

Junhao Zhu

Atlanta, GA 30308 | (470) 652-9837 | junhao.zhu@gatech.edu | github.com/JunhaoZhu0220 | [Personal Page](#)

Availability: May 2027 – Dec 2027

EDUCATION

Georgia Institute of Technology

M.S. in Computer Science, Machine Learning Specialization

Atlanta, GA

Aug 2026 – May 2028 (Expected)

Relevant Coursework: Machine Learning, Natural Language Processing, Deep Learning for Text Data

New York University Shanghai

B.S. in Computer Science, Minor in Interactive Media Arts

Shanghai, China

Sep 2022 – May 2026

Cumulative GPA: 3.787 **Dean's List** (2023–2026) | **Founders' Day Award**

PUBLICATIONS

Learning to Be Guided: Shaping Latent Geometry for Classifier-Free Guidance

Haoming Liu*, **Junhao Zhu***, Shaoyi Zheng*, Yuanhe Guo, Jinnuo Liu, Shenji Wan, Hongyi Wen. (*Equal contribution)

arXiv preprint, 2026 | [\[Preprint\]](#) [\[Project Page\]](#)

Coherence-Aware Distributional Evaluation of Open-Ended Text Generation

Jinnuo Liu*, **Junhao Zhu***, Weifeng Jiang, Haoming Liu, Hongyi Wen. (*Co-first authors)

arXiv preprint, 2026 | [\[Preprint\]](#) [\[Code\]](#)

RESEARCH EXPERIENCE

Research Assistant

NYU Shanghai Center for Data Science

Oct 2024 – Present

Shanghai, China

CHORD: Coherence-Aware Evaluation of LLM Text Generation

May 2026 – Sep 2026

Advisor: Prof. Hongyi Wen

- Developed CHORD, a distributional metric for document-level coherence in LLM-generated text, using discourse-sensitive representations from frozen pretrained language models.
- Built a diagnostic evaluation suite with **9 coherence-degrading perturbations** and matched meaning-preserving controls to distinguish coherence failures from benign rewrites.
- Detected **all 9 failure types**, compared with 3–5 for MAUVE, FBD, and perplexity; achieved **Spearman $\rho = 0.98$** with human judgments, with results consistent across 4 LLM backbones (Qwen, Gemma, Mistral, Llama).
- Led distillation design for **0.8B and 2B Qwen-3.5 models**, retaining most of the 27B model's detection capability while reducing computational cost for lightweight deployment.

Learning to Be Guided: Shaping Latent Geometry for Classifier-Free Guidance

May 2026 – Present

Advisor: Prof. Hongyi Wen

- Proposed Prototype Regularization, a lightweight objective that aligns latent diffusion representations with learned mode prototypes, improving guided ImageNet generation and extending to video and text-to-image tasks.
- Led integration into the LeanVAE pipeline for class-conditional video generation on UCF101; reduced Fréchet Video Distance (FVD) by **16%** (167 to 140) and increased generated-video action-recognition accuracy from 92.9% to **98.9%** relative to the official baseline; achieved a **19% FVD reduction** without guidance.

Syntax-Guided Cascade Discrete Diffusion for Text Generation [\[Code\]](#) [\[Thesis\]](#)

Jun 2025 – Dec 2025

Advisor: Prof. Shengjie Wang & Prof. Hongyi Wen

- Designed a **3-stage diffusion framework** that generates structural words before completing sentences, using a parameter-free conditioning mechanism.
- Compared syntax-based and semantic anchor selection; the syntax-guided variant **reduced perplexity by $\sim 21\%$** relative to a standard single-stage diffusion baseline.
- Received the **Capstone Presentation Distinction Award**.

SELECTED PROJECTS

Better Anchors, Better Text: Anchor Selection in Diffusion Language Models [\[Code\]](#) [\[Report\]](#)

May 2026

Course project, CSCI-SHU 376 Natural Language Processing | New York University Shanghai

- Built a modular anchor-selection framework for diffusion language models, fixing informative words to guide generation and enabling interchangeable scoring methods within the same training pipeline.
- Achieved **6.5% lower perplexity** and **12.2% higher MAUVE** than the unanchored baseline on OpenWebText while preserving output diversity.

Open-Source Contribution: ml5.js Feature Extractor Module [\[Code\]](#) [\[Docs\]](#)

Jan 2026 – Aug 2026

- Collaborated with digital media faculty and open-source maintainers to ship a MobileNet-based transfer-learning module in **ml5.js v1.4.0**, enabling in-browser training of custom image classifiers and regressors; built two real-time webcam demos with model save/load support.

TECHNICAL SKILLS

Systems & Infrastructure: Linux/Unix, Bash, Git, Docker, Singularity, Slurm (HPC), Weights & Biases, $\LaTeX$

Languages: Python, C, JavaScript, R

Frameworks & Libraries: PyTorch, Hugging Face Transformers, Accelerate, NumPy, pandas, spaCy, NLTK

Areas: Large Language Models, Diffusion Models, NLP Evaluation, Generative AI