HAI-GEN 2024: 5th Workshop on Human-AI Co-Creation with Generative Models

Werner Geyer
IBM Research AI
Cambridge, Massachusetts, USA
werner.geyer@us.ibm.com

Mary Lou Maher
UNC Charlotte
Charlotte, North Carolina, USA
m.maher@uncc.edu

Justin D. Weisz
IBM Research AI
Yorktown Heights, New York, USA
jweisz@us.ibm.com

Daniel Buschek
University of Bayreuth
Germany
daniel.buschek@uni-bayreuth.de

Lydia B. Chilton
Columbia University
New York City, New York, USA
chilton@cs.columbia.edu

ABSTRACT

Generative AI has rapidly entered the public consciousness with the development of applications such as ChatGPT, Midjourney, and GitHub Copilot. Nielsen recently argued that we have entered a new era of human-computer interaction in which users need only specify *what* they want and not *how* it should be produced [1]. This paradigm of intent-based outcome specification shifts control over from people to AI, enabling new forms of co-creativity and co-creation. Although these systems are capable of holding fluent conversations and producing high-fidelity images, difficulties remain regarding their ability to produce outputs that satisfy their users' needs. Our workshop will bring together researchers and practitioners from both the HCI and AI disciplines to explore the implications of this shift in control, deepen our understanding of the human-AI co-creative process, and examine how we can design, build, use, and evaluate human-AI co-creative systems that are both effective and safe.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**; *Interaction design*; • **Computing methodologies** → **Artificial intelligence**; • **Applied computing** → *Arts and humanities.*

KEYWORDS

Generative modelling, artificial intelligence, generative design, user experience, collaboration, creativity

ACM Reference Format:
Werner Geyer, Mary Lou Maher, Justin D. Weisz, Daniel Buschek, and Lydia B. Chilton. 2024. HAI-GEN 2024: 5th Workshop on Human-AI Co-Creation with Generative Models. In *29th International Conference on Intelligent User Interfaces - Companion (IUI Companion '24), March 18–21, 2024, Greenville, SC, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3640544.3645254

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
IUI Companion '24, March 18–21, 2024, Greenville, SC, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0509-0/24/03
https://doi.org/10.1145/3640544.3645254

1 DESCRIPTION

Generative AI technologies have introduced a paradigm shift in machine learning and are rapidly being adopted by consumers and enterprises alike. Consumer applications like ChatGPT, DALL-E, and Midjourney are making access to these technologies widely available and have fueled the rapid commercialization of generative AI. At the same time, generative models have enabled a radically new way for people to interact with computing technology by allowing them to specify *what they want* rather than *how it should be produced.* Users are able to create specifications of the kinds of outputs they desire in natural language, through sketches and gestures, novel UI controls or other means. In a recent article, Nielsen [1] called this form of interaction *intent-based outcome specification.* He argues that it is the first new interaction paradigm in 60 years, fundamentally different from batch processing and command-based interaction. Intent-based outcome specification shifts control over how computation is performed from the user towards the generative AI models, enabling new forms of co-creativity and co-creation while also creating new challenges of how to best support users with creating outcome specifications that yield the desired results and are effective and safe to use.

The central questions of our workshop are: How do we design, implement, and evaluate intent-based co-creative experiences that enhance human creativity in their work, play, and education across a range of media types (text, images, audio, code, and video)? How will user needs for creativity-support drive the development of generative AI algorithms? How can the capabilities of generative models be leveraged for positive and effective co-creative user experiences? The goal of our workshop is to bring together researchers and practitioners from the domains of HCI & AI to establish a joint community to deepen our understanding of the human-AI co-creative process and to explore the opportunities and challenges of creating meaningful, effective and safe user experiences for intent-based generative systems. Submissions are encouraged, but not limited to, the following topics:

- Novel, AI-augmented user experiences across a range of media types (text, images, audio, code, and video) that support the creation of physical and/or digital artifacts
- Novel ways for creating, managing, and interacting with outcome specifications
- Use cases & novel applications for co-creativity including business use cases

- Techniques, methodologies, & algorithms that enable new user experiences and interactions with generative models and allow for directed and purposeful manipulation of the model output
- Usage patterns, design guidelines, and design principles for intent-based outcome specification that support effective and safe use
- Issues of governance, privacy, and ownership of AI-generated or human-AI co-created content
- Security, including forensic tools and approaches for deep fake detection
- Evaluations of human-AI co-creative processes and quality metrics of AI-generated or human-AI co-created content
- User research on needs & algorithmic requirements for co-creative systems, perceptions of human-AI co-creative systems, trust of co-creative tools & artifacts, and/or implications for HCI theories
- Lessons learned from computational art & design and generative design, and how these impact research
- Potential risks and harms of working with generative models and generated content, and strategies to detect or mitigate those harms

2 PREVIOUS HISTORY

- **HAI-GEN 2023** (https://hai-gen.github.io/2023). The 4th HAI-GEN workshop was held in conjunction with IUI 2023. We held an online workshop with approximately 50 participants, 8 paper presentations, and 1 keynote. Proceedings were published in CEUR.
- **HAI-GEN 2022** (https://hai-gen.github.io/2022). The 3rd HAI-GEN workshop was held in conjunction with IUI 2022. We held an online workshop with approximately 54 participants, 8 paper presentations, and 1 keynote. Proceedings were published in CEUR.
- **HAI-GEN 2021** (https://hai-gen2021.github.io). The 2nd HAI-GEN workshop was held in conjunction with IUI 2021. We held an online workshop with approximately 45 participants, 12 paper presentations, and 2 keynotes. Proceedings were published in CEUR.
- **HAI-GEN 2020** (https://hai-gen2020.github.io). The 1st HAI-GEN workshop was held in conjunction with IUI 2020. As IUI 2020 was canceled due to COVID-19, we held an online workshop with approximately 35 participants, 7 paper presentations, and 2 keynotes. Proceedings were published in CEUR.

Prior to HAI-GEN 2020, the organizers were involved in several other workshops and events that inspired the creation of HAI-GEN, including **Machine Learning for Creativity and Design** (https://neurips2019creativity.github.io/), **GANocracy** (http://ganocracy.csail.mit.edu/), and **Art in the Age of Algorithms** (https://engineering.columbia.edu/news/art-age-algorithms).

3 ORGANIZERS

Werner Geyer (contact author) is a Principal Research Staff Member and Senior Research Manager at IBM Research in Cambridge, MA, where he is leading a research team centered around AI Interaction technologies. He is also strategy lead for Human-Centered AI at IBM Research and he has held various roles as co-chair at ACM RecSys, including general chair as well as a series of workshops and tutorials on Social Recommender Systems. More recently, his team is exploring generative modeling techniques in business settings. https://research.ibm.com/people/werner-geyer, *werner.geyer@us.ibm.com*

Mary Lou Maher is a Professor and Director of the Human Centered Computing Lab in the College of Computing and Informatics at UNC Charlotte. Her research in computational creativity extends genetic algorithms, analogical reasoning, and more recently deep learning models to automate the generation of creative artifacts to enhance human creativity. She has served on the organizing committees for workshops at CHI and IUI on interaction with generative AI. http://maryloumaher.net, *m.maher@uncc.edu.*

Justin D. Weisz is a Research Manager at IBM Research in Yorktown Heights, NY. He is a co-lead for Human-Centered AI strategy and he leads the Human-AI Collaboration team. He was a co-organizer of HAI-GEN 2021-2023 and is the PI of a project that is exploring enterprise use cases of generative AI technologies. https://research.ibm.com/people/justin-weisz, *jweisz@us.ibm.com.*

Daniel Buschek is a Professor of Mobile Intelligent User Interfaces at the University of Bayreuth, Germany. His group works at the intersection of Human-Computer Interaction and AI, exploring interaction with models for text and images to shape the future of AI tools in a human-centred way. https://www.hciai.uni-bayreuth.de/, *daniel.buschek@uni-bayreuth.de.*

Lydia B. Chilton is an Assistant Professor in the Computer Science Department at Columbia University. For ten years she was a leader in the crowdsourcing research space of HCI, now she breaks down problems for a combination of people and AI to solve. She organized the first CHI Workshop on Crowdsourcing and Human Computation, which had over 100 attendees. She has lead the 2-day crowdsourcing workshop and hackathon, CrowdCamp 3 times. https://www.cs.columbia.edu/~chilton/, *chilton@cs.columbia.edu*

4 WORKSHOP PROGRAM COMMITTEE

- John Chung, Midjourney
- Hai Dang, University of Bayreuth
- Florian Fischer, University of Bayreuth
- Sebastian Gehrmann, Bloomberg
- Kazjon Grace, The University of Sydney
- Yiren Liu, University of Illinois at Urbana-Champaign
- Ryan Louie, Stanford University
- Brian Magerko, Georgia Institute of Technology
- Nikolas Martelaro, Carnegie Mellon University
- Michael Muller, IBM Research
- Jeba Rezwana, Towson University
- Steven Ross, IBM Research
- Koray Tahiroğlu, Aalto University
- Dan Ventura, Brigham Young University
- Tim Zindulka, University of Bayreuth

5 WORKSHOP PROGRAM

The workshop (https://hai-gen.github.io/2024) was organized in the form of a **mini conference** with a keynote speaker, paper presentations, discussion sessions, and a poster and demo session. Our keynote speaker was Mauro Martino, an artist, designer, and researcher renowned for pioneering work at the intersection of art, data, and artificial intelligence. In his talk, Mauro Martino presented Urbana Verba - Generative Urban Cinema, a take on generative cinema to reinterpret and recount our cities in entirely novel ways.

We accepted 13 papers, 2 demos, and 3 posters spanning multiple thematic areas. Authors proposed novel tools and approaches based on LLMs for group ideation and brainstorming, the use of generative AI in visual arts, poetry, and music, novel conceptual models for co-creativity and evaluation of co-creativity, conversational and agent-based approaches to video editing, the application of generative AI for learning and co-learning, novel ways of interacting and discovering with generative AI, ethical considerations, and risks and harms of generative co-creative AI. The workshop included ample discussion time to enable participants to learn, collaborate, and plan for future collaboration.

REFERENCES

[1] Jakob Nielsen. 2023. AI: First New UI Paradigm in 60 Years. *Nielsen Norman Group* (18 06 2023). https://www.nngroup.com/articles/ai-paradigm/