

Foundational AI Literacy, Skill Labs

Course Specifications

Course Number: ACI76-027SL_rev1.0

Lab Length: Approximately 6 hours

Fundamentals of AI

Introduction

Objective

SecAI Domain

1.0: Basic AI Concepts Related to Cybersecurity

4.0: AI Governance, Risk, and Compliance

SecAI+ Objectives

1.1: Compare and contrast various AI types and techniques used in cybersecurity.

1.2: Explain the importance of data security in relation to AI.

1.3: Explain the importance of security throughout the life cycle of AI.

4.2: Explain risks associated with AI.

Overview

Fundamentals of AI Lab is a hands-on training course that introduces the essential capabilities and responsible applications of artificial intelligence in today’s workplace. This immersive lab guides you through using AI tools, exploring ethical scenarios, and analyzing real-world tasks that require critical thinking. Over the two-hour session, you will participate in guided exercises, interactive demonstrations, and practical workplace simulations. By completing this course, you will develop the foundational skills needed to work efficiently, assess AI outputs for bias and accuracy, and integrate AI workflows into your daily operations.

Learning Objectives

- Identify the capabilities and limitations of AI tools.
- Apply ethical standards to professional AI use.
- Evaluate AI-generated outputs for accuracy and bias.
- Integrate AI workflows into business processes.
- Develop strategies for responsible and sustainable AI adoption.

| | Key Term | Description |
|---|------------------------------|---|
| 1 | Artificial Intelligence (AI) | The simulation of human intelligence processes by machines, especially computer systems, to perform tasks such as learning, reasoning, and self-correction. |
| 2 | Machine Learning | A subset of AI that enables systems to learn and improve from experience |

Course Outline

| | Key Term | Description |
|----|-----------------------------------|--|
| | (ML) | without explicit programming, using algorithms to analyze data and identify patterns. |
| 3 | Neural Network | A computational model inspired by the human brain, consisting of interconnected nodes ('neurons') that process information in layers to recognize relationships in data. |
| 4 | Algorithm | A step-by-step procedure or set of rules used by computers to perform tasks, solve problems, or make decisions. |
| 5 | Natural Language Processing (NLP) | A branch of AI focused on enabling computers to understand, interpret, and respond to human language in a meaningful way. |
| 6 | Bias | Unfair or prejudiced outcomes produced by AI systems caused by incomplete, inaccurate, or imbalanced training data, leading to systematic errors. |
| 7 | Ethics in AI | Guidelines and principles for designing, implementing, and using AI responsibly, ensuring fairness, transparency, privacy, and accountability. |
| 8 | Automation | The use of technology, including AI, to perform tasks automatically, reducing the need for human intervention and improving efficiency. |
| 9 | AI Workflow | The series of steps that outline how AI systems are implemented and integrated into business processes, from data input to decision-making and output. |
| 10 | Responsible AI Adoption | Strategies and practices to deploy AI in ways that are sustainable and ethical, considering long-term impacts on society, business, and individuals. |

Everyday AI Productivity Lab

Introduction

Objective

SecAI+ Domain

1.0: Basic AI Concepts Related to Cybersecurity

4.0: AI Governance, Risk, and Compliance

SecAI+ Objectives

1.1: Compare and contrast various AI types and techniques used in cybersecurity.

1.2: Explain the importance of data security in relation to AI.

4.2: Explain risks associated with AI.

Overview

Everyday AI Productivity Lab is a hands-on learning experience focused on using artificial intelligence tools to boost your workplace efficiency. In this interactive course, you will discover practical ways to leverage AI for everyday tasks, improve your workflow, and adopt responsible use practices. Through guided exercises and real-world scenarios, you'll build the skills needed to maximize productivity while ensuring accuracy, fairness, and ethical decision-making. The two-hour session includes demonstrations and workplace simulations designed for deeper exploration and immediate application.

Learning Objectives

- Identify AI tools that improve productivity and workflow.

Course Outline

- Apply responsible AI practices in common workplace scenarios.
- Evaluate AI outputs for accuracy, bias, and usefulness.
- Integrate AI-powered solutions into daily work tasks.
- Develop sustainable strategies for long-term AI adoption.

| | Key Term | Description |
|----|------------------------------|--|
| 1 | Artificial Intelligence (AI) | Technology that enables machines to perform tasks requiring human-like intelligence, such as learning, reasoning, and problem-solving. |
| 2 | Productivity Tools | Software or applications designed to help users complete tasks efficiently and improve work output. |
| 3 | Workflow Automation | The process of using technology to automatically perform repetitive or routine tasks within a business process. |
| 4 | Responsible AI | The practice of developing and using AI systems ethically, with attention to safety, fairness, and societal impact. |
| 5 | Bias and Fairness | Refers to the presence of systematic favoritism or discrimination in AI outputs, and ensuring algorithms treat all groups equitably. |
| 6 | Data Accuracy | The degree to which information produced or analyzed by AI matches real-world facts and is free from errors. |
| 7 | Ethical Decision-Making | Making choices about AI use that align with professional standards and societal values, considering potential consequences. |
| 8 | Workplace Scenarios | Realistic examples or simulations of situations encountered in professional environments, used for training or testing skills. |
| 9 | Sustainable Adoption | Implementing AI in ways that support long-term organizational goals and avoid negative consequences over time. |
| 10 | AI Integration | The process of embedding AI systems or capabilities into everyday tools, workflows, or business processes to enhance performance. |

AI Ethics in Action

Introduction

Objective

SecAI+ Domain

4.0: AI Governance, Risk, and Compliance
2.0: Securing AI Systems

SecAI+ Objectives

4.2: Explain risks associated with AI.
4.3: Summarize the impact of compliance on business use and development of AI.
2.5: Given a scenario, implement monitoring and auditing for AI systems.

Course Outline

Overview

Ethics in Action is an interactive, scenario-based lab designed to build essential skills for responsible AI use in professional environments. In this lab, you'll learn how to audit AI tool outputs for fairness, accuracy, and transparency, and develop practical strategies for handling digital ethics challenges. Through guided practice, you will examine real-world examples of AI bias, misinformation, and privacy risks, gaining confidence in making ethical decisions and recommending safe use policies for yourself or your organization.

Learning Objectives

- Recognize signs of bias or stereotyping in AI-generated content
- Detect and fact-check misinformation in AI outputs
- Apply best practices to protect data privacy when using AI
- Determine when human oversight is needed for AI-generated results
- Draft guidelines for ethical and transparent use of AI in your workflow

| | Key Term | Description |
|----|----------------|--|
| 1 | AI Ethics | The study and practice of how artificial intelligence systems should be designed and used in morally responsible ways. |
| 2 | Bias | Systematic favoritism or unfairness in AI outputs caused by skewed data or algorithms. |
| 3 | Transparency | The principle of making AI processes, decisions, and data sources clear and accessible to users. |
| 4 | Fairness | The concept of ensuring that AI systems treat all individuals and groups equitably without discrimination. |
| 5 | Misinformation | False, misleading, or inaccurate content created or spread by AI models or users. |
| 6 | Oversight | Human supervision to review, verify, and approve AI-generated decisions or outputs to prevent errors and misuse. |
| 7 | Fact-Checking | The process of validating the truthfulness and accuracy of AI-created information using reliable sources. |
| 8 | Data Privacy | Protecting sensitive personal and organizational information from unauthorized access or exposure when using AI tools. |
| 9 | Stereotyping | Assigning generalized traits or characteristics to individuals or groups in AI outputs, often leading to unfair representation. |
| 10 | Guidelines | Written standards or recommendations that define acceptable and ethical practices for building and deploying AI systems in everyday workflows. |

AI Prompting 101

Introduction

Objective

SecAI+ Domain

1.0: Basic AI Concepts Related to Cybersecurity

SecAI+ Objectives

1.1: Understand important AI concepts

Overview

Prompting 101

Prompting 101 introduces learners to the art and science of crafting effective prompts for AI tools, focusing on clarity, tone, and context. Through hands-on exercises and guided practice, participants learn to write structured prompts, control the format and personality of AI output, and analyze the impact of prompt specificity. This lab empowers users to recognize how context and structure influence output accuracy, and to assemble a starter prompt library for professional or everyday use.

Lab Objectives

In this lab, you will learn to:

- Write structured prompts with clear roles and instructions.
- Adjust prompts to control tone and format of AI output.
- Compare vague vs. precise prompts and analyze results.
- Recognize how prompt context influences accuracy.

| | Key Term | Description |
|----|--------------------|---|
| 1 | Prompt Engineering | The practice of designing clear and specific instructions or questions for an AI system to produce useful, targeted responses. |
| 2 | Structured Prompt | A prompt that uses a defined format, including roles or guidance, to direct how the AI should respond (e.g., "You are a cybersecurity analyst..."). |
| 3 | Context | Background information or details included in a prompt that help the AI understand the setting, audience, or purpose of the request. |
| 4 | Tone | The style or attitude (formal, informal, friendly, neutral, etc.) specified in a prompt to adjust how the AI communicates. |
| 5 | Role Assignment | Telling the AI to take on a certain persona or expertise in the response such as teacher, analyst, or support agent. |
| 6 | Vague Prompt | A request or question that lacks detail, often resulting in general or unfocused AI responses. |
| 7 | Precise Prompt | A highly detailed prompt that focuses the AI's response on specific facts, style, or requirements. |
| 8 | Iteration | The process of refining a prompt multiple times until you get the intended quality and clarity in responses. |
| 9 | Prompt Library | A personal or organizational collection of effective prompts saved for repeated use across similar tasks or workflows. |
| 10 | Audience Awareness | Adjusting the prompt's language or structure to meet the needs or understanding level of your target reader or user. |

AI Prompt Iteration Lab

Introduction

Objective

SecAI+ Domain

1.0: Basic AI Concepts Related to Cybersecurity

SecAI+ Objectives

- 1.1: Compare and contrast various AI types and techniques used in cybersecurity (e.g., prompt engineering, model training, validation, iterative prompting).
- 1.2: Explain the importance of data security in relation to AI (e.g., output refinement and safeguarding sensitive information during prompt iteration).
- 1.3: Explain the importance of security throughout the life cycle of AI (e.g., feedback and iteration, human-centric AI design principles).

Overview

The AI Prompt Iteration Lab teaches learners the art of refining AI-generated responses through an iterative process to achieve higher accuracy, usefulness, and alignment with task requirements. Participants will engage in hands-on practice using follow-up prompts to clarify, correct, and expand AI outputs, while also learning to break down complex tasks and critically evaluate results. The lab emphasizes the development of a repeatable methodology for prompt refinement that can be applied in professional and educational settings.

Lab Objectives

- Apply iterative prompting techniques to systematically improve the quality of AI-generated results.
- Use follow-up prompts to clarify, correct, or expand upon initial AI output.
- Explore strategies for breaking down complex tasks to make them manageable for AI-based solutions.
- Evaluate AI responses critically against the specific requirements of the task.
- Develop a repeatable process for refining prompts that enhances both efficiency and effectiveness when working with AI systems.
- These objectives support building foundational skills in prompt engineering and help learners become proficient in guiding AI towards delivering targeted, accurate, and contextually appropriate responses.

| | Key Term | Description |
|---|------------------------------|--|
| 1 | Artificial Intelligence (AI) | Technology that enables machines to simulate human intelligence by learning, reasoning, and problem-solving |
| 2 | Prompt Engineering | Designing, structuring, and refining the input (prompt) given to an AI system to optimize the quality and accuracy of its output |
| 3 | Iterative Prompting | Repeatedly refining and adjusting prompts based on previous AI responses to achieve better results, greater clarity, or higher relevance |
| 4 | Model Validation | Process of evaluating an AI model's responses to ensure they meet accuracy, reliability, and security standards |

Course Outline

| | Key Term | Description |
|----|----------------------------|--|
| 5 | Large Language Model (LLM) | An advanced AI model trained on vast datasets; capable of generating human-like responses to prompts |
| 6 | Follow-up Prompt | An additional or revised prompt used to clarify, correct, or expand the output generated by an AI system |
| 7 | Human-in-the-loop | Incorporating human review and intervention in the AI workflow to improve outcomes, provide oversight, and ensure alignment with task requirements |
| 8 | Hallucination | When an AI system generates plausible-sounding but incorrect or fabricated information in its output |
| 9 | Task Decomposition | Breaking down a complex problem or assignment into smaller, manageable sub-tasks to make it easier for an AI system to process |
| 10 | Validation | Critically assessing AI-generated output for accuracy, completeness, and alignment to specific objectives and requirements |

Few-Shot vs Zero-Shot AI Prompting

Introduction

Objective

SecAI+ Domain

1.0: Basic AI Concepts Related to Cybersecurity

SecAI+ Objectives

1.1: Compare and contrast various AI types and techniques used in cybersecurity (e.g., prompt engineering, model training, validation, iterative prompting).

1.2: Explain the importance of data security in relation to AI (e.g., output refinement and safeguarding sensitive information during prompt iteration).

1.3: Explain the importance of security throughout the life cycle of AI (e.g., feedback and iteration, human-centric AI design principles).

Overview

Introduction

This lab explores zero-shot and few-shot prompting—essential techniques for guiding AI behavior. You'll discover how using examples (few-shot) or providing minimal context (zero-shot) changes the way AI interprets and answers your queries. Through hands-on activities, you will design, test, and analyze prompts in practical cybersecurity scenarios. Along the way, you'll consider not just effectiveness, but also data security and responsible AI practices.

By the end, you'll be ready to tailor prompts for automation, security operations, and compliance, directly supporting the skills outlined in SecAI+ Domain 1.0.

Learning Objectives

- Explain the difference between zero-shot and few-shot prompting and when each approach is best used.
- Construct and test few-shot prompts to control and guide AI outputs.

Course Outline

- Evaluate trade-offs between brevity (zero-shot) and added context (few-shot) in prompt design—especially for accuracy vs. efficiency.
- Design prompts for quality and consistency across repeated tasks or outputs.
- Apply both prompting approaches to realistic cybersecurity cases, such as alert automation, policy compliance, and incident summarization.
- Integrate principles of data security and human oversight when leveraging prompts in sensitive environments.

| | Key Term | Description |
|----|---------------------|--|
| 1 | Zero-Shot Prompting | A technique where an AI model is given a task or question without any prior examples or demonstrations. The model relies solely on its pre-trained knowledge and the instruction provided in the prompt to generate a response |
| 2 | Few-Shot Prompting | A prompting technique that provides the AI model with a small number of example inputs and outputs before presenting the actual task. These examples help guide the model toward the desired response format and style |
| 3 | Prompt Engineering | The practice of designing, refining, and optimizing text inputs (prompts) to effectively communicate with AI models and achieve desired outputs. It involves understanding model capabilities, structuring instructions clearly, and iterating based on results |
| 4 | Context Window | The maximum amount of text (measured in tokens) that an AI model can process at one time, including both the input prompt and the generated output. This limit affects how much information can be provided in few-shot examples |
| 5 | In-Context Learning | The ability of AI models to learn and adapt their behavior based on examples or instructions provided within the prompt itself, without requiring additional training or fine-tuning. This is the mechanism that enables few-shot prompting to work |
| 6 | System Prompt | A set of persistent instructions or context provided to an AI model that defines its role, behavior, and constraints throughout a conversation or task. System prompts establish baseline behavior before user inputs are processed |
| 7 | Prompt Template | A reusable, structured format for prompts that includes placeholders for variable information. Templates ensure consistency across similar tasks and make it easier to apply prompting best practices at scale in security operations or automation workflows |
| 8 | Iterative Prompting | The process of refining and improving prompts through multiple cycles of testing, evaluation, and adjustment. This involves analyzing AI outputs, identifying weaknesses, and modifying prompts to achieve better accuracy, relevance, or security compliance |
| 9 | Hallucination | When an AI model generates information that appears plausible but is factually incorrect, fabricated, or not supported by its training data or the provided context. In cybersecurity contexts, hallucinations can lead to dangerous misinformation or flawed security recommendations |
| 10 | Token | The basic unit of text that AI models process. A token can be a word, part of a word, or even a character, depending on the model. Token limits determine how much text can be included in a prompt and response, directly impacting the feasibility of few-shot prompting strategies |