

# Adam Fourney

✉ adamfo@microsoft.com  
🌐 <https://www.adamfourney.com>

🔗 <https://github.com/afourney>  
📞 0000-0002-4986-7794

I build capable AI agents and study how people can collaborate with and oversee them. My work combines technical leadership in agent harnesses and multi-agent architectures with empirical research on human-AI interaction. My work is also accessible through widely adopted open-source tools with more than ¼ million GitHub stars collectively. Before the rise of LLMs, I designed, built, and studied interactive experiences for information retrieval systems, voice assistants, and AI features in productivity software.

## Experience

---

### Microsoft

*Senior Principal Researcher*

|                                                                                              |             |
|----------------------------------------------------------------------------------------------|-------------|
| Humanist Superintelligence Team, Microsoft AI                                                | Present     |
| <i>Multi-agent Team (Magentic)</i> , Microsoft Research, AI Frontiers                        | 2023 – 2026 |
| <i>Human-AI eXperiences Team (HAX)</i> , Microsoft Research, Redmond                         | 2021 – 2023 |
| <i>Information and Data Sciences Team (IDEAS)</i> , Microsoft Research, Redmond              | 2017 – 2021 |
| <i>Context Learning and User Experience for Search (CLUES)</i> , Microsoft Research, Redmond | 2015 – 2017 |

## Project Highlights

---

- 2024 **Magentic-One** — *Principal investigator*  
Led development of a category-defining general-purpose agent harness that pioneered many now-standard features, including: planning, orchestration, code execution, file system access, web browsing, and context compaction. Released as open source in 2024, roughly a year ahead of the 2025 mainstream wave of CLI agent harnesses. A core component of AutoGen and Microsoft Agent Framework.
- 2025 **Magentic-UI** — *Core developer and research collaborator*  
Core developer and co-author of an open-source research platform for studying human-agent collaboration and oversight. Magentic-UI evolves Magentic-One into a modern user-facing application supporting collaborative planning, user intervention, and action authorization.
- 2024 **MarkItDown** — *Project lead and main developer*  
Lead development of MarkItDown, another open-source project to emerge from Magentic-One. Provides agent-focused tooling that converts diverse file formats into Markdown. Ranked 2<sup>nd</sup> among Microsoft repositories and 40<sup>th</sup> globally by GitHub stars (September 2026), with approximately 500,000 package downloads per day.
- 2019 **Guidelines for Human-AI Interaction** — *Co-author*  
Co-developed the Guidelines for Human-AI Interaction through a closely collaborative research effort. With more than 4,500 research citations, the Guidelines for Human-AI Interaction are now recommended design practice in the Microsoft Responsible AI Standard, v2 (A5.3.1).

## Education

---

- 2015 Computer Science, PhD — University of Waterloo  
Thesis: *Web Search, Web Tutorials & Software Applications: Characterizing and Supporting the Coordinated Use of Online Resources for Performing Work in Feature-Rich Software*  
Award: *Bill Buxton Dissertation Award* for best Canadian HCI dissertation in 2015  
Advisor: Michael Terry
- 2009 Computer Science, Master of Mathematics — University of Waterloo  
Advisor: Richard Mann
- 2007 Computer Science, BSc (Honours) — University of Ottawa  
Award: *Computer Science Faculty Plaque* for highest cumulative GPA in an honours program

## Selected Publications

---

I have authored over 70 papers and patents. For a complete list, see: <https://www.adamfourney.com/publications/>

### Agent Systems and Societies

- 2026 *SentinelBench: A Benchmark for Long-Running Monitoring Agents*  
Matheus Kunzler Maldaner, **Adam Fourney**, Amanda Swearngin, Hussein Mozannar, Gagan Bansal, et al.  
Under review at AAAI. arXiv: 2606.05342
- 2026 *The Collaboration Gap: Exploration and Benchmarking of Open-World Agentic Cooperation*  
Tim R. Davidson, **Adam Fourney**, Saleema Amershi, Robert West, Eric Horvitz, Ece Kamar  
To appear in Proceedings of COLM 2026.
- 2025 *Magentic Marketplace: An Open-Source Environment for Studying Agentic Markets*  
Gagan Bansal, Wenyue Hua, Zezhou Huang, **Adam Fourney**, Amanda Swearngin, et al.  
In arXiv e-prints. arXiv: 2510.25779
- 2024 *Magentic-One: A Generalist Multi-Agent System for Solving Complex Tasks*  
**Adam Fourney**, Gagan Bansal, Hussein Mozannar, Cheng Tan, Eduardo Salinas, et al.  
In arXiv e-prints. arXiv: 2411.04468

### Human-Agent Interaction and Oversight

- 2026 *Overseeing Agents Without Constant Oversight: Challenges and Opportunities*  
Madeleine Grunde-McLaughlin, Hussein Mozannar, Maya Murad, Jingya Chen, Saleema Amershi, **Adam Fourney**  
Under review at TOCHI. arXiv: 2602.16844
- 2025 *Magentic-UI: Towards Human-in-the-loop Agentic Systems*  
Hussein Mozannar, Gagan Bansal, Cheng Tan, **Adam Fourney**, Victor Dibia, et al.  
In arXiv e-prints. arXiv: 2507.22358
- 2025 *Navigating Rifts in Human-LLM Grounding: Study and Benchmark*  
Omar Shaikh, Hussein Mozannar, Gagan Bansal, **Adam Fourney**, and Eric Horvitz  
In Proceedings of ACL 2025. doi: 10.18653/v1/2025.acl-long.1016
- 2025 *Interactive Debugging and Steering of Multi-Agent AI Systems*  
Will Epperson, Gagan Bansal, Victor Dibia, **Adam Fourney**, Jack Gerrits, et al.  
In Proceedings of CHI 2025.
- 2024 *Challenges in Human-Agent Communication*  
Gagan Bansal, Jennifer Wortman Vaughan, Saleema Amershi, Eric Horvitz, **Adam Fourney**, et al.  
In arXiv e-prints. arXiv: 2412.10380

### AI-Assisted Programming

- 2024 *Generation Probabilities Are Not Enough: Uncertainty Highlighting in AI Code Completions*  
Helena Vasconcelos, Gagan Bansal, **Adam Fourney**, Q. Vera Liao, and Jennifer Wortman Vaughan  
In ACM Trans. Comput.-Hum. Interact. doi: 10.1145/3702320
- 2024 *Reading Between the Lines: Modeling User Behavior and Costs in AI-Assisted Programming*  
Hussein Mozannar, Gagan Bansal, **Adam Fourney**, and Eric Horvitz  
In Proceedings of CHI 2024. doi: 10.1145/3613904.3641936
- 2024 *When to Show a Suggestion? Integrating Human Feedback in AI-Assisted Programming*  
Hussein Mozannar, Gagan Bansal, **Adam Fourney**, and Eric Horvitz  
In Proceedings of AAAI 2024. doi: 10.1609/aaai.v38i9.28878

### AI Guidelines and Recommendations

- 2021 *Planning for Natural Language Failures with the AI Playbook*  
Matthew K. Hong, **Adam Fourney**, Derek DeBellis, and Saleema Amershi  
In Proceedings of CHI 2021. doi: 10.1145/3411764.3445735
- 2019 *Guidelines for Human-AI Interaction*  
Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, **Adam Fourney**, Besmira Nushi, et al.  
In Proceedings of CHI 2019. doi: 10.1145/3290605.3300233