Linbo Tang

+1 608-886-7491 | linbo0204@outlook.com | LinkedIn | GitHub | Portfolio

Education

Harvard University — MA, USA
S.M in Data Science; ***GPA: 4.00/4.00*** — *Sep 2024 – May 2026 (expected)*

University of Wisconsin - Madison — WI, USA
B.S in Computer Science and Data Science; ***GPA: 4.00/4.00*** — *Sep 2020 – Dec 2023*

Research Interests

Computer vision, generative modeling, and medical imaging — with a focus on integrating geometric, temporal, and physical constraints into deep learning frameworks for interpretable and domain-adaptive perception. Broadly, I aim to develop models that unify low-level reconstruction and high-level reasoning, enabling consistent visual understanding across spatial, temporal, and semantic scales.

Research Experience

Zickler's Vision Lab — Dec 2024 – Present
Harvard University — *Advisor: Prof. Todd Zickler*

- Developing a diffusion-based framework for stereo vision that jointly models disparity and geometric scene structure from synthetic datasets.
- Adapted an **hourglass diffusion transformer** for stereo disparity estimation, employing image-conditioned geometric embeddings and progressive diffusion refinement to achieve higher depth fidelity and structural consistency.
- Developed a large-scale synthetic data generation pipeline producing stereo image pairs, disparity maps, and geometric annotations for model training, evaluation, and ablation studies.

Psychiatry and Radiology Image Computing Group — Sep 2024 – Present
Harvard Medical School — *Advisor: Prof. Yogesh Rathi*

- Developed a novel **scanner-adaptive, coil-level** denoising framework that conditions the network on measured noise statistics to achieve robust performance across scanners and protocols, submitted to **ISBI 2026**
- Designed a supervised training strategy with synthetic coil-wise noise realizations to ensure physically consistent noise modeling and improved generalization across acquisition settings.
- Introduced an attention-level conditioning mechanism via feature-wise modulation to dynamically adjust denoising strength according to local noise context.
- Evaluated the framework on in-vivo multi-coil diffusion MRI data, achieving up to a **23%** reduction in residual noise compared to the next best-performing method

MSC Computer Vision Group — Jan 2023 – Feb 2024
University of Wisconsin - Madison — *Advisor: Prof. Vikas Singh*

- Curating **EmotionNet5k**, a novel selfie video dataset containing 5000 samples of facial expressions, with over 300 stars on GitHub;
- Implementing a SOTA learnable wavelet model to remove co-founders and achieve a **3.13%** increase in accuracy on the downstream Facial Emotion Recognition (FER) task.
- Utilized **Category Theory** to improve compositionality of **StableDiffusion** by learning morphisms between embeddings to ensure reflexivity, transitivity, and neighborhood constraints;
- Developed metrics to verify Transitivity Axioms compliance and correlate with CLIP scores for model assessment.

MIT CSAIL — Apr 2025 – Present
MIT EECS — *Advisor: Dr. Giannis Daras and Prof. Jon Tamir*

- Developed **ContextMRI**, a diffusion-based MRI reconstruction framework that leverages **fine-grained metadata embeddings** to condition on imaging parameters such as TR, TE, and flip angle.
- Designed a **structured embedding architecture** separating categorical and numerical attributes, using discrete embeddings and sinusoidal encodings respectively for improved quantitative representation.
- Introduced an **out-of-distribution (OOD) training strategy** by incorporating cardiac MRI data with distinct parameter ranges, forcing reliance on metadata priors and improving generalization.

- Demonstrated superior reconstruction accuracy and interpretability on the fastMRI brain dataset, particularly under diverse or uncommon acquisition settings.

Ubisoft La Forge Apr 2024 – Aug 2024

Research Student; La Forge *Advisor: Alexis Rolland*

- Developed a generative base-material creation pipeline utilizing **StableDiffusion XL** and **inverse rendering** models to obtain rendered textures from text prompts and physically based rendering (PBR) maps from rendered textures.
- Refined a SwinIR-based super-resolution model for PBR maps by introducing a multi-branch architecture that separately processes base color, normal, roughness, and height maps while enforcing inter-map consistency.
- Achieved higher perceptual fidelity and improved structural coherence across PBR maps compared to standard single-branch and texture-domain super-resolution baselines.

Teach Experience

Teaching Fellow Aug 2025 – Present

Harvard University *Advisor: Prof. Michael Brenner*

- Led weekly discussion sections for **AM 215: Mathematical Modeling**, covering topics such as dimensional analysis, Monte-Carlo simulation, and extreme-value statistics.
- Collaborated with the teaching staff on the design of homework assignments and course structure, providing individualized guidance to students throughout the semester.

Undergraduate Teaching Assistant Sep 2022 – May 2023

University of Wisconsin–Madison *Advisor: Prof. Frederic Sala*

- Assisted in teaching three core courses: **CS 200 (Programming I)**, **CS 300 (Programming II)**, and **CS 540 (Introduction to Artificial Intelligence)**.
- Guided students through fundamental and intermediate programming concepts, data structures, and AI algorithms, emphasizing debugging and problem-solving strategies.
- Conducted tutoring sessions twice weekly, supporting over **400 students** across two semesters.
- Created and maintained a public teaching portfolio showcasing section materials, code examples, and instructional resources.

Technical Skills

Programming Languages: Java, Python, C/C++, R, JavaScript, CSS, HTML, Kotlin;

Frameworks: Pytorch, Tensorflow, Scikit-Learn, MySQL, OpenCV, Flask, React, Selenium;

Platforms: Isaac Sim, UE4, AWS, GCP, Docker, Linux, Git.