Pradeep Sundaram Lead MLOps Engineer

Purchase, NY · pradeep.sundaram@gmail.com · +1 914-555-0143 · linkedin.com/in/pradeepsundaram

PROFILE SUMMARY

- **Lead MLOps Engineer** with **11 years** in ML engineering and platforms at **Mastercard** and BlackRock, running the **model platform for fraud and risk scoring** in a regulated payments business.
- Works with **Python** and Scala, **Kubeflow Pipelines** and Airflow, **MLflow**, **Apache Spark** and Feast, KServe on **Kubernetes**, Terraform, Argo CD, and Prometheus with Evidently.
- Strongest in **model risk controls**, **feature store design**, **real-time scoring** and **audit-ready lineage**, with approval steps built into the pipeline instead of kept in a spreadsheet.
- Works with **data science leads**, model risk management, security and the payment network teams, and presents platform health and cost to the **engineering director** every month.
- Leads **7 MLOps engineers**, sets the platform roadmap with the director, hires for the team, and chairs the **production readiness review** that every new model goes through.

TECHNICAL SKILLS

Languages & Scripting:	Python, Scala, Bash, SQL
ML Platforms & Tracking:	MLflow, Kubeflow, Databricks, Amazon SageMaker
Pipelines & Orchestration:	Kubeflow Pipelines, Apache Airflow, Argo Workflows
Model Serving:	KServe, NVIDIA Triton Inference Server, ONNX Runtime
Feature & Data Tools:	Feast, Apache Spark, Apache Kafka, Delta Lake, Great Expectations
Cloud, Containers & CI/CD:	Kubernetes, Red Hat OpenShift, AWS, Terraform, Argo CD, Jenkins
Observability & Drift:	Prometheus, Grafana, Evidently, Splunk

EDUCATION

New York University, Courant Institute — **M.S. in Computer Science** — New York, NY · Sep 2013 to May 2015

WORK EXPERIENCE

Mastercard — **Lead MLOps Engineer** — Purchase, NY · May 2021 to Present

- Lead the **7-engineer** team running the model platform behind fraud and risk scoring, which serves **60** production models for 9 data science teams on Kubernetes across 2 data centers and AWS.
- Built approval steps into **Kubeflow Pipelines** and the **MLflow** registry, so every release carries its validation report and model risk sign-off, cutting audit evidence prep from 6 weeks to **4 days**.
- Launched a **Feast** feature store over Spark and Kafka with matching online and offline features, which ended training and scoring mismatches and cut new feature delivery from 5 weeks to **8 days**.
- Moved real-time fraud scoring to **KServe** with Triton and blue-green releases, holding p99 latency under **25 ms** through holiday peaks while release frequency went from quarterly to every 2 weeks.
- Rolled out drift and stability monitoring with **Evidently** on every scoring model, with alerts routed to model owners, catching **11** shifts in merchant or device data before scores degraded.

BlackRock — **Senior Machine Learning Engineer** — New York, NY · Jul 2015 to Apr 2021

- Moved **45** portfolio risk and NLP models from on-prem batch servers to containers on Kubernetes with Jenkins builds, cutting deploy time from 2 weeks to **1 day**.
- Set up the first **MLflow** tracking server for the research teams, so model runs, parameters and data versions were logged for **all** models going to production.
- Rebuilt nightly feature jobs in **Apache Spark** with data checks before each run, cutting failed overnight model runs from about 8 a month to **1**.