Xinyan Wang
Personal Website: **https://xinyan-wang-stat.github.io/**
Email : **xwang2587@wisc.edu**
Linkedin: **xinyan-wang-34b462206**

Research Interests

LLM post-training, agentic RL, efficient reasoning, safety and red-teaming, representation analysis.

Education

- **University of Wisconsin–Madison** — Madison, WI, USA
  Ph.D. in Statistics; Advisor: Prof. Jun Shao — *Sep 2023 – May 2028 (Expected)*
  M.S. in Computer Science — *Sep 2023 – May 2026*
  M.S. in Statistics — *Sep 2022 – May 2023*
  Visiting Student in Statistics and Computer Science — *Sep 2021 – May 2022*
- **East China Normal University** — Shanghai, China
  B.S. in Statistics — *Sep 2018 – June 2022*

Experience

- **SaFo Lab** — Johns Hopkins University
  Research Intern; Mentor: Prof. Chaowei Xiao — *Aug 2025 – Present*
  - Research on safe and efficient large reasoning models (inference-time overthinking mitigation, robustness).

Research Experience

- **ROM: Real-time Overthinking Mitigation in Large Reasoning Models** — UW–Madison
  First author; advised by Prof. Chaowei Xiao (arXiv 2026)
  - Framed overthinking in large reasoning models (LRMs) as a latent productive-to-redundant transition that surfaces in late-layer hidden states around first-correct-solution (FCS) boundaries.
  - Built ROM, a model-agnostic streaming framework that monitors a frozen LRM with a lightweight hidden-state detector and intervenes at well-formed reasoning boundaries in real time; designed Counterfactual Self-Correction (CSC) to keep useful pre-FCS self-correction while labeling only post-FCS continuation as redundant.
  - Cut response length by ∼24–27% on Qwen3-8B and DeepSeek-R1-Distill-Qwen-32B across MATH500, GSM8K, AIME25, and MMLU-Pro at no accuracy loss; stacked with L1 for ∼21% further token reduction and 46.5% lower wall-clock latency.
- **ReasoningBomb: A Denial-of-Service Attack on Large Reasoning Models** — UW–Madison
  With Prof. Chaowei Xiao et al. (ACM CCS 2026)
  - Formalized prompt-induced inference-time denial-of-service (PI-DoS) for LRMs and proved any practical attack must be high-amplification, stealthy, and optimizable.
  - Built a reinforcement-learning attacker, guided by a constant-time surrogate reward ($4.39 \times 10^5\times$ training speedup), that crafts short, natural prompts driving victim LRMs into pathologically long, often non-terminating reasoning.
  - Across 7 open-source and 3 commercial models, induced 286.7× input-to-output amplification and 6–7× more tokens than benign queries, with bypass rates of 99.8%, 98.7%, and 98.4% against input-, output-, and dual-stage detectors.
- **MLE with Datasets from Populations Having Shared Parameters** — UW–Madison
  Advisor: Prof. Jun Shao (Statistical Theory and Related Fields, 2023)
  - Studied maximum likelihood estimation with multiple datasets from heterogeneous populations sharing parameters; established asymptotic normality, estimated variability via bootstrap, and reduced the estimator's standard deviation by nearly 57% versus single-dataset MLEs.

Publications

Xiaogeng Liu, **Xinyan Wang**, Yechao Zhang, Sanjay Kariyappa, Chong Xiang, Muhao Chen, G. Edward Suh, Chaowei Xiao. *ReasoningBomb: A Stealthy Denial-of-Service Attack by Inducing Pathologically Long Reasoning in Large Reasoning Models.* ACM Conference on Computer and Communications Security (CCS), 2026.

Xinyan Wang, Xiaogeng Liu, Chaowei Xiao. *ROM: Real-time Overthinking Mitigation via Streaming Detection and Intervention.* arXiv:2603.22016, 2026.

Jun Shao, **Xinyan Wang**. *MLE with Datasets from Populations Having Shared Parameters.* Statistical Theory and Related Fields, 7(3):213–222, 2023.

Teaching

STAT 240 (Data Science Modeling I), STAT 309 (Introduction to Probability and Mathematical Statistics I), STAT 340 (Data Science Modeling II), STAT 610 (Statistical Inference), STAT 611 (Statistical Models for Data Science), STAT 613 (Statistical Methods for Data Science).

Professional Service

Reviewer: ACL 2026, ECCV 2026

Skills

Languages: Python, R, SQL
Frameworks: PyTorch, Scikit-learn, Flask