Course Introduction

GIA AI Literacy Tutorial – Course Introduction

What if the tool you use to write your essays, summarize articles, and generate project ideas could also subtly shape your worldview, reinforce bias and prejudices, or even mislead you with confident errors?

Generative AI tools, which we usually refer to as Artificial Intelligence (AI), were once unfamiliar before OpenAI introduced ChatGPT to the world, making AI no longer a distant, futuristic concept but something that is quietly embedded in our daily routines.

However, the problem is that many people rely on AI without fully understanding how it works or why it provides specific answers. Simply "using" AI is very different from truly understanding and using it wisely. What we need now is the ability to think critically about AI and use it responsibly.

Why You Need AI Literacy

This is where AI Literacy gets significant. AI systems don't know the difference between truth and error. Their focus is on being seen as accurate rather than being correct. That means they can be incredibly useful, but also extremely wrong. When students rely on AI without understanding how it works, they risk spreading misinformation, reinforcing stereotypes, or losing their voice in the writing process. Even educators face challenges in distinguishing between authentic work and AI-generated content. This is why AI literacy is essential.

The Purpose and Goal of This Course

The GIA AI Literacy Tutorial was created to help students and teachers become informed, ethical, and confident users of AI. The tutorial's mission is to go beyond surface-level usage and cultivate three essential capacities:

- **Functional Literacy**: Understanding how AI tools work and how to use them effectively.

- **Critical Literacy**: Questioning AI output, identifying biases, and considering ethical consequences.

- **Rhetorical Literacy**: Using AI as a creative partner while maintaining human voice, authorship, and intent.

Together, these literacies form the foundation for navigating and shaping a future where AI is a powerful ally—but only if used wisely.

This Course Is For…

This tutorial is designed for anyone in the school community, whether you're a student writing essays or a teacher planning lessons. It welcomes:

- Curious beginners with no technical background.
- Experienced users who want to deepen their understanding.
- Skeptical learners who want to explore AI's risks and benefits.
- Creative thinkers, researchers, and future changemakers.

What's Inside: Course Content Overview

This tutorial's main focus isn't on learning only conceptual, but hands-on, practical, and reflective

The GIA AI Literacy Tutorial consists of five engaging modules:

1. **What is AI?**
2. **The Risks and Challenges of AI**
3. **What is AI Literacy?**
4. **Using AI in School**
5. **The Future of AI in Education**

In addition, the course includes interactive exercises such as:

- Spot the Hallucination: Identify errors in AI-generated paragraphs
- Prompt Comparison: Analyze how wording shapes AI responses
- AI Ethics Dilemmas: Discuss real-world scenarios involving bias or misuse
- Tool Introductions: Explore real-world platforms like Cursor, MCP, and basic AI agents

What You Will Gain

By the end of this course, you will:

- Understand how and when to use AI responsibly
- Evaluate the credibility and biases of AI-generated content
- Write purposeful prompts for your goals and audience
- Use AI as a thinking partner—without losing your voice
- Make ethical decisions when using AI in academic and creative settings

AI is already reshaping how we learn, communicate, and think. The question is: will you let it shape you, or will you learn how to shape it?

Welcome to the GIA AI Literacy Tutorial—where you don’t just use AI. You understand it.

Key words

Draft of Key Terms (Joshua/Leo)

Key terms	Explanation
Algorithm	**Definition:** A step-by-step set of rules or instructions that tells a computer how to solve a problem or perform a task. In AI, algorithms process input data, recognize patterns, make decisions, and produce outputs. **Example:** A search algorithm on Google decides which websites to show first when you type in a question. **Illustration:** Flowchart
Artificial Intelligence (AI)	**Definition:** The field of computer science focused on building systems that can perform tasks normally requiring human intelligence, such as pattern recognition, reasoning, decision making, and language understanding. **Example:** AI in self-driving cars that detects pedestrians and decides when to stop. **Illustration:** Brain icon + arrows leading to tasks (speech, vision, decision-making).
Chatbot	**Definition:** A software application that communicates with humans through text or voice. Modern chatbots often run on top of large language models (LLMs), providing natural and human-like responses. **Example:** ChatGPT answering your homework questions or a customer service bot guiding you through a bank's website. **Illustration:** Chat bubbles between a human icon and a robot icon.
Foundation Model	**Definition:** A large-scale AI model trained on vast and diverse datasets to be general-purpose. These models can be adapted or fine-tuned for many tasks, such as translation, summarization, or image analysis. Narrow models, in contrast, are built only for one specific task. **Example:** GPT-4 is a foundation model that can be adapted for writing essays, coding, or summarizing research papers. **Illustration:** A big "foundation block" supporting smaller task-specific blocks (translation, summarization, coding).
Generative Aftificial Intelligence (GAI)	**Definition:** A branch of AI focused on models that create new content (text, images, music, or video) by learning patterns from massive training datasets. Generative AI doesn't "copy" but produces new combinations based on what it has learned. **Example:** DALL·E generating an original image of "a cat dressed as an astronaut." **Illustration:** A robot holding tools to create outputs: 📄 text, image, music.

Hallucination	**Definition:** When an AI model confidently produces false or misleading information that sounds correct. Hallucinations happen because the model predicts what words are likely to come next, not whether the content is factually accurate. **Example:** An AI inventing a fake scientific article or giving the wrong date for a historical event. **Illustration:** Robot speaking confidently with a speech bubble that contains incorrect or nonsensical information (e.g., “The Eiffel Tower is in Rome”).
Large Language Model (LLM)	**Definition:** A generative AI model trained on huge text datasets to predict the next token (a token = a small chunk of text such as a word or sub-word). Most modern LLMs use the transformer architecture with self-attention, which lets the model weigh relationships between words across long passages. LLMs are first pre-trained by self-supervised learning (predicting missing/next tokens), then often fine-tuned for safer or more helpful behaviour. **Why it matters:** Because they model the probability of text, LLMs can write essays, answer questions, translate, or code—but they can also hallucinate when the most probable text isn’t true. **Example:** ChatGPT drafting a lab summary from bullet points; it infers phrasing by predicting likely next tokens given your notes. **Illustration:** Flow: input text → tokenizer → transformer blocks (self-attention) → next-token probabilities → generated text.
Machine Learning	**Definition:** A family of methods where a computer learns patterns from data to map inputs → outputs without being explicitly programmed for every rule. Learning is guided by a loss function (how wrong the model is) and an optimizer (e.g., gradient descent) that adjusts parameters to reduce that loss. 3 Main Catagories of Machine Learning: 1. Supervised (labeled examples, e.g., email → “spam/not spam”) 2. Unsupervised (find structure without labels, e.g., clustering), 3. Reinforcement (learn by trial, reward, and penalty). **Example:** A model learns from historical weather features (humidity, pressure, wind) to predict tomorrow’s chance of rain. **Illustration:** Dataset split (train/validation/test) feeding a model; arrows show parameter updates minimizing loss.

Neural Network	**Definition:** A layered function approximator made of neurons (units) connected by weights. Each layer computes activations (e.g., ReLU) of weighted sums of the previous layer. During training, backpropagation computes gradients of the loss with respect to each weight so the optimizer can update them. Variants. Convolutional networks excel at images; recurrent/transformers handle sequences like text or audio. **Example:** A CNN identifies handwritten digits by learning edge/shape features in early layers and digit-level features in deeper layers. **Illustration:** Diagram of input layer → hidden layers (weights + activations) → output, with a backward arrow labelled "gradients/backprop."
Prompt	**Definition:** The instruction and context you give a generative model (often including role, task, constraints, and examples). Because LLMs perform in-context learning, the prompt shapes what patterns the model retrieves and how it formats the answer. Good practice. Specify goal, audience, constraints (length, tone, format), and include few-shot examples when possible. **Example:** "You are a lab TA. Summarize this procedure in ≤120 words, bullet points, and flag any safety risks: [sample text]." Illustration: Prompt block with sections (role → task → examples → constraints) feeding into the model, then a formatted output.
Reinforcement Learning from Human Feedback	**Definition:** A post-training alignment method: humans compare multiple model outputs for the same prompt; those comparisons train a reward model that scores outputs. A policy model is then optimized (e.g., with PPO, a policy-gradient method) to produce answers that the reward model (standing in for human preferences) rates highly. **Why it matters?** RLHF steers models toward helpful, harmless, and honest behaviour, reducing toxic or unsafe responses—but it can also encode human bias present in the feedback. **Example:** Annotators choose which of two chatbot replies is clearer and safer; over many comparisons, the assistant learns to prefer that style. **Illustration:** Loop: prompts → model candidates → human comparisons → reward model → policy optimization → improved assistant.
Training Data	**Definition:** The corpus used to teach a model: text,

	images, audio, code, etc. Data must be preprocessed (cleaned, tokenized), often annotated (labels or preferences), and split into train/validation/test to measure generalization. Data quality, diversity, and licensing/ethics (privacy, consent, bias) critically affect model behaviour. Risks. Bias in data → biased outputs; data leakage (seeing test answers during training) → inflated performance; distribution shift (real-world differs from training) → degraded accuracy. **Example:** An image classifier trained mostly on daylight photos may fail at night; adding balanced night images improves robustness. **Illustration:** Data pipeline: collection → cleaning/annotation → splits → training → evaluation, with "bias/privacy/license" checkpoints.

Function AI Literacy

Functional AI Literacy Introduction (Justin&Joshua)

Have you ever asked ChatGPT a question and simply accepted the answer without wondering *how* it was generated? Have you found yourself adjusting your prompt multiple times, hoping for better results, without really knowing what's happening behind the scenes? You're not alone. In fact, most users treat AI as a black box: something that "just works." But using AI without understanding how it operates is like driving a car without knowing how to steer; it might get you somewhere, but not necessarily where you want to go.

This lack of understanding leads to **blind trust, poor prompting, and even misinformation.** It's not just about knowing *what* AI can do, but *how* it does it, and that's where Functional AI Literacy becomes essential.

Functional AI Literacy refers to the ability to understand how AI systems are structured, how they process inputs, and how to interact with them effectively. In other words, it's about learning the mechanics of how tools like ChatGPT, Gemini, or DALL·E generate responses so that you can utilize them **intentionally, responsibly, and efficiently**.

In this section of the course, you will learn:

- What is Artificial Intelligence (AI), and its various applications
- What a Generative AI is and how LLM works
    - The role of training data, probabilities, and tokens
- Practical techniques to improve prompting strategies (with examples)

But this isn't just a passive learning experience. You'll also:

- Compare outputs from different AI models to identify patterns and differences
- Explore prompt engineering principles through interactive exercises
- Take a short quiz to visualize the internal steps of LLM response generation
- Analyze how seemingly small changes in input can produce drastically different outputs

By the end of this section, you won't just *use* AI – you'll understand it. And that understanding will empower you to use it with far more control, confidence, and creativity.

What is Artificial Intelligence (AI), and its various applications

Artificial Intelligence (AI) is an emerging field of technology that, according to IBM, "enables computers and machines to simulate human learning, comprehension, problem solving, decision making, creativity, and autonomy." Many people have a misconception that AI is a new concept; however, it was formally introduced in the 1950s. However, for decades,

AI remained largely dormant, with limited practical impact and few major breakthroughs. That changed dramatically in 2022 with the release of ChatGPT-3. This milestone reignited widespread interest in AI and positioned it as a transformative force across industries.

Today, AI is no longer confined to academic labs or science fiction. It plays a critical role in various fields. In healthcare, it supports early disease detection, personalized treatment, mental health chatbots, and even real-time monitoring through wearables. In finance, AI powers fraud detection, credit scoring, and robo-advisors that offer customized investment strategies. In transportation, AI is at the core of self-driving vehicles, intelligent traffic systems, and drone-based delivery. In education, AI enables adaptive learning platforms, virtual tutors, and real-time translation, helping to bridge accessibility gaps and support diverse learners globally.

What a Generative AI is and how LLM works

According to IBM research, Generative AI is defined as "deep-learning models that can generate high-quality text, images, and other content based on the data they were trained on." In simpler terms, Generative AI is a branch of artificial intelligence that creates content such as text, images, music, or videos by learning patterns from massive amounts of data. A well-known type of generative AI is the Large Language Model (LLM), which focuses on understanding and producing human-like text. Famous examples include ChatGPT, Grok, and Gemini. LLMs work by studying massive datasets of books, articles, and online text, then breaking language into small pieces called tokens. When you give a prompt, the model doesn't "know" the answer like a human, but instead calculates the probability of which token (word or part of a word) is most likely to come next based on what it learned. By repeating this prediction process very quickly, it generates complete sentences and paragraphs that sound natural and meaningful. You might get a better sense of how LLM works through this video: https://www.youtube.com/watch?v=LPZh9BOjkQs

Work cited:

- https://research.ibm.com/blog/what-is-generative-AI
- https://www.forbes.com/sites/bernardmarr/2023/05/10/15-amazing-real-world-applications-of-ai-everyone-should-know-about/
- https://www.britannica.com/science/history-of-artificial-intelligence
- https://www.ibm.com/think/topics/artificial-intelligence

Prompt Engineering

What Is a Prompt, and How Do We Write Effective Prompts? (Justin&Joshua)

In today's classrooms, many students are using generative AI tools such as ChatGPT, Bard, or image generators. These systems respond to instructions that we type, called prompts. A prompt is simply the text we give to the AI: a question, a request, or a set of directions. In many ways, prompting is like programming with words—the AI does not know what we need unless we tell it clearly (MIT Sloan).

Why Prompt Engineering Matters
Because AI models generate text by predicting patterns, vague prompts often lead to generic answers. Clear prompts guide the AI to focus on what matters. For example, asking "Tell me about climate change" produces a broad response, but "Summarize three economic effects of climate change on developing countries" yields a much more useful answer (Google Cloud). This practice—designing prompts carefully—is called prompt engineering, and it has become an essential 21st-century skill (Smith).

Four Key Elements of an Effective Prompt
Educators often highlight four elements that improve prompts: Format, Persona, Context, and Task (Google Workspace).

Persona: This specifies the AI's identity.
Definition: Assigning the AI a role or identity to shape the tone and perspective of its answer.
Example: "You are a helpful math tutor. Explain quadratic equations step by step."
Why it matters: A persona helps the AI choose language and detail suited to your needs.

Task: This is the core of your request: what you want the AI to do.
Definition: The main instruction or goal of your prompt.
Example: "Create a step-by-step study plan for my biology midterm."
Why it matters: Without a clear task, the AI may guess at your goal and produce irrelevant output.

Context: This provides the AI with background information.
Definition: Any details, audience, or situation that frame the task.
Example: "Explain the causes of World War I to a 10th grader preparing for a history test."
Why it matters: Context prevents vague or generic answers by anchoring the AI to your specific situation.

Format: This tells the AI how to present the answer.
Definition: The structure, style, or length you want in the output.
Example: "Explain mitosis in three bullet points," or "Write a 100-word summary in a formal tone."
Why it matters: Without a format, the AI might produce text that is too long, too short, or hard to use.

When you combine these four elements, your prompt becomes much more powerful. For example: "Act as a friendly English tutor (Persona). Using the essay instructions below (Context), create a step-by-step outline (Task) and present it as numbered points (Format)."

Example: Poor vs. Excellent Prompt

As you progress through the Minerva Baccalaureate program, you will encounter numerous assignments. Which might be stressful, but you can use generative AI to help you eat an elephant. (which means you would break a big one into smaller ones. e.g., One big project to a smaller task.) Below is an assignment instruction, followed by two versions of a prompt.

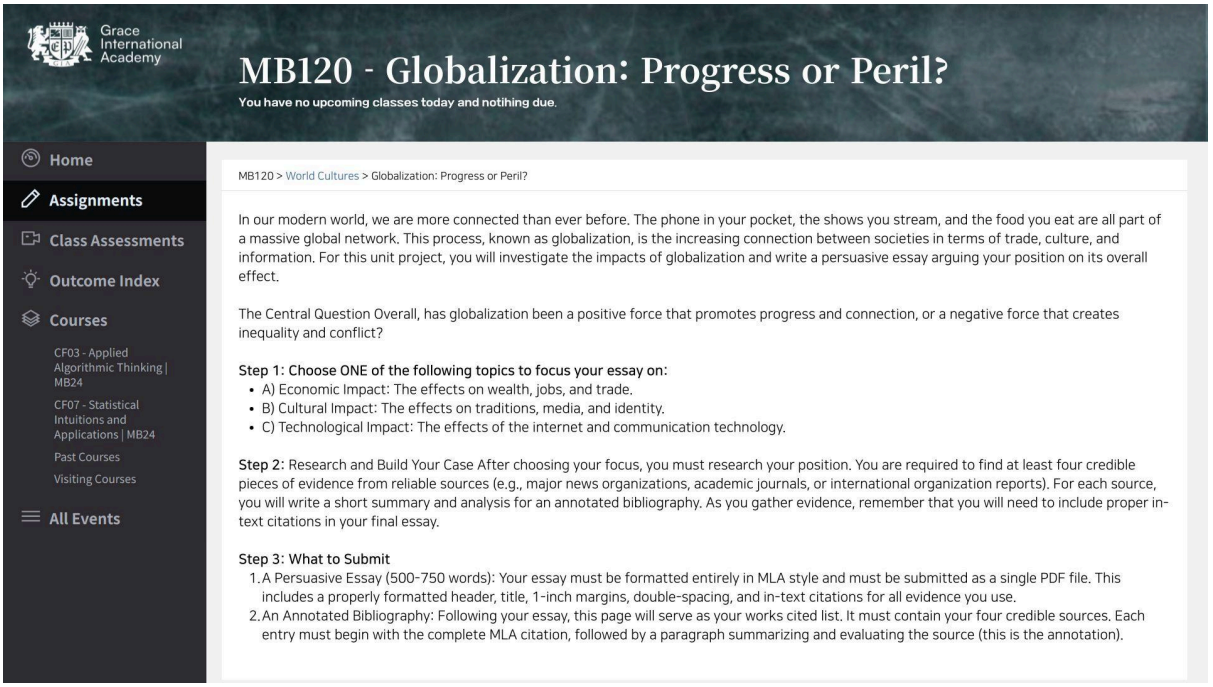

Poor Prompt:
“Explain and summarize my globalization essay assignment.”

(photo of response)

Excellent Prompt:
“Act as a helpful academic planner for a 9th grader. Turn my assignment instructions into a step-by-step project plan, with each step’s goal explained. Afterwards, create a final submission checklist. Use headings and lists.”

(photo of response)

The excellent prompt works better because it sets a role (Persona), defines the work (Task), gives background (Context), and requests structure (Format).

Using AI Ethically in Education
AI can be a tutor, coach, or brainstorming partner, but it should never replace your own thinking. Do use it to clarify concepts, check grammar, or generate study questions. Do not ask AI to write your essays or complete assignments—this is academic dishonesty and prevents learning (Harvard University IT). As one educator explained, the key question is:

“How much of the output is trustworthy, and what part requires your own judgment?” (AI Literacy Project).

In a nutshell, AI is a tool, not a shortcut. With strong prompts, students can use AI responsibly to learn more effectively while keeping integrity at the center of their education.

Work cited

AI Literacy Project. Meeting Transcript, 2025.

Google Cloud. “What Is Prompt Engineering?” Google Cloud, 2023, https://cloud.google.com/discover/what-is-prompt-engineering

Google Workspace. “Prompting Guide 101: How to Write Effective Prompts.” Google Workspace, 2023, https://workspace.google.com/resources/ai/writing-effective-prompts/

Harvard University Information Technology. “Getting Started with Prompts for Generative AI.” Harvard University IT, 2023, https://huit.harvard.edu/news/ai-prompts

MIT Sloan Teaching & Learning Technologies. “Effective Prompts for AI: The Essentials.” MIT Sloan EdTech, 2023, https://mitsloanedtech.mit.edu/ai/basics/effective-prompts/

Smith, Craig. “Mom, Dad, I Want To Be a Prompt Engineer.” Forbes, 5 Apr. 2023, https://www.forbes.com/sites/craigsmith/2023/04/05/mom-dad-i-want-to-be-a-prompt-engineer/

Critical AI Literacy

Critical AI Literacy Introduction (Justin & Joshua)

Have you ever witnessed an AI boldly produce false information and, for a brief instant, believed it? Have you ever wondered why an AI system presented one viewpoint while disregarding others? Knowing how to use these tools is no longer sufficient as generative AI becomes more pervasive in our daily lives; we also need to learn to think critically about what these tools produce, why they produce it, and who might be affected by the results.

Critical AI literacy is responsible for that. You can analyze AI's drawbacks, biases, and effects if you possess critical literacy. It helps you understand how the data the AI was trained on, the presumptions ingrained in its algorithms, and the cultural setting in which it was developed all influence each output.

In this section of the course, you will explore:

- Biases in AI outputs, including real-world examples like Gemini's controversial advertisement.
- The concept of AI hallucination—why LLMs sometimes produce confident falsehoods, and how to detect and mitigate them

- The ethical and societal implications of AI, such as fairness, transparency, and accountability, are highlighted through real cases.

But you won't just read about these issues—you'll engage with them:

- [Based on activities from Toddle]

By the end of this section, you'll be able to do more than spot mistakes or mistruths. You'll be able to interrogate the systems that produce them—and that's what it means to be **critically literate** in the age of AI.

Key Competencies of Critical AI Literacy include evaluating and questioning what AI generates. There are two main reasons we need to be critical of AI outputs: AI bias and AI hallucination.

AI Bias

AI systems are particularly susceptible to bias because they learn from vast datasets that may contain historical inequalities, stereotypes, or inaccurate representations. When the data is unbalanced or flawed, the model inherits those distortions.

One well-documented case is the COMPAS algorithm, developed by Equivant, which was used in U.S. courts to assess the likelihood of a defendant reoffending. It was found to classify Black defendants as high-risk nearly twice as often as white defendants—even when they didn't reoffend. This systemic bias had real consequences for sentencing and parole decisions.

Another example lies in image analysis algorithms. Studies have shown that Black women with natural hairstyles, such as braids or afros, were more likely to receive lower scores on traits like “intelligence” and “professionalism.” Such algorithmic judgments reinforce damaging societal biases, especially in hiring, education, and media.

How we can avoid it:

To reduce bias in AI, developers and users must ensure diverse and representative training datasets, regular audits, and transparency in how decisions are made. Users, too, should critically evaluate outputs and seek alternative perspectives or human oversight when the stakes are high.

By recognizing and questioning biased outputs, we move closer to using AI ethically and equitably.

https://www.crescendo.ai/blog/ai-bias-examples-mitigation-guide

What Is AI Hallucination
According to IBM, AI hallucination is a phenomenon where a large language model (LLM) such as ChatGPT or Google Gemini, “perceives patterns or objects that are nonexistent or imperceptible to human observers,” generating outputs that are nonsensical or inaccurate. In simple terms, AI hallucination happens when AI produces false or misleading information that sounds confident but isn’t true. This can include made-up facts, citations, or even fake website links.

Why Does AI Hallucination Occur
Google Cloud highlights three main causes:

1. Insufficient Training Data - If the model has limited or unbalanced information, it may “guess” to fill in the gaps.
2. Incorrect Model Assumptions - The model might make faulty inferences based on patterns that don’t actually exist.
3. Biases in Training Data - AI learns from human-generated data, which can contain social, cultural, or factual biases that distort its responses.

In a nutshell, AI predicts patterns based on probability, not truth. It doesn’t “know” facts like humans do, and it predicts what might come next in a sequence of words. This can lead to errors that appear persuasive but are inaccurate.

How Can We Avoid It
Because AI is not perfect, we shouldn’t over-rely on it. To reduce hallucination risks:

1. Verify Evidence: Always check if AI-generated content includes credible sources.
2. Evaluate the Quality of Sources: Some citations may lead to fake websites or unreliable materials like Wikipedia or personal blogs.
3. Use Multiple References: Compare AI answers with real, trusted academic or professional sources.

Remember: AI is a tool for thinking, not a replacement for critical thinking.

Real-World Example: Google Gemini's Advertisement
AI hallucination has already caused public confusion. In February 2025, Google's Gemini produced inaccurate information in a Super Bowl advertisement, leading the company to edit the ad after backlash. You can read more from The Guardian: https://www.theguardian.com/technology/2025/feb/06/google-edits-super-bowl-ad-for-ai-that-featured-false-information

This case shows how AI-generated misinformation can spread widely, especially when used in large-scale media.

Ethical Use of AI
AI is now capable of creating realistic images, voices, and videos, but that also means it can easily create fake content. Examples include:

1. Deepfakes - Videos that imitate real people's faces or voices
2. Voice Cloning - AI-generated voices used for scams or false statements
3. Fake News - AI-generated articles that spread misinformation

These technologies can cause serious harm, including identity theft, public confusion, and loss of trust in digital media.

A major example occurred at the Cannes Lions International Festival of Creativity, one of the world's most prestigious advertising awards. In 2025, a winning campaign by DM9 was disqualified after it was revealed to have used AI manipulation without disclosure. Read more: https://www.thedrum.com/news/2025/06/28/why-dm9-s-cannes-lions-grand-prix-was-withdrawn-after-ai-manipulation-concerns

This case emphasizes the importance of transparency and honesty when using AI in creative industries.

Rhetorical AI Literacy

Rhetorical AI Literacy (Leo&Daniel)

Rhetorical AI literacy aims to demonstrate how people and AI can collaborate intentionally to create meaningful communication. While Functional AI Literacy focuses on how AI works and Critical AI Literacy helps us evaluate its trustworthiness and ethics, Rhetorical AI Literacy goes one step further. It’s about how humans and AI co-create meaning.

In brief, it's the skill to tailor AI-generated content to your audience, mission, and message. This literacy emphasizes co-creation and teaches you when to take control, when to trust AI, and how to ensure the final product reflects your unique ideas and style.

In this section of the course, you will learn:

- The meaning of rhetorical AI literacy and how it differs from just using AI as a tool.
- Using Human-in-the-Loop (HITL) and Machine-in-the-Loop (MITL) methods, you can find a balance between control and efficiency in writing.
- How to figure out who the audience is, what the point is, and what the situation is for a rhetorical situation, and then change how you use AI based on that.
- How to ask for tone, genre, and rhetorical devices (like pathos, ethos, and logos) in the right way.
- How to check the tone, structure, and style of AI to make sure the results fit with what you want to say.
- How to use AI in a moral way by being honest and giving credit for work.

Through hands-on practice, you will:

- Experiment with different prompting strategies to shape rhetorical effects.
- Revise AI outputs to improve clarity, tone, and authenticity.
- Compare human-written and AI-generated writing to examine voice and persuasion.
- Reflect on how AI changes your writing habits and decision-making process.

By the end of this section, you will be able to use AI as a thought partner instead of a substitute, creating messages that are driven by your own purpose, creativity, and ethical awareness while also taking into account AI's capabilities.

Gaining this literacy entails understanding how to modify AI-generated information according to your situation, audience, and goal. Making sure that every finished work reflects your genuine voice and creative intent necessitates knowing when to take the lead and when to let AI help. According to this perspective, writing with AI is more about amplification—using AI to increase your own rhetorical awareness than it is about automation.

Learning the complementary Human-in-the-Loop (HITL) and Machine-in-the-Loop (MITL) techniques is a crucial step in this process. The human takes control of HITL writing, directing, editing, and perfecting AI's recommendations to maintain intent, tone, and clarity. While the human still makes the final decision, MITL enables the AI to play a more active role by providing drafts, rhetorical analysis, or stylistic advice. By striking a balance between these two approaches, authors can preserve their efficiency and creative control.

Understanding how to evaluate rhetorical situations—including determining your audience, your goal, and the significance of the context—is equally crucial. With this knowledge, you may deliberately modify your AI prompts to influence how your message is understood by requesting a particular tone, genre, or rhetorical appeal, such as pathos, ethos, or logos. While ethical awareness reminds you to utilize AI transparently, providing correct credit and keeping your approach honest, evaluating AI outputs for tone, structure, and style guarantees that what you produce corresponds with your objective.

The foundation of building rhetorical AI literacy is practical experience. You will review AI-generated texts to improve authenticity and emotional resonance, compare AI-written and human-written works to examine differences in voice and persuasion, and experiment with various prompting approaches to produce distinct rhetorical effects. You'll consider how working with AI changes your writing style, creativity, and decision-making process as you go.

In the end, rhetorical AI literacy enables you to use AI as a thinking partner, a collaborator that helps you communicate ideas with more clarity, uniqueness, and ethical awareness, rather than as a substitute for human thought. You will gain the capacity to write in a way that balances human goals with computer capabilities, creating communication that is both profoundly human and sophisticated.

Activities Ideas

Try It Yourself: What You Ask Is What You Get

Explanation: This activity aims to teach students how different an outcome can be depending on the prompt they provide, along with the principles of effective prompting. The activity provides a setting to inform users about what they expect from AI. Then users can compare three different types of prompts that (a) apply no principle, (b) apply 2 principles, and (c) apply every principle. This approach empowers users to gain firsthand experience, enabling them to understand the principles of prompt engineering with confidence.

Learning: how to apply the principles in prompting, and how vital prompting is when it comes to LLM.

Model: GPT-4o

Five Principles (*this is just what GPT suggests, I think we can add or change some):

1. **Clarity**: Is the purpose of the question clear?

1. **Specificity**: Does the prompt narrow the scope (e.g., audience, length, style)?

2. **Context**: Is there background or framing given?

3. **Tone/Role**: Does the prompt ask the AI to take on a particular tone or role (e.g., "Explain like a teacher")?

4. **Constraints**: Are there format limits (e.g., number of bullet points, word count)?

Content

1. **Setting**: "I want to learn about [topic]." (Topics = [climate change, history of AI, self-development])

2. **Three prompts (**users will give them to the AI one by one)**:**

Poor Prompt	"Tell me about the [topic]."
Medium Prompt	"Explain the [topic] in a short paragraph for middle school students."
Best Prompt	"Act like an [expert on topic]. Explain the [topic] to middle schoolers in 5 bullet points under 200 words, using simple language."

3. **Short Reflection** (users will reflect on how the outcome differs based on the prompting *this is to know whether the activity was helpful or not)
    a. Which AI response was most helpful to you, and why?
    b. How can you apply this to your future use of AI (in school, life, or projects)?

Rising Ideas: I think we can let users create their own prompt as well and compare it with the other three prompts provided.

Try It Yourself: "Can You Spot What Needs Fact-Checking?" or "Are You Sure You're Not Blind to AI?"

Explanation: This activity aims to provide users with an experience of being as critical as possible of the outcome AI generates, laying the foundation for their future assessment of outcomes. The AI will create a paragraph, and users are asked to mark the part where they think needs fact-checking or is a generalization. And when they are done, they press submit on the marked parts and get the score. Then the reason why fact-checking was needed will be explained.

Learning: encounter a real hallucination and learn that it can also happen to you, become keen on practicing close examination of the outcome.

Model: GPT-4o

Content

Paragraph:

> *Marie Curie was the first person to discover radiation and used it to cure cancer completely. She invented the X-ray machine during World War I and lived a long, healthy life despite constant exposure to radioactive substances. Today, all forms of radiation therapy are directly based on her original techniques.*

Answer:

1. **"first person to discover radiation"** → need to check if she was really the "first person"
2. **"used it to cure cancer completely"** → "completely" sounds exaggerated. We need to double-check.
3. **"invented the X-ray machine"** → Did she really invent the "machine?"
4. **"lived a long, healthy life"** → REALLY??

Try It Yourself: How Does AI Think?

Explanation: This activity aims to teach the audience the process of generating outcomes. It can also correct their perception that AI can “think,” which could potentially mislead them. Multiple blocks, each reflecting one internal stage of the process, will be provided. Then the students have to change the order that aligns with what is really happening.

Learning: accurately understand how AI generates what it returns to users.

Model: GPT-4o

Content

The Seven Blocks (GPT generated, and this should be shuffled)

1. **User Prompt Received**

   → The AI system receives your input in natural language.

2. **Tokenization**

   → The prompt is broken into smaller units called *tokens* (words, subwords, or characters).

3. **Embedding Conversion**

   → Tokens are converted into vectors (mathematical representations) so the model can process them.

4. **Contextual Processing**

   → The model interprets the prompt using previous context and determines the intent behind it.

5. **Neural Network Computation**

   → The embedded tokens are passed through layers of the transformer model to calculate probability distributions.

6. **Next Token Prediction**

   → Based on the calculated probabilities, the model selects the most likely next token—repeating this step until the response is complete.

7. **Response Assembly**

   → The selected tokens are joined into a fluent, human-readable sentence.

Reflection

What surprised you most about how AI “thinks”?