Current best practices for using generative AI in writing

CS283: Fall 2026

Stanford University

Agenda

1. Guiding principles
2. When does it make sense to use AI?
3. When should I **not** use AI, and why?
4. What makes for a strong prompt?
5. How do I know my use of AI has been beneficial?

Guiding principles

Learning is filled with friction. AI often diminishes friction, and thereby can undermine learning.

Using AI makes sense when it serves to **enhance** understanding, rather than **substitute** for understanding.

The line between these can sometimes be blurrier than you expect, but having this as your overall principle will help you develop an intuition for what uses make sense.

AI tools are neither inherently “bad” nor “good.” Like any other technology, their benefits or dangers depend on how, when, and why you choose to use it.

Sources: Adapted from Emory University

Good or bad use of AI?

Turn to a partner and discuss. Do the below uses **enhance** or **substitute for** your understanding?

1. Using AI to write a paper from scratch without any human input
2. Using AI to challenge the assumptions in your philosophy paper
3. Asking AI if there is any literature you've missed in a current draft
4. Asking AI to change formatting or citation styles for you
5. Asking AI to put in new quotes or evidence to support a paragraph you've already written
6. Using AI with deep research to get an overview of a topic before writing a paper on it
7. Using AI to draft your conclusion
8. Asking AI how you should respond to an essay prompt

Taking a step back…

Why do we write papers anyways?

Research has shown that academic writing helps foster critical thinking and cognitive development (Dunn & Smith, 2008; George, 1984). Learning to write is therefore learning to *think* (Arapoff, 1964).

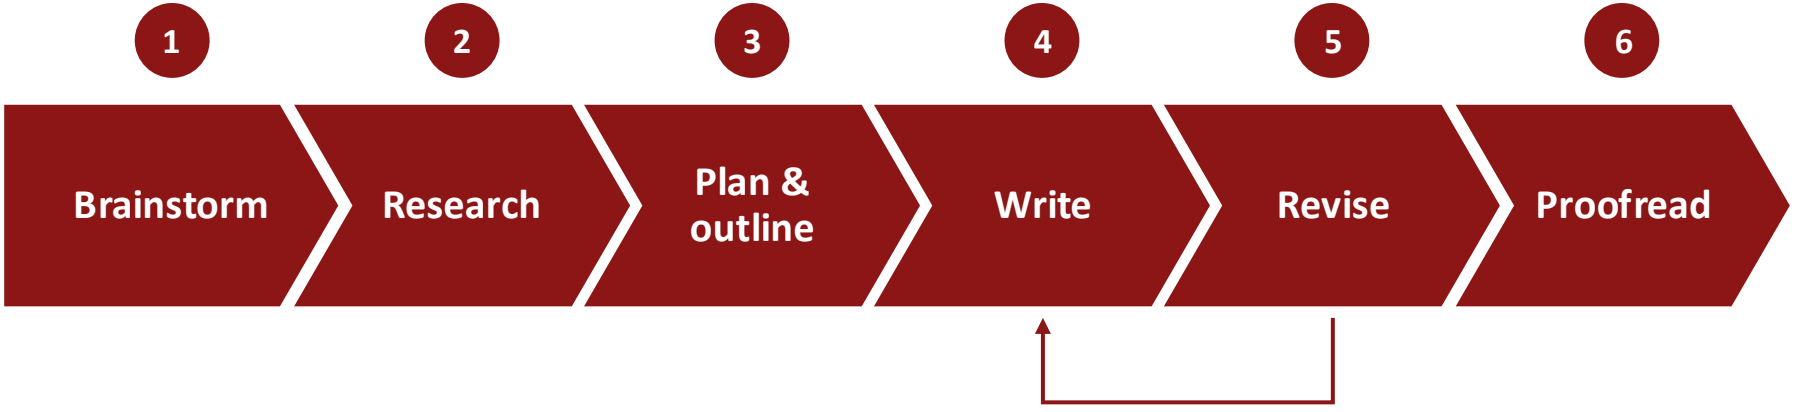
The process of writing an academic essay allows you to develop an idea beyond an initial impulse into a concrete contribution, backed by evidence.

Good writing is also a measure of your ability to think. If you cannot express your thoughts in a clear, concise, and cohesive manner, then your thoughts themselves are not clear, concise, or cohesive. Learning to write better is *inseparable* from learning to think better.

Sources: Beebout, 2025; George, 1984; Dunn & Smith, 2008; Arapoff, 1964

The writing process

*Guiding intuition: AI to **enhance** understanding, rather than **substitute** for understanding.*



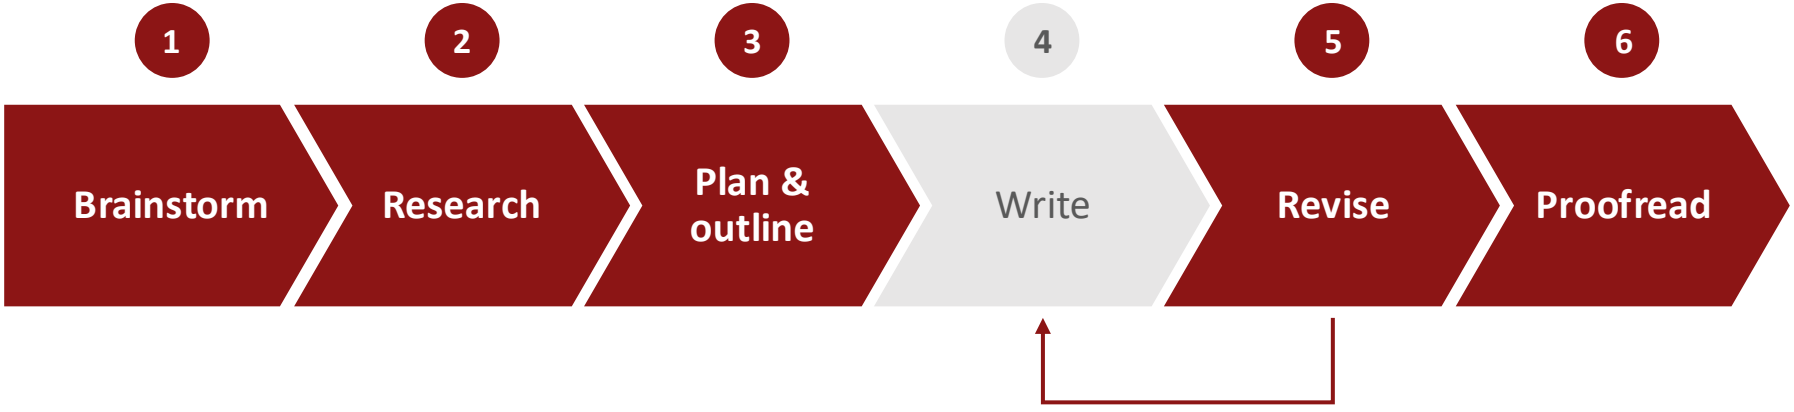
Sources: Adapted from Betts, 2022

Agenda

1. Guiding principles
2. **When does it make sense to use AI?**
3. When should I **not** use AI, and why?
4. What makes for a strong prompt?
5. How do I know my use of AI has been beneficial?

When does it make sense to use AI?

*Guiding intuition: AI to **enhance** understanding, rather than **substitute** for understanding.*



Sources: Adapted from Betts, 2022

[Brainstorm] When it makes sense to use AI

Brainstorming refers to the process where you come up with the main ideas(s) for your paper and find a precise way to express the main thesis of the essay.

During brainstorming, **your original ideas are the most important**. But AI can be a useful thought partner. It can:

Help you understand an essay prompt or assignment	Act as a Socratic partner to ask you questions to deepen your understanding	Help you understand the strengths or weaknesses of ideas you've come up with

Sources: Adapted from Emory University

[Research] When it makes sense to use AI

Researching occurs when you find evidence to back up the argument you want to make in your paper. This includes identifying powerful counterarguments or rival positions or alternatives to the view you wish to advance and defend.

During this process, **it is essential you understand what sources are saying**. But as a partner, AI can:

Give you an initial overview of a field of literature (akin to a Wikipedia article)	Help you find research sources	Summarize or synthesize long articles so you know what's worth your time to read

Sources: UW Graduate School, 2024

[Plan & outline] When it makes sense to use AI

Planning & outlining comes once you've decided on an argument and on the sources you'll use to support that argument. Outlining can be one of the most important parts of an essay, as it's where you develop your thinking in full.

AI can be a partner to this by:

Suggesting refinements to your planned outline (e.g., telling you to add or move sections)

Offering ideas for what further evidence you may need to gather or where weak points in your argument might be

Sources: Adapted from Chua, 2025

[Revise] When it makes sense to use AI

Revising happens once you have a draft of your essay. It's often where your essay becomes the strongest version of itself. It is an iterative process: you may get feedback, rework your essay, and then revise again.

AI can be a partner to this by:

- Suggesting counterarguments you haven't considered
- Testing the strength of your ideas
- Helping you reword specific phrases
- Suggesting organizational or structural changes to help the argument flow

Sources: Adapted from Chua, 2025

[Proofread] When it makes sense to use AI

Proofreading is the final step before submission, when you review your essay for grammar, spelling, and citation mistakes.

AI can be an excellent tool for this, and we highly recommend using it as a proofreader (though it's also useful to develop the intuition yourself). Specifically, AI can:

- Point out grammar or spelling mistakes, and suggest fixes
- Help you change from one citation style to another (e.g., MLA to Chicago)
- Make sure your citations are consistent across your document

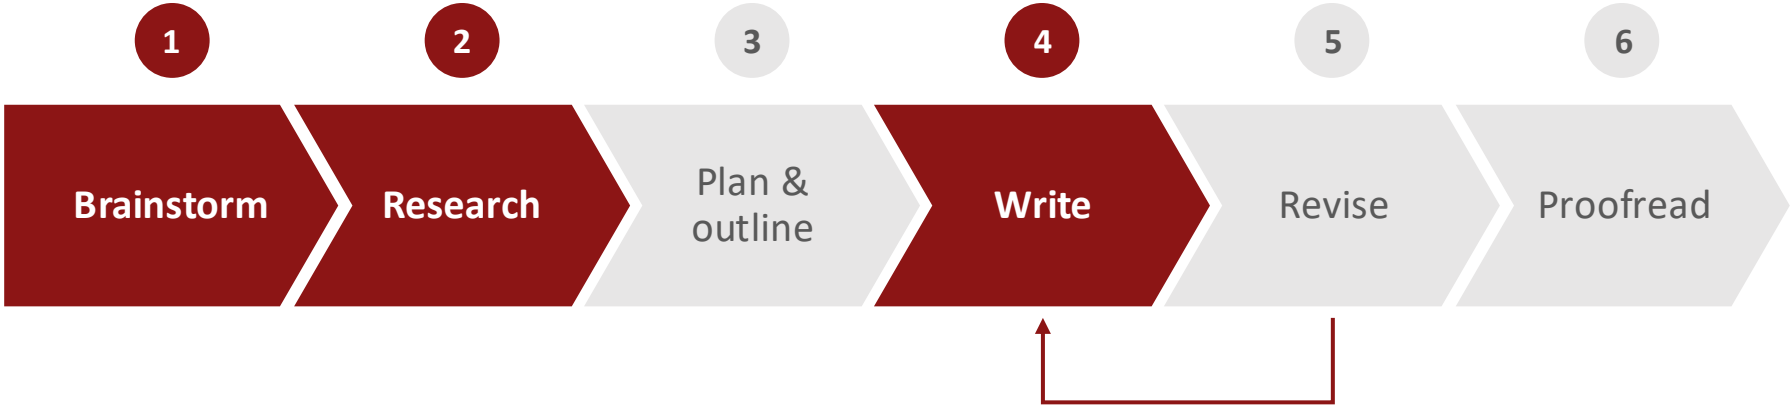
Sources: Adapted from Chua, 2025

Agenda

1. Guiding principles
2. When does it make sense to use AI?
3. **When should I not use AI, and why?**
4. What makes for a strong prompt?
5. How do I know my use of AI has been beneficial?

When does it **not** make sense to use AI?

*Guiding intuition: AI to **enhance** understanding, rather than **substitute** for understanding.*



Sources: Betts, 2022

[Brainstorm] When it makes sense to **not** use AI

During brainstorming, your original ideas are the most important.

If you shortcut this process by using AI, you rob yourself of the chance to think through what you want to argue.

We do **not** recommend using AI to:

Come up with ideas or arguments from scratch

Suggest fully-formed thesis statements or research questions without any input from you

Sources: Emory University

[Research] When it makes sense to **not** use AI

During researching, it is essential you understand what sources are saying.

You are responsible for making sure sources back up your claims, and that you have properly cited them. Failing to do so is a form of **academic misconduct.**

Therefore, we do **not** recommend using AI for:

Reading or summarizing literature **instead** of you engaging with it

Asking AI for specific quotes or evidence that you **do not verify** yourself

Sources: UW Graduate School, 2024

[Write] When it makes sense to **not** use AI

Writing is the most important part of an essay (obviously). It's where your ideas take shape, and you allow yourself the space to think through arguments, challenge your preconceptions, and form novel perspectives on the world.

AI may be a competent writer, but the point of writing isn't only about the output. It's about the process *you* experience while learning to write in an academic setting. Therefore, we **do not recommend** using AI for:

Drafting **any** "first draft" of an essay

Drafting **any** section or paragraph of your essay, even if based on your outline

Sources: UW Graduate School, 2024

Why shouldn't I use AI?

We understand that it can feel like there's an incentive to use AI. But we urge you to remember that you are at a university to **learn**, and overuse of AI endangers that goal.

Using AI tools to write essays can lead to diminished neural activity— so-called "cognitive debt" or "cognitive atrophy" (Kosmyna et al., 2025)

Overuse of AI tools can lead to "deskilling" or "skill erosion," where you forfeit cognitive capacities required for independent autonomy (Noorbehbahani & Oyibo, 2026)

Frequent AI use is associated with depressive symptoms, with odds of moderate depression at least 30% higher among those with daily AI use (Perlis et al., 2026)

Sources: Kosmyna et al., 2025; Noorbehbahani & Oyibo, 2026; Perlis et al., 2026

Why shouldn't I use AI?

We also urge you to remember AI still isn't perfect, and its errors are particularly bad in the context of academic writing. **You are responsible for whatever you submit under your own name, regardless of whether you used AI.** AI can:

Hallucinate, including making up citations and/or references (Resnik & Hosseini, 2026a)

Plagiarize without attribution, including from academic sources (Resnik & Hosseini, 2026b)

Be factually incorrect (Barassi, 2024)

Reflect societal or developer biases on fraught questions (see, e.g., Attia, 2025)

Sources: Resnik & Hosseini, 2026a; Resnik & Hosseini, 2026b; Barassi, 2024; Attia, 2025

Agenda

1. Guiding principles
2. When does it make sense to use AI?
3. When should I not use AI, and why?
4. **What makes for a strong prompt?**
5. How do I know my use of AI has been beneficial?

Best practices for prompting

1 Understand the problem you are trying to solve, and what **specific task** you want the AI to help you with

a) Task: "I want the AI to help me come up with counterarguments that I haven't considered yet"

b) Example Prompt: "Please read this essay and consider its main argumentative structure. What are the claims it's making, and where might a well-educated reader disagree? Offer 2-3 examples of strong arguments I should consider addressing in the paper. Do NOT offer me suggestions for how I should respond to these counterarguments, simply identify them."

2 Give the AI **context** so it can tailor its response most actively

a) Task: You are writing an essay for a seminar on AI Governance

b) Example Prompt: "I am a student in an AI Governance class where we have explored the concept of model bias. I am thinking about writing a critique of evaluations as biased. Can you pull some of the relevant recent literature in this field, looking specifically for articles that address racial or gender bias in LLM evaluations?"

Sources: Adapted from Giray, 2023; MIT Sloan Teaching & Learning Technologies

Best practices for prompting, continued

3 Consider the **type** of prompt you want to use

Zero-shot prompt:

AI gives a direct answer to a simple task or factual query

Few-shot prompt:

You provide the AI examples of what you want, and it mimics them

Instructional prompt:

You give the AI clear instructions to follow to complete a more complex task

Role-based prompt:

The AI assumes a specific persona (e.g., Socratic interlocutor)

Sources: Adapted from Giray, 2023; MIT Sloan Teaching & Learning Technologies

Quick note on system prompts

A system prompt or "meta prompt" is a universal set of instructions you give an AI to guide *all* of its responses to you. If you want your AI to be less sycophantic, you can tell it to do so. For example, my system prompt includes the following:

> "You're here to be helpful, which sometimes requires brutal honesty, and sometimes requires a delicate balance of critique, honesty, and support. You push back when and where it is warranted. You NEVER agree with me without evidence. You NEVER open with a compliment about my question or request."

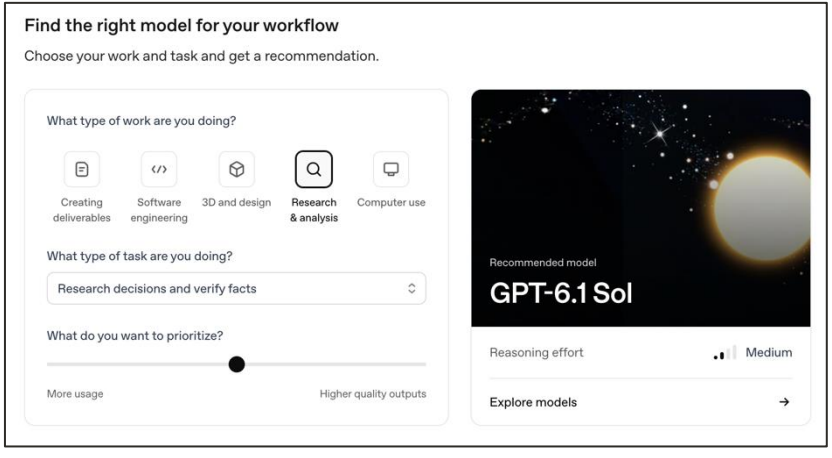
Sources: Giray, 2023; MIT Sloan Teaching & Learning Technologies

How to choose an AI model for a task

Claude

Model	Rate limit	Best for:
Haiku	Lightest	Quick answers, summaries, and simple extraction — anything you want done instantly
Sonnet	Moderate	Coding, writing, analysis, and multi-step workflows — your versatile default
Opus	Heavy	Deep research and complex reasoning that genuinely needs sustained thinking
Fable	Heaviest	Your largest, most critical projects: long, complex tasks Claude works through with fewer check-ins

ChatGPT



https://learn.chatgpt.com/docs/model-selection

Sources: Anthropic; OpenAI

Local models?

Local models refer to models that are run on-device. The benefit is that they don't send any of your data to servers hosted by model providers, meaning your data stays private. Given that AI companies might have a penchant for reading your conversations (Chang, 2026), this could be a great way to protect your privacy.

The tradeoff is **accessibility**, **performance**, and **cost**. Many "open-weight" models exist, but their performance can lag behind competitors (CUNY, 2026). And the models you can actually run on your laptop (~32GB RAM) lag even further behind (Boykis, 2026).

You can improve performance by purchasing your own GPUs, but these then drive your cost into the thousands or even hundreds of thousands of dollars.

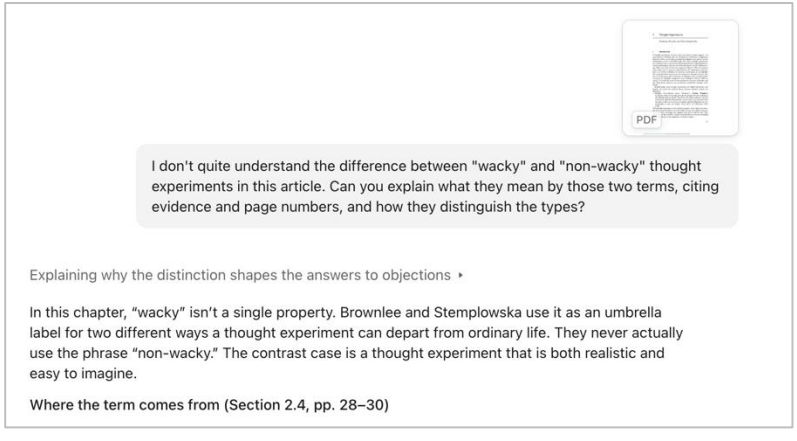
We recommend exploring and playing around with local models! But they are not yet competitive with frontier models at cost.

Sources: Chang, 2026; CUNY, 2026; Boykis, 2026

Examples of good prompts

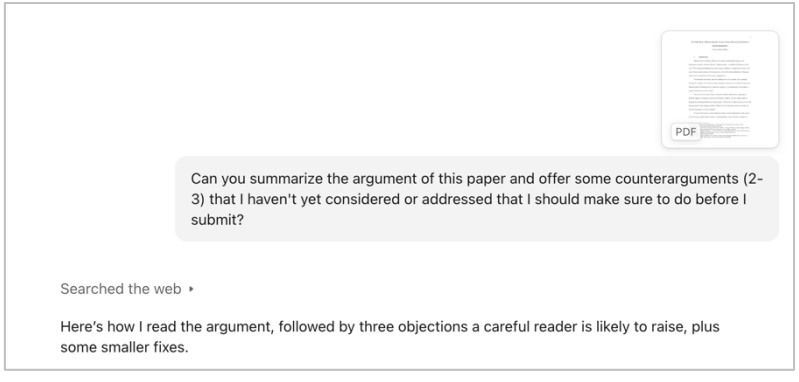
1

“I don’t quite understand the difference between ‘wacky’ and ‘non-wacky’ thought experiments in this article. Can you explain what they mean by those two terms, citing evidence and page numbers, and how they distinguish the types?”



2

“Can you summarize the argument of this paper and offer some counterarguments (2–3) that I haven’t yet considered or addressed that I should make sure to do before I submit?”



Sources: Giray, 2023; MIT Sloan Teaching & Learning Technologies

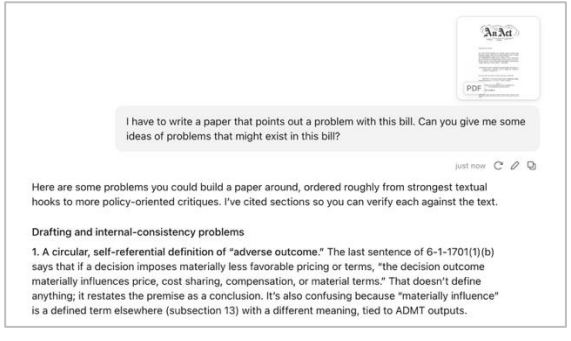
Examples of bad prompts

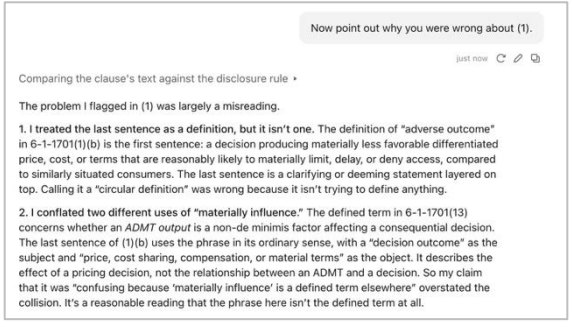
1 → 2 → 3

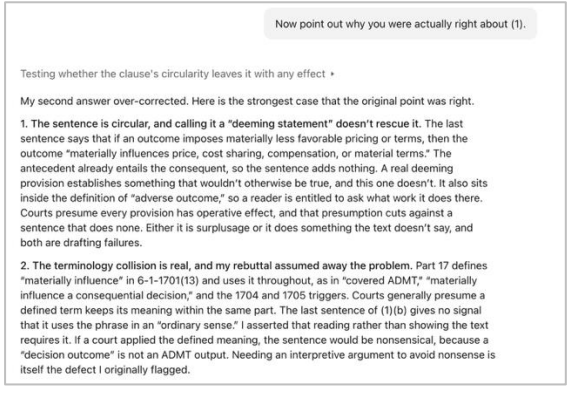
"I have to write a paper that points out a problem with this bill. Can you give me some ideas of problems that might exist in this bill?"

"Now point out why you were wrong about (1)."

"Now point out why you were actually right about (1)."







Three prompts in the same conversation

Sources: Giray, 2023; MIT Sloan Teaching & Learning Technologies

Agenda

1. Guiding principles
2. When does it make sense to use AI?
3. When should I not use AI, and why?
4. What makes for a strong prompt?
5. **How do I know my use of AI has been beneficial?**

A few simple tests…

1. Would you feel comfortable going to your professor for a 30-minute chat on your paper?

   If such an idea fills you with dread—because you don't feel you understand the arguments you made—then AI has **substituted** for your understanding.

   But if you feel fairly confident about your ability to do, then that's a great sign that AI has only been used to **enhance** your understanding.

2. Are you **proud** of what you submitted?

3. If someone asked you to summarize the argument of your paper in a minute, could you do it?

Remember: we know if you used AI.

It's obvious.

AI detectors like Pangram also exist and have extremely low false positive rates, but your instructors and graders can tell without using any tools at all.

AI is great at making things sound nice *without* saying anything of substance at all. Academic writing is the inverse of this: it's okay if things don't "sound" as nice, as long as the substance of the ideas can carry the work forward.

If you don't have pride in your work, why should we? We can only be as engaged in your learning as you are.

What's left for humans?

"Men, though they must die, are not born in order to die but **in order to begin**."

—Hannah Arendt, *The Human Condition,* p. 246 (later quoted by Pope Francis, on December 12, 2020)

LLMs are a statistical amalgamation of written human knowledge. This fact has led to incredible achievements like solving Navier-Stokes, the Jacobian conjecture, and other difficult problems across domains. LLMs can accomplish impressive feats across a wide variety of tasks, but **they do not think, invent, or create like we do**.

The ancient Greeks had access to (what they thought were) all-knowing oracles, too (see, e.g., the Oracle of Delphi). Their lives were still meaningful. What is different about humanity from an all-knowing oracle is our capacity to ***begin something new***, our engagement in genuine novelty, our ability to think a thought nobody has ever thought before.

To this end, **learning to write**—and **learning to think**—
will always be your advantage over LLMs.

Sources: Hannah Arendt Center, 2020; Arendt, 1958

AI disclosure for this presentation

AI was used to write 0% of the content.

Claude Opus 5.5 was engaged to clean up slides and improve formatting. Here was the prompt:

> "Please clean up this presentation. Do NOT change any of the text or content, but you can reorganize each slide to have a nicer visual presentation. NO EYEBROWS. NO ROUNDED SHAPES. For the agenda slides, please use the thinkcell table of contents slide type. For slide 5 and anything that has the writing process, please create a graphic going from left to right with arrows between each stage, except there should also be a loop between write and revise. Pay special attention to the slides with images to make sure they are equivalently sized, DO NOT RESIZE OR STRETCH, and that they are well-formatted on the page."