

Kshitiz Neupane

Hattiesburg, MS | kneupane999@gmail.com | 601-913-8554 | LinkedIn | github.com/kshitiz-tech

EDUCATION

University of Southern Mississippi

B.S. Computer Science & Mathematics GPA: 4.0 / 4.0

– Relevant Coursework: Probability & Statistics, Linear Algebra, Differential Equations, Discrete Mathematics, Data Structures, Operating Systems

Hattiesburg, MS

Expected May 2028

RESEARCH EXPERIENCE

ShiftGuard – Distribution Shift Early-Warning Benchmark | *Independent Research, Python, SciPy, Scikit-Learn*

GitHub

- Designed a reproducible benchmark asking whether label-free input-distribution changes can provide early warning of downstream classifier failure.
- Ran **12,000 model-window evaluations** across 3 datasets, 5 classifier families, 5 seeds, Wasserstein variants, KL divergence, PSI, and MMD under multiple shift mechanisms.
- Maximum marginal Wasserstein achieved **0.799 failure ROC-AUC** (95% CI 0.760–0.832), while analysis showed no metric dominated every dataset or shift type.

Siloed Exact Unlearning (SEU) | *Machine Unlearning, PyTorch, BERT, Transformers*

Private Repo

- Contributed to a machine-unlearning architecture that trains independent bottleneck adapter silos over a frozen BERT backbone so targeted data shards can be removed without complete retraining.
- Helped evaluate IMDb, AG News, and DBpedia against oracle retraining, SISA-style baselines, exactness tests, scalability experiments, and membership-inference analysis.
- Documented low-millisecond deletion (about **1.6 ms** in reported runs), oracle-matching post-unlearning accuracy, and exact retained CPU parameters for unaffected silos.

Proof-Carrying Cognition – Mathematical LLM Revision | *Experimental Design, LLM APIs, Benchmark Auditing*

- Designed a controlled protocol testing whether mathematical LLM revision changes when identical feedback is attributed to self, user, independent verifier, or tool sources.
- Built stateless, resumable experiment infrastructure with strict JSON schemas, controlled provenance, usage/latency logging, reproducible scorers, and **238 passing protocol tests**.
- Audited benchmark integrity before execution, identifying invalid questions, incorrect keys, ambiguous wording, and unusable feedback pairs rather than overstating smoke-test evidence.

Research Assistant – ML, NLP & Cybersecurity

Aug 2024 – Present

University of Southern Mississippi

Hattiesburg, MS

- Developed NLP/graph pipelines for social-engineering anomaly detection using communication structure, message metadata, contextual embeddings, and BERT-based representations.
- Tested sentiment, topic, embedding, and anchor-comparison approaches; recorded failure modes and limitations alongside stronger experimental directions.
- Produced diagnostic analyses and reproducible experiment artifacts to support hypothesis iteration and research communication.

INDUSTRY EXPERIENCE

Software Developer Intern – Butwal Technologies

Feb 2024 – Aug 2024

Software Engineering

Butwal, Nepal

- Developed reusable automated web/API test workflows and technical reports supporting iterative software releases.
- Contributed to REST API integration, debugging, validation, Git collaboration, and CI/CD-oriented engineering workflows.
- Produced technical documentation, communicated implementation progress, and delivered assigned work against project requirements and deadlines.

SELECTED TECHNICAL STUDIES

Bayesian Linear Regression (MAP/posterior inference)

| **PCA** (SVD/maximum variance)

| **LoRA & Attention**

| **F1**

Predictive Modeling (ROC-AUC 0.943)

TECHNICAL SKILLS

Research Areas: Distribution shift, machine unlearning, mathematical LLM evaluation, mechanistic interpretability, NLP/security, representation analysis

ML & Math: PyTorch, Scikit-Learn, NumPy, Pandas, SciPy, Bayesian inference, PCA/SVD, transformers, embeddings, GNNs, probability/statistics

Experimentation: Hypothesis design, benchmark auditing, ablations, bootstrap confidence intervals, calibration, reproducibility, failure analysis, pytest