

SWETA PATI

USA • +1-571-567-0267 • spati@gmu.edu • [linkedin.com/in/sweta-pati/](https://www.linkedin.com/in/sweta-pati/) • github.com/swetapati22 • [Portfolio](#)

ML and NLP engineer with 4+ years as a Data Scientist, building scalable ML solutions for Fortune 500 clients. Specializes in productionizing ML models, fine-tuning LLMs, and information extraction tasks. Skilled in ML, DL, NLP, Tableau, Gen AI, Agentic AI and RAG, Docker, Kubernetes, AWS. Submitted research on instruction-tuned LLMs for event extraction.

EDUCATION

George Mason University

Aug. 2023 – May 2025

Masters' in Computer Science; GPA: 3.93

Virginia, USA

Kalinga Institute of Industrial Technology

May 2016 – May 2020

Bachelors in Computer Science; GPA: 3.60

Odisha, India

PROFESSIONAL EXPERIENCE

George Mason University

May. 2024 – Present

Summer Research Intern and Research Assistant, Natural Language Processing Lab

Virginia, USA

- Submitted research paper [Instruction-Tuning LLMs for Event Extraction with Annotation Guidelines](#) to **ACL 2025**.
- Boosted F1-score by 10% for EE tasks by synthesizing guidelines for 500+ events, 4000+ arguments and fine-tuning LLaMA-3.1 8B with LoRA and regularization techniques, using a custom-designed evaluation framework.
- Optimized evaluation by transforming text into code prompts, reducing hallucination via implicit prompt optimization.

Highradius Technologies Private Limited

Sep. 2020 – Aug. 2023

Data Scientist

Odisha, India

- Increased dollar recovery rate by 12% implementing imbalanced classification ML model (DVP) for fraud detection with 85% average accuracy for 10+ fortune 500 clients.
- Increased sales acceleration and reduced business disruptions by 20% through the optimization of credit decisioning for order release using predictive models.
- Reduced Days Sales Outstanding (DSO) by 30% by building a regression model for Payment Date Prediction (PDP).
- Accelerated DVP and PDP deployment by 48% and 57%, automating the model building and deployment pipelines.
- Productionized ML models for 10+ Fortune 500 clients like Nestle, Kraft Heinz, Clorox, Uber, Kellogg's, etc.
- Enhanced stakeholder decision-making by visualizing model insights and business impact using Tableau.
- Mentored over 15 interns in data science and business concepts during the 2021 and 2022 internship programs.

Research Intern

Jun. 2018 – Jun. 2018

Indian Institute of Technology, Kharagpur

West Bengal, India

- Implemented BCH error correction code and compared existing test suites in OpenSSL with a proposed authentication protocol. Published research paper [An OpenSSL Extension for PUF based Authentication](#) in **IEEE DSP 2018**.

PROJECTS

Medi-Buddy: Multi-Agent Framework for Medical Insights ([Source Code](#))

Tech Stack: Python, LLaMA-3.3-70B, Phidata, Groq Inference API, PubMed API, Wikipedia API, DuckDuckGo, GitHub

- Developed a multi-agent framework using LLaMA-3.3 70B to extract and summarize evidence-based medical insights from PubMed and Wikipedia.

Authentic Counterhate Argument Detection (EMNLP 2023 Implementation). ([Source Code](#)) | ([Presentation](#))

Tech Stack: Python, RoBERTa, Longformer, XLM-RoBERTa, Hugging Face, Transformers, Google Translate API, GitHub, Shell

- Implemented EMNLP paper using **RoBERTa** and **Longformer** for classifying authentic counterhate arguments.
- Extended the paper by training multilingual models with **XLM-RoBERTa** on English, Japanese, and Spanish data by automated complete NLP pipeline with robust evaluation and error analysis.

TECHNICAL SKILLS

Programming Languages: Python, SQL, PyTorch, Numpy, Pandas, Scikit-learn, PySpark

Large Language Models: LLMs (GPT, LLaMA, BERT), Fine-Tuning, PEFT, Instruction Tuning, Prompt Engineering, RAG, Agentic AI, LangChain, Vector DBs

NLP and Deep Learning: Transformers, Event Extraction, Neural Networks, RNN, CNN, Computer Vision

Data Science and Machine Learning: Statistics, Supervised, Unsupervised, Classification, Regression, Ensemble, Bagging, Boosting, SVM, Random Forest, A/B Testing, Feature Engineering, Feature Selection, Model Monitoring

MLOps and Deployment: CI/CD, Docker, Kubernetes, Terraform, AWS, Google Cloud (GCP), Weights & Biases, MLflow

Data Visualization: Data Mining, Data Analysis, Tableau, Seaborn, Plotly, Matplotlib, Streamlit

Tools and Platforms: GitHub, Jira, Hugging Face, OpenAI API, Groq API, HPC Environments

AWARDS AND CERTIFICATIONS

- Google AI Essentials**, Google, [[Credential](#)], Mar 2025
- Event-Driven Agentic Document Workflows**, Deep Learning.AI and LlamaIndex, [[Credential](#)], Mar 2025
- Oracle Cloud Infrastructure 2024 Generative AI Certified Professional**, Oracle, [[Credential](#)], Aug 2024