

# HAI-GEN 2020: Workshop on Human-AI Co-Creation with Generative Models

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

Ranjitha Kumar  
University of Illinois at Urbana-Champaign  
Urbana, Illinois  
ranjitha@illinois.edu

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

Adam Tauman Kalai  
Microsoft Research  
Cambridge, Massachusetts  
adam.kalai@microsoft.com

## ABSTRACT

Recent advances in generative modeling will enable new kinds of user experiences around content creation, giving us “creative superpowers” and move us toward co-creation. This workshop brings together researchers and practitioners from both fields HCI and AI to explore and better understand both the opportunities and challenges of generative modelling from a Human-AI interaction perspective for the creation of both physical and digital artifacts.

## 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, Ranjitha Kumar, and Adam Tauman Kalai. 2020. HAI-GEN 2020: Workshop on Human-AI Co-Creation with Generative Models. In *25th International Conference on Intelligent User Interfaces Companion (IUI'20 Workshops)*, March 17, 2020, Cagliari, Italy. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 DESCRIPTION

Recent advances in generative modeling through deep learning approaches such as generative adversarial networks (GANs) [5], variational autoencoders (VAEs) [8], and sequence-to-sequence models [6] will enable new kinds of user experiences around content creation, giving us “creative superpowers” and move us toward co-creation and curation. While the areas of computational design, generative design, and computational art have existed for some time, content with unprecedented fidelity is now being produced

due to breakthroughs in generative modeling using deep learning. Ian Goodfellow’s work on face generation [5] and StyleGan [7], OpenAI’s GPT-2 [9], or recent deep fake videos of Mark Zuckerberg [4] and Bill Gates [10] are prominent examples of content generated by AI that is almost indistinguishable from human-generated content. These examples also highlight some of the significant societal, ethical and organizational challenges generative AI is posing including security, privacy, ownership, quality metrics and evaluation of generated content.

The goal of this workshop is to bring together researchers and practitioners from both fields HCI and AI to explore the opportunities and challenges of generative modelling from an HCI perspective. We envision that the user experience of creating both physical and digital artifacts will become a partnership of humans and AI: Humans will take the role of specification, goal setting, steering, high-level creativity, curation, and governance. AI will augment human abilities through inspiration, low level creativity and detail work, and the ability to test ideas at scale.

Submissions in form of short papers, long papers and demos following the IUI paper and demo guidelines are encouraged but not limited to the following topics:

- Novel user experiences supporting the creation of both physical and digital artifacts in an AI augmented fashion
- Business use cases of generative models
- 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
- Governance, privacy, content ownership
- Security including forensic tools and approaches for deep fake detection
- Evaluation of generative approaches and quality metrics
- User studies
- Lessons learned from computational art and design, and generative design and how these impact research

Generative Design, Computational Design, or Computational Art are topics that have been around for a while (e.g. [1],[2],[3]) but for the most part have not been grounded in generative deep learning approaches combined with a strong HCI perspective and theory on co-creation. This workshop aims to lay the groundwork for bringing this exciting area deeper into the field of HCI 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. His website is <https://researcher.watson.ibm.com/researcher/view.php?person=us-Werner.Geyer>. He can be reached at [werner.geyer@us.ibm.com](mailto:werner.geyer@us.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. Her website is <http://www.cs.columbia.edu/~chilton/> She can be reached at [chilton@cs.columbia.edu](mailto:chilton@cs.columbia.edu)

**Ranjitha Kumar** is an Assistant Professor in the Computer Science Department at the University of Illinois at Urbana-Champaign. She develops data-driven design techniques for creating effective user experiences, tying interface, interaction, and algorithmic design choices to user-centered goals. Her research has received best paper awards and nominations at premiere conferences in HCI, and been recognized by the machine learning community through invited papers at IJCAI and ICML. She received her PhD from the Computer Science Department at Stanford University in 2014, and was formerly the Chief Scientist at Apropose, Inc., a data-driven design company she founded that was backed by Andreessen Horowitz and New Enterprise Associates. Her website is <http://ranjithakumar.net/>, and she can be reached at [ranjitha@illinois.edu](mailto:ranjitha@illinois.edu).

**Adam Tauman Kalai** is a Principal Researcher at Microsoft Research working on machine learning and crowdsourcing. He has co-organized the conference on Crowdsourcing and Human Computation (HCOMP 2017), New England Machine Learning Day (NEML 2012-2018, with about 300 participants each), the Conference on Learning Theory (COLT 2010), and hackathons on Crowdsourcing and AI Fairness. He can be reached at [adum@microsoft.com](mailto:adum@microsoft.com).

## 3 PROGRAM COMMITTEE

- Nancy Baym, Microsoft Research
- Zoya Bylinskii, Adobe Research
- Carrie Cai, Google
- Elizabeth Clark, University of Washington
- Sebastian Gehrmann, Harvard School of Engineering
- Katy Gero, Columbia University
- Per Ola Kristensson, University of Cambridge
- Jacquelyn Martino, IBM Research AI
- Mauro Martino, IBM Research AI
- Michael Mateas, University of California, Santa Cruz
- Antti Oulasvirta, Aalto University
- Dafna Shahaf, Hebrew University of Jerusalem
- Akash Srivastava, IBM Research AI
- Hendrik Strobelt, IBM Research AI
- Michael Terry, Google

- Steven Wu, University of Minnesota
- Haiyi Zhu, Carnegie Mellon University

## 4 WORKSHOP PROGRAM

- 9:00 - 9:25 Opening & Introductions
- 9:25 - 10:10 Challenges in Building ML Algorithms for the Creative Community  
Keynote 1 by Douglas Eck, Google AI  
*Session Chair: Werner Geyer*
- 10:10 - 10:50 Session 1 - Generative Music  
*Session Chair: Lydia Chilton*  
  
[20 mins] Paper 1: Cococo: AI-Steering Tools for Music Novices Co-Creating with Generative Models Ryan Louie, Andy Coenen, Cheng-Zhi Anna Huang, Michael Terry and Carrie Cai  
  
[20 mins] Paper 2: Latent Chords: Generative Piano Chord Synthesis with Variational Autoencoders Agustin Macaya, Manuel Cartagena, Rodrigo Cadiz and Denis Parra
- 10:50 - 11:20 Coffee Break
- 11:20 - 12:30 Session 2 - Generative Text, Images, and Drawing  
*Session Chair: Adam Tauman Kalai*  
  
[20 mins] Paper 3: How Novelists Use Generative Language Models: An Exploratory User Study Alex Calderwood, Katy Ilonka Gero and Lydia B. Chilton  
  
[10 mins] Demo Paper 4: Literary Style Transfer with Content Preservation Katy Gero, Chris Kedzie and Lydia B. Chilton

### Invited Guest Speakers

- [20 mins] Invited Talk 1: Draw with Me: Human-in-the-Loop for Image Restoration Thomas Weber, Zhiwei Han, Stefan Matthes, Yuanting Liu and Heinrich Hussmann
- [20 mins] Invited Talk 2: Creative Sketching Partner: An Analysis of Human-AI Co-Creativity Pegah Karimi, Jeba Rezwana, Safat Siddiqui, Mary Lou Maher, Nasrin Dehbozorgi
- 12:30 - 14:00 Lunch
- 14:00 - 14:45 Visual Human-AI collaboration tools  
Keynote 2 by Hendrik Strobelt  
*Session Chair: Ranjitha Kumar*
- 14:45 - 15:05 Session 3 - The Dark Side of Generative Approaches

*Session Chair: Ranjitha Kumar*

[20 mins] Paper 5: Business (mis)Use Cases of Generative AI Stephanie Houde, Vera Liao, Jacquelyn Martino, Michael Muller, David Piorkowski, John Richards, Justin Weisz and Yunfeng Zhang

- 15:05 - 15:30 Discussion / Wrap-Up

## 5 SUMMARY OF WORKSHOP DISCUSSION

This workshop took place virtually on March 17 because of the global corona virus pandemic in 2020. Despite its virtual nature, the 30+ participants had very stimulating discussions after each paper and also an extended discussion at the end of the workshop. The discussion evolved around tools, ideas, goals and grand challenges that will help us, as a community make progress. The following topics evolved during this discussion:

- While generative approaches are an emerging topic in HCI communities, for many years there have been existing communities around computational creativity (including a conference series) as well as AI & Creativity events (e.g. NeurIPS workshop) at AI events. For us to bring this topic into the HCI communities, we need to make sure to establish and maintain connections into the other communities.
- During Doug Eck's keynote the term fluency came up and he was stating that while our AI systems can achieved fluency now in creating content, our systems still do not understand intent or can properly deal with the meaning of the content created. A major topic for future research is meeting the needs of humans in the creation process and providing the abilities to properly control algorithmic output.
- Another topic of interest that evolved during our discussion was the sense of ownership of the artifact created and the threat of loosing that sense of ownership. Research has shown that the sense of ownership positively influences the

creative process. However, there is a danger that with generative methods, creators will feel threatened by AI systems, which in turn might negatively influence productivity.

- The relationship between the Human and the AI system came up as another emerging topic. In particular, the question when we think about AI as a collaborator versus a tool we use. One aspect is how this relationship is formed, like we might form an anthropomorphic relationship with a musical instrument, another aspect is how the AI system is represented and how the interaction is modeled. For example, a conversational interface that is being used during the co-creation activity between Human and AI System might strongly influence how this relationship is perceived by the human during the creation activity. The perception of the AI system may have an influence on the creative output.

## REFERENCES

- [1] [n.d.]. Introduction. <https://neurips2019creativity.github.io/>
- [2] [n.d.]. MIT Quest for Intelligence. <http://ganocracy.csail.mit.edu/>
- [3] 2019. Making Art in the Age of Algorithms. <https://engineering.columbia.edu/news/art-age-algorithms>
- [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](https://www.vice.com/en_us/article/ywyxex/deepfake-of-mark-zuckerberg-facebook-fake-video-policy)
- [5] 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*.
- [6] 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>
- [7] Tero Karras, Samuli Laine, and Timo Aila. 2018. A Style-Based Generator Architecture for Generative Adversarial Networks. *ArXiv abs/1812.04948* (2018).
- [8] Diederik P. Kingma and Max Welling. 2013. Auto-Encoding Variational Bayes. *CoRR abs/1312.6114* (2013).
- [9] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. Language Models are Unsupervised Multitask Learners. (2018). <https://d4mucfpksyww.cloudfront.net/better-language-models/language-models.pdf>
- [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>