

Chris Donahue

chrisdonahue.com
chrisdonahue@cmu.edu
+1 512 825 8619

Office Address

Gates & Hillman Centers (GHC) 7127
Computer Science Department
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213

Currently

ASSISTANT PROFESSOR, Computer Science Department, Carnegie Mellon University
RESEARCH SCIENTIST (part-time), Google DeepMind and Magenta

Research interests

Machine learning for music and audio, foundation models, AI for human creativity

Education

- 2022 POSTDOC in Computer Science, Stanford University
Advised by [Percy Liang](#)
Research in foundation models for music and natural language processing.
- 2019 PHD in Music, University of California San Diego
Co-advised by [Miller Puckette](#) (music) and [Julian McAuley](#) (computer science)
Thesis: “Enabling new musical interactions with machine learning”
- 2016 MA in Music, University of California San Diego
Advised by [Miller Puckette](#)
Thesis: “Extensions to convolution for generalized cross-synthesis”
- 2014 BS in Computer Science (with high honors), The University of Texas at Austin
Co-advised by [Peter Stone](#) (computer science) and [Russell Pinkston](#) (music)
Thesis: “Applications of genetic programming to digital audio synthesis”

Selected papers

* Indicates equal contribution. For a complete list, see my [Google Scholar profile](#).

- 2023 John Thickstun, David Hall, **Chris Donahue**, and Percy Liang. “Anticipatory music transformer”, [arXiv 2023](#).
- 2023 **Chris Donahue***, Antoine Caillon*, Adam Roberts*, Ethan Manilow, Philippe Esling, Andrea Agostinelli, Mauro Vezetti, Ian Simon, Olivier Pietquin, Neil Zeghidour, and Jesse Engel. “SingSong: generating musical accompaniments from singing”, [arXiv 2023](#).
- 2022 **Chris Donahue**, John Thickstun, and Percy Liang. “Melody transcription via generative pre-training”, in [ISMIR 2022](#).
- 2022 Karan Goel, Albert Gu, **Chris Donahue**, and Christopher Ré. “It’s raw! Audio generation with state-space models”, in [ICML 2022](#) (Oral presentation).
- 2021 Rodrigo Castellon*, **Chris Donahue***, and Percy Liang. “Codified audio language modeling learns useful representations for music information retrieval”, in [ISMIR 2021](#) (**Best Paper Runner-up**).
- 2021 Rishi Bommasani, . . . , **Chris Donahue**, . . . , and Percy Liang (100+ authors). “On the opportunities and risks of foundation models”, [arXiv 2021](#).
- 2021 Hao-Wen Dong, **Chris Donahue**, Taylor Berg-Kirkpatrick, and Julian McAuley. “Towards automatic instrumentation by learning to separate parts in symbolic multitrack music”, in [ISMIR 2021](#).
- 2021 Mina Lee*, **Chris Donahue***, Robin Jia, Alexander Iyabor, and Percy Liang. “Swords: A benchmark for lexical substitution with improved data coverage and quality”, in [NAACL 2021](#).
- 2020 **Chris Donahue**, Mina Lee, and Percy Liang. “Enabling language models to fill in the blanks”, in [ACL 2020](#).
- 2019 **Chris Donahue**, Huanru Henry Mao, Yiting Ethan Li, Garrison W. Cottrell, and Julian McAuley. “LakhNES: Improving multi-instrumental music generation with cross-domain pre-training”, in [ISMIR 2019](#).
- 2019 Paarth Neekhara*, **Chris Donahue***, Miller Puckette, Shlomo Dubnov, and Julian McAuley. “Expediting TTS synthesis with adversarial vocoding”, in [INTERSPEECH 2019](#).
- 2019 **Chris Donahue**, Ian Simon, and Sander Dieleman. “Piano Genie”, in [IUI 2019](#).
- 2019

- Chris Donahue**, Julian McAuley, and Miller Puckette. “Adversarial audio synthesis”, in [ICLR 2019](#).
- 2019 Jesse Engel, Kumar Krishna Agrawal, Shuo Chen, Ishaan Gulrajani, **Chris Donahue**, and Adam Roberts. “GANSynth: Adversarial neural audio synthesis”, in [ICLR 2019](#).
- 2018 **Chris Donahue**, Huanru Henry Mao, and Julian McAuley. “The NES Music Database: A multi-instrumental dataset with expressive performance attributes”, in [ISMIR 2018](#).
- 2018 **Chris Donahue**, Zachary C. Lipton, Akshay Balsubramani, and Julian McAuley. “Semantically decomposing the latent spaces of generative adversarial networks”, in [ICLR 2018](#).
- 2018 **Chris Donahue**, Bo Li, and Rohit Prabhavalkar. “Exploring speech enhancement with generative adversarial networks for robust speech recognition”, in [ICASSP 2018](#) (Oral presentation).
- 2017 **Chris Donahue**, Zachary C. Lipton, and Julian McAuley. “Dance Dance Convolution”, in [ICML 2017](#).

Professional experience

- 2023- ASSISTANT PROFESSOR, Computer Science Department, Carnegie Mellon University
- 2023- RESEARCH SCIENTIST (part-time), Google DeepMind and Magenta
- 2022-2023 RESEARCH SCIENTIST, Google DeepMind and Magenta
Built [SingSong](#), a generative AI system which creates music to accompany user singing.
- 2020- CO-FOUNDER AND INVENTOR, Beat Sage
Created [Beat Sage](#), a free service which converts music audio into rich interactive game content. Used millions of times by thousands of daily active users.
- 2018 INTERN, Google
Built [Piano Genie](#) (IUI 2019), an intelligent instrument which allows non-musicians to improvise. Work on the [Magenta](#) team with [Ian Simon](#) and [Sander Dieleman](#).
- 2017 INTERN, Google
Explored speech enhancement as a pre-processing procedure for speech recognition (ICASSP 2018). Work with [Bo Li](#) and [Rohit Prabhavalkar](#) on the acoustic modeling research team.
- 2016 INTERN, Google
Developed techniques for semantic clustering of URLs on Google’s web crawling team.
- 2015 INTERN, Google
Trained music autotagging models and used them to make predictions on a large music

catalogue. Work with [Nicolas Boulanger-Lewandowski](#) on the Google Play Music team.

- 2011-2014 MENTOR, UT Freshman Research Initiative
Mentored for UT's Freshman Research Initiative program in the Computational Intelligence in Game Design lab under [Joel Lehman](#) and [Risto Mikkulainen](#).
- 2011-2014 INTERNSHIPS at UT Applied Research Laboratories, Qualcomm, and two startups
Various software engineering internships mostly involving full stack web development.

Honors & awards

- 2021 **Best paper runner-up at ISMIR 2021** (top 3 papers of over 200 submissions)
2019 Named one of the Best Reviewers for NeurIPS 2019
2018 **Unity Global Graduate Fellowship recipient (\$30000)**
2018 **UCSD Chancellor's Research Excellence Scholarship recipient (\$25000)**
2018 NVIDIA hardware grant recipient (\$3000 value)
2018 Support award for the 19th ISMIR Conference (\$500)
2018 Support award for the 7th ICLR Conference (\$500)
2018 HPC @ UC Allocation for machine learning research (20000 GPU hours)
2017 Sponsored Scholar award at the 34th International Conference on Machine Learning (\$1800)
2017 NVIDIA hardware grant recipient (\$1500 value)
2016 XSEDE Startup Allocation for machine learning research (6250 GPU hours)
2015 **UCSD Academic Senate COR Grant for graduate research (\$12000)**
2011 Microsoft Summer Research Fellowship (\$1250)
2011 Chevron University of Texas Computer Science Scholarship (\$2000)

Supervised students

- 2023- Irmak Bukey, PhD student, Computer Science Department, CMU
2023- Wayne Chi, PhD student, Computer Science Department, CMU
2020 Alexander Iyabor, Undergraduate, Computer Science Department, Stanford
2020 Rodrigo Castellon, Undergraduate, Computer Science Department, Stanford

Invited talks & seminars

- 2023 Oct Stanford HAI: New Horizons in Generative AI. "Music generation with precise control and composable outputs" (host: Percy Liang).
- 2023 Oct Guest lecture for CMU 11-667: Large Language Models. "LLMs beyond text: Music" (host: Daphne Ippolito).
- 2023 Sep Guest lecture for CMU 07-300: Research Overview. "Unlocking musical expression with generative models" (host: Ruben Martins).
- 2023 Apr

- Guest lecture for CMU Music Technology Course. “Demystifying music generative modeling: from Markov chains to AI Drake” (host: Annie Hui-Hsin Hsieh).
- 2023 Apr CMU Music & Technology Seminar. “Frontiers in controllable music generation” (host: Riccardo Schulz).
- 2023 Feb AAAI Workshop on Creative AI Generation. “Frontiers in controllable music generation” (host: Haw-Shiuan Chang).
- 2022 Jul ICML Workshop on Machine Learning for Audio Synthesis. “Frontiers and challenges in music audio generation” (host: Sander Dieleman).
- 2022 Feb ACMI Lab. “Unlocking musical expression with machine learning” (host: Zachary C. Lipton).
- 2019 Dec ASA San Diego. “Unlocking musical expression with machine learning” (host: Scott Hawley).
- 2019 Aug Bish Bash (Dolby). “Neural Loops: A factorized generative model for musical loops ” (host: Jordi Pons).
- 2019 Feb P-lambda Seminar. “Machine learning methods for enriching musical interaction” (host: Percy Liang).
- 2019 Feb Stanford DESINST 240—Designing Machine Learning. “Pairing human control with generative models for creative content synthesis” (host: Abhay Agarwal).
- 2019 Jan Georgia Tech music tech seminar. “Music generation with language models” (host: Jason Freeman).
- 2018 Oct Unity Unite Conference. “Low- to high-level learning problems in game audio” (host: Diana Ford).

Service

COMMITTEE MEMBER

- Area chair, NAACL: Annual Conference of the North American Chapter of the Association for Computational Linguistics
- Metareviewer, ISMIR: International Society for Music Information Retrieval Conference

CONFERENCE REVIEWS

- (ISMIR) International Society for Music Information Retrieval Conference
- (ICLR) International Conference on Learning Representations
- (ICML) International Conference on Machine Learning
- (NeurIPS) Conference on Neural Information Processing Systems

- (CHI) ACM Conference on Human Factors in Computing Systems
- (ICASSP) IEEE International Conference on Acoustics, Speech, & Signal Processing

JOURNAL REVIEWS

- Transactions of the International Society for Music Information Retrieval
- ACM Computing Survey
- IEEE Transactions on Signal Processing
- Journal of Selected Topics in Signal Processing
- IET Computer Vision

OTHER

- Co-organizer, ISMIR 2020 Tutorial, “Designing generative models for interactive co-creation”
- Mentor, Women in Music Information Retrieval Mentorship Program, 5 years
- PhD Thesis Committee, Yinghao Ma, Queen Mary University of London, PhD Center for Digital Music
- MS Thesis Committee, Alexander Wang, Carnegie Mellon University, MS Music & Technology

Media coverage

Coverage of [Piano Genie](#), an intelligent instrument which enables amateur improvisation, and **[Fruit Genie](#)**, a live performance involving Piano Genie and The Flaming Lips.

[BUSINESS INSIDER](#) [A Google intern helped build an AI tool inspired by 'Guitar Hero' to let rookies play piano](#)

[THE VERGE](#) [Google's AI-powered Piano Genie lets anyone improvise perfectly by bashing buttons](#)

[ENGADGET](#) [Google's Piano Genie lets anyone improvise classical music](#)

[EVENING STANDARD](#) [Piano Genie: Google's AI programme is like Guitar Hero for the piano world](#)

[STEREOGUM](#) [Watch The Flaming Lips Play A Bowl Of Fruit At Google I/O](#)

Coverage of [Dance Dance Convolution](#), a system for converting audio into interactive game content, and **[Beat Sage](#)**, a free service based on this work.

[MIT TECH REVIEW](#) [Machine-Learning Algorithm Watches Dance Dance Revolution, Then Creates Dances of Its Own](#)

[THE VERGE](#) [Scientists have taught a neural network to choreograph Dance Dance Revolution levels](#)

[VICE](#) [This Machine Learned to Choreograph by Watching Dance Dance Revolution](#)

THE REGISTER [Yet another job menaced by AI! Uh, wait, it says here . . . Dance Dance Revolution designers](#)

UPLOADVR [New AI Tool Turns Any Song Into A Custom Beat Saber Map, And It Really Works](#)

ROAD TO VR [This 'Beat Saber' Project Uses AI to Generate Custom Beat Maps for Any Song](#)