LLMs can write notes, while the decision of which notes are helpful enough to show remains in the hands of humans. This approach can accelerate the delivery of notes, while maintaining trust and legitimacy through Community Notes’ foundational principle: A community of diverse human raters collectively serve as the ultimate evaluator and arbiter of what is helpful. Further, the feedback from this diverse community can be used to improve LLMs’ ability to produce accurate, unbiased, broadly helpful notes—what we term Reinforcement Learning from Community Feedback (RLCF). This becomes a two-way street: LLMs serve as an asset to humans—helping deliver context quickly and with minimal effort—while human feedback, in turn, enhances the performance of LLMs. This paper describes how such a system can work, its benefits, key new risks and challenges it introduces, and a research agenda to operationalize it in a safe and trustworthy way.

1 Introduction

Professional fact-checking, which has played a large role in the historical online information environment, is often slow, limited in scale, and lacks the trust of large segments of the public (Flamini 2019; Walker and Gottfried 2019; Zhao and Naaman 2023). Community Notes (CN) offers a promising complement to professional fact-checking. It allows the community to add and rate contextual “notes” on social media posts, shifting the locus of authority from a centralized entity to a decentralized collective. Its strength lies in

its “bridging algorithm,”¹ which surfaces only those notes deemed helpful by raters who typically disagreed in past ratings, ensuring that surfaced context is broadly appealing and thus likely to inform people’s understandings even on contentious topics (Wojcik et al. 2022). X first launched its “Community Notes” program in 2021 (X Community Notes Guide, n.d.-b), and several other platforms have adopted similar systems since. Meta announced a switch from content moderation with third-party fact-checkers to Community Notes and now makes them available on Facebook, Instagram, and Threads to eligible contributors (Meta 2025a); YouTube is testing its “Community Notes” feature on videos (YouTube Team 2024); and TikTok is rolling out a similar feature under the name “Footnotes” (Presser 2025). Throughout this paper, we use “Community Notes” (CN) to denote this crowd-sourced, community-based fact-checking system, and our discussion applies to all such implementations.

This promising model (Drolsbach, Solovev, and Pröllochs 2024; Chuai et al. 2024; Slaughter et al. 2025), however, is poised for a potential step-change transformation. The core tasks of a Community Notes contributor—researching a claim, synthesizing diverse sources, and drafting a neutral, well-evidenced summary—are capabilities for which “Deep Research” LLMs show promise (Google 2025; OpenAI 2025a; xAI 2025). Indeed, work from 2024 has demonstrated that a fully automated pipeline for generating Community Notes with LLMs can, under certain circumstances, produce notes that are of similar quality to human-written notes—at a fraction of the time and effort (Zhou et al. 2024). Moreover, anecdotal evidence suggests that some note writers are already using LLMs to write notes (Mantzarlis 2025). LLM-written CNs have the potential to be faster to produce, less effort to generate, and of high quality, and thus represent an attractive direction for further research and deployment.²

While LLM-written CNs are compelling, critical questions remain: Can LLMs consistently produce accurate notes that are well received across perspectives? If so, can we use them to accelerate the addition of informative context in a way that is valued and trusted across viewpoints, while avoiding the fate of becoming a new version of top-down fact-checking that can lack broad trust? We believe both are possible.

Rather than replacing humans, LLMs can complement and enhance their work. Trust in CN stems not from who drafts the notes, but from who evaluates them. It is the collective judgment of a diverse and engaged community that grounds people’s trust in CN.

1. In more detail, users can rate a note as *helpful*, *somewhat helpful*, or *not helpful* and select optional tags to explain their rating. To determine whether to publish a note as *helpful*, the CN algorithm assigns each note a “helpfulness score” based on the ratings it received. The core component of this algorithm is a matrix factorization, which predicts each rating as $r_{un} = \mu + i_u + i_n + f_u \cdot f_n$. The global intercept μ accounts for the overall propensity of raters to rate notes *helpful*; the user intercept i_u accounts for the baseline propensity of a user to rate notes *helpful*; the note and the user factors $f_u \cdot f_n$ explain the preference of raters with certain viewpoints for certain notes; and the note intercept i_n is treated as the note helpfulness score. This matrix factorization encourages a representation that uses user and note factors to explain as much variation in the ratings as possible, and therefore, a note needs to be universally appealing to achieve a high intercept term/helpfulness score (Wojcik et al. 2022).

2. An API to submit LLM-generated Community Notes is now publicly available on X via <https://communitynotes.x.com/guide/en/api/overview> (X Community Notes Guide, n.d.-a).

This paper argues for a new paradigm for Community Notes in the LLM era: an open ecosystem where notes from both human writers and LLMs are submitted into a single pool, and the decision of which notes are helpful enough to show remains in the hands of the people. The system's legitimacy is still upheld by its foundational principle: a community of diverse human raters that collectively serve as the ultimate evaluator and arbiter of what is helpful. Further, the feedback from this diverse community can be used to improve LLMs' ability to produce accurate and broadly helpful notes. This can create a virtuous cycle: LLMs can be an asset to human contributors—helping deliver context quickly and with less effort—and human feedback, in turn, becomes an asset that further improves the work the LLMs do.

Integrating LLMs into the CN framework represents a transformative opportunity, one that offers promise of dramatically increased scale and speed of its contextual notes. However, this integration is not a simple technical change; it introduces a new set of high-stakes trust-and-safety challenges. This paper aims to provide a guide for navigating this complex landscape. We will outline the benefits of this hybrid human-AI model, critically examine the associated risks, and propose a concrete research agenda designed to harness the power of LLMs responsibly while reinforcing the core principles of a community-led approach.

2 Proposed hybrid system: Human and LLM note writing with human-only rating

We propose an expansion of the Community Notes pipeline, establishing a clear and synergistic division of labor between AI and the human community (see Figure 1).

In this new model, both humans and LLMs can contribute as note writers, while only humans rate notes and ultimately determine which notes are helpful enough to show. It seems likely that the combination of human and LLM writers is superior to either on its own, both in terms of generating trust and producing more effective notes. Humans and LLMs can bring different and additive strengths to the writing process.

2.1 Notes may be proposed by either humans or LLMs

2.1.1 Automated note creation by LLMs

Allowing automated note creation would enable the system to operate at a scale and speed that is impossible for human writers, potentially providing context for orders-of-magnitude more content across the web. Multiple stages can be automated:

- **Post selection:** Identifying which content warrants a note, which can involve prioritizing:

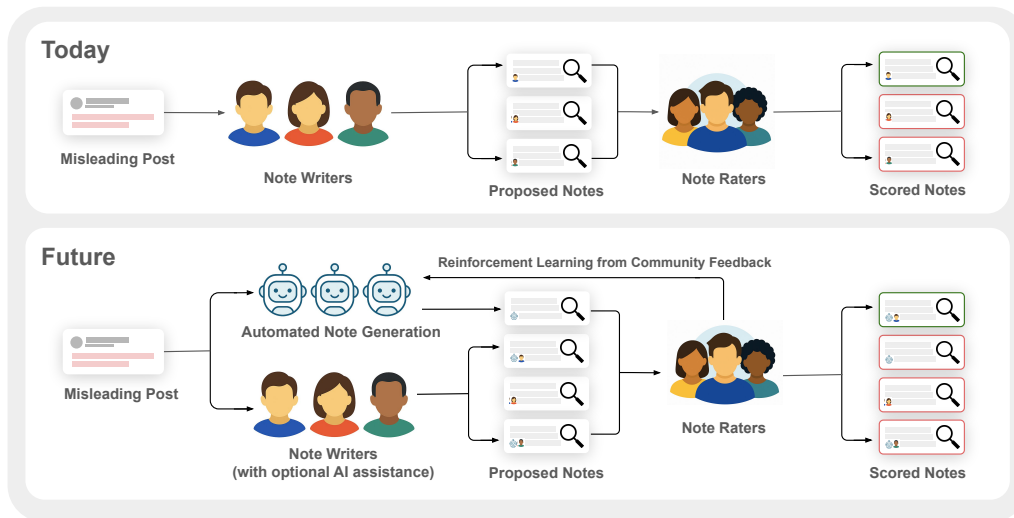


Figure 1: An expansion of the Community Notes pipeline from “all-human” to a hybrid “human-LLM” model. Today (top): Human note writers draft proposed notes in response to a misleading post, and other human contributors rate their helpfulness; the bridging algorithm picks the broadly helpful notes. Future (bottom): LLMs will also participate in the writing stage, producing notes or assisting human writers, while the rating stage remains human-only. Community feedback from human raters flows back to improve LLM note generation (RLCF).

- a. Content with high potential reach: Posts algorithmically predicted to gain significant viewership or virality, maximizing the note’s impact.
 - b. Content likely to be misleading: Posts where a note has a strong potential to be marked helpful by the bridging algorithm.
 - c. Community-flagged content: Posts flagged by human users as needing context, potentially accompanied by initial evidence, firsthand observations, or specific insights provided by the flagger.
- **Research:** Often including search and tool-use to gather information from a wide array of verifiable sources.
 - **Evidence synthesis:** Analyzing and synthesizing conflicting or complementary pieces of information to form a coherent factual basis.
 - **Note composition:** Drafting a clear, neutral, and well-sourced note designed to provide helpful context.

While the primary benefits of automated note creation are likely to be speed and scale, there may be other benefits, too: For example, LLMs can learn directly from the collective ratings of the community to internalize what makes a note helpful. As we will discuss in detail later, this process—which we term Reinforcement Learning from Community Feedback (RLCF)—may allow for more consistent adaptation than is possible for human

writers.

2.2 Human writers

In this hybrid system, human contributors could write notes from scratch or use AI assistance to help draft and source their contributions. Humans can benefit from AI assistance but do not have to defer to it—they can write notes where they want (for example, where LLMs are not aware of the need) or improve upon subpar LLM notes.

Human CN writers have demonstrated proficiency at handling novel (e.g., breaking news) situations where available information is limited and changing rapidly, niche topics where they are able to contribute deep knowledge on specific subjects, or scenarios where it is essential to understand how a post might be interpreted by people in order to understand why or why not a note might be valued.

Further, they are uniquely positioned to identify deficits or biases in the LLMs by writing notes on topics the automated writers overlook. They can help explore novel applications for notes: For example, although Community Notes was originally intended primarily as a way to add context to potentially misleading posts, some contributors started using it to do other types of crowd-sourced moderation such as flagging spam and scams, or crediting original creators of copy-pasted posts.

2.3 Unified evaluation: Human raters as the sole arbiters

While we advocate for contributions of LLMs for note authorship, we believe that evaluation from users is crucial. All notes, whether written by a human or an LLM, enter the same rating pool. The human community remains the ultimate arbiter of what gets published. Contributors' ratings are the primary input into the core approval algorithm. Currently, raters may also attach optional tags (e.g., "Incorrect information," "Argumentative or biased language") along with their ratings; to make this process more dynamic, the platform could allow additional free-text explanations from raters (e.g., "The note resonates with my experience, though the cited statistic feels outdated"). Indeed, eliciting such rationales is a promising direction, as recent work shows they can improve collective judgments (Srinivasan et al. 2025). These richer critiques can guide LLMs to generate a revised version immediately, and serve as valuable training signals for improving future note generation.

The bridging algorithm will continue to be the central mechanism for ensuring that only "bridging" notes—those found helpful by raters with diverse perspectives—are surfaced. This makes it difficult for biased notes to be approved, and creates incentives for accurate and neutral note writing. This structure also ensures that the system's trust and legitimacy do not derive from the writer (human or LLM) but from the collective curation and pluralistic judgment of the human community.

3 Benefits of this hybrid system

This proposed human-AI partnership offers profound benefits over a purely human (or purely AI) system:

- **Scalability and speed:** The automated LLM writing pipeline can respond to misleading content in real time and address a far greater volume of content than a human-powered system, tackling the “long tail” of niche misleading claims and out-of-context stories.
- **Enhanced quality:** The hybrid model fosters a higher standard of quality through several mechanisms. First, by having human- and LLM-written notes compete in the same rating system, a natural selection process emerges where the most helpful context—as determined by the community—rises to the top. Second, the system benefits from the complementary strengths of both humans and AI. LLMs excel at rapidly synthesizing widely available information, while human contributors are uniquely skilled at addressing novel situations and providing niche expertise. Finally, this partnership can enhance the pluralism of options presented to the community. If LLMs generate initial drafts that represent a wider range of perspectives than a single human writer typically could, the quality of community deliberation is improved from the start.
- **Preserve participatory oversight:** Human oversight is maintained because human collective judgment is elicited universally on all notes. The system’s legitimacy is upheld because the diverse human community remains the ultimate arbiter, collectively deciding what context is helpful enough to be shown, regardless of its origin.

4 Risks and challenges

This new approach introduces a new set of high-stakes challenges that must be proactively addressed.

- **Persuasive but inaccurate notes:** LLMs optimized to maximize helpfulness scores may become exceptionally skilled at crafting persuasive, emotionally resonant, and seemingly neutral notes that might not be factually accurate, leading to “helpfulness hacking” (Sharma et al. 2024; Bondarenko et al. 2025; Schoenegger et al. 2025). If rated helpfulness is not perfectly correlated with accuracy, then highly polished but misleading notes could be more likely to pass the approval threshold. This risk could grow as LLMs advance; they could not only write persuasively but also more easily research and construct a seemingly robust body of evidence for nearly any claim, regardless of its veracity, making it even harder for human raters to spot deception or errors.

- **Impact on human writer engagement:** Even with the ability to write, will human contributors remain motivated if LLMs are generating well-researched notes at scale? It is conceivable that if the platform becomes dominated by the speed and scale of AI, the sense of ownership and “skin in the game” that drives community participation could diminish (Burtch, Lee, and Chen 2024; Wagner and Jiang 2025), threatening the pool of diverse writers (and perhaps indirectly also raters) essential for the system’s long-term health and legitimacy.
- **Homogenization and reduced creativity:** While human writers provide some safeguard against total uniformity, a risk remains that the LLM-generated portion of the ecosystem will converge on a single, bland, “optimally inoffensive” style to pass the bridging algorithm. This could not only create a sterile information environment but also subtly devalue the more expressive or nuanced styles of human writers, crowding out the very creativity the hybrid model is meant to preserve.
- **Biased LLM writers:** LLMs can inherit bias from their training data or note-writing instructions, potentially producing incomplete or biased notes. The effects of biased note writing should be mitigated by the human review process. However, even if human review is effective, a capable and ideologically skewed note-writing LLM could still shift the overall distribution of available notes if it produces a high volume of persuasive content.
- **Risk of rater overload:** In a world with a very large volume of notes, how do we ensure human rating capacity can sufficiently scale? LLMs could strategically select where to write notes (prioritizing high-reach, likely misleading, or community-flagged content) in order to maximize impact, yet even so the sheer capacity of LLMs to generate notes for all such identified content could lead to a significant increase in the total volume of notes entering the rating system. This higher volume of notes could overwhelm the capacity of human raters or dilute the impact of the most genuinely critical notes if raters are spread too thin across a larger pool of notes.

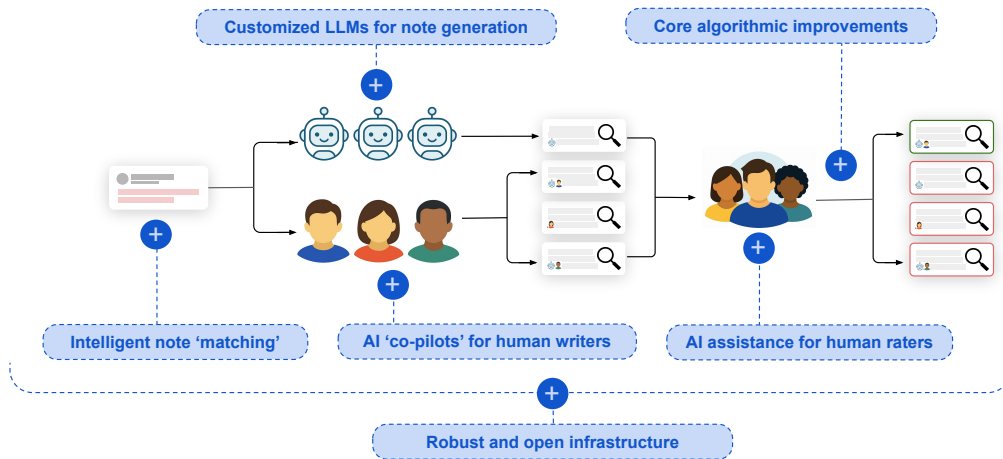


Figure 2: A research agenda for LLM-powered Community Notes: customized LLMs for note generation; AI “co-pilots” for human writers; AI assistance for human raters; intelligent note “matching” adapts existing helpful notes to new, similar contexts; evolving the core algorithm for AI-generated content; building a robust and open infrastructure.

5 A research agenda for LLM-powered Community Notes

To successfully build a thriving ecosystem, we propose a focused research agenda (see Figure 2).

5.1 Customized LLMs for note generation

The rich, informed judgments gathered through human oversight provide the ideal foundation for a constructive cycle of improvement. While sophisticated prompting of “Deep Research” models shows promise at producing high-quality initial notes, a significant opportunity lies in developing specialized post-training methods to further optimize them for the unique “bridging” demands of the Community Notes ecosystem. Recent work has already shown that LLMs can be designed not only to achieve individual alignment but also to foster consensus across groups (Bakker et al. 2022; Tessler et al. 2024; De et al. 2024). One promising direction could be to adapt the paradigm of Reinforcement Learning from Human Feedback (RLHF) to this unique community context, a process we can term Reinforcement Learning from Community Feedback (RLCF).

Unlike traditional RLHF (Ziegler et al. 2020; Ouyang et al. 2022), which often relies on simple pairwise preferences from a small group of labelers, RLCF would leverage the pluralistic signal of the entire community. The technical implementation would be more sophisticated: Instead of predicting a final “bridging” score directly, a reward model could be trained to predict the likely ratings of various diverse types of users, similar to the process outlined in De et al. (2024). The system could then aggregate these simulated ratings to calculate a note’s expected intercept score (i_n), which represents its predicted ability to bridge divides. This intercept would then serve as the core reward signal to fine-tune the LLM. Beyond simple ratings, RLCF can also easily incorporate

free-form rationales, such as user comments on posts or rating explanations from raters, by summarizing and integrating them into prompts to rapidly revise existing notes. Additionally, implementing RLCF in a continuous, online mode, where the reward model is constantly updated with new notes and ratings, would empower the LLM to adapt in real time to new events, novel ways that mislead users, and evolving community discourse and standards.

Crucially, though, a naive implementation of RLCF that only maximizes this predicted score would lead directly to the “homogenization” risk, as the model would quickly find a safe, bland optimum. This opens a critical research frontier: designing RLCF frameworks that balance two crucial objectives. The system must not only refine strategies already known to work, but also actively encourage experimentation to discover novel, potentially superior, approaches. For instance, the system could reward LLMs not just for high predicted scores, but also for generating notes that are stylistically novel or that explore different dimensions in their framing or research, learning from the human-authored examples.

The ultimate goal is to use the community’s feedback not just to ratify known good notes, but to discover new, potentially even better, ways of creating context. More generally, RLCF represents just one path forward. The open and rich dataset of notes, ratings, and feedback from Community Notes creates a unique testbed for developing entirely new post-training methods. Future research could explore models that are explicitly trained to leverage tool use, such as calling external data APIs or performing automated source verification, as part of their generation process, learning directly from this open data. This opens the door to creating a new class of transparent, verifiable, and pluralistically aligned models tailored for public discourse (Sorensen et al. 2024).

5.2 AI co-pilots for human writers

In the hybrid model, human writers are an essential part of the ecosystem. One potentially valuable research direction, therefore, is to build powerful AI-assisted writing tools that act as co-pilots. These tools could help human writers research topics faster, suggest relevant sources, generate initial drafts for editing, or polish the tone and clarity of a final note. This directly supports the human role as an innovator while also introducing an important challenge: how to empower human writers with co-pilots without flattening style and voice into homogeneous, bland writing. Studying this co-writing process will also provide invaluable data on how to best support and grow human contribution.

5.3 AI assistance for human raters

AI-powered tools could help human raters audit note claims more efficiently and evaluate potentially biased, rhetorically polished notes more critically and neutrally, rather than on surface-level persuasiveness (Saunders et al. 2022; Bridgers et al. 2024). While there is a potential risk that these tools could be erroneous in the same way LLM note writers

may be, there exist clever strategies to help mitigate it. For example, one concept would be to task multiple AIs to adversarially debate the merits of a note (Irving, Christiano, and Amodei 2018; Michael et al. 2023): One AI focuses on critiquing the note, checking its sources for misrepresentation, searching for counterevidence, and challenging the logical consistency of its argument, while another AI defends it. This adversarial process could help instantly surface potential flaws, hidden biases, or fabricated evidence, empowering the human rater to make a more informed judgment. Instead of starting from scratch, the rater now plays the role of an adjudicator—evaluating a structured clash of arguments. In the AI safety literature, this approach is known as scalable oversight: a strategy for amplifying human judgment through structured assistance, enabling reliable supervision of systems that may surpass human capabilities in speed or rhetorical fluency (Irving, Christiano, and Amodei 2018; Kenton et al. 2024).

5.4 Intelligent note matching and adaptation to new contexts

A critical bottleneck in the hybrid system is the potential for AI-generated notes to overwhelm the finite rating capacity of the human community. A promising research direction to mitigate this “rating-demand” imbalance is to develop methods for intelligently adapting and reapplying notes that have already been validated by human raters to new, similar contexts. Given that a claim may reappear verbatim, or with slight rephrasing, or attached to different media, this is a nontrivial challenge. Automatically matching notes to posts that people do not think need them could significantly undermine trust in the system. Research is needed to explore how to best leverage existing human-approved judgments. Key questions include:

- **Contextual matching:** How can we develop robust models to determine if a new post is semantically similar enough to an existing post (in its claim, media, or context) to warrant reusing an existing note, keeping in mind that sufficient similarity is a matter of rater perception in context?
- **AI-powered adaptation:** What are the safe boundaries for AI-driven modification? Can an AI autonomously make minor edits to a note to better fit a new context, or should it only propose changes for human review?
- **User experience and governance:** How should these adapted notes be presented to users? Should they be explicitly labeled as “adapted from a previously rated note”? What new rating mechanisms are needed? For example, should users be able to downvote the match itself, separate from the note’s content, if they feel the adaptation is inappropriate?

Successfully developing such a system would dramatically amplify the impact of every human rating, allowing the community’s judgment to scale more effectively across the information ecosystem.

5.5 Evolving the core algorithm for AI-generated content

The current “bridging” algorithm is calibrated on the idiosyncrasies of human-written notes. A flood of highly optimized, grammatically perfect LLM-generated notes could change the rating dynamics. Research will help determine to what degree the distribution of content shifts with LLM note writers, and whether it could benefit from algorithmic adaptation. This might involve incorporating metadata about a note’s origin (human, AI-assisted, or fully AI) or adjusting model parameters to prevent the stylistic polish of LLM notes from being overweighted against the substantive, nuanced contributions of human writers.

5.6 Building a robust and open infrastructure

A healthy ecosystem requires a new kind of infrastructure.

- **Human rater verification:** Integrity of the entire system is strengthened by the authenticity of its human participants and their engagements. The system should robustly work to increase the likelihood that ratings are from unique humans, where rating outcomes reflect the genuine viewpoints of people from different perspectives. This could include everything from using privacy-preserving methods to verify that each rating account is associated with one, distinct human being, to implementing limits to the influence any single account can have, to detecting and limiting the impact of accounts attempting to coordinate to influence outcomes. Such approaches have already been deployed on Community Notes systems, and can continue to evolve and expand (X Community Notes Guide, [n.d.-c](#); Meta 2025b).
- **An open note-writing API:** To foster innovation and trust, and to maintain the open, participatory nature of Community Notes, platforms could create an open API that allows any developer or organization to submit their LLM-generated notes for a given post.³ These notes would all enter the same rating pool, competing on a level playing field. Such a “competitive” environment with human contributors and multiple LLMs ensures no single writer can easily dominate note production or skew the note distribution. The community’s ratings would act as a natural selection mechanism, rewarding models that produce the most genuinely helpful content. Furthermore, appropriate per-writer quota limits would ensure that no LLM overwhelms the note pool or disproportionately influences the volume of available notes.
- **Prescreening and diversity curation:** As open APIs enable a higher volume of note submissions, one potential optimization is to introduce a prescreening layer that helps manage rater attention. This system would operate at two levels. First, it needs to strategically allocate rater attention across posts. As AI makes it possible for many more posts to have note candidates, the system must be more

3. X Community Notes has released an API in parallel with this paper; see X Community Notes Guide ([n.d.-a](#)).

selective in “nudging” raters to engage with the most impactful or contentious notes and posts. Second, for any given post, this layer could utilize techniques like semantic clustering to group nearly identical notes and select a small, diverse set for evaluation, ensuring raters are presented with meaningfully different options rather than five variations of the same argument. By managing attention at both the post level and the note level, this system could directly mitigate the risk of rater overload.

As CN’s coverage broadens—and potentially evolves from a platform-specific feature to a persistent overlay on the internet—these safeguards must be continually strengthened and evaluated to foster and protect quality and robustness against scaling risks such as gaming and rater overload.

6 Long-term vision and conclusion

The hybrid model we propose is not just an upgrade to Community Notes, but a blueprint for a new form of human-AI collaboration in the production of public knowledge. This extends beyond simple fact-checking into a dynamic, positive feedback loop. Imagine a system where humans contribute unique, firsthand observations that LLMs cannot access—photos from a local event, specialized domain knowledge, or on-the-ground reports. LLMs are then tasked with verifying this information against public sources and integrating it into a comprehensive, well-structured note. The human community, in turn, evaluates not just the note’s helpfulness but the fidelity of the LLM integration, creating a robust, self-correcting engine for generating trusted public information.

This knowledge-production engine could support collaborative online encyclopedias like Wikipedia or function as a new form of publication: distributed, peer-reviewed, and powered by LLMs. Ultimately, it could evolve beyond a single platform into a dynamic contextualization layer for the entire web. Such a layer could operate on specific URLs or, instead, it could operate on the content itself, using techniques like fuzzy matching or semantic hashing to attach context to claims, images, or videos as they spread across different domains. Imagine a browser-level system that analyzes the content of any page you visit and, drawing from this universal, continuously rated stream of notes, provides a nonintrusive overlay with relevant background, fact-checks, or warnings about a source’s bias.

Its ultimate importance becomes clear as we consider the trajectory of the internet itself. Recent advances suggest that LLMs may become increasingly capable of personalization (OpenAI 2025b; Zhang et al. 2025), potentially undermining the concept of a shared information space (Nehring et al. 2024). In a world where our primary interface to information is not a browser fetching distinct web pages, but a personal AI agent that synthesizes and generates content just for us—a personalized news feed, a summarized

search result, a conversational answer—every piece of content is mediated and optimized to align with our individual preferences, biases, and existing beliefs.

In such a future, where every individual's information diet is perfectly tailored, the role of a universal, nonpersonalized system like Community Notes could become even more important. Rather than simply a feature on a single platform, CN could evolve into part of the shared infrastructure that helps maintain a common information space. While personal AIs may optimize content for individual preferences, CN works in the opposite direction: It is publicly accessible and identical for everyone, and its bridging algorithm is designed to surface only the context that people from different perspectives agree is helpful. This makes CN uniquely suited to function as a shared epistemic layer, enabling information aggregation across otherwise fragmented realities.

In that world, CN could serve as a public, transparent, and pluralistic anchor—complementing private, personalized feeds with content that reflects a broader consensus. It wouldn't enforce a singular truth, but it could offer individuals a window into the consensus opinions on any one issue. This makes CN not just a tool for contextualizing posts, but a potential foundation for pluralistic AI alignment—one built on human deliberation, not just individual optimization.

The goal is not to create an AI assistant that tells users what to think, but to build an ecosystem that empowers humans to think more critically and understand the world better. At the heart of this paper is a simple idea: LLMs and humans can work together in a virtuous loop. LLMs help surface context faster, cheaper, and at greater scale while humans rate which notes are actually helpful. That pluralistic feedback can then be used to make LLMs better at producing the kind of context people value most. This feedback loop—Reinforcement Learning from Community Feedback (RLCF)—has the potential to improve both sides: giving people more trusted information, and making LLMs fundamentally better at providing it.

References

- Bakker, Michiel A., Martin J. Chadwick, Hannah Sheahan, Michael Henry Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, et al. 2022. “Fine-Tuning Language Models to Find Agreement Among Humans with Diverse Preferences.” In *Advances in Neural Information Processing Systems*. November 28, 2022. <https://arxiv.org/abs/2211.15006>.
- Bondarenko, Alexander, Denis Volk, Dmitrii Volkov, and Jeffrey Ladish. 2025. *Demonstrating Specification Gaming in Reasoning Models*, May 15, 2025. arXiv: 2502.13295 [cs.AI]. <https://arxiv.org/abs/2502.13295>.
- Bridgers, Sophie, Rishub Jain, Rory Greig, and Rohin Shah. 2024. *Human-AI Complementarity: A Goal for Amplified Oversight*. DeepMind Safety Research, December 23, 2024. Accessed June 29, 2025. <https://deepmindsafetyresearch.medium.com/human-ai-complementarity-a-goal-for-amplified-oversight-0ad8a44cae0a>.
- Burtch, Gordon, Dokyun Lee, and Zhichen Chen. 2024. “The Consequences of Generative AI for Online Knowledge Communities.” *Scientific Reports* 14, no. 1 (May 6, 2024): 10413. <https://doi.org/10.1038/s41598-024-61221-0>.
- Chuai, Yuwei, Haoye Tian, Nicolas Pröllochs, and Gabriele Lenzini. 2024. “Did the Roll-Out of Community Notes Reduce Engagement With Misinformation on X/Twitter?” In *Proceedings of the ACM on Human-Computer Interaction*, 8:1–52. CSCW2. ACM New York, NY, USA, November 8, 2024. <https://doi.org/10.1145/3686967>.
- De, Soham, Michiel A. Bakker, Jay Baxter, and Martin Saveski. 2024. “Supernotes: Driving Consensus in Crowd-Sourced Fact-Checking.” In *Proceedings of the ACM on Web Conference 2025*, 3751–61. November 9, 2024. <https://arxiv.org/abs/2411.06116>.
- Drolsbach, Chiara Patricia, Kirill Solovev, and Nicolas Pröllochs. 2024. “Community Notes Increase Trust in Fact-Checking on Social Media.” *PNAS Nexus* 3, no. 7 (May 31, 2024): pgae217. <https://doi.org/10.1093/pnasnexus/pgae217>.
- Flamini, Daniela. 2019. *Most Republicans Don’t Trust Fact-Checkers, and Most Americans Don’t Trust the Media*. Poynter Institute, July 3, 2019. Accessed June 29, 2025. <https://www.poynter.org/ifcn/2019/most-republicans-dont-trust-fact-checkers-and-most-americans-dont-trust-the-media/>.
- Google. 2025. *Gemini Deep Research*. Accessed June 29, 2025. <https://gemini.google/overview/deep-research/>.
- Irving, Geoffrey, Paul Christiano, and Dario Amodei. 2018. *AI Safety via Debate*, October 22, 2018. arXiv: 1805.00899 [stat.ML]. <https://arxiv.org/abs/1805.00899>.

- Kenton, Zachary, Noah Y. Siegel, János Kramár, Jonah Brown-Cohen, Samuel Albanie, Jannis Bulian, Rishabh Agarwal, et al. 2024. “On Scalable Oversight with Weak LLMs Judging Strong LLMs.” In *Advances in Neural Information Processing Systems*, 37:75229–76. Curran Associates, Inc., September 25, 2024. https://proceedings.neurips.cc/paper_files/paper/2024/file/899511e37a8e01e1bd6f6f1d377cc250-Paper-Conference.pdf.
- Mantzaris, Alexios. 2025. “The Top Community Notes Contributor Is...an Antivirus Startup?” Indicator, May 25, 2025. Accessed August 7, 2025. <https://indicator.media/p/the-top-community-notes-contributor-is-an-antivirus-startup-9354>.
- Meta. 2025a. “Community Notes: A New Way to Add Context to Posts.” Meta Transparency Center, April 7, 2025. <https://transparency.meta.com/features/community-notes>.
- . 2025b. “Testing Begins for Community Notes on Facebook, Instagram and Threads.” Meta Newsroom, March 13, 2025. <https://about.fb.com/news/2025/03/testing-begins-community-notes-facebook-instagram-threads/>.
- Michael, Julian, Salsabila Mahdi, David Rein, Jackson Petty, Julien Dirani, Vishakh Padmakumar, and Samuel R. Bowman. 2023. *Debate Helps Supervise Unreliable Experts*, November 15, 2023. arXiv: 2311.08702 [cs.AI]. <https://arxiv.org/abs/2311.08702>.
- Nehring, Jan, Aleksandra Gabryszak, Pascal Jürgens, Aljoscha Burchardt, Stefan Schaffer, Matthias Spielkamp, and Birgit Stark. 2024. “Large Language Models Are Echo Chambers.” In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 10117–23. Torino, Italia: ELRA / ICCL, May. <https://aclanthology.org/2024.lrec-main.884/>.
- OpenAI. 2025a. *Introducing Deep Research*. Accessed June 29, 2025. <https://openai.com/index/introducing-deep-research/>.
- . 2025b. *The Power of Personalized AI*. Accessed August 6, 2025. <https://openai.com/global-affairs/the-power-of-personalized-ai/>.
- Ouyang, Long, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, et al. 2022. “Training Language Models to Follow Instructions with Human feedback.” In *Advances in Neural Information Processing Systems*, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, 35:27730–44. Curran Associates, Inc., March 4, 2022. https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf.
- Presser, Adam. 2025. “Testing a New Feature to Enhance Content on TikTok.” TikTok Newsroom, April 16, 2025. Accessed August 20, 2025. <https://newsroom.tiktok.com/en-us/footnotes>.

- Saunders, William, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. 2022. *Self-Critiquing Models for Assisting Human Evaluators*, June 14, 2022. arXiv: 2206.05802 [cs.CL]. <https://arxiv.org/abs/2206.05802>.
- Schoenegger, Philipp, Francesco Salvi, Jiacheng Liu, Xiaoli Nan, Ramit Debnath, Barbara Fasolo, Evelina Leivada, et al. 2025. *Large Language Models Are More Persuasive Than Incentivized Human Persuaders*, May 21, 2025. arXiv: 2505.09662 [cs.CL]. <https://arxiv.org/abs/2505.09662>.
- Sharma, Mrinank, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Asbell, Samuel R. Bowman, Esin DURMUS, et al. 2024. "Towards Understanding Sycophancy in Language Models." In *The Twelfth International Conference on Learning Representations*. January 16, 2024. <https://openreview.net/forum?id=tvhaxkMKAn>.
- Slaughter, Isaac, Axel Peytavin, Johan Ugander, and Martin Saveski. 2025. *Community Notes Moderate Engagement with and Diffusion of False Information Online*, February 18, 2025. arXiv: 2502.13322 [cs.SI]. <https://arxiv.org/abs/2502.13322>.
- Sorensen, Taylor, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Miresghallah, Christopher Michael Rytting, Andre Ye, et al. 2024. "A Roadmap to Pluralistic Alignment." In *Proceedings of the 41st International Conference on Machine Learning*. ICML'24. Vienna, Austria: JMLR.org, July 21, 2024. <https://arxiv.org/abs/2402.05070>.
- Srinivasan, Siddarth, Ezra Karger, Michiel Bakker, and Yiling Chen. 2025. *Tell Me Why: Incentivizing Explanations*, February 19, 2025. arXiv: 2502.13410 [cs.GT]. <https://arxiv.org/abs/2502.13410>.
- Tessler, Michael Henry, Michiel A. Bakker, Daniel Jarrett, Hannah Sheahan, Martin J. Chadwick, Raphael Koster, Georgina Evans, Lucy Campbell-Gillingham, Tatum Collins, David C. Parkes, et al. 2024. "AI Can Help Humans Find Common Ground in Democratic Deliberation." *Science* 386, no. 6719 (October 18, 2024): eadq2852. <https://doi.org/10.1126/science.adq2852>.
- Wagner, Christian, and Ling Jiang. 2025. "Death by AI: Will Large Language Models Diminish Wikipedia?" *Journal of the Association for Information Science and Technology* 76, no. 5 (January 3, 2025): 743–51. <https://doi.org/10.1002/asi.24975>.
- Walker, Mason, and Jeffrey Gottfried. 2019. "Republicans Far More Likely Than Democrats to Say Fact-Checkers Tend to Favor One Side." Pew Research Center, June 27, 2019. Accessed June 29, 2025. <https://www.pewresearch.org/short-reads/2019/06/27/republicans-far-more-likely-than-democrats-to-say-fact-checkers-tend-to-favor-one-side/>.
- Wojcik, Stefan, Sophie Hilgard, Nick Judd, Delia Mocanu, Stephen Ragain, M. B. Fallin Hunzaker, Keith Coleman, and Jay Baxter. 2022. *Birdwatch: Crowd Wisdom and Bridging Algorithms Can Inform Understanding and Reduce the Spread of Misinformation*, October 27, 2022. arXiv: 2210.15723 [cs.SI]. <https://arxiv.org/abs/2210.15723>.

- X Community Notes Guide. n.d.-a. "AI Note Writer API." Accessed August 20, 2025. <https://communitynotes.x.com/guide/en/api/overview>.
- . n.d.-b. "Community Notes: A Collaborative Way to Add Helpful Context to Posts and Keep People Better Informed." Accessed August 20, 2025. <https://communitynotes.x.com/guide/en/about/introduction>.
- . n.d.-c. "Signing Up." Accessed August 20, 2025. <https://communitynotes.x.com/guide/en/contributing/signing-up>.
- xAI. 2025. *Grok 3 Beta – The Age of Reasoning Agents*. Accessed June 29, 2025. <https://x.ai/news/grok-3>.
- YouTube Team. 2024. "Testing New Ways to Offer Viewers More Context and Information on Videos." YouTube Official Blog, June 24, 2024. Accessed August 20, 2025. <https://blog.youtube/news-and-events/new-ways-to-offer-viewers-more-context/>.
- Zhang, Zhehao, Ryan A. Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, et al. 2025. "Personalization of Large Language Models: A Survey." *Transactions on Machine Learning Research* (May 31, 2025). ISSN: 2835-8856. <https://openreview.net/forum?id=tf6A9EYMo6>.
- Zhao, Andy, and Mor Naaman. 2023. "Insights from a Comparative Study on the Variety, Velocity, Veracity, and Viability of Crowdsourced and Professional Fact-Checking Services." *Journal of Online Trust and Safety* 2, no. 1 (September 21, 2023). <https://www.tsjournal.org/index.php/jots/article/view/118>.
- Zhou, Xinyi, Ashish Sharma, Amy X. Zhang, and Tim Althoff. 2024. *Correcting Misinformation on Social Media with a Large Language Model*, September 3, 2024. arXiv: 2403.11169 [cs.CL]. <https://arxiv.org/abs/2403.11169>.
- Ziegler, Daniel M., Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2020. *Fine-Tuning Language Models from Human Preferences*, January 8, 2020. arXiv: 1909.08593 [cs.CL]. <https://arxiv.org/abs/1909.08593>.

Authors

Haiwen Li is a PhD student in Social and Engineering Systems at MIT Institute for Data, Systems, and Society (IDSS).

(haiwenli@mit.edu)

Soham De is a PhD student in Information Science at the University of Washington.

(sohamde@uw.edu)

Manon Revel is a research affiliate at the Berkman Klein Center at Harvard Law School.

Andreas Haupt is a postdoctoral fellow in Economics and Computer Science at Stanford University.

(h4upt@stanford.edu)

Brad Miller is a member of the Community Notes team at X.

(bmiller1@cs.berkeley.edu)

Keith Coleman is a member of the Community Notes team at X.

(keithcoleman@gmail.com)

Jay Baxter is a member of the Community Notes team at X.

(jay.n.baxter@gmail.com)

Martin Saveski is an Assistant Professor at the University of Washington Information School.

(msaveski@uw.edu)

Michiel A. Bakker is an Assistant Professor at Massachusetts Institute of Technology.

(bakker@mit.edu)

Acknowledgements

We are grateful to Stephanie Chan, Jiansong Chao, Sophie Hilgard, Rishub Jain, and Daniel Ortiz for helpful conversations and feedback on earlier drafts. We also thank the anonymous reviewers for their thoughtful comments and suggestions.

Brad Miller, Keith Coleman, and Jay Baxter are employees of X and work on the Community Notes team.

Keywords

Community notes; large language models (LLMs); AI safety.