Chase Lambuth

MACHINE LEARNING • APPLIED AI • DATA SCIENCE

M.S. Artificial Intelligence candidate at Johns Hopkins (Dec 2026) with a B.S. in Mathematics. Builds and evaluates end-to-end Python and PyTorch applications such as RAG, computer vision, multimodal learning, and forecasting.

CONTACT

chaselambuth@gmail.com
919-273-4945
https://github.com/chaselambuth
LinkedIn: Chase Lambuth
North Carolina

TECHNICAL SKILLS

PROGRAMMING LANGUAGES
Python · SQL

ML & DEEP LEARNING
PyTorch · scikit-learn · XGBoost
Hugging Face · Transformers · OpenCV
Computer Vision · Semantic Segmentation · Multimodal Learning
Time-Series Forecasting · Reinforcement Learning · Transfer Learning

GENERATIVE AI & NLP
LLMs · RAG · LoRA/PEFT (familiarity)
Prompt Engineering · Tool Calling · Orchestration
BM25 · Dense/Hybrid Retrieval · RRF · Reranking · Embeddings · Semantic Search

ENGINEERING & MLOPS
FastAPI · Flask · Streamlit · REST APIs · OpenAPI
Azure · Azure DevOps · Azure AI Search · Azure OpenAI
Docker · PostgreSQL · Git · pytest · MLflow · AI-assisted development: Codex · Claude · Cursor

DATA & EVALUATION
NumPy · Pandas · Matplotlib · EDA · Data Visualization
Feature Engineering · Statistical Modeling · Experimental Design
Statistics · Linear Algebra · Multivariable Calculus · Optimization
Regression · PCA · Clustering · Calibration · Uncertainty
Cross-Validation · Hyperparameter Tuning · Robustness · Explainability
Excel · Power BI

HONORS & AWARDS

★ Outstanding Mathematics Major
★ Presidential Scholarship
★ Phi Theta Kappa

RELEVANT COURSEWORK

★ Advanced Applied Machine Learning
★ Theory of Machine Learning
★ Deep Learning Using Transformers
★ Applied Generative AI
★ Creating AI-Enabled Systems
★ AI Algorithm Design & Analysis
★ Data Structures & Algorithms
★ Intelligent Systems
★ Statistical Methods
★ Rapid AI Systems: Prototype to Production
★ AI in Biology
★ AI for Robotics

EDUCATION

Johns Hopkins University
M.S. Artificial Intelligence · Expected Dec 2026 · **GPA: 3.96**

University of Mount Olive • 2023–2025
B.S. Mathematics · Summa Cum Laude · **GPA: 4.0**

Wayne Community College • 2019–2023
Associate in Science · Early College High School Program · **GPA: 3.94**

PROJECTS

COMPUTER VISION • EFFICIENT AI • Research Paper ↗

Task-Constrained Token Routing for Vision Transformers

- Compared task-agnostic and task-conditioned token reduction in a PVTv2-style model on seven LoveDA classes and six foreground masks.
- Adaptive task embeddings raised macro target IoU from 0.276 (fixed task-embedding routing) to 0.307. Random pruning led at 0.312 with 75% of Stage 2 tokens retained.

TIME SERIES • FORECASTING • GitHub ↗

NOAA Temperature Anomaly Forecasting

- Built a Flask app for NOAA ingestion, feature engineering, training, and interactive forecasts. Compared Ridge, Huber, XGBoost, LSTM, PatchTST, iTransformer, and NHiTS using chronological splits and horizon-specific RMSE/MAE.
- Seasonal deviations from climatology reduced one-month test RMSE from tuned PatchTST's 1.1705 to 1.1013; final residual MAE was 0.8245.

NLP SYSTEMS • RETRIEVAL ENGINEERING • GitHub ↗

Adaptive RAG Under Technical Domain Shift

- Built a Flask RAG app for two 5,000-document corpora. Compared BM25, dense retrieval, hybrid RRF, and reranking for answer quality and retrieval latency.
- On 894 paired TechQA cases, hybrid RRF raised automatically scored task success from 2.9% with direct generation to 14.3%, including correct abstentions.
- Developed and evaluated a hierarchical routing policy; it converged on static hybrid retrieval, revealing a limitation in query-level adaptation.

AGENTIC AI • SERVICE ORCHESTRATION

Oncology Research Copilot (In Progress)

- Two-person team project developing an Azure assistant for breast cancer and non-small cell lung cancer research and authorized drug-response prediction.
- Integrates document ingestion, hybrid retrieval, and tool orchestration via OpenAPI, with role-based access controls and reproducible research/prediction integration tests.

MODEL ROBUSTNESS • MEDICAL AI • Research Paper ↗

Automated Stress Testing for Medical Imaging Classifiers

- Stress-tested calibrated ResNet-18 under noise, blur, brightness/contrast, JPEG compression, and affine distortions using five search methods, including PPO.
- On 2,048 validation images, Bayesian optimization lowered AUROC from 0.834 to 0.591 and raised high-confidence errors from 2.4% to 12.1%.

MULTIMODAL AI • MEDICAL IMAGING • GitHub ↗

Dual-View Chest X-ray Classification with Clinical Text Fusion

- Combined frontal/lateral X-ray views and radiology reports using PyTorch, dual ResNet-18 encoders, and Bio_ClinicalBERT. Built a Streamlit app, tested preprocessing, and analyzed false-negative predictions using Grad-CAM and confusion-matrix analysis.
- Compared image-only, text-only, and multimodal models. Full-image gated fusion achieved 73.7% accuracy and 0.420 macro F1 on 472 held-out studies.

SECURITY ANALYTICS • AGENTIC AI • GitHub ↗

CyberShield Network Analysis and Agentic Investigation (In Progress)

- Developing Python/PostgreSQL analytics with Power BI and an agentic investigator that selects read-only evidence tools.
- Across four validation days, CatBoost detected 17/21 positive windows with 18 false alerts among 2,464 benign windows.