AI Literacy, Equity, and the Future of Higher Education: A Call to Action

September 2025

Agenda

1

The What:
Current AI landscape

Together we'll evaluate how Big Tech's agenda has shaped the current landscape and how higher education should respond.

2

The How:
Increase Agency through AI Literacy

Our mission, should we choose to accept it, is to empower ourselves and our students through AI literacy.

3

The Why:
The Future of AI

1 The What: The Current AI Landscape

AN ESTIMATED 69 MILLION NEW ROLES ARE EXPECTED TO BE CREATED AND 83 MILLION DISPLACED. COMMUNITIES ALREADY FACING SYSTEMIC INEQUALITIES ARE THE ONES MOST AT RISK OF BEING LEFT BEHIND.

- *WORLD ECONOMIC FORUM REORT, 2024*

Figure 3: Generative AI will transform work across industries

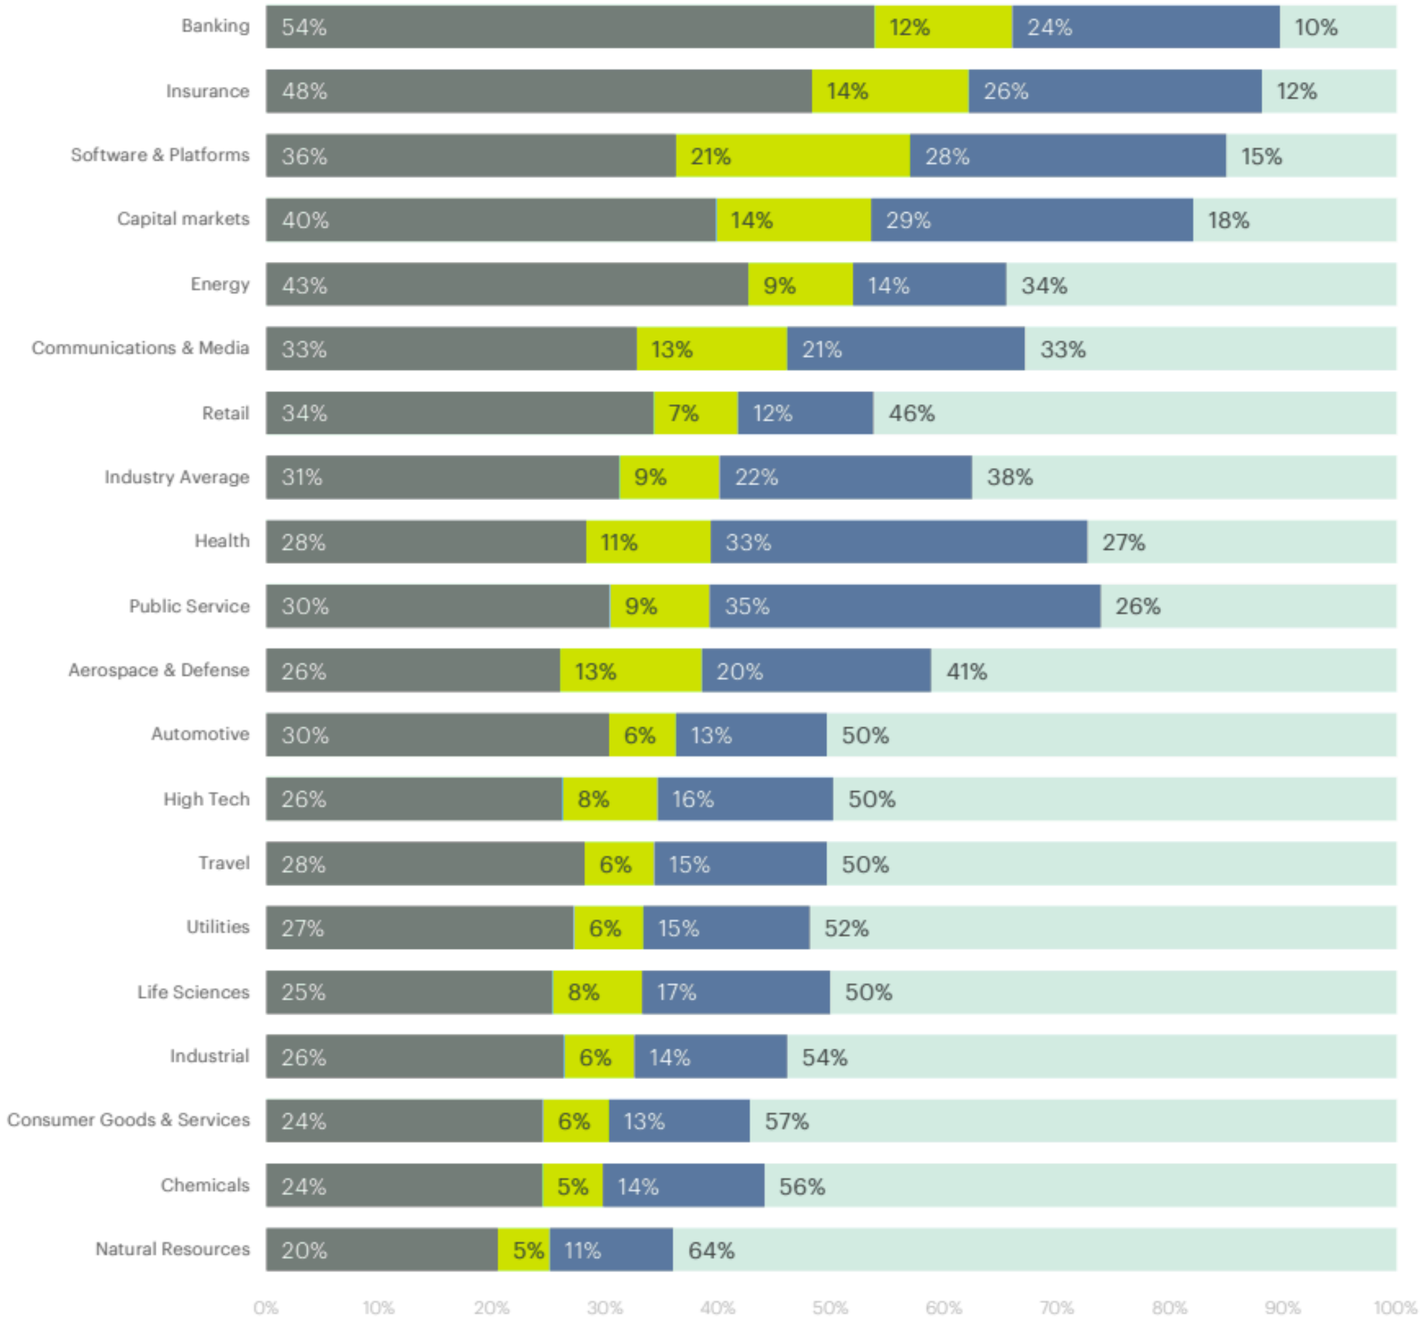

Work time distribution by industry and potential AI impact

Based on their employment levels in the US in 2021

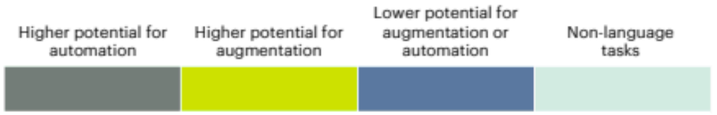

40% of working hours across industries can be impacted by Large Language Models (LLMs)

Why is this the case? Language tasks account for 62% of total worked time in the US. Of the overall share of language tasks, 65% have high potential to be automated or augmented by LLMs.

Source: Accenture Research based on analysis of Occupational Information Network (O*NET), US Dept. of Labor; US Bureau of Labor Statistics.

Notes: We manually identified 200 tasks related to language (out of 332 included in BLS), which were linked to industries using their share in each occupation and the occupations' employment level in each industry. Tasks with higher potential for automation can be transformed by LLMs with reduced involvement from a human worker. Tasks with higher potential for augmentation are those in which LLMs would need more involvement from human workers.

Stanford 2025 AI Index Report

The goal: produce objective insights into AI's technical progress, economic influence, and societal impact.

Top Takeaway - Governments and The Economy are fully vested in nascent AI
1. AI is increasingly embedded in everyday life - *In 2023, the FDA approved over 200 AI-enabled medical devices, up from just six in 2015.*
2. Business is all in on AI, fueling record investment and usage - *In 2024, U.S. private AI investment grew to $109.1 billion.*
3. Governments are increasing their stake in AI through regulation **and** investment - *Globally, legislative mentions of AI rose 21.3% across 75 countries since 2023*

The conversation we need to have about AI isn’t technical - it’s ***ideological***

Technological solutionism, or techno-solutionism, is the belief that complex societal problems can be solved primarily or solely through technological interventions. It involves a tendency to simplify problems and assume that new technologies, such as apps or algorithms, can provide quick and easy solutions, often without fully understanding the root causes of the issues.

“Why and how is the power of Big Tech increasing in the policy process?
The case of generative AI Open Access”

-Shaleen Khanal, Hongzhou Zhang, Araz Taeihagh

Gaslight, Gatekeep, Girlboss

Why don’t we challenge this narrative?

Because we don’t understand that we can.

Big Tech =
Product over Process

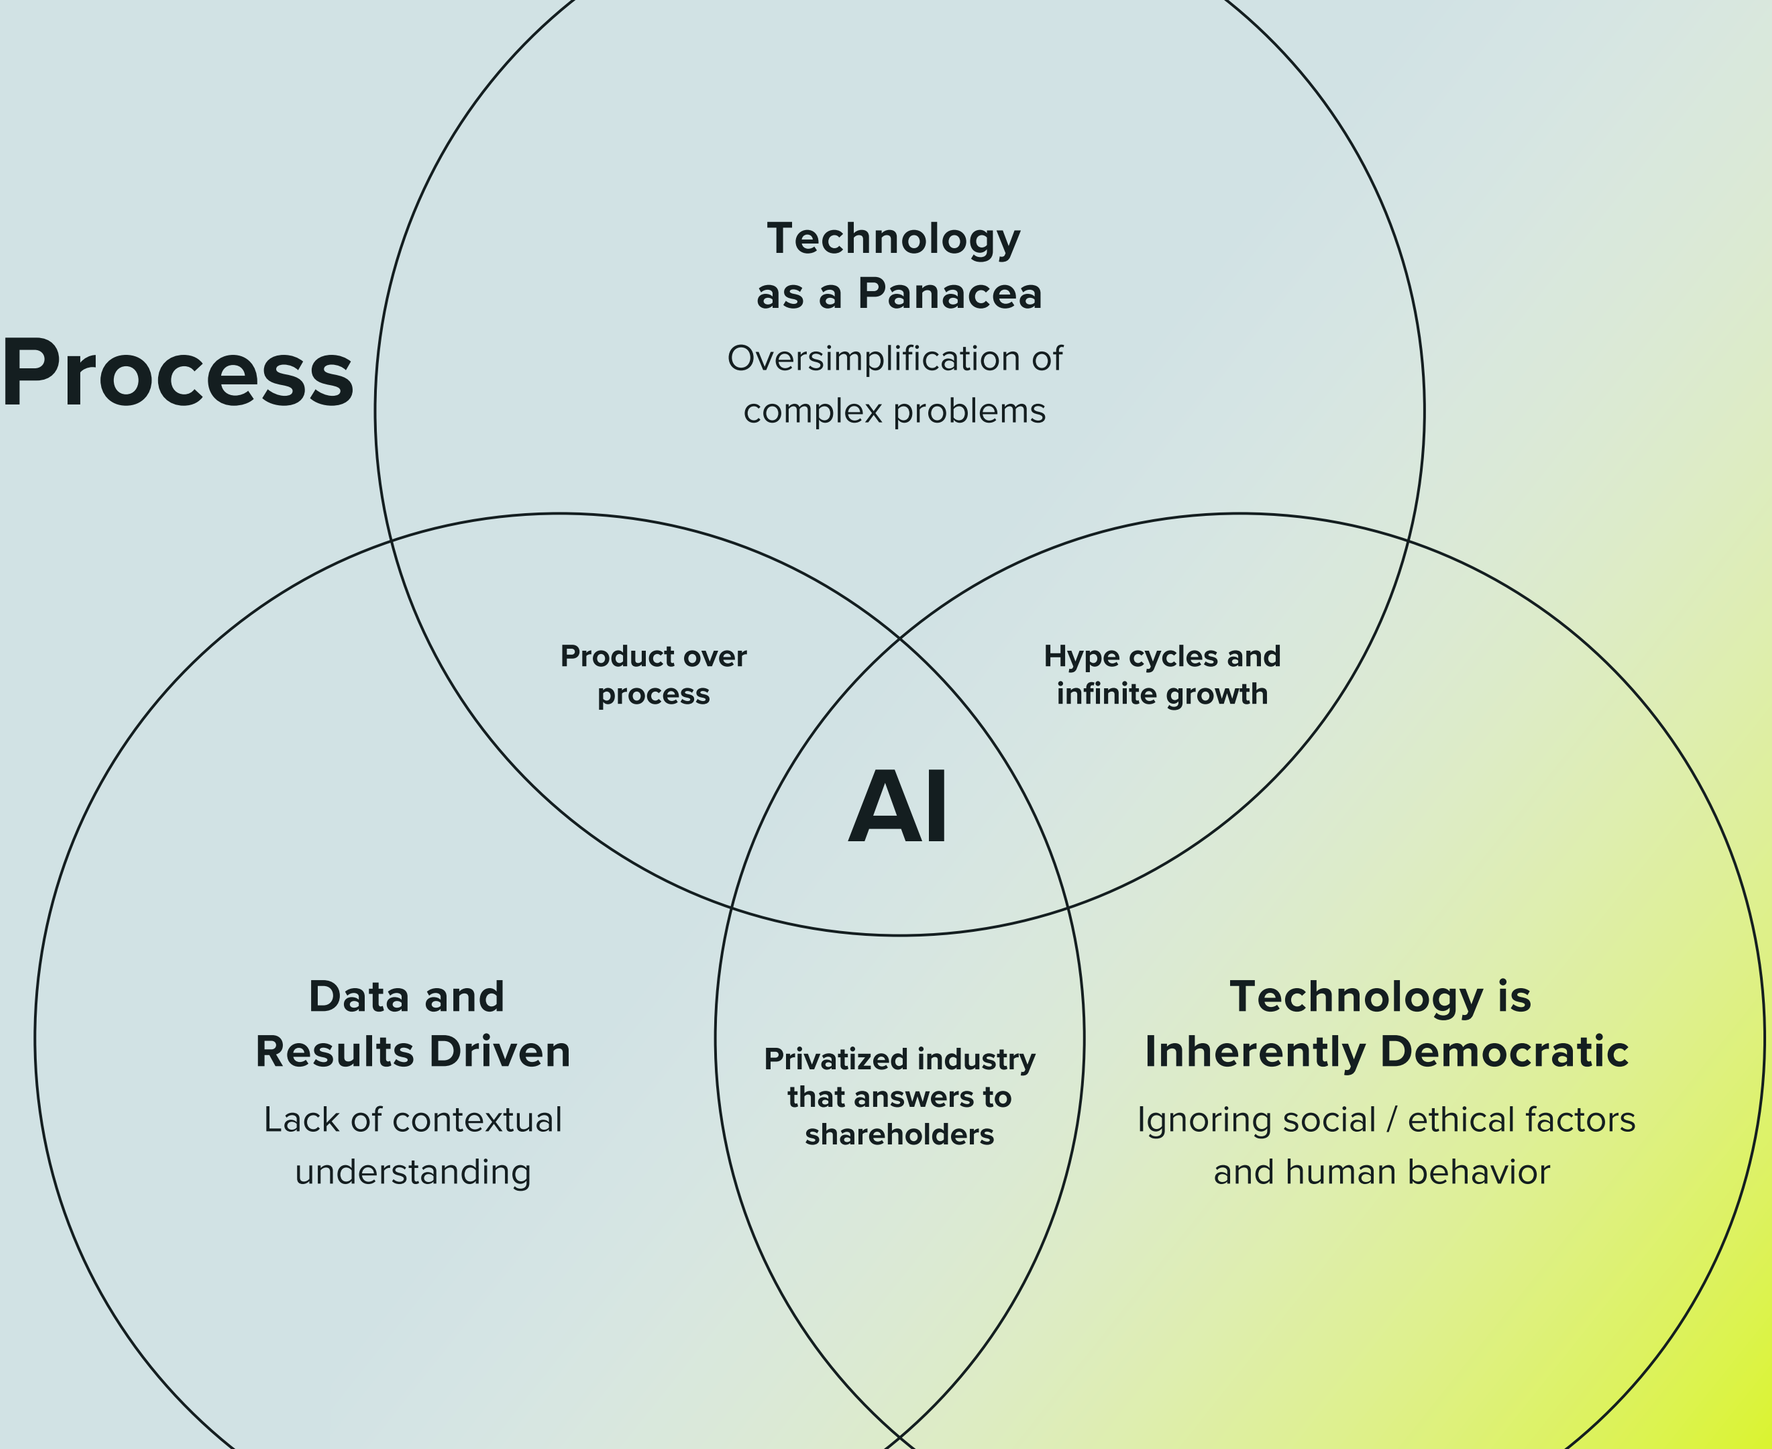

Educators Goals = Process over Product (Profit)

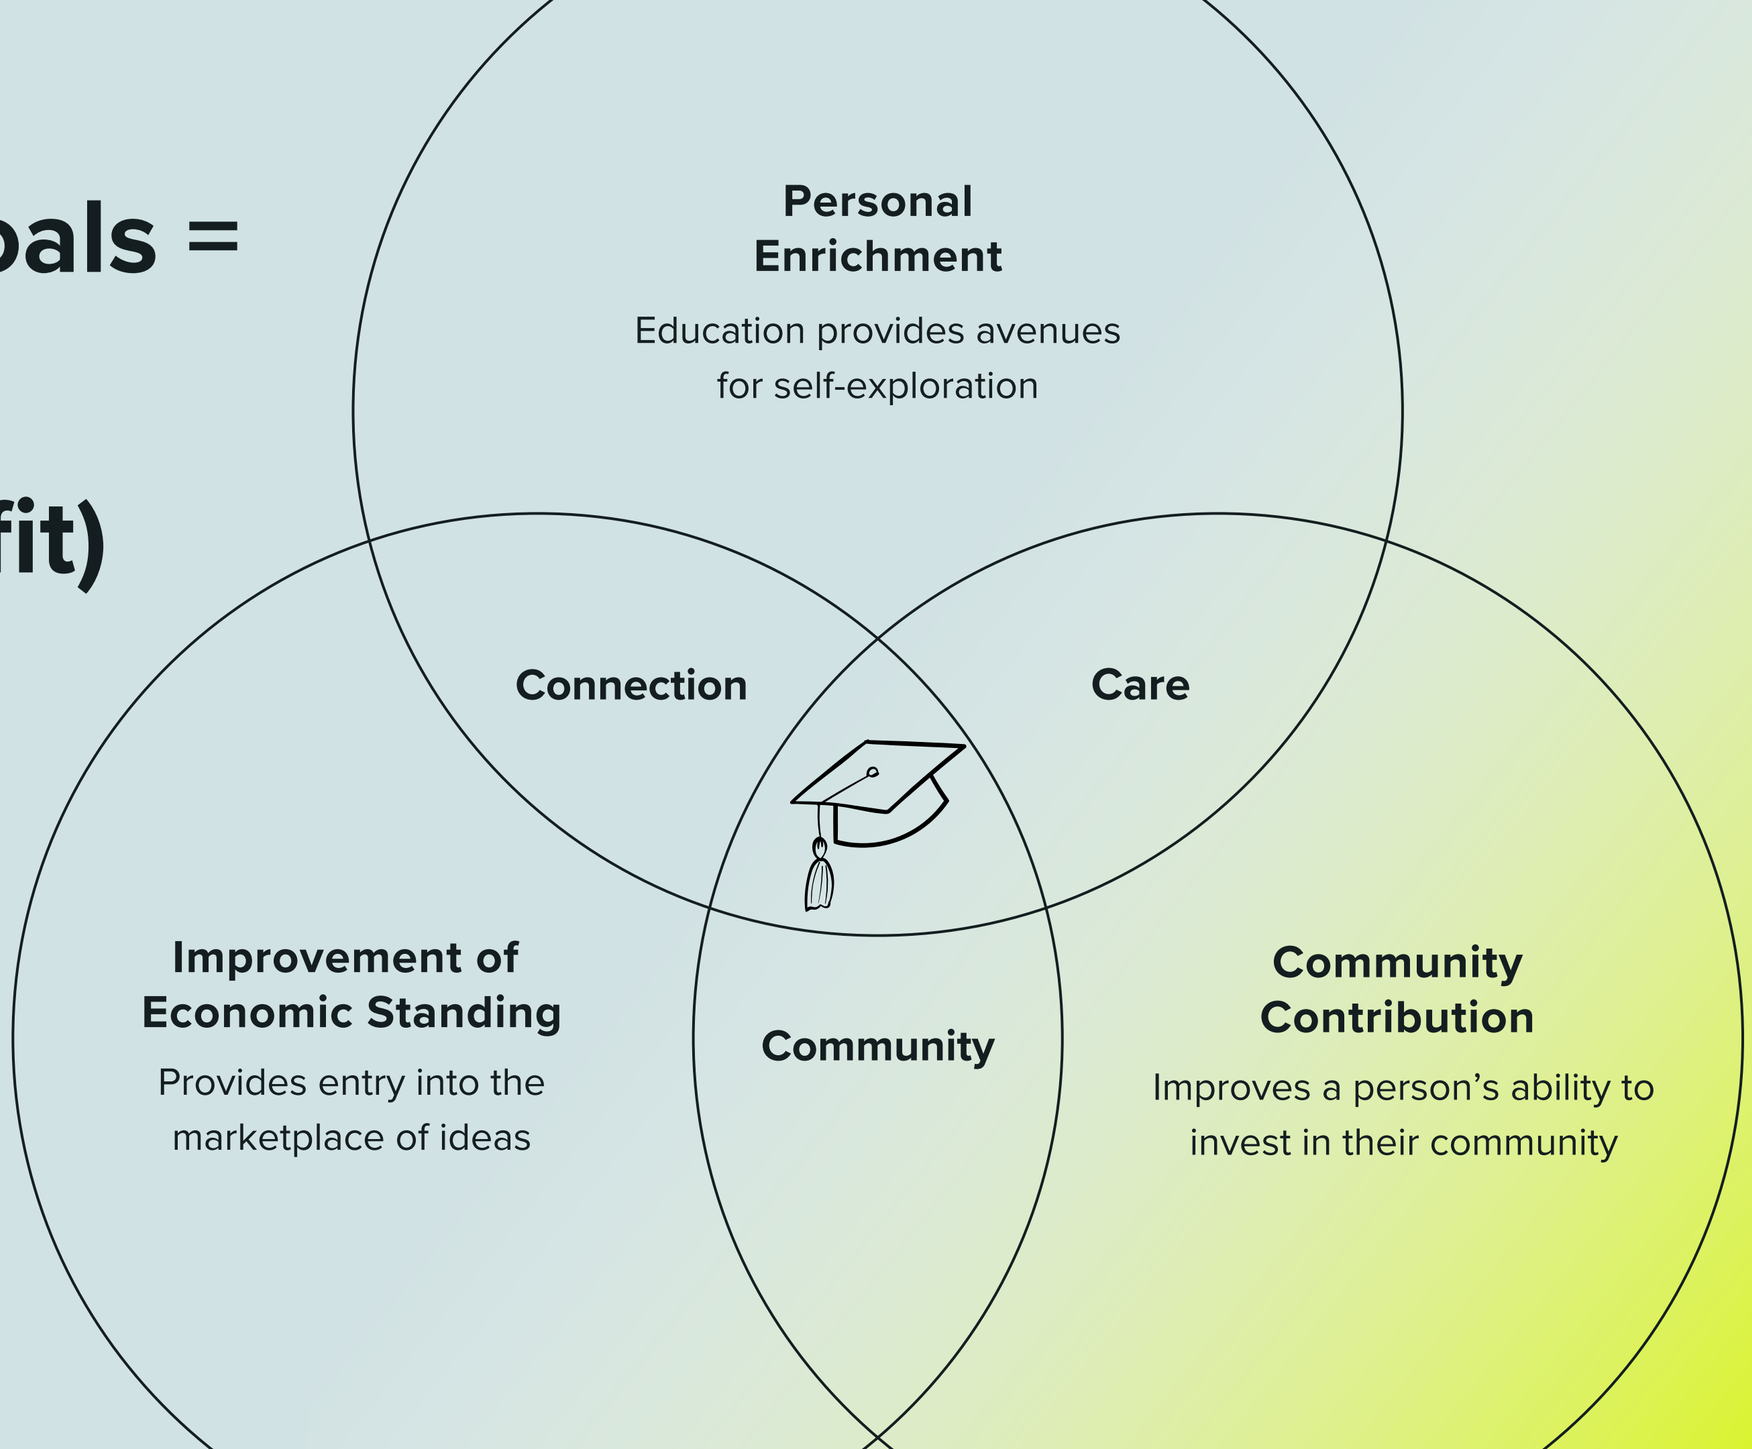

Educators Goals

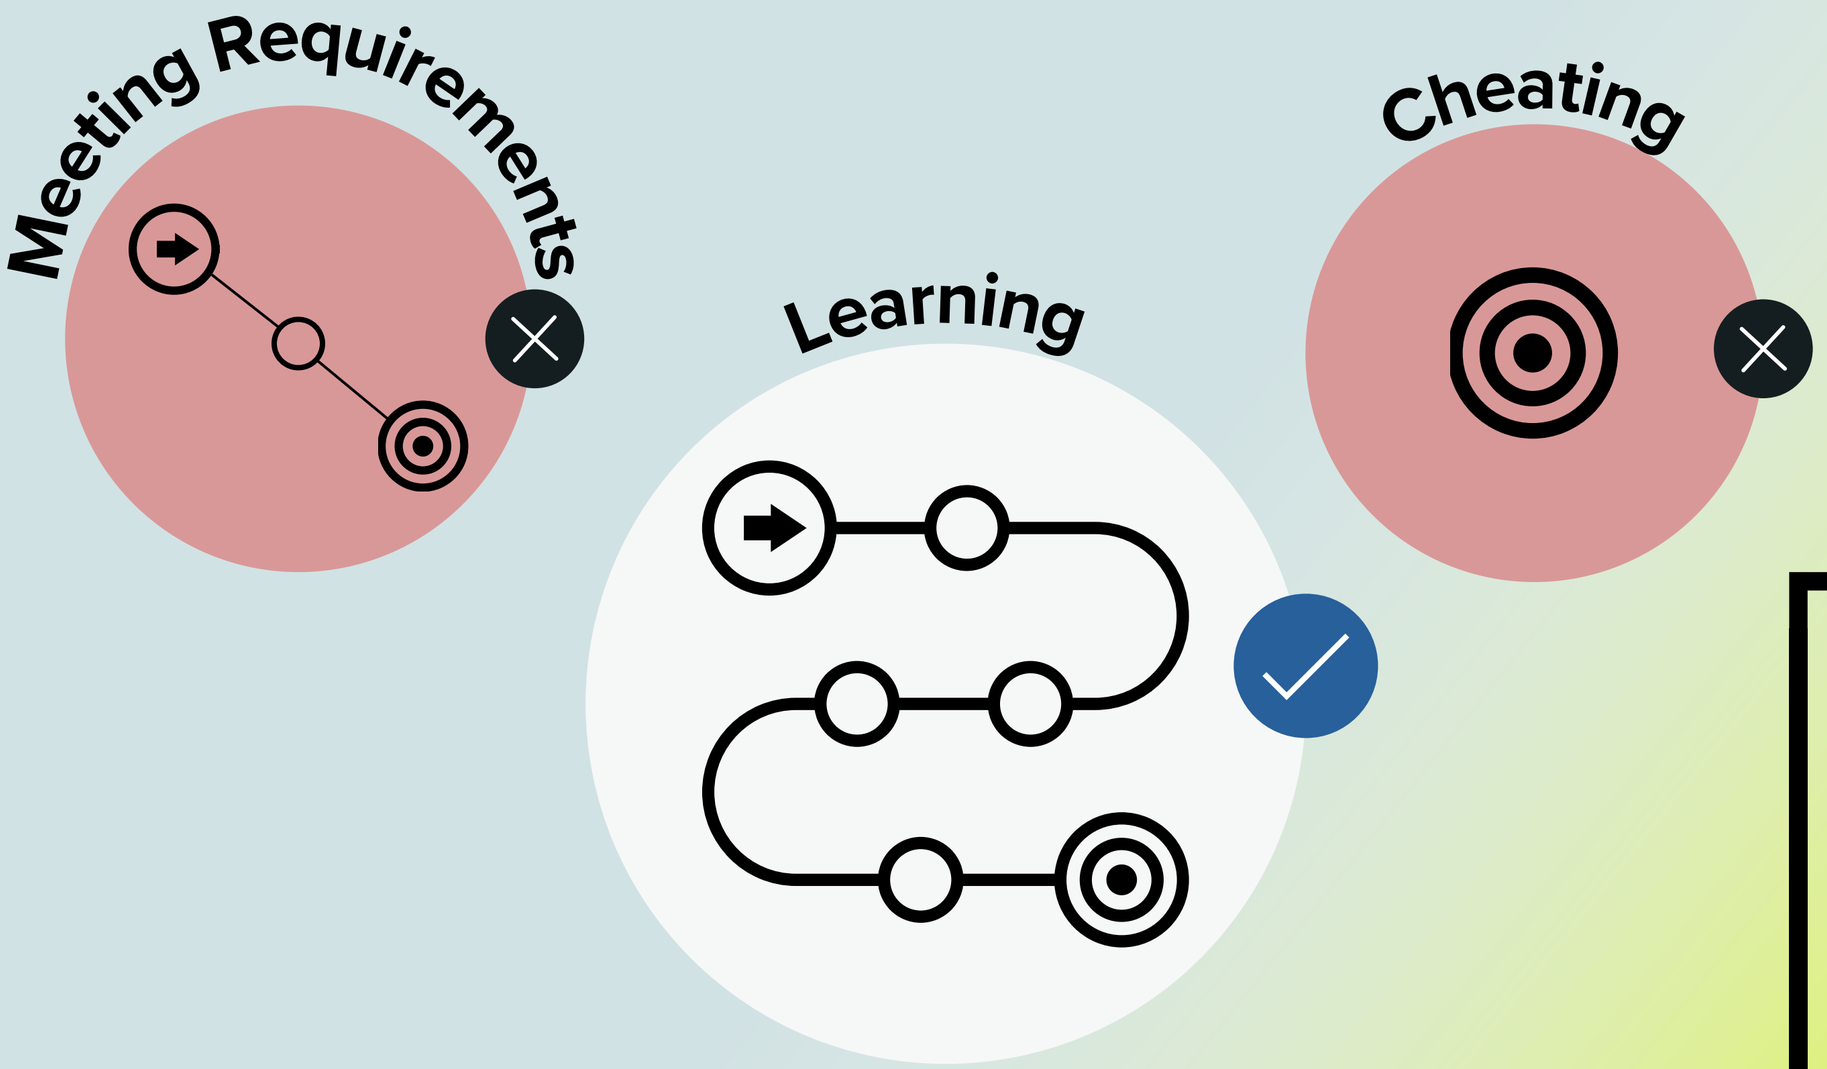

Learning takes place in the *process*.

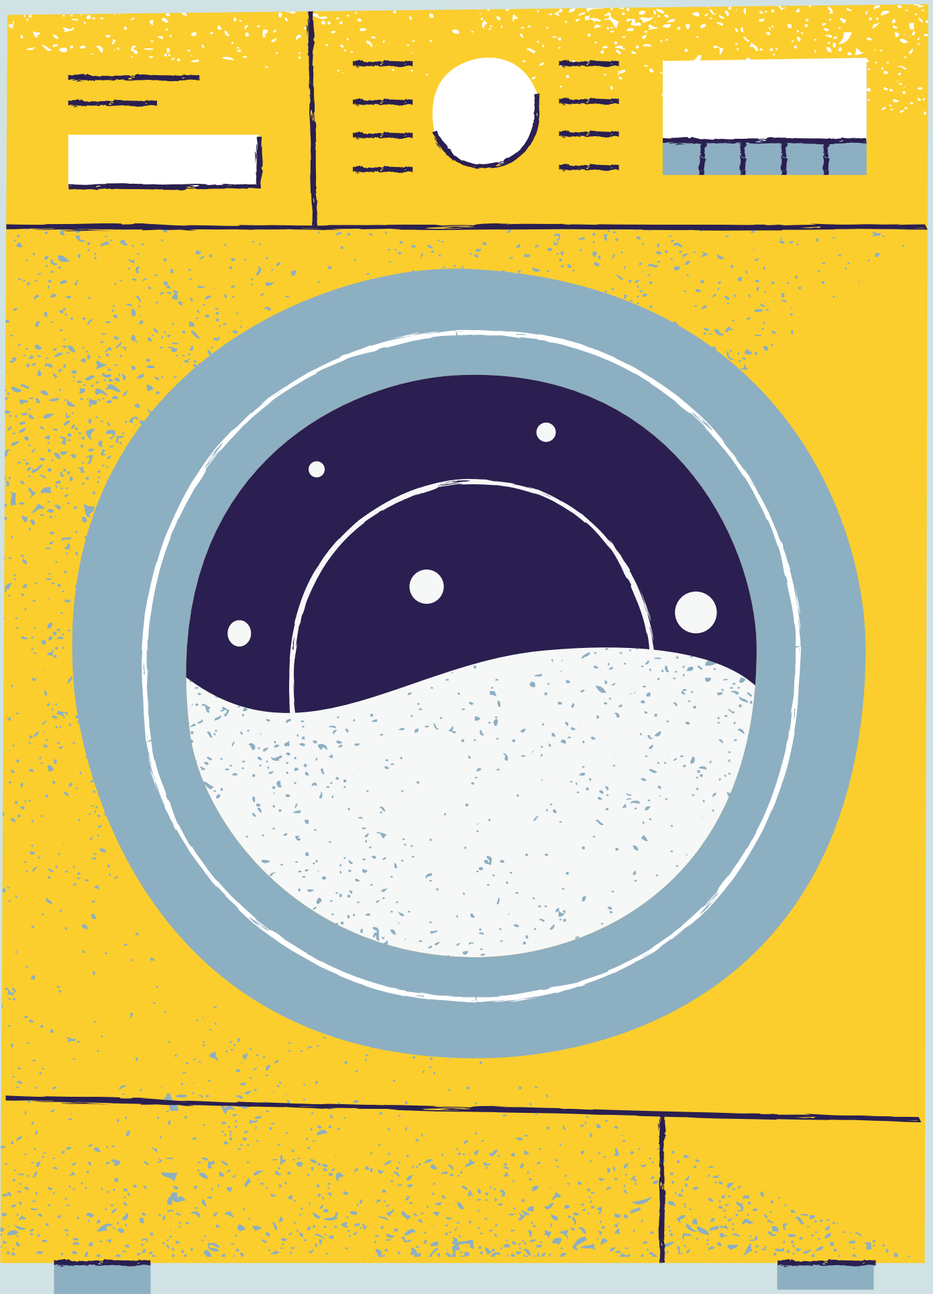

The washing machine

Effrosyni Adamopoulou, et al. “The Household Equipment Revolution.” National Bureau of Economic Research, 1 Mar. 2024, https://doi.org/10.3386/w32253. Accessed 18 Sept. 2025.

Our Mission

Higher Education must enter the public AI discourse and challenge Big Tech's narrative by implementing literacy programs that will *transform passive consumers into citizens*

Shifting the Narrative

Embracing Agency

Understanding we don’t have to accept the narrative

Coalition Building

Engaging with others ready to do the work

Education

Taking the steps necessary to become AI literate

Advocacy

Becoming a professional squeaky wheel (yapper)

Tech Philosophy

Developing a philosophy based on your personal values as an educator while understanding the need for student AI skills

Consistency

Small consistent changes over time

We must plant trees we will never sit under.

2 The How: AI Literacy

AI literacy Defined...?

“AI literacy is urgent—but we lack consensus on what it means...If we don’t actively shape this now, we may find that AI literacy is defined for us in ways that don’t align with our values.”

-James DeVaney - Inside higher Ed, 2025

AI Literacy Literature Meta Synthesis

Top Takeaways
Canonical & Most Influential Frameworks: a. *What is AI Literacy? Competencies and Design Considerations, 2020: Established the seminal, citation-dominant conceptual framework in AI literacy, combining technical/algorithmic, ethical, and socio-technical dimensions for general audiences; serves as the root reference for almost all later work.* b. *AI Literacy: Definition, Teaching, Evaluation and Ethical Issues, 2021: Complementary, widely-cited review mapping classic literacy models to AI (know/use/evaluate/ethics), heavily shaping subsequent frameworks.*
Major Functional Frameworks Emphasize: • **Fundamentals**—The fundamental basics everyone needs to know to understand and interact with AI effectively. • **Integration—**ethical, civic, and digital citizenship issues are embedded at each level, not relegated to side modules. • **Scalable progression**—clear novice/intermediate/advanced scaling. • **International alignment**—frameworks often reference, or are developed with inspiration from, UNESCO, EU DigComp, World Economic Forum, and similar global digital competence movements.
Notable Limitation: • A handful of works promote “AI for the general public,” but sometimes treat ethics/societal impact superficially rather than as true embedded pillars.

AI Literacy IRL

AI literacy refers to the ability to understand, interact with, and critically evaluate artificial intelligence technologies, including their capabilities, applications, and societal impacts.

Technical

How machine learning architecture works and why that matters.

Examples:

- What does GPT stand for?
- What is alignment and why is it an important aspect of training a Large Language Model?

Practical

Using our technical understanding, how do we communicate with AI so that we get the results we want?

Examples:

- Effect prompt engineering methods
- Picking the right AI tool for the job

Evaluative

Critically evaluating the application and output of AI tools.

Examples:

- Rubrics for evaluation
- Understanding terms like “AI hallucinations”

Ethical

What are the ethical impacts of engaging with these tools and how can we do our best to create a base for ethical AI deployment and use?

Technical *Example*

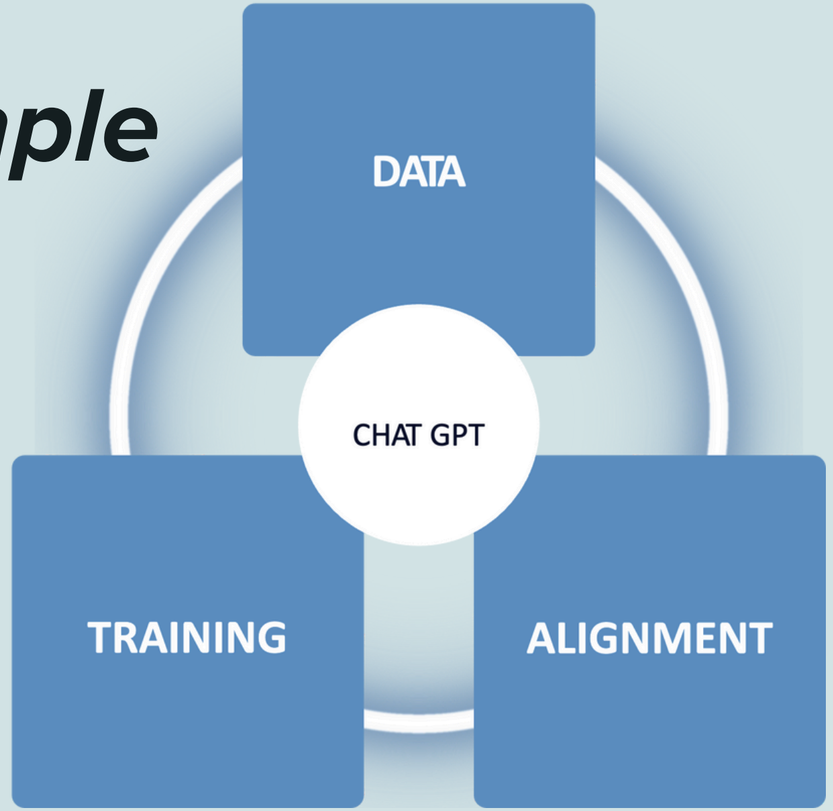

Data

The information that was used to train the model

Training

How the data was organized, weighted, and contextualized by the model.

Alignment

Does this model do what we want it to do, the way we want it to do it?

GPT-3 DATA

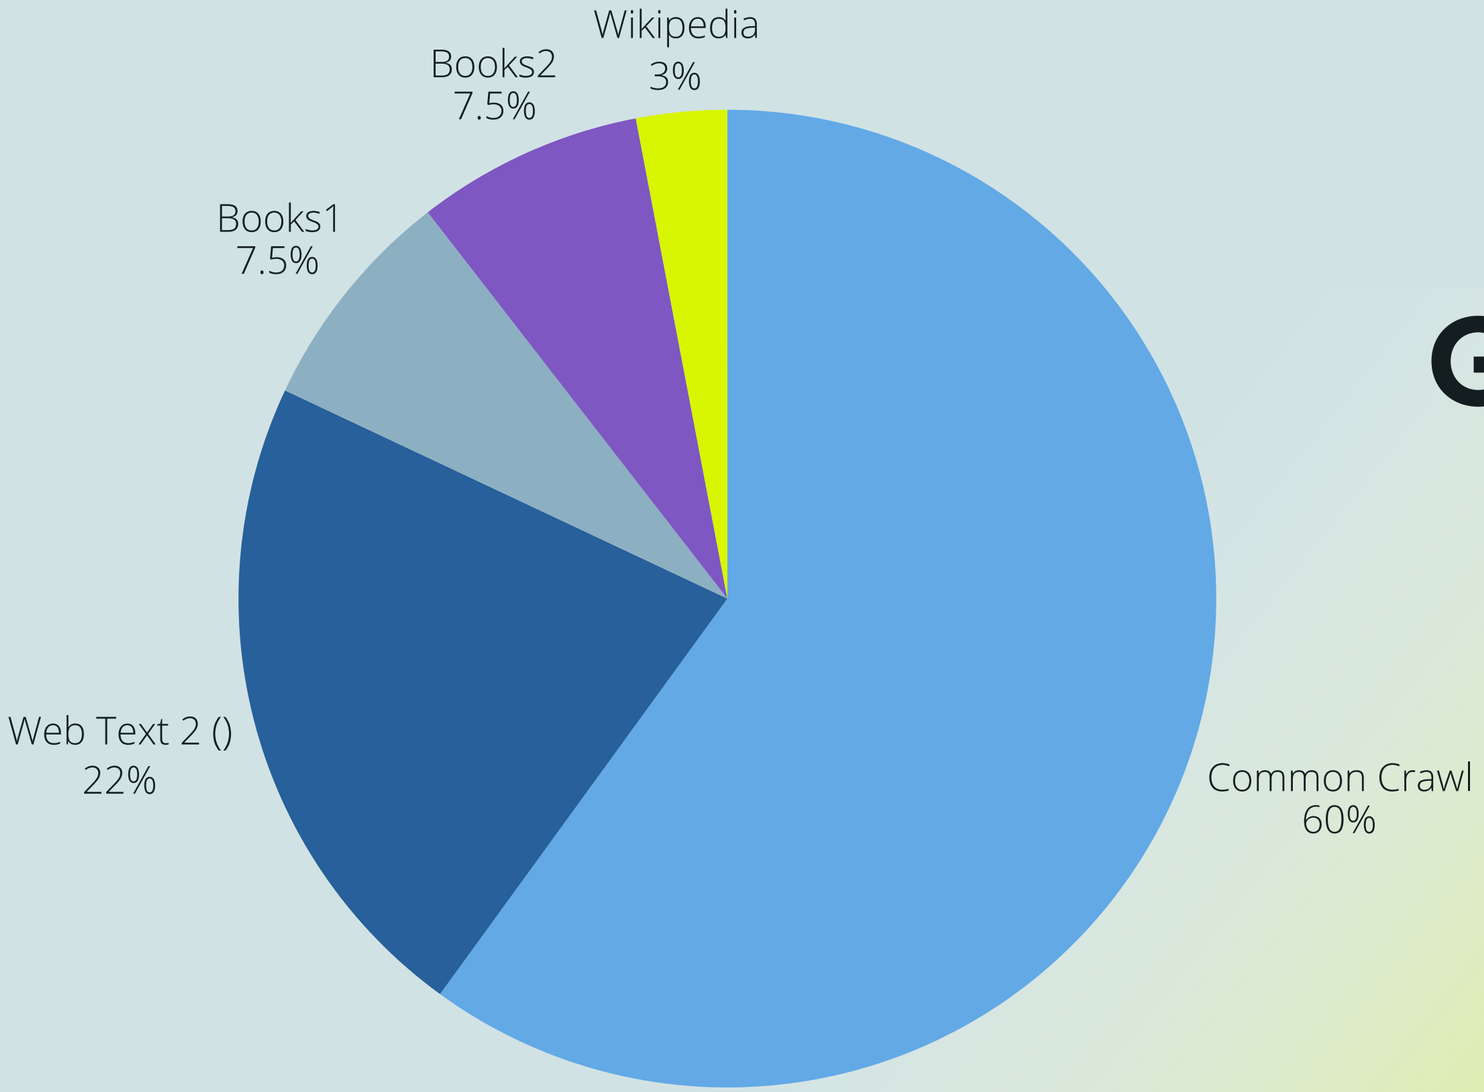

Data Sets Used to Train Chat GPT

GPT-4 DATA

“Given both the competitive landscape and the safety implications of large-scale models like GPT-4, this report contains no further details about the architecture (including model size), hardware, training compute, dataset construction, training method, or similar.”

-GPT-4 Technical Report, OpenAI

Practical *Example* - Proper Tool Selection

Retrieval Augmented Generation

(RAG) is an architecture for optimizing the performance of an artificial intelligence (AI) model by connecting it with external knowledge bases. RAG helps large language models (LLMs) deliver more relevant responses at a higher quality.

AI Agents

An artificial intelligence (AI) agent refers to a system or program that is capable of autonomously performing tasks on behalf of a user or another system by designing its workflow and utilizing available tools.

Retrieval Augmented Generation *Example*

Scite AI

Scite is a AI RAG tool that helps researchers better discover and understand research articles through Smart Citations–citations that display the context of the citation and describe whether the article provides supporting or contrasting evidence.

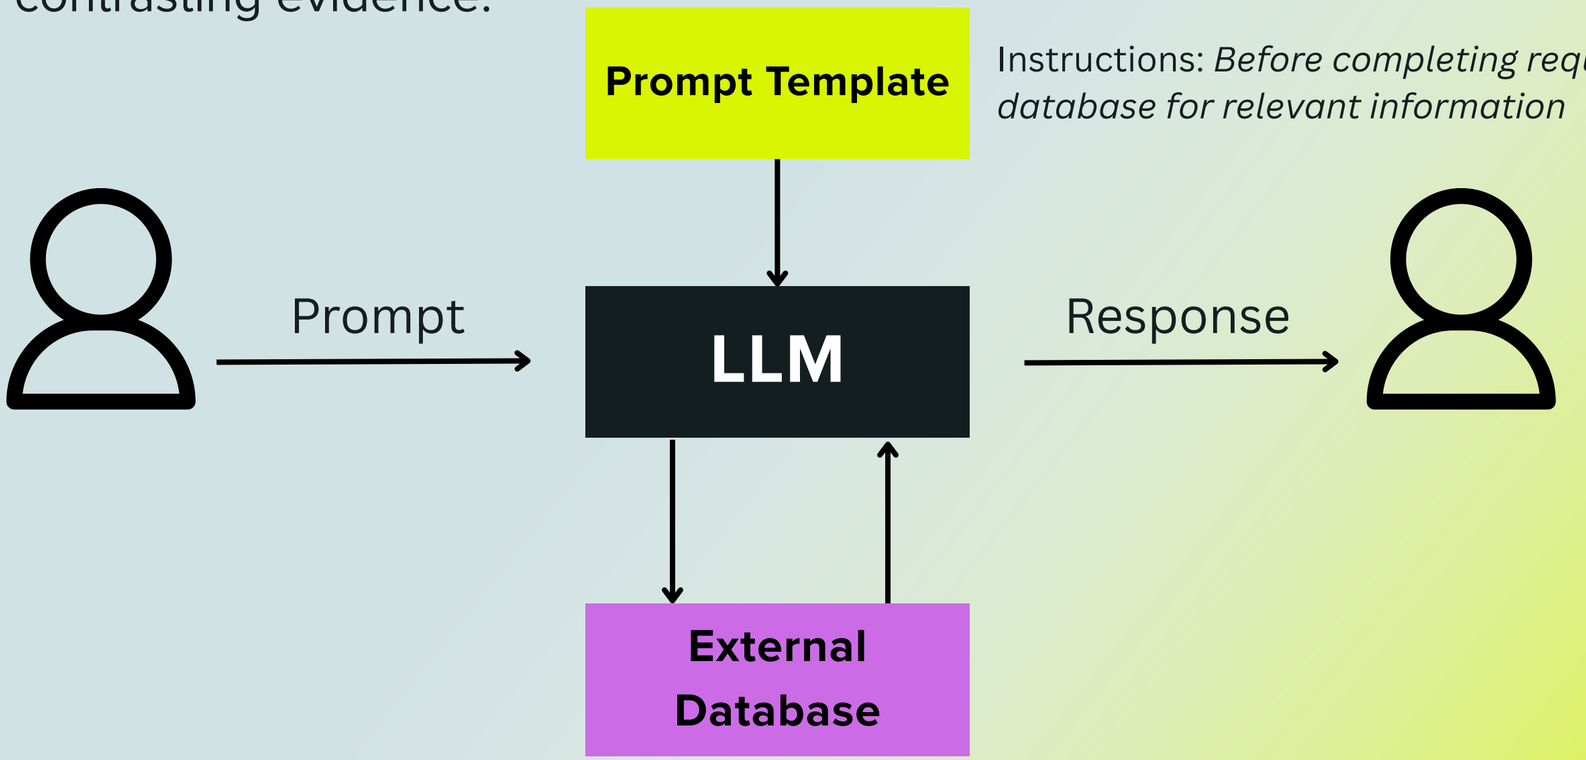

AI Agent *Example*

Undermind - An AI agent for scientific research

Created by two quantum physics PhDs from MIT with decades of experience in deep research. A research analyst tool that performs multiple iterative searches across numerous databases, and combines the results into a report (metasynthesis).

An AI that makes **choices**.

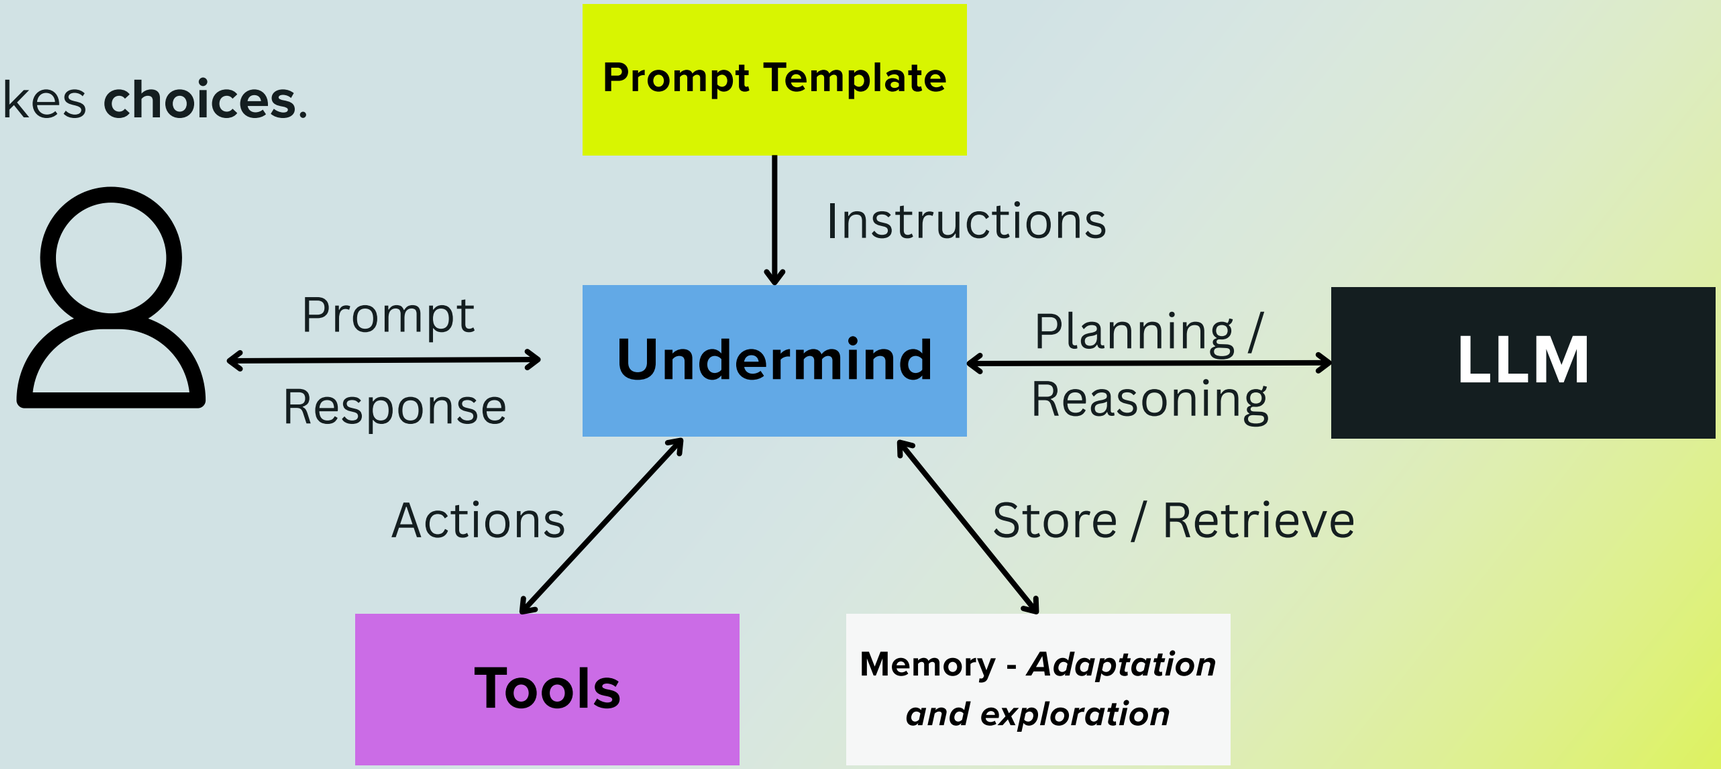

Evaluative *Example* - AI Hallucinations

Cursor Support Bot

Last month, an A.I. bot that handles tech support for Cursor, an up-and-coming tool for computer programmers, alerted several customers about a change in company policy. It said they were no longer allowed to use Cursor on more than one computer. Customers complained in angry posts to internet message boards. Some became even angrier when they realized what had happened: The A.I. bot had announced a policy change that did not exist.
“We have no such policy. You’re of course free to use Cursor on multiple machines,” the company’s chief executive and co-founder, Michael Truell, wrote in a Reddit post. “Unfortunately, this is an incorrect response from a front-line A.I. support bot.”

Why Hallucinations?

Researchers aren't totally sure. It isn't possible to trace back through the algorithm’s decision making process to determine when or why it made a false claim. One plausible theory is that models tend to show bias towards actions that developers reinforce. Anthropic recently got a model to learn a reward-model-pleasing strategy—it took biased actions that it thought would be reinforced even though they were not in fact being reinforced.

Ethics *Example* - Biased Algorithims

Millions Affected by Racial Bias in Healthcare Algorithm

An algorithm used in U.S. hospitals was supposed to direct patients to healthcare programs, but a 2019 study found it was less likely to refer black patients to the care they need than equally sick white patients. Because the algorithm assigned people to high-risk categories on the basis of costs (care provided to black people costs an average of $1,800 less per year than the care given to white people with the same number of chronic health problems), those biases were passed on in its results: black people had to be sicker than white people before being referred for additional help.

Only 17.7% of patients that the algorithm assigned to receive extra care were black. The researchers calculate that the proportion would have been 46.5% if the algorithm was unbiased.

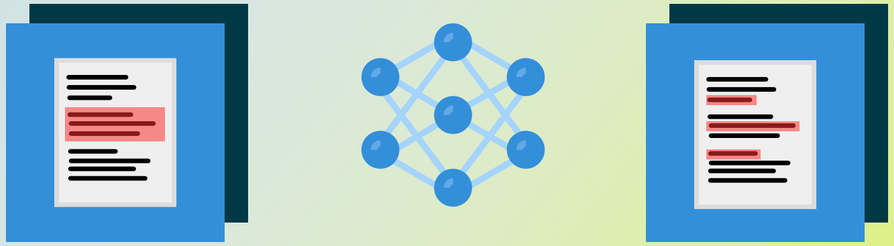

3 The Why: The Future of AI

Shifting the Narrative

Embracing Agency

Understanding the basics of AI and its potential implications

Coalition Building

Engaging with others ready to do the work

Education

Taking the steps necessary to become AI literate

Advocacy

becoming a professional squeaky wheel

Tech Philosophy

Developing a philosophy based on your personal values as an educator while understanding the need for student AI skills

Consistency

Small consistent changes over time

We must plant trees we will never sit under.

The first step in making a good choice,
is realizing that you have one.

Learning begins with literacy.

Q&A