HAI-GEN 2023: 4th Workshop on Human-AI Co-Creation with Generative Models

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

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

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

Hendrik Strobelt
IBM Research AI
Cambridge, Massachusetts, USA
hendrik.strobelt@ibm.com

ABSTRACT

Recent advances in generative AI have resulted in a rapid and dramatic increase in the fidelity of created artifacts, from realistic-looking images of faces [10] to antimicrobial peptide sequences that treat diseases [5] to faked videos of prominent business leaders [4, 11]. We believe that people skilled within their creative domain can realize great benefits by incorporating generative models into their own work: as a source of inspiration, as a tool for manipulation, or as a creative partner. Our workshop will bring together researchers and practitioners from both the HCI and AI disciplines to explore and better understand the opportunities and challenges in building, using, and evaluating human-AI co-creative systems.

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:
Mary Lou Maher, Justin D. Weisz, Lydia B. Chilton, Werner Geyer, and Hendrik Strobelt. 2023. HAI-GEN 2023: 4th Workshop on Human-AI Co-Creation with Generative Models. In *28th International Conference on Intelligent User Interfaces Companion (IUI '23 Companion).* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3581754.3584166

1 DESCRIPTION

Recent advances in generative AI through deep learning approaches such as generative adversarial networks (GANs) [7], multi-modal generation [8, 9], and large language models [1, 2] will enable new kinds of user experiences around content creation, across a range of media types (text, images, audio, code, and video). These advances have enabled content to be produced with an unprecedented level of fidelity, for tasks such as generating faces [7], prose and poems [2], deep fake videos of celebrities [4, 11], music [6], and code [3]. In many cases, content generated by generative models is either indistinguishable from human-generated content or could not be produced by human hands. These examples also highlight some of the significant societal, ethical, and organizational challenges generative AI is posing around issues such as security, privacy, ownership, quality metrics, and evaluation of generated content.

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 user experiences with deep generative models. We envision that the user experience of creating both physical and digital artifacts will become a partnership between people and AI: people will take the role of specification, goal setting, steering, high-level creativity, curation, and governance, whereas AI will augment human abilities through inspiration, creative contributions, refinement, and the ability to design at scale. The central questions of our workshop are: How do we design, implement, and evaluate co-creative systems that enhance human creativity in their work, play, and education? How will user needs drive the development of generative AI algorithms? How can the capabilities of generative models be leveraged for positive and effective co-creative user experiences?

Submissions were encouraged, but not limited to, the following topics:

- Novel, AI-augmented user experiences that support the creation of physical and/or digital artifacts.
- Business use cases & novel applications of generative models.
- Techniques, methodologies, & algorithms that enable new user experiences and interactions with generative models and allow for directed and purposeful manipulation of the model output.
- Interaction with generative models for educational purposes.
- 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.

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 '23 Companion, March 27–31, 2023, Sydney, Australia
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0107-8/23/03.
https://doi.org/10.1145/3581754.3584166

- 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 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

Mary Lou Maher (Primary Contact) 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 & 2022 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.*
Hendrik Strobelt is the Explainability Lead at the MIT-IBM Watson AI Lab and Senior Research Scientist at IBM Research. His recent research is on visualization for and human collaboration with AI models to foster explainability and intuition. His work involves NLP models and generative models. He is advocating to utilize a mix of data modalities to solve real-world problems. His research is applied to tasks in machine learning, in NLP, in the biomedical domain, and in chemistry. http://hendrik.strobelt.com, *hendrik.strobelt@ibm.com.*
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.*
Werner Geyer is a Principal Research Scientist and Senior Research Manager at IBM Research in Cambridge, MA, where he is leading a research team centered around AI Interaction technologies. He is leading the Human-Centered AI strategy at IBM Research. Werner 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.*

4 WORKSHOP PROGRAM COMMITTEE

- Maya Ackerman, Santa Clara University
- Daniel Buschek, University of Bayreuth
- Katy Gero, Columbia University
- Kazjon Grace, The University of Sydney
- Anna Huang, Google
- Ryan Louie, Northwestern University
- Brian Magerko, Georgia Tech
- Charles Martin, Australian National University
- Michael Muller, IBM Research AI
- Jeba Rezwana, University of North Carolina, Charlotte
- Steven Ross, IBM Research AI
- Dan Ventura, Brigham Young University

5 WORKSHOP PROGRAM

The workshop was organized in the form of a **half-day mini conference** with a keynote speaker, paper presentation and discussion sessions, and panel discussions based on position statements from the Program Committee. Our keynote speaker was Philipp Schmitt, an artist and designer based in Brooklyn, NY, who presented a talk on generative AI and art.

We accepted 8 papers spanning multiple thematic areas. Authors proposed design guidelines for interacting with generative AI systems and examined ethical dilemmas in human-AI co-creation. Paper presentations and discussions investigated prompt engineering with a conversational programming assistant, new methods for interacting with text-to-image models, and examined the use of generative models to aid in the creation of music and poetry, as well as ideation in generating scientific research questions. The workshop included ample discussion time to enable participants to learn, collaborate, and plan for future collaboration.

REFERENCES

[1] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma

Brunskill, et al. 2021. On the Opportunities and Risks of Foundation Models. *arXiv preprint arXiv:2108.07258* (2021).

[2] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).

[3] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde, Jared Kaplan, Harri Edwards, Yura Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).

[4] Samantha Cole. 2019. This Deepfake of Mark Zuckerberg Tests Facebook's Fake Video Policies. https://www.vice.com/en_us/article/ywyxex/deepfake-of-mark-zuckerberg-facebook-fake-video-policy

[5] Payel Das, Kahini Wadhawan, Oscar Chang, Tom Sercu, Cicero Dos Santos, Matthew Riemer, Vijil Chenthamarakshan, Inkit Padhi, and Aleksandra Mojsilovic. 2018. Pepcvae: Semi-supervised targeted design of antimicrobial peptide sequences. *arXiv preprint arXiv:1810.07743* (2018).

[6] Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. 2020. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341* (2020).

[7] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *NIPS*.

[8] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. https://doi.org/10.48550/ARXIV.2204.06125

[9] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-Shot Text-to-Image Generation. https://doi.org/10.48550/ARXIV.2102.12092

[10] James Vincent. 2018. These faces show how far AI image generation has advanced in just four years. https://www.theverge.com/2018/12/17/18144356/ai-image-generation-fake-faces-people-nvidia-generative-adversarial-networks-gans

[11] James Vincent. 2019. Listen to this AI voice clone of Bill Gates created by Facebook's engineers. https://www.theverge.com/2019/6/10/18659897/ai-voice-clone-bill-gates-facebook-melnet-speech-generation