CONCEPT

AI-Enabled Threat Detection: Machine Learning, UEBA and SOAR in Security Operations



CONCEPT

AI-Enabled Threat Detection: Machine Learning, UEBA and SOAR in Security Operations

Technical depth: Practitioner · **Practical mode:** Lab

Introduction

Security operations teams face more telemetry than analysts can read, while attackers use stolen credentials and legitimate tools that rule-based detection misses. Machine learning, behavioural analytics and generative AI assistants promise relief, but poorly trained models flood the queue with false positives, drift silently and can themselves be attacked. This Core Concept course takes practitioners through building, evaluating and governing AI-enabled detection and response in hands-on labs, using MITRE ATT&CK, MITRE ATLAS and NIST AI guidance. Participants leave with an AI-Enabled Detection Use-Case Pack for their own SOC.

Course Objectives

- Assess where AI adds detection value in a SOC by baselining ATT&CK coverage, data sources and analyst workload
- Select supervised, unsupervised and behavioural analytics techniques that fit specific threat use cases
- Engineer features from security logs and train, evaluate and threshold detection models on real datasets
- Integrate model outputs into SIEM correlation and SOAR playbooks with human approval for high-impact actions
- Identify model drift, adversarial manipulation and generative AI risks, and apply guardrails from NIST and OWASP guidance
- Produce an AI-Enabled Detection Use-Case Pack ready for SOC engineering review

Target Audience

- SOC analysts responsible for triage and investigation who use or will use AI-assisted tooling
- Detection engineers who build and tune correlation rules and analytics content
- Security data engineers who manage log pipelines, schemas and SIEM data platforms
- Security automation engineers responsible for SOAR playbooks and integrations
- Threat hunters who apply anomaly and behavioural analytics to find hidden activity

Course Outline

Day 1: AI in Security Operations and Current State

- AI Use Cases in the SOC: Machine Learning, UEBA and Generative AI Assistants
- SOC Automation Current-State Assessment Using SOC-CMM
- Security Data Pipelines and Log Normalisation with OCSF
- Detection Coverage Baseline Against MITRE ATT&CK
- Alert Fatigue and Analyst Workload Metrics Dashboard

Day 2: Detection Models and AI Security Frameworks

- Supervised Classification for Phishing and Malware Detection
- Unsupervised Anomaly Detection: Isolation Forest, k-Means and DBSCAN
- UEBA Behavioural Baselines and Entity Risk Scoring
- NIST AI RMF 1.0 and ISO/IEC 42001:2023 Controls for Security AI
- Attacks on Detection Models: MITRE ATLAS and NIST AI 100-2 E2025

Day 3: Building and Integrating AI Detection

- Feature Engineering from Authentication and Network Flow Logs in Python
- Model Evaluation with Precision, Recall and ROC Curves
- SIEM Integration of Model Scores into Correlation Rules
- SOAR Playbook Automation for Enrichment and Containment
- Generative AI Assistant Prompting for Alert Summaries and Hunt Queries

Day 4: Tuning, Trust and AI Risk

- False Positive Tuning and Concept Drift Monitoring
- Prompt Injection and Data Leakage Controls from the OWASP Top 10 for LLM Applications 2025
- Human-in-the-Loop Approval Gates for Automated Response
- Model Explainability with SHAP for Analyst Trust
- Detecting AI-Enabled Attacks: Deepfake Phishing and Automated Reconnaissance

Day 5: Detection Labs and the Use-Case Pack

- Credential Abuse Detection Lab on a Supplied Dataset
- Insider Data Exfiltration Lab with UEBA Scoring
- SOAR Playbook Build and Dry Run
- AI-Enabled Detection Use-Case Pack Drafting
- Peer Review of Model Metrics and Pack Defence

Skills You Will Gain

- Security Data Engineering
- Anomaly Detection Modelling
- Detection Model Evaluation
- Behavioural Analytics
- Security Orchestration and Automation
- AI Model Risk Control
- Detection Use-Case Design

Why Attend This Course

- Return to work with an AI-Enabled Detection Use-Case Pack for your SOC, with model metrics already reviewed by peers
- Tell vendor claims about AI detection apart from measurable improvement in precision and response time
- Cut repetitive triage work through automation while keeping analysts in control of high-impact actions
- Compare AI adoption in security operations with practitioners from other sectors and countries

Conclusion

AI improves security operations only when models are fed the right data, measured honestly and kept under human control. This course moves from where AI fits in the SOC, through detection models and AI security frameworks, to feature engineering, evaluation, SIEM and SOAR integration, drift, explainability and attacks on AI itself. The final day applies that work in detection labs and turns it into an AI-Enabled Detection Use-Case Pack that participants take back to their team, ready for engineering review and phased deployment.