Brian Chao

brianchc@stanford.edu • bchao1.github.io • Google Scholar • github.com/bchao1 • LinkedIn

Research Interests

I study how visual generative models should represent information and allocate computation. My recent work explores adaptive representations for diffusion models across space, scale, and time, spanning pre-training, post-training, and inference (Foveated Diffusion, Spectral Progressive Diffusion, Level-of-Token Diffusion), with broader interests in long-context autoregressive video generation, world models, robotics, and information-theoretic aspects of diffusion.

Education

Stanford University · *Ph.D. Candidate in Electrical Engineering* 2022 – Present
Advisor: Prof. Gordon Wetzstein. Supported by the NSF GRFP and the Stanford Graduate Fellowship.

National Taiwan University · *B.S. in Electrical Engineering* 2017 – 2021

Publications

Visual Generation

[1] **Level-of-Token Diffusion**. Kiyohiro Nakayama, **Brian Chao**, Jan Ackermann, Hansheng Chen, Federico Tombari, Leonidas Guibas, Lior Yariv, Gordon Wetzstein. *arXiv*, 2026.
A generative level-of-detail diffusion framework that allocates token budget based on input control signals, achieving up to 4× *speedup in image and video generation.*

[2] **Spectral Progressive Diffusion for Efficient Image and Video Generation**. Howard Xiao, **Brian Chao**, Lior Yariv, Gordon Wetzstein. *NeurIPS*, 2026.
A plug-and-play acceleration framework that exploits the intrinsic spectral autoregression of diffusion.
Adopted by SGLang, with >2× speedup across FLUX.1 & 2, Z-Image, Qwen-Image, Ideogram, and Wan.

[3] **Foveated Diffusion: Efficient Spatially Adaptive Image and Video Generation**. **Brian Chao***, Lior Yariv*, Howard Xiao, Gordon Wetzstein. *arXiv*, 2026.
Spatially adaptive tokenization that concentrates compute on selected regions for up to 4× *faster generation.*
Invited talks: NVIDIA Research, Black Forest Labs, Cohere, Decart, Rhoda AI.

Other Selected Publications

[4] Textured Gaussians for Enhanced 3D Scene Appearance Modeling (first author) *CVPR 2025*
[5] Gaussian Wave Splatting for Computer-generated Holography (co-first author) *SIGGRAPH 2025*
[6] Fast Wave-optics Rendering of Multiplane Images for 3D Holographic Displays (first author) *arXiv 2026*
[7] Random-phase Wave Splatting of Translucent Primitives for Holography (co-first author) *arXiv 2025*
[8] Large Étendue 3D Holographic Display with Dynamic Fourier Modulation (first author) *SIGGRAPH Asia 2024*
[9] Full-colour 3D Holographic AR Displays with Metasurface Waveguides *Nature 2024*
[10] High Brightness Holographic Projection (first author) *Optics Letters 2023*
[11] Time-Division Multiplexing Light Field Display with Learned Coded Aperture (co-first author) *IEEE TIP 2022*
[12] Robust Light Field Synthesis from Stereo Images with Left-Right Geometric Consistency (first author) *ICIP 2021*
[13] Self-Supervised Deep Learning for Fisheye Image Rectification (co-first author) *ICASSP 2020*

* denotes equal contribution. Full list on Google Scholar.

Industry Experience

Waymo Research, AI Foundations Team · *Research Scientist Intern* Jun 2026 – Oct 2026

- 4D world and video models with Titas Anciukevičius, Max Jiang, Vincent Casser, Philipp Henzler, Thomas Funkhouser.

Meta Reality Labs · *Research Scientist Intern* Jun 2025 – Jan 2026

- Neural rendering for computational 3D displays (Fast Wave-optics MPI Rendering) with Grace Kuo and Changwon Jang.

Meta Reality Labs · *Research Scientist Intern* Jun 2024 – Jan 2025

- Neural 3D scene reconstruction (Textured Gaussians) with Changil Kim, Hung-Yu Tseng, Jia-Bin Huang, Johannes Kopf.