Guorui Sang

Ph.D. Student (2nd year), Computer Science, University of Illinois Chicago
gsang@uic.edu | Chicago, IL | guoruisang.github.io | Google Scholar | LinkedIn

Seeking Summer 2027 research internship – available May 9 – Aug 21, 2027, full-time

Generative-modeling researcher: I build **diffusion and consistency models** for robot trajectories, singing voice, and images, and transformer architectures that work on continuous data. **1 ICASSP 2025 paper** (co-first author), **1 ICLR 2026 workshop paper**, **2 first-author submissions** (ICLR 2027, ICRA 2027).

EDUCATION

University of Illinois Chicago — Ph.D., Computer Science | GPA 4.00/4.00 Aug 2025 – May 2030 (expected)
Beijing University of Technology — M.S., Computer Science | GPA 3.94/4.00 Sep 2022 – Jul 2025
Hainan University — B.S., Software Engineering | GPA 3.66/4.00 Sep 2017 – Jul 2021

RESEARCH EXPERIENCE

University of Illinois Chicago — Ph.D. research, advisor Pedram Rooshenas Aug 2025 – Present
Graduate Research Assistant, May 2026 – Present; Teaching Assistant (MATLAB course: labs, grading, proctoring), Aug 2025 – May 2026

Diffusion Fields with Hamiltonian Correction — robot trajectories *ICLR 2026 Workshop; submitted to ICRA 2027*

- Built an **action-conditioned diffusion model of robot dynamics**: from an initial state and a torque sequence, it generates the arm's motion as a continuous-time field (diffusion probabilistic field), so one model covers **variable horizons** and can be **queried at any time**.
- Added a latent Hamiltonian (energy-based) correction for physical consistency: **position/momentum RMSE down 76%/73%** vs. the uncorrected model; **47/50 (94%)** three-target reaching sequences on a 3-DoF MuJoCo arm.

Function-Space Transformer — architecture for continuous data *arXiv:2609.38348; submitted to ICLR 2027*

- Designed a transformer that works on **continuous functions instead of fixed grids**: it learns where to place latent anchors from the input and refines them recursively, so one model handles PDEs and images at any resolution.
- **Outperforms Perceiver IO and matches FNO** on PDEBench (Burgers rel. L2 0.0182 vs. 0.0185; on par on Darcy); on ImageNet-1K, **+1.7 pts top-1 over a ViT baseline** (69.5% vs. 67.8%) with fewer parameters (∼3M).

Beijing University of Technology — Graduate Researcher Oct 2023 – Feb 2025

ConSinger — few-step singing voice synthesis *IEEE ICASSP 2025, co-first author*

- Co-developed a **consistency-model** singing-voice synthesizer that generates singing from a music score (lyrics, duration, pitch) in a few denoising steps, plus an **FAD-based scorer** that picks the best denoising start point.
- On PopCS: **2.8× faster than DiffSinger** (RTF 0.018 vs. 0.052), **better FAD** (0.32 vs. 0.88), MOS 3.88 vs. 3.81.

M.S. Thesis — pose-guided person image generation

- Diffusion model (three-branch U-Net, mixed CFG): **best SSIM and PSNR** vs. PIDM, UPGPT, NTED, DPTN.

PUBLICATIONS

- Y. Song*, **G. Sang***, J. Yu, C. Xiao. ConSinger: Efficient High-Fidelity Singing Voice Generation with Minimal Steps. *IEEE ICASSP 2025*. (*equal contribution) [arXiv] [demo]
- **G. Sang**, P. Rooshenas. Hamiltonian-Guided Diffusion Fields for Variable-Length Rigid-Arm Trajectory Generation. *ICLR 2026 Workshop (ReALM-GEN)*.
- **G. Sang**, P. Rooshenas. Function-Space Transformer with Adaptive Anchors. *Submitted to ICLR 2027*. [arXiv]
- **G. Sang**, P. Rooshenas. Diffusion Fields with Hamiltonian Correction for Variable-Horizon Robot Arm Trajectories. *Submitted to ICRA 2027*.

SKILLS

- **Generative modeling:** diffusion / score-based models, latent diffusion, DiT, VAEs, consistency models, diffusion policies, classifier-free guidance
- **Deep learning:** transformers (ViT), U-Nets, neural operators (FNO), physics-informed ML, multi-GPU training
- **Applications:** image generation, audio / singing-voice synthesis, robot trajectory generation, PDE surrogates
- **Tools:** Python, PyTorch, PyTorch Lightning, NumPy, MuJoCo, Weights & Biases, Git, Linux, C/C++