

Hana Kobayashi-Reid

Machine Learning Engineer II | Machine Learning Engineer

CONTACT

hana.kobayashireid@email.example
(514) 555-0153
Montréal, QC

SKILLS

- Python
- Go
- PyTorch
- XGBoost
- Feast
- Kafka
- Spark
- Kubernetes
- BentoML
- MLflow
- Airflow
- Prometheus and Grafana
- Vertex AI
- Bilingual (English and French)

EDUCATION

BEng, Software Engineering
École Polytechnique des Rivières
2021

CERTIFICATIONS

· Google Cloud Professional Machine Learning Engineer

PROFESSIONAL SUMMARY

Machine learning engineer with five years building real-time scoring systems for consumer lending and insurance. Owns the path from feature pipelines to online inference: a Feast feature store, models served on Kubernetes, shadow deployments and drift alerts. Works in Python and Go, partners with data scientists on what to ship, and keeps on-call quiet by testing models like software.

PROFESSIONAL EXPERIENCE

Machine Learning Engineer II

Apr 2023 — Present

Saint-Laurent Lending · Montréal, QC

- Built the online feature pipeline for credit decisions on Kafka and Feast, giving underwriting models fresh transaction features within a few seconds of each event.
- Moved model serving from a single Flask monolith to BentoML services on Kubernetes, holding p99 latency under 120 milliseconds during month-end application peaks.
- Introduced shadow and canary releases for every new credit model, catching a feature mismatch between training and serving before it touched a single applicant.
- Added drift and calibration dashboards in Prometheus and Grafana, with alerts routed to the owning model team rather than the general platform pager.

Machine Learning Engineer

Jun 2021 — Mar 2023

Fleuve Assurance Numérique · Québec City, QC

- Productionized a claims fraud model built by the data science team, rewriting notebook code into a tested Python package trained on versioned data snapshots.
- Scheduled retraining in Airflow with an automatic comparison against the live model, promoting a challenger only when it beat the incumbent on a frozen holdout.
- Cut nightly batch scoring from five hours to forty minutes by moving feature joins into Spark and caching static policy attributes.