Last updated: April 17, 2026

Jiatao Gu

Assistant Professor | Ph.D.

Homepage: https://jiataogu.me/ | Phone: +1 (650) 885-1592 | Email: jgu32@cis.upenn.edu, jgu32@apple.com
Social: Google Scholar | LinkedIn | GitHub | X

Research Interests

My research goal is to advance the capabilities of AI agents to *perceive, reason about, and act within* the physical world by harnessing *Generative ML* approaches for *robust world modeling*, *rigorous reasoning*, and *effective decision-making*.

Education

Ph.D in Electrical and Electronic Engineering — Aug 2014 – Sep 2018
The University of Hong Kong — Hong Kong, China
Supervisor: *Prof. Victor O.K. Li*
Dissertation: *Efficient Neural Machine Translation*

Bachelor in Electronic Engineering — Sep 2010 – Jun 2014
Tsinghua University — Beijing, China

Exchange Program in Information Science — Oct 2012 – Mar 2013
The University of Tokyo — Tokyo, Japan
Advisers: *Prof. Hitoshi Iba*, *Prof. Danushka Bollegala*

Appointments

Assistant Professor — July 2025 – Present
Computer and Information Science, University of Pennsylvania — Philadelphia, PA, US

Staff Research Scientist — Jun 2022 – Present
Machine Learning Research (MLR), Apple Inc. — New York City, NY, US

Senior Research Scientist — Aug 2018 – Jun 2022
Facebook AI Research (FAIR), Meta Platforms Inc. — New York City, NY, US

Research Intern — Sep 2017 – Dec 2017
Salesforce Research, Salesforce Inc. — Palo Alto, CA, US
Mentor: *Dr. Richard Socher*

Research Intern — Apr 2017 – Aug 2017
Microsoft Research, Microsoft Corporation — Redmond, WA, US
Mentor: *Dr. Hany Hassan*

Visiting Scholar — Jun 2016 – Jan 2017
Courant Institute of Mathematical Science, New York University — New York City, NY, US
Adviser: *Prof. Kyunghyun Cho*

Grants & Funding

MURI: NEON: Neuro-symbolic Evaluation and Observation of Neural Networks — 2026-2029
ONR · $7,500,000 · Co-PI — (Awarded)

Modeling the 4D World: Dataset, Evaluation, and Understanding — 2026-2027
Apple · $150,000 · PI — (Awarded)

Real-Time, Long-Context, Physical World Models through Data-Free Learning — 2026
NVIDIA · 32K A100 GPU-Hours · PI — (Awarded)

AWS ASSET Fellow — 2026-2027
Amazon · A calendar year — (Awarded)

Professional Activities

Professional Service

Action Editor: ACL Rolling Review (ARR), since 2022

Area Chair: 3DV 2026, COLM 2026, 2025, NeurIPS 2026, 2025, ICLR 2026, 2025, 2024, ICML 2026, 2024, ACL 2021, EMNLP 2020, 2021

Reviewer (Conference): NeurIPS, ICML, ICLR, ACL, EMNLP, NAACL, IJCNLP, CVPR, ECCV, SIGGRAPH, SIGGRAPH Asia, AAAI, IJCAI

Reviewer (Journal): TMLR, TACL

Invited Talks

Scalable Normalizing Flows for Visual Generation

UNC Chapel Hill, Invited Lecture — April 12
University of Minnesota, Industrial Problems Seminar — March 27
CVPR 2025 Tutorial on Scalable Generative Models in Computer Vision, invited talk — Jun 11, 2025
CVPR 2025 Workshop on Visual Generative Modeling, invited talk — Jun 12, 2025

End-to-End Generative Video Modeling with Scalable Normalizing Flows

ICCV 2025 Workshop on GeoFreeNVS: Geometry-Free Novel View Synthesis and Controllable Video Models
October 19

Learning Reliable and Efficient Generative Intelligence

CVPR 2025 Workshop on Generative Models for Computer Vision, invited talk — Jun 11, 2025
Nanyang Technological University, invited talk — Apr 26, 2025
University of Pennsylvania, ASSET Seminar talk — Apr 16, 2025

Towards Flexible, Scalable and Knowledgeable Generative Intelligence

New York University, guest lecture — Apr 16, 2024
University of Pennsylvania, guest lecture — Oct 16, 2024

BOOT: data-free distillation of denoising diffusion models with bootstrapping

ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling — Jul 28, 2023

Towards Efficient Diffusion Models for 3D Generation

BIRS Workshop, 3D Generative Models, virtual talk — Jul 14, 2023

Non-autoregressive Sequence Generation

ACL 2022 Tutorial on Non-autoregressive Sequence Generation — May 22, 2022

StyleNeRF: A style-based 3d-aware generator for high-resolution image synthesis

VALSE Webinar, virtual talk (Chinese) — Dec 29, 2021

Multilingual Denoising Pre-training for Neural Machine Translation

Google Translate, virtual talk — Jul 17, 2020
ACL2020 WNGT Workshop, virtual talk — Jul 10, 2020

Understanding Knowledge Distillation in Neural Sequence Generation

Tsinghua University, Beijing, China — Jan 03, 2020
Microsoft Translate, Redmond, WA, USA — Dec 04, 2019

Non-autoregressive Neural Machine Translation

New York University, New York, NY, USA — Dec 19, 2017
Google Brain, Mountain View, CA, USA — Nov 20, 2017

Learning to Share: towards Zero-Resource Neural Machine Translation without Pivots

Microsoft Research, Redmond, WA, USA — Jul 27, 2017

Trainable Decoding for Neural Machine Translation

NYU Shanghai, Shanghai, China — Apr 14, 2017

Learning to Translate in Real-time with Neural Machine Translation

Carnegie Mellon University, Pittsburgh, PA, USA — Oct 27, 2016
New York University, New York, NY, USA — Oct 19, 2016

Teaching Experience

Instructor

Advanced Topics in Deep Generative Models, CIS 7000, UPenn — Fall 2025

Teaching Assistant

Human Computer Interaction, ELEC 3641, HKU — Spring 2018

Embedded System, ELEC 3226/3442, HKU — Spring 2016

S/W Engineering and O.S., ELEC 2501/3541, HKU — Spring 2015

Selected Media Coverage

VentureBeat and 9to5mac reported **STARFlow**

Marktechpost reported **Matryoshka Diffusion**

MIT Technology Review reported **Data2vec**

SyncedReview reported **StyleNeRF**

SyncedReview reported **mBART**

ZDNET reported **Levenshtein Transformer**

CNBC and VentureBeat reported on **Non-Autoregressive Neural Machine Translation**

Slator reported on **Learning to Translate in Real-time with Neural Machine Translation**

Publications

*=equal contribution

Conference Papers

[C1] **Jiatao Gu**, Ying Shen, Tianrong Chen, Laurent Dinh, Yuyang Wang, Miguel Angel Bautista, David Berthelot, Joshua Susskind, and Shuangfei Zhai. “STARFlow-V: End-to-End Video Generative Modeling with Normalizing Flows”. In: *CVPR 2026*. arXiv:2511.20462. 2026.

[C2] Yongyuan Liang, Wei Chow, Feng Li, Ziqiao Ma, Xiyao Wang, Jiageng Mao, Jiuhai Chen, **Jiatao Gu**, Yue Wang, and Furong Huang. “ROVER: Benchmarking Reciprocal Cross-Modal Reasoning for Omnimodal Generation”. In: *ICLR 2026*. arXiv:2511.01163. 2026.

[C3] Huangjie Zheng, Shansan Gong, Ruixiang Zhang, Tianrong Chen, **Jiatao Gu**, Mingyuan Zhou, Navdeep Jaitly, and Yizhe Zhang. “Continuously Augmented Discrete Diffusion Model for Categorical Generative Modeling”. In: *ICLR 2026*. arXiv:2510.01329. 2026.

[C4] **Jiatao Gu**, Tianrong Chen, David Berthelot, Huangjie Zheng, Yuyang Wang, Ruixiang Zhang, Laurent Dinh, Miguel Angel Bautista, Joshua Susskind, and Shuangfei Zhai. “STARFlow: Scaling Latent Normalizing Flows for High-resolution Image Synthesis”. In: *NeurIPS 2025*. arXiv:2506.06276. 2025.

[C5] Chen Wang, Chuhao Chen, Yiming Huang, Zhiyang Dou, Yuan Liu, **Jiatao Gu**, and Lingjie Liu. “PhysCtrl: Generative Physics for Controllable and Physics-Grounded Video Generation”. In: *NeurIPS 2025*. arXiv:2509.20358. 2025.

[C6] Ruixiang Zhang, Shuangfei Zhai, **Jiatao Gu**, Yizhe Zhang, Huangjie Zheng, Tianrong Chen, Miguel Angel Bautista, Joshua Susskind, and Navdeep Jaitly. “Flexible Language Modeling in Continuous Space with Transformer-based Autoregressive Flows”. In: *NeurIPS 2025*. arXiv:2507.00425. 2025.

[C7] Tianrong Chen, Huangjie Zheng, David Berthelot, **Jiatao Gu**, Joshua Susskind, and Shuangfei Zhai. “TADA: Improved Diffusion Sampling with Training-free Augmented Dynamics”. In: *NeurIPS 2025*. arXiv:2506.21757. 2025.

[C8] Shuangfei Zhai, Ruixiang Zhang, Preetum Nakkiran, David Berthelot, **Jiatao Gu**, Huangjie Zheng, Tianrong Chen, Miguel Angel Bautista, Navdeep Jaitly, and Joshua Susskind. “Normalizing Flows are Capable Generative Models”. In: *ICML 2025*. arXiv:2412.06329. 2025.

[C9] Qihang Zhang, Shuangfei Zhai, Miguel Angel Bautista, Kevin Miao, Alexander Toshev, Joshua Susskind, and **Jiatao Gu**. “World-consistent Video Diffusion with Explicit 3D Modeling”. In: *CVPR 2025*. arXiv:2412.01821. 2025.

[C10] Chuhao Chen, Zhiyang Dou, Chen Wang, Yiming Huang, Anjun Chen, Qiao Feng, **Jiatao Gu**, and Lingjie Liu. “Vid2Sim: Generalizable, Video-based Reconstruction of Appearance, Geometry and Physics for Mesh-free Simulation”. In: *CVPR 2025*. arXiv:2506.06440. 2025.

[C11] **Jiatao Gu**, Yuyang Wang, Yizhe Zhang, Qihang Zhang, Dinghuai Zhang, Navdeep Jaitly, Joshua Susskind, and Shuangfei Zhai. “DART: Denoising autoregressive transformer for scalable text-to-image generation”. In: *ICLR 2025*. arXiv:2410.08159. 2025.

[C12] **Jiatao Gu**, Ying Shen, Shuangfei Zhai, Yizhe Zhang, Navdeep Jaitly, and Joshua Susskind. “Kaleido Diffusion: Improving Conditional Diffusion Models with Autoregressive Latent Modeling”. In: *NeurIPS 2024*. arXiv:2405.21048. 2024.

[C13] Yizhe Zhang, He Bai, Ruixiang Zhang, **Jiatao Gu**, Shuangfei Zhai, Joshua Susskind, and Navdeep Jaitly. “How Far Are We from Intelligent Visual Deductive Reasoning?” In: *COLM 2024*. arXiv:2403.04732. 2024.

[C14] Zhuofeng Wu, He Bai, Aonan Zhang, **Jiatao Gu**, VG Vydiswaran, Navdeep Jaitly, and Yizhe Zhang. “Divide-or-Conquer? Which Part Should You Distill Your LLM?” In: *EMNLP 2024*. arXiv:2402.15000. 2024.

[C15] Jing Nathan Yan, **Jiatao Gu**, and Alexander M. Rush. “Diffusion Models Without Attention”. In: *CVPR 2024*. arXiv:2311.18257. 2024.

[C16] **Jiatao Gu**, Shuangfei Zhai, Yizhe Zhang, Joshua Susskind, and Navdeep Jaitly. “Matryoshka Diffusion Models”. In: *ICLR 2024*. arXiv:2310.15111. 2024.

[C17] Tianrong Chen, **Jiatao Gu**, Laurent Dinh, Evangelos A. Theodorou, Joshua Susskind, and Shuangfei Zhai. “Generative Modeling with Phase Stochastic Bridges”. In: *ICLR 2024*. arXiv:2310.07805. 2024.

[C18] **Jiatao Gu**, Shuangfei Zhai, Yizhe Zhang, Lingjie Liu, and Joshua Susskind. “BOOT: Data-free Distillation of Denoising Diffusion Models with Bootstrapping”. In: *ICML 2024*. arXiv:2306.05544. 2024.

[C19] **Jiatao Gu**, Qingzhe Gao, Shuangfei Zhai, Baoquan Chen, Lingjie Liu, and Joshua Susskind. “Control3Diff: Learning Controllable 3D Diffusion Models from Single-view Images”. In: *3DV 2024*. arXiv:2304.06700. 2024.

[C20] Yizhe Zhang, **Jiatao Gu**, Zhuofeng Wu, Shuangfei Zhai, Joshua Susskind, and Navdeep Jaitly. “PLANNER: Generating Diversified Paragraph via Latent Language Diffusion Model”. In: *NeurIPS 2023*. arXiv:2306.02531. 2023.

[C21] Hansheng Chen, **Jiatao Gu**, Anpei Chen, Wei Tian, Zhuowen Tu, Lingjie Liu, and Hao Su. “Single-Stage Diffusion NeRF: A Unified Approach to 3D Generation and Reconstruction”. In: *ICCV 2023*. arXiv:2304.06714. 2023.

[C22] **Jiatao Gu**, Alex Trevithick, Kai-En Lin, Joshua Susskind, Christian Theobalt, Lingjie Liu, and Ravi Ramamoorthi. “Nerfdiff: Single-image view synthesis with nerf-guided distillation from 3d-aware diffusion”. In: *ICML 2023*. arXiv:2302.10109. 2023.

[C23] Shuangfei Zhai, Tatiana Likhomanenko, Etai Littwin, Dan Busbridge, Jason Ramapuram, Yizhe Zhang, **Jiatao Gu**, and Joshua Susskind. “Stabilizing Transformer Training by Preventing Attention Entropy Collapse”. In: *ICML 2023*. arXiv:2303.06296. 2023.

[C24] **Jiatao Gu**, Shuangfei Zhai, Yizhe Zhang, Miguel Angel Bautista, and Joshua Susskind. “f-DM: A Multi-stage Diffusion Model via Progressive Signal Transformation”. In: *ICLR 2023*. arXiv:2210.04955. 2023.

[C25] Peiye Zhuang, Samira Abnar, **Jiatao Gu**, Alex Schwing, Joshua Susskind, and Miguel Angel Bautista. “Diffusion probabilistic fields”. In: *ICLR 2023*. arXiv:2303.00165. 2023.

[C26] Chen Huang, Hanlin Goh, **Jiatao Gu**, and Joshua Susskind. “MAST: Masked Augmentation Subspace Training for Generalizable Self-Supervised Priorse”. In: *ICLR 2023*. arXiv:2303.03679. 2023.

[C27] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, **Jiatao Gu**, and Michael Auli. “Data2vec: A general framework for self-supervised learning in speech, vision and language”. In: *ICML 2022*. arXiv:2202.03555. 2022.

[C28] **Jiatao Gu**, Lingjie Liu, Peng Wang, and Christian Theobalt. “Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis”. In: *ICLR 2022*. arXiv:2110.08985. 2022.

[C29] Ann Lee, Peng-Jen Chen, Changhan Wang, **Jiatao Gu**, Sravya Popuri, Xutai Ma, Adam Polyak, Yossi Adi, Qing He, Yun Tang, Juan Pino, and Wei-Ning Hsu. “Direct Speech-to-Speech Translation With Discrete Units”. In: *ACL 2022*. arXiv:2107.05604. 2022.

[C30] Yun Tang, Hongyu Gong, Ning Dong, Changhan Wang, Wei-Ning Hsu, **Jiatao Gu**, Alexei Baevski, Xian Li, Abdelrahman Mohamed, Michael Auli, and Juan Pino. “Unified Speech-Text Pre-training for Speech Translation and Recognition”. In: *ACL 2022*. arXiv:2204.05409. 2022.

[C31] Ann Lee, Hongyu Gong, Paul-Ambroise Duquenne, Holger Schwenk, Peng-Jen Chen, Changhan Wang, Sravya Popuri, Juan Pino, **Jiatao Gu**, and Wei-Ning Hsu. “Textless Speech-to-Speech Translation on Real Data”. In: *NAACL 2022*. arXiv:2107.05604. 2022.

[C32] Zhuofeng Wu, Sinong Wang, **Jiatao Gu**, Rui Hou, Yuxiao Dong, VG Vydiswaran, and Hao Ma. “IDPG: An instance-dependent prompt generation method”. In: *NAACL 2022*. arXiv:2204.04497. 2022.

[C33] **Jiatao Gu** and Xiang Kong. “Fully Non-autoregressive Neural Machine Translation: Tricks of the Trade”. In: *ACL 2021*. arXiv:2012.15833. 2021.

[C34] Chunting Zhou, Graham Neubig, **Jiatao Gu**, Mona Diab, Paco Guzman, Luke Zettlemoyer, and Marjan Ghazvininejad. “Detecting hallucinated content in conditional neural sequence generation”. In: *ACL 2021*. arXiv:2011.02593. 2021.

[C35] Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, **Jiatao Gu**, and Angela Fan. “Multilingual translation from denoising pre-training”. In: *ACL 2021*. arXiv:2008.00401. 2021.

[C36] Hang Le, Juan Pino, Changhan Wang, **Jiatao Gu**, Didier Schwab, and Laurent Besacier. “Lightweight Adapter Tuning for Multilingual Speech Translation”. In: *ACL 2021*. arXiv:2106.01463. 2021.

[C37] Jungo Kasai, James Cross, Marjan Ghazvininejad, and **Jiatao Gu**. “Non-autoregressive machine translation with disentangled context transformer”. In: *ICML 2021*. arXiv:2001.05136. 2021.

[C38] Lior Yariv, **Jiatao Gu**, Yoni Kasten, and Yaron Lipman. “Volume Rendering of Neural Implicit Surfaces”. In: *NeurIPS 2021*. arXiv:2106.12052. 2021.

[C39] Jiajun Shen, Peng-Jen Chen, Matthew Le, Junxian He, **Jiatao Gu**, Myle Ott, Michael Auli, and Marc’Aurelio Ranzato. “The source-target domain mismatch problem in machine translation”. In: *EACL 2021*. arXiv:1909.13151. 2021.

[C40] Xiang Kong, Adithya Renduchintala, James Cross, Yuqing Tang, **Jiatao Gu**, and Xian Li. “Multilingual neural machine translation with deep encoder and multiple shallow decoders”. In: *EACL 2021*. 2021.

[C41] Lingjie Liu, **Jiatao Gu**, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. “Neural sparse voxel fields”. In: *NeurIPS 2020*. arXiv:2007.11571. 2020.

[C42] Chau Tran, Yuqing Tang, Xian Li, and **Jiatao Gu**. “Cross-lingual retrieval for iterative self-supervised training”. In: *NeurIPS 2020*. arXiv:2006.09526. 2020.

[C43] Saining Xie, **Jiatao Gu**, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. “PointContrast: Unsupervised pre-training for 3d point cloud understanding”. In: *ECCV 2020*. arXiv:2007.10985. 2020.

[C44] Chunting Zhou, Graham Neubig, and **Jiatao Gu**. “Understanding knowledge distillation in non-autoregressive machine translation”. In: *ICLR 2020*. arXiv:1911.02727. 2020.

[C45] Xutai Ma, Juan Pino, James Cross, Liezl Puzon, and **Jiatao Gu**. “Monotonic multihead attention”. In: *ICLR 2020*. arXiv:1909.12406. 2020.

[C46] Junxian He, **Jiatao Gu**, Jiajun Shen, and Marc’Aurelio Ranzato. “Revisiting self-training for neural sequence generation”. In: *ICLR 2020*. arXiv:1909.13788. 2020.

[C47] Maha Elbayad, **Jiatao Gu**, Edouard Grave, and Michael Auli. “Depth-adaptive transformer”. In: *ICLR 2020*. arXiv:1910.10073. 2020.

[C48] Changhan Wang, Kyunghyun Cho, and **Jiatao Gu**. “Neural machine translation with byte-level subwords”. In: *AAAI 2020*. arXiv:1909.03341. 2020.

[C49] Changhan Wang, Juan Pino, and **Jiatao Gu**. “Improving Cross-Lingual Transfer Learning for End-to-End Speech Recognition with Speech Translation”. In: *Interspeech 2020*. arXiv:2006.05474. 2020.

[C50] Anne Wu, Changhan Wang, Juan Pino, and **Jiatao Gu**. “Self-supervised representations improve end-to-end speech translation”. In: *Interspeech 2020*. arXiv:2006.12124. 2020.

[C51] Arya D McCarthy, Xian Li, **Jiatao Gu**, and Ning Dong. “Addressing Posterior Collapse with Mutual Information for Improved Variational Neural Machine Translation”. In: *ACL 2020*. arXiv:1909.09237. 2020.

[C52] Hang Le, Juan Pino, Changhan Wang, **Jiatao Gu**, Didier Schwab, and Laurent Besacier. “Dual-decoder transformer for joint automatic speech recognition and multilingual speech translation”. In: *COLING 2020*. arXiv:2011.00747. 2020.

[C53] **Jiatao Gu**, Changhan Wang, and Jake Zhao. “Levenshtein Transformer”. In: *NeurIPS 2019*. arXiv:1905.11006. 2019.

[C54] **Jiatao Gu**, Yong Wang, Kyunghyun Cho, and Victor OK Li. “Improved Zero-shot Neural Machine Translation via Ignoring Spurious Correlations”. In: *ACL 2019*. arXiv:1906.01181. 2019.

[C55] **Jiatao Gu**, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. “Non-Autoregressive Neural Machine Translation”. In: *ICLR 2018*. arXiv:1711.02281. 2018.

[C56] **Jiatao Gu**, Yong Wang, Yun Chen, Kyunghyun Cho, and Victor OK Li. “Meta-Learning for Low-Resource Neural Machine Translation”. In: *EMNLP 2018*. arXiv:1808.08437. 2018.

[C57] **Jiatao Gu**, Yong Wang, Kyunghyun Cho, and Victor OK Li. “Search Engine Guided Neural Machine Translation”. In: *AAAI 2018*. arXiv:1705.07267. 2018.

[C58] **Jiatao Gu**, Daniel Jiwoong Im, and Victor OK Li. “Neural machine translation with gumbel-greedy decoding”. In: *AAAI 2018*. arXiv:1706.07518. 2018.

[C59] **Jiatao Gu**, Hany Hassan, Jacob Devlin, and Victor OK Li. “Universal Neural Machine Translation for Extremely Low Resource Languages”. In: *NAACL 2018*. arXiv:1802.05368. 2018.

[C60] **Jiatao Gu**, Graham Neubig, Kyunghyun Cho, and Victor OK Li. “Learning to translate in real-time with neural machine translation”. In: *EACL 2017*. arXiv:1610.00388. 2017.

[C61] **Jiatao Gu**, Kyunghyun Cho, and Victor OK Li. “Trainable greedy decoding for neural machine translation”. In: *EMNLP 2017*. arXiv:1702.02429. 2017.

[C62] **Jiatao Gu**, Zhengdong Lu, Hang Li, and Victor OK Li. “Incorporating copying mechanism in sequence-to-sequence learning”. In: *ACL 2016*. arXiv:1603.06393. 2016.

Journal Papers

[J1] Qiao Feng, Yiming Huang, Yufu Wang, **Jiatao Gu**, and Lingjie Liu. “PhysHMR: Learning Humanoid Control Policies from Vision for Physically Plausible Human Motion Reconstruction”. In: *SIGGRAPH Asia 2025* (2025). arXiv:2510.02566.

[J2] Dinghuai Zhang, Yizhe Zhang, **Jiatao Gu**, Ruixiang Zhang, Joshua Susskind, Navdeep Jaitly, and Shuangfei Zhai. “Improving GFlowNets for Text-to-Image Diffusion Alignment”. In: *TMLR* (2025). arXiv:2406.00633.

[J3] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, **Jiatao Gu**, and Christian Theobalt. “Neural Actor: Neural Free-view Synthesis of Human Actors with Pose Control”. In: *SIGGRAPH ASIA 2021* (2021). arXiv:2106.02019.

[J4] Yinhan Liu, **Jiatao Gu**, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. “Multilingual denoising pre-training for neural machine translation”. In: *TACL 2020* (2020). arXiv:2001.08210.

[J5] **Jiatao Gu**, Qi Liu, and Kyunghyun Cho. “Insertion-based decoding with automatically inferred generation order”. In: *TACL 2019* (2019). arXiv:1902.01370.

Preprints

[P1] Yunzhou Song, Long Le, Yong-Hyun Park, Jie Wang, Junyao Shi, Lingjie Liu, **Jiatao Gu**, Eric Eaton, Dinesh Jayaraman, and Kostas Daniilidis. “OmniGuide: Universal Guidance Fields for Enhancing Generalist Robot Policies”. In: *arXiv preprint arXiv:2603.10052* (2026). arXiv:2603.10052.

[P2] David Berthelot, Tianrong Chen, **Jiatao Gu**, Marco Cuturi, Laurent Dinh, Bhavik Chandna, Michal Klein, Joshua Susskind, and Shuangfei Zhai. “The Coupling Within: Flow Matching via Distilled Normalizing Flows”. In: *arXiv preprint arXiv:2603.09014* (2026). arXiv:2603.09014.

[P3] Yao Tang, Li Dong, Yaru Hao, Qingxiu Dong, Furu Wei, and **Jiatao Gu**. “Multiplex Thinking: Reasoning via Token-wise Branch-and-Merge”. In: *arXiv preprint arXiv:2601.08808* (2026). arXiv:2601.08808.

[P4] Yizhe Zhang, Richard Bai, Zijin Gu, Ruixiang Zhang, **Jiatao Gu**, Emmanuel Abbe, Samy Bengio, and Navdeep Jaitly. “Reversal Blessing: Thinking Backward May Outpace Thinking Forward in Multi-Choice Questions”. In: *arXiv preprint arXiv:2502.18435* (2025). arXiv:2502.18435.

[P5] Shansan Gong, Ruixiang Zhang, Huangjie Zheng, **Jiatao Gu**, Navdeep Jaitly, Lingpeng Kong, and Yizhe Zhang. “DiffuCoder: Understanding and Improving Masked Diffusion Models for Code Generation”. In: *arXiv preprint arXiv:2506.20639* (2025). arXiv:2506.20639.

[P6] Georgia Gabriela Sampaio, Ruixiang Zhang, Shuangfei Zhai, **Jiatao Gu**, Joshua Susskind, Navdeep Jaitly, and Yizhe Zhang. “TypeScore: A Text Fidelity Metric for Text-to-Image Generative Models”. In: *arXiv preprint arXiv:2411.02437* (2024). arXiv:2411.02437.

[P7] Chen Wang, **Jiatao Gu**, Xiaoxiao Long, Yuan Liu, and Lingjie Liu. “GECO: Generative Image-to-3D within a SECOnd”. In: *arXiv preprint arXiv:2405.20327* (2024). arXiv:2405.20327.

[P8] Ying Shen, Yizhe Zhang, Shuangfei Zhai, Lifu Huang, Joshua Susskind, and **Jiatao Gu**. “Many-to-many Image Generation with Auto-regressive Diffusion Models”. In: *arXiv preprint arXiv:2404.03109* (2024). arXiv:2404.03109.

[P9] Liangchen Song, Liangliang Cao, **Jiatao Gu**, Yifan Jiang, Junsong Yuan, and Hao Tang. “Efficient-NeRF2NeRF: Streamlining Text-Driven 3D Editing with Multiview Correspondence-Enhanced Diffusion Models”. In: *arXiv preprint arXiv:2312.08563* (2023). arXiv:2312.08563.

[P10] Samira Abnar, Omid Saremi, Laurent Dinh, Shantel Wilson, Miguel Angel Bautista, Chen Huang, Vimal Thilak, Etai Littwin, **Jiatao Gu**, Joshua Susskind, and Samy Bengio. “Adaptivity and Modularity for Efficient Generalization Over Task Complexity”. In: *arXiv preprint arXiv:2310.08866* (2023). arXiv:2310.08866.

[P11] Peng Wang, Yuan Liu, Guying Lin, **Jiatao Gu**, Lingjie Liu, Taku Komura, and Wenping Wang. “Progressively-connected Light Field Network for Efficient View Synthesis”. In: *arXiv preprint arXiv:2207.04465* (2022). arXiv:2207.04465.

[P12] Sravya Popuri, Peng-Jen Chen, Changhan Wang, Juan Pino, Yossi Adi, **Jiatao Gu**, Wei-Ning Hsu, and Ann Lee. “Enhanced Direct Speech-to-Speech Translation Using Self-supervised Pre-training and Data Augmentation”. In: *arXiv preprint arXiv:2204.02967* (2022). arXiv:2204.02967.

[P13] Zhuofeng Wu, Sinong Wang, **Jiatao Gu**, Madian Khabsa, Fei Sun, and Hao Ma. “Clear: Contrastive learning for sentence representation”. In: *arXiv preprint arXiv:2012.15466* (2021). arXiv:2012.15466.

Workshop Papers

[W1] Yuan Gao, Chen Chen, Tianrong Chen, and **Jiatao Gu**. “One Layer Is Enough: Adapting Pretrained Visual Encoders for Image Generation”. In: *arXiv preprint arXiv:2512.07829* (2026). arXiv:2512.07829.

[W2] Changhan Wang, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Ann Lee, Peng-Jen Chen, **Jiatao Gu**, and Juan Pino. “fairseq S^2: A Scalable and Integrable Speech Synthesis Toolkit”. In: *arXiv preprint arXiv:2109.06912* (2021). arXiv:2109.06912.

[W3] Changhan Wang, Juan Pino, Anne Wu, and **Jiatao Gu**. “CoVoST: A diverse multilingual speech-to-text translation corpus”. In: *arXiv preprint arXiv:2002.01320* (2020). arXiv:2002.01320.

[W4] Xutai Ma, Mohammad Javad Dousti, Changhan Wang, **Jiatao Gu**, and Juan Pino. “Simuleval: An evaluation toolkit for simultaneous translation”. In: *arXiv preprint arXiv:2007.16193* (2020). arXiv:2007.16193.

[W5] Juan Pino, Liezl Puzon, **Jiatao Gu**, Xutai Ma, Arya D McCarthy, and Deepak Gopinath. “Harnessing indirect training data for end-to-end automatic speech translation: Tricks of the trade”. In: *arXiv preprint arXiv:1909.06515* (2019). arXiv:1909.06515.

[W6] Changhan Wang, Anirudh Jain, Danlu Chen, and **Jiatao Gu**. “VizSeq: A Visual Analysis Toolkit for Text Generation Tasks”. In: *arXiv preprint arXiv:1909.05424* (2019). arXiv:1909.05424.