

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

Research Assistant – ML, NLP & Cybersecurity

University of Southern Mississippi

Aug 2024 – Present

Hattiesburg, MS

- Developed end-to-end pipelines integrating NLP feature extraction with graph-based anomaly detection for organizational communication data.
- Engineered contextual embeddings, sentiment/topic features, email parsing, and metadata representations for social-engineering detection experiments.
- Maintained systematic experiment logs and evaluated both promising and unsuccessful modeling directions, including limitations of anchor-comparison methods.

INDUSTRY EXPERIENCE

Software Developer Intern – Butwal Technologies

Software Engineering

Feb 2024 – Aug 2024

Butwal, Nepal

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

PROJECTS

ShiftGuard – Distribution Shift Early-Warning Benchmark | *Python, SciPy, Scikit-Learn, Wasserstein, MMD*

GitHub

- Designed a reproducible benchmark testing whether label-free distribution distances can warn of classifier degradation under covariate, support, class-prior, and temporal shifts.
- Evaluated **3 datasets, 5 classifier families, 5 seeds, 400 experimental sequences, and 12,000 model-window evaluations** across Wasserstein, KL, PSI, and MMD metrics.
- Maximum marginal Wasserstein reached **0.799 failure ROC-AUC** (95% CI 0.760–0.832); built configuration-driven experiments, bootstrap analysis, tests, and reproducible artifacts.

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

Private Repo

- Contributed to a machine-unlearning framework using independent bottleneck adapter silos over a frozen BERT backbone, enabling shard deletion without full-model retraining.
- Evaluated the method across IMDb, AG News, and DBpedia using oracle retraining, SISA-style comparisons, exactness checks, scalability studies, and membership-inference analysis.
- Verified low-millisecond unlearning (about **1.6 ms** in documented runs) with post-unlearning accuracy matching oracle evaluation and exact retained CPU parameters for unaffected silos.

F1 Race Analysis & Predictive Modeling | *Python, Pandas, Scikit-Learn*

GitHub

- Built leakage-aware ML pipelines from historical race, qualifying, pit-stop, driver, and constructor data while separating pre-race prediction features from post-race explanations.
- Engineered qualifying delta, grid position, and rolling driver/constructor form features and persisted preprocessing-plus-model joblib artifacts for reuse.
- Achieved **ROC-AUC 0.943 / Accuracy 88.7%** for podium classification and **MAE 3.554** for finish-position regression on held-out data.

ADDITIONAL ML WORK

Bayesian Linear Regression (posterior inference, uncertainty) | **PCA on Digits** (SVD, reconstruction) | **LoRA & Self-Attention** | **Neural Network Regression**

TECHNICAL SKILLS

ML & Data: Python, NumPy, Pandas, SciPy, Scikit-Learn, PyTorch, Hugging Face, Jupyter, Matplotlib

Methods: Classification, Regression, Bayesian Inference, PCA/SVD, Transformers, Embeddings, GNNs, Distribution Shift, Uncertainty Quantification

Evaluation & Engineering: ROC-AUC, F1, MAE/RMSE, bootstrap CIs, calibration, pytest, Git/GitHub, Docker, SQL, Linux