

Yuer Tang

(510) 953-0647 | yuertang17@g.ucla.edu | Los Angeles, CA | [Github](#) | [Website](#)

EDUCATION

University of California, Los Angeles (UCLA)

Los Angeles, CA

Major: Data Theory, BS

Sep. 2022–Expected Mar. 2026

Minor: Data Science Engineering Minor

Major GPA: 3.77

- Core Courses: Machine Learning (grad), Reinforcement Learning (grad), Cognitive Science
- Awards: Susan Elizabeth Baumgarten Award (2024-2025), Dean's Honors list

RESEARCH EXPERIENCE

Leslie Kaelbling Lab @ MIT

Cambridge, MA

MIT CSAIL Summer Intern (mentor: Dr. Jiayuan Mao)

Jun. 2025 - Present

Project: beta-RankVAE for Disentangled, continuous, and Interpretable Scale Control in Fine-Grained Robotic Policies

- Developed trajectory collection pipeline using MetaWorld simulation environment to assist model building and validation and used it to collect 6 DoF manipulation trajectories with targeted scale features.
- Built and implemented a novel beta-VAE architecture where Convolutional and Pooling layers were strategically integrated into the Encoder/Decoder to effectively preserve the spatial and temporal structure of the trajectories, mitigating the inherent trade-off of the beta-VAE.
- Engineered a custom loss function by integrating a marginal pairwise rank loss and a masked KL divergence on the targeted latent dimension, explicitly promoting it to capture the full, continuous spectrum of policy scale (e.g., degree of door openness).
- Developed an Inverse Kinematics visualization tool powered by a Random Forest model to transform reconstructed/generated 6-DoF policies into executable actions, enabling real-time visualization and high-fidelity evaluation in the simulated environment.
- Currently pioneering an LLM-Assisted Scale Perception module to automate the labeling of the scale feature and the generation of pairwise rank labels, streamlining the training data pipeline for generalized policy learning.

Yingnian Wu Lab @ UCLA

Los Angeles, CA

Undergraduate Researcher (directly with Professor)

Mar. 2025 - Present

Project: Meta-Adaptive Latent Planning (MALP)

- Improved the multi-task performance of Latent Plan Transformer) by integrating Model-Agnostic Meta-Learning (MAML) leading to MPI-MAML (Meta-Planning as Inference).
- Rebuilt LPT codebase for D4RL Kitchen dataset compatibility to enable MAML integration and streamline validation, improving reproducibility across meta-learning benchmarks.
- Engineered full MAML training loop with inner-loop K-shot adaptation and outer-loop meta-gradient accumulation to achieve rapid task generalization, enabling 5-shot adaptation to novel manipulation tasks.
- Currently Implementing ANIL (Almost-No-Inner-Loop) in the codebase to test reduced inner loop gradient computation, focusing on improving training speed, stability, and memory usage.

Andrea Bertozzi Lab @ UCLA

Los Angeles, CA

Undergraduate Researcher (mentor: Dr. Justin Baker)

Jan. 2025 - Oct. 2025

Project: On the Dynamics of Coherent Memory Structures in Neural Fields

- Develop intrinsically explainable AI architectures using biologically inspired neural field representations.
- Contributed to drafting figures, handling periodic, Dirichlet and free-flow boundary conditions, and performed a robust neuroscience background and literature review **for work under review at ICLR 2026**.

Project: : Operator-Theoretic Tools for Modeling Conscious-Unconscious Neural Dynamics

- Applied Mori-Zwanzig projection operator formalism into the neuroscience application of propofol induced loss of consciousness to explicitly separate neural dynamics into resolved (unconscious) and unresolved (conscious) states.
- Validated the model's predictive capability on neural data, achieving a major reduction in Mean Squared Error (MSE) compared to the baseline DMD (e.g., awake state MSE improved from 3.140 to 0.101).
- **Under review at AAAI 2026 workshop on NeuroAI.**

Project: When Linear Models Aren't Enough: Forcing as Closure in Chaotic Time Series.

- Clarified through mathematical derivation that HAVOK forcing term emerges when the truncated delay-embedded coordinates fail to span a Koopman-invariant subspace.
- Validated the framework on the Lorenz and Rössler systems to localize nonlinear transitions, showing the forcing signal activates sparsely and with high magnitude exclusively during regime shifts (e.g., Lorenz lobe-switching), confirming its role in capturing deviations from linear evolution.
- **Will present at 2026 Joint Mathematics Meeting** in Washington DC, January 6th, 2026.

Tao Gao Lab @ UCLA
Undergraduate Researcher

Los Angeles, CA
Mar. 2023 - May. 2024

Project: Online Bayesian Inverse Planning

- Engineered a foundational Markov Decision Process (MDP) solver utilizing Value Iteration and the Bellman Update operator to pre-compute optimal value functions and Q-values for multiple hypothetical goals.
- Developed a gamma-Discounted Softmax Policy Extractor to derive the probabilistic policy from the learned Q-values, incorporating a temperature parameter to control action stochasticity and provide the necessary action likelihoods for Bayesian updating.
- Implemented a robust Bayesian Goal Inference Mechanism that uses observed agent actions and the pre-computed policies to iteratively update the posterior probability of each latent goal in real-time.
- Architected and deployed a full-stack, low-latency online visualization system using Python (for the RL inference backend) and Node.js/Socket.io (for the frontend/server), establishing a live data pipeline for real-time demonstration of goal inference results.

PUBLICATIONS AND PRESENTATIONS

1. **Tang, Yuer.**, Baker, Justin. “*Operator-Theoretic Tools for Conscious and Unconscious Brain Activity.*” Under Review at AAAI NeuroAI workshop (2026).
<https://openreview.net/group?id=AAAI.org%2F2026%2FWorkshop%2FNeuroAI%2FAuthors&referrer=>
2. **Tang, Yuer.** *When Linear Models Aren't Enough: Forcing as Closure in Chaotic Time Series.* Poster presentation at Joint Mathematics Meeting, January 6, 2026, Washington, D.C.
<https://meetings.ams.org/math/jmm2026/meetingapp.cgi/Paper/59043>
3. Baker, J.M., Hardiman-Mostow, H., **Tang, Y.**, Bertozzi, A.L. “*On the Dynamics of Coherent Memory Structures in Neural Fields.*” Under Review at ICLR (2026).
<https://openreview.net/forum?id=RanDLuKUQT>
4. Jiang, E.H., Zhang, Z., Zhang, D., Lizarraga, A., Xu, C., Zhang, Y., Zhao, S., Xu, Z., Yu, P., **Tang, Y.**, Kong, D., Wu, Y.N. “*DODT: Enhanced Online Decision Transformer Learning through Dreamer's Actor-Critic Trajectory Forecasting.*” arXiv (2024).
<https://arxiv.org/abs/2410.11359>

INDUSTRY EXPERIENCE

1585 Healthcare (*Value-based Healthcare Investment firm*)
Data Science & Data Engineering Intern

Brooklyn, NYC
Jun. 2024 - Sep. 2024

- Built reliable data pipelines that cleaned, matched, and organized large healthcare datasets, including more than one hundred thousand patient records and three hundred thousand insurance claims.
- Automated repeated workflows such as PDF extraction, geocoding, and case matching, reducing manual work and ensuring consistent, error-free data for downstream analysis.