

ZHUOYAN (TERRY) TAO

213-312-7936 | terrytao@andrew.cmu.edu | linkedin.com/in/zhuoyan-tao | github.com/ZhuoyanTao | zhuoyantao.com

EDUCATION

Carnegie Mellon University

Master of Science in Intelligent Information Systems, Language Technologies Institute

Pittsburgh, PA

Dec 2027

Relevant Coursework: Advanced NLP, Speech Generation, Deep Learning Systems

University of Southern California

Bachelor of Science in Computer Science; Bachelor of Science in Applied and Computational Mathematics

Los Angeles, CA

May 2026

GPA: 3.84/4.0 | Magna Cum Laude | Relevant Coursework: Computer Systems, Operating Systems, Algorithms, Artificial Intelligence, Embedded Systems, Numerical Methods, Optimization, Linear Algebra, Probability Theory

SKILLS

ML & Speech: PyTorch, torchaudio, Hugging Face Transformers, ESPnet, VERSA; Whisper, MMS, WavLM, HuBERT

Data: NumPy, pandas, Matplotlib, Gaussian-process regression

Systems: Python, C++, Go, Bash/Shell, Docker, Linux, SLURM, Git

HPC & Performance: profiling (py-spy/cProfile), vectorization, distributed training (DDP), NCSA Delta / PSC Bridges-2 GPU clusters

EXPERIENCE

Apple Inc.

Austin, TX

Software Engineering Intern

May 2026 – Aug 2026

- Built (sole author) a sweep-orchestration tool for the team's satellite network simulator, scaling experiment throughput **~1000×** to **~2000 concurrent experiments** on an internal compute cluster; presented org-wide
- Integrated the new tooling to drive analysis the team couldn't run before: swept projected demand over years of growth to model how capacity and performance evolve, using Gaussian-process sensitivity analysis to rank the most influential parameters
- Delivered a **~12× speedup** to the simulator's interference-computation hot path via batched vectorization; contributed reproducibility fixes unifying random-number generation and the simulation clock
- Quantified the impact of timing bugs and edge cases on overall system behavior, root-causing a Python/Go clock-scheduler race that silently skipped **~half** of scheduling cycles and under-reported capacity by up to **~2×** under load
- Reviewed teammates' PRs and helped validate changes across the 29K-line shared Python/Go codebase, catching a latent numerical bug in review that would have corrupted output logs

PUBLICATIONS

Z. Tao, J. Shi, H. Shim, S. Watanabe. "ANCHOR: Autoregressive Non-intrusive Chunk-Ordered Refinement for Joint Multi-Resolution Speech Quality Modeling." Proc. Interspeech 2026 (**Oral**). [**First Author**] arXiv:2606.10233

Z. Tao, A. Kommineni, S. Narayanan. "Off-Target LID Posterior Mass Predicts Probe Gains in Spoken Language Diarization: A Pre-Registered Validation." EMNLP, under review. [**First Author**]

F. Jiang, R. Sanabria, S. Deshmukh, ..., **Z. Tao** et al. "TurnBench: A Multi-Domain Benchmark for Turn-Taking Dynamics in Spoken Dialogue." IEEE SLT 2026, under review. arXiv:2608.25218

RESEARCH PROJECTS

CMU WavLab | Speech & Audio ML Research

Pittsburgh, PA

Advisor: Prof. Shinji Watanabe

Jun 2025 – Present

- Proposed **ANCHOR**, extending ARECHO (NeurIPS 2025 Spotlight) to predict speech quality from partial audio; introduced dual-resolution query tokens with chunk-first decoding in a shared Transformer decoder, cutting **PLCMOS MAE by 48%** on 2 s prefixes and establishing a **4–6 s perceptual context horizon**
- Built HuBERT- and prosody-based speech-to-speech translation systems preserving pragmatic intent (stance, emotion, dialog function) across English–Spanish for the Interspeech 2026 S2ST Challenge
- Merged multiple PRs into **ESPnet (16k+ stars)**, including a **≈50×** speedup for speaker-embedding extraction and new TTS recipes

USC SAIL Lab | Multilingual Speech Processing Research

Los Angeles, CA

Advisor: Prof. Shrikanth Narayanan

Aug 2023 – May 2026

- Introduced **off-target LID posterior mass** as a diagnostic measure predicting downstream Spoken Language Diarization (SLD) probe gains on pre-trained speech encoders (Whisper, MMS, classical LID); first-author paper under review at EMNLP
- Designed and pre-registered an evaluation protocol on unseen language pairs; tested generalization on two held-out real-world corpora (DISPLACE; MUCS Hindi–English code-switching), with observed probe gains falling within the frozen **95% prediction interval**
- Validated the diagnostic across all three encoder families, with the mapping to gains differing by encoder; isolated which utterances benefit most from fine-tuning in low-resource and code-switching settings

LEADERSHIP

Supplemental Instruction Leader | USC MATH 229 (Multivariable Calculus)

Aug 2023 – May 2026

- Led 3–4 weekly sessions for 98 students; review sessions reached **~300 attendees**, the highest in program history

VP & Executive Producer | Unity 2D Platformer, Open Alpha Game Design Club

Jan 2023 – Apr 2024

- Directed a Unity 2D platformer with a 10-member audio team (Ableton Live & Wwise); shipped on Steam (**85% positive**)
- Designed and taught a semester-long Unity & SFX curriculum for 40 students; recruited and mentored incoming cohorts

HONORS

Phi Beta Kappa (2025) | **Trustee Scholar** (full tuition), USC | **Grand Challenges Scholar**, NAE / USC Viterbi (2026)