TORONTO
AI SCHOOL

Toronto AI School

Master Program in Cloud AI & MLOps Engineering

Week 1 – Foundations of Cloud AI Systems

Session	Topics Covered
Day 1 – Cloud AI Landscape	- Evolution of cloud computing for AI - Overview of AWS, Azure, and Google Cloud AI ecosystems - Cloud-native architecture for AI systems - Containerization fundamentals with Docker - Microservices architecture for machine learning - Case Study: Netflix cloud infrastructure - Case Study: Uber real-time ML infrastructure
Day 2 – Data Infrastructure for AI	- Data pipelines for machine learning - Streaming data platforms - Data lake vs data warehouse - Feature engineering pipelines - Data versioning strategies - Case Study: Amazon recommendation system data pipelines - Case Study: Shopify customer analytics pipelines
Day 3 – Cloud Storage for ML	- Object storage architecture - Distributed storage systems - Data governance in cloud environments - Data security for AI systems - Large scale dataset management - Case Study: Tesla autonomous driving datasets - Case Study: Spotify music recommendation data architecture

Week 2 – Machine Learning Infrastructure

Session	Topics Covered

Day 1 – ML Infrastructure Design	- Designing scalable ML infrastructure - Training pipelines architecture - GPU infrastructure for ML training - Experiment tracking systems - Feature store architecture - Case Study: OpenAI model training infrastructure - Case Study: Meta AI research infrastructure
Day 2 – ML Data Engineering	- Batch vs real-time ML pipelines - Data preprocessing automation - ETL pipelines for AI systems - Workflow orchestration - ML dataset quality monitoring - Case Study: LinkedIn data pipeline architecture - Case Study: DoorDash delivery optimization models
Day 3 – Infrastructure Automation	- Infrastructure as code - Kubernetes basics for ML workloads - CI/CD pipelines for ML systems - Model packaging strategies - Cloud deployment pipelines - Case Study: Airbnb ML infrastructure - Case Study: Twitter recommendation model deployment

Week 3 – MLOps Fundamentals

Session	Topics Covered
Day 1 – ML Lifecycle Management	- Machine learning lifecycle overview - Model experimentation workflows - Model version control - Reproducibility in ML pipelines - Experiment tracking tools - Case Study: Google ML pipeline architecture - Case Study: Netflix recommendation ML lifecycle
Day 2 – Model Deployment	- Batch vs real-time model serving - Model APIs and inference pipelines - Edge AI deployment - Scalable model serving - Latency optimization strategies - Case Study: Amazon Alexa ML serving architecture - Case Study: Tesla real-time inference pipelines
Day 3 – Monitoring ML Systems	- Model performance monitoring - Data drift detection - Model drift monitoring - ML observability frameworks - Incident response for ML failures - Case Study: Uber surge pricing ML monitoring - Case Study: Spotify ML monitoring systems

Week 4 – Cloud Platforms for AI

Session	Topics Covered

Day 1 – AWS AI Services	- AWS SageMaker architecture - Serverless AI deployment - AWS data pipelines - Scaling ML workloads - AWS monitoring tools - Case Study: Amazon product recommendation ML - Case Study: Zillow property valuation ML
Day 2 – Google Cloud AI	- Vertex AI platform overview - AutoML systems - BigQuery ML - Dataflow pipelines - AI model training on GCP - Case Study: Waymo autonomous systems - Case Study: YouTube recommendation ML
Day 3 – Azure AI	- Azure Machine Learning services - Azure AI pipelines - MLOps using Azure DevOps - Enterprise AI architecture - AI governance in Azure - Case Study: Microsoft Copilot architecture - Case Study: LinkedIn AI talent matching

Week 5 – Distributed Machine Learning

Session	Topics Covered
Day 1 – Distributed Training	- Parallel training strategies - Distributed GPUs - Model sharding - Parameter servers - Training optimization - Case Study: GPT model training - Case Study: Meta LLaMA model training
Day 2 – Scalable Data Processing	- Apache Spark for ML - Distributed data processing - Streaming ML pipelines - Large-scale feature engineering - Data scaling strategies - Case Study: Uber data platform - Case Study: Airbnb data processing systems
Day 3 – High Performance ML	- Hardware acceleration - GPU optimization - ML performance tuning - Scaling training workloads - Memory optimization - Case Study: NVIDIA AI infrastructure - Case Study: Tesla training supercomputers

Week 6 – AI Model Deployment

Session	Topics Covered

Day 1 – API-Based AI Systems	- REST APIs for ML models - GraphQL APIs for AI services - Model serving architecture - AI microservices - Containerized model deployment - Case Study: OpenAI API architecture - Case Study: Stripe fraud detection API
Day 2 – Edge AI Systems	- Edge inference systems - IoT AI deployment - Edge ML pipelines - Model optimization for devices - Edge security considerations - Case Study: Tesla vehicle AI deployment - Case Study: Apple on-device AI
Day 3 – Global AI Infrastructure	- Global load balancing - Multi-region ML deployments - AI service reliability - Fault tolerant ML systems - Disaster recovery strategies - Case Study: Google global AI infrastructure - Case Study: Amazon global cloud systems

Week 7 – Data Engineering for AI

Session	Topics Covered
Day 1 – Data Pipelines	- ETL architecture for ML - Real-time streaming pipelines - Data ingestion frameworks - Data validation pipelines - Pipeline orchestration - Case Study: Snowflake data architecture - Case Study: Shopify analytics pipelines
Day 2 – Feature Stores	- Feature engineering pipelines - Feature storage architecture - Online vs offline features - Feature store monitoring - Feature version control - Case Study: Uber Michelangelo feature store - Case Study: Airbnb feature engineering platform
Day 3 – Data Governance	- Data privacy regulations - Data compliance in Canada - Secure data pipelines - Data lineage tracking - AI data ethics - Case Study: TD Bank data governance - Case Study: Canadian healthcare data systems

Week 8 – AI System Security

Session	Topics Covered

Day 1 – Cloud Security for ML	- Securing ML pipelines - Cloud IAM architecture - Data encryption strategies - Secure model deployment - Security monitoring - Case Study: Capital One cloud breach lessons - Case Study: AWS security architecture
Day 2 – AI Security Threats	- Model theft attacks - Data poisoning - Adversarial ML attacks - Prompt injection threats - AI risk mitigation - Case Study: AI chatbot security failures - Case Study: enterprise AI risk management
Day 3 – Compliance and Governance	- Responsible AI frameworks - Enterprise AI governance - Model auditing systems - AI compliance strategies - Risk management frameworks - Case Study: healthcare AI compliance - Case Study: banking AI compliance

Week 9 – AI System Optimization

Session	Topics Covered
Day 1 – Model Optimization	- Model compression techniques - Quantization - Pruning strategies - Inference optimization - Latency reduction - Case Study: mobile AI optimization - Case Study: autonomous driving inference
Day 2 – Performance Monitoring	- ML observability systems - System performance dashboards - Monitoring ML pipelines - Resource optimization - AI cost optimization - Case Study: Netflix infrastructure monitoring - Case Study: AWS cost optimization strategies
Day 3 – AI Reliability	- High availability AI systems - Failover architecture - Service resilience - ML disaster recovery - System reliability engineering - Case Study: Google reliability engineering - Case Study: Amazon infrastructure resilience

Week 10 – Capstone Cloud AI Projects

Session	Topics Covered

Day 1 – AI Platform Design	- Design enterprise AI platform - Build scalable ML pipeline - Deploy real-time inference API - Implement monitoring system - Integrate security frameworks - Project: Retail recommendation platform - Project: Fraud detection system
Day 2 – Industry Projects	- Build AI for healthcare diagnostics - Build telecom network prediction model - Build logistics demand forecasting - Build e-commerce recommendation system - Deploy models in cloud - Project: Supply chain optimization - Project: Financial fraud detection
Day 3 – Startup AI Infrastructure	- Building AI SaaS infrastructure - Scaling AI startups - Cloud architecture for AI companies - Cost optimization for startups - Global AI product deployment - Project: AI automation SaaS - Project: AI analytics platform