Yu (Zoey) Zhu

Email : yzhu201@ucsc.edu
Linkedin: https://www.linkedin.com/in/yu-zoey-zhu-8a330113a/
Mobile : +1-(530)-304-6852
Personal Web: https://yuzoeyzhu.github.io/

Education

- **University of California, Santa Cruz** — Santa Cruz, CA
  PhD in Statistics; GPA: 4.0 — *Sep 2020 - Jun 2025*
- **University of California, Davis** — Davis, CA
  MS in Statistics; GPA: 3.9 — *Sep 2017 - Jun 2019*

Skills Summary

- **Languages**: R, Python, SQL, Java, C++, JavaScript, Machine Learning Frameworks: Keras, PyTorch, Tensorflow
- **Statistics and Experimentation**: A/B Testing, Experimental Design, Causal Inference, Bayesian Optimization
- **AI & Learning Systems**: Reinforcement Learning (Policy Optimization, Q-learning, PPO/GRPO-style Optimization, RLHF), Sequence-level RL for Ranking & Personalization, LLM Post-training (SFT, Distillation, Preference Modeling), LLM Inference & Serving (vLLM, Low-latency Deployment), Reward Modeling & Alignment under Distribution Shift, Agentic Workflow Design & Feedback Loops, Transfer Learning & Cross-Domain Generalization, Online Experimentation & Evaluation Harness Design

Work Experience

- **Meta** — Menlo Park, CA
  Research Scientist, Machine Learning (Feed Recommendation System) — *June 2025 - Present*
  - **Reinforcement Learning Sequence Generator with Transfer Learning and Policy Optimization**:
    - Contributed to the development of an RL-based sequence generator that produces candidate feed sequences from late-stage ranked item lists.
    - Designed a transfer learning structure for the recurrent policy model to improve training efficiency and accelerate convergence under large-scale personalization data.
    - Improved recurrent training stability through reward function refinement and hyperparameter tuning, reducing variance under noisy engagement signals.
    - Enhanced reward alignment by refining debiased Viewport View-based objective to better match downstream online engagement metrics.
    - Supported multi-objective personalization at sequence level, balancing engagement quality and content diversity.
  - **LLM Evaluation for Blender Sequence with Distilled Judge Model**:
    - Developed an LLM-as-Judge framework to evaluate RL-generated feed recommendation sequence candidates using structured hierarchical latent reasoning.
    - Bootstrapped supervision by leveraging Gemini to generate comparative judgments and reasoning traces over candidate sequences.
    - Performed supervised fine-tuning (SFT) and knowledge distillation to train a lightweight evaluator model suitable for low-latency production inference. Modeled distilled evaluator outputs as sequence-level preference signals integrated into the blender system.
    - Validated evaluator impact through controlled online A/B experiments, delivering measurable lift in topline metrics compared to existing evaluator with value model structure.
    - Used LLM reasoning traces to identify reward mis-specification and sequence-level failure modes, informing subsequent generator reward refinement.
- **Microsoft** — Redmond, WA
  Research Data Scientist Intern (Azure Compute) — *June 2024 - Sept 2024*
  - **Adaptive Reinforcement Learning under Distribution Shift for Azure Configuration Allocation [Report]**:
    - Formulated dynamic configuration allocation as a constrained Markov Decision Process (MDP) under non-stationary reward distributions and combinatorial resource constraints.
    - Modeled evolving configuration failure probabilities as distribution shift over time, including gradual stochastic drift and abrupt regime changes.
    - Designed a hybrid model-based / model-free RL framework integrating rolling statistical estimators with Q-learning policy updates.
    - Engineered a reward shaping mechanism combining simulated coverage signals and real-time testing feedback to improve sample efficiency.
    - Implemented an online–offline training pipeline with adaptive exploration to maintain policy robustness under probability drift, consistently outperforming static and optimization-based baselines.

- **Human-Aligned LLM Agent for Root Cause Reasoning**:
  - Built a human-aligned LLM reasoning agent using Gemini to automate root cause identification and mitigation planning from Sev incident discussions.
  - Constructed a supervised fine-tuning (SFT) corpus from hierarchically labeled root cause taxonomy, incorporating semantic conditioning to improve fine-grained classification consistency.
  - Engineered structured reasoning prompts enabling interpretable reasoning traces and actionable mitigation outputs.
  - Designed a human-in-the-loop feedback interface capturing engineer preference signals and established feedback logging to enable future RLHF-style post-training.
  - Developed offline evaluation harness (accuracy / precision / recall) and significantly outperformed prior transformer-based classification baseline, reducing false positives and false negatives in production workflow

- **Tencent America** Palo Alto, CA
  Research Data Scientist Intern (IEG Global) *Oct 2023 - May 2024*
  - **Online Game Experimentation [CODE@MIT 2024]**:
    - Explored the challenge of A/B testing in online gaming, where dynamic and ephemeral network interference among players existed in user-randomized experiments, which compromised the validity of causal effects due to the SUTVA violations.
    - Proposed an innovative framework for treatment effect estimation tailored for scenarios where a completely randomized experimental design is implemented without explicit knowledge of network structures.
    - Developed the proposed method into an integrated pipeline for implementation within a real-game application.
  - **Media Channel Modeling for Game Advertising [Report]**:
    - Developed a Bayesian hierarchical model framework to evaluate the effectiveness of different media channels for a PC game's online ads campaigns, showing a great fit with an extremely short-term advertising data set.
    - Incorporated the Carryover and Shape Effects to correlate advertising costs with installs, validated on the methodologies of the mixture of weighted functions from Jin et al., 2017.
    - Assessed the impact of various priors and effect functions on model performance with limited data. Proposed the optimized advertising budget allocation plan for the incoming ads campaign.

Academic Research and Publications

- **Bayesian Latent Variable Model for Quantification Learning in Mortality Surveillance** Santa Cruz, CA
  Research Assistant - Prof. Zehang (Richard) Li *Sep 2020 - June 2025*
  - **Bayesian Federated Learning under Distribution Shift [AOAS in revision]**:
    - Proposed a modular Bayesian federated learning framework enabling cross-domain classification and population-level estimation without data sharing, achieving robust generalization under heterogeneous and shifting data distributions.
  - **Bayesian Latent Structure Modeling under Distribution Shift [JRSSA]**:
    - Developed hierarchical latent variable models for partially labeled cross-domain data, addressing non-ignorable selection bias and improving robustness under domain heterogeneity.
  - **Flexible Bayesian Tensor Decomposition for Verbal Autopsy Data [Stats in Medicine]**:
    - Developed hierarchical latent variable models for partially labeled cross-domain data, addressing non-ignorable selection bias and improving robustness under domain heterogeneity.