

ALEKSANDR KOLMAKOV — ML / DATA ENGINEER

Porto, Portugal | +351 932 125 222 | aleksandr.kolmakov@pm.me | [LinkedIn](#)

SUMMARY

ML/LLM + data engineer — 2yr LLM production + 6yr ML + 9yr engineering. Specialist in LLM agents, MLOps, and modern data platforms (BigQuery, dbt, ClickHouse, Kafka).

CORE SKILLS

- **Data stack:** Python, SQL, BigQuery, ClickHouse, Snowflake, DuckDB, PostgreSQL, dbt, Airflow, Mage, Dagster, dlt
- **Processing:** Spark, Apache Beam, Kafka, Redpanda, Kinesis, Parquet, Delta Lake
- **ML & LLM:** PyTorch, MLflow, DVC, LLM agents (MCP), LLM-as-a-judge, trajectory evaluations
- **Infra:** GCP, AWS (EC2/Kinesis/Redshift), Docker, Kubernetes, Terraform, CloudFormation
- **Observability:** Prometheus, Grafana, Phoenix Arize (LLM), OpenTelemetry, Loki
- **Security:** CVE remediation, enterprise SSO (WorkOS/SAML), FIDO2/WebAuthn

EXPERIENCE

Senior Backend Engineer (B2B contract ends Oct 2026) | Termius | SF-based, remote | 2026

- Deleted 70M+ legacy rows and freed 140 GB on a live production MySQL database with zero downtime via an adaptive-window Celery cleanup pipeline — kill-switch controlled and throughput-validated on a prod replica before enabling in production
- Remediated 10 critical/high CVEs (incl. OpenSSL) across production ECR images via zero-downtime ECS redeploy, with auditable security-dated image tags traced to AWS Inspector findings
- Restored a broken Terraform deploy pipeline after GitHub Actions' Node deprecation — root-caused two native-library failures (libexpat symbol mismatch) and upgraded Terraform 0.14 → 1.12

LLM / Data Engineer | Helloworld | Porto, Portugal | 2024 – 2026

- Architected and implemented an end-to-end data platform (dbt, dlt, Mage), processing 200,000+ daily API requests into ClickHouse, with embedded PII anonymization → +5% company revenue
- Transitioned from 140+ static dashboards to an LLM-powered analytics agent based on ClickHouse MCP, enabling natural-language reporting → 40% active-user adoption for ad-hoc queries, unlocking 10–12 senior developer hours per week
- Engineered an LLM observability platform (Prometheus, Grafana, Phoenix Arize, OpenTelemetry): resource/load monitoring, prompt versioning, prompt tracing, output evaluation via trajectory evaluation and LLM-as-a-judge

Senior ML Engineer | Sensemitter | Dubai, UAE | 2023 – 2024

- Moved ML code from Celery tasks to Dataflow; optimized emotion storage in Delta Lake parquets for artifact versioning and time travel
- Improved a PyTorch ML pipeline by transitioning to 3D gaze interpolation, removing interpolation inconsistency
- Upgraded the BigQuery pipeline to dbt, cutting processing costs 15%; robust data-quality checks prevented wasteful ML runs, increasing model reliability

Senior ML Engineer (Reinforcement Learning) | Arloid LLC | Dubai, UAE | 2020 – 2023

- Developed a proof-of-concept with DQN and Q-learning; later optimized the algorithm with PCA dimensionality reduction
- Drove a 12% client energy reduction leading an 8-engineer team delivering an AI SaaS platform; owned data architecture design → production
- Cut monthly AWS costs 30% migrating ML workloads from SageMaker to open-source MLOps (MLflow/DVC) with CloudFormation
- Achieved 80% accuracy in automated fault detection (data pipeline + ML models on Snowflake)

Backend Developer | Distillery | St. Petersburg, Russia | 2017 – 2020

- Engineered sub-second data delivery: <0.3s complex queries via a Neo4j migration; <1s media downloads via AWS HLS streaming pipelines
- As the sole backend engineer in a 17-person team, took a legacy project to production: CI/CD, test coverage, Blue/Green deployments on AWS from scratch

I hereby authorise the processing of my personal data in accordance with GDPR EU 2016/679. PADI Divemaster · Reef conservation specialist