First International Workshop on Conversational Interfaces for Agentic Web browsing: Are We Ready to Move Beyond GUIs?

LUDOVICA PIRO, Politecnico di Milano, Italy
HELOISA CANDELLO, Inteli, Brazil
MARISTELLA MATERA, Politecnico di Milano, Italy
CHRISTINE MURAD, Department of Computer Science, Carleton University, Canada and Department of Systems Design Engineering, University of Waterloo, Canada

As AI features are being integrated in web browsers and search engines, web browsing is shifting from being a paradigm relying on rich Graphical User Interfaces to one that integrates more and more conversational interaction. However, Conversational Web Browsing (CWB), still needs to be formalized from an interaction design standpoint. In this workshop we want to discuss what are the pressing issues, or the possible gaps in current conversational interaction design guidelines that still need to be addressed to account for the complexities offered by web interfaces. The workshop aims to advance a research agenda for understanding how to design inclusive conversational web browsing, identifying key design challenges, opportunities, and open questions for the CUI community.

CCS Concepts: • **Human-centered computing** → **Natural language interfaces**; **Interaction design process and methods**.

Additional Key Words and Phrases: Natural language interfaces, agentic AI, conversational web browsing

ACM Reference Format:
Ludovica Piro, Heloisa Candello, Maristella Matera, and Christine Murad. 2018. First International Workshop on Conversational Interfaces for Agentic Web browsing: Are We Ready to Move Beyond GUIs?. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 8 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Motivation

The web has been, for decades, primarily a visual experience. Modern websites have evolved to include rich dynamic content, interactive components, and multimedia content that users access visually through structured Graphical User Interfaces (GUIs).

During the first wave of chatbots of the 2010s, there was the proliferation of customer care chatbots and other specialized bots that provided users with additional help and features through dedicated chats [3]. Nonetheless, due to the difficulty in generalizing models and the variability of web tasks, the web has still remained mostly experienced through a GUI, with only a few works exploring multimedia features for accessibility [5, 15].

Authors' Contact Information: Ludovica Piro, ludovica.piro@polimi.it, Politecnico di Milano, Milano, Italy; Heloisa Candello, candelloh@acm.org, Inteli, São Paulo, Brazil; Maristella Matera, maristella.matera@polimi.it, Politecnico di Milano, Milan, Italy; Christine Murad, c2murad@uwaterloo.ca, Department of Computer Science, Carleton University, Ottawa, Ontario, Canada and Department of Systems Design Engineering, University of Waterloo, Waterloo, Ontario, Canada, cpalmer@prl.com.

Manuscript submitted to ACM

In recent years, the rise of LLM applications has brought forward a different paradigm for web browsing: natural language interaction. This is not an entirely new concept, as various works in web accessibility research have explored various approaches to enable voice interaction with the web [2, 5, 8, 15]. The work by Borodin et al. [5], for example, proposed *HearSay* a non-visual web browser that could handle navigation and form filling. Skipping banners and menus, the browser would segment the page into semantically related sections and read aloud the page starting from the most relevant section in order to save browsing time. Similarly, in [2], the authors propose a paradigm to have a goal-oriented interaction with the web. In their approach, they propose a model to make the website *bot-friendly* without needing to develop a custom bot: the website is annotated with custom annotations to instruct the bot, that can access the content present in the page itself. In this paradigm users can access content in a random access fashion: they can submit specific requests and commands without the need to navigate the whole page to get the information they need.

Based on this approach, Pucci et al. [17] identified design patterns for *conversational web browsing* (CWB) through co-design sessions and interviews involving blind and vision impaired participants. The user research led to the definition of four classes of patterns:

- **Orientation:** this class includes patterns that help users in creating a mental model of the website and navigate it. The patters inform how to keep the user informed of what the possible actions on the page are, how to convey through dialogue an overview of the page, or how to provide navigational context upon landing on a webpage or before leaving it.
- **Navigation:** this class supports users browsing and navigation actions, providing patterns to support both structured browsing and direct access to information. Users are supported in issuing pointed questions bypassing the hierarchy of a webpage, or in exploring the hierarchy both in-depth (diving deeper into a topic) and in-breath (exploring different topics at the same level of detail).
- **Content reading:** this class pertains pattern that suggest how to summarize, index, or segment content. Its patterns inform how to present web content at different levels of detail to support both skimming mechanisms and the possibility to access the full original content of the webpage.
- **Scaffolding:** this last class pertains patterns that support conversation control mechanisms such as default actions like invoking help, repeating a piece of information, or starting and stopping the conversational agent. Through these patterns, users control the interaction with the conversational agent itself.

Now, the unprecedented flexibility in dialogue offered by LLMs is making natural-language interaction with the Web far more practical and generalizable than previously possible. Indeed, LLMs can interpret users' high-level goals across a variety of intents spanning different domains and varying in complexity [4, 10, 20]. This introduces a new layer of abstraction between users and web interfaces, where tasks can be expressed in natural language and delegated to intelligent agents that navigate, retrieve, and synthesize information on the user's behalf. Conversational web browsing is gaining momentum as a paradigm for accessing and manipulating web content through natural language interaction, either via voice or text [16, 17].

We can see this in Google AI mode or OpenAI's Deep Research [14], where the agent retrieves information from the web and provides a synthesis of multiple sources. Even more, applications such as Browser Use and agentic Web browsers such as Dia[1], Comet[2], or BrowserOS[3] promise to automate complex workflows, such as planning trips, and to operate almost completely independently on the web in place of human users [7].

[1] Dia browser
[2] Comet
[3] BrowserOS

However, CWB still needs to be formalized from an interaction design standpoint. In fact, while the CUI community has produced many informative guidelines [13], very few works focus on browsing content via conversation on the web [17]. Thus, existing conversational design guidelines may not sufficiently account for the structure, interactivity, and complexity of web browsing and searching environments or address issues like privacy and trust[11], user agency [1], and inclusivity [16]. For the first issue, privacy and trust perspective, the literature highlights several open challenges [18, 19]. The first challenge is related to managing sensitive data. In fact, to perform tasks such as making purchases or planning events, agents often require access to credit card numbers and other highly private data. In this context it is fundamental to determine how this data should be managed by the LLM-based agent, and when it agent should disclose such information during interactions [19]. A second challenge concerns the influence of web interfaces on agent behavior: for example, GUIs dark patterns may bias the agent's decision-making, potentially steering it toward undesirable outcomes, such as selecting premium options [18]. Understanding how to manage user data and the interplay between agents and web interfaces is therefore essential to achieve trustworthy agents.

Understanding how to ensure user agency, instead, is essential to maintain alignment between agents actions and user goals. In fact, users input can be ambiguous and their preference may change over time [18]. Thus, accommodating for user intervention is necessary to let users retain control over the interaction. Moreover the work by Chan et al. [9], shows how supporting user agency is also fundamental for establishing accountability, as the responsibility for preventing harmful outcomes ultimately remains with human users.

Lastly, when it comes to inclusivity, while early studies on conversational interaction with the web were rooted in web accessibility research, the needs of people with disabilities and of other underrepresented user groups, like children and elderly people, are currently underexplored in this context [6, 12].

Therefore, the aim of this workshop is to promote research on the topic to better understand the challenges of designing inclusive conversational web browsing. Through this workshop, we want to connect the conversational interaction and the web accessibility communities, engaging them in discussions and activities to envision how to design a web that is inclusive and conversation-first.

Our focus is articulated into four main themes:

- **Theme 1: Designing for the new paradigm for web browsing.** What are pressing design challenges for the new conversational Web and agentic Web browsing? How should conversational systems integrate with existing web interfaces and workflows?
- **Theme 2: Inclusivity and Accessibility.** How can we develop an inclusive conversational web browsing experience? What are possible new barriers that may be introduced with this new paradigm?
- **Theme 3: Interaction Models for the Agentic Web.** How should interaction paradigms evolve when intelligent agents can autonomously navigate and act on the Web on behalf of users?
- **Theme 4: Limitations of Current Solutions.** What aspects of web interaction are not covered by current solutions (e.g., AI-assisted browsing tools and agentic browsers)?
- **Theme 5: The Use of Voice and Speech in the Agentic Web.** How can the modality of speech further enhance the interaction capabilities of the agentic web?

2 Expected contributions and dissemination

We invite participants to submit short position papers (2-3 pages) describing research ideas, design explorations, systems, or critical perspectives related to conversational and agentic web browsing. Submissions may include:

- Theoretical perspectives on conversational web interaction.
- Design frameworks and interaction models for conversational web browsing.
- Applications, systems, or prototypes enabling conversational access to web content.
- Design guidelines and methodologies for conversational interaction with complex web environments.
- Empirical studies evaluating conversational browsing experiences.
- Critical reflections on limitations, risks, and societal implications of agentic web systems.

Through these contributions, the workshop aims to advance a research agenda for inclusive conversational web browsing, identifying key design challenges, opportunities, and open questions for the CUI community.

Prospective participants will be invited to submit either a position paper (2-3 pages) or a statement of interest (1-2 pages). We highly recommend that all contributors follow the ACM Guide for Accessible Submissions during the drafting process and perform a comprehensive accessibility audit before final submission.

To ensure participation, the authors will promote the workshop by setting up the workshop website and through relevant mailing lists, social media platforms, and through relevant SIG Chapters.

3 Workshop format and schedule

We would like to propose a half-day, in-person, interactive workshop, as a long session (180 minutes), with an expected attendance of 20 to 30 participants.

The workshop will follow this tentative schedule:

- **Introduction from the organizers (15 min)**: the organizers will provide a brief introduction of themselves and the workshop's aims.
- **Lightning presentations (40 min)**: participants will present themselves and, in 1 minute, briefly introduce their position paper. This moment will be used to sort people into groups for the following discussion. We aim to group people interested in the same theme while trying to preserve multidisciplinary.
- **Discussion in groups (30 min):** in thematic groups of 5 to 8 people, participants will discuss current gaps in the design of CWB and possible approaches to address them. The discussion will be guided by provocative questions supplied by the workshop organizers, for example, on current conversational interaction guidelines and their adequacy for emerging Web browsing paradigms.
- **Group activity (30 min):** practical design activity to speculate about applications of conversational web browsing and the future of agentic browsing. Attendees will be asked to envision how the gap identified during the discussion could be addressed in practice. They will be provided with a pen and paper to create a storyboard illustrating a scenario in which an agent or user interacts with the web to accomplish a CWB task. The storyboard should serve as a practical enactment of a best practice or solution identified by the group, showing how they would like agents to interact with web content and users in the future.
- **Plenary reporting and future plans:** all groups will present the outcome of their activity to share with everybody how they envision the future of CWB.

3.1 Accessibility and inclusion

The organizers will work to make the session as inclusive as possible. We intend to create a discussion environment where everybody will feel valued and comfortable in sharing their opinion. Groups will be created to try to balance different backgrounds and encourage multidisciplinarity. We will work together with Diversity & Accessibility Chairs,

Workshop Chairs, and Student Volunteers to meet participants' accessibility needs before and during the workshop and conference.

3.2 Post workshop plans

Accepted submissions will be published on the workshop website in advance and archived in the proceedings for long-term reference. The website will serve as an archive for the materials produced during the workshop. Attendees will be invited to author a collaborative white paper presenting the outcome of the workshop. The proposers are also considering organizing a journal special issue (*ACM Transactions on the Web* and *International Journal of Human-Computer Interaction* are a candidate journal).

4 Organizers

The workshop is led by an interdisciplinary team of experts whose specialized knowledge ensures a deep alignment with our core objectives. The organizers have widely published in domains related to Human-Computer Interaction, Conversational User Interfaces, and Web Engineering.

- **Ludovica Piro** is a PostDoc researcher at Politecnico di Milano, specializing in Conversational User Interfaces and web accessibility. She has a multidisciplinary background having obtained a master degree in Communication Design and a PhD in Computer Science. Her work explores the intersection of conversational interaction and accessible web browsing.
- **Heloisa Candello** is an Associate Professor at Inteli, and worked as a Senior Research Scientist for more than 10 years in IBM Research. She is specializing in Conversational User Interfaces and responsible AI. Her work explores the design and evaluation of ethical, engaging AI interactions, with contributions published in top HCI venues and patented innovations in conversational systems. She has organized numerous ACM (CHI, CSCW, CUI, IUI) workshops and serves as a steering committee co-chair of the CUI conference series. Heloisa will serve as the web chair for this workshop.
- **Maristella Matera** is a full professor of Computer Engineering at Politecnico di Milano and leads the Human-Centric Interactive Technology (HINT) Lab. She has 25+ years of experience in HCI research, with contributions at the intersection of user interface design and web engineering. Her current research interests focus on AI technologies for multimodal systems, with an emphasis on ensuring that emerging technologies are aligned with human needs and values. The impact of her research on society and inclusivity has been recognized by two Google Awards for research on Society-centered AI (2023 and 2025).
- **Christine Murad** is a Post-Doctoral Fellow in the SPIRL lab at Carleton University, and a Research Associate at the University of Waterloo, in the Technologies for Aging Gracefully lab. Christine is currently based in Waterloo, Ontario (Canada). Christine earned her PhD (2024), MSc (2019), and her Honours B.Sc in Computer Science from the University of Toronto. Christine's research involves exploring the design of human-centered AI, with a focus on conversational voice interfaces and the development of design heuristics and tools to assist in conversational voice interface design. Christine is a founding member of the ACM Conversational User Interfaces (CUI) Steering Committee since its inception in 2019.

5 Call for participation

For decades, the web has been a primarily visual experience. However, the rise of Large Language Models (LLMs) is shifting this paradigm toward Conversational Web Browsing (CWB). While intelligent agents now automate complex workflows, we lack formal interaction design frameworks that ensure these experiences are private, ethical, and inclusive.

This workshop bridges the Conversational Interaction and Web Accessibility communities to envision a "conversation-first" web. We invite participants to submit short position papers (2-3 pages) describing research ideas, design explorations, systems, or critical perspectives related to conversational and agentic web browsing. Papers will be evaluated based on relevance to the topics of the workshop, which are:

- Theme 1: Designing for the new paradigm for web browsing.
- Theme 2: Inclusivity and Accessibility.
- Theme 3: Interaction Models for the Agentic Web.
- Theme 4: Limitations of Current Solutions.
- Theme 5: The Use of Voice and Speech in the Agentic Web.

Submissions may include:

- Theoretical perspectives on conversational web interaction.
- Design frameworks and interaction models for conversational web browsing.
- Applications, systems, or prototypes enabling conversational access to web content.
- Design guidelines and methodologies for conversational interaction with complex web environments.
- Empirical studies evaluating conversational browsing experiences.
- Critical reflections on limitations, risks, and societal implications of agentic web systems.

Through these contributions, the workshop aims to advance a research agenda for inclusive conversational web browsing, identifying key design challenges, opportunities, and open questions for the CUI community.

Prospective participants will be invited to submit either a position paper (2-3 pages) or a statement of interest (1-2 pages). Preference will be given to position paper submissions. Authors may use the ACM single column template to draft their submissions. We highly recommend that all contributors follow the ACM Guide for Accessible Submissions during the drafting process and perform a comprehensive accessibility audit before final submission. Accepted contributions will be published on our Website.

Key dates below:

- *Submission deadline:* May 29th 2026
- *Acceptance Notification:* June 12th 2026
- *Send position papers to:* this Google form
- *Website:* Workshop on Conversational Interaction for Agentic Web Browsing
- *Workshop's Date and time:* July 21st 2026

6 References

References

[1] Sanju Ahuja and Jyoti Kumar. 2022. Assistant or Master: Envisioning the User Autonomy Implications of Virtual Assistants. In *Proceedings of the 4th Conference on Conversational User Interfaces* (Glasgow, United Kingdom) *(CUI '22)*. Association for Computing Machinery, New York, NY, USA, Article 4, 5 pages. doi:10.1145/3543829.3544514

[2] Marcos Baez, Florian Daniel, and Fabio Casati. 2019. *Conversational Web Interaction: Proposal of a Dialog-Based Natural Language Interaction Paradigm for the Web*. Springer-Verlag, Berlin, Heidelberg, 94–110. doi:10.1007/978-3-030-39540-7_7

[3] Marcos Baez, Florian Daniel, Fabio Casati, and Boualem Benatallah. 2021. Chatbot Integration in Few Patterns. *IEEE Internet Computing* 25, 3 (2021), 52–59. doi:10.1109/MIC.2020.3024605

[4] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2022. On the Opportunities and Risks of Foundation Models. arXiv:2108.07258 [cs.LG] https://arxiv.org/abs/2108.07258

[5] Yevgen Borodin, Jalal Mahmud, I. V. Ramakrishnan, and Amanda Stent. 2007. The HearSay non-visual web browser. In *Proceedings of the 2007 International Cross-Disciplinary Conference on Web Accessibility (W4A)* (Banff, Canada) *(W4A '07)*. Association for Computing Machinery, New York, NY, USA, 128–129. doi:10.1145/1243441.1243444

[6] Stacy M. Branham and Antony Rishin Mukkath Roy. 2019. Reading Between the Guidelines: How Commercial Voice Assistant Guidelines Hinder Accessibility for Blind Users. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh PA USA, 2019-10-24). ACM, 446–458. doi:10.1145/3308561.3353797

[7] BrowserOS. [n. d.]. *BrowserOS - Open-Source AI Browser | Privacy-First Alternative*. https://www.browseros.com/

[8] Julia Cambre, Alex C Williams, Afsaneh Razi, Ian Bicking, Abraham Wallin, Janice Tsai, Chinmay Kulkarni, and Jofish Kaye. 2021. Firefox Voice: An Open and Extensible Voice Assistant Built Upon the Web. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 250, 18 pages. doi:10.1145/3411764.3445409

[9] Alan Chan, Rebecca Salganik, Alva Markelius, Chris Pang, Nitarshan Rajkumar, Dmitrii Krasheninnikov, Lauro Langosco, Zhonghao He, Yawen Duan, Micah Carroll, Michelle Lin, Alex Mayhew, Katherine Collins, Maryam Molamohammadi, John Burden, Wanru Zhao, Shalaleh Rismani, Konstantinos Voudouris, Umang Bhatt, Adrian Weller, David Krueger, and Tegan Maharaj. 2023. Harms from Increasingly Agentic Algorithmic Systems. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) *(FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 651–666. doi:10.1145/3593013.3594033

[10] Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. 2023. Mind2Web: Towards a Generalist Agent for the Web. arXiv:2306.06070 [cs.CL] https://arxiv.org/abs/2306.06070

[11] Trung Dong Huynh, William Seymour, Luc Moreau, and Jose Such. 2023. Why Are Conversational Assistants Still Black Boxes? The Case For Transparency. In *Proceedings of the 5th International Conference on Conversational User Interfaces* (Eindhoven, Netherlands) *(CUI '23)*. Association for Computing Machinery, New York, NY, USA, Article 52, 5 pages. doi:10.1145/3571884.3604319

[12] Cosmin Munteanu, Sayan Sarcar, Jaisie Sin, Christina Ziying Wei, Sergio Sayago, Wei Zhao, Jenny Waycott, Rezvan Boostani, Moojan Ghafurian, and Cristina Getson. 2025. Designing Age-Inclusive Interfaces: Emerging Conversational and Generative AI to Support Interactions across the Life Span. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces (CUI '25)*. Association for Computing Machinery, New York, NY, USA, Article 37, 5 pages. doi:10.1145/3719160.3728626

[13] Christine Murad, Heloisa Candello, and Cosmin Munteanu. 2023. What's The Talk on VUI Guidelines? A Meta-Analysis of Guidelines for Voice User Interface Design. In *Proceedings of the 5th International Conference on Conversational User Interfaces* (Eindhoven, Netherlands) *(CUI '23)*. Association for Computing Machinery, New York, NY, USA, Article 19, 16 pages. doi:10.1145/3571884.3597129

[14] OpenAI. [n. d.]. *Deep research | OpenAI API*. https://developers.openai.com/api/docs/guides/deep-research/

[15] Javier Pérez. 2019. Handsfree for Web: A Google Chrome extension to browse the web via voice commands. In *Proceedings of the 16th International Web for All Conference* (San Francisco, CA, USA) *(W4A '19)*. Association for Computing Machinery, New York, NY, USA, Article 12, 2 pages. doi:10.1145/3315002.3332443

[16] Ludovica Piro, Himanshu Verma, and Maristella Matera. 2025. Conversational Web Browsing: Where Do Existing DesignGuidelines Fall Short?. In *Proceedings of the 16th Biannual Conference of the Italian SIGCHI Chapter (CHItaly '25)*. Association for Computing Machinery, New York, NY, USA, Article 21, 9 pages. doi:10.1145/3750069.3750161

[17] Emanuele Pucci, Isabella Possaghi, Claudia Maria Cutrupi, Marcos Baez, Cinzia Cappiello, and Maristella Matera. 2023. Defining Patterns for a Conversational Web. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 118, 17 pages. doi:10.1145/3544548.3581145

[18] Yingxuan Yang, Mulei Ma, Yuxuan Huang, Huacan Chai, Chenyu Gong, Haoran Geng, Yuanjian Zhou, Ying Wen, Meng Fang, Muhao Chen, Shangding Gu, Ming Jin, Costas Spanos, Yang Yang, Pieter Abbeel, Dawn Song, Weinan Zhang, and Jun Wang. 2025. Agentic Web: Weaving the Next Web with AI Agents. arXiv:2507.21206 [cs.AI] https://arxiv.org/abs/2507.21206

[19] Arman Zharmagambetov, Chuan Guo, Ivan Evtimov, Maya Pavlova, Ruslan Salakhutdinov, and Kamalika Chaudhuri. 2025. AgentDAM: Privacy Leakage Evaluation for Autonomous Web Agents. arXiv:2503.09780 [cs.AI] https://arxiv.org/abs/2503.09780

[20] Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. 2024. WebArena: A Realistic Web Environment for Building Autonomous Agents. arXiv:2307.13854 [cs.AI] https://arxiv.org/abs/2307.13854

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009