HAI-GEN 2022: 3rd Workshop on Human-AI Co-Creation with Generative Models

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

Mary Lou Maher
University of North Carolina, Charlotte
Charlotte, NC, USA
m.maher@uncc.edu

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

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

David Bau
Northeastern University
Boston, MA, USA
d.bau@northeastern.edu

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

ABSTRACT

Recent advances in generative AI have resulted in a rapid and dramatic increase to the fidelity of created artifacts, from realistic-looking images of faces [9] to antimicrobial peptide sequences that treat diseases [5] to faked videos of prominent business leaders [4, 10]. 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:
Justin D. Weisz, Mary Lou Maher, Hendrik Strobelt, Lydia B. Chilton, David Bau, and Werner Geyer. 2022. HAI-GEN 2022: 3rd Workshop on Human-AI Co-Creation with Generative Models. In *27th International Conference on Intelligent User Interfaces (IUI '22 Companion), March 22–25, 2022, Helsinki, Finland*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3490100.3511166

1 DESCRIPTION

Recent advances in generative AI through deep learning approaches such as generative adversarial networks (GANs) [7], variational autoencoders (VAEs) [8], 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, 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, 10], music [6], and even 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 powerful 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, creativity, low-level detail work, and the ability to design at scale.

The central question of our workshop is *how can we build co-creative systems that make people feel that they have "creative superpowers"?* How will user needs drive the development of generative AI algorithms, and how can the capabilities of generative models be leveraged to create effective co-creative user experiences?

Submissions took the form of papers and demos, following the IUI paper and demo guidelines. 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
- Issues of governance, privacy, and ownership of AI-generated or human-AI co-created content

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 '22 Companion, March 22–25, 2022, Helsinki, Finland
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9145-0/22/03.
https://doi.org/10.1145/3490100.3511166

- 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

Workshop information, including our program and call for participation was hosted on our web site: http://hai-gen.github.io/2022/.

2 PREVIOUS HISTORY

- **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 via WebEx 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 via WebEx 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

- **Justin D. Weisz** (Primary Contact) 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 and is the PI of a project that is exploring enterprise use cases of generative AI technologies. https://researcher.watson.ibm.com/researcher/view.php?person=us-jweisz, *jweisz@us.ibm.com*.
- **Mary Lou Maher** is a Professor and Director of the HCI Lab in the College of Computing and Informatics at UNC Charlotte. She has a joint appointment in the Department of Software and Information Systems and the School of Data Science. Her research in computational creativity extends genetic algorithms, analogical reasoning, and more recently deep learning models to automate the generation of creative designs and respond to the question: Can computers be creative? Her research in HCI has lead to a complementary focus in computational creativity to design and evaluate interaction models for co-creative systems and respond to the question: Can computational creativity enhance human creativity? http://maryloumaher.net, *m.maher@uncc.edu*.
- **Hendrik Strobelt** is the Explainability Lead at the MIT-IBM Watson AI Lab and 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.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*
- **David Bau** is an Assistant Professor at the Northeastern University Khoury School of Computer Science. David has pioneered methods for understanding and controlling the internal structure of deep networks. His work has enabled new creative applications such as semantic paint using neurons of a generative network, and direct rewriting of rules within deep generative models through user interactions. http://davidbau.com/research/, *d.bau@northeastern.edu*
- **Werner Geyer** is a Principal Research Staff Member and Research Manager at IBM Research in Cambridge, MA, where he is leading a research team centered around AI Interaction technologies. 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://researcher.watson.ibm.com/researcher/view.php?person=us-Werner.Geyer, *werner.geyer@us.ibm.com*

4 WORKSHOP PROGRAM COMMITTEE

- Maya Ackerman, Santa Clara University
- Daniel Buschek, University of Bayreuth
- Sebastian Gehrmann, Google
- Katy Gero, Columbia University
- Kaz Grace, The University of Sydney
- Anna Huang, Google
- Narendra Nath Joshi, IBM Research AI
- Ryan Louie, Northwestern University
- Michael Mateas, University of California, Santa Cruz
- Michael Muller, IBM Research AI
- Jeba Rezwana, University of North Carolina, Charlotte
- Steven Ross, IBM Research AI

5 WORKSHOP PROGRAM

The workshop was organized in the form of a **half-day mini conference** with a keynote speaker and multiple paper and demo sessions. Our keynote speaker was Sarah Schwettmann, a neuroscientist and machine learning researcher at MIT CSAIL. She presented a talk titled, "AI Models as a Shared Creative Resource: The Human Use of New Machines."

We accepted 8 papers and demos spanning multiple thematic areas. This work explored frameworks for interacting with and controlling generative models, as well as evaluating their expressive range. It also included examinations of the use of generative AI

technologies in specific application areas including writing, music, and software engineering. Each session included ample discussion time to enable participants to learn, collaborate, and plan for future collaboration.

5.1 Accepted Papers

- Stephanie Houde, Steven Ross, Michael Muller, Mayank Agarwal, Fernando Martinez, John Richards, Kartik Talamadupula and Justin D. Weisz. *Opportunities for Generative AI in UX Modernization*
- Max Kreminski, Isaac Karth, Michael Mateas and Noah Wardrip-Fruin. *Evaluating Mixed-Initiative Creative Interfaces via Expressive Range Coverage Analysis*
- Charles Martin. *Performing with a Generative Electronic Music Controller*
- Daijin Yang, Yanpeng Zhou, Zhiyuan Zhang, Toby Li and Ray Lc. *AI as an Active Writer: Interaction strategies with generated text in human-AI collaborative fiction writing*
- Michael Muller, Steven Ross, Stephanie Houde, Mayank Agarwal, Fernando Martinez, John Richards, Kartik Talamadupula and Justin D. Weisz. *Drinking Chai with Your (AI) Programming Partner: A Design Fiction about Generative AI for Software Engineering*
- Imke Grabe, Miguel González-Duque, Sebastian Risi and Jichen Zhu. *Towards A Framework for Human-AI Interaction Patterns in Co-Creative GAN Applications*
- Halley Young, Vincent Dumoulin, Pablo Castro, Jesse Engel and Cheng-Zhi Anna Huang. *Composing Features: Compositional Model Augmentation for Steerability of Music Transformers*
- John Chung, Minsuk Chang and Eytan Adar. *Gestural Inputs as Control Interaction for Generative Human-AI Co-Creation*

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] Diederik P. Kingma and Max Welling. 2013. Auto-Encoding Variational Bayes. *CoRR* abs/1312.6114 (2013).

[9] 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

[10] 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