Samantha Pease, Ph.D. (She/Her)

(330) 940-9424 • sam@walking-stick.com • sampease.github.io • github.com/SamPease • linkedin.com/in/sam-pease

SUMMARY

Empirical ML researcher with a Ph.D. in mathematics, focused on AI safety, robustness, and failure analysis. Experienced building and evaluating end-to-end ML systems under uncertainty, including retrieval-augmented generation and Bayesian ranking models.

TECHNICAL SKILLS

Python, PyTorch, PyTorch Geometric, NumPy/SciPy, SQL, Git, Docker, LaTeX
Bayesian modeling, uncertainty estimation, retrieval-augmented generation, empirical safety evaluation, failure analysis
Research focus: agentic oversight, mechanistic interpretability, robustness, privacy-aware system design

EXPERIENCE

Machine Learning Engineer Intern Summer 2024
Covar LLC *Durham, NC*

- Synthesized 20+ research papers on 3D reconstruction and computer vision to identify performance and failure trade-offs; translated theoretical limits into concrete evaluation criteria.
- Built an end-to-end differentiable rendering pipeline (video → 3D Gaussian Splatting) integrating COLMAP SfM and Segment Anything (SAM) across 10+ complex scenes.
- Benchmarked open-source systems and documented **failure modes, edge cases, and robustness gaps** in large-scale video datasets (2K–10K frames/scene).

Mathematics Instructor & Researcher 2017–2025
Rutgers University & Duke University

- Communicated abstract mathematical concepts clearly to diverse audiences; praised for stepwise reasoning and clarity in explaining complex systems.

SELECTED PROJECTS

Safety-Focused RAG Agent 2025

- Designed a retrieval-augmented system to faithfully synthesize community knowledge in a sensitive domain, prioritizing safety, accuracy, and provenance.
- Analyzed and mitigated **LLM refusal and moralizing biases**, improving reliability on frequently suppressed but safety-critical queries.
- Implemented citation-first retrieval using **FAISS IVFPQ** with full execution tracing via **LangSmith** to reduce hallucinations, support auditability, and enable privacy-aware monitoring.

Bayesian Ranking & Uncertainty (BARSElo) 2025

- Developed a Bradley–Terry ranking model with margin-of-victory extensions for 400+ players, explicitly addressing sparse-data regimes.
- Incorporated Bayesian priors and conducted out-of-distribution-style generalization tests for new entrants under partial observability.
- Replaced heuristic ranking rules with a principled **expected head-to-head win probability** metric to improve calibration and interpretability.

EDUCATION

Rutgers University–Newark Ph.D. Mathematics
Defended Fall 2025; Degree conferred May 2026
Research: Langlands Program (Advisor: Chen Wan)

Duke University B.S. Mathematics & Computer Science, with Distinction