

# Jay Whang

<https://jaywhang.com>

jaywhang 'at' utexas.edu

---

EDUCATION	<b>University of Texas at Austin</b> , Austin, TX – Ph.D. in Computer Science <b>Stanford University</b> , Stanford, CA – M.S. in Computer Science <b>University of Southern California</b> , Los Angeles, CA – B.S. in Computer Science, B.A. in Mathematics	2019 – 2023 (expected)  2017 – 2019  2010 – 2014
RESEARCH INTERESTS	My research interests lie broadly in <b>deep generative modeling</b> with the goal of enabling it to work well in practice. In particular, I'm interested in using <b>likelihood-based models</b> to perform useful downstream tasks such as <b>image/video generation</b> , <b>inverse problems</b> and <b>compression</b> .	
CONFERENCE PAPERS	<ul style="list-style-type: none"><li>[1] Imagen: Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. (<b>NeurIPS 2022 Outstanding Paper</b>) <a href="#">[pdf]</a> <a href="#">[project page]</a> Chitwan Saharia*, William Chan*, Saurabh Saxena†, Lala Li†, <b>Jay Whang†</b>, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho†, David J Fleet†, Mohammad Norouzi*.</li><li>[2] Deblurring via Stochastic Refinement. (<b>CVPR 2022 Oral Presentation</b>) <a href="#">[pdf]</a> <b>Jay Whang</b>, Mauricio Delbracio, Hossein Talebi, Chitwan Saharia, Alex Dimakis, Peyman Milanfar.</li><li>[3] Composing Normalizing Flows for Inverse Problems. (<b>ICML 2021</b>) <a href="#">[pdf]</a> <b>Jay Whang</b>, Erik Lindgren, Alexandros Dimakis.</li><li>[4] Solving Inverse Problems with a Flow-based Noise Model. (<b>ICML 2021</b>) <a href="#">[pdf]</a> <b>Jay Whang</b>, Qi Lei, Alexandros Dimakis.</li><li>[5] Training Variational Autoencoders with Buffered SVI. (<b>AISTATS 2019</b>) <a href="#">[pdf]</a> Rui Shu, Hung Bui, <b>Jay Whang</b>, Stefano Ermon.</li></ul>	
PREPRINTS & WORKSHOP PAPERS	<ul style="list-style-type: none"><li>[6] Imagen Video: High Definition Video Generation with Diffusion Models. <a href="#">[pdf]</a> <a href="#">[project page]</a> Jonathan Ho*, William Chan*, Chitwan Saharia*, <b>Jay Whang*</b>, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, Tim Salimans*</li><li>[7] Neural Distributed Source Coding. <a href="#">[pdf]</a> <b>Jay Whang</b>, Anish Acharya, Hyeji Kim, Alexandros Dimakis.</li><li>[8] Model-based Deep Learning: Key Approaches And Design Guidelines. <a href="#">[pdf]</a> Nir Shlezinger, <b>Jay Whang</b>, Yonina Eldar, Alex Dimakis. – <b>Audience Choice Award</b> at <i>IEEE Data Science Learning Workshop (DSLW 2021)</i></li><li>[9] Approximate Probabilistic Inference with Composed Flows <a href="#">[pdf]</a> <b>Jay Whang</b>, Erik Lindgren, Alexandros Dimakis. – <b>Best Paper Award</b> at <i>UAI 2021 Workshop on Tractable Probabilistic Modeling</i>. – <i>NeurIPS 2020 Workshop on Deep Learning and Inverse Problems</i>.</li><li>[10] Compressed Sensing with Invertible Generative Models and Dependent Noise <a href="#">[pdf]</a> <b>Jay Whang</b>, Qi Lei, Alexandros G. Dimakis. – <i>NeurIPS 2020 Workshop on Deep Learning and Inverse Problems</i>.</li><li>[11] Fast Exploration with Simplified Models and Approximately Optimistic Planning in Model Based Reinforcement Learning. (<i>Preprint</i>) <a href="#">[pdf]</a> Ramtin Keramati*, <b>Jay Whang*</b>, Patrick Cho* and Emma Brunskill. – <i>ICML 2018 Workshop on Exploration in Reinforcement Learning</i>. <a href="#">[pdf]</a></li></ul>	

WORK EXPERIENCE	<b>Google Brain (Toronto)</b> – <i>Student Researcher</i>	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).	
	<b>Google Research (Luma)</b> – <i>Research Intern</i>	Summer - Fall 2021
	– Leverage diffusion model for stochastic blind image deblurring (published at CVPR 2022).	
	<b>DeepMind</b> , Mountain View, CA – <i>Research Engineer Intern</i>	Summer 2018
	– Investigated ways to improve sampling and training speed of WaveNet with progressive training.	
	<b>YouTube</b> , Mountain View, CA – <i>Software Engineer</i>	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.	
TEACHING	<b>Facebook</b> , Menlo Park, CA – <i>Software Engineer Intern</i>	Summer 2014
	– Designed and implemented a physics-based layout engine for contextual dialog boxes in JavaScript.	
	<b>Microsoft</b> , Redmond, WA – <i>Software Development Engineer (SDE) Intern</i>	Summer 2012
	– Created a web UI for remote configuration and deployment of Windows 8 on bare metal machines.	
	<b>Microsoft</b> , Redmond, WA – <i>Software Development Engineer in Test (SDET) Intern</i>	Summer 2011
	– Designed and implemented functional and stress tests for a cluster manager on Windows HPC.	
	<b>Stanford University</b>	
	– 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
SERVICE	– CS 148: Computer Graphics by Prof. Ron Fedkiw	Fall 2017
	<b>University of Southern California</b>	
	– CSCI 103: Introduction to Programming by Prof. Mark Redekopp	Fall 2013
	– CSCI 271: Discrete Mathematics by Prof. David Kempe	Spring 2013
	Reviewer for ICLR 2022-2023, CVPR 2023, ICML 2020-2022, NeurIPS 2020-2022, MLSys 2021	
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	<b>ML Frameworks:</b> Proficient in Jax, PyTorch, TensorFlow, and other scientific packages.	
	<b>Languages:</b> Proficient in Python and C++ (over 30k lines of production code at Google).	
	<b>Spoken Languages:</b> English (fluent), Korean (native).	