Emma Ryla Barnes

New York, NY | +1 (301) 300-8714 | emmabarnsse@gmail.com | LinkedIn | GitHub

EXPERIENCE

Machine Learning Research Scientist, High Assurance Group June 2025 – Present
Two Six Technologies *Promoted from Research Intern, May 2024 – May 2025*

- Researched novel application of inductive logic programming (ILP) to automated reasoning; conducted literature review spanning ILP, answer set programming, and symbolic RL integration; designed experimental framework with benchmarks for reasoning coverage and logical consistency.
- Designed and evaluated LLM integration for hybrid symbolic-neural document understanding, benchmarking against rule-based extraction baselines; achieved 29% improvement in extraction accuracy.
- Extended FST-based knowledge extraction pipeline to capture logical precedence and reasoning chains from linguistic structures, improving inference coverage.
- Researched sliding-window approaches for zero-shot temporal anomaly detection in streaming data; analyzed statistical properties and distribution shift thresholds.
- Implemented hallucination detection model in LLM RAG outputs and integrated with pipeline for 93% detected hallucination accuracy and +12% accuracy in system outputs.
- Extended document modeling to extract tabular and graphical data into Neo4j graph representations; formulated canonical entity resolution as an optimization problem over embedding space with graph clustering constraints.
- Productized research prototypes into production systems through refactoring, documentation, system design specifications, and Git-based workflows.
- Supported formal verification research including SAT/SMT solver integration, contributed to symbolic pipeline bridging informal hardware documentation and logical representations.

Machine Learning Engineer & Consultant November 2023 – Present
Self-Employed

- Computer Vision — Anomaly Detection Pipeline: Designed and built a spatiotemporal autoencoder architecture ground-up for infrastructure fault detection in large-scale AV systems (LED display fault recognition) in security-critical contexts. Defined problem scope and requirements, identified key research, gathered and curated training data, engineered the model architecture, and deployed the pipeline for real-time inference on live AV infrastructure.
- Information Retrieval — RAG-based semantic search system: Built vector embedding pipeline with retrieval-augmented generation for travel recommendation.
- IoT Security: Architected security layer for router firmware implementing OAuth 2.0/JWT authentication, API key rotation, and role-based access control with encrypted channel communication.

Data Science Researcher, Computational Social Science September 2022 – November 2023
The Network Group *Promoted from Data Engineering Intern, March 2022 – August 2022*

- Led quantitative analysis of >1M illicit marketplace reviews to identify behavioral patterns, network structures, and market dynamics; developed statistical models and ML algorithms for pattern detection in underground economy data.
- Applied statistical inference and ML to estimate demand metrics and revenue for opaque economic activity; developed methodology for quantitative intelligence gathering in data-sparse environments.
- Spearheaded development of web scrapes supporting OSINT collection with IP rotation and secure credential management; built ETL processes for large-scale dataset integration into production databases.
- Secured $8,000 in research funding; authored proposal and presented to executive leadership and donor events.

INDEPENDENT RESEARCH

Mechanistic Interpretability in Neural Translation — *Tools: PyTorch, TransformerLens, Gemini*

- Applied sparse autoencoders (SAEs) to decompose transformer activations during translation; analyzed how semantic features map across linguistic boundaries; evaluated layer-wise evolution and identified distinct circuits for syntactic vs. semantic transfer.

Cross-Cultural Emotion Classification in Poetic Texts — *Tools: Python, scikit-learn, spaCy, multilingual BERT*

- Developed emotion detection models across English, French, Russian, and Chinese poetry; demonstrated that cultural variation in emotion encoding requires poetry translation to be approached as creative re-composition rather than semantic transfer.

DISSERTATION

Compiler-Inspired Fault Recovery in Symbolic LLM Verification — *Tools: Python, PyTorch, SymPy, Pydantic, FastAPI, React, Ollama*

University of London BSc Thesis · First Class Honours (1st)

- Built a neurosymbolic verification pipeline applying a three-way compiler-inspired fault taxonomy (syntactic / execution / semantic) to LLM-generated proof steps; deterministic fault classifier attributes verification failures to source components with 94.3% accuracy, enabling asymmetric targeted repair without secondary model calls.
- Typed fault repair recovers 63% of failures vs. 38% for untyped retry on matched budget — largest gains on localised code faults (+35 pp syntactic, +29 pp execution) — directly attributable to fault classification rather than repair mechanism.
- Designed repair trajectories as typed RL rollouts structurally analogous to DeepSeek-Prover-V2's Lean proof-search sequences; intermediate fault signals provide shaped reward without human annotation, with automatic curriculum-compatible difficulty signals derived from attempt count.

EDUCATION

BSc Computer Science — Machine Learning & Artificial Intelligence Graduated September 2025
University of London · First Class Honours

BS Mathematics Graduated May 2022
University of Maryland & Montgomery College

TECHNICAL SKILLS & INDEPENDENT STUDY

Languages & Frameworks: Python, PyTorch, Rust, Haskell, JavaScript/TypeScript, React, FastAPI, MERN stack
ML & Research: Transformers, RL, SAEs, ILP, Answer Set Programming, Neuro-symbolic systems, FSTs, Anomaly Detection, RAG, Vector Embeddings
Infrastructure & Tools: Neo4j, MongoDB, Git, OAuth 2.0/JWT, Docker, Ollama, HuggingFace
Active self-study areas: Mechanistic interpretability (SAEs, feature geometry, circuit analysis) · Logic programming (Prolog, ASP) · Dynamic Systems Theory & Control Theory · Formal methods & verified systems (Rust, Haskell) · Computational linguistics · Computational social science theory

LEADERSHIP

Speaker Series Organizer October 2025 – Present

- Curate and coordinate invited speaker series featuring researchers across interpretable AI, formal methods, quantum computing, network analysis, and neuroscience.
- Coordinate with faculty and research scientists from MIT, Y Combinator, and Dartmouth College; lead technical discussions targeting cross-disciplinary collaboration.

CONFERENCES & CERTIFICATIONS

Presentations: "Computational Analysis of Illicit Marketplace Networks," Human Trafficking National Convening, Washington D.C., 2023
Attended: NVIDIA GTC 2025 (AI & HPC) · HCSS 2025 (Formal Verification & High-Assurance Systems)
Certifications: IBM Data Science Professional (2023) · IBM AI Engineering Professional (2025)