KUSH PATEL

AML / Financial Crime Analyst | Transaction Monitoring · KYC / CDD · Investigations · FINTRAC Reporting | SQL · Python · Power BI

Vancouver, BC +1 (778) 637-1602 dharma.patel552@gmail.com linkedin.com/in/kush-patel-48885719b github.com/KushPatel29
kushpatel29.github.io

Open to remote and relocation across Canada · Legally eligible to work in Canada · Available immediately

SUMMARY

- Finance and controls analyst transitioning toward AML analytics, with 4+ years of paid reconciliation, risk reporting, data-quality and access-control work across financial reporting and distribution operations.
- Built synthetic transaction monitoring, investigator briefs and KYC/CDD workflows, plus measured screening against a real public sanctions list. Alerts prioritise review; the demonstration records no completed dispositions or regulatory filings.

SKILLS

Financial Crime (project methods): KYC/CDD (customer risk rating, EDD triggers, PEP handling, beneficial ownership at 25% direct and indirect), sanctions name screening (fuzzy matching, transliteration, threshold tuning), structuring and smurfing typologies, rule design and tuning, alert triage and prioritisation, investigation case files, analytic briefs, PCMLTFA and FINTRAC reporting (LCTR including the 24-hour rule, EFTR, suspicious transaction report timing), threshold versus suspicion

Detection Analytics: SQL window functions for velocity and aggregation rules, entity-relative z-scores, anomaly detection (IsolationForest, Mahalanobis distance, gradient boosting), precision and recall by typology, cost curves for alert volume

Controls & Audit: reconciliation and control totals, journal-entry testing, segregation of duties, exception logs, evidence integrity, RLS/RBAC and column masking

Tools: SQL (T-SQL, SQLite), Python (pandas, scikit-learn), Power BI, Excel, Streamlit, Tableau, Git; Kafka streaming

PROFESSIONAL EXPERIENCE

Two Rivers Specialty Meats

Dec 2023 – Jun 2026 · North Vancouver, BC

Data Analyst — Operations, Logistics & Enterprise Reporting

- Built source-to-mart GL/P&L and inventory controls, cutting discrepancies an estimated ~45% and reporting errors ~30% (internal directional estimates, not audited).
- Automated schema, completeness, uniqueness and referential-integrity checks, with anomaly thresholds, exception logs and pipeline alerts.
- Delivered AP/AR aging, invoice-match exceptions and credit-exposure reporting for Finance.
- Secured self-service reporting with RLS/RBAC and column masking; implemented FEFO and lot traceability for recall readiness.
- As sole analyst, integrated ERP, Square POS, workforce and REST feeds into Bronze/Silver/Gold datasets using Microsoft Fabric, Azure Data Factory/Synapse, SQL and PySpark.

Shivam Investments

Dec 2020 – Jul 2022 · Vadodara, India (hybrid)

Financial Analyst — Data & Reporting

- Gathered requirements from portfolio and risk stakeholders; delivered Power BI dashboards with auditable measure definitions and reconciled GL/P&L figures across sources.
- Automated recurring preparation with Python and SQL views, replacing spreadsheet consolidation and cutting manual effort an estimated ~40% (internal directional estimate).

SELECTED PROJECTS — public code; all entities and transactions are synthetic

Transaction Monitoring — SQL, Python, scikit-learn, SQLite, Streamlit · synthetic data

github.com/KushPatel29/aml-transaction-monitoring · live: aml-transaction-monitoring.streamlit.app

Layered detection over 100,299 transactions and 1,500 entities: five explainable rules plus an anomaly model find 60 of 60 planted cases where a flat $10,000 rule finds 38, alerting on 366 entities instead of 555. The alerts become 366 controlled investigations, 78 of them P0 with a four-hour review clock, worked against five typology hypotheses that each name a lawful lookalike and the evidence that tells them apart. Scores prioritise evidence; they never decide suspicion, and the record shows zero dispositions and zero filings.

FINTRAC mapping and an analytic brief — PCMLTFA, LCTR, EFTR, STR · synthetic case

regulatory mapping · analytic brief INV-ENT-01216

Threshold reports versus suspicion: each rule mapped to FINTRAC's indicator categories and the report it relates to, including why a $10,000 threshold rule cannot see structuring. The brief works one case: six cash deposits on six consecutive days, each just under $10,000, totalling $53,374.51 in 4.9 days, nine times the subject's usual weekly cash; assessed as likely sized to avoid the threshold (moderate confidence), with the three facts needed before a reporting decision.

Sanctions screening on Canada's real list — Python, rapidfuzz · real public list, synthetic customers

github.com/KushPatel29/aml-transaction-monitoring

Name screening that is measured, not assumed: customer names screened against the Consolidated Canadian Autonomous Sanctions List (5,708 entries, 9,604 names with aliases, Cyrillic transliterated). The alert threshold is set on one half of the list by a rule fixed in advance and reported on the other: in October 2026 it caught 97.2% of the name variants customers type, at 10 alerts per 10,000 synthetic customers. The list is downloaded at run time and never stored.

KYC: customer risk rating and beneficial ownership — Python, pandas · synthetic customers

github.com/KushPatel29/aml-transaction-monitoring

Who stands behind the account: 1,500 customers rated at onboarding and at a yearly refresh on stated factors, each rating with its reasons; a foreign PEP is high risk whatever the score, and every high-risk client is flagged for enhanced due diligence. For 450 businesses, beneficial owners traced through holding companies by multiplying shares down every chain to the 25% line, with opaque jurisdictions, loops and no-25%-owner cases flagged. Only 11 of the 60 planted-case customers rated high: a rating decides scrutiny, not suspicion.

GL / P&L Reconciliation — T-SQL, SQLite, Power BI, Tableau · synthetic data

github.com/KushPatel29/gl-reconciliation-dashboard

Controls that refuse to fake a sign-off: a close-certification engine tests all 5,400 journal entries and enforces segregation of duties; at month-end 8 of 60 account-periods certify, 9 go to a named reviewer and 43 are blocked.

A PIPEDA-structured project assessment classifies 105 columns across 14 files, tests masked outputs and identifies organizational controls required before real-client use. GL close and streaming-recovery evidence: kushpatel29.github.io.

EDUCATION

Northeastern University — Master of Professional Studies in Analytics (Applied Machine Intelligence), GPA 3.76/4.0 · Vancouver, BC — **Sep 2022 – Jun 2024**

Gujarat Technological University — Bachelor of Engineering, Computer Science · India — **Aug 2017 – Aug 2021**

CERTIFICATIONS

Google Data Analytics Professional Certificate · UC Davis SQL for Data Science · IBM Python for Data Science, AI & Development · Databricks Fundamentals.