

Gagan Bansal

Principal Researcher, Microsoft Research AI Frontiers

Agentic AI • Human-Agent Collaboration • Multi-Agent Systems

gaganbansal1993@gmail.com • gagb.github.io • github.com/gagb • [Google Scholar](#)

I build AI agent systems and study how to keep humans in control of them. My work spans shipping widely-adopted frameworks (AutoGen, Magentic-One) to research on agent failure modes, manipulation vulnerabilities, and effective human oversight.

Education

Ph.D. and M.S., Computer Science & Engineering, University of Washington 2015 – 2021
Ph.D.: Three New Maxims for Developing Human-Centered AI (Advisor: [Daniel S. Weld](#))

B.Tech., Computer Science & Engineering, IIT Delhi 2011 – 2015
Thesis: Multi-Paragraph Multi-Document Summarization (Advisor: [Mausam](#))

Experience

Microsoft Research AI Frontiers, Redmond

Principal Researcher Nov 2025 – Present
Senior Researcher Jan 2022 – Nov 2025

Microsoft Research, Redmond – Research Intern Summer 2018, 2019
Mentors: [Besmira Nushi](#), [Ece Kamar](#), [Eric Horvitz](#)

University of Washington, Seattle – Research Intern Summer 2014
Mentor: [Daniel S. Weld](#)

Open Source Leadership (Co-Lead)

AutoGen	37k★	Multi-agent AI framework; core of Microsoft Agent Framework
MarkItDown	85k★	File-to-Markdown converter for LLM pipelines
Magentic-One	–	SOTA multi-agent system for complex web and file tasks
Magentic-UI	9k★	Human-centered web agent with co-planning and action guards
Magentic Marketplace	–	Simulation environment for AI-powered two-sided markets

Selected Blogs (Co-Lead Author)

Magentic Marketplace: An open-source simulation environment for agentic markets	Nov 2025
Magentic-UI: An experimental human-centered web agent	May 2025
AutoGen v0.4: Reimagining the foundation of agentic AI	Jan 2025
Magentic-One: A generalist multi-agent system for solving complex tasks	Nov 2024
AutoGen Studio: A low-code interface for building multi-agent workflows	Jun 2024
AutoGen: Enabling next-generation large language model applications	Sep 2023

Selected Press

Microsoft built a fake marketplace to test AI agents (TechCrunch)	Nov 2025
AutoGen v0.4: A turning point toward intelligent AI agents (VentureBeat)	Jan 2025
AutoGen: Multiple AI agents talk to each other (VentureBeat)	Oct 2023

Publications and Pre-Prints

- **Magentic Marketplace: An Open-Source Environment for Studying Agentic Markets**
Gagan Bansal, W. Hua, Z. Huang, A. Fourney, A. Swearngin, W. Epperson, T. Payne, J. Hofman, B. Lucier, C. Singh, M. Mobius, A. Nambi, A. Yadav, K. Gao, D. Rothschild, A. Slivkins, D.G. Goldstein, H. Mozannar, N. Immerlica, M. Murad, M. Vogel, S. Kambhampati, E. Horvitz, S. Amershi
MSR-TR-2025-50, October 2025
[\[PDF\]](#)
- **Magentic-UI: Towards Human-in-the-loop Agentic Systems**
H. Mozannar, Gagan Bansal, C. Tan, A. Fourney, V. Dibia, J. Chen, J. Gerrits, T. Payne, M. Kunzler Maldaner, M. Grunde-McLaughlin, E. Zhu, G. Bassman, J. Alber, P. Chang, R. Loynd, F. Niedtner, E. Kamar, M. Murad, R. Hosn, S. Amershi
MSR-TR-2025-40, July 2025
[\[PDF\]](#)
- **Interactive Debugging and Steering of Multi-Agent AI Systems**
W. Epperson, Gagan Bansal, V. Dibia, A. Fourney, J. Gerrits, E. Zhu, S. Amershi
CHI 2025
[\[PDF\]](#)
- **Generation Probabilities Are Not Enough: Uncertainty Highlighting in AI Code Completions**
H. Vasconcelos, Gagan Bansal, A. Fourney, Q.V. Liao, J.W. Vaughan
ToCHI, April 2025, Vol 32(1)
[\[PDF\]](#)
- **Navigating Rifts in Human-LLM Grounding: Study and Benchmark**
O. Shaikh, H. Mozannar, Gagan Bansal, A. Fourney, E. Horvitz
ACL 2025
[\[PDF\]](#)
- **Challenges in Human-Agent Communication**
Gagan Bansal, J.W. Vaughan, S. Amershi, E. Horvitz, A. Fourney, H. Mozannar, V. Dibia, D.S. Weld
ArXiv 2025
[\[PDF\]](#)
- **Magentic-One: A Generalist Multi-Agent System for Solving Complex Tasks**
A. Fourney, Gagan Bansal, H. Mozannar, C. Tan, E. Salinas, F. Niedtner, G. Proebsting, G. Bassman, J. Gerrits, J. Alber, P. Chang, R. Loynd, R. West, V. Dibia, A. Awadallah, E. Kamar, R. Hosn, S. Amershi
MSR-TR-2024-47, November 2024
[\[PDF\]](#)
- **AutoGen Studio: A No-Code Developer Tool for Building and Debugging Multi-Agent Systems**
V. Dibia, J. Chen, Gagan Bansal, S. Syed, A. Fourney, E. Zhu, C. Wang, S. Amershi
EMNLP 2024, pp. 72-79
[\[PDF\]](#)
- **Measuring AI Agent Autonomy: Towards a Scalable Approach With Code Inspection**

P. Cihon, M. Stein, **Gagan Bansal**, S. Manning, K. Xu
Workshop on Socially Responsible Language Modelling Research, 2024

- **AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation Framework** ★
Q. Wu, **Gagan Bansal**, J. Zhang, Y. Wu, S. Zhang, E. Zhu, B. Li, L. Jiang, X. Zhang, A. Awadallah, R.W. White, D. Burger, C. Wang
COLM 2024 • Best Paper, LLM Agents Workshop ICLR'24
[\[PDF\]](#)
- **Reading Between the Lines: Modeling User Behavior and Costs in AI-Assisted Programming** ★
H. Mozannar, **Gagan Bansal**, A. Fourney, E. Horvitz
CHI 2024 • Honorable Mention
[\[PDF\]](#)
- **When to Show a Suggestion? Integrating Human Feedback in AI-Assisted Programming**
H. Mozannar, **Gagan Bansal**, A. Fourney, E. Horvitz
AAAI 2024
[\[PDF\]](#)
- **Understanding the Role of Human Intuition on Reliance in Human-AI Decision-Making with Explanations**
V. Chen, Q.V. Liao, J.W. Vaughan, **Gagan Bansal**
CSCW 2023
[\[PDF\]](#)
- **Aligning Offline Metrics and Human Judgments of Value of AI-Pair Programmers**
V. Dibia, A. Fourney, **Gagan Bansal**, F. Poursabzi-Sangdeh, H. Liu, S. Amershi
ACL 2023
[\[PDF\]](#)
- **Embedded Attributes for Modifying Behaviors of Generative AI Systems**
S.A. Amershi, A. Fourney, V.C. Dibia, **Gagan Bansal**
US Patent App. 18/197,878, 2024
- **Three Maxims for Developing Human-Centered AI for Decision Making**
Gagan Bansal
Ph.D. Thesis, University of Washington, 2021
- **Human Evaluation of Spoken vs. Visual Explanations for Open-Domain QA**
A.V. Gonzalez, **Gagan Bansal**, A. Fan, Y. Mehdad, R. Jia, S. Iyer
ACL 2021
- **Is the Most Accurate AI the Best Teammate? Optimizing AI for Teamwork**
Gagan Bansal, T. Wu, B. Nushi, E. Kamar, E. Horvitz, D.S. Weld
AAAI 2021
[\[PDF\]](#)
- **Does the Whole Exceed its Part? The Effect of Explanations on Complementary Team Performance**
Gagan Bansal, T. Wu, J. Zhou, R. Fok, B. Nushi, E. Kamar, M.T. Ribeiro, D.S. Weld
CHI 2021
[\[PDF\]](#)
- **Beyond Accuracy: On the Role of Mental Models in Human-AI Teams**
Gagan Bansal, B. Nushi, E. Kamar, W. Lasecki, D.S. Weld, E. Horvitz
HCOMP 2019
[\[PDF\]](#)

- **Updates in Human-AI Teams: Understanding and Addressing the Performance/Compatibility Tradeoff**
Gagan Bansal, B. Nushi, E. Kamar, W. Lasecki, D.S. Weld, E. Horvitz
AAAI 2019
[\[PDF\]](#)
- **The Challenge of Crafting Intelligible Intelligence**
D.S. Weld, Gagan Bansal
CACM 2018
[\[PDF\]](#)
- **A Coverage-Based Utility Model for Identifying Unknown Unknowns**
Gagan Bansal, D.S. Weld
AAAI 2018
- **Hierarchical Summarization: Scaling Up Multi-Document Summarization**
J. Christensen, S. Soderland, Gagan Bansal, Mausam
ACL 2014

Awards and Achievements

- Best Paper, LLM Agents Workshop ICLR'24 (AutoGen)
- Honorable Mention, CHI 2024 (Reading Between the Lines)
- Department rank 5 in graduating class of Computer Science and Engineering, IIT Delhi 2015
- Gold medal at Indian National Chemistry Olympiad 2011

Talks, Panels & Media

- Guest instructor for agents capstone project, University of Washington (2025)
- Invited talk on human-centered AI, UW RAISE Seminar (2025)
- Panel on Human-Centered Agents for Healthcare, World Salon (2025)
- Mind and Machine Intelligence Summit, University of California, Santa Barbara (2023)
- Human-AI Interaction course at KAIST University, Seoul, Korea (2021)
- Invited Speaker at BitHacks Speaker Series, USA (2021)

Academic Service

- **Thesis committees:** Will Epperson (CMU), Tae Soo Kim (KAIST)
- **Workshop organization:** CHI 23 Workshop on Trust and Reliance in AI-Human Teams, Hamburg, Germany; CHI 22 Workshop on Trust and Reliance in AI-Human Teams, New Orleans, USA; CHI 19 Workshop on Human-Centered Machine Learning, Glasgow, UK

Mentoring (Interns)

- Wayne Chi (Ph.D. Student at CMU, Summer 2025)
- Omar Shaikh (Ph.D. Student at Stanford, Summer 2024)
- Will Epperson (Ph.D. Student at CMU, Summer 2024; now Researcher at Microsoft)
- Shannon Shen (Ph.D. Student at MIT, Summer 2023)
- Hussein Mozannar (Ph.D. Student at MIT, Summer 2022; now Researcher at Microsoft)