Jay Whang

https://jaywhang.com jwhang 'at' google 'dot' com

Education

University of Texas at Austin, Austin, TX 2019 – 2023
- Ph.D. in Computer Science (advised by Prof. Alex Dimakis)

Stanford University, Stanford, CA 2017 – 2019
- M.S. in Computer Science

University of Southern California, Los Angeles, CA 2010 – 2014
- B.S. in Computer Science, B.A. in Mathematics

Research

I am a senior research scientist at Google DeepMind, working on Gemini. My research interests lie broadly in **deep generative modeling** with the goal of enabling it to work well at scale. In particular, I'm interested in using **likelihood-based models** for useful downstream tasks such as **image** and **video generation**.

Conference Papers

[1] Imagen: Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. (**NeurIPS 2022 Outstanding Paper**) [pdf] [project page]
Chitwan Saharia*, William Chan*, Saurabh Saxena†, Lala Li†, **Jay Whang**†, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho†, David J Fleet†, Mohammad Norouzi*.

[2] Deblurring via Stochastic Refinement. (**CVPR 2022 Oral Presentation**) [pdf]
Jay Whang, Mauricio Delbracio, Hossein Talebi, Chitwan Saharia, Alex Dimakis, Peyman Milanfar.

[3] Neural Distributed Source Coding. (**JSAIT 2024**) [pdf]
Jay Whang*, Alliot Nagle*, Anish Acharya, Hyeji Kim, Alexandros Dimakis.

[4] Composing Normalizing Flows for Inverse Problems. (**ICML 2021**) [pdf]
Jay Whang, Erik Lindgren, Alexandros Dimakis.

[5] Solving Inverse Problems with a Flow-based Noise Model. (**ICML 2021**) [pdf]
Jay Whang, Qi Lei, Alexandros Dimakis.

[6] Training Variational Autoencoders with Buffered SVI. (**AISTATS 2019**) [pdf]
Rui Shu, Hung Bui, **Jay Whang**, Stefano Ermon.

Preprints & Workshop Papers

[7] Imagen Video: High Definition Video Generation with Diffusion Models. [pdf] [project page]
Jonathan Ho*, William Chan*, Chitwan Saharia*, **Jay Whang***, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, Tim Salimans*

[8] Model-based Deep Learning: Key Approaches And Design Guidelines. [pdf]
Nir Shlezinger, **Jay Whang**, Yonina Eldar, Alex Dimakis.
– ***Audience Choice Award*** *at IEEE Data Science Learning Workshop (DSLW 2021)*

[9] Approximate Probabilistic Inference with Composed Flows [pdf]
Jay Whang, Erik Lindgren, Alexandros Dimakis.
– ***Best Paper Award*** *at UAI 2021 Workshop on Tractable Probabilistic Modeling.*
– *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems.*

[10] Compressed Sensing with Invertible Generative Models and Dependent Noise [pdf]
Jay Whang, Qi Lei, Alexandros G. Dimakis.
– *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems.*

[11] Fast Exploration with Simplified Models and Approximately Optimistic Planning in Model Based Reinforcement Learning. (*Preprint*) [pdf]
Ramtin Keramati*, **Jay Whang***, Patrick Cho* and Emma Brunskill.
– *ICML 2018 Workshop on Exploration in Reinforcement Learning.* [pdf]

Work Experience

Google Brain (Toronto) – *Student Researcher* Jan. 2022 – Jan. 2023
- Co-authored Imagen Video, Google's state-of-the-art text-to-video diffusion model.
- Core contributor on Imagen, Google's latest text-to-image model (published at NeurIPS 2022).

Google Research (Luma) – *Research Intern* Summer - Fall 2021
- Leverage diffusion model for stochastic blind image deblurring (published at CVPR 2022).

DeepMind, Mountain View, CA – *Research Engineer Intern* Summer 2018
- Investigated ways to improve sampling and training speed of WaveNet with progressive training.

YouTube, Mountain View, CA – *Software Engineer* Dec. 2014 – July 2017
- Trained and productionized various classifiers for detecting abusive videos and users.
- Wrote a real-time data processing backend pipeline for aggregating user activities on YouTube.

Facebook, Menlo Park, CA – *Software Engineer Intern* Summer 2014
- Designed and implemented a physics-based layout engine for contextual dialog boxes in JavaScript.

Microsoft, Redmond, WA – *Software Development Engineer (SDE) Intern* Summer 2012
- Created a web UI for remote configuration and deployment of Windows 8 on bare metal machines.

Microsoft, Redmond, WA – *Software Development Engineer in Test (SDET) Intern* Summer 2011
- Designed and implemented functional and stress tests for a cluster manager on Windows HPC.

Teaching

Stanford University
- CS 234: Reinforcement Learning by Prof. Emma Brunskill Winter 2018
- CS 230: Deep Learning by Prof. Andrew Ng Spring & Fall 2018, Spring 2019
- CS 224N: NLP with Deep Learning by Prof. Richard Socher Winter 2017
- CS 148: Computer Graphics by Prof. Ron Fedkiw Fall 2017

University of Southern California
- CSCI 103: Introduction to Programming by Prof. Mark Redekopp Fall 2013
- CSCI 271: Discrete Mathematics by Prof. David Kempe Spring 2013

Service

Reviewer for ICLR, CVPR, ICML, NeurIPS, MLSys, and JSAIT.

Awards

Member of Phi Beta Kappa National Honor Society 2013 – Present
Three-time USA Mathematics Olympiad (USAMO) qualifier 2007 – 2009
Mathematical Olympiad Summer Program (MOSP) participant 2007

Skills

ML Frameworks: Proficient in Jax, PyTorch, TensorFlow, and other relevant scientific packages.
Languages: Proficient in Python and C++.
Spoken Languages: English (fluent), Korean (native).

Last updated: January 1, 2025