

# The Promise and Peril of SE Education in the Age of AI

Michael Hilton  
Carnegie Mellon University



Software and Societal  
Systems Department

# Michael Hilton



A.S. Grossmont Community College



B.S. San Diego State University



Software Engineer at DoD



M.S. Cal Poly San Luis Obispo



PhD at Oregon State



Internship at Microsoft Research




Assistant/Associate/Full Teaching Professor at CMU



OPINION  
GUEST ESSAY

# You're a Computer Science Major. Don't Panic.

Nov. 12, 2025



Erik Carter

▶ Listen to this article · 4:26 min [Learn more](#)  Share full article    356

By **Mary Shaw and Michael Hilton**

Dr. Shaw and Dr. Hilton teach software engineering at Carnegie Mellon University.



**Ask:**

**Discuss  
“Integrating AI  
into Software  
Engineering  
Education” for the  
CSEE&T audience**

# Thanks to everyone who discussed with me

- I appreciate the time many people took to talk with me
- This is a rapidly changing world
- Disclaimer: My expertise is in SE/Education, not in AI



## Two Different Questions

**How should we teach:** How can we effectively deliver Software Engineering education in a world where GenAI exists?

**What should we teach:** How should we think about educating student for a career in Software Engineering in a world where GenAI exists?

**How can we effectively  
[teach | | learn] Software  
Engineering in a world  
where GenAI exists?**

**“ Learning results from what the student does and thinks and only from what the student does and thinks. The teacher can advance learning only by influencing what the student does to learn”**

**-Herb Simon**



Turing Award 1975  
Nobel in Econ 1978

# Principle:

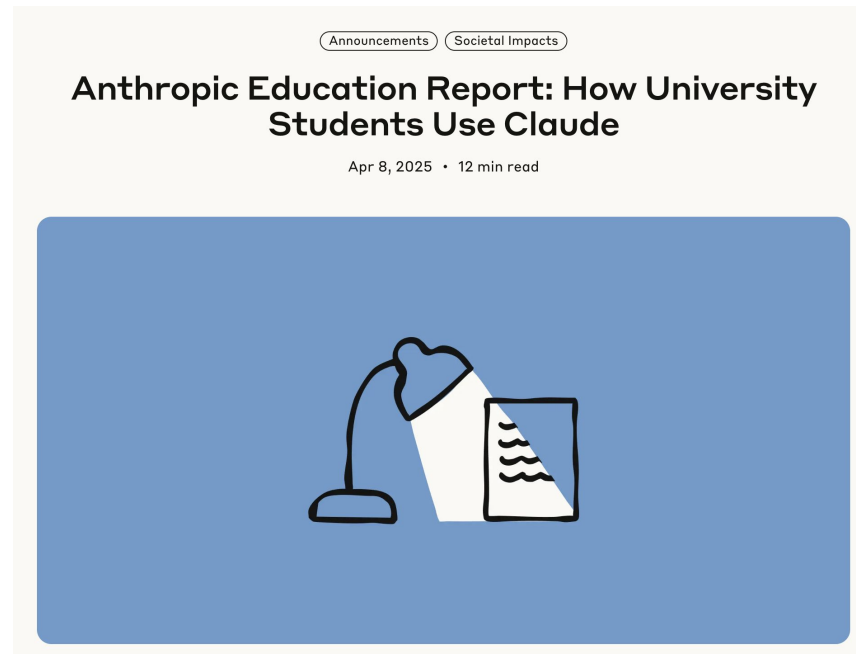
AI will result in learning when it influences students to “do and think”

# How are students using AI?

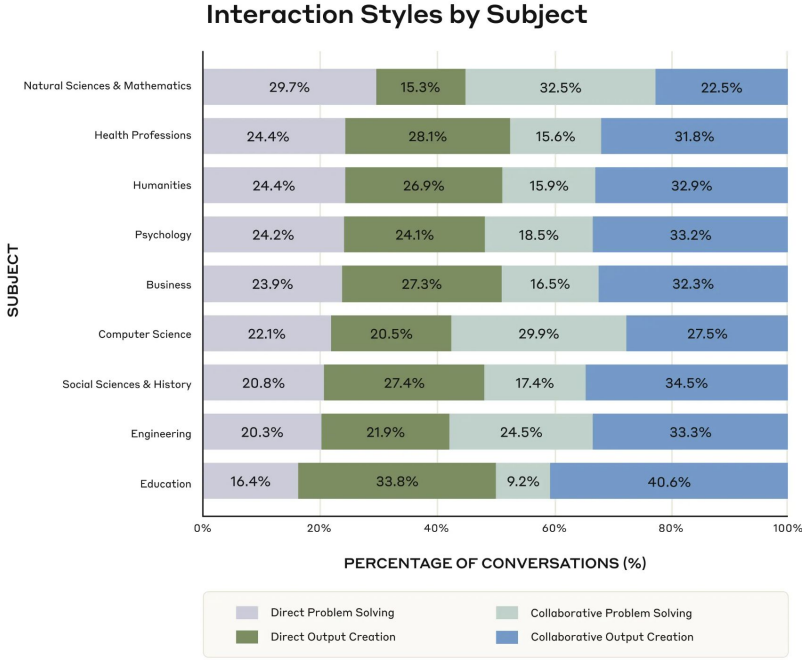
# Report on student AI use

“one of the first large-scale studies of real-world AI usage patterns in higher education, analyzing one million anonymized student conversations on Claude.ai”

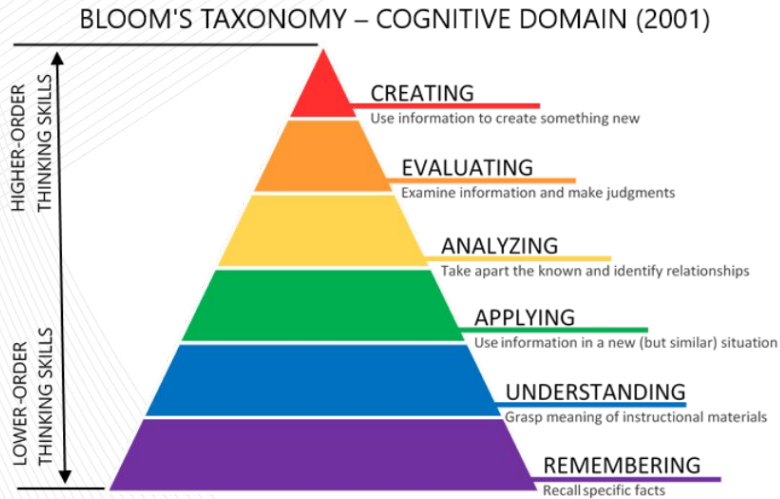
“Computer Science students particularly overrepresented (accounting for 36.8% of students’ conversations while comprising only 5.4% of U.S. degrees)”




# ~50% of requests are Direct Problem Solving || Output Creation



# Cognitive tasks students delegate to AI




<https://citt.ufl.edu/resources/the-learning-process/designing-the-learning-experience/blooms-taxonomy/>



# Has AI helped your learning?

Do you think AI has been a net positive in your learning experience?


100 responses



# Has AI helped your PEER'S learning?


Do you think AI has been a net positive in your learning experience?

100 responses

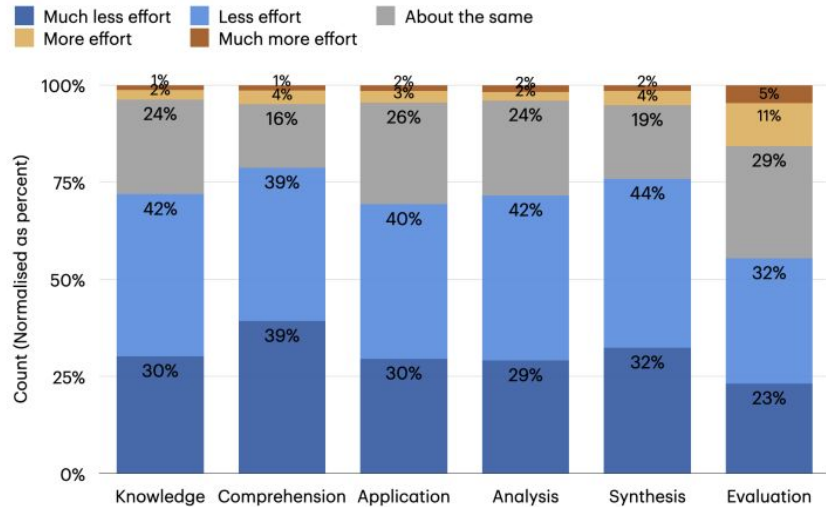


Do you think AI has been a net positive in your peers' learning experience?

100 responses



# Similar effects observed in practice



Distribution of perceived effort (%) in cognitive activities (based on Bloom's taxonomy) when using a GenAI tool compared to not using one.

## The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers

Hao-Ping (Hank) Lee  
Carnegie Mellon University  
Pittsburgh, Pennsylvania, USA  
haoping@cs.cmu.edu

Ian Drosos  
Microsoft Research  
Cambridge, United Kingdom  
i-androsos@microsoft.com

Advait Sarkar  
Microsoft Research  
Cambridge, United Kingdom  
advait@microsoft.com

Sean Rintel  
Microsoft Research  
Cambridge, United Kingdom  
serintel@microsoft.com

Lev Tankelevitch  
Microsoft Research  
Cambridge, United Kingdom  
levt@microsoft.com

Richard Banks  
Microsoft Research  
Cambridge, United Kingdom  
rbanks@microsoft.com

Nicholas Wilson  
Microsoft Research  
Cambridge, United Kingdom  
nwilson@microsoft.com

### Abstract

The rise of Generative AI (GenAI) in knowledge workflows raises questions about its impact on critical thinking skills and practices. We survey 319 knowledge workers to investigate 1) when and how they perceive the enactment of critical thinking when using GenAI, and 2) when and why GenAI affects their effort to do so. Participants shared 936 first-hand examples of using GenAI in work tasks. Quantitatively, when considering both task- and user-specific factors, a user's task-specific self-confidence and confidence in GenAI are predictive of whether critical thinking is enacted and the effort of doing so in GenAI-assisted tasks. Specifically, higher confidence in GenAI is associated with less critical thinking, while higher self-confidence is associated with more critical thinking. Qualitatively, GenAI shifts the nature of critical thinking toward information verification, response integration, and task stewardship. Our insights reveal new design challenges and opportunities for developing GenAI tools for knowledge work.

Confidence Effects From a Survey of Knowledge Workers. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 23 pages. <https://doi.org/10.1145/3706598.3713778>

### 1 Introduction

Generative AI (GenAI) tools, defined as any “end user tool [...] whose technical implementation includes a generative model based on deep learning”, are the latest in a long line of technologies that raise questions about their impact on the quality of human thought: a line that includes writing (objected to by Socrates), printing (objected to by Trithemius), calculators (objected to by teachers of arithmetic), and the Internet.

Such consternation is not unfounded. Used improperly, technologies can and do result in the deterioration of cognitive faculties that ought to be preserved. As Bainbridge [7] noted, a key irony of automation is that by mechanising routine tasks and leaving




# Challenge: building component skills

# Developing Mastery

Principle: To develop mastery, students must acquire component skills, practice integrating them, and know when to apply what they have learned.

Ambrose, Susan A., et al. *How Learning Works : Seven Research-Based Principles for Smart Teaching*, John Wiley & Sons, Incorporated, 2010.



# Component skills matter

“For students to achieve mastery within a domain, whether narrowly or broadly conceived, they need to develop a set of key component skills, practice them to the point where they can be combined fluently and used with a fair degree of automaticity, and know when and where to apply them appropriately.”

Mastery



# Example Component Skills

- Research with introductory statistics students identified two key skills involved in statistical data analysis:
  - The ability to recognize the relevant variables
  - The ability to categorize them according to types.
- 45 minutes of direct instruction and feedback on identifying statistical problem types boosted beginners' ability (Lovett 2002)

Consider a student writing a linked list.  
Component skills include:


- Syntax
- Abstraction of pointers
- Abstract reasoning about data structure
- ...

We don't need more linked lists in the world,  
but do students need to develop those skills?


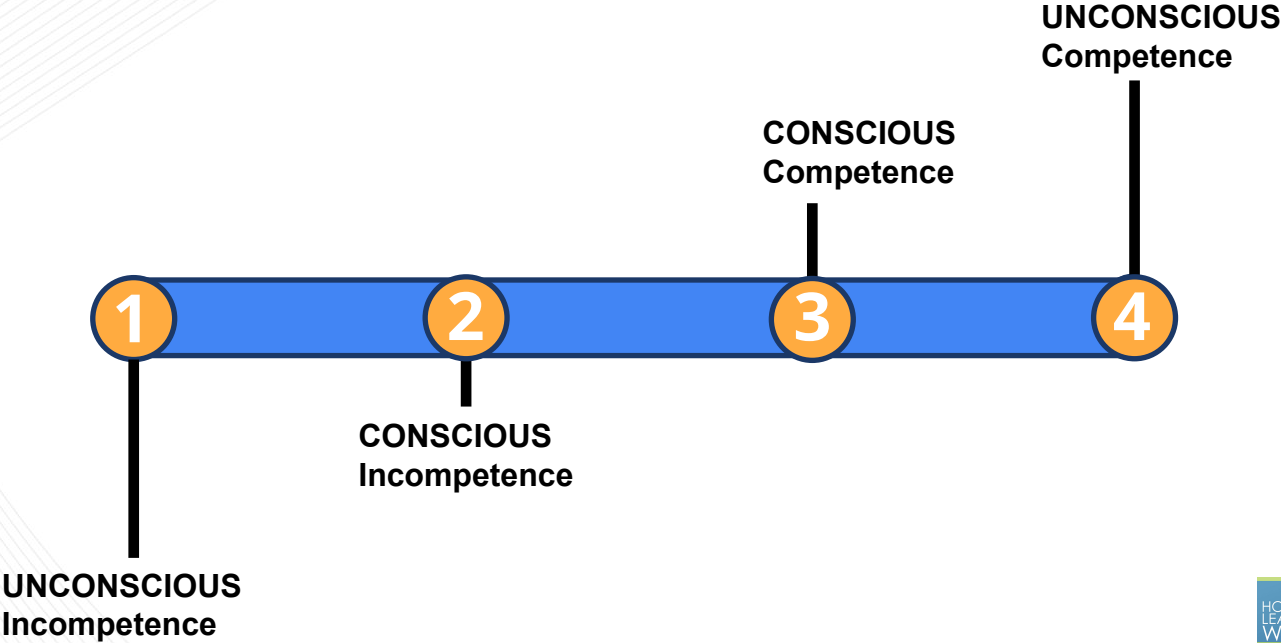


# Analogy: Calculators?

- We teach math to people without using calculators, is that similar to using LLMs to program?
- LLMs can program faster than a human
- But calculators are always right. We don't need to teach people to evaluate the output of a calculator
- Similar analogy, compilers? Power tools?



# What is hard about this?



**Challenge: How do we motivate  
students to persevere in component  
skill learning?**

# Student motivation

Before LLMs:

- Often motivated with Extrinsic Motivation (Grades)
- More difficult to find answers (or easier to detect)
- Often easiest path forward is to do the work


After LLMs:

- Correct answers available with little to no effort
- LLM answers are impossible to detect with high confidence
- Easiest path forward is not doing the work

# Students know this

What concerns do you have about using AI in software development?

100 responses



**“People don’t care how much you know,  
till they know how much you care”  
-My Father**

**How do we get students to care?**

# Open question, IMHO key to good teaching

Suggestions:

- Explain why component skills matter
- Tie high level skills to long term goals
- Provide authentic tasks when possible
- Identify and reward what you value (be explicit about what you want students to do!)

**What if we are causing students to think/do (aka learn) the wrong thing?**

# GenAI in-class experience

- Apr 6, 2023
- Wanted to experiment with GenAI Tools
- Asked students to create a Tic Tac Toe React app
  - Class didn't teach react
- Ask them to not manually edit code (Vibe Code)

README

## Tic Tac Toe React App

This repository only contains this Readme. Your task is to build a tic tac toe react app using generative AI tools like ChatGPT or GitHub Copilot to assist you in **every** aspect of the process.

Before you begin working on the TO-DO, please take note of the following questions which we will be discussing after the activity:

1. What are the advantages and disadvantages of using generative AI tools like ChatGPT and GitHub Copilot?
2. What kind of task did you find worked well with ChatGPT and GitHub Copilot?
3. What kind of task did you find did not work well with ChatGPT and GitHub Copilot?
4. What is something you did not expect when using ChatGPT and GitHub Copilot?

# Experience

- Students' proficiency didn't seem to correlate with their experience.
- Did seem to correlate with demographic groups that we would expect to be high/lower self-efficacy
- A student got bad code from AI, tried to follow the (incorrect) directions 6 times in a row, comment to me was "I can't figure it out, I'm not sure what I'm doing wrong"
- Another student said "programming with AI is like yelling at stupid person"

# Follow up Experiences

- In subsequent classes, instead of just asking students to try it out, I instead first demo the exercise, intentionally demonstrating the GenAI tool failing for me (while it is claiming to be correct)
- Anecdotally, once I demonstrate this, I get very different answers from students
  - "It is also failing for me " "The AI isn't able to create the code I need"

# Anti-pattern

“slot machine” programming  
No mental model, just repeated prompting



# Last semester: Vibe Coding Exercise

- Ask AI to answer explain feature in NodeBB (open source project), without code
- Ask AI where to fix (without code)
- Ask AI for patch (without providing code)
- Ask AI for patch (with providing code)
- Switch to Cursor/Windsurf (Context aware AI)
- Repeat steps above
- Asked students for their perceptions

# This semester: Vibe Coding Exercise



- Feedback:

- "I'm probably going to use ai tools equally but be more cautious of their response."
- "told me to change a file that i can't find: src/controllers/post.js"
- "I feel like I am probably less likely to use AI unless I have very specific questions. It's not very good at understanding and working on a complex codebase, and it's just better for me to understand it myself."
- "The most time consuming thing about AI is double checking the work of AI since most of the time it may not be fully correct."
- "might use tools like cursor more (mainly just did chatgpt + copy/paste in past) but i also don't want to drain my wallet with a billion subscriptions. also im not sure about the quality of the changes"
- "In the future, I will continue to use AI tools, but I will be more careful on how much I use them, making sure to use them less"
- "I think I will equally likely to use AI tools as before, I will continue to use it when getting stuck and just take it as suggestion instead of fully believe in it"

# Last semester: Vibe Coding Exercise

- Feedback:
  - “I’m probably going to use ai tools equally but **be more cautious** of their response.”
  - “The most time consuming thing about AI is **double checking** the work of AI since most of the time it may not be fully correct.”
  - “might use tools like cursor more (mainly just did chatgpt + copy/paste in past) but i also don’t want to **drain my wallet** with a billion subscriptions. also im not sure about the quality of the changes”
  - “I think I will equally likely to use AI tools as before, I will continue to use it when getting stuck and just take it **as suggestion instead of fully believe in it**”

**What does this matter?**


# Why Attribution matters?

**Students who feel like they don't belong will attribute mistakes to their not belonging.**

**Students who do feel like they belong will attribute mistakes to the inherent challenge of the task.**

# Changing Attribution

- Changing Attribution can be a powerful tool.
- “Walton and Cohen (2011) ... intervention particularly helped the African American students to reframe their interpretation of the daily adversities of being in college rather than seeing adversity as symbolizing that they do not belong, they reframed the setbacks as simply setbacks.”
- Intervention didn’t help students who already felt they belonged



# Positive use of GenAI in education

# Requirements Engineering with GenAI



## Before GenAI:

- Students would interview humans as part of RE
- Students would find subjects (often friend, relative, etc)
- High cost in human time of subjects
- Course staff didn't have visibility into interviews
- Subjects would often be overly agreeable, don't really know what they need, talk in circles, etc

## After GenAI:

- Students practice interviews with GenAI agent before interviewing humans
- Course staff has transcript of interview, can review, give feedback.
- Low cost, students can repeat interviews till they achieve mastery
- GenAI characteristics match human subject tendencies

# Principle:



AI will result in learning when it influences students to “do and think”

If students “do and think” more (e.g., do more practice interviews, get good feedback) they will learn more

Practice interviews should not replace human interviews, but can serve as simulator to help students practice at scale before interviewing humans.

**How should we think  
about educating student  
for a career in Software  
Engineering in a world  
where GenAI exists?**

**Will SE exist in 10 years?**

# Jevons paradox

- The Jevons Paradox describes a situation where technological advancements that improve the efficiency of resource use actually lead to increased overall consumption. This occurs because efficiency gains make the resource cheaper, which in turn stimulates increased demand and usage.



NEWSLETTER

LISTEN & FOLLOW   


**Why the AI world is suddenly obsessed with a 160-year-old economics paradox**

FEBRUARY 4, 2025 · 6:30 AM ET

 Greg Rosalsky




**Maybe in 10  
years, everyone  
will be a  
software  
developer?**



## Knoll's law of media accuracy

“everything  
is absolutely  
story of whi  
firsthand kn



spapers  
e rare  
ave

# Should we teach students how to use GenAI?

- Horses used to be essential transportation
- Some people loved horses
- Horses are no longer efficient for transportation, have been relegated to niche (and expensive) hobby
- Zen coding used to be the way, but it is realistic it could be relegated to a hobby




Metcalf corner of Sparks, 1895, Ottawa CA

# Most developers spend <1hr coding a day

Code time per day  
Percentage of developers by code time segment

14%

**Reminder: Software Engineering is a lot MORE than Programming**



## My experience:

One page app

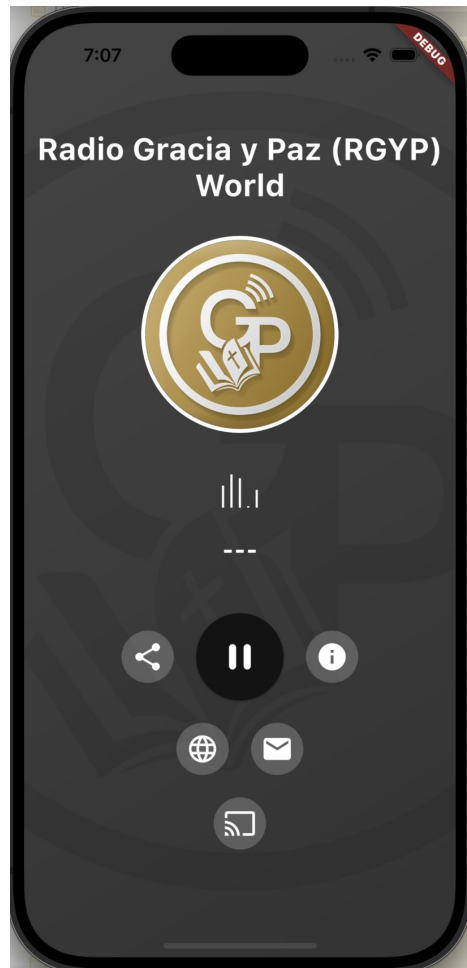
Play audio stream

Decided to only program with LLM

“Naive” Prompting strategy would not work

Trying to guide agent as if it was a student,  
worked quite well

My Takeaway: Tools are very valuable, but  
you need to have SE knowledge to use  
successfully




**So where do we  
go from here?**

# What would it look like to rethink SE education?

- Should we teach reading before writing?
- Should we teach "judging" of AI as a separate skill?
- What are the skills students will need to do SE in a world where code writing is "free"?
- Should we focus on defining and identifying quality attributes such as readability, maintainability, etc?
- How can we teach students validation and verification for code that they didn't write?

# Jeff Goldblum Principle



**Software Engineering should  
be about people building  
software for people**

**“ Learning results from what the student does and thinks and only from what the student does and thinks. The teacher can advance learning only by influencing what the student does to learn”**

**-Herb Simon**




Turing Award 1975  
Nobel in Econ 1978



## Developing Mastery


Principle: To develop mastery, students must acquire component skills, practice integrating them, and know when to apply what they have learned.

Ambrose, Susan A., et al. *How Learning Works : Seven Research-Based Principles for Smart Teaching*, John Wiley & Sons, Incorporated, 2010.



## Changing Attribution

- Changing Attribution can be a powerful tool.
- “Walton and Cohen (2011) ... intervention particularly helped the African American students to reframe their interpretation of the daily adversities of being in college rather than seeing adversity as symbolizing that they do not belong, they reframed the setbacks as simply setbacks.”
- Intervention didn't help students who already felt they belonged



**Software Engineering should be about people building software for people**

