Justin D. Weisz, Ph.D.

Human-Centered AI Research Leader

hello@jweisz.org

SUMMARY

Research leader in human-centered AI focused on designing **useful**, **safe**, and **personal** AI systems that **amplify and augment human capabilities**. ACM Senior Member with numerous top-tier HCI & AI publications in AAAI, NeurIPS, CHI, IUI, AIES, and more. Research topics span human-AI collaboration & co-creativity; agentic AI; responsible & human-centered explainable AI (HCXAI); AI risk, governance, and evaluation. Versed in qualitative, quantitative, prototyping, crowdsourcing, and speculative design methods to understand how to build effective AI user experiences. Extensive collaboration experience with AI & ML researchers, engineers, product managers, and designers to prototype the future and **deliver research insights into shipped products**.

↪ *Selected Product Contributions & Systems Built*: watsonx AutoAI, watsonx AI Risk Atlas, watsonx Code Assistant, Programmer's Assistant, Agentic Code Explorer, TJBot, Idea Gauntlet

EDUCATION

Carnegie Mellon University Pittsburgh, PA

- **Ph.D. in Computer Science**, Computer Science Department 2009
  Committee: Sara Kiesler (co-chair), Hui Zhang (co-chair), Luis von Ahn, Wendy Kellogg (IBM)
  Dissertation: *Collaborative Online Video Watching*
- **M.S. in Computer Science**, Computer Science Department 2007
- **B.S. in Computer Science**, School of Computer Science 2003
  Graduated with University and SCS College Honors
  Senior Honors Thesis: *Detecting Cheaters in a Distributed Multiplayer Game*

SKILLS

AI/ML • Agentic AI • LLM Frameworks • LLM memory • Post-training • LLM-as-a-Judge • Evals • RAG • Context Engineering • **Human-Computer Interaction** • Experimental Design • Statistical Analysis • Qualitative Methods • Survey Design • Prototyping • Crowdsourcing • Design Fiction • Interaction Design • User Experience (UX) Evaluation • **Development** • Python • R • TypeScript • Agile • Open Source • AI-Assisted Coding • **Leadership** • Research Program Design • Research Portfolio Management • Research Product Impact • Cross-Functional Technical Leadership • Mentorship & Coaching

PROFESSIONAL EXPERIENCE

AI Research & Product Advisor 2026 –

Brilliance Foundation

- Incubating new human-centered AI research lab focused on developing AI systems that sharpen people's critical thinking skills and help them avoid cognitive surrender and overreliance.

Snorkle Labs

- Advising on business, IP, tech stack, and funding strategy for health care data visualization startup. Creating new AI product to educate patients on medical conditions and treatment options.

SplitStay

- Advising on research, product, and IP strategy for group travel coordination startup. Focus on using AI for preference elicitation & negotiation.

IBM T.J. Watson Research Center Yorktown Heights, NY

Senior Research Scientist, Manager, and Strategy Lead, Human-Centered Trustworthy AI 2022 – 2025

- Co-led global research strategy on human-centered trustworthy AI, overseeing 30+ research scientists, engineers, and designers with backgrounds spanning AI, machine learning, and human-computer interaction. Research topics spanned human-AI collaboration & co-creation; personal AI & mutual theory of mind; AI risk identification & mitigation; contextual privacy in LLMs; agentic AI development & multi-agent collaboration under differing value alignments; factuality assessment; disclosure of AI use in co-creative processes.
- HCAI research portfolio recognized by IBM as having an Outstanding scientific impact from 2017-2025: 50+ papers garnering 3,500+ citations, 10 papers with > 100 citations, 14 organized workshops, and multiple open source projects (e.g. QUX360, EvalAssist, Contextual Privacy Toolkit, GANDissect, GLTR, LMDiff, Transformer Explainer, Diffusion Explainer).
- Led cross-organization and cross-disciplinary teams to develop & evaluate prototype systems for human-AI collaboration & co-creativity in software engineering and design.
- Worked hands-on with product teams to translate design prototypes and research findings into new product features and measure productivity impact of AI.
- Educated 3,000+ design practitioners on designing effective and safe generative AI user experiences by developing training materials, delivering seminars, and crafting new design thinking methods.
- Added new risks from generative AI to the IBM AI Risk Atlas.
- Directly managed team of 10+ researchers, software engineers, and designers. Provided career development, mentoring & coaching, and performance evaluation.

Manager & Strategy Lead, Human-AI Collaboration 2019 – 2022

- Led research strategy on human-AI collaboration focused on trustworthy AI systems (fairness, explainability), human-centered automation of data science, and human-AI co-creation with generative models.

Research Staff Member, HCI of AI 2018 – 2019

- Developed research agenda to train users to have effective interactions with conversational AI

Senior Manager, Emerging Technology Research 2016 – 2018

- Led team of researchers and software engineers to develop novel methods for interacting with and programming emerging technologies, including quantum computers, drones, and other IoT devices.

Senior Manager, Mobile Solutions & Infrastructure 2015 – 2016

- Led cross-disciplinary team to produce novel mobile & wearable solutions in healthcare and UAV solutions in agriculture.

Manager, Mobile Solutions & Infrastructure 2013 – 2015

- Led cross-disciplinary team to produce novel mobile & wearable solutions for field service technicians, middleware for migrating data to the cloud, and transformative mobile applications for multiple industries.

Research Staff Member (RSM), Next Generation Technologies 2010 – 2013

- Developed analytics pipelines for identifying emotional content in email and instant message data, designed visualizations for expertise-building recommender system, and developed location-based services for WebSphere Commerce.

IBM T.J. Watson Research Hawthorne, NY

Intern, Social Computing Group (2005), Enhanced Web Experience Group (2004)

- Across two internships, applied qualitative, quantitative, and visualization methods to examine use of IBM Community Tools (ICT), an instant messaging platform supporting topical group chats.

Carnegie Mellon University Pittsburgh, PA

Graduate Research Assistant 2003 – 2009

- Examined the experience of conversing while watching online videos via controlled experiments and field studies and identified social benefits of group-based online video watching for both friends and strangers, despite distractive aspects of chat.
- Designed new visual representation of large-scale online audiences.
- Built applications for viewing and broadcasting video on End System Multicast, a novel peer-to-peer overlay network.

PUBLICATIONS

ACM: **DBLP:** **GOOGLE SCHOLAR:** **ORCID:**

Refereed Conference Papers

P.36 Anderson, A., Piorkowski, D., **Weisz, J. D.**, Burnett, M., & Varshney, K. R. (2026, October). Can LLM Code Explanations Adapt to Diverse Problem-Solvers' Needs?. In Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (Vol. 9, No. 1, pp. 192-205).

P.35 Morrison, K., Gonzalez, G. E., Ashktorab, Z., Riemer, M., Anderson, A., Bouneffouf, D., & **Weisz, J. D.** (2026, August). Wavelength.AI: Extending the collaborative game Wavelength as a testbed for studying shared understanding in human–agent collaboration. In D. Calvanese (Ed.), Proceedings of the Thirty-Fifth International Joint Conference on Artificial Intelligence, IJCAI-26 (pp. 8466–8470). International Joint Conferences on Artificial Intelligence Organization.

P.34 He, J., Akella, A., Do, H. J., & **Weisz, J. D.** (2026, June). Overcoming Barriers to AI Disclosure: The Role of Granularity and Sociocultural Enablers in Reducing Perceived Stigmas. In Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT '26). Association for Computing Machinery, New York, NY, USA, 2117–2143.

P.33 Do, H. J., Ostrand, R., Geyer, W., Murugesan, K., Wei, D., & **Weisz, J.** (2025, October). Highlight All the Phrases: Enhancing LLM Transparency through Visual Factuality Indicators. In Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (Vol. 8, No. 1, pp. 772-784).

P.32 Ngong, I. C., Kadhe, S. R., Wang, H., Murugesan, K., **Weisz, J. D.**, Dhurandhar, A., & Ramamurthy, K. N. (2025, July). Protecting users from themselves: Safeguarding contextual privacy in interactions with conversational agents. In Findings of the Association for Computational Linguistics: ACL 2025 (pp. 26196-26220).

P.31 Muller, M., He, J., & **Weisz, J. D.** (2025, June). Exploration and Optimization of Generative Variability in Future Work: A Mixed-Initiative Analysis. In Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (pp. 1-13).

P.30 Riemer, M., Ashktorab, Z., Bouneffouf, D., Das, P., Liu, M., **Weisz, J. D.**, Campbell, M. (2025, May). Position: Theory of Mind Benchmarks are Broken for Large Language Models. In Proceedings of the 42nd International Conference on Machine Learning (ICML 2025) Position Paper Track.

P.29 He, J., Houde, S., & **Weisz, J. D.** (2025, April). Which contributions deserve credit? perceptions of attribution in human-ai co-creation. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (pp. 1-18).

P.28 Houde, S., Brimijoin, K., Muller, M., Ross, S. I., Silva Moran, D. A., Gonzalez, G. E., Kunde, S., Foreman, M. A., & **Weisz, J. D.** (2025, March). Controlling AI Agent Participation in Group Conversations: A Human-Centered Approach. In Proceedings of the 30th International Conference on Intelligent User Interfaces (pp. 390-408).

P.27 Muller, M., **Weisz, J. D.**, Houde, S., & Ross, S. I. (2024, June). Drinking Chai with Your (AI) Programming Partner: Value Tensions in the Tokenization of Future Human-AI Collaborative Work. In Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work (pp. 1-15).

P.26 He, J., Houde, S., Gonzalez, G. E., Silva Moran, D. A., Ross, S. I., Muller, M., & **Weisz, J. D.** (2024, June). AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas. In Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work (pp. 1-14). ***Best Paper Honorable Mention award***

P.25 Muller, M., Houde, S., Gonzalez, G., Brimijoin, K., Ross, S. I., Moran, D. A. S., & **Weisz, J. D.** (2024, June). Group Brainstorming with an AI Agent: Creating and Selecting Ideas. In International Conference on Computational Creativity.

P.24 **Weisz, J. D.**, He, J., Muller, M., Hoefer, G., Miles, R., & Geyer, W. (2024, May). Design principles for generative AI applications. In Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (pp. 1-22). ***37k+ downloads on ACM DL (rank #9 for CHI papers)***

P.23 He, J., Piorkowski, D., Muller, M., Brimijoin, K., Houde, S., & **Weisz, J.** (2023, June). Rebalancing worker initiative and ai initiative in future work: Four task dimensions. In Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work (pp. 1-16).

P.22 Muller, M., Candello, H., & **Weisz, J.** (2023, June). Interactional co-creativity of human and ai in analogy-based design. In International Conference on Computational Creativity.

P.21 Ross, S. I., Martinez, F., Houde, S., Muller, M., & **Weisz, J. D.** (2023, March). The programmer's assistant: Conversational interaction with a large language model for software development. In Proceedings of the 28th International Conference on Intelligent User Interfaces (pp. 491-514).

P.20 Muller, M., and **Weisz, J**. 2022. Extending a Human-AI Collaboration Framework with Dynamism and Sociality. In 2022 Symposium on Human-Computer Interaction for Work (pp. 1-12).

P.19 Sun, J., Liao, Q. V., Muller, M., Agarwal, M., Houde, S., Talamadupula, K., and **Weisz, J. D.** 2022. Investigating Explainability of Generative AI for Code through Scenario-based Design. In 27th International Conference on Intelligent User Interfaces (pp. 212-228).

P.18 **Weisz, J. D.**, Muller, M., Ross, S. I., Martinez, F., Houde, S., Agarwal, M., Talamadupula, K., and Richards, J. T. 2022. Better together? an evaluation of ai-supported code translation. In 27th International Conference on Intelligent User Interfaces (pp. 369-391). ***IBM Pat Goldberg 2022 Best Paper Honorable Mention award***

P.17 Wang, D., Andres, J., **Weisz, J. D.**, Oduor, E., and Dugan, C. 2021. Autods: Towards human-centered automation of data science. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (pp. 1-12).

P.16 Ehsan, U., Liao, Q. V., Muller, M., Riedl, M. O., and **Weisz, J. D.** 2021. Expanding explainability: Towards social transparency in ai systems. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (pp. 1-19). ***Best Paper Honorable Mention award***

P.15 Narkar, S., Zhang, Y., Liao, Q. V., Wang, D., and **Weisz, J. D.** 2021. Model LineUpper: Supporting Interactive Model Comparison at Multiple Levels for AutoML. In 26th International Conference on Intelligent User Interfaces (pp. 170-174).

P.14 **Weisz, J. D.**, Muller, M., Houde, S., Richards, J., Ross, S. I., Martinez, F., Agarwal, M., and Talamadupula, K. 2021. Perfection not required? Human-AI partnerships in code translation. In 26th International Conference on Intelligent User Interfaces (pp. 402-412).

P.13 Drozdal, J., **Weisz, J.**, Wang, D., Dass, G., Yao, B., Zhao, C., Muller, M., Ju, L., and Su, H. 2020. Trust in AutoML: exploring information needs for establishing trust in automated machine learning systems. In Proceedings of the 25th International Conference on Intelligent User Interfaces (pp. 297-307).

P.12 Weidele, D. K. I., **Weisz, J. D.**, Oduor, E., Muller, M., Andres, J., Gray, A., & Wang, D. 2020. AutoAIViz: opening the blackbox of automated artificial intelligence with conditional parallel coordinates. In Proceedings of the 25th International Conference on Intelligent User Interfaces (pp. 308-312).

P.11 Wang, D., **Weisz, J. D.**, Muller, M., Ram, P., Geyer, W., Dugan, C., Tausczik, Y., Samulowitz, H., and Gray, A. 2019. Human-AI collaboration in data science: Exploring data scientists' perceptions of automated AI. Proceedings of the ACM on Human-Computer Interaction, 3(CSCW), 1-24.

P.10 Ashktorab, Z., Jain, M., Liao, Q. V., and **Weisz, J. D.** 2019. Resilient chatbots: Repair strategy preferences for conversational breakdowns. In Proceedings of the 2019 CHI conference on human factors in computing systems (pp. 1-12).

P.9 Ashktorab, Z., **Weisz, J. D.**, and Ashoori, M. 2019. Thinking Too Classically: Research Topics in Human-Quantum Computer Interaction. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (pp. 1-12).

P.8 **Weisz, J. D.**, Jain, M., Joshi, N. N., Johnson, J., and Lange, I. 2019. BigBlueBot: teaching strategies for successful human-agent interactions. In Proceedings of the 24th International Conference on Intelligent User Interfaces (pp. 448-459).

P.7 **Weisz, J. D.**, Ashoori, M., and Ashktorab, Z. 2018. Entanglion: A board game for teaching the principles of quantum computing. In Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play (pp. 523-534).

P.6 Gowda, M., Manweiler, J., Dhekne, A., Choudhury, R. R., and **Weisz, J. D.** 2017. Integrating GLONASS with GPS for drone orientation tracking. In International Conference on Communication Systems and Networks (pp. 77-92). Springer, Cham.

P.5 Gowda, M., Manweiler, J., Dhekne, A., Choudhury, R. R., and **Weisz, J. D.** 2016. Tracking drone orientation with multiple GPS receivers. In Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking (MobiCom '16). ACM, New York, NY, USA, 280-293.

P.4 Lai, J., Lu, J., Pan, S., Soroker, D., Topkara, M., **Weisz, J.**, Boston, J., and Crawford, J. 2014. Expediting expertise: supporting informal social learning in the enterprise. In Proceedings of IUI 2014. ACM, New York, NY, 133-142.

P.3 **Weisz, J. D.** and Kiesler, S. 2008. How text and audio chat change the online video experience. In Proceedings of uxTV'08. ACM, New York, NY, 9-18.

P.2 **Weisz, J. D.**, Kiesler, S., Zhang, H., Ren, Y., Kraut, R. E., and Konstan, J. A. 2007. Watching together: integrating text chat with video. In Proceedings of SIGCHI 2007. ACM, New York, NY, 877-886.

P.1 **Weisz, J. D.**, Erickson, T., and Kellogg, W. A. 2006. Synchronous broadcast messaging: the use of ICT. In Proceedings of SIGCHI 2006. ACM, New York, NY, 1293-1302. ***Best Paper award***

Refereed Workshop Papers & Extended Abstracts

EA.30 Morrison, K., Gonzalez, G. E., Ashktorab, Z., Riemer, M., Anderson, A., & **Weisz, J. D.** (2026, April). "On the Same Wavelength": Exploring Human–Agent Collaboration Strategies through the Cooperative Game Wavelength. In the CHI 2026 Workshop on Human-Agent Collaboration, Barcelona, Spain, April 2026.

EA.29 **Weisz, J. D.**, Muller, M., & Varshney, K. (2025, August). Story Arena: A Multi-Agent Environment for Envisioning the Future of Software Engineering. In the 2025 Workshop on The End of Programming

(as we know it): Envisioning Radical Re-Conceptualizations of Co-Coding with AI, held in conjunction with the Aarhus 2025 Decennial Conference, August 18-22, 2025.

EA.28 Brachman, M., Goldberg, A., Anderson, A., Houde, S., Muller, M., & **Weisz, J. D.** (2025, June). Towards Personalized and Contextualized Code Explanations. In Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (pp. 120-125).

EA.27 **Weisz, J. D.**, Kumar, S. V., Muller, M., Browne, K. E., Goldberg, A., Heintze, K. E., & Bajpai, S. (2025, April). Examining the use and impact of an ai code assistant on developer productivity and experience in the enterprise. In Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (pp. 1-13).

EA.26 Moran, D. A. S., Brimijoin, K., Gonzalez, G. E., Houde, S., Muller, M., Brachman, M., & **Weisz, J. D.** (2025, March). ACE: Moving toward Co-Investigation with the Agentic Code Explorer. In Joint Proceedings of the ACM IUI 2025 Workshops, co-located with the 30th Annual ACM Conference on Intelligent User Interfaces (IUI 2025), Cagliari, Italy, March 24th, 2025, CEUR-WS.org/Vol-3957.

EA.25 Morrison, K., Ashktorab, Z., Bouneffouf, D., Gonzalez, G. E., & **Weisz, J. D.** (2025, March). Establishing the Cooperative Game Wavelength as a Testbed to Explore Mutual Theory of Mind. In Proceedings of the International Workshop on Advancing AI Through Theory of Mind (ToM4AI) held in conjunction with AAAI 2025, Philadelphia, PA, March 3, 2025.

EA.24 Varshney, K. R., Ashktorab, Z., Bouneffouf, D., Riemer, M., & **Weisz, J. D.** (2025, February). Scopes of Alignment. In Proceedings of the 2nd International Workshop on AI Governance (AIGOV) held in conjunction with AAAI 2025, Philadelphia, PA, USA, March 3, 2025.

EA.23 Do, H. J., Ostrand, R., **Weisz, J. D.**, Dugan, C., Sattigeri, P., Wei, D., Murugesan, K., & Geyer, W. (2024, May). Facilitating human-LLM collaboration through factuality scores and source attributions. In Workshop on Trust and Reliance in Evolving Human-AI Workflows at CHI 2024 (TREW).

EA.22 **Weisz, J. D.**, Muller, M., Goldberg, A., & Moran, D. A. S. (2024, May). Expedient Assistance and Consequential Misunderstanding: Envisioning an Operationalized Mutual Theory of Mind. In Workshop on Theory of Mind in Human-AI Interaction at CHI 2024 (ToMinHAI).

EA.21 Gonzalez, G. E., Moran, D. A. S., Houde, S., He, J., Ross, S. I., Muller, M. J., Kunde, S., & **Weisz, J. D.** (2024, March). Collaborative Canvas: A Tool for Exploring LLM Use in Group Ideation Tasks. In Joint Proceedings of the ACM IUI 2024 Workshops, co-located with the 29th Annual ACM Conference on Intelligent User Interfaces (IUI 2024), Greenville, South Carolina, USA, March 18, 2024, CEUR-WS.org/Vol-3660.

EA.20 He, J., Piorkowski, D., Muller, M. J., Brimijoin, K., Houde, S., & **Weisz, J. D.** (2023, April). Understanding How Task Dimensions Impact Automation Preferences with a Conversational Task Assistant. In Proceedings of the Workshop on Intervening, Teaming, Delegating (AutomationXP), co-located with the ACM Conference on Human Factors in Computing Systems (CHI 2023), Hamburg, Germany, April 23, 2023, CEUR-WS.org/Vol-3394.

EA.19 Ross, S. I., Houde, S., Martinez, F., Muller, M., & **Weisz, J. D.** (2023, March). The Programmer's Assistant User Experience. In Companion Proceedings of the 28th International Conference on Intelligent User Interfaces (pp. 102-104). 🏆 ***Best Demo Honorable Mention award***

EA.18 Ross, S. I., Muller, M., Martinez, F., Houde, S., & **Weisz, J. D.** (2023). A Case Study in Engineering a Conversational Programming Assistant's Persona. In Joint Proceedings of the IUI 2023 Workshops: HAI-GEN, ITAH, MILC, SHAI, SketchRec, SOCIALIZE, co-located with the ACM International Conference on Intelligent User Interfaces (IUI 2023), Sydney, Australia, March 27-31, 2023, CEUR-WS.org/Vol-3359.

EA.17 **Weisz, J. D.**, Muller, M., He, J., & Houde, S. (2023). Toward general design principles for generative AI applications. In Joint Proceedings of the IUI 2023 Workshops: HAI-GEN, ITAH, MILC,

SHAI, SketchRec, SOCIALIZE, co-located with the ACM International Conference on Intelligent User Interfaces (IUI 2023), Sydney, Australia, March 27-31, 2023, CEUR-WS.org/Vol-3359.

EA.16 Houde, S., Ross, S., Muller, M., Agarwal, M., Martinez, F., Richards, J., Talamadupula, K., and **Weisz, J. D.** 2022. Opportunities for Generative AI in UX Modernization. In Joint Proceedings of the IUI 2022 Workshops: APEx-UI, HAI-GEN, HEALTHI, HUMANIZE, TExSS, SOCIALIZE co-located with the ACM International Conference on Intelligent User Interfaces (IUI 2022), Virtual Event, Helsinki, Finland, March 21-22, 2022, CEUR-WS.org/Vol-3124.

EA.15 Muller, M., Ross, S., Houde, S., Agarwal, M., Martinez, F., Richards, J., Talamadupula, K., and **Weisz, J. D.** 2022. Drinking Chai with Your (AI) Programming Partner: A Design Fiction about Generative AI for Software Engineering. In Joint Proceedings of the IUI 2022 Workshops: APEx-UI, HAI-GEN, HEALTHI, HUMANIZE, TExSS, SOCIALIZE co-located with the ACM International Conference on Intelligent User Interfaces (IUI 2022), Virtual Event, Helsinki, Finland, March 21-22, 2022, CEUR-WS.org/Vol-3124.

EA.14 Muller, M., Wang, A. Y., Ross, S., **Weisz, J. D.**, Agarwal, M., Talamadupula, K., Houde, S., Martinez, F., Richards, J., Drozdal, J., Liu, X., Piorkowski, D. and Wang, D. 2021. How Data Scientists Improve Generated Code Documentation in Jupyter Notebooks. In Joint Proceedings of the ACM IUI 2021 Workshops co-located with 26th ACM Conference on Intelligent User Interfaces (ACM IUI 2021), College Station, United States, April 13-17, 2021, CEUR-WS.org/Vol-2903.

EA.13 Muller, M., **Weisz, J. D.**, & Geyer, W. (2020). Mixed Initiative Generative AI Interfaces: An Analytic Framework for Generative AI Applications. In Proceedings of the Workshop The Future of Co-Creative Systems-A Workshop on Human-Computer Co-Creativity of the 11th International Conference on Computational Creativity (ICCC 2020).

EA.12 Houde, S., Liao, V., Martino, J., Muller, M., Piorkowski, D., Richards, J., **Weisz, J.**, and Zhang, Y. 2020. Business (mis)Use Cases of Generative AI. In Joint Proceedings of the Workshops on Human-AI Co-Creation with Generative Models and User-Aware Conversational Agents co-located with 25th International Conference on Intelligent User Interfaces (IUI 2020), Cagliari, Italy, March 17, 2020, CEUR-WS.org/Vol-2848.

EA.11 Wang, D., Ram, P., Weidele, D. K. I., Liu, S., Muller, M., **Weisz, J. D.**, Valent A. N., Chaudhard, A., Torres, D., Samulowitz, H., and Amini, L. 2020. Autoai: Automating the end-to-end ai lifecycle with humans-in-the-loop. In Proceedings of the 25th International Conference on Intelligent User Interfaces Companion (pp. 77-78).

EA.10 Andres, J., Wolf, C. T., Muller, M., **Weisz, J. D.**, Joshi, N. N., Sharma, A., Brimijoin, K., Desmond, M., Ashktorab, Z., Pan, Q., Duesterwald, E., and Dugan, C. 2020. Cultivating Human Expertise Through AI-Assisted Data Science. 2020. In Proceedings of the 1st Workshop on Data Science with Human in the Loop (DASH) co-located with KDD 2020.

EA.9 Wang, D., Andres, J., **Weisz, J. D.**, Oduor, E., Khurana, U., Samulowitz, H., Chaudhary, A., Valente, A. N., Torres, D., and Dugan, C. 2020. IBM AutoAI: Human-in-the-Loop Automated Machine Learning Supports Data Scientists to Build Better Models. 2020. In Proceedings of the 1st Workshop on Data Science with Human in the Loop (DASH) co-located with KDD 2020.

EA.8 **Weisz, J. D.**, and Muller, M. The Next Decade of Data Science. 2020. In Proceedings of the 1st Workshop on Data Science with Human in the Loop (DASH) co-located with KDD 2020.

EA.7 Agarwal, M., Talamadupula, K., Houde, S., Martinez, F., Muller, M., Richards, J., Ross, S., and **Weisz, J. D.** 2020. Quality Estimation & Interpretability for Code Translation. In Proceedings of the NeurIPS 2020 Workshop on Computer-Assisted Programming (CAP@NeurIPS 2020).

EA.6 Uchidiuno, J. O., Manweiler, J., and **Weisz, J. D.** 2018. Privacy and fear in the drone era: Preserving privacy expectations through technology. In Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (pp. 1-6).

EA.5 Dibia, V. C., Ashoori, M., Cox, A., and **Weisz, J. D.** 2017. TJBot: An Open Source DIY Cardboard Robot for Programming Cognitive Systems. In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17). ACM, New York, NY, USA, 381-384.

EA.4 Danis, C., Vogt, K., **Weisz, J. D.**, Ma, Y., and Olsen, R. 2017. Factors that Help and Hinder a Daily Weighing and Reporting Behavior. In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17). ACM, New York, NY, USA, 2480-2486.

EA.3 Ashoori, M., Bellamy, R. K. E., and **Weisz, J. D.** 2016. The Impending Ubiquity of Cognitive Objects. In In Workshops at the Thirtieth AAAI Conference on Artificial Intelligence.

EA.2 Ashoori, M., Bellamy, R. K. E., and **Weisz, J. D.** 2015. Creating the Mood: Design for a Cognitive Meeting Room. In Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '15). ACM, New York, NY, USA, 2001-2006.

EA.1 Topkara, M., **Weisz, J. D.**, Pan, S., Lu, J., and Lai, J. 2014. Dare to compare: motivating expertise building in the enterprise through intelligent user modeling interfaces. In 3rd Workshop on Data-driven User Behavioral Modeling and Mining from Social Media (CIKM DUBMOD 2014).

Pre-prints

PP.9 Ngong, I. C., Murugesan, K., Kadhe, S., **Weisz, J. D.**, Dhurandhar, A., & Ramamurthy, K. N. (2026). AgentSCOPE: Evaluating Contextual Privacy Across Agentic Workflows. arXiv preprint arXiv:2603.04902.

PP.8 Anderson, A., Piorkowski, D., Burnett, M., & **Weisz, J.** (2025). An LLM's Attempts to Adapt to Diverse Software Engineers' Problem-Solving Styles: More Inclusive & Equitable?. arXiv preprint arXiv:2503.11018.

PP.7 Ngong, I., Kadhe, S., Wang, H., Murugesan, K., **Weisz, J. D.**, Dhurandhar, A., & Ramamurthy, K. N. (2025). Protecting users from themselves: Safeguarding contextual privacy in interactions with conversational agents. arXiv preprint arXiv:2502.18509.

PP.6 Riemer, M., Ashktorab, Z., Bouneffouf, D., Das, P., Liu, M., **Weisz, J. D.**, & Campbell, M. (2024). Position: Theory of Mind Benchmarks are Broken for Large Language Models. arXiv preprint arXiv:2412.19726.

PP.5 Agarwal, M., Talamadupula, K., Martinez, F., Houde, S., Muller, M., Richards, J., Ross, S., and **Weisz, J. D.** 2021. Using document similarity methods to create parallel datasets for code translation. arXiv preprint arXiv:2110.05423.

PP.4 Muthusamy, V., Unuvar, M., Völzer, H., & **Weisz, J. D.** 2020. Do's and Don'ts for Human and Digital Worker Integration. arXiv preprint arXiv:2010.07738.

PP.3 Ashoori, M., and **Weisz, J. D.** 2019. In AI we trust? Factors that influence trustworthiness of AI-infused decision-making processes. arXiv preprint arXiv:1912.02675.

PP.2 Ismail, S., Manweiler, J. G., and **Weisz, J. D.** 2018. CardKit: A Card-Based Programming Framework for Drones. arXiv preprint arXiv:1804.08458.

PP.1 Dibia, V., Cox, A., and **Weisz, J.** 2018. Designing for Democratization: Introducing Novices to Artificial Intelligence Via Maker Kits. arXiv preprint arXiv:1805.10723.

Organized Workshops

W.10 Dinakar, K., **Weisz, J. D.**, Lieberman, H., & Geyer, W. (2026, March). AgentCraft: Workshop on Developing Trustworthy Agentic AI Systems. In Companion Proceedings of the 31st International Conference on Intelligent User Interfaces (pp. 240-242).

W.9 Wang, Q., Wester, J., Pafla, M., Lee, M., **Weisz, J. D.**, & Si, M. (2025, July). ToMinHAI at CUI'2025: Theory of Mind in Human-CUI Interaction. In Proceedings of the 7th ACM Conference on Conversational User Interfaces (pp. 1-6).

W.8 Ehsan, U., Watkins, E. A., Wintersberger, P., Manger, C., Hubig, N., Savage, S., **Weisz, J. D.**, & Riener, A. (2025, April). New Frontiers of Human-centered Explainable AI (HCXAI): Participatory Civic AI, Benchmarking LLMs, XAI Hallucinations, and Responsible AI Audits. In Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (pp. 1-6).

W.7 Mokryn, O., Shaer, O., Geyer, W., Maher, M. L., **Weisz, J. D.**, Buschek, D., & Chilton, L. B. (2025, March). HAI-GEN 2025: 6th Workshop on Human-AI Co-Creation with Generative Models. In Companion Proceedings of the 30th International Conference on Intelligent User Interfaces (pp. 179-182).

W.6 Wang, Q., Walsh, S., Si, M., Kephart, J., **Weisz, J. D.**, & Goel, A. K. (2024, May). Theory of mind in human-ai interaction. In Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (pp. 1-6).

W.5 Geyer, W., Maher, M. L., **Weisz, J. D.**, Buschek, D., & Chilton, L. B. (2024, March). HAI-GEN 2024: 5th Workshop on Human-AI Co-Creation with Generative Models. In Companion Proceedings of the 29th International Conference on Intelligent User Interfaces (pp. 122-124).

W.4 Ehsan, U., Wintersberger, P., Watkins, E. A., Manger, C., Ramos, G., **Weisz, J. D.**, Daumé III, H., Riene A., & Riedl, M. O. (2023, April). Human-centered explainable AI (HCXAI): coming of age. In Extended abstracts of the 2023 CHI conference on human factors in computing systems (pp. 1-7).

W.3 Maher, M. L., **Weisz, J. D.**, Chilton, L. B., Geyer, W., & Strobelt, H. (2023, March). HAI-GEN 2023: 4th Workshop on Human-AI Co-Creation with Generative Models. In Companion Proceedings of the 28th International Conference on Intelligent User Interfaces (pp. 190-192).

W.2 **Weisz, J. D.**, Maher, M. L., Strobelt, H., Chilton, L. B., Bau, D., & Geyer, W. (2022, March). HAI-GEN 2022: 3rd Workshop on Human-AI Co-Creation with Generative Models. In 27th International Conference on Intelligent User Interfaces (pp. 4-6).

W.1 Geyer, W., Chilton, L. B., **Weisz, J. D.**, & Maher, M. L. (2021, April). HAI-GEN 2021: 2nd Workshop on Human-AI Co-Creation with Generative Models. In 26th International Conference on Intelligent User Interfaces-Companion (pp. 15-17).

Journal Articles

J.2 Wang, A. Y., Wang, D., Drozdal, J., Muller, M., Park, S., **Weisz, J. D.**, Liu, X., Wu, L, and Dugan, C. (2022). Documentation Matters: Human-Centered AI System to Assist Data Science Code Documentation in Computational Notebooks. ACM Transactions on Computer-Human Interaction, 29(2), 1-33.

J.1 Coden, A., Lin, W. S., Houck, K., Tanenblatt, M., Boston, J., MacNaught, J. E., Soroker D., **Weisz, J. D.**, Pan, S., Lai, J-H., Lu, J., Wood, S., Xia, Y., and Lin, C. Y. (2016). Uncovering insider threats from the digital footprints of individuals. IBM Journal of Research and Development, 60(4), 8-1.

Book Chapters

BC.2 **Weisz, J. D.**, He, J., Muller, M. (2025). Human-centered AI design principles for generative AI applications. In W. Xu (eds) Handbook of Human-Centered Artificial Intelligence. Springer, Singapore.

BC.1 **Weisz, J. D.** (2009). Online video as a social activity. In P. Cesar, D. Geerts, and K. Chorianopoulos (Eds.), Social Interactive Television: Immersive Shared Experiences and Perspectives. Hershey, PA: IGI Global.

Patents

US9342436B2 Capture and display of historical run-time execution traces in a code editor
US9633003B2 System support for evaluation consistency
US9699131B2 User interface with recipient status indication
US9747760B2 Safety equipment criteria verification
US9792576B1 Operating a plurality of drones and trucks in package delivery
US9858794B2 Detecting and notifying of various potential hazards
US10081426B2 Drone-based mosquito amelioration based on risk analysis and pattern classifiers
US10268458B1 Prospective voice user interface modality identification
US10390752B2 Cognitive chair for monitoring Parkinson's disease progression
US10393355B2 Lighting pattern optimization for a task performed in a vicinity
US10528888B2 Identifying locations of potential user errors during manipulation of multimedia content
US10529446B2 Continuous health care plan coordination between patient and patient care team
US10565982B2 Training data optimization for voice enablement of applications
US10585640B2 Automated voice enablement of applications
US10620911B2 Machine learning to identify a user interface trace
US10620912B2 Machine learning to determine and execute a user interface trace
US10628337B2 Communication mode control for wearable devices
US10693827B2 Determining intended electronic message recipients via linguistic profiles
US10733286B2 Composite propensity profile detector
US10832160B2 Predicting user attentiveness to electronic notifications
US10887396B2 Sensor data manipulation using emulation
US10904488B1 Generated realistic representation of video participants
US10915118B2 System and method to operate a drone
US10967972B2 Vehicular alert system
US10970051B2 Systems and methods for programming drones
US10976446B2 Acceleration of real time computer vision processing on UAVs through GPS attitude estimation
US11032606B2 System, method, and recording medium for providing notifications in video streams to control video playback
US11165722B2 Cognitive messaging with dynamically changing inputs
US11245639B2 Composition of persistent object instances linking resources across multiple, disparate systems
US11322031B2 Method of optimized path planning for UAVs for the purpose of ground coverage
US11360576B1 Movement pattern-based mobile device user interfaces
US11429360B1 Computer assisted programming with targeted visual feedback
US11663851B2 Detecting and notifying for potential biases in artificial intelligence applications
US11747970B2 Interactive graphical display of multiple overlapping hypotheses or document versions
US12380283B2 Summarization for attention management and recall

OPEN SOURCE

- **Entanglion**. World's first open source board game to teach fundamental principles of quantum computing. Used in hundreds of high school and university classrooms as part of curriculum.
- **TJBot**. Paper / 3D-printed robot designed to teach people how to program with IBM's Watson AI services. Used by hundreds of thousands of students & makers worldwide.
- **TJBot: Community Edition**. Modernized TJBot software stack with TypeScript and Python libraries and support for on-device STT/TTS/Vision models.
- **Idea Gauntlet**. Web-based game for debating ideas against AI agents, designed to showcase how AI can help improve critical-thinking skills.
- **CardKit**. Card-based programming method to enable business users to program IoT devices with AI services.

AWARDS & HONORS

IBM Technical Achievement Awards
Awarded for technical contributions with substantial business impact.

- AI Governance Innovations to watsonx.gov and Open Source (Outstanding) 2025
- Leadership in Human-AI Co-Creation (Outstanding) 2024
- Generative AI for Designers 2023
- Leadership in Human-AI Co-Creation 2022
- Research contributions to IBM Watson AutoAI product (Outstanding) 2019

IBM Volunteer Excellence Award 2021
Awarded for contributions to TJBot classroom learning activities in support of Vietnam Teacher's Day.

IBM Master Inventor 2016
Appointed for IP mentorship activities and development of IP in areas strategic to IBM.

PROFESSIONAL SERVICE

Conferences & Workshops

Workshop Co-chair, Workshop on Agentic AI Systems Development (AgentCraft) 2026
Workshop Co-chair, Theory of Mind in Human-AI Interaction (ToMinHAI, ToMinCUI) 2024, 2025
Workshop Co-chair, Human-Centered Explainable AI (HCXAI) 2023, 2025, 2026
AC, Computer Supported Cooperative Work (CSCW) 2021, 2023
Workshop Co-chair, Human-AI Co-Creation with Generative Models (HAI-GEN) 2021 – 2025
Reviewer, IUI 2014 – present
Reviewer, CSCW 2006 – present
Reviewer, CHI 2006 – present

Human-Centered AI Community

Founder & Editor-in-Chief, Human-Centered AI publication on Medium 2023 – present

Advising

Referee, M.S. Thesis, Jule Ebeling, Bauhaus-Universität Weimar, Germany 2024 – 2025
Advisor, HCI M.S. Capstone, Amin El Asery & Lakshya Sharma, Georgia Tech University 2024 – 2025

IBM

Invention Disclosure Team (IDT) Member 2018 – 2025
Member of multiple patent disclosure review committees for patent disclosures related to AI and HCI.

Pat Goldberg Memorial Best Paper Award 2017 – 2021 (Reviewer); 2022 – 2025 (Co-chair)
Member of award committee to identify the most technically significant and impactful papers published by IBM Research each year.

Pillar Lead, Cognitive Professional Interest Communities (PICs) 2017 – 2019
Oversaw activities that strengthen our research community, including recruiting and student outreach activities at academic conferences and invited speakers.

Chair, Human-Computer Interaction Professional Interest Community (PIC) 2013 – 2016
Maintained a strong HCI community within IBM and promoting HCI research externally via conference sponsorships, hosting external speakers, and recruiting talented students.

Reviewer, IBM Ph.D. Fellowship 2010 – 2025
Identified top Ph.D. candidates whose research interests align with the IBM business.

Carnegie Mellon University

Board Member, Alumni Advisory Board, School of Computer Science 2011 – 2020
Provided feedback to the Dean of Computer Science on new department initiatives, faculty hiring, curriculum, and student career opportunities.