HAI-GEN 2021: 2nd Workshop on Human-AI Co-Creation with Generative Models

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

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

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

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

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 [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:
Werner Geyer, Lydia B. Chilton, Justin D. Weisz, and Mary Lou Maher. 2021. HAI-GEN 2021: 2nd Workshop on Human-AI Co-Creation with Generative Models. In *26th International Conference on Intelligent User Interfaces (IUI '21 Companion), April 14–17, 2021, College Station, TX, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3397482.3450707

1 DESCRIPTION

Recent advances in generative AI through deep learning approaches such as generative adversarial networks (GANs) [6], variational autoencoders (VAEs) [9], and language models [7] will enable new kinds of user experiences around content creation. These advances have enabled content to be produced with an unprecedented level of fidelity. Goodfellow's work on face generation [6] and StyleGan [8], OpenAI's GPT-3 [3] model, and recent deep fake videos of Mark Zuckerberg [4] and Bill Gates [11] are all prominent examples of the state of the art in the generation of text, image, and video content. 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.

While the areas of computational creativity, generative design, and computational art have existed for some time (e.g. [1, 2]), these communities have not been grounded at the intersection of generative deep learning approaches and human-computer interaction.

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 short papers, long 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

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 '21 Companion, April 14–17, 2021, College Station, TX, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8018-8/21/04.
https://doi.org/10.1145/3397482.3450707

- 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

2 ORGANIZERS

- **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's been holding 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 modelling techniques in business settings. https://researcher.watson.ibm.com/researcher/view.php?person=us-Werner.Geyer
- **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/
- **Justin D. Weisz** is a Research Staff Member and Research Manager at IBM Research in Yorktown Heights, NY. He leads the Human-AI Collaboration team, whose mission is to design, build, and rigorously investigate new forms of human-AI partnerships that enhance and extend human capabilities. He also leads the HCI research strategy within the AI organization, as well as a new project exploring enterprise use cases of generative AI technologies. https://researcher.watson.ibm.com/researcher/view.php?person=us-jweisz
- **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

3 HISTORY & RELATED WORKSHOPS

- **HAI-GEN 2020**. We organized the first HAI-GEN Workshop at IUI'20. Although IUI was canceled due to COVID-19, we held a successful online workshop with two keynotes via WebEx. https://hai-gen2020.github.io.
- **Machine Learning for Creativity and Design**. Draws 200-300 participants and has been around since 2017, but is focused more on art and design. https://neurips2019creativity.github.io/
- **GANocracy** First held in 2019 and drew 50 participants. Members of our OC and PC were involved in its organization. The content specifically focused on algorithmic approaches (e.g. GAN). http://ganocracy.csail.mit.edu/
- **Art in the Age of Algorithms**. Local group at Columbia focused on computational art and design. https://engineering.columbia.edu/news/art-age-algorithms

4 WORKSHOP PROGRAM COMMITTEE

- Maya Ackerman, Santa Clara University
- Carrie Cai, Google
- Nicholas Davis, University of North Carolina, Charlotte
- Sebastian Gehrmann, Google
- Katy Gero, Columbia University
- Kaz Grace, The University of Sydney
- Anna Kantosalo, Aalto University
- Michael Muller, IBM Research AI
- Hendrik Strobelt, IBM Research AI
- Haiyi Zhu, Carnegie Mellon University

5 WORKSHOP PROGRAM

The workshop was organized in form of a **full-day mini conference** with two keynote speakers and multiple paper and demo sessions. One keynote was presented by Refik Anadol, a media artist, director, and pioneer in the aesthetics of machine intelligence. His work addresses the challenges and possibilities that ubiquitous computing has imposed on humanity. A second keynote was presented by Dr. Maya Ackerman, Assistant Professor at the Department of Computer Science and Engineering at Santa Clara University, and founder of WaveAI, a musical AI startup.

We accepted 12 papers and demos in total spanning multiple thematic sessions. The papers consisted of general research on designing and building generative co-creative systems, as well as studies and demos of real-world systems for specific co-creative experiences with visual, auditory, textual, and other artifacts. Each session included ample discussion time to enable participants to learn, collaborate, and plan for future collaboration.

5.1 Accepted Papers

- Kenneth Arnold, April Volzer and Noah Madrid. *Generative Models can Help Writers without Writing for Them*
- Daniel Buschek, Lukas Mecke, Florian Lehmann and Duong Hai Dang. *Nine Potential Pitfalls when Designing Human-AI Co-Creative Systems*
- Michelle Carney, Chong Li, Edwin Toh, Nida Zada, Ping Yu and Jesse Engel. *Tone Transfer: In-Browser Interactive Neural Audio Synthesis*

- Doron Friedman and Dan Pollak. *Image Co-Creation by Non-Programmers and Generative Adverserial Networks*
- Eyal Gruss. *Avatars for All!*
- Yucheng Jin, Ge Gao, Tiffany Ong and Xi Su. *CreativeSVG: A Creativity Support Tool for Abstract Background Design with Generative Vector Graphics*
- Jingoog Kim, Safat Siddiqui and Mary Lou Maher. *Measuring Ideation in AI-based Co-Creative Design*
- Vivian Liu and Lydia Chilton. *Semantically Guided Generation of 3D Animal Shapes*
- Kyle McDonell and Laria Reynolds. *Multiversal views on language models*
- Paolo Morettin, Andrea Passerini and Stefano Teso. *Co-creating Platformer Levels with Constrained Adversarial Networks*
- Michael Muller, April Yi Wang, Steven I. Ross, Justin D. Weisz, Mayank Agarwal, Kartik Talamadupula, Stephanie Houde, Fernando Martinez, John Richards, Jaimie Drozdal, Xuye Liu, David Piorkowski and Dakuo Wang. *How Can Generative Models Ask Humans for Help?*
- Jeba Rezwana, Mary Lou Maher and Nicholas Davis. *Creative PenPal: A Virtual Embodied Conversational Agent to Improve User Engagement and Collaborative Experience in Human-AI Co-Creative Design Ideation*

REFERENCES

[1] 2019. Machine Learning for Creativity and Design. https://neurips2019creativity.github.io/

[2] 2020. International Conference on Computational Creativity. http://computationalcreativity.net/iccc20/

[3] 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).

[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] 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*.

[7] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Comput.* 9, 8 (Nov. 1997), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

[8] Tero Karras, Samuli Laine, and Timo Aila. 2018. A Style-Based Generator Architecture for Generative Adversarial Networks. *ArXiv* abs/1812.04948 (2018).

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

[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