OSINT Research Platform as Accountability Infrastructure for Collaborative Sensemaking

Anirban Mukhopadhyay
Virginia Tech
Department of Computer Science
Virginia, USA

Kelechi Eze
Virginia Tech
Department of Computer Science
Virginia, USA

Kurt Luther
Virginia Tech
Department of Computer Science
Virginia, USA

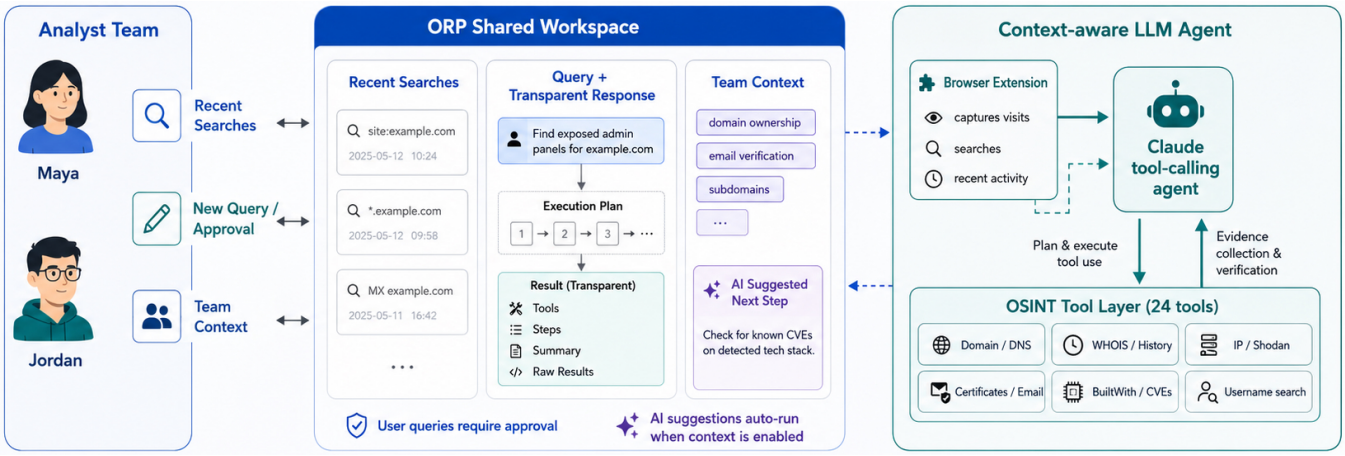


Figure 1: Overview of the OSINT Research Platform (ORP). The system has three main components for supporting collaborative investigations. (1) User interaction: team members submit queries, review prior searches, approve user-initiated tool plans, and observe what others are working on. (2) ORP Shared Workspace: the interface combines recent searches, transparent LLM agent execution plans and results, and team context, where browser activity is summarized and AI-generated next steps are shown. (3) Context-Aware LLM Agent: a browser extension supplies recent visits and searches to the agent, which uses passive external tools to run queries. Context-based queries run automatically when Context mode is enabled.

Abstract

As AI agents move from answering questions to observing teams, calling tools, and proposing next steps, it becomes difficult to say what the agent contributed, and how the team leveraged those contributions in their final result. Here, we present the OSINT Research Platform (ORP) as a design provocation for accountability infrastructure. ORP is a shared workspace for small teams conducting Open-Source Intelligence (OSINT) vulnerability assessments. A browser extension provides lightweight team context; a tool-calling LLM agent uses a bounded set of passive OSINT tools; and the interface exposes plans, tool calls, results, sources, and proactive suggestions. Our provocation is that accountability should be designed into the division of labor while work unfolds, rather than reconstructed from a final output or an AI-use log. ORP shows how systems can observe team members, record it, and provide more contextual support and metadata. However, these observations create privacy, surveillance, and interpretive risks.

Keywords

Human–AI collaboration; Accountability; OSINT; Context-Aware Agents

ACM Reference Format:
Anirban Mukhopadhyay, Kelechi Eze, and Kurt Luther. 2026. OSINT Research Platform as Accountability Infrastructure for Collaborative Sensemaking. In *Proceedings of CSCW 2026 Workshop on Contribution, Originality, and Accountability in Agent-Mediated Collaboration (CSCW '26).* ACM, New York, NY, USA, 3 pages. https://doi.org/XX.XXX

1 Introduction

AI systems are moving beyond answering prompts toward observing activity, accessing external data, calling tools, and completing multi-step tasks with greater autonomy [8]. In complex and collaborative work, a final result may combine human goals, AI-generated plans and responses, and human edits [2, 9]. As these elements become intertwined, it becomes harder to determine what the AI contributed, what people contributed, and who is accountable for the outcome.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CSCW '26, Salt Lake City, UT
© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
https://doi.org/XX.XXX

This problem is particularly relevant in collaborative sensemaking [6]. Teams must work through large amounts of information, follow several leads at once, and revise their understanding as new evidence appears. AI can support this process through multimodal interpretation, external tool use, summarization, and proactive suggestions. However, these same capabilities make the resulting division of labor difficult to inspect and have meaningful human oversight.

We present the OSINT Research Platform(ORP), as a design provocation for accountability infrastructure. ORP makes parts of team and agent activity visible within a shared workspace. Our central argument is that accountability should be designed into the evolving workflow, not added only through a final disclosure that AI was used.

2 Background: OSINT as Collaborative Sensemaking

Open-source intelligence (OSINT) serves as the task context for this system and it involves collecting, verifying, and analyzing publicly available information for an investigative purpose [7]. It is used in journalism, fact-checking, human rights investigations, cybersecurity, and crisis response. The work requires analysts to move across search engines, websites, databases, and technical tools, while judging the quality and relevance of raw evidence.

OSINT is often collaborative because teams can divide a broad problem into smaller areas, pursue multiple leads in parallel, and combine different forms of expertise. Accountability is especially important in this domain. Findings may affect real people or organizations, and incorrect associations can cause harm. Investigations must also remain within an appropriate ethical scope. Therefore, systems that support OSINT need meaningful human oversight over AI use.

3 OSINT Research Platform (ORP)

Our earlier studies suggest that AI support for collaborative OSINT needs both action and context [3, 4]. Tool-calling agents can move beyond templates and suggestions by collecting real-world information through tools, while context awareness can help them understand what team members are already doing and when support may be useful [1, 5]. These capabilities also allow agents to take different roles, such as assisting with a specific task, suggesting a next step, or supporting coordination. However, their value depends on accuracy and how well their contributions fit the team's evolving work.

ORP was designed around this need. It is a collaborative web application that combines a shared workspace, a tool-calling LLM agent, and a browser extension that captures investigation-related activity. Team members can submit natural-language questions, inspect recent searches, review agent plans, examine evidence, and see a summary of team context. The schematic diagram is presented in Figure 1.

For user-initiated questions, the agent first creates an execution plan. The interface shows the intended actions and allows the user to approve, edit, or cancel them. After execution, ORP presents the response along with tools used, intermediate steps, and raw outputs from tools. This supports a task-oriented agent role with transparent actions and user control.

The browser extension provides context from work that happens outside the platform, such as relevant searches and visited pages. ORP uses these events to support team awareness and to generate suggestions that complement current work. In this way, the agent can shift between assisting with a specific query and offering lightweight coordination support. The system remains bounded to passive information gathering based on tools, while human analysts define scope, interpret evidence, and decide what is part of the final report.

4 Scenario: Collaborative Investigation with ORP

Consider Maya and Jordan, two analysts examining an organization's public online presence. Maya focuses on ownership and infrastructure, while Jordan examines public employee information and visible website technologies. They divide the work themselves and enable context sharing within their ORP workspace.

Maya asks the agent what public information is available about the history and ownership of a domain. ORP proposes a plan and shows the actions it intends to take. Maya reviews and approves the plan. The agent then gathers public records, organizes the results, and returns source-linked evidence. The result distinguishes Maya's question, the agent's planning, external tool calls, and the agent's summary.

Meanwhile, Jordan searches external websites. Selected browsing events appear in the shared context panel, allowing Maya to see Jordan's current focus without a separate status update. These events are treated as team activity. Visiting a page does not mean Jordan accepted its content or intends to use it.

As the investigation develops, ORP uses this shared context to identify connections across the team's work. For example, after observing Jordan examining a technology associated with the organization's website and Maya investigating related infrastructure, the agent may suggest checking whether the same technology appears across other public-facing assets. The analysts can follow, modify, or ignore the suggestion. If they pursue it, the subsequent searches, tool calls, and evidence become part of the shared investigation history.

This interaction illustrates how collaborative work in ORP produces more than a final AI response. The system captures traces of human questions, browser activity, agent plans, approvals, tool observations, and suggestions through which an investigative direction develops. This creates a partial account of how the investigation evolved and where the AI entered the team's workflow.

5 Accountability Infrastructure

ORP supports accountability first by bounding the agent's actions. Tool calling gives an AI access to external information, but the available tools define the scope of that access. A known set of passive tools makes it easier for team members to understand what the agent can do and reduces the risk of actions that exceed the investigation's ethical boundaries.

Visible plans and execution records provide a second layer. Rather than showing only a response summary, ORP exposes proposed actions, selected tools, parameters, intermediate outputs, and sources. These records help analysts inspect how a response was produced and allows meaningful human oversight.

Browser events and shared context provide a third layer by representing work outside the AI interface. This can reveal that a person was already following a lead before the agent suggested it. Context therefore supports a more accurate account of how ideas and investigative directions emerged. The system could go further by recording whether suggestions were accepted, rejected, or followed up on.

6 Challenges and Open Questions

Capturing browser activity creates privacy risks. Team members may open unrelated pages, expose sensitive personal details, or pursue tentative leads they do not wish to share. Users need to understand exactly what the system records and should be able to pause or withhold activity. Without such controls, a coordination feature could become workplace surveillance.

Context interpretation also introduces failure modes. A visited page may not reflect a person's current intent, and an old search may describe a lead that has already been abandoned. The agent may infer the wrong team focus or connect unrelated events. ORP should therefore separate observed events from AI-generated interpretations to allow participants to form their own opinions.

Accountability records like browsing events and tool use should be treated as partial representations rather than complete accounts of the collaboration. This is because of the over representation of work that produces technical traces in the system compared to team process oriented contributions such as coordination or domain judgment.

Future studies should evaluate key questions including whether people understand how context-aware suggestions are generated, whether records help teams understand their AI usage better, and whether the infrastructure makes some actions more visible while obscuring others.

7 Conclusion

This paper proposes ORP as a design provocation for accountability infrastructure in collaborative sensemaking. Its bounded tools, visible plans, shared context, and activity records show how systems might make evolving human and AI contributions more tractable. At the workshop, I would like to discover other accountability infrastructures that are aimed at addressing the challenges across other domains and forms of agent-mediated work.

References

[1] Michelle S Lam, Omar Shaikh, Hallie Xu, Alice Guo, Diyi Yang, Jeffrey Heer, James A Landay, and Michael S Bernstein. 2026. Just-In-Time Objectives: A General Approach for Specialized AI Interactions. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. 1–26.

[2] Yaxi Lu, Shenzhi Yang, Cheng Qian, Guirong Chen, Qinyu Luo, Yesai Wu, Huadong Wang, Xin Cong, Zhong Zhang, Yankai Lin, Weiwen Liu, Yasheng Wang, Zhiyuan Liu, Fangming Liu, and Maosong Sun. 2024. Proactive Agent: Shifting LLM Agents from Reactive Responses to Active Assistance. arXiv:2410.12361 [cs]. doi:10.48550/arXiv.2410.12361

[3] Anirban Mukhopadhyay and Kurt Luther. 2025. OSINT Clinic: Co-designing AI-Augmented Collaborative OSINT Investigations for Vulnerability Assessment. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 868, 22 pages. doi:10.1145/3706598.3713283

[4] Anirban Mukhopadhyay and Kurt Luther. 2025. Tailoring Generative AI to Augment Creative Leadership in Capture-The-Flag Development. In *CHI'2025 Workshop on Tools for Thought: Research and Design for Understanding, Protecting, and Augmenting Human Cognition with Generative AI. Yokohama, Japan.*

[5] Omar Shaikh, Shardul Sapkota, Shan Rizvi, Eric Horvitz, Joon Sung Park, Diyi Yang, and Michael S Bernstein. 2025. Creating general user models from computer use. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. 1–23.

[6] Sukrit Venkatagiri, Anirban Mukhopadhyay, David Hicks, Aaron Brantly, and Kurt Luther. 2023. CoSINT: Designing a Collaborative Capture the Flag Competition to Investigate Misinformation. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) *(DIS '23)*. Association for Computing Machinery, New York, NY, USA, 2551–2572. doi:10.1145/3563657.3595997

[7] Heather J Williams and Ilana Blum. 2018. Defining second generation open source intelligence (OSINT) for the defense enterprise. (2018).

[8] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, Rui Zheng, Xiaoran Fan, Xiao Wang, Limao Xiong, Yuhao Zhou, Weiran Wang, Changhao Jiang, Yicheng Zou, Xiangyang Liu, Zhangyue Yin, Shihan Dou, Rongxiang Weng, Wensen Cheng, Qi Zhang, Wenjuan Qin, Yongyan Zheng, Xipeng Qiu, Xuanjing Huang, and Tao Gui. 2023. The Rise and Potential of Large Language Model Based Agents: A Survey. arXiv:2309.07864 [cs]. http://arxiv.org/abs/2309.07864

[9] Shuning Zhang, Hui Wang, and Xin Yi. 2025. Exploring Collaboration Patterns and Strategies in Human-AI Co-creation through the Lens of Agency: A Scoping Review of the Top-tier HCI Literature. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (Oct. 2025), 1–43. doi:10.1145/3757594