Applied AI Skills & Workflow Design, Skill Labs

Course Specifications

Course Number: ACI76-028SL_rev1.0
Lab Length: Approximately 7 hours

AI Research Analyst Lab

Introduction

Objective

SecAI Domain

1.0: Basic AI Concepts Related to Cybersecurity

SecAI+ Objectives

1.1: Understand important AI concepts.
1.2: Recognize limitations and potential errors in AI-generated outputs.
1.3: Evaluate the trustworthiness of AI findings by cross-referencing with credible sources.
1.4: Apply critical thinking skills to identify gaps, conflicts, and inconsistencies in information.
1.5: Synthesize and communicate findings clearly in a workplace-ready research brief.

Overview

In today's fast-paced, information-rich workplaces, professionals are often required to make decisions and deliver insights based on multiple sources of information. However, not all sources are equally reliable, and sometimes the available data contains contradictions or gaps. This lab places you in the role of an AI Research Analyst, where you will leverage artificial intelligence tools to summarize documents, verify facts, spot inconsistencies, and synthesize findings into a clear, actionable research brief.

Learning Objectives

- Generate summaries from multiple sources using AI.
- Cross-check AI-generated claims against trusted references.
- Identify gaps, inconsistencies, or contradictions in information sets.
- Synthesize insights into a cohesive report.
- Deliver a research brief suitable for workplace communication.

	Key Term	Description
1	Artificial Intelligence (AI)	The simulation of human intelligence processes by computer systems, including learning, reasoning, and self-correction
2	Machine Learning	A subfield of AI that focuses on algorithms and statistical models that enable computer systems to learn and improve from experience without being explicitly programmed

	Key Term	Description
3	Natural Language Processing (NLP)	The branch of AI that enables computers to understand, interpret, and generate human language
4	Fact-Checking	The process of verifying the accuracy and truthfulness of information by cross-referencing with credible sources
5	Synthesis	The act of combining insights from multiple sources or data points to create a unified, coherent report or summary
6	Bias (in AI)	Predisposition or prejudice present in AI outputs, often stemming from the data the AI was trained on
7	Hallucination (AI)	The generation of false or misleading information by an AI system, especially in text or responses
8	Source Credibility	The trustworthiness and reliability of the origin of information, which impacts its acceptance in analysis
9	Contradiction	The state of having information that conflicts or disagrees, requiring critical analysis and resolution
10	Research Brief	A concise document that summarizes findings, verifies facts, highlights gaps, and provides actionable insights suitable for decision-making in the workplace

Data Storytelling Lab: Turning Messy Data into Insights with AI

Introduction
Objective

SecAI+ Domain

1.0 Basic AI Concepts Related to Cybersecurity
3.0 AI-assisted Security
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.
3.1: Given a scenario, use AI-enabled tools to facilitate security tasks.
3.3: Given a scenario, use AI to automate security tasks.
4.2: Explain risks associated with AI.
4.3: Summarize the impact of compliance on business use and development of AI.

Overview

Learners are given access to a messy, real-world dataset. Using generative AI tools (e.g., ChatGPT, Gemini, Copilot), they will clean data, analyze for trends and anomalies, generate visualizations, and critically compare AI outputs with their own interpretations. The outcome is a short data visualization report telling the "story" of the data.

Learning Objectives

- Import and clean data using AI-powered tools for suggestion and automation.
- Identify trends and anomalies in raw and partially cleaned datasets.
- Generate and refine visualizations (charts, graphs, infographics) with AI support.
- Evaluate the accuracy and value of AI's analysis and recommendations versus manual insights.
- Produce a professional, polished report presenting findings and visual stories.

	Key Term	Description
1	Messy Data	Raw, unstructured data containing errors, inconsistencies, missing values, or formatting issues that require cleaning before analysis
2	Data Cleaning	The process of identifying and correcting errors, removing duplicates, handling missing values, and standardizing data formats to improve data quality
3	Generative AI	AI systems that can create new content, including text, code, images, or analysis outputs, by learning patterns from training data and generating original responses
4	Trend Analysis	The examination of data patterns over time to identify consistent directions, changes, or movements that reveal meaningful insights and predictions
5	Anomaly Detection	The process of identifying unusual patterns, outliers, or deviations in data that do not conform to expected behavior and may indicate errors or significant events
6	Data Visualization	The graphical representation of data using charts, graphs, maps, and other visual elements to make complex information easier to understand and communicate
7	AI-powered Tools	Software applications and platforms that integrate artificial intelligence capabilities to automate, enhance, or assist with tasks such as data processing, analysis, and decision-making
8	Critical Evaluation	The systematic assessment and judgment of AI outputs, data analysis, or information quality by comparing results against known facts, logical reasoning, and alternative sources
9	Data Integrity	The accuracy, consistency, and reliability of data throughout its lifecycle, ensuring it remains complete and unchanged unless authorized modifications are made
10	Visualization Report	A document or presentation that combines data visualizations with narrative explanations to communicate findings, insights, and recommendations to stakeholders

Audience-Aware Communication: Tailoring Technical Messages with AI

Introduction
Objective

SecAI+ Domain

1.0: Basic AI Concepts Related to Cybersecurity
2.0: AI Tools and Prompt Engineering
3.0: AI Ethics and Responsible Use
4.0: AI in Professional Workflows

SecAI+ Objectives

1.1: Understand important AI concepts.
1.2: Recognize limitations and potential errors in AI-generated outputs.
1.3: Evaluate the trustworthiness of AI findings by cross-referencing with credible sources.
1.4: Apply critical thinking skills to identify gaps, conflicts, and inconsistencies in information.
1.5: Synthesize and communicate findings clearly in a workplace-ready research brief.

2.1: Craft effective prompts for AI tools.
2.2: Iterate and refine prompts based on output quality.
2.3: Understand prompt components and their impact.

3.1: Maintain transparency about AI use in professional work.
3.2: Avoid misleading simplification that could harm decision-making.
3.3: Ensure accuracy when adapting technical information.
3.4: Consider stakeholder impact of communication choices.

4.1: Integrate AI tools into existing workflows.
4.2: Maintain quality control over AI-generated outputs.
4.3: Document AI-assisted processes for organizational accountability.
4.4: Deliver professional work products that meet workplace standards.

Overview

In today's workplace, professionals must communicate the same information to multiple audiences—each with different technical backgrounds, priorities, and needs. This lab teaches learners how to use AI to adapt technical messages for executives, technical staff, and clients while maintaining accuracy and clarity. Learners will practice identifying communication risks, evaluating AI-generated adaptations, and delivering professional communication packages.

Learning Objectives

By the end of this lab, learners will be able to:

- Analyze audience needs and identify key differences in communication requirements for executives, technical staff, and clients.
- Use AI effectively to adjust tone, technical complexity, and format for different stakeholder groups.
- Identify and mitigate risks of oversimplification or over-technicality in adapted messages.
- Evaluate AI-generated content critically for accuracy, appropriateness, and effectiveness.
- Compare multiple versions of the same message side-by-side to ensure consistency of core information.
- Deliver a professional communication package with audience-specific messages ready for real-world use.

	Key Term	Description
1	Audience-Aware Communication	The practice of tailoring messages to fit the knowledge level, priorities, and communication preferences of different stakeholders, ensuring information is accessible and relevant to each specific group
2	Stakeholder Analysis	The process of identifying and understanding the needs, concerns, technical expertise, and communication preferences of different groups who will receive a message, allowing for more targeted and effective communication

	Key Term	Description
3	Message Adaptation	The act of modifying a core message to suit different audiences by adjusting language, detail level, tone, and focus, while maintaining the same essential information and meaning across all versions
4	Tone Adjustment	Changing the style and attitude of written content to match the expectations and preferences of a specific audience, ranging from formal and serious to conversational and accessible, depending on the reader's role and context
5	Technical Complexity	The level of specialized knowledge, terminology, and detail required to understand technical information, which must be carefully calibrated based on each audience's expertise and familiarity with the subject matter
6	Oversimplification	The risk of removing too much detail or nuance from a message, resulting in incomplete information that could lead to misunderstandings or poor decision-making by the audience
7	Prompt Engineering	The skill of crafting clear, specific instructions for AI tools that produce desired outputs, including defining the task, specifying the audience, setting the tone, and providing context to guide the AI effectively
8	Consistency Verification	The process of checking that different versions of the same message maintain the same core facts, key points, and critical information across all audience adaptations to ensure no contradictions or inaccuracies emerge
9	Professional Communication Package	A complete set of audience-specific communications for a single topic or situation, typically including multiple versions of the same core message tailored for different stakeholder groups, ready for distribution and use in professional settings

Decision-Making with AI: Scenario Planning and Risk Tradeoffs

Introduction
Objective

SecAI+ Domain

1.0: Basic AI Concepts Related to Cybersecurity

SecAI+ Objectives

1.1: Understand important AI concepts

Overview
Prompting 101

In this lab, learners work through a realistic workplace scenario and use a generative AI assistant to propose options, surface risks, and stress-test different courses of action. Learners compare AI-generated recommendations with their own professional judgment, identify blind spots and assumptions in AI output, and build a simple decision matrix that incorporates security, ethical, and operational risk factors inspired by SecAI+ governance and risk objectives.

Lab Objectives

By the end of this lab, learners will be able to:

- Define a workplace problem or scenario with clear constraints, stakeholders, and success criteria.
- Use an AI tool to generate multiple solution options and associated risks for a given scenario.
- Evaluate AI-generated recommendations against professional judgment, organizational policy, and security considerations.
- Identify possible blind spots, biases, or unrealistic assumptions in AI reasoning and outputs.
- Construct a basic decision matrix that weighs benefits, risks, and tradeoffs across at least three AI-supported options.
- Describe when AI should inform, not replace, critical decisions in security-sensitive or high-impact contexts.

	Key Term	Description
1	Decision-support AI	AI tools used to inform and enhance human decisions by generating options, insights, or risk signals, while humans retain final authority.
2	Scenario planning	A structured process of exploring multiple "what-if" situations to test strategies and understand how different choices change risk and outcomes.
3	Risk evaluation	The practice of assessing the likelihood and impact of potential negative outcomes so that different options can be compared and prioritized.
4	Decision matrix	A table that scores or ranks options against criteria such as impact, cost, security, and compliance to support transparent, repeatable decisions.
5	Generative AI	A type of AI that can create new content (such as text, code, or images) based on patterns learned from training data, often used through chat-style interfaces.
6	Prompt engineering	The practice of crafting and refining inputs (prompts) to guide AI systems toward more accurate, relevant, and safe outputs.
7	Human-in-the-loop	A design approach where humans review, adjust, or approve AI-supported actions before they are implemented in real systems or processes.
8	Bias in AI output	Systematic skew or unfairness in AI responses caused by data, design, or usage patterns that can distort analysis and decision quality.
9	Blind spot	A risk, stakeholder, or constraint that is missing or underrepresented in either AI-generated analysis or a human's own reasoning about a scenario.
10	Risk appetite	The level and type of risk an organization is willing to accept in pursuit of its objectives, which should guide how AI-informed options are selected.

Prompt Chaining & Workflow Design

Introduction
Objective

SecAI+ Domain

1.0: Basic AI Concepts Related to Cybersecurity

SecAI+ Objectives

1.1: Understand important AI concepts.
1.2: Recognize limitations and potential errors in AI-generated outputs.
1.3: Evaluate the trustworthiness of AI findings by cross-referencing with credible sources.
1.4: Apply critical thinking skills to identify gaps, conflicts, and inconsistencies in information.
1.5: Synthesize and communicate findings clearly in a workplace-ready research brief.

Overview

This lab introduces learners to the concept of prompt chaining, where each AI-generated output becomes the input to the next prompt in a multi-step workflow. Through guided activities, learners build structured AI workflows to perform complex tasks such as summarizing data, generating content, refining style, or automating reasoning steps. Participants will plan, implement, and document a fully functional chained prompt sequence while considering efficiency, consistency, and risk in automated decision paths.

Learning Objectives

By the end of this lab, learners will be able to:

- Explain the purpose and benefits of prompt chaining.
- Design multi-step AI workflows where outputs feed subsequent prompts.
- Apply structured input/output formatting to enhance consistency between steps.
- Evaluate efficiency gains and potential error propagation risks in chained systems.
- Create and document a reusable AI workflow that performs a real-world task.

	Key Term	Description
1	Prompt Chaining	Designing a multi-step sequence where the output of one AI prompt is deliberately used as the input to the next step to accomplish a larger task.
2	AI Workflow	A structured series of AI-assisted steps, often combining prompts, tools, and validation checks, to support a cybersecurity or business objective.
3	Intermediate Output	A partial result produced at an intermediate step in a chain that shapes or constrains later prompts.
4	Structured Input	Information passed between steps using a consistent format (such as lists, tables, or JSON-like schemas) to reduce ambiguity.
5	Error Propagation	The way inaccuracies, hallucinations, or biases in one AI output can be amplified when reused in subsequent steps.
6	Hallucination Risk	The likelihood that an AI model produces confident but incorrect or fabricated details that may enter a chain as if they were facts.
7	Verification Step	A specific point in the workflow where AI outputs are checked against trusted tools, logs, policies, or documentation.
8	Guardrail Prompting	Using constraints, instructions, and formatting rules in prompts to reduce unsafe, irrelevant, or low-quality outputs within a chain.
9	Overreliance on AI	Depending too heavily on AI-generated outputs without sufficient human review, especially risky in security workflows.
10	Workflow Documentation	A concise artifact that captures the purpose, steps, inputs/outputs, risks, and validation methods of an AI prompt chain so others can safely reuse it.

AI-Powered Code Assistant

Introduction

As AI models gain the ability to process enormous amounts of code, developers can now offload complex understanding, refactoring, and documentation tasks to intelligent assistants. Google's Gemini 1.5 Flash model, with its extended context window, enables developers to work with entire applications—spanning hundreds of thousands of lines of code—in a single session. This lab demonstrates how to use Google AI Studio to transform Gemini into a "Senior Full-Stack Developer" capable of understanding, generating, and improving code across multiple languages.

In this hands-on exploration, you'll use Gemini's Chat prompt type to perform real-world software engineering tasks: creating boilerplate structures, explaining advanced functions to less experienced developers, and refactoring legacy code with stylistic or architectural constraints. This exercise not only develops practical AI tool fluency but also builds intuition about prompt design, model context management, and iterative code collaboration between human and machine.

Learning Objectives

By the end of this lab, you will be able to:

- Use Google AI Studio's chat prompt type to engage Gemini as a code assistant.
- Generate structured boilerplate code such as a FastAPI CRUD service using Gemini 1.5 Flash.
- Analyze and explain complex code snippets in plain language for educational or review purposes.
- Document Gemini's process and rationale as part of an AI-assisted code review.
- Reflect on the ethical and practical implications of AI-augmented software engineering.

	Key Term	Description
1	Gemini 1.5 Flash	A high-speed variant of Google's multimodal AI model optimized for reasoning and code tasks within long context windows
2	Long context window	The capacity of a model to process and remember large text or code inputs (often 1M+ tokens) simultaneously
3	Google AI Studio	A web-based interface for experimenting with Google's Gemini models through prompt testing, file uploads, and API prototyping
4	System instructions	A configuration feature in Gemini prompts that defines the model's role, personality, or specialized behavior
5	Chat prompt type	A mode that allows for conversational interaction with the AI model, supporting iterative back-and-forth exchanges
6	Boilerplate code	Standardized code sections used as templates for setting up common application structures such as APIs or data layers
7	FastAPI	A modern Python framework for creating APIs efficiently with data validation, async support, and auto-generated docs
8	Refactoring	The process of restructuring existing code without changing its external behavior, aimed at improving readability, maintainability, or performance
9	Functional programming	A coding paradigm emphasizing pure functions, immutability, and declarative logic over imperative statements

	Key Term	Description
10	Code validation	The act of testing or running code to confirm that new or modified sections function as expected and adhere to given requirements

AI in Test-Driven Development (TDD)

Introduction

Test-Driven Development (TDD) emphasizes writing tests before implementing application logic, ensuring reliability, maintainability, and early bug detection. With recent advances in AI coding assistants, this process can be accelerated through automated generation of comprehensive test suites. In this lab, you'll explore how AI models—specifically Gemini 1.5 Pro and Gemini 1.5 Flash—can assist in creating unit tests and identifying coverage gaps across JavaScript modules. You will also evaluate how effectively each model handles edge cases and subtle logic issues that might escape standard testing approaches.

By integrating Google AI Studio's Structured Prompt feature, you'll guide the AI to produce detailed Input (Code) and Output (Test Cases) scenarios. This workflow not only demonstrates AI's capacity for fast test generation but also deepens your understanding of how models reason about code behavior and quality.

Learning Objectives

By the end of this lab, you will be able to:

- Use Google AI Studio's Structured Prompt feature to generate and refine unit test cases.
- Employ AI models (Gemini 1.5 Pro and Flash) to identify coverage gaps in automated test generation.
- Quantify test coverage using Jest tools.
- Analyze the completeness and accuracy of AI-generated tests for different logic scenarios.
- Compare model performance in handling complex or edge-case behaviors in code.

	Key Term	Description
1	Test-driven development (TDD)	A software development approach where tests are written before the actual code, guiding implementation
2	Unit test	A short code segment that verifies whether a specific function or component behaves as expected
3	Jest	A JavaScript testing framework used to test frontend and backend applications efficiently
4	Edge case	A scenario at the extreme limits of input conditions, often where software bugs are most likely to occur
5	Code coverage	A measure indicating how much of the source code is executed when a test suite runs
6	Automated test generation	The process of using tools or AI to automatically create test cases based on code structure or examples
7	Logic bug	An error in the code's reasoning that causes incorrect outputs despite

	Key Term	Description
		syntactically valid code