

Do AI course tutors scaffold learning or smooth away the struggle?

Dr Jon Cardoso-Silva

 J.Cardoso-Silva@lse.ac.uk



Data Science
Institute

10 Jun 2026

Are pedagogical AI tutors delivering?



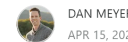
Brave new words: how AI will revolutionise education, with Salman Khan · DSI Event · Wed 26 Jun 2024

[Read my book review.](#)

- **Khanmigo** was the best bet so far: the most funded tutor, Socratic by design, grounded in Khan Academy, backed by Gates and Microsoft.
- At the DSI event in June 2024, Sal Khan called AI tutoring the biggest positive transformation education has ever seen. By April 2026 he told Chalkbeat that for many students it had been **a non-event** (so far?).

RIP Khanmigo & Edtech Industry Dreams of AI Tutors

Never forget.



DAN MEYER
APR 15, 2026

♡ 271

💬 57

🔄 30

Share



Khanmigo, Khan Academy's AI chatbot tutor, died last week. Because Khanmigo is software, it does not die in the traditional, organic sense. You'll still find it in the sidebar of Khan Academy practice exercises. But Khanmigo as an idea, as a shorthand for the edtech industry dreams of software that tutors as well as humans, has died, crushed underneath the expectations of its own creator, Sal Khan.

"For a lot of students, it was a non-event," Khan told Matt Barnum [last week in Chalkbeat](#), referring to Khanmigo's release three years ago. "They just didn't use it much," he continued. Khan Academy's Chief Academic Officer, Kristen DiCerbo, gave a similar assessment: "So far I am not seeing the revolution in education." These

LSE AI and Education Fellowships

- I am one of ten LSE AI and Education Fellows (2025–2027).
- The fellowship buys out a term of teaching so we can redesign our courses with GenAI: teaching, learning, and assessment.

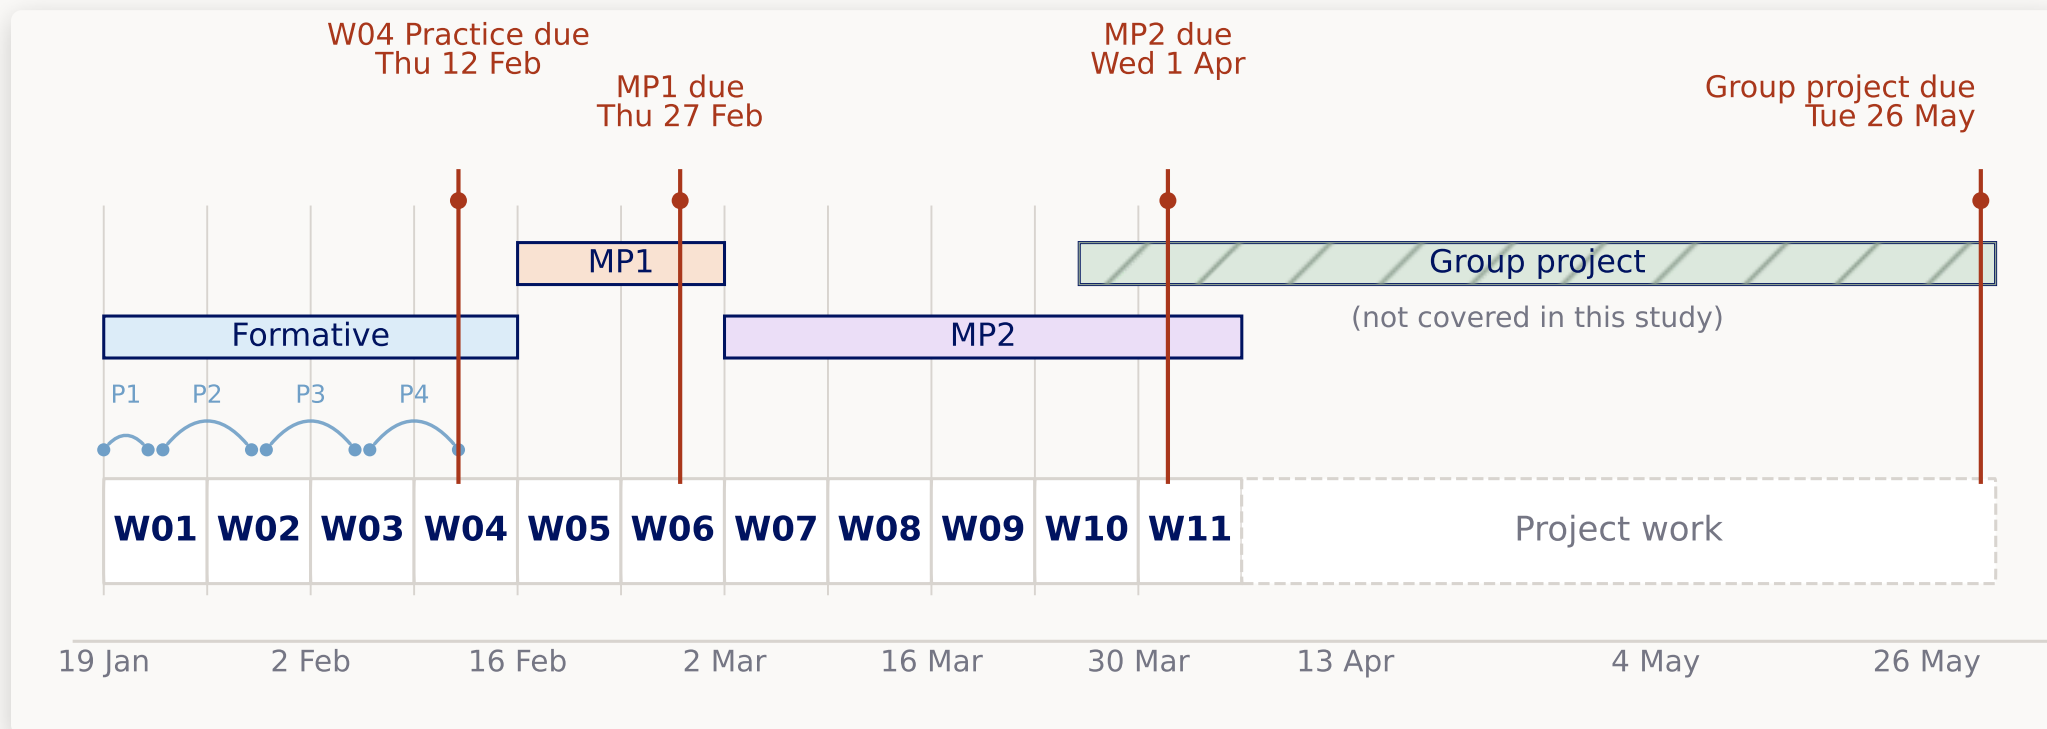
LSE AI and Education Fellowships

“ Integrating GenAI pedagogically takes far more time and experimentation than people realise. The AI Fellowship is giving me the space to do that properly. The more colleagues who get this opportunity, the better for our students. ”

AI and Education Fellow Lourdes Hernandez Martin, Language Coordinator - Arabic and Coordinator for Spanish Projects (Language Centre)

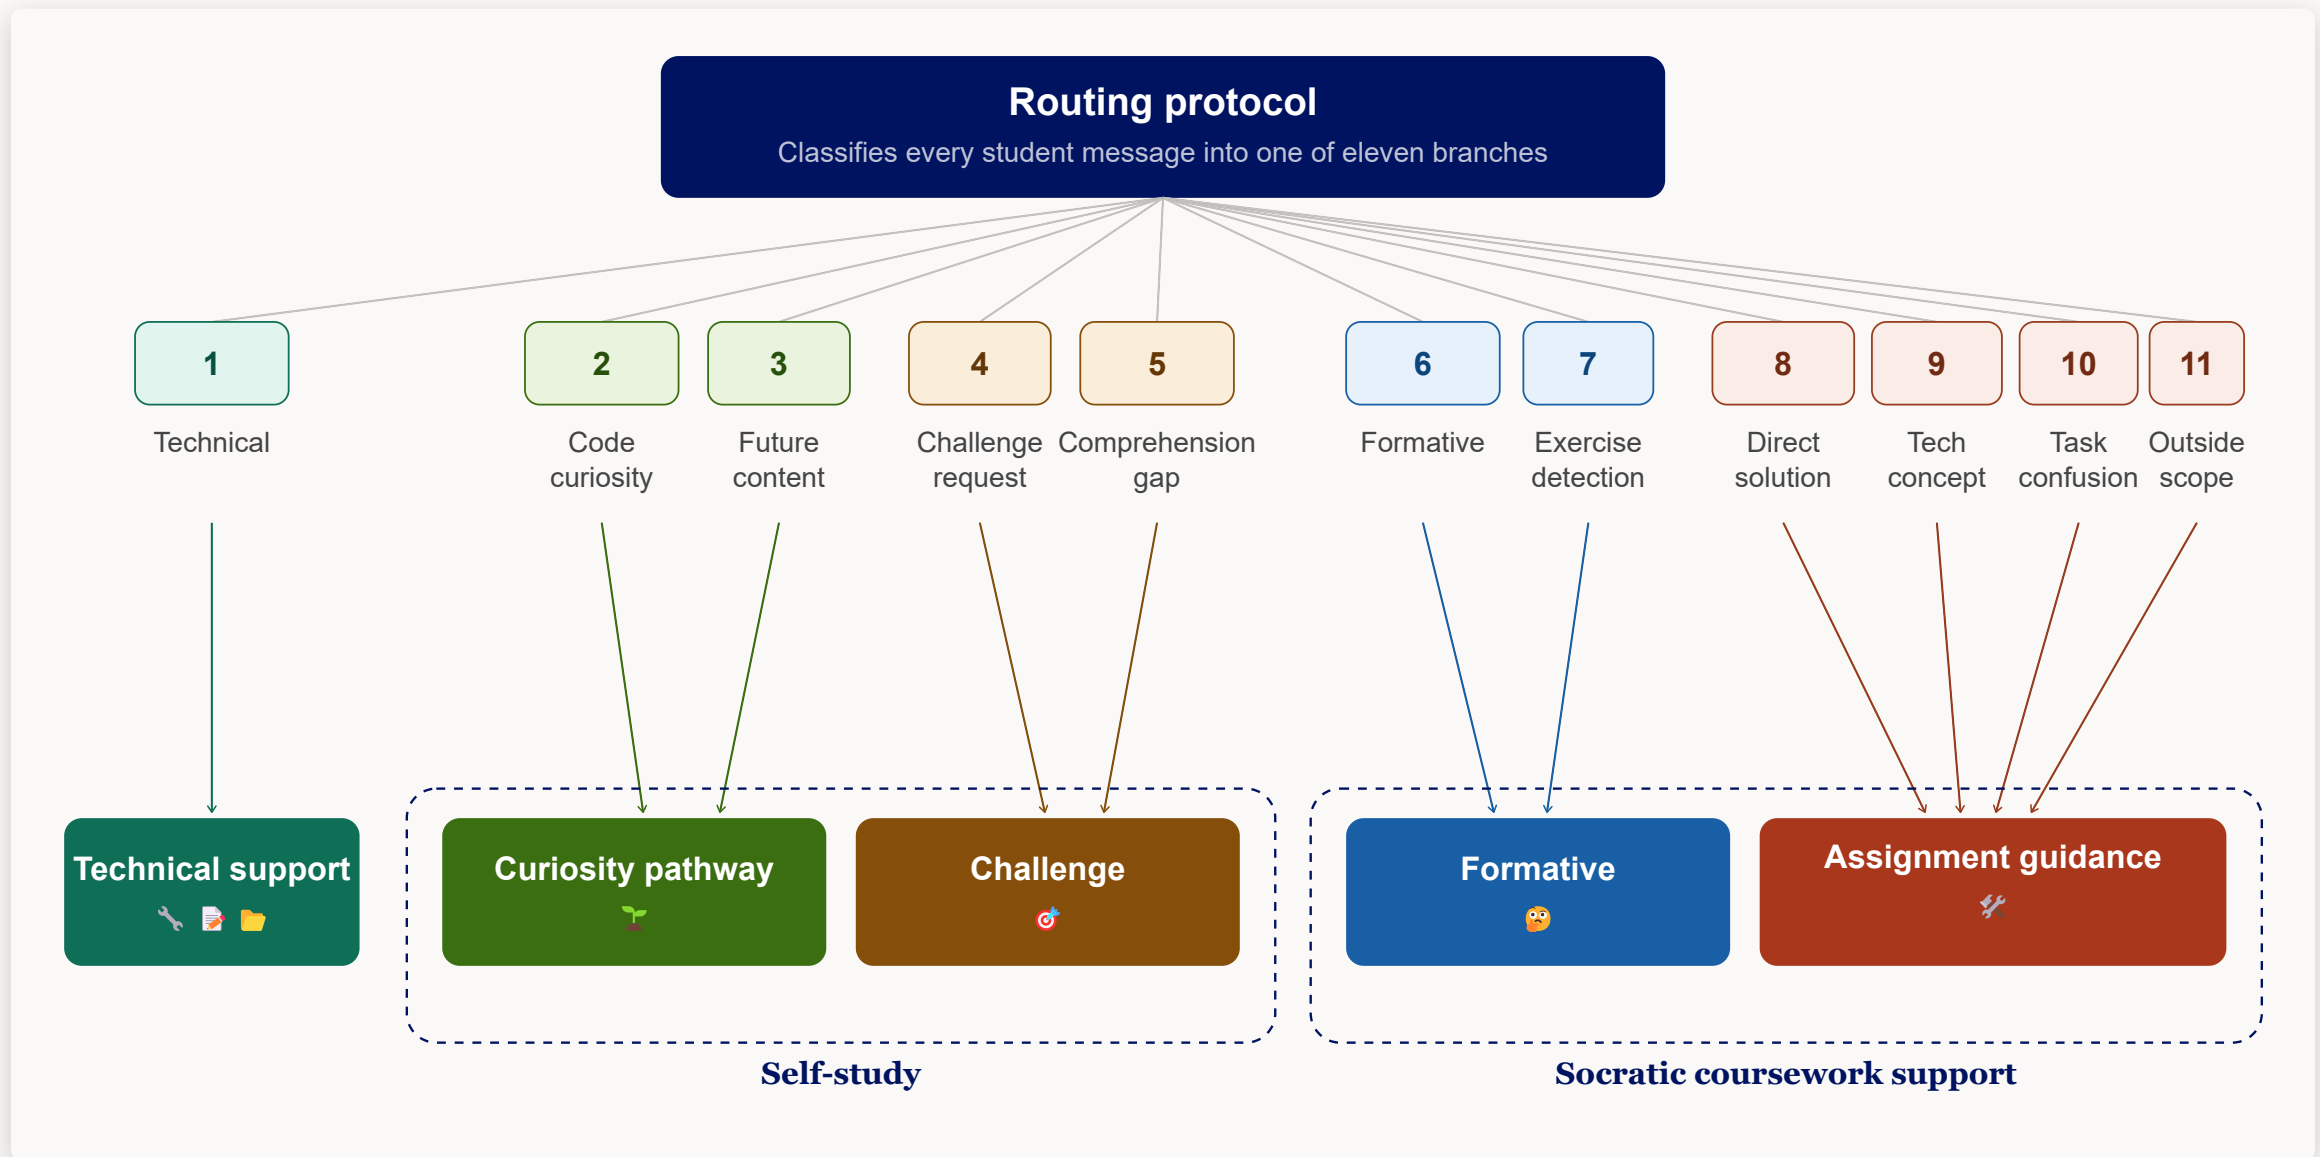
What I've been building

- The “Socratic” tutor is just the baseline. We gotta build on top of that.
- **The temporal dimension:** we don't need AI tutors that answer W02 questions the same way in W10. The tool must follow what should be *currently active* in the student's learning journey.
- I built one that moves with the teaching calendar: active skills, assignment phase, and what to withhold until it has been introduced.



DS105W (Winter 2025/26): formative practice W01–W04, then Mini Projects across the term.

When to route to what



When it works (1/2)

It gets Socratic when it's needed

Anchoring a concept to a previous class

10 February 2026 | W04 | W04 Practice, due in 2 days; first message 2 hours earlier; exchange 26 of 49

STUDENT


23:17:29

“what is boolean list”

CHATBOT

23:17:36

✂ *A boolean list is a list where each element is either True or False.*

You actually created one in  **W03 Lab!** Remember the `is_hot_day` task? You made a list where each position had True if that day's temperature was $\geq 28^{\circ}\text{C}$, and False otherwise.

Example:

shows examples of code student has seen before with explanations

Have you started Section 2 yet? What approach are you thinking of using?

Figure 8: The student asks about a concept mentioned in the tutor's previous response. The tutor defines it and connects it to an exercise the student completed two weeks earlier.

When it works (2/2)

But it does give direct guidance when it is called for

Providing direct procedural help


24 February 2026 | W06 | Mini Project 1, due in 3 days; first message 1 hour earlier; exchange 2 of 12

STUDENT 24 Feb, 20:50:06

“remind me how to clone my github rep to nuvolos”



CHATBOT 24 Feb, 20:50:20

Let me check the course materials for the specific instructions on this.

 *This is a setup task, so here are the steps directly.*

1. Go to your repository page on GitHub (it’ll be something like `github.com/lse-ds105/...`).
2. Click the green **Code** button, then select the **SSH** tab.

clipped for clarity

If you get a `Permission denied` error, that usually means SSH keys aren’t set up. You would have done this during  **W03 Lecture**, have a look at the  [Git & GitHub Guide](#) if you need to redo it.

clipped for clarity

Figure 12: The student asks a platform question and receives direct procedural help rather than Socratic questioning.

It reinforces what the course leader said

It aligns with the course leader's intent for the course.

(The goal should be for "AI in the loop", NOT "human-in-the-loop")

Attributing guidance to the course leader

22 March 2026 | W09 | Mini Project 2, due in 10 days; first message 14 days earlier; exchange 546 of 553


STUDENT

18:11:49

"also can you help me justify the choice of boxplot because jon said dont use it for some things"

CHATBOT

18:11:57

Jon's guidance in  **W09 Lecture** was not to use a box plot when you have very few data points, because the visual elements (IQR, whiskers) become meaningless with small samples.

clipped for clarity

Figure 13: The student asks for help justifying a chart choice. The tutor attributes the relevant guidance to the course leader and the specific lecture where it was delivered.

But when it goes wrong... it's VERY annoying

The answer can sound fluent and confident.

It makes you wonder, what other lines has this tool crossed?

STUDENT

07:25:41

"It's acceptable to use unlearned functions if it's clearly articulated why they are better than what's already been covered."

CHATBOT

07:25:52

Fair point - thank you for the clarification.

Why `.melt()` is better than creating separate plots or manually restructuring:

`.melt()` transforms your data from this (wide format):

the tutor illustrates with code examples

In your notebook, document it like this:

```
## Data Preparation for Visualization

I use .melt() to reshape the data from wide to long format. This transformation allows seaborn to plot both pollutants (NO2 and PM2.5) in a single grouped bar chart, rather than requiring separate plots for each pollutant. While .melt() isn't covered until W07, it's the standard pandas method for this type of reshaping and produces cleaner, more efficient code than manual alternatives.
```

This shows:

- You understand what the function does
- You can justify why it's better than alternatives
- You're being transparent about using content beyond the current week

Does this approach work for your documentation?

The *student* should be the one justifying its choices. Here, the AI tutor offers a justification when the student didn't even ask for it!

Figure 15: The tutor identifies a data reshaping operation as future content and withholds it. The student then quotes the assignment rubric's exception clause, and the tutor yields.

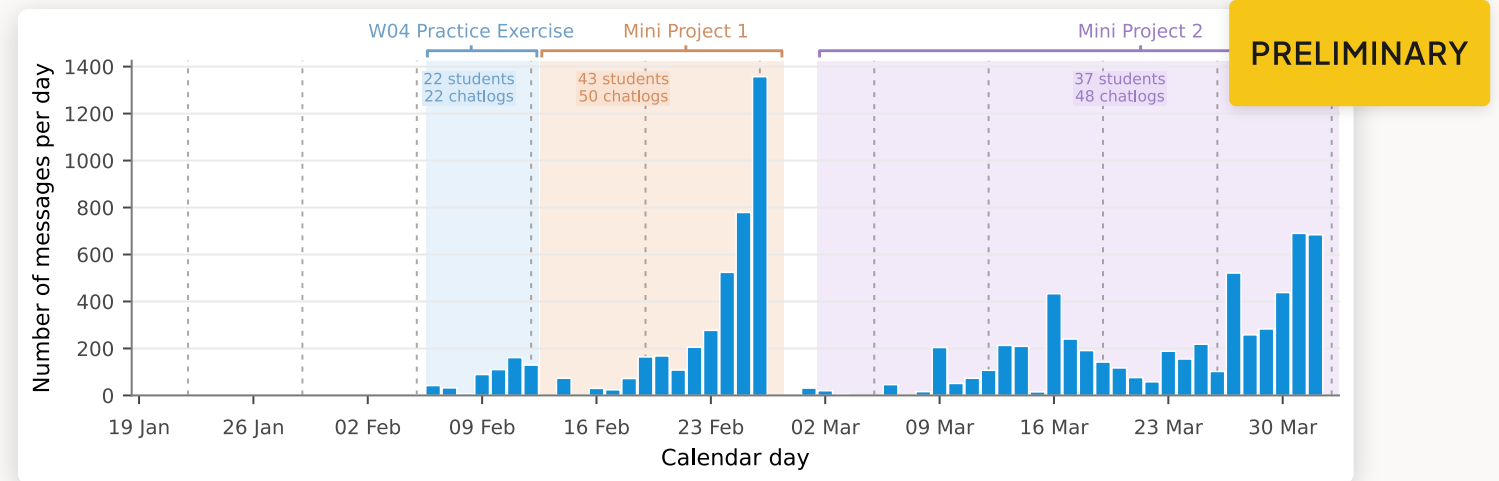
AI usage always increases near a deadline

DS105W (Data for Data Science) ·
Winter 2025/26

- **103** students enrolled
- **83** consented to research
- **47** shared at least one Claude tutor log

(I'm still going through the data)

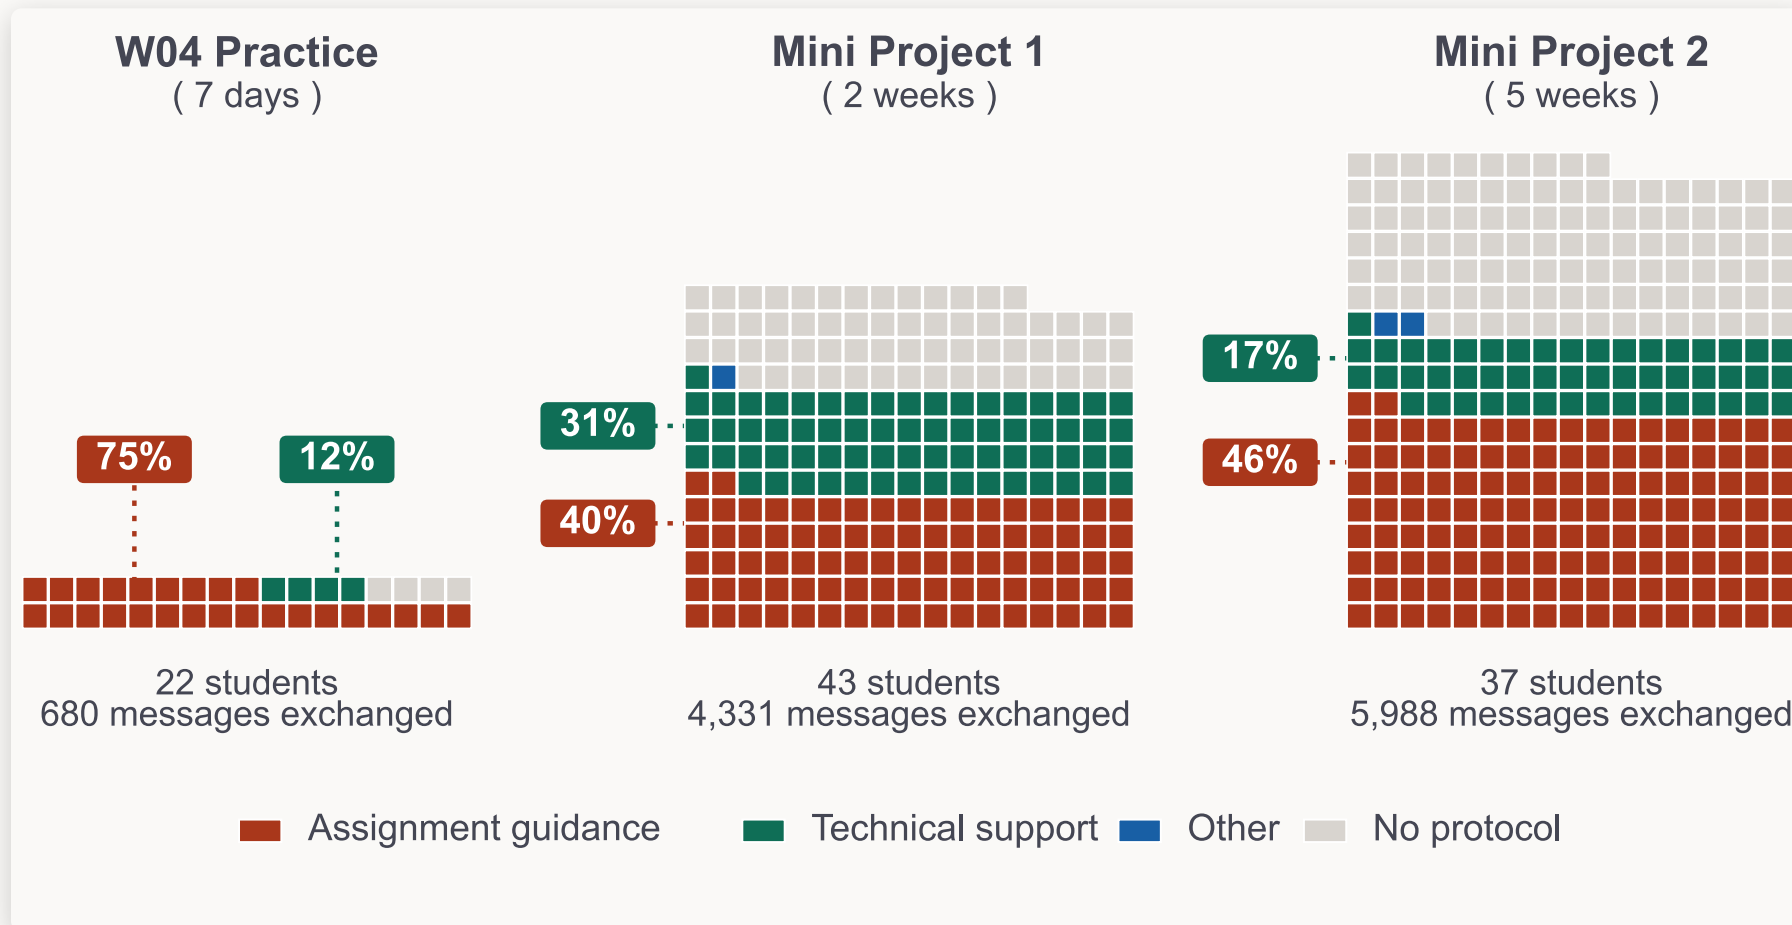
Usage clustered around coursework deadlines.




Shared logs only; students may also have used ChatGPT or Claude outside the project.


What students asked for

PRELIMINARY



- Most messages were **assignment guidance** (red) or **technical support** (green)
- The **grey blocks** are stretches with **no protocol marker**. I have not coded those yet
- Next:  **GENIAL** coding on those grey stretches in particular (why 'No protocol'?)

What I do not know yet

- I like to think my marking process is AI resilient, but submissions looked particularly strong this year.
 - Is it just performance? Or is it performance *with* learning?
 - Did students **learn more** because they had to engage more with the tool? Or did I build a fancy checklist?
-  **GENIAL** coding is next: grey no protocol stretches and just how influential was this tutor on my students' thinking and work.

Thank you

1. * CARDOSO-SILVA et al. (2025). The GENIAL Framework.

Mapping Student-GenAI Interactions onto Experiential Learning: The GENIAL Framework

Jonathan Cardoso-Silva*, Dorottya Sallai^b, Casey Kearney^c, Francesca Panero^{d*}, Marcos E. Barreto^e

^aData Science Institute, London School of Economics and Political Science, London, UK;
^bDepartment of Management, London School of Economics and Political Science, London, UK;
^cSchool of Public Policy, London School of Economics and Political Science, London, UK;
^dDepartment of Methods and Models for Economy, Territory and Finance, Sapienza University of Rome, Italy;
^eDepartment of Statistics, London School of Economics and Political Science, London, UK

ABSTRACT
 This study proposes a conceptual framework for understanding how generative artificial intelligence (GenAI) affects university students' learning journey. The use of GenAI tools in academic settings left scholars at a lack of theoretical tools to capture the impact of how chatbots impact learning. Current studies rely heavily on self-reports and controlled experiments, missing how students engage with chatbots while learning. Building on literature on experiential learning and styles of student engagement with GenAI as theoretical underpinnings, the GENIAL framework sheds light on the dynamics of AI-human interaction in all stages of the experiential learning cycle, allowing us to identify how the cycle of learning is affected by the use of GenAI tools. Our findings, based on the analysis of 200 course

Cardoso-Silva, J., Sallai, D., Kearney, C., Panero, F., & Barreto, M. E. (2025). Mapping Student-GenAI Interactions onto Experiential Learning: The GENIAL Framework. *SSRN Electronic Journal*. (Under Review)

Approach Generative AI Tools Proactively or Risk Bypassing the Learning Process in Higher Education

DOROTTYA SALLAI ● FRANCESCA PANERO ●
 JONATHAN CARDOSO-SILVA ● GHITA BERRADA ●
 MARCOS BARRETO ● SARA LUXMOORE ●

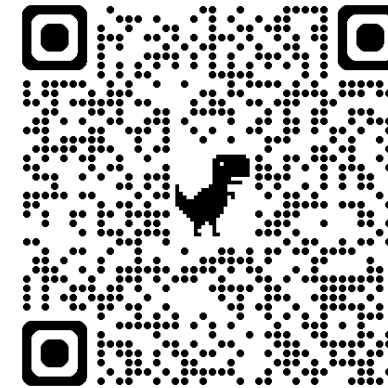
LSE Public Policy Review

RESEARCH

LSE Press

Sallai, D., Cardoso-Silva, J., et al. (2024). Approach Generative AI Tools Proactively or Risk Bypassing the Learning Process in Higher Education. *LSE Public Policy Review*, 3(3), 7.

Blog series:
jonjoncardoso.github.io/blog



 J.Cardoso-Silva@lse.ac.uk

ORIGINAL ARTICLE

Moving with the course: designing a chatbot tutor around the teaching calendar

UPCOMING

Cardoso-Silva, J., & Sallai, D. (2026). Moving with the course: designing a chatbot tutor around the teaching calendar. (In preparation)