Algorithmic Thinking in a World of Abundant Computation

VSC User Day 17/12/25

crunch

Who am I?

LP Kerkhove

Founder, CTO

+32 477 25 63 51
lp@crunchanalytics.be

Academic

MS Business Engineering
PhD Operations Research
Professor at Ghent University

Founder, CTO @ Crunch Analytics (°2016)

Specialists in data-driven innovation for fashion retailers
My focus = how to improve things with data and algorithms

Author

Data-driven Retailing: A Non-technical Practitioners' Guide (2022)
Springer, Management for Professionals

crunch

Our solutions often combine multiple technologies

Data platform

Creating or extending the data-architecture to deliver stable and cost-effective solutions

Optimization

What is the best compromise on portfolio level? How can we adhere to operational constraints?

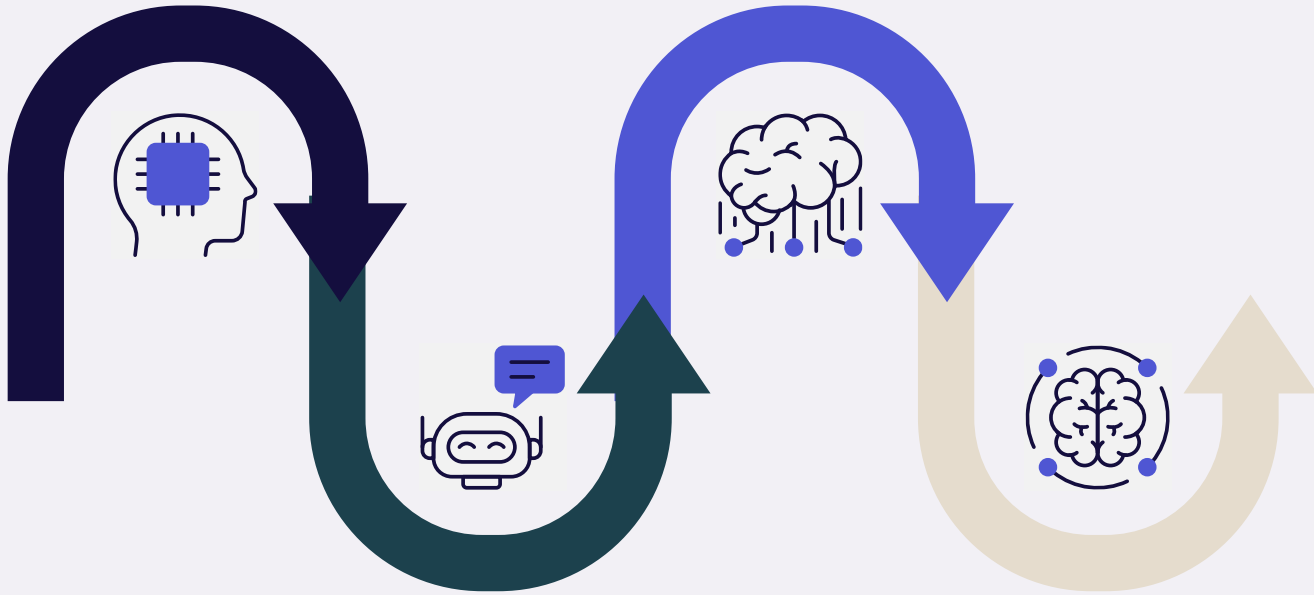

Predictive models

What demand can we expect? How will it react to price changes?

User interface

Integrate the tools into business processes

crunch

Who here teaches students?

crunch

Who has noticed students becoming better at programming the last 2-3 years?

crunch

Who has noticed students becoming better at critical thinking the last 2-3 years?

crunch

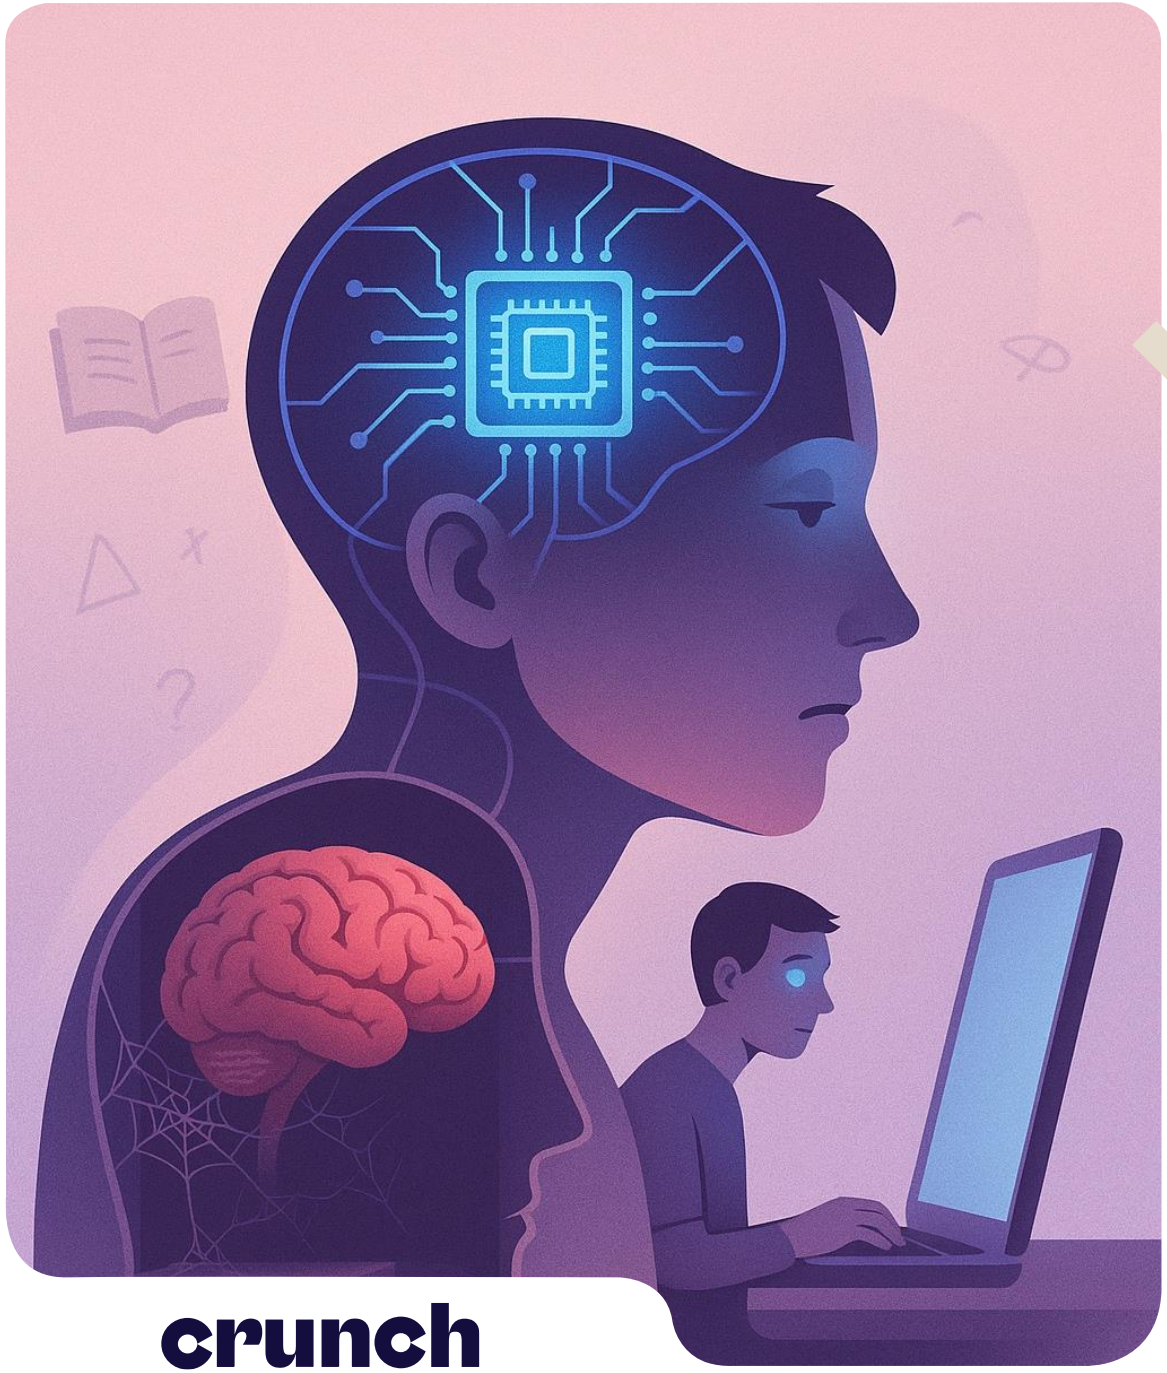

crunch

Is gen AI making idiots of us all?

01

Entry exam for studying medicine in Belgium

47%
Students passed in 2025

19%
Students passed in 2024

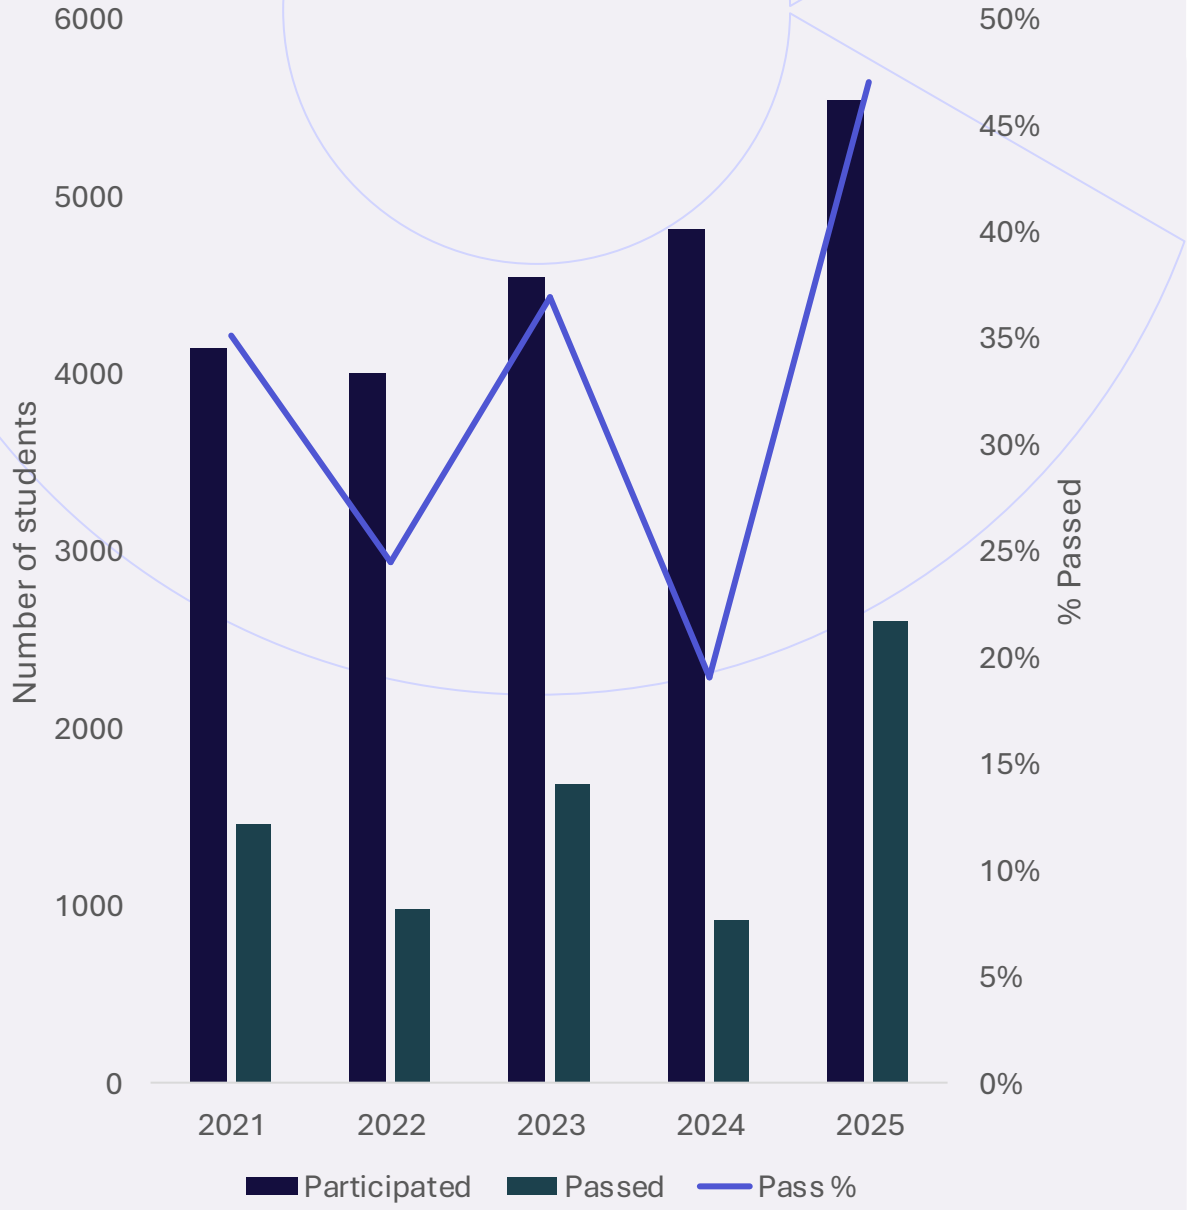

crunch

Thought for 4m 51s

Hieronder staan de slaagcijfers (Vlaanderen, **toelatingsexamen arts**), laatste 5 jaar. Ik toon per jaar het aantal **deelnemers** (geldige deelnemers) en het aantal **geslaagden** (gunstig gerangschikten), met het **slaag%** ter referentie.

Jaar	Deelnemers	Geslaagd	Slaag%
2025	5 544	1 740	31,4% Vlaanderen.be
2024	4 814	1 723	35,8% Vlaanderen.be
2023	4 545	1 598	35,2% Vlaanderen.be
2022	4 005	1 279	31,9% Vlaanderen.be
2021	4 144	1 284	31,0% Vlaanderen.be

Context bij vermeende ChatGPT-fraude

- In **juli 2025** werden drie kandidaten betrapt op het gebruik van ChatGPT tijdens de toelatingsexamens; ze werden uitgesloten. Daarna startte een breder onderzoek en stapten meerdere studenten naar de rechter. Nieuwsblad +2
- Media berichtten over klachten en een doorlichting van het examen en procedures door de Vlaamse minister van Onderwijs. HLN +1

Entry exam for studying medicine in Belgium

31%
Students passed in 2025

32%
Students passed in 2024

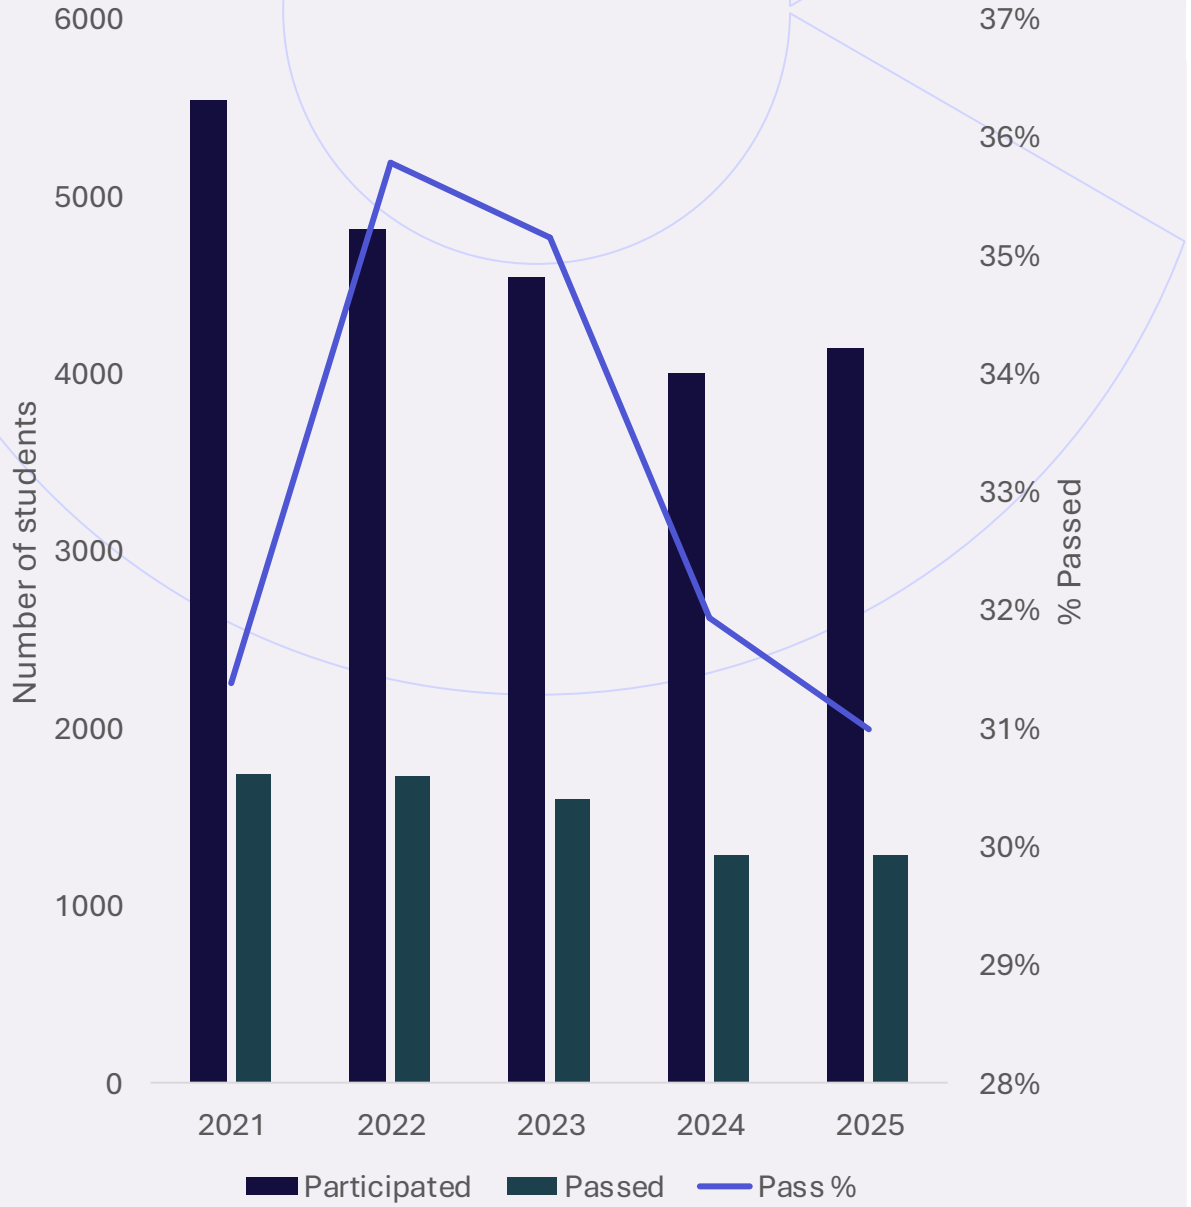

crunch

Classic methods of learning are being bypassed

- Asking questions in class
- Doing research assignments
- Writing assignments
- Book reports
- Math homework

In basic education, but also – and dangerously – in university education.

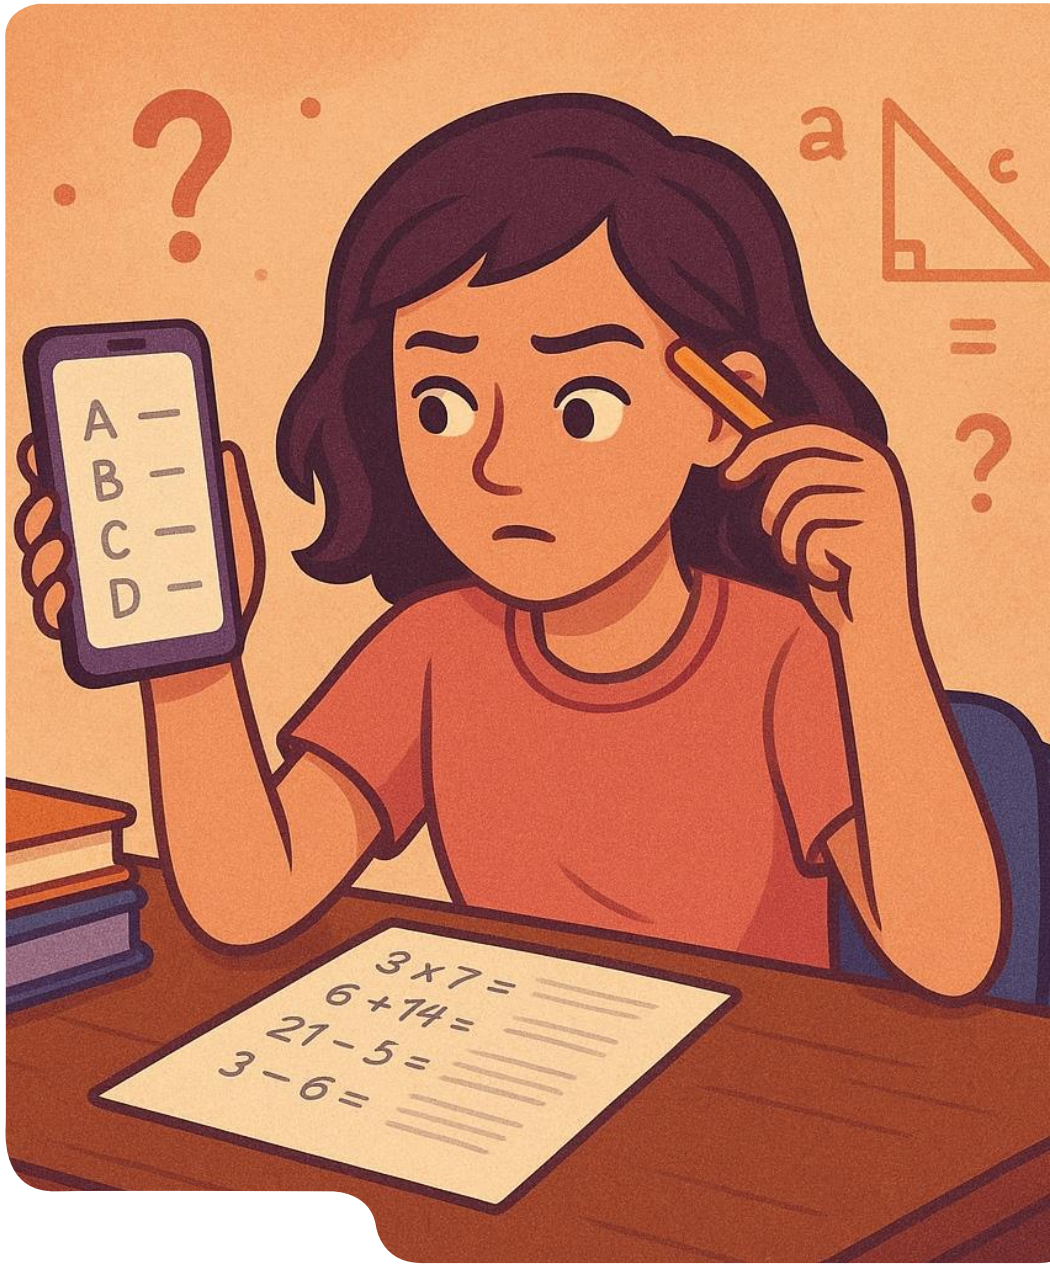

crunch

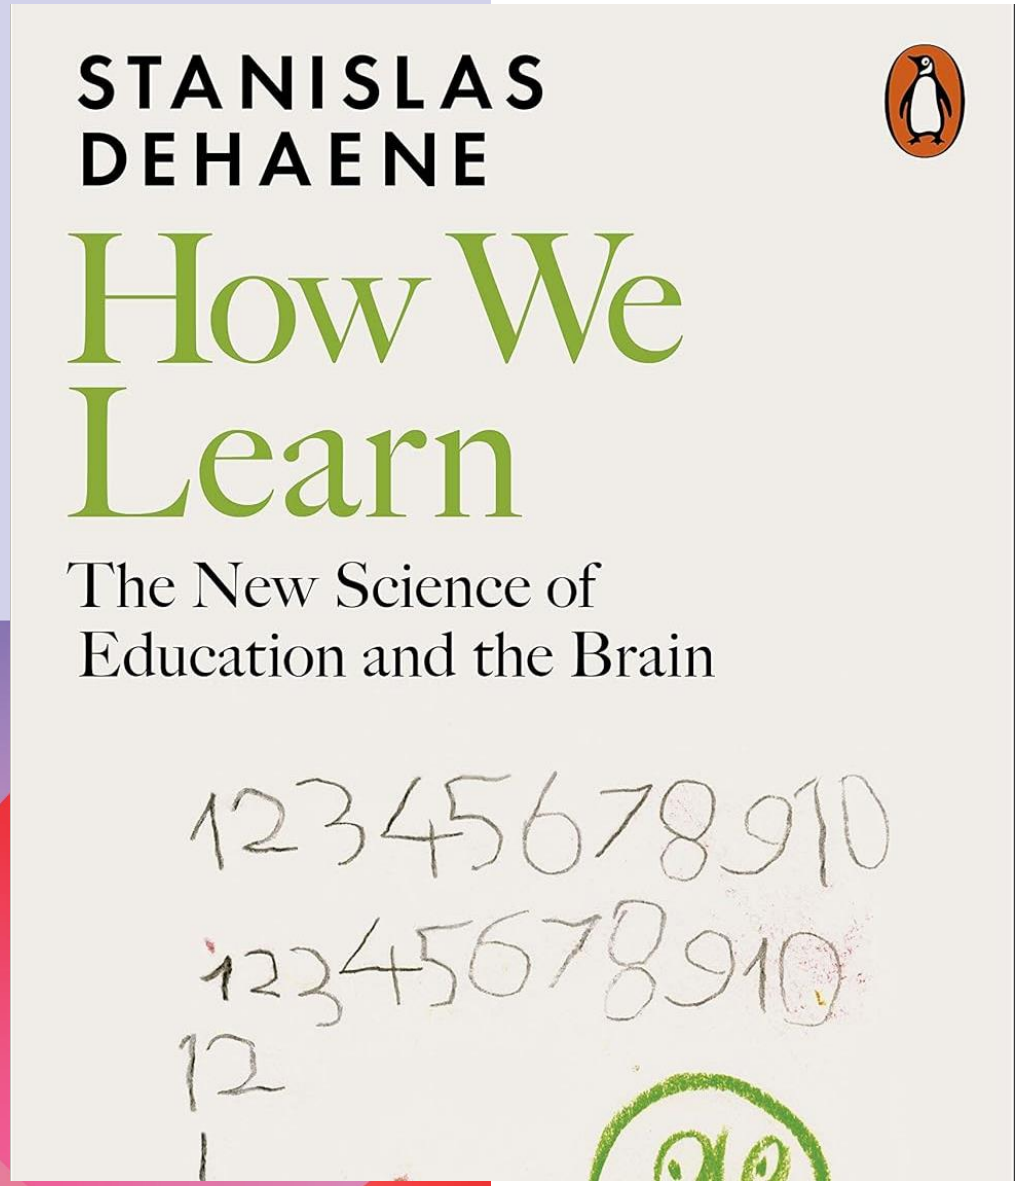

What is learning really?

- Committing information to long term memory
- Making it possible to combine multiple pieces of knowledge to form new ideas, the idea of **interleaving**
- Growth comes from **deliberate practice**, revisiting problems and challenging ourselves
- Neural paths are strengthened by **retrieval practice** – using what you know

crunch

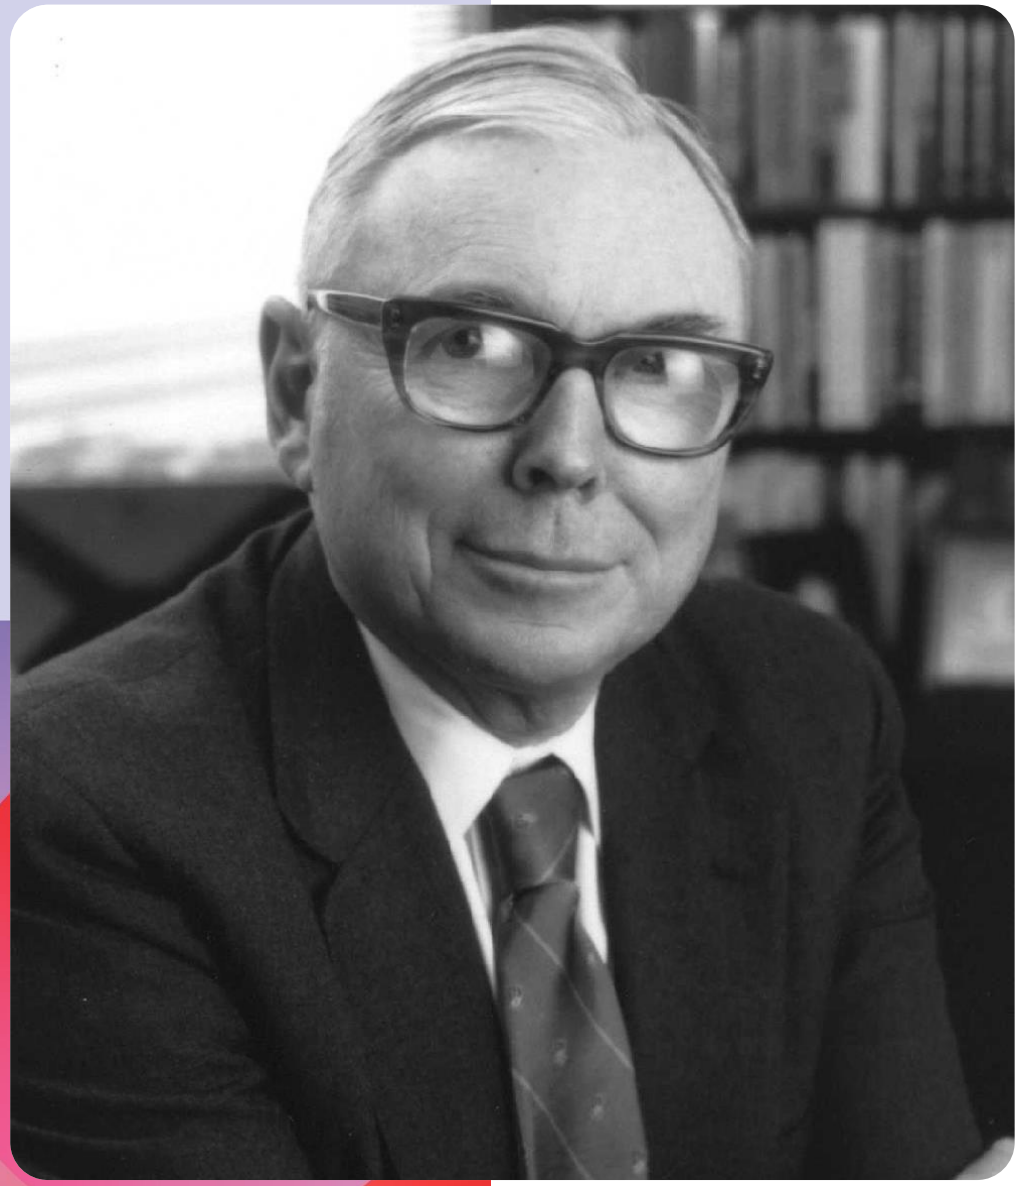

Latticework of mental models

You must know the big ideas in the big disciplines and use them routinely — all of them, not just a few. Most people are trained in one model — economics, for example — and try to solve all problems in one way. You know the saying: To the man with only a hammer, every problem looks like a nail. This is a dumb way of handling problems.

Charlie Munger

crunch

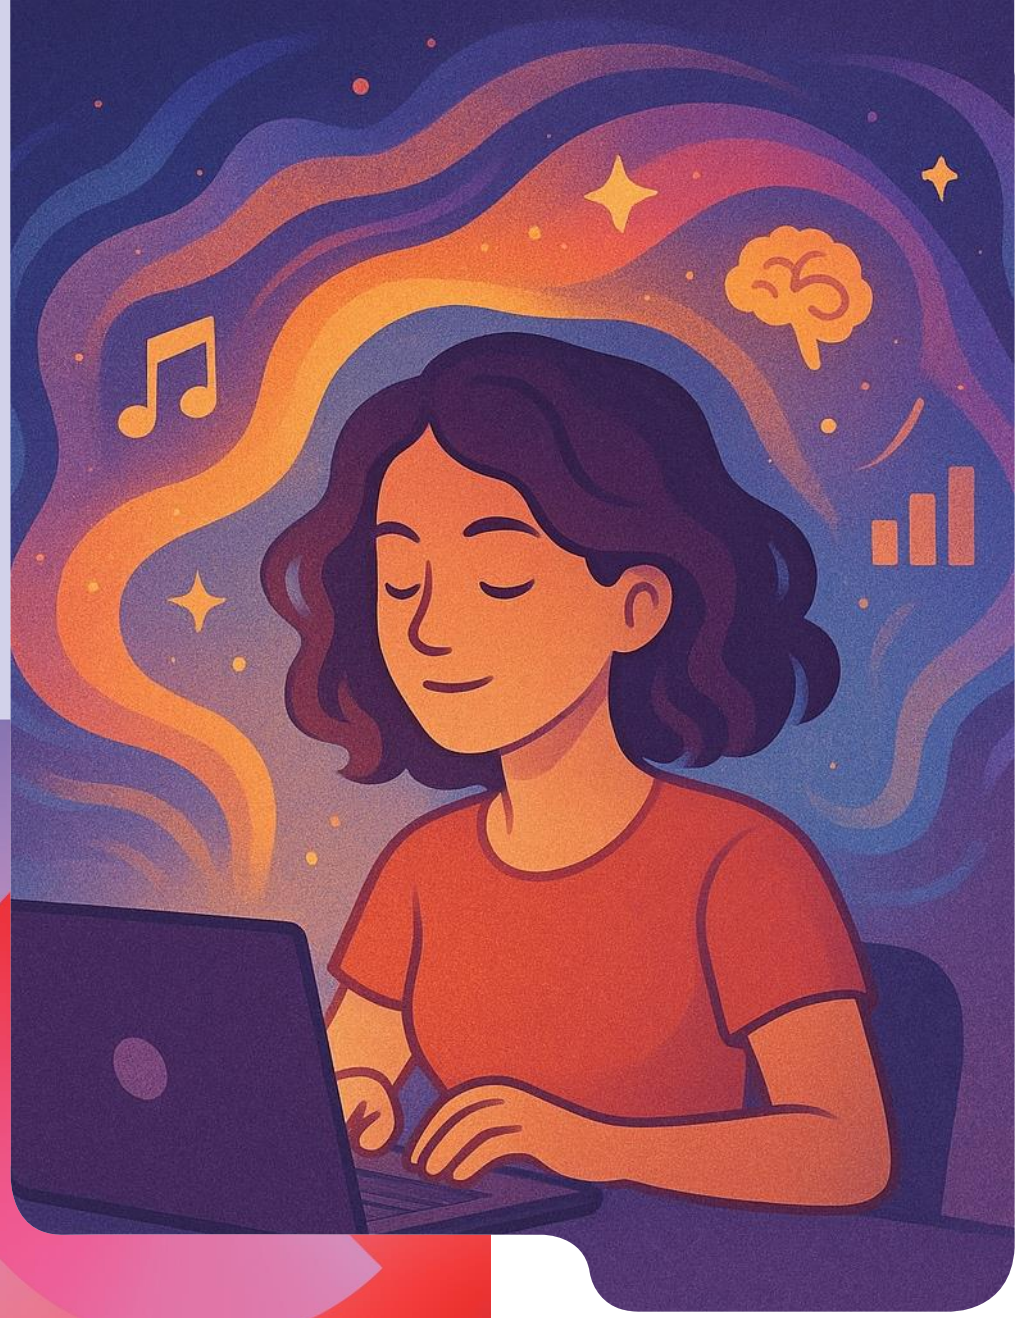

Flow state

- Complete concentration
- Altered time perception
- A balance between capabilities and challenge

… and generally a feeling of enjoyment in the task at hand

But the very nature of the question – wait – response interaction (“**vibe coding**”) with a **LLM model breaks up this flow state**. Further encouraging multi-tasking and lessening overall output (and enjoyment).

crunch

The Chinese room

- A person who does not understand Chinese
- Detailed rule books (syntax) to translate Chinese into English

But does that mean that the room “understands”?
(≈ semantics) Opinions differ

- **John Searle**: The fact that the person does not understand, means that we cannot say that the system understands.
- **Daniel Dennett**: As a whole there is no fundamental difference between the Chinese room and the processes in a biological brain.

crunch

The Chinese room

- A person who does not understand Chinese
- Detailed rule books (syntax) to translate Chinese into English

But does that mean that the room “understands”?
(≈ semantics) Opinions differ

- **John Searle**: The fact that the person does not understand, means that we cannot say that the system understands.
- **Daniel Dennett**: As a whole there is no fundamental difference between the Chinese room and the processes in a biological brain.

Perhaps more importantly, does over-use of LLMs risk use becoming much like the man in the box? Doing our job successfully, but not understanding what we do? Should we care as an organization?

crunch

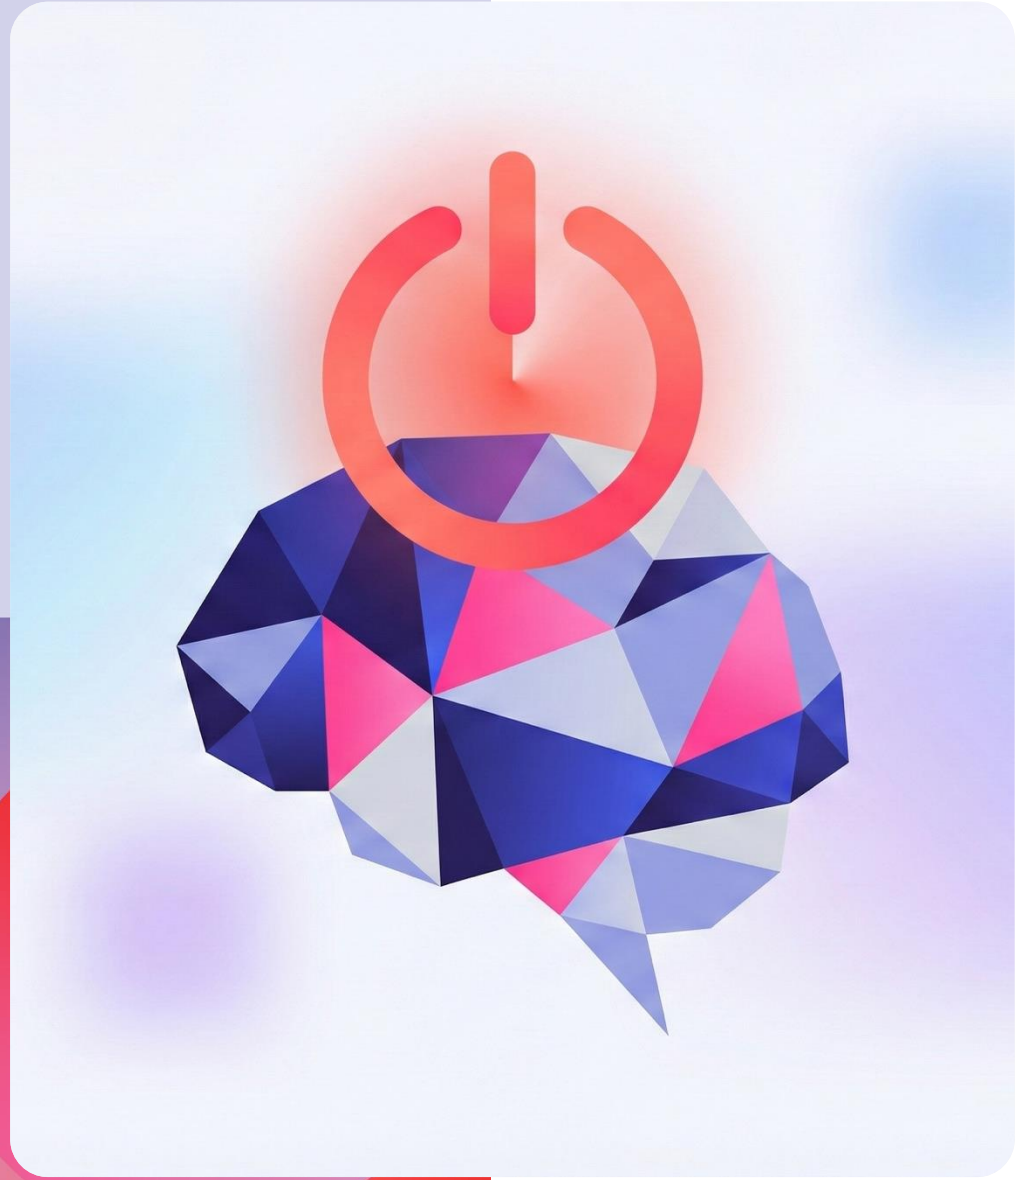

Abstraction or laziness

As the **black box** becomes more capable,
we appear to be less inclined to try and open the black box.

Some may argue that this is a luddite's point of view, and this is programming continuing its progress of abstraction further away from machine code.

But – I think it more likely that it is just people becoming intellectually lazy.

crunch

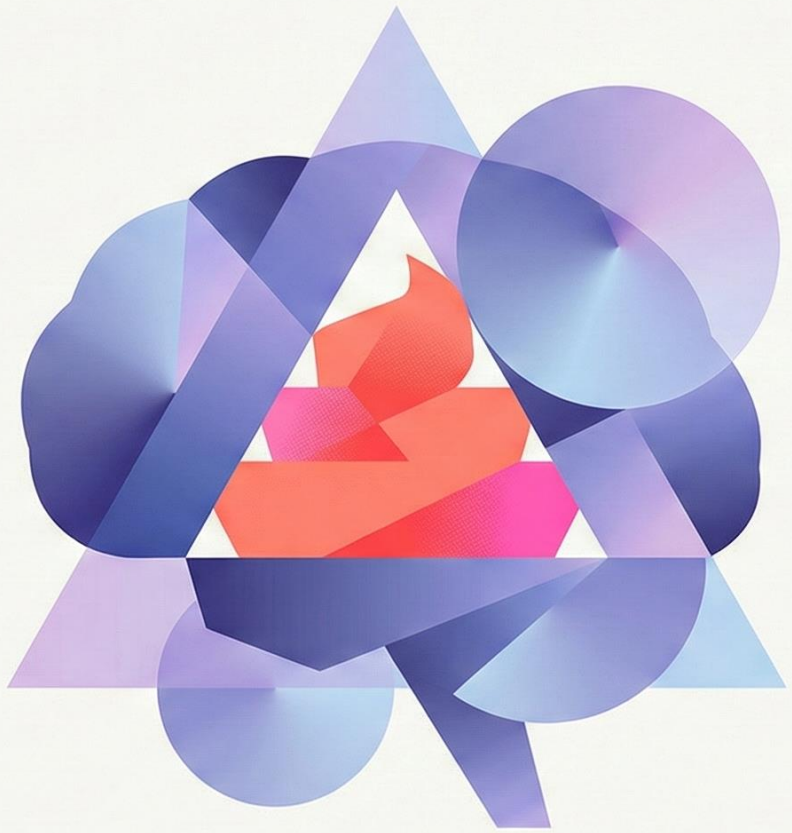

The workplace aka - idiots at scale

02

crunch

[...] I also don't understand why it works, but somehow it does [...]

On the logic used to set prices for more than 100,000 products

crunch

[...] I don’t see how it can be so expensive, we are already using AI to do our purchasing. I just upload my spreadsheets and ask ChatGPT how much to order. [...]

Purchasing director of organization with >200mio€ turnover

crunch

Project owner: Can you explain how this related to our revenue streams?
Analyst: I do not know anything about the revenue streams of XXX

Conversation after 2 months of data analysis had taken place

crunch

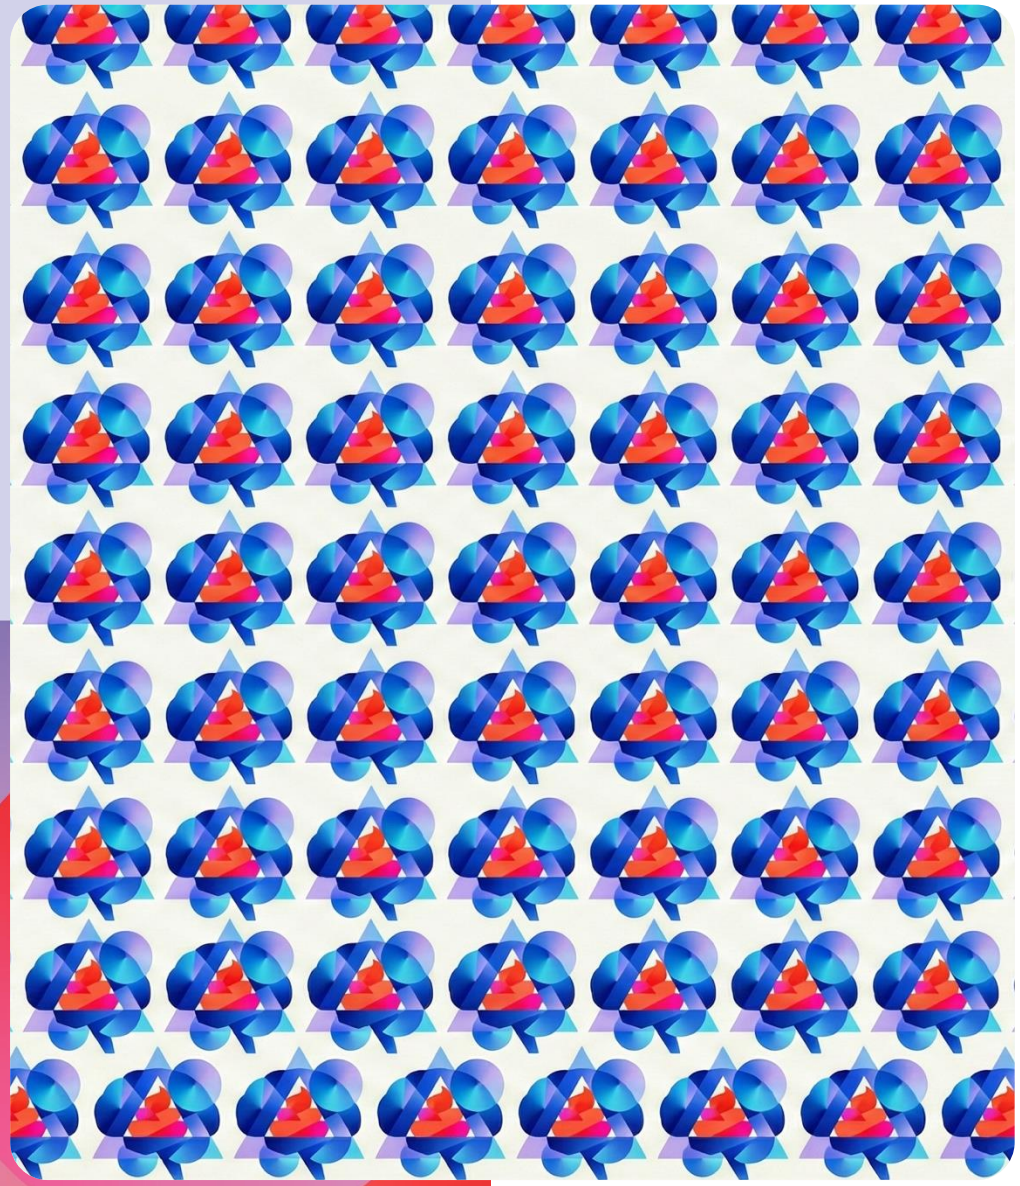

Automating stupidity

- A cry for **automation used to be the ideal justification to inject more intelligence** into a process.
- Come for the automation – stay for the intelligence
  - Statistically sound forecasting
  - Optimization models
  - …
- **Now we have tools that enable the automation of stupidity**

crunch

“Software gets slower faster than hardware gets faster.”

Wirth’s Law

Who here uses an IDE?

Does it use less than 500MB of RAM?

lkl2050

How come the Logi options + takes so much RAM in the background?

Questions

I am on a M1 Pro Mac. As a background app that changes a few costumed buttons on my mx master mouse, how can it cost so much RAM?

Code Helper (GPU)	477.0 MB
https://www.reddit.com	459.5 MB
Logi Options+	409.3 MB
Pages	
python3.9	

69 175 Share

Enshittification

Your computer probably contains many examples...

crunch

Junior roles a

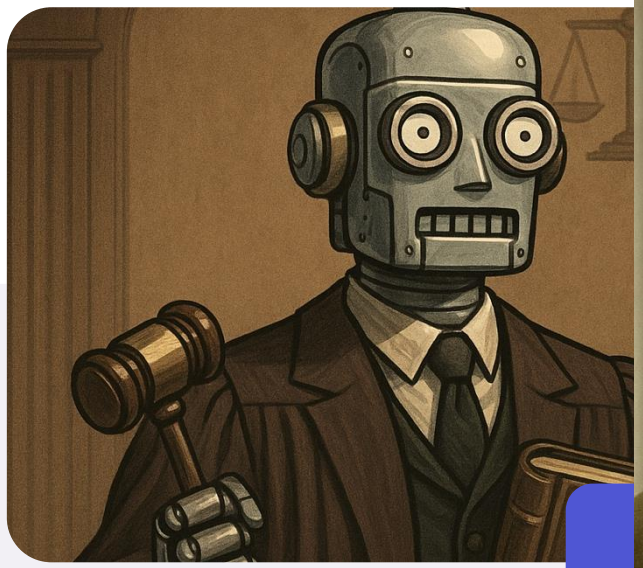

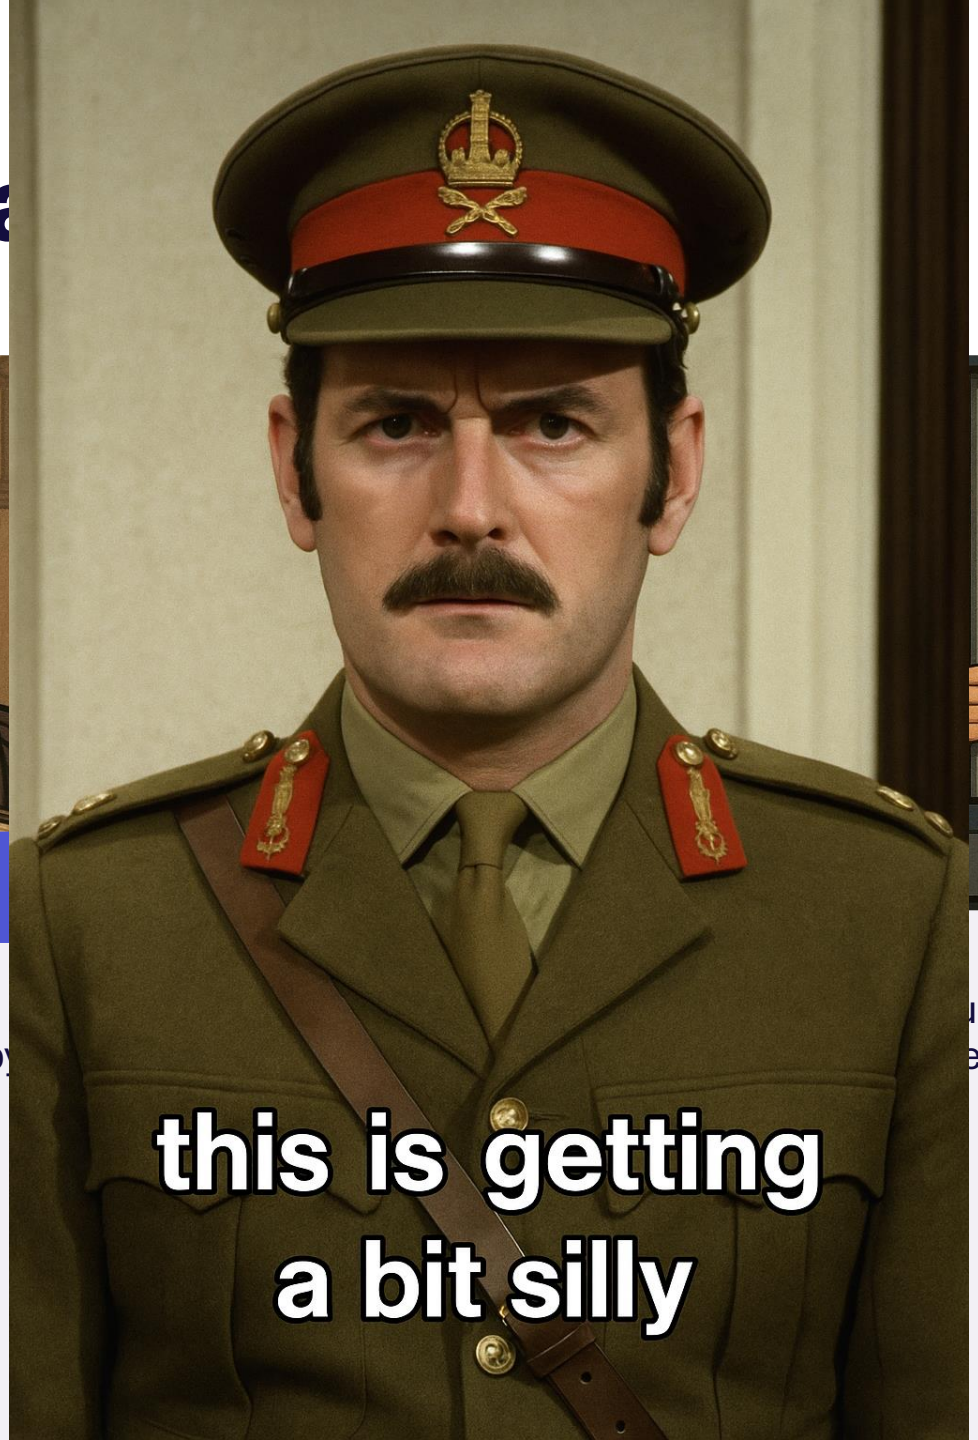

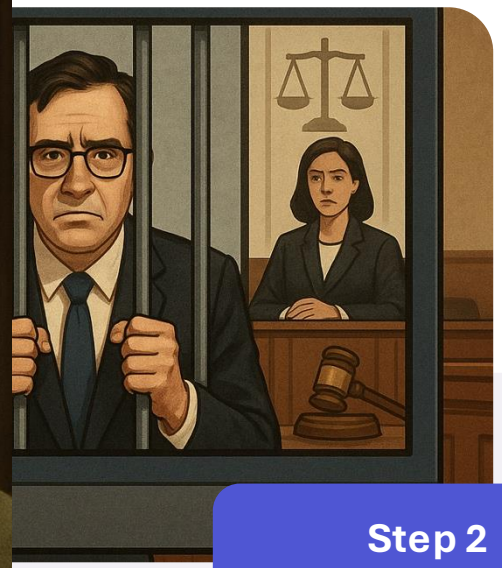

Step 2

Law firms are using AI to perform
and research traditionally done b

ulated cases and research to train
et them experience

crunch

Lower-skilled employees benefit most from genAI tools

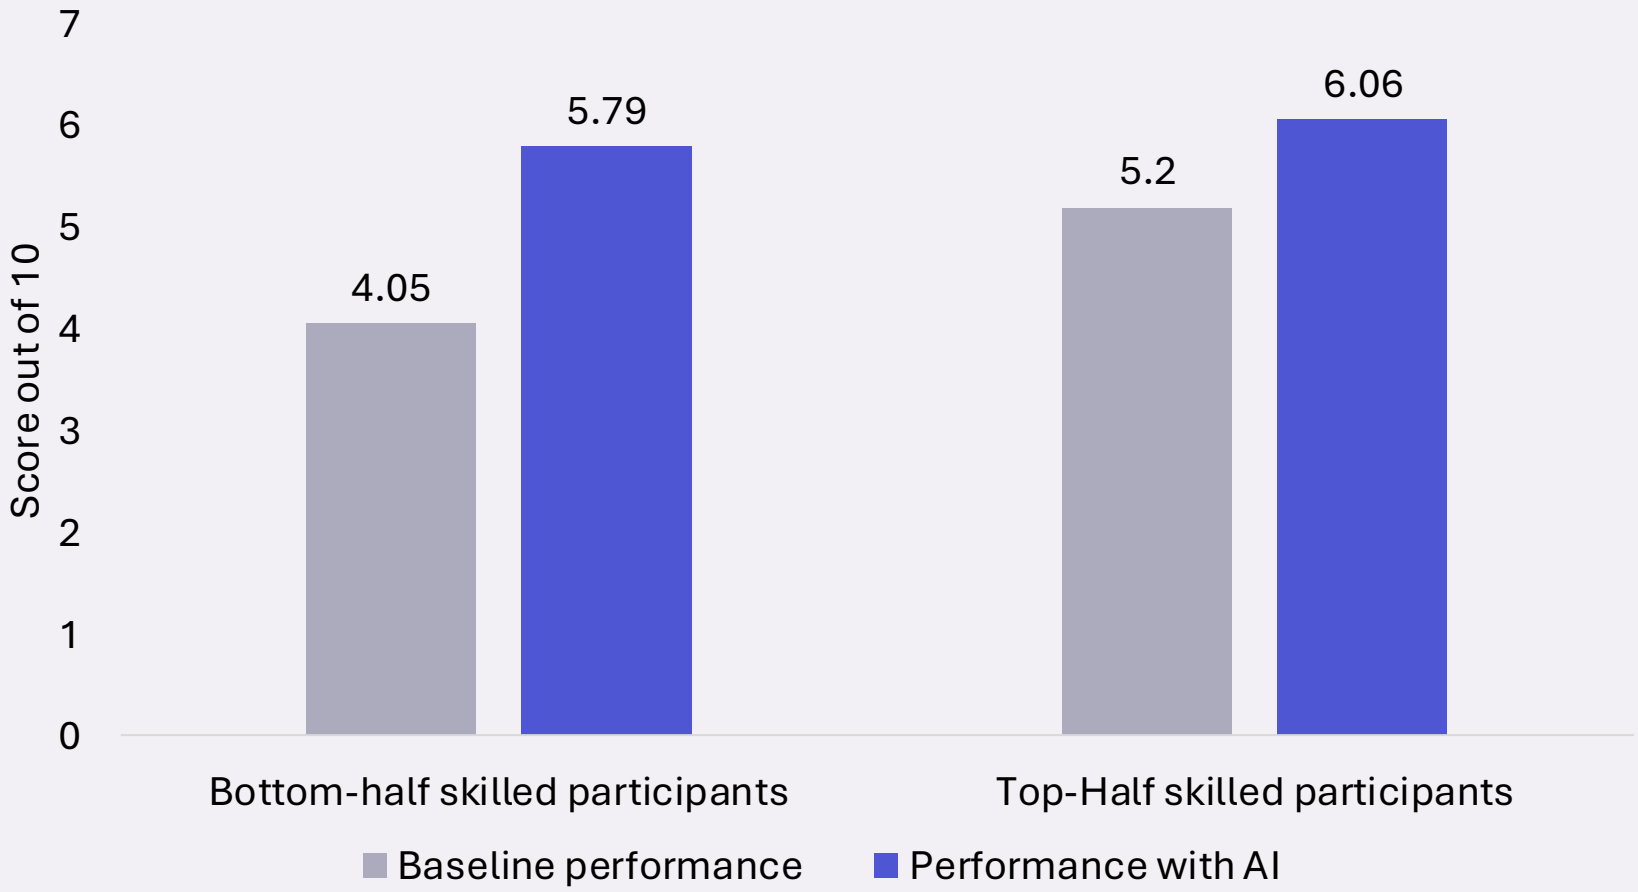

For **tasks well suited for the use of generative AI**, the performance increase is noticeably larger for lower-skilled participants.

Step outside your comfort zone!

crunch

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., ... & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).

… and using generative AI for complex problem solving can be problematic

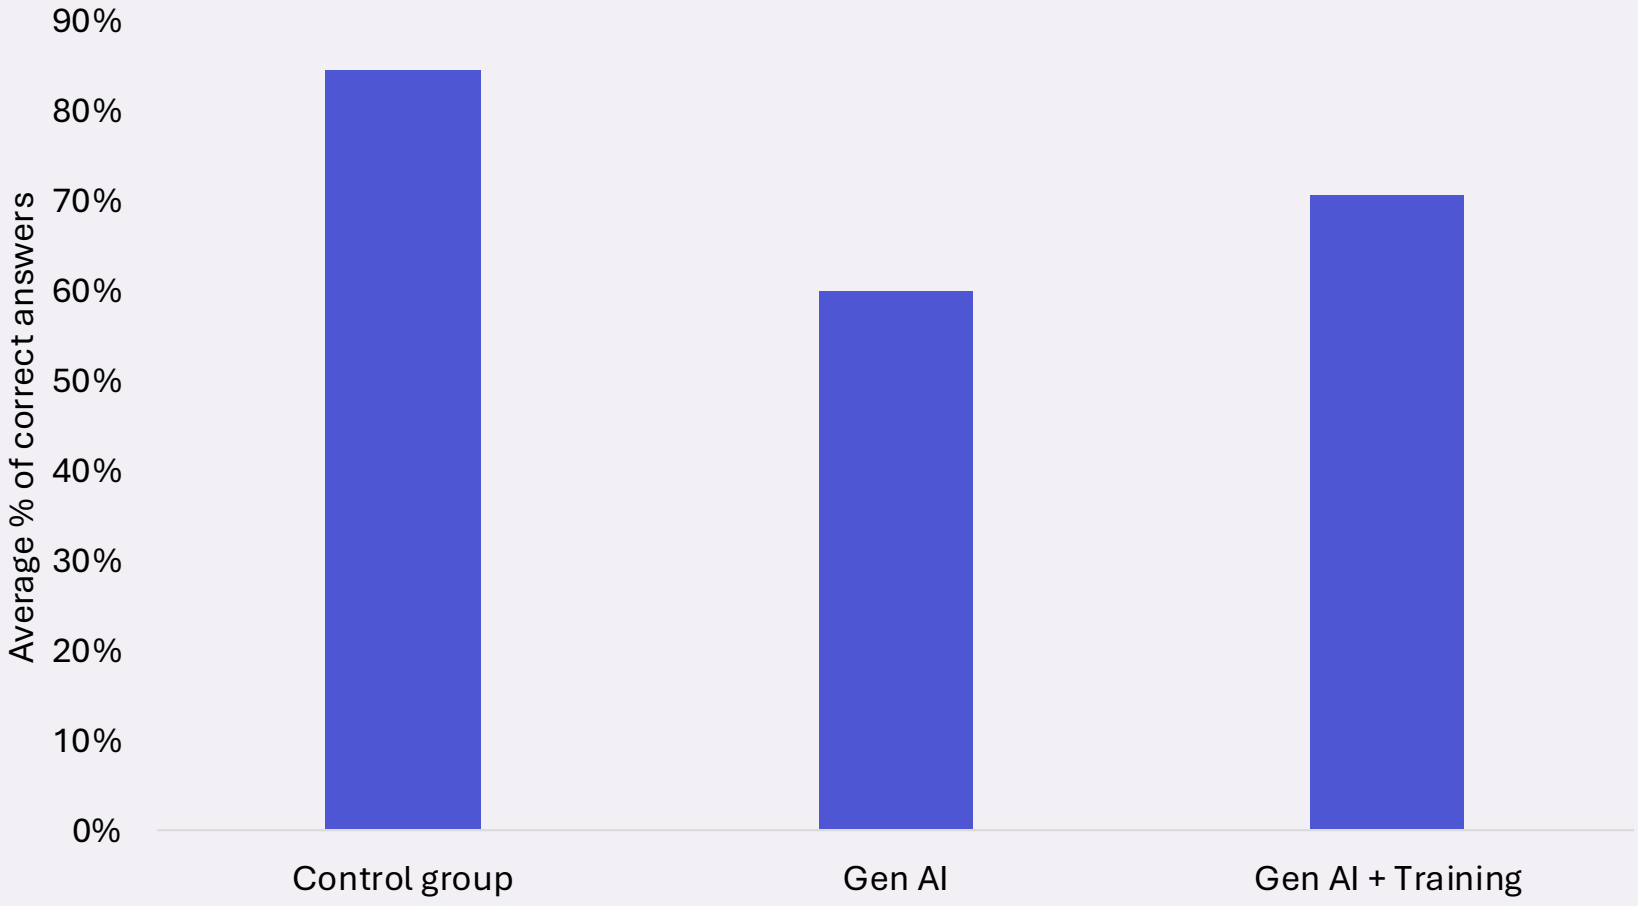

- When tasked with a complex problem: The analysis of a complex business case with multiple important details and several sources of information.
- **The use of generative AI was observed to decrease the quality of problem solving.**
- Starting complex problem solving by asking a question to Chat GPT appears to be a poor strategy.

crunch

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., … & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).

So what now?

03

crunch

I. Train competent people

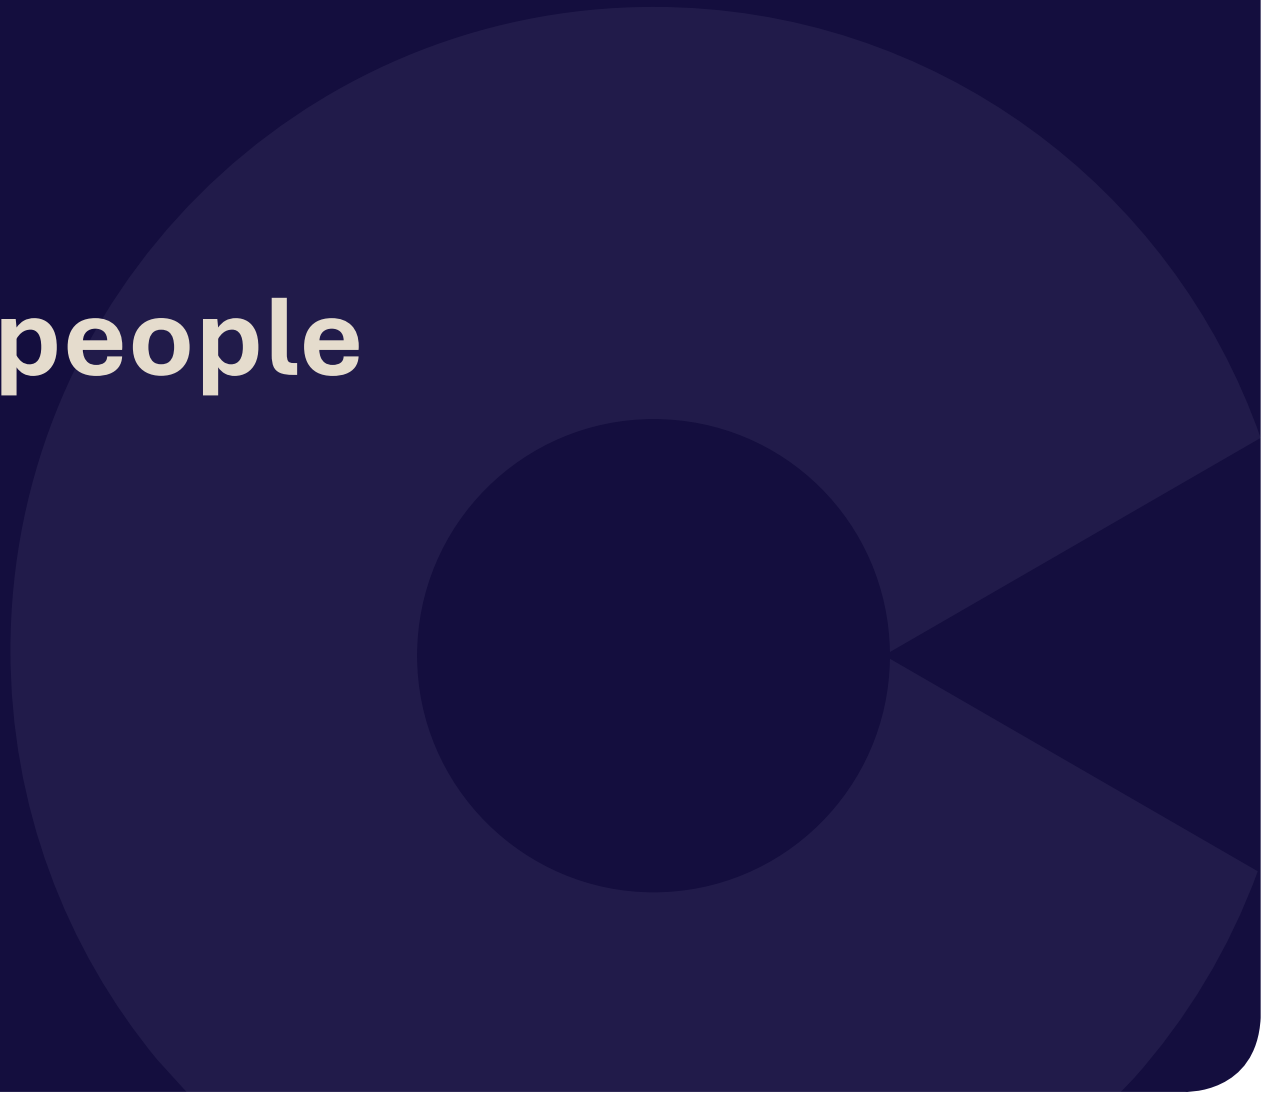

crunch

The Oxford tutorial

- Students meet bi-weekly with a tutor - One on one
- The assignment is writing an essay on a topic
- Writing (and thinking) is discussed in depth:
  Challenge arguments, refine reasoning,
  and push the student's thinking
- Dialogical - Socratic method

... but this is exponentially more work than the current 'submission culture'

crunch

Teaching should focus on application

... But beware of bullshit 'active learning'

- Find ways to get students to apply material on problems that have real-world complexity
- If a task is too easy to automate with LLMs, it might no longer be a skill worth acquiring
- The task should be **uncomfortable**, requiring judgement calls, extra data to be collected, trade-offs that are not self-evident,...
- How much a student learns will inevitably be controlled more by the student than in the past

crunch

Deliberate practice

Build skills with the end in mind

- Learning Spanish to be able to do basic communication and translation
- Learning Spanish to understand fundamental cultural differences and thinking patterns

Spend time on deliberate practice

- Highly structured, goal-oriented training involving focused effort and tasks pushed just beyond one’s comfort zone.

Know when to disconnect

- When thinking about how to approach complex problems

crunch

II. Design with domain knowledge

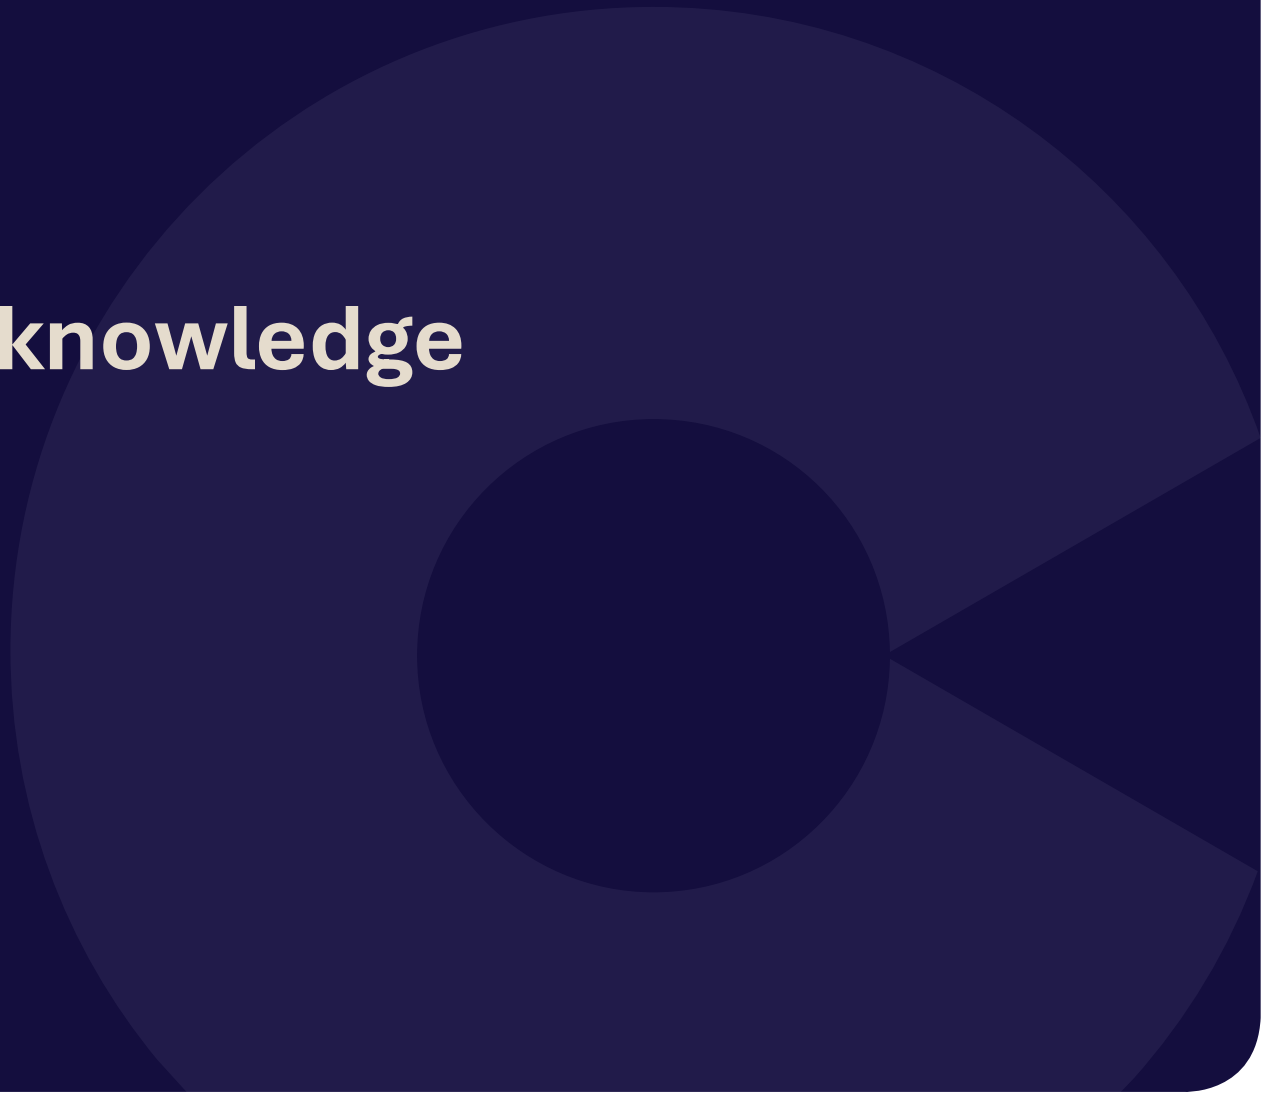

crunch

Democratizing math
... not skipping it altogether

- Give users the ability to formulate their own constraints, and as soon as the third constraint is entered you have a conflict
- This is a typical context, where intelligently designed applications could use the right tools for the calculation, and softer methods for the interpretation

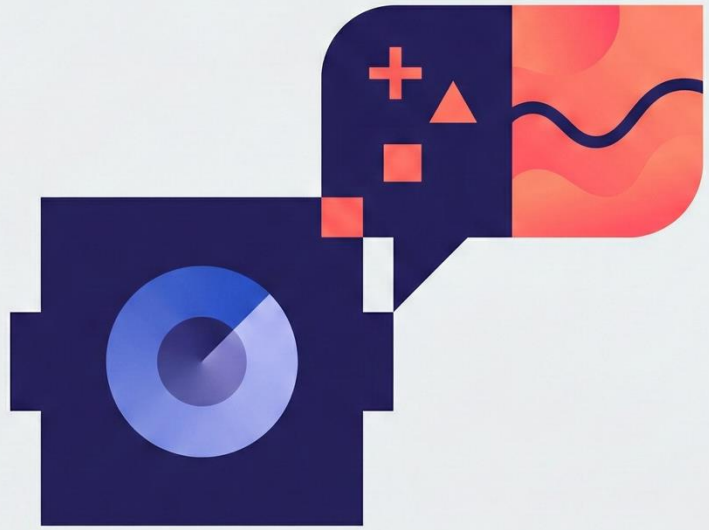

User: "Why is the Bosch Pro Drill not in the final assortment? It has the highest margin!"

LLM: "It was excluded because you hit the 'Maximum Inventory Capital' constraint for Power Tools.”

Action: I recommend shifting €5,000 budget from the 'Gardening Gloves' category (where the shadow price is only 0.2) to Power Tools. This would allow us to stock the Bosch Drill and increase total profit by €11,500."

crunch

Dynamic pricing

Example architecture

- Keep LLM models and “vibe coded code” away from the core functionality of the specific application.

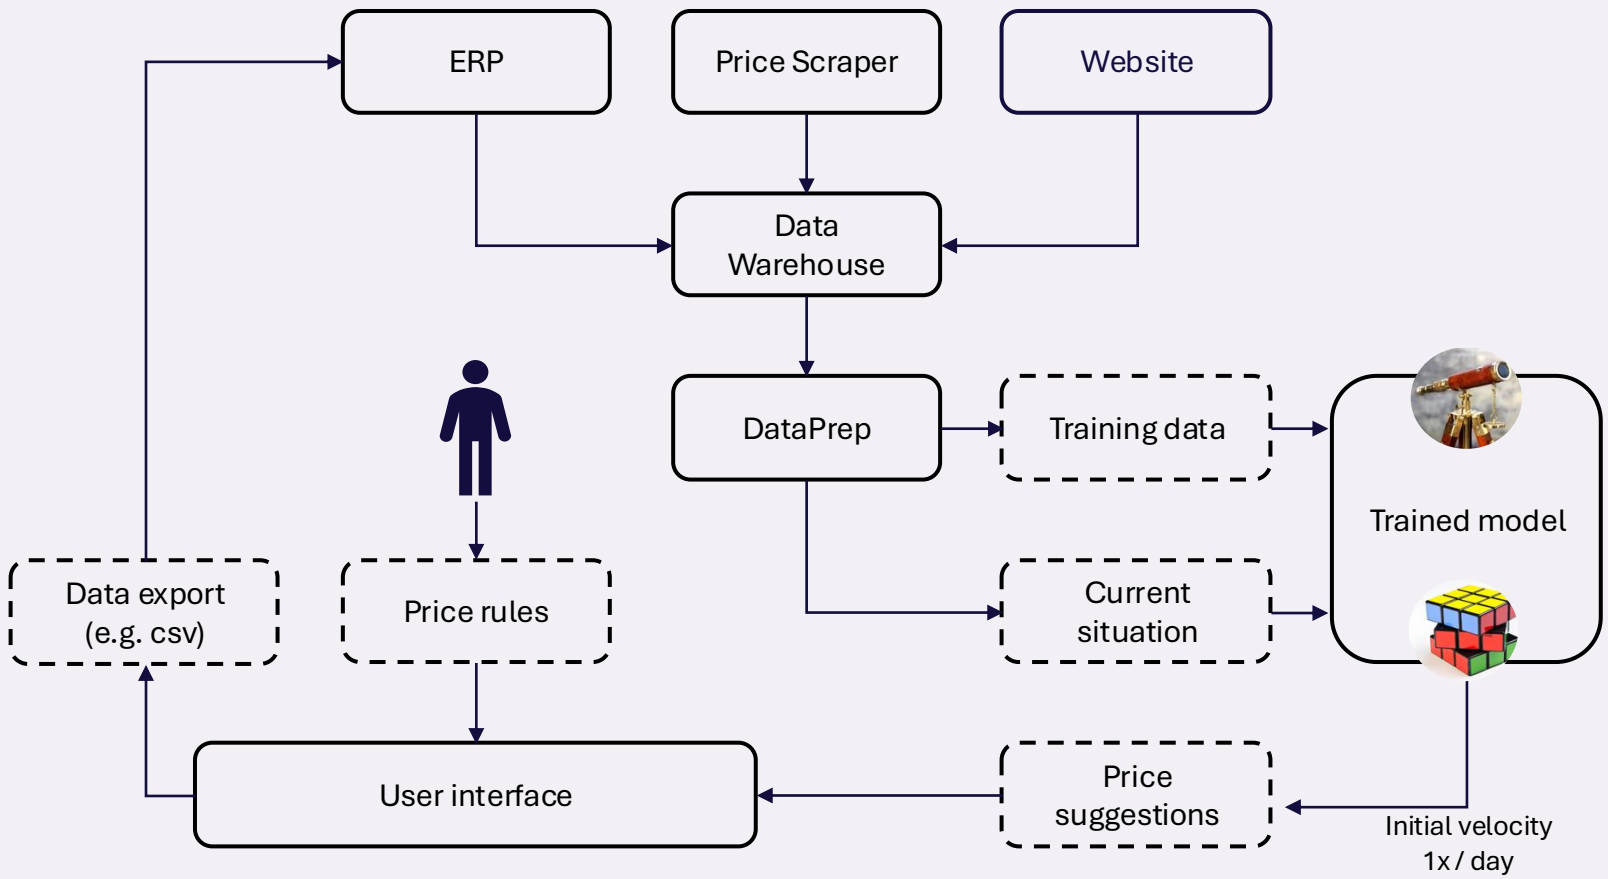

crunch

Deploy

About This Platform

This application leverages **LangGraph multi-agent systems** to perform intelligent sensitivity analysis on optimization models.

- **Planner Agent:** Strategically proposes scenarios
- **Coder Agent:** Translates ideas into code
- **Executor:** Runs modified models
- **Analyzer:** Synthesizes insights
- **Final analysis agent:** Provides a complete report

OpenAI API Configuration

AI-Powered Sensitivity Analysis

Harness the power of multi-agent AI systems to automatically explore and analyze optimization model sensitivities with intelligence and efficiency.

About & Research | Sensitivity Analysis | Results & Visualization | Advanced Settings

Master's Thesis Research

Research Title

crunch

The architect’s responsibility

The convenience of genAI temps us to become passive passengers in our own research.

We must resist this laziness.

Our responsibility is threefold

- **To ourselves**: Refuse to be the person in the Chinese room, do your own thinking
- **To our students**: Raise people who want understand the details
- **To our code**: Create software that is well designed, with code that is well written – especially in the heart of the applications we create. Assume that users will never read the manual.

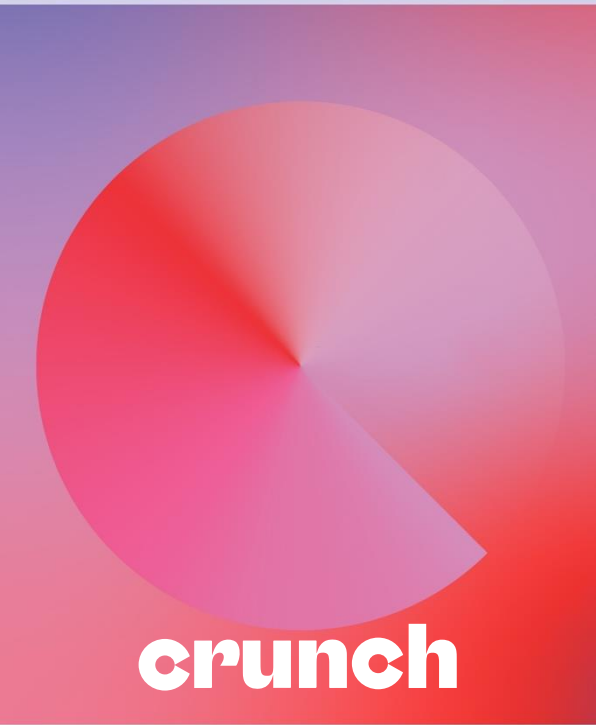

Thank you!

crunch