

AI in Telecom

CONTENTS

This course gives engineers and architects a technically rigorous grounding in how Artificial Intelligence and Machine Learning are deployed across modern telecom networks. It moves beyond vendor marketing to examine the real data pipelines, model architectures, API interfaces, and operational constraints that determine whether an AI system succeeds or fails in a production network environment.

The modules are ordered to build understanding progressively: from the industry context and programmable-network foundations, through domain-specific use cases in RAN, Core, and BSS, to Generative AI and the governance disciplines required for safe deployment. Lab exercises are integrated into each module.

The course covers both classical ML models (anomaly detection, churn prediction, reinforcement learning) and Large Language Model applications (NOC copilots, agentic operations, configuration generation). Hands-on lab exercises are included throughout.

TARGET AUDIENCE

The target group is engineers, architects, and technical managers who need to understand, design, or evaluate AI systems operating on live telecom networks including network operations engineers, RAN and Core architects, BSS developers, and data scientists working in the telecom domain.

PREREQUISITES

- Familiarity with RESTful / gRPC APIs
- Basic understanding of telecom network architecture (RAN, Core, OSS/BSS)
- Exposure to ML concepts (helpful but not required)

Note: Lab exercises require a Python 3.11+ environment.

Introduction

- The four macro-trends converging in telecom AI: margin pressure, software-defined networks, democratised ML tooling, and Generative AI
- Market investment signals and deployment maturity: Customer Experience AI, Revenue Assurance, AIOps, RAN Optimisation, Generative AI
- The Autonomous Network vision: TM Forum ANL maturity scale (L0-L5) and where operators stand today
- Hype vs. reality: NOC engineer replacement, zero-touch networks, federated learning at scale, LLM configuration push
- The API foundation: 3GPP SBA, O-RAN interfaces (A1/E2/O1/O2), CAMARA/GSMA Open Gateway, TM Forum Open APIs

Network Anomaly Detection & AIOps

- The problem with threshold-based operations: alert storms, blind spots, slow MTTR, excessive false positives
- Data sources and ingestion: SNMP traps, 3GPP PM counters, Syslogs, NetFlow/IPFIX, TM Forum TMF642 alarm feeds, signalling traces
- ML techniques for anomaly detection: Isolation Forest, Autoencoders, LSTM-based forecasting
- Supervised methods: Gradient Boosted Trees (XGBoost/LightGBM), Graph Neural Networks for root cause analysis
- Alarm correlation and Root Cause Analysis (RCA) topology-aware graph scoring
- Closed-loop remediation: low-risk automated actions, medium-risk configuration within safe bounds, high-risk human approval gate

AI-Driven RAN Optimisation

- Why the RAN is the hardest optimisation problem: rule-based SON limitations and multi-objective trade-offs
- O-RAN architecture and AI integration points: Non-RT RIC, Near-RT RIC, xApps and rApps
- O-RAN A1/E2 control loops timescales and primary use cases
- Handover optimisation with Reinforcement Learning: state, action, reward formulation; PPO agents
- AI-driven energy saving: traffic prediction, sleep candidate identification, coverage verification, O1 NETCONF/YANG command issuance
- Beam management in 5G NR mmWave: CNN/LSTM on CSI feedback, position-assisted beam prediction, federated learning across gNBs

Intelligent Core: Traffic Steering & Network Slicing

- The 5G Core as a programmable system: Service Based Architecture, NRF service discovery, key NFs relevant to AI (PCF, NWDAF, SMF, UDM, AMF)
- NWDAF – the native AI node of the 5G Core: analytics types, Nnwdaf_AnalyticsInfo and Nnwdaf_DataManagement services (3GPP TS 23.288 / TS 29.520)
- AI-powered network slicing: demand forecasting per slice, admission control models for URLLC, resource reallocation via NSSMF
- Intent-Based Networking: intent parsing (NLP/LLM), intent fulfilment engine, actuator layer (NETCONF/YANG, 5GC SBI, SDN controller APIs)

AI in BSS: Churn, Revenue Assurance & Customer Experience

- Why BSS data is uniquely valuable for AI: CDRs, CRM data, billing and payment data, customer care interaction transcripts
- Churn prediction: feature engineering with ARPU trend, data usage delta, care contacts, network quality signals; LightGBM/XGBoost model; SHAP explanations
- Revenue assurance with AI: anomaly detection on CDR streams, Wangiri fraud detection, graph-based fraud ring detection via GNN
- AI-powered customer experience: intelligent care routing and next-best-action (NBA), network experience indexing (QoE index), proactive care

Generative AI & LLMs in Telecom Operations

- Where Generative AI fits and where it does not: NOC copilots, configuration drafting, customer care chatbots vs. real-time closed-loop control
- NOC Copilot architecture: Retrieval-Augmented Generation (RAG) with runbook vector store, alarm context, and topology data
- Network configuration generation from natural-language intent: LLM drafting + deterministic validation agent + human approval gate
- Agentic AI in network operations: ReAct framework, tool-equipped agents for multi-step diagnostic workflows
- Telecom domain adaptation of LLMs: RAG over 3GPP/O-RAN specs, domain fine-tuning, structured output with schema validation

Architecture Patterns, Governance & Deployment

- Reference architecture: AI/ML Platform layer, AI Orchestration & Application layer, Network Exposure & API layer, Physical/Virtual Network
- MLOps for telecom: data drift monitoring, model versioning and staged rollout, shadow mode testing, feedback loop monitoring
- Explainability and auditability: SHAP values for tabular models, attention visualisation for neural networks, full audit trails for configuration changes
- Privacy and data governance: GDPR requirements, telecom-specific regulation, federated learning, differential privacy, data minimisation in feature pipelines

- Standards references: 3GPP TS 23.501/23.502/23.288/29.520/32.500, O-RAN WG1/WG2/WG3 specs, TM Forum IG1230/IG1251, ETSI GR ENI 001, ETSI GS ZSM 002

Practical Exercises During the Course

- Anomaly detection with Isolation Forest (Python / scikit-learn) on synthetic PM counter dataset
- Reinforcement learning for handover optimisation (PPO agent, Stable-Baselines3) using E2 telemetry dataset
- NWDAF slice load monitoring and proactive resource reallocation via mock NSSF API (Python / Flask)
- Churn prediction with LightGBM and SHAP feature attribution on synthetic CDR + CRM dataset
- NOC Copilot with Anthropic API: incident summary, root cause, and ranked remediation steps from alarm context