

# Yue Yu

+1 (734) 882-8778 | [khrisyu@outlook.com](mailto:khrisyu@outlook.com) | [linkedin.com/in/khrisyu](https://linkedin.com/in/khrisyu) | [github.com/khrisyu9](https://github.com/khrisyu9)

## EDUCATION

<b>Indiana University</b>	Bloomington, IN
<i>Ph.D. in Statistical Science &amp; M.S. in Computer Science—Advisor: Prof. David Crandall</i>	<i>Dec. 2026 (Expected)</i>
<b>University of Michigan</b>	Ann Arbor, MI
<i>M.S. in Applied Statistics—Advisor: Prof. Ji Zhu</i>	<i>May 2020</i>
<b>Sun Yat-sen University</b>	Guangzhou, China
<i>B.E. in Theoretical &amp; Applied Mechanics, with Minor in Statistics</i>	<i>Jun. 2018</i>

## SELECTED PUBLICATIONS

<b>On the Limits of Test-Time Compute: Sequential Reward Filtering for Better Inference</b>	<i>Under Review of ICRL 2026</i>
<i>Yue Yu, Qiwei Di, Quanquan Gu, Dongruo Zhou</i>	
<b>Sample and Computationally Efficient Continuous-Time Reinforcement Learning with General Function Approximation</b>	<i>UAI 2025</i>
<i>Runze Zhao*, Yue Yu*, Adams Yiyue Zhu, Chen Yang, Dongruo Zhou</i>	

## SELECTED PROJECTS

<b>Group-Concatenated RL Fine-Tuning for LLMs</b>   <i>Project Leader</i>	<i>Aug. 2025 – Present</i>
• Designed novel RL fine-tuning algorithm improving sample efficiency by 20% for LLM reasoning tasks, reducing training costs while maintaining high accuracy on mathematical and logical reasoning benchmarks	
• Built evaluation pipeline using vLLM and VERL for distributed training, achieving 5% higher accuracy on MATH and improved training stability across 100+ diverse reasoning tasks	
<b>Reward-Filtered Test-Time Compute for LLMs</b>   <i>Project Leader</i>	<i>Mar. 2025 – Sept. 2025</i>
• Developed inference-time optimization reducing LLM serving costs by 30% while improving reasoning accuracy by up to 20%, enabling cost-effective deployment of complex problem-solving AI systems	
• Conducted large-scale evaluation across 10K+ test cases with multiple model families (Qwen3, GPT-5) and verification methods, demonstrating robust performance improvements across various science Q-A benchmarks	
<b>Computationally Efficient Continuous-Time RL</b>   <i>Project Co-Leader</i>	<i>Jul. 2024 – Feb. 2025</i>
• Achieved state-of-the-art sample and computational efficiency (50% improvement) in continuous control and diffusion fine-tuning tasks by developing a model-based continuous-time reinforcement learning (CTRL) algorithm	
• Demonstrated comparable or superior performance with substantially fewer policy updates and rollouts, validated through extensive benchmarking against Guidance, Online PPO, Non-adaptive and Ensemble ODE baselines	
<b>Michigan Pharmacy Fraud Analysis</b>   <i>Lab Project Member</i>	<i>May. 2020 – Dec. 2020</i>
• Performed feature engineering and data pre-processing (e.g., dimension reduction, exploratory data analysis, missing data imputation) and improved traditional ML performance by developing positive–unlabeled learning models such as one-class SVM and Isolation Forest, achieving 0.88 AUC	
• Delivered an engaging and comprehensive presentation on explainable, well-performed statistical models and visualized data for the Office of the Inspector General of the Michigan State Government	

## SELECTED PROFESSIONAL EXPERIENCE

<b>SAS Programmer I</b>	
<i>Bain Institute for Clinical Research</i>	<i>Nov. 2020 – Jun. 2021</i>
• Built automated statistical analysis pipelines processing 10K+ patient records across 5+ cardiovascular trials, reducing analysis time from weeks to days through optimized SAS macro programming	
• Developed GLMs, survival models, and mixed-effects models collaborating with biostatisticians and clinicians for FDA regulatory submission standards	

## TECHNICAL SKILLS

<b>Machine Learning/Artificial Intelligence:</b> PyTorch, vLLM, Transformers, Diffusers, Keras, scikit-learn
<b>Programming Languages:</b> Python, R, C++, SQL, MATLAB, FORTRAN
<b>Tools:</b> Git, NumPy, pandas, Matplotlib, Jupyter, Wandb, L <sup>A</sup> T <sub>E</sub> X