

# Kabiir Krishna

Waterloo, Ontario | +1 (437) 669 6092 | [k7krishn@uwaterloo.ca](mailto:k7krishn@uwaterloo.ca) | [Website](#) | **Work Permit Approved**

## SKILLS

- **Programming Languages:** Python, SQL
- **Tools & Frameworks:** LangChain, Ollama, Streamlit, PyTorch, HuggingFace, Scikit-learn, PowerBI, Tableau, MySQL, Databricks (Snowflake), Git, Docker, Kubernetes, Apache Kafka, **PySpark**, JIRA, AWS (VPC, EC2, S3, CloudWatch, Lambda)
- **LLM/NLP:** LLM Fine-tuning (RLHF), Prompt Engineering, Sentiment Analysis
- **Skills:** Data Engineering, Data Science, AI/ML Forecasting, Feature Engineering, Data Cleaning, Data Visualization, Microservice Orchestration, Agile Software Development
- **Certifications:** AWS Cloud 101, Databricks Fundamentals, Data Analysis with Databricks, Data ingestion with Databricks

## PROJECTS

**End-to-end RLHF pipeline for safe response optimization using prompt pairs, PPO, and reward modeling with PyTorch & HuggingFace.** *PyTorch | LangChain | HuggingFace* ([Repository](#))

**Custom GPT-style language model for next-word prediction using Bigram & Transformer (GPT) architectures using PyTorch.** *PyTorch* ([Repository](#))

**Proof-of-concept - enhancing CERN's CLI workflow with LLM-powered task automation by integrating Microsoft Phi-2 into Ganga CLI using PyTorch.** *PyTorch | Microsoft Phi-2 | [Ganga CLI](#)* ([Proof of concept](#))

**Sentiment dashboard for IMDb reviews using LSTM trained on weak/strong labels (VADER & DistilBERT) with Streamlit UI.** *PyTorch | Streamlit | DistilBERT | VADER* ([Repository](#))

**Cricket Data ETL and performance analytics pipeline using PySpark, SQL & Seaborn on Databricks/AWS to support match strategy planning.** *PySpark (Apache Spark) | Databricks (Snowflake) | AWS | Seaborn (Matplotlib) | SQL* ([Repository](#))

## EXPERIENCE

**LLM Trainer** Nov'24 – Present  
Outlier.AI, Canada

- Fine-tune LLMs (GPT-X etc.) using Reinforcement Learning from Human Feedback to improve response relevance for coding tasks by scoring base model's response w.r.t accuracy, truthfulness etc. & writing new responses to tune Reward Model.
- Hands-on experience training agents via demonstration and feedback across real-world web-based workflows, annotating step-wise actions to improve training.

**Full Time, Cyber Security Specialist – One Cyber** Jul'21 – Aug'23  
PricewaterhouseCoopers (PwC)

- Developed NLP pipelines for automated threat analysis of security telemetry logs (>1 million rows) through anomaly detection (LSTMs and Random Forest classifiers) with an F1-score around 0.9.
- Developed and deployed interactive dashboards (Power BI, Tableau) to visualize real-time data security policy violations, offering actionable insights & leading to reduction in enterprise-wide cyber risks by ~25% (number of ISO 27001 controls).

**Co-Op, SWE - DevOps & Data-Engineering** Apr'20 – Jun'20  
IBM

- Built scalable ETL pipelines and microservices (Docker, Kafka, Kubernetes) for real-time incident reporting.
- Created & annotated custom video datasets for training, leading to better mean average precision for models (TinyYOLO etc).

**Co-Op, ML - Data Analyst** May'19 – Jul'19  
Deloitte (Touche Tohmatsu LLC)

- Developed sentiment & trend analysis engine for scraped (selenium) FMCG data using VADER, PySpark, and Sklearn.
- Built Power BI dashboards to visualize key metrics & deliver actionable insights to drive data-driven strategies.

## EDUCATION

**Master of Engineering**, *University of Waterloo, Waterloo, Ontario* Sep'23 – Oct'24  
M.Eng in Electrical and Computer Engineering (specialization in Software Engineering); CGPA 3.9/4

**Bachelors of Technology**, *Vellore Institute of Technology, Vellore, Tamil Nadu* Jul'17 – Jul'21  
B.Tech in Electronics & Comm. Engineering (specialization in IoT & Sensors); CGPA 8.39/10