

AI Foundations & AI for Cybersecurity

Target Audience

This course is designed for cybersecurity professionals, SOC analysts, IT and network engineers, security architects, risk and governance practitioners, and technical managers who want to understand how artificial intelligence works and how to apply it to security. It suits learners at a beginner to intermediate level who want a structured path from AI and data foundations through machine learning, generative AI and LLMs, to AI-driven security operations and the security and governance of AI systems themselves. No prior experience with machine learning is assumed.

Course Outcomes

Participants will be able to:

- Explain core AI, machine learning, deep learning and generative AI concepts and where they apply in cybersecurity.
- Prepare and explore security data using Python, NumPy and Pandas.
- Understand how machine learning models are built and evaluated for malware, phishing, intrusion and anomaly detection.
- Understand neural networks, transformers, LLMs and Retrieval-Augmented Generation (RAG).
- Describe how AI supports security operations, threat intelligence, network, endpoint, identity and vulnerability management.
- Recognise key attacks on AI, ML, LLM, RAG and AI agent systems and the controls that mitigate them.
- Apply governance and risk frameworks such as the NIST AI RMF, MITRE ATLAS and OWASP guidance.
- Design an AI-powered cybersecurity solution as a capstone.

Course Objectives

- Build a solid foundation in AI, data and machine learning concepts.
- Develop practical Python skills for working with security data.
- Show how machine learning, deep learning and generative AI strengthen threat detection and response.
- Introduce LLMs and RAG as accelerators for security analysis and reporting.
- Establish awareness of the attack surface of AI systems and how to secure it.
- Embed responsible AI, governance and risk management practices.
- Reinforce learning through demonstrations, guided activities and a final capstone.

Course Outline

The course comprises 40-hours of instructor-led theory, practical demonstrations and implementation activities, and is divided into 10 different modules. Each module will be followed by guided exercises to reinforce learning and

gauge understanding of the topics covered. The course level is beginner to intermediate.

Prerequisites

Participants should have basic computer proficiency and a general awareness of IT and cybersecurity concepts. Prior programming or machine learning experience is beneficial but not mandatory. Each participant needs a workstation with Python installed and internet access from the machine.

Table of Contents

Module 1 – Artificial Intelligence and Data Foundations

- Introduction to AI: evolution, types of AI, AI vs Machine Learning vs Deep Learning, predictive vs generative AI
- AI applications in cybersecurity
- AI system and project lifecycle
- Data fundamentals: structured and unstructured data, features and labels, training, validation and test datasets
- Data collection, labeling and preprocessing: missing data, outliers, normalization, data imbalance
- Cybersecurity datasets: security logs as AI data, data quality and governance

Module 2 – Python and Machine Learning Fundamentals

- Python for AI: working with files (CSV, JSON), NumPy and Pandas
- Data manipulation and visualization
- Types of machine learning: supervised, unsupervised, semi-supervised, reinforcement learning
- Core ML tasks: classification, regression, clustering, anomaly detection
- Machine Learning workflow: training and inference
- Machine Learning applications in cybersecurity

Module 3 – Machine Learning for Cybersecurity

- Key algorithms: linear and logistic regression, decision trees, random forest, SVM, KNN, Naive Bayes, gradient boosting, K-Means, Isolation Forest
- Feature engineering: feature selection, extraction and importance
- Model tuning: hyperparameter tuning, overfitting and underfitting, cross-validation
- Model evaluation: accuracy, precision, recall, F1 score, confusion matrix, ROC and AUC, false positives and negatives
- Use cases: malware classification, phishing detection, intrusion detection, user behavior analysis, security anomaly detection

Module 4 – Deep Learning and Generative AI

- Neural network fundamentals: neurons, weights, activation functions, loss functions, backpropagation, gradient descent
- CNN, RNN and LSTM networks, sequence modeling
- Transformer architecture and the attention mechanism
- Deep Learning applications in cybersecurity
- Introduction to Generative AI: foundation models, Generative AI vs traditional AI, multimodal AI

- Generative AI applications in cybersecurity

Module 5 – Generative AI, LLMs and RAG

- LLM fundamentals: architecture, tokens, embeddings, context windows, inference
- Pre-training, fine-tuning and instruction tuning
- Prompt engineering: zero-shot and few-shot prompting, structured prompting, context management
- Hallucinations and limitations of LLMs
- LLM applications in cybersecurity: alert summarization, threat intelligence analysis, incident investigation, report generation
- Retrieval-Augmented Generation (RAG): architecture, document chunking, vector databases, retrieval mechanisms
- RAG applications in cybersecurity

Module 6 – AI for Security Operations and Threat Intelligence

- SOC fundamentals: SOC architecture and workflows, SIEM and security monitoring
- AI-assisted alert analysis: classification, prioritization and correlation
- AI-based anomaly detection and User and Entity Behavior Analytics (UEBA)
- AI-assisted threat hunting, incident investigation and incident response
- Cyber Threat Intelligence: strategic, tactical and operational intelligence, Indicators of Compromise, threat intelligence feeds
- AI for threat intelligence: IOC extraction, threat actor and malware analysis, report summarization, correlation
- AI-powered security copilots

Module 7 – AI for Network, Endpoint and Threat Detection

- AI for network security: traffic analysis, anomaly detection, AI-based intrusion detection systems, traffic classification
- AI for endpoint security: malware and ransomware detection
- Phishing, malicious URL and email security
- Identity and access security: account compromise and fraud detection
- Vulnerability management: AI-assisted prioritization and risk scoring
- AI-assisted digital forensics
- Limitations of AI in cybersecurity

Module 8 – Security of AI, ML and LLM Systems

- AI and Machine Learning attack surfaces, AI supply-chain attacks
- Adversarial Machine Learning: adversarial examples, evasion attacks, data and model poisoning
- Model extraction, theft and inversion, membership inference
- LLM security risks: prompt injection (direct and indirect), jailbreaking, system prompt leakage
- Sensitive information disclosure, improper output handling, excessive agency, model denial of service, vector and embedding security
- OWASP GenAI/LLM security guidance

Module 9 – Securing RAG, AI Agents and Enterprise AI

- RAG security: knowledge-base security, data poisoning, retrieval manipulation, unauthorized retrieval, access control
- AI agent architecture and security: agent identity and permissions, tool calling, human approval and oversight, agent monitoring
- Model Context Protocol (MCP) fundamentals and security considerations
- AI application security: API, identity and access, data, model and workload security
- AI logging, monitoring and security testing
- Secure AI development lifecycle and AI supply-chain security

Module 10 – AI Risk Management, Governance and Cybersecurity Capstone

- Responsible AI: principles, privacy and data protection, transparency, fairness, accountability, human oversight
- AI governance and risk: risk identification and assessment, third-party AI risks, shadow AI, enterprise AI and acceptable-use policies
- Frameworks: NIST AI Risk Management Framework (Govern, Map, Measure and Manage), NIST Cybersecurity Framework, MITRE ATLAS, OWASP AI security guidance
- Capstone: identify a cybersecurity problem suitable for AI, select AI/ML techniques, prepare data, develop a threat-detection approach
- Capstone: integrate Generative AI/LLMs and RAG, design an AI-powered SOC workflow
- Capstone: identify AI-specific threats, develop security controls, perform risk assessment, design human-in-the-loop controls
- Capstone: develop an enterprise AI cybersecurity architecture, document and present the solution