AI For Improved Learning

Andres Espinosa
aespinosa31@gatech.edu

1 INTRODUCTION

This paper describes a course built for students to improve their confidence in using AI—specifically LLMs like ChatGPT. But this course does not—unlike many courses with AI in the title—provide students with a set of prompts to increase the output quality of their AI. Instead, it creates an outline for students to compile a set of AI-related metacognitive tools, and use it to improve the rate and depth at which they learn new tasks.

LLMs have been lauded for their ability to reduce effort and increase productivity. Certain tasks that would have previously taken hours can be done now in seconds with the help of an LLM. Can we say the same about learning? Does AI help us turn hours of learning into seconds? Or is it a shortcut that helps us now at the cost of harming us later?

1.1 Problem

AI is frequently cited as a major disruptor and replacer of existing human practices and social interaction patterns. The wealthiest and most powerful leaders consistently launch claims highlighting the transformative nature of the technology. Jensen Huang: "Every company will become an AI company." (Huang, 2023) Dario Amodei: "AI could compress a century of scientific progress into a decade." (Amodei, 2024) Sam Altman: "AI will probably most likely lead to the end of the world, but in the meantime, there'll be great companies." (Altman, 2023)

In the literature of AI education, AI usage can be split into two categories of which all interactions fall on a spectrum between the two: substitutive use and constructive use. Substitutive use, which the current state of media seems to broadcast the most, is associated with tasks that have been previously accomplished solely by humans, but are now capable of being solved with AI. And so, we have seen the brutal collapse of the standard American High-School English essay on the significance of the green light in *The Great Gatsby*. Conversely, constructive use denotes patterns of usage that enhance the capabilities of the LLM user. It can be incredibly difficult to recognize when usage is substitutive or constructive, and

it is often even more difficult to employ the self-control to steer usage from the immediate gratification of substitutive use to the more enriching constructive use.

1.2 Motivation

A common theme among works in the area of LLM-Human interaction is the design to help people use AI to reduce the amount of time it takes to accomplish a task and improve productivity. This perspective is especially prevalent in articles geared toward an academic or corporate audience—oftentimes with a title adhering to the pattern "How to use ChatGPT to accomplish [**menial task**] in [**less time**]". The motivation for this course is to provide the student with a transformative and unique experience, focusing not on immediate increased output, but on increased self-sufficiency and problem-solving efficiency. This course differentiates itself from the current landscape of blogposts and workshops associated with LLMs in a few ways.

Firstly, a historical analysis of AI has been performed in order to analyze what aspects of AI have remained constant through the release of stronger models, and which have improved or seen significant changes over time. From this, the limitations of AI have been categorized into those either temporary or fundamental. By detailing this analysis, the course helps students in identifying how to incorporate technological advancements into a sustainable improved learning process—preparing them for the everchanging landscape of AI. Secondly, this course pushes the user to consider two versions of AI simultaneously. The current version, the latest and best LLM we have access to (defined by user preference), and the instruction-sensitive quasi-omniscient oracle. This oracle, an analogue for Artificial General Intelligence, serves as a thought experiment to have users consider the following question: when we have the entirety of human knowledge and experience at our disposal—with the ability to have that information synthesized and presented to us at will—, in any form we desire, how can we learn best? Finally, and overall, this course aims to develop an intuition in students that enables them to identify constructive and substitutive AI use, and know when to use which.

The purpose of this course is to provide students with a foundational understanding of contemporary AI systems and to develop the metacognitive skills necessary to evaluate, guide, and integrate these systems into their own learning processes and improve the comprehensiveness at which they learn.

2 BACKGROUND

2.1 Literature Review

Recent innovations in AI, and the consequent rise in wary education headlines, demonstrate that there is much need for content to guide students on how to effectively utilise the tools. Research on AI-supported learning has shifted from a primary investigation into whether AI improves performance—approximated by increased output—to examining how it influences cognitive engagement. Early work on intelligent tutoring systems demonstrated that adaptive feedback can improve learning outcomes when students remain actively involved in the problem-solving process (Pang and Wei, 2025). With the emergence of LLMs, however, the nature of AI assistance has changed. Instead of guiding learners through structured steps, these systems can generate complete responses—raising concerns about whether this increased output reflects genuine learning or delegated thinking.

Several studies have found that students using LLMs produce higher-quality work with reduced effort, suggesting that AI can function as an exploratory cognitive offloading mechanism (Lehmann, Cornelius, and Sting, 2025). These studies discuss how students interact with AI and build an awareness of the problems it solves reliably, and they mold their workflows to accommodate this efficient personal assistant. While this may improve efficiency, it risks compromising the productive struggle associated with deeper learning. Kosmyna et al. show that students who relied on LLMs for essay writing exhibited reduced neural engagement, worse recall, and lower ability to reflect on their work compared to peers using search engines or unaided reasoning (Kosmyna et al., 2025). Similarly, Melumad et al. find that participants who learned through LLM-generated explanations retained less information and produced lower-quality advice than those who participated in a more active information filtering and synthesis process through web search (Melumad and Yun, 2025). Still, structured AI integration—such as guiding students to request explanations or feedback instead of final answers—has been shown to enhance motivation and critical thinking by reducing extraneous cognitive load (Ma and Zhong, Zhu, Liu, and Zhao, 2025, 2025). This aligns with broader findings in educational psychology that tools which scaffold (and not overlay) reasoning can improve metacognitive awareness.

At a structural level, the decentralization of responsibility for AI education poses a significant equity risk. When decisions about AI use are left to individual instruc-

tors or institutions, disparities in funding, training, and infrastructure translate directly into disparities in educational quality. Firth et al.'s broader analysis of internet-mediated cognition suggests that technological shifts reshape learning environments most profoundly during critical developmental periods (Firth et al., 2019). If generative AI is incorporated unevenly across educational systems, students in under-resourced contexts may experience its cognitive costs without receiving its pedagogical benefits. Such responses are understandable but unevenly distributed: institutions with greater financial and administrative capacity are better positioned to experiment with alternative assessment models, oral examinations, process-based evaluation, and AI transparency policies. Schools lacking these resources are more likely to rely on prohibitive or reactive measures that fail to address underlying pedagogical challenges. When such guidance is absent, students may confuse successful output with genuine understanding, a confusion that disproportionately harms learners who rely on formal education as their primary site of skill development.

Understanding how to consult AI so it strengthens transferable skills, rather than simply improving immediate performance, is essential for designing effective learning interventions. Despite this progress, a key gap remains: there is little guidance for students on the concrete steps that can be taken to improve individual learning.

3 METHODOLOGY

3.1 Course Structure and Objectives

The course follows a three module structure, with modules 1 and 3 covering AI topics and metacognitive practices, and module 2 serving as a set of practical exercises to build the main skills of the course. These modules build to these course objectives:

1. **Identify and Explain** substitutive use and constructive use of AI.
2. **Describe** how large language models generate responses and the dependence on input data and user prompting.
3. **Recognize** instances and patterns of AI undermining and aiding learning.
4. **Use** AI to clarify concepts, test understanding, generate alternative perspectives, find learning paths, and solicit feedback.
5. **Analyze** an AI-assisted work process to determine how and when learning is enhanced or replaced.

6. **Demonstrate** responsible AI use that aligns with academic integrity.
7. **Explore** through various subjects how AI can be used for learning in each area, and in an area of student choice.
8. **Apply** constructive AI-use strategies to new, unfamiliar learning environments.
9. **Demonstrate** autonomy in using AI for learning through completing a task with AI as a constructive tool without substituting their own reasoning.
10. **Develop** improvement in their ability to explain their reasoning, revise their thinking, and engage critically with material after AI-supported practice.

These ten course objectives have been identified as the set of metacognitive tools that will enable students to build a stronger and healthier relationship with AI, and new technological advancements throughout their life. The following sections describe the structure and design of the modules that contribute to these objectives.

3.2 Module 1

The first module covers the first three course objectives through a set of three lectures—outlined as academic blog posts. The following subsections outline the message of each lecture and the connection between the lecture and the corresponding learning objective(s) they build toward.

3.2.1 *Learning with AI*

AI is here to stay. As a society, we are figuring out how to integrate LLMs and AI bots into our businesses, our workflows, and our daily tendencies of living. The term "integration" fails to capture the nature of AI-human collaboration; "replacement" tends to be the more apt alternative. The denomination of substitutive vs. constructive use of AI aids LLM users to recognize the patterns associated with AI integration and those with replacement. See the appendix for a table containing examples of substitutive and constructive usage of AI.

This lecture covers the first course objective and provides some helpful context for the 8th and 9th objectives. This lecture also contributes to the second learning objective directly and sets a foundational understanding for the history and mechanisms behind AI.

3.2.2 *An Omniscient Oracle*

In order to identify how to learn best with an LLM, we must first analyze what occurs when AI performs at its absolute best. This leads to the omniscient oracle.

Barring the terrifying ethical implications of that, we can use it to imagine how we could use an LLM, functioning perfectly, to learn. Given a task the user would like to learn, what is the best question to ask to optimize the learning of the task?

This lecture provides context for much of module 2 and reflections that occur in module 3. This lecture does not directly contribute to the learning objectives, but instead challenges students to think deeper about how the technology they have access to today will change.

3.2.3 *Caveat Prompter*

A fundamental property, and limitation, of AI is the reliance on accurate and comprehensive data. LLMs do not possess an inherent mechanism for verifying truth; instead, they generate responses based on patterns in their training data, often using linguistic confidence and statistical likelihood as proxies for correctness. The most important single piece of data that an AI will use at inference time, is the prompt provided. It is therefore crucial to recognize the steps one can take to best extract value from an LLM.

This lecture covers the second learning objective directly and, by showing examples, provides context for module 2's covering of learning objectives 4-9. The lecture specifically discusses the limitations of AI and those that can be categorized as either fundamental or temporary.

3.3 Module 2

Module 2 comprises a set of practical exercises, each building toward a specific subject material, with the goal of improving students' efficacy and confidence in using AI tools for learning. The emphasis of each assignment is on the metacognitive practices associated with LLM usage. Students are guided in interfacing with an LLM to solve tasks and, after, reflect on their process and identify the behaviors that contributed to increased learning. LLMs are incredibly powerful tools. In consequence, the usage of those very tools in the pursuit of a refinement of learning begs restraint and foresight from the user—in other words, with great power comes great responsibility.

The following subsections communicate the contents and purpose of each assignment.

3.3.1 *English Literature*

This assignment requires students to read the book "The Great Gatsby". This book, a classic of American literature, has extensive essays, reflections, and analyses available on the internet—and correspondingly, in the training data of any LLM. AI, frequently used to replace the writing and reflection process, is repositioned here as a crucial piece in a deeper writing and revision process. Students are challenged: to analyze how LLM summaries of text leave out valuable meaning, to rewrite passages of the book from different perspectives, to criticize LLM output, and to reflect on the overall process of using AI for writing.

3.3.2 *Mathematics*

This submodule brings the student through a PSAT preparatory exam (math section) using AI. Students take a practice test twice—both without the aid of AI during— and use AI to design a study plan document to improve their score in a week. Students are challenged to design prompts that accurately target the individual need of the student. The student is then required to iterate on the study plan to submit a document describing the efficacy of AI in improving their score.

3.3.3 *History*

One of the History student learning objectives by the UC Boulder College of Arts and Science is: Conduct historical research, which includes: navigating libraries, databases, and archives; identifying, locating, and managing sources; and summarizing significant amounts of information. From this objective alone, it is clear to see how History skills are impacted by LLMs. This assignment walks students through using an LLM to research and compile a short but thorough historical analysis on the USA's 1999 intervention in the Kosovo-Serbian conflict.

3.3.4 *Foreign Language*

This submodule challenges students to use an LLM to learn the language of Italian. Foreign-language study is a field that benefits greatly from personalization. Students are challenged to use AI to identify their current level of Italian—students, depending on their disposition for foreign languages or knowledge of another language, such as Spanish, will have drastically different aptitudes for this task. This assignment will provide students with experience in strategically relaying information to an AI bot to augment its pedagogical capabilities.

3.3.5 *Student-Designed*

This assignment challenges students to build their own assignment in a material of their choosing. Students will engage with AI in a deeper level than the previous assignments by questioning how to create a set of guidelines for other students to work with AI constructively.

3.4 Module 3

The final module focuses on transfer and generalization of the metacognitive strategies developed throughout the course. While Module 2 situates AI use within structured subject areas, Module 3 removes these material-specific scaffolds and asks students to evaluate their understanding in general contexts.

Students are presented with open-ended reflection tasks that require them to learn a new concept or skill using AI, without explicit instruction on how to proceed. They are required to document their process in detail, including:

- The prompts they used
- The reasoning behind those prompts
- Points at which AI use was substitutive or constructive
- Revisions made after identifying ineffective strategies

The primary objective of this module is to assess whether students have internalized the ability to regulate their own AI use. Rather than demonstrating proficiency in a specific domain, students demonstrate their ability to learn effectively with AI across domains. This module contributes most directly to objectives 8, 9, and 10, emphasizing autonomy and sustained critical engagement.

4 ANALYSIS

The design of this course is grounded in the distinction between substitutive and constructive AI use, and aims to shift student behavior along this spectrum. A fundamental belief here is that students can be trained to recognize and regulate their own AI usage, resulting in improved learning outcomes despite the presence of tools capable of completing tasks on their behalf.

The structure of Module 2 plays a critical role in this process. By exposing students to repeated cycles of interaction, critique, and reflection, the course creates conditions for metacognitive development. Students are not only using AI, but actively

analyzing how their use of AI affects their understanding. This aligns with prior research suggesting that learning improves when students engage in reflection and self-explanation during problem-solving tasks (Ma and Zhong, Zhu, Liu, and Zhao, 2025, 2025).

The deliberate inclusion of substitutive tasks—where AI completes work on behalf of the student—is essential. Rather than prohibiting these behaviors, the course surfaces them and asks students to evaluate their consequences. Students should be able to better identify the learning value present in a task and correspondingly employ a more constructive or subsitutive use of AI consistent with their learning goals. This approach reduces the likelihood of uncritical reliance on AI by making the trade-offs visible.

The use of multiple subject areas further strengthens the design by preventing students from associating effective AI use with a single domain. Instead, students develop generalized strategies that can be applied across contexts. This is particularly important given the rapid evolution of AI tools, as domain-specific techniques may quickly become obsolete.

The inclusion of the omniscient oracle thought experiment serves as a conceptual anchor for—and an attempt, albeit inevitably imperfect, at future-proofing—the material. By considering an idealized version of AI, students are encouraged to think beyond current limitations and focus on the underlying principles of learning. This abstraction helps prepare students against further technological change.

5 DISCUSSION

While the course presents a structured approach to improving AI-assisted learning, several limitations must be acknowledged.

First, the effectiveness of the course depends heavily on student engagement and honesty. Because many activities involve self-reflection and process documentation, students who do not fully participate may fail to develop the intended metacognitive skills. This challenge is not unique to AI-based learning, but it is amplified by the ease with which AI can obscure a student's true level of understanding.

Second, the course assumes access to reliable AI tools. Students in under-resourced environments may face limitations in access to high-quality models, stable internet connections, or sufficient time to engage deeply with the material. This raises

broader concerns about equity, as the benefits of effective AI use may disproportionately favor students with greater access to technological resources.

Third, the distinction between substitutive and constructive use, while conceptually useful, may be difficult to manifest in practice. Many interactions with AI fall somewhere between the two extremes, and students may struggle to accurately categorize their own behavior. Although this course is built with it in mind, the facility of AI substitutive use—a siren beckoning the user to relinquish their autonomy in the learning process—presents a consistent conflict with the self-control of the user. Future iterations of the course could benefit from more precise rubrics or examples to support this distinction.

Despite these limitations, the course addresses a critical gap in current educational practice. Most existing approaches to AI in education focus either on prohibition or on productivity optimization. This course instead positions AI as a tool for developing deeper learning skills, emphasizing process over output.

6 CONCLUSION

Generative AI represents a fundamental shift in the relationship between humans and information. Unlike previous technologies, LLMs are capable of performing complex cognitive tasks, raising important questions about the nature of learning in an AI-mediated environment.

This paper has presented a course designed to address these challenges by developing students' metacognitive awareness and ability to use AI constructively. Through a combination of theoretical presentation, practical exercises, and reflective activities, the course encourages students to move beyond passive reliance on AI and toward intentional, strategic use.

The central contribution of this work is the internalization and application of the substitutive versus constructive framework. By making this distinction explicit, and giving students practice in the framework, the course provides students with a lens through which to evaluate their own behavior and adapt to new technologies.

As AI continues to evolve, the ability to learn effectively with these systems will become increasingly important. Educational approaches that focus solely on content acquisition or productivity gains are unlikely to prepare students for this reality. Instead, there is a need for frameworks that emphasize adaptability, critical thinking, and self-regulation.

This course represents one step toward that goal, offering a model for how AI can be integrated into education in a way that supports, rather than replaces, human learning.

REFERENCES

1. Firth, Joseph, Torous, John, Stubbs, Brendon, Firth, Josh A, Steiner, Genevieve Z, Smith, Lee, Alvarez-Jimenez, Mario, Gleeson, John, Vancampfort, Davy, Armitage, Christopher J, et al. (2019). "The "online brain": how the Internet may be changing our cognition". In: *World psychiatry* 18.2, pp. 119–129.
2. Altman, Sam (2023). *Interviews on Artificial Intelligence Risk and Economic Transformation*. Multiple interviews. Statement widely reported: AI may pose existential risks while also enabling major economic growth.
3. Huang, Jensen (2023). *NVIDIA GTC 2023 Keynote Address*. Conference keynote. Statement widely reported: "Every company will become an AI company".
4. Amodei, Dario (2024). *Public Statements on AI and Scientific Acceleration*. Interviews and essays. Paraphrased idea: AI could dramatically accelerate scientific progress, potentially compressing decades of discovery into much shorter timelines.
5. Kosmyna, Nataliya, Hauptmann, Eugene, Yuan, Ye Tong, Situ, Jessica, Liao, Xian-Hao, Beresnitzky, Ashly Vivian, Braunstein, Iris, and Maes, Pattie (2025). *Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task*. arXiv: 2506.08872 [cs.AI]. URL: https://arxiv.org/abs/2506.08872.
6. Lehmann, Matthias, Cornelius, Philipp B., and Sting, Fabian J. (2025). *AI Meets the Classroom: When Do Large Language Models Harm Learning?* arXiv: 2409.09047 [cs.CY]. URL: https://arxiv.org/abs/2409.09047.
7. Ma, Nan and Zhong, Zhiyong (2025). "A Meta-Analysis of the Impact of Generative Artificial Intelligence on Learning Outcomes". In: *Journal of Computer Assisted Learning* 41.5. e70117 JCAL-24-1657.R3, e70117. DOI: https://doi.org/10.1111/jcal.70117. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jcal.70117. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/jcal.70117.
8. Melumad, Shiri and Yun, Jin Ho (2025). "Experimental Evidence of the Effects of Large Language Models versus Web Search on Depth of Learning". In: *PNAS Nexus* 4.10. DOI: 10.1093/pnasnexus/pgaf316.

9. Pang, Weina and Wei, Zhe (2025). "Shaping the Future of Higher Education: A Technology Usage Study on Generative AI Innovations". In: *Information* 16.2. ISSN: 2078-2489. DOI: 10.3390/info16020095. URL: https://www.mdpi.com/2078-2489/16/2/95.
10. Zhu, Y., Liu, Q., and Zhao, L. (July 2025). "Exploring the impact of generative artificial intelligence on students' learning outcomes: a meta-analysis". In: *Education and Information Technologies* 30. Received 23 September 2024; Accepted 28 January 2025; Published 26 February 2025, pp. 16211–16239. DOI: 10.1007/s10639-025-13420-z. URL: https://doi.org/10.1007/s10639-025-13420-z.

7 APPENDIX: CONSTRUCTIVE VS. SUBSTITUTIVE USE

Task	Constructive AI Use	Substitutive AI Use
Writing	Student drafts their own thesis and asks AI for feedback, counterarguments, or clarity improvements.	Student prompts AI to "write a full essay on X topic" and submits it with minimal revision.
Reading	Student reads a text and uses AI to explore themes, alternative interpretations, or clarify confusing passages.	Student relies on AI summaries instead of reading the text.
Math	Student attempts problem first, then asks AI to explain mistakes or provide step-by-step reasoning.	Student inputs the problem and copies the solution without engaging with the steps.
Coding	Student writes in pseudo-code and prompts AI to reduce syntax memorization load.	Student asks AI to generate an entire program and runs it without understanding the logic.
Foreign Lang.	Student practices conversation with AI, requests corrections, and asks for explanations of grammar mistakes.	Student uses AI to translate entire assignments without attempting to form sentences independently.
History	Student gathers sources and uses AI to compare perspectives or identify gaps in reasoning.	Student asks AI to "write a report on X" and accepts it as complete research.
Creative Work	Student generates ideas with AI and iterates on them, refining their own creative direction.	Student asks AI to produce a finished story/artwork and presents it as their own.

Table 1—Examples of Constructive vs. Substitutive AI Use Across Tasks