

# Machine Learning Engineer (Contract)

**Onshore – Toronto, Canada | Hybrid | Start: Early February**

At Budhhi, ML Engineering is about more than deploying models.

It's about building intelligent systems that scale, adapt, and deliver real impact.

We are looking for a Senior Machine Learning Engineer to design, build, and operate production-grade ML and GenAI systems across Budhhi's internal platforms and client-facing solutions. This role is for engineers who are comfortable owning the full ML lifecycle — from training pipelines to inference, monitoring, and platform-level thinking.

## **What You'll Do**

### *Build End-to-End ML Systems*

- Design and implement training pipelines, feature pipelines, and orchestration workflows
- Build real-time and batch inference systems for ML and GenAI workloads
- Integrate LLMs, RAG pipelines, and agentic AI patterns into production systems

### *Engineer for Scale & Reliability*

- Implement CI/CD for ML including versioning, rollbacks, and safe deployments
- Monitor models for performance, drift, and regressions
- Design systems with security, compliance, and cost awareness

### *Platform & Product Ownership*

- Contribute to shared ML and GenAI platforms, tooling, and infrastructure
- Partner with Product Managers to embed ML into user-facing features
- Work closely with Data Scientists and AI Architects to operationalize models

## *Lead by Doing*

- Write clean, testable, production-quality code
- Review and elevate code quality across teams
- Help define and enforce engineering and MLOps best practices

## **What We're Looking For**

### *Experience & Seniority*

- 8–12+ years in ML engineering, software engineering, or applied AI
- Proven experience deploying ML systems into production
- Comfortable owning systems end-to-end in high-ambiguity environments

### *Core ML & AI Expertise (Non-Negotiable)*

- Hands-on experience with:
- Training pipelines and orchestration
- Model serving and inference (batch & real-time)
- LLMs, GenAI, RAG, and agentic AI patterns
- Classical ML and deep learning
- Strong understanding of model evaluation and monitoring
- Strong experience in building applications with best practices of Predictable AI, Generative AI, Explainable AI and Ethical AI

### *Platform & Infrastructure Strength*

- Experience with MLOps (CI/CD for ML, versioning, monitoring)
- Cloud-native ML experience (Azure, AWS, GCP)
- Familiarity with data infrastructure (event streams, feature stores, vector DBs)
- Exposure to IaC (Terraform, Bicep, or equivalent)

## *Engineering Excellence*

- Strong Python skills
- Solid backend engineering fundamentals (APIs, services)
- Commitment to testing, observability, and maintainability

## **Contract Details**

- Engagement Type: Contract (extendable based on performance)
- Location: Toronto, ON (Hybrid)
- Start Date: Early February
- Compensation: Market-competitive, premium for the right talent
- Future: Long-term partnership and extensions encouraged

## **Why Budhhi?**

- Build real ML and GenAI systems at scale
- Work across internal IP and enterprise client platforms
- Collaborate with senior AI Architects and Product leaders
- High-ownership culture with minimal bureaucracy
- Focus on clarity, engineering rigor, and impact

## **Ready to Engineer Intelligence?**

If you enjoy building production-grade ML systems, thinking in platforms, and shipping meaningful AI — we want to talk.

Apply now and help power the next generation of intelligent systems at Budhhi.