

AI Red Teaming

Special thanks to "Learn Prompting" for their help in making this roadmap

learnprompting.org

Find the detailed version of this roadmap and other similar roadmaps

roadmap.sh

- ✓ AI and Data Scientist Roadmap
- ✓ AI Engineer Roadmap
- ✓ Data Analyst Roadmap
- ✓ MLOps Roadmap

Introduction

AI Security Fundamentals

Why Red Team AI Systems?

Ethical Considerations

Role of Red Teams

Foundational Knowledge

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Neural Networks

Generative Models

Large Language Models

Prompt Engineering

AI / ML Fundamentals

Cybersecurity Principles

Confidentiality, Integrity, Availability

Threat Modeling

Risk Management

Vulnerability Assessment

Prompt Hacking

Jailbreak Techniques

Safety Filter Bypasses

Prompt Injection

Direct

Indirect

Countermeasures

System Security

Code Injection

Insecure Deserialization

Remote Code Execution

Infrastructure Security

API Protection

Authentication

Authentication

Testing Methodologies

Black Box Testing

White Box Testing

Grey Box Testing

Automated vs Manual

Continuous Testing

Professional Development

Tools and Frameworks

Model Weight Stealing

Unauthorized Access

Model Extraction

Data Poisoning

Adversarial Examples

Model Inversion

Model Manipulation

Adversarial Training

Robust Model Design

Continuous Monitoring

Defense Strategies

Conferences

Research Groups

Forums

Community Engagement

Lab Environments

CTF Challenges

Red Team Simulations

Practical Experience

Specialized Courses

Industry Credentials

Certifications

Real-world Applications

LLM Security Testing

Agentic AI Security

Responsible Disclosure

Testing Platforms

Monitoring Solutions

Benchmark Datasets

Custom Testing Scripts

Reporting Tools

Future Directions

Emerging Threats

Advanced Techniques

Research Opportunities

Industry Standards

Visit the following relevant tracks to keep learning

[AI Engineer](#)

[AI & Data Scientist](#)

[Data Analyst](#)