Course Syllabus: AI Security Engineer

Course Title: Fortifying AI: Securing Systems Against Threats

Target Audience: Ideal for cybersecurity professionals, engineers, and students interested in protecting AI systems. Basic programming knowledge (e.g., Python) and familiarity with security concepts are helpful but not required.

Course Level: Comprehensive program covering Basic, Intermediate, and Advanced levels.

Duration: 12 weeks (flexible for self-paced learning).

Course Description:
This course trains students to become AI Security Engineers, protecting AI systems from threats like prompt injection, data leaks, and adversarial attacks. You'll learn to secure models, data, and infrastructure, ensuring platforms like Zomato's billing system are safe and reliable. From foundational cybersecurity to advanced AI-specific defenses, you'll build skills to safeguard AI in production environments.

Learning Objectives:

Upon completion, students will be able to:

- Understand AI-specific security threats and vulnerabilities.
- Secure AI models, data pipelines, and infrastructure.
- Mitigate attacks like prompt injection and adversarial inputs.
- Implement secure AI deployment practices.
- Address compliance and ethical issues in AI security.
- Develop a portfolio of AI security projects.

Course Structure:

Part 1: Basic Foundations (Weeks 1-4)

This section introduces AI security and cybersecurity basics.

- Week 1: Introduction to AI Security
    - Role of an AI Security Engineer.
    - AI vulnerabilities: Prompt injection, data leaks, adversarial attacks.
    - Case Study: Securing Zomato's DynamoDB billing system.
    - Exercise: Identify risks in a sample AI system.
- Week 2: Cybersecurity Fundamentals
    - Basics: Encryption, authentication, access controls.
    - Tools: OpenSSL, AWS IAM for security.
    - Hands-on: Set up basic encryption for data.
- Week 3: AI System Components
    - AI pipeline: Data, models, inference, deployment.
    - Security risks at each stage (e.g., data poisoning).
    - Exercise: Map security risks for an AI pipeline.
- Week 4: Introduction to AI Attacks
    - Common attacks: Prompt injection, model inversion.
    - Examples: Misleading a chatbot, stealing model data.
    - Hands-on Project: Simulate a prompt injection attack on a sample AI.

Part 2: Intermediate Concepts (Weeks 5-8)

This section focuses on securing AI models and data.

- Week 5: Securing AI Data
    - Data protection: Anonymization, encryption.
    - Secure storage: NoSQL databases like DynamoDB.
    - Hands-on: Secure a dataset for AI training.

- Week 6: Model Security
    - Protecting models: Watermarking, model encryption.
    - Defending against adversarial attacks (e.g., image perturbations).
    - Case Study: Securing Zomato's recommendation model.
- Week 7: Secure Deployment
    - Secure APIs: Rate limiting, authentication.
    - Tools: AWS API Gateway, OAuth.
    - Hands-on: Deploy a secure AI model API.
- Week 8: Testing for Security
    - Penetration testing for AI systems.
    - Tools: Burp Suite, custom scripts for AI attacks.
    - Hands-on Project: Conduct a security audit on an AI system.

Part 3: Advanced & Expert-Level Application (Weeks 9-12)

This section prepares students for enterprise-grade AI security.

- Week 9: Advanced Adversarial Defenses
    - Robustness techniques: Adversarial training, defensive distillation.
    - Tools: CleverHans, ART for adversarial testing.
    - Exercise: Defend a model against adversarial inputs.
- Week 10: Secure Infrastructure
    - Securing cloud infrastructure: Kubernetes, serverless.
    - Monitoring: Intrusion detection for AI systems.
    - Hands-on: Secure an AI infrastructure on AWS.
- Week 11: Compliance and Ethics
    - Compliance: GDPR, CCPA for AI security.
    - Ethical issues: Privacy, transparency in AI security.

    - Exercise: Develop a security compliance plan for AI.

- Week 12: Capstone Project & Trends
    - Capstone Project: Secure an AI system for a Zomato-like platform (e.g., billing or recommendations).
    - Trends: Quantum-resistant AI security, federated learning security.
    - Career paths: Cybersecurity, AI ops, ethical hacking.

Assignments & Grading:

- Weekly Security Labs & Exercises: 25%
- Intermediate Projects (Weeks 4 & 8): 30%
- Capstone Project: 35%
- Class Participation & Peer Reviews: 10%