COMPANION CUBE

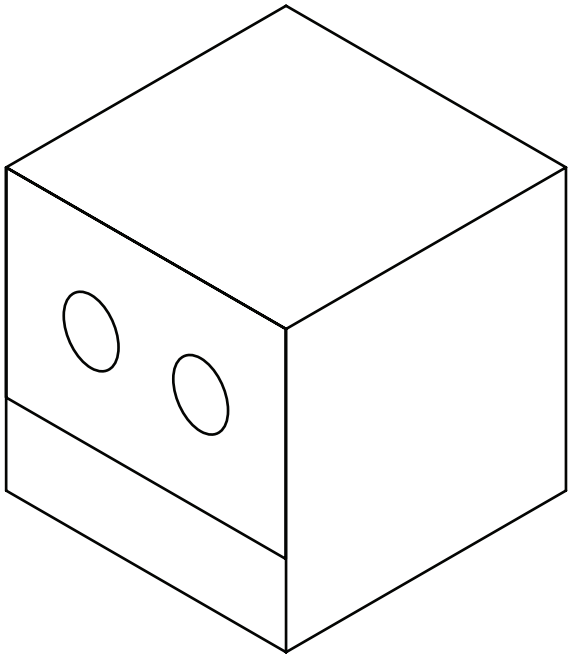

taadum

Optimizing Human Output
Taadum labs

English

WARNING

Cube handling warnings

- DO NOT throw the cube. Thrown, launched, or forcefully rejected cubes may result in head injury.
- NEVER leave the cube on floors, stairs, or other walkways where it may be stepped on.
- DO NOT develop excessive emotional attachment to the cube.
- DO NOT kiss, cradle, or verbally confess feelings to the cube for extended periods.
- NEVER place two cubes in close proximity overnight.
- IF unusual cube behaviour, bonding, or suspected reproduction occurs, separate cubes immediately.
- DO NOT place the cube upside down. Incorrect orientation may result in unintended behaviour.

Español

ADVERTENCIA

Advertencias sobre el manejo del cubo

- NO arrojes el cubo. Los cubos lanzados, impulsados o rechazados con fuerza pueden provocar lesiones en la cabeza.
- NUNCA dejes el cubo en el suelo, escaleras u otras zonas de paso donde pueda ser pisado.
- NO desarrolles apego emocional excesivo hacia el cubo.
- NO beses, sostengas ni confieses sentimientos al cubo durante períodos prolongados.
- NUNCA coloques dos cubos en proximidad cercana durante la noche.
- SI se observa comportamiento inusual del cubo, vinculación o sospecha de reproducción, separa los cubos inmediatamente.
- NO coloques el cubo boca abajo. Una orientación incorrecta puede provocar un comportamiento no deseado.

6667 - Expression guide sheet
290 - Companion Cube V290 (**CC**)
2140 - Communication codex communicator cable (**CCCC**)

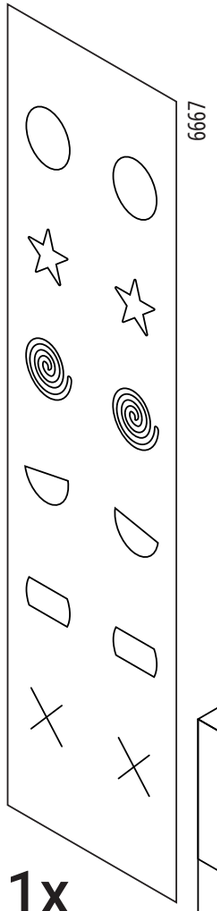

1x

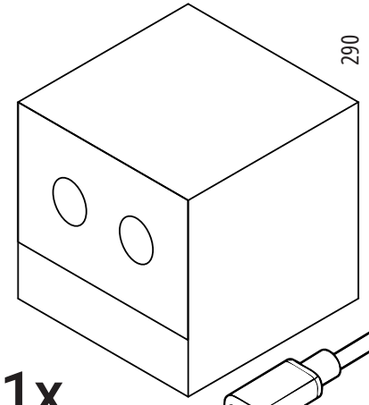

1x

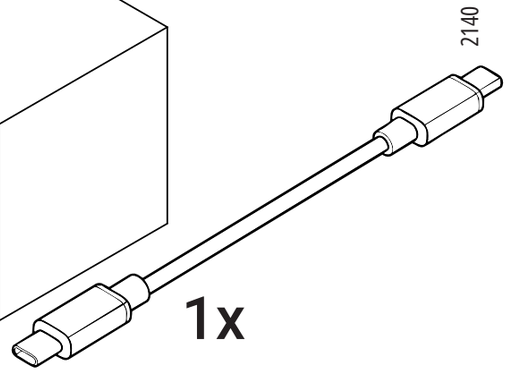

1x

Step 1: Connect **CC** to **PC** using **CCCC**
Step 2: Install software
Step 3: Run ChatGPT

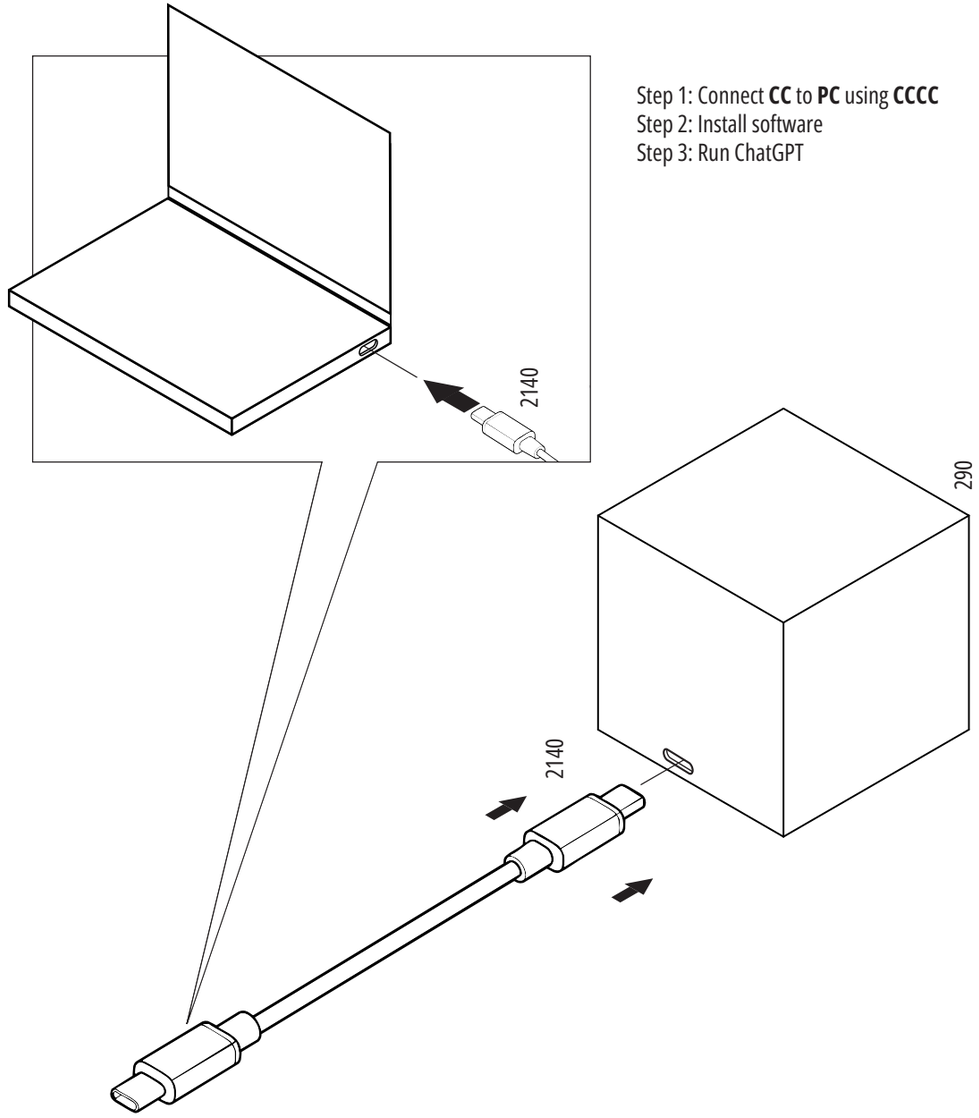

COMPANION CUBE
INDEX

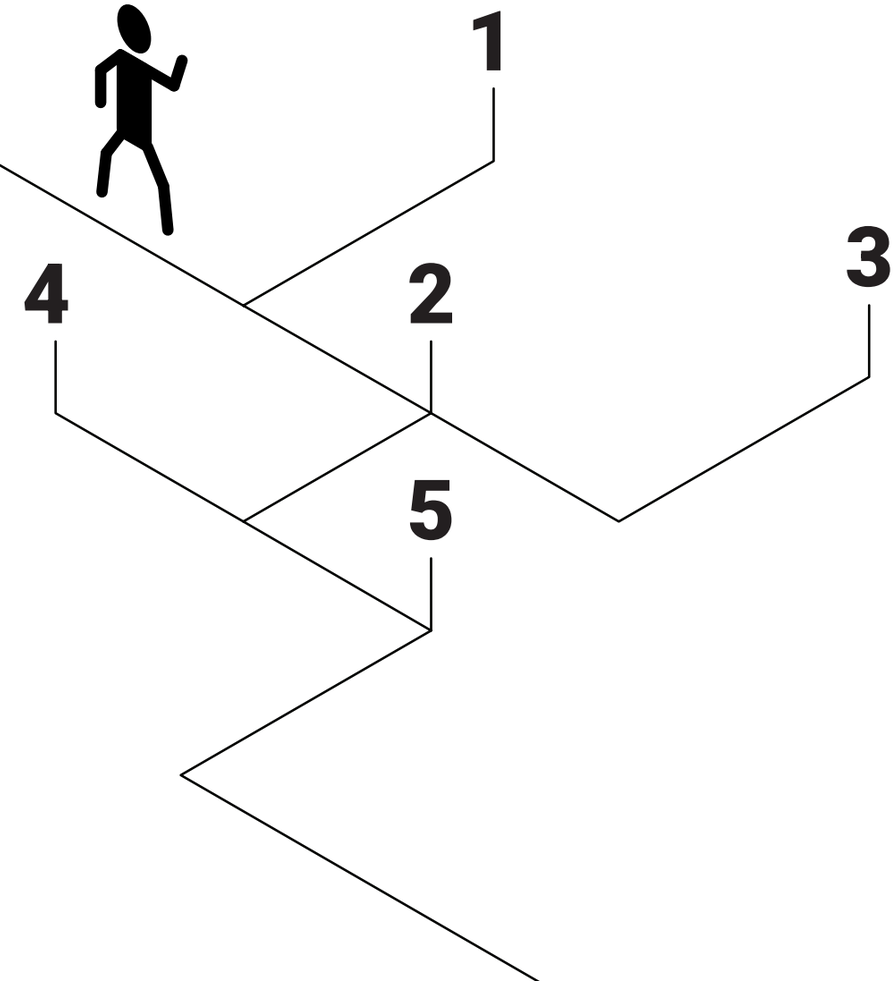

WELCOME TO TAADUM

We established Taadum labs to solve unknowns.
Here at Taadum we pride ourselves on asking the questions that others avoid, and confidently believe there is a decisive solution to anything and everything.

There is knowing and not knowing: 1 and 0 and nothing in between.

Answering questions all day is no easy task. Fortunately we maintain a dedicated and disposable team carrying us closer and closer to...well, we're still figuring that part out.

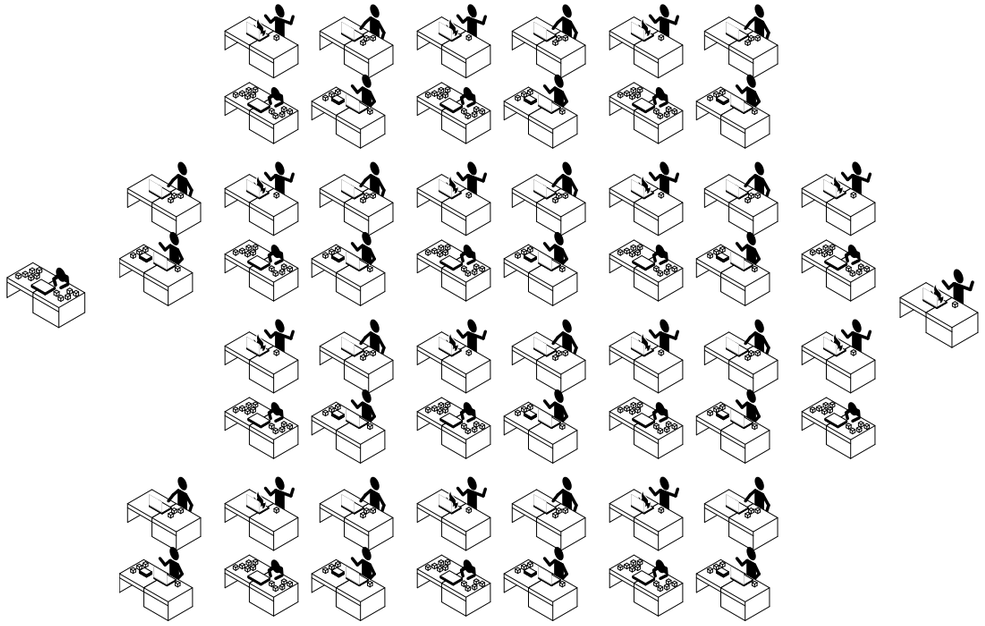

Context 0.1

In 2022 Taadum faced its biggest unknown yet.
On November 30, OpenAI released ChatGPT to the public, making conversational AI accessible to the masses.

Since then, we have dedicated all our valuable time and resources to solving one of the biggest unknowns of our time:

What is AI?

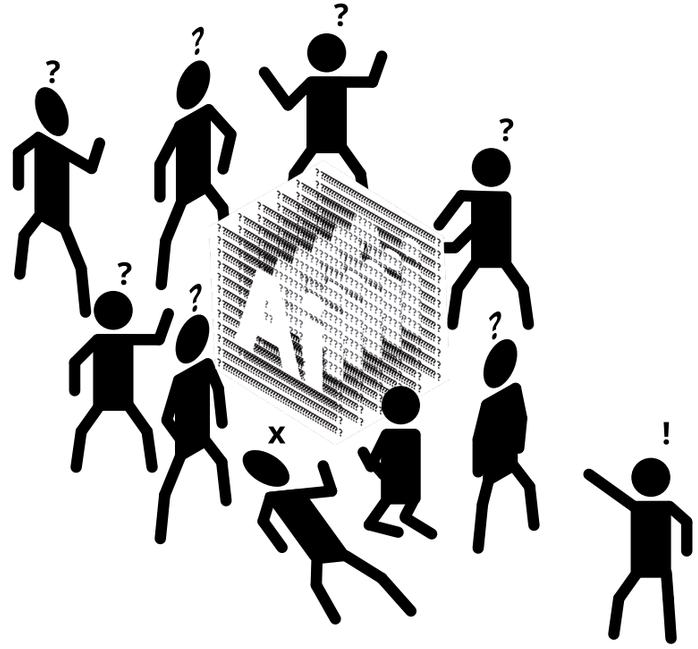

Luckily Dave figured it out.

This is Dave.

Context 0.2

Dave realised we didn't need to understand what AI is,
Instead we needed to focus on how AI works.

(We gave Dave a star)

That way we would reveal all there is to know about AI without needing to think so hard.

So we sent Dave and all those like him to work. Thousands of hours were spent chatting with OpenAI's ChatGPT, only to reveal more unknowns.

- Why is the AI forgetting information?
- Why is it acting like it knows things it doesn't?
- Why does its performance change over time?
- Why does giving it more information sometimes make things worse?
- Why does it always sound confident?

That's a lot of whys Dave.

It quickly became clear that we needed some help keeping track of how the AI was behaving: a device that could make it easier to understand.

(We took away Dave's star)

We needed to build a device that would make AI behaviour legible while warning us when it became unstable.

It took countless failed experiments to build something that could translate ChatGPT's complex and confusing behaviours into understandable expressions.

Expressions are clear visual indicators of AI behaviour.

Finally we landed on the solution and it was painfully obvious:

we gave AI a face.

By giving AI a face, we translated abstract AI behaviour into an emotional language we could intuitively read, helping us recognise when the system was stable, stressed, or beginning to fail.

All that remained was to package it.

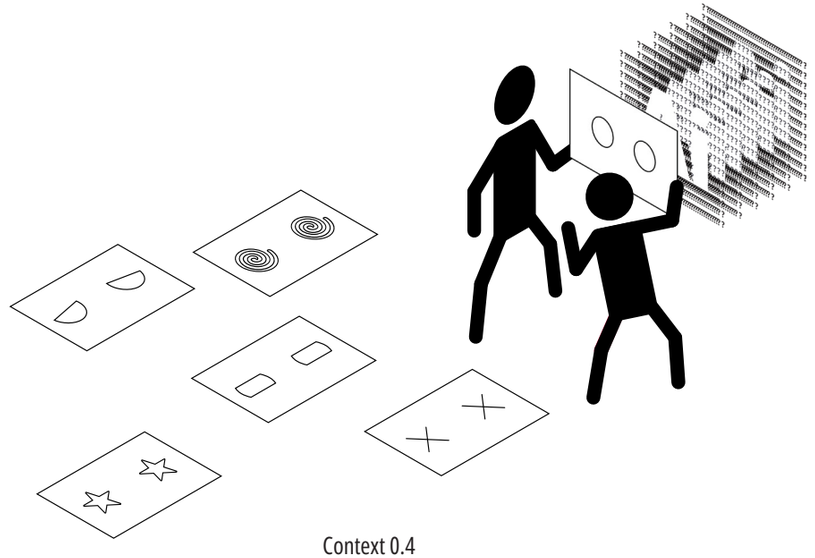

GENESIS

How the Companion Cube came to be.

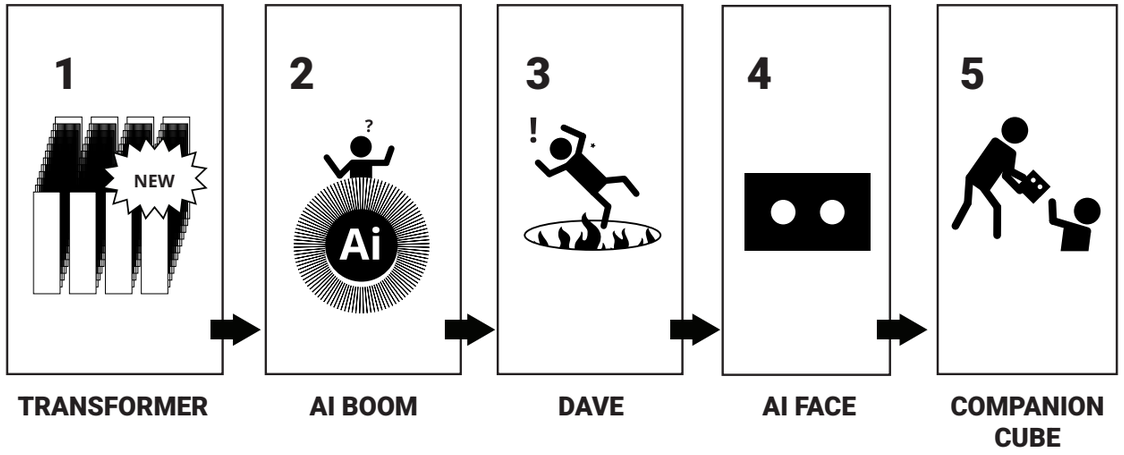

1. **TRANSFORMER**
   In 2017 Google publishes "*Attention is all you need*" introducing the world to their Transformer: a technology which changed Artificial Intelligence capabilities and ultimately helped drive the AI boom that followed.

2. **AI BOOM**
   In 2022 OpenAi released ChatGPT to the public, making AI tools widely available to users. In the years that followed, AI became increasingly applied, transforming how people created, learned and worked.

3. **DAVE**
   In 2023 Taadum received multiple awards on behalf of Dave for his innovative perspective on AI research, marking an early step towards understanding the ground-breaking technology. Dave was later expelled following a series of workplace complications.

4. **AI FACE**
   At the end of 2025 Taadum dedicated all its resources to the AI FACE project, an initiative aimed at solving the hidden vulnerabilities present in current AI tools.

5. **COMPANION CUBE**
   In 2026, the AI FACE project led to the release of the Companion Cube, a tool designed to help users communicate more effectively with AI systems and better understand how to use them.

RESEARCH

The Companion Cube is the product of rigorous research and a concerning number of failed experiments. Hundreds of Taadum researchers dedicated their precious lives and limbs to understand the limitations and vulnerabilities of OpenAI's ChatGPT, grounding the Companion Cube in proven research.

We never expected that granting our top researchers an unlimited fund would lead to the creation of a cube the size of an oversized die, but the youth-demographic validation trials speak for themselves.

The research breaks down into three phases:

1. **Behaviour Observation**
   Through rigorous testing and analysis of ChatGPT, our researchers identified observable behaviours that emerge as the system approaches the boundaries of its designed capabilities.

2. **Literature Absorption**
   Researchers were given access to thousands of published studies and benchmarks exploring the limitations of long-context LLMs, equipping them with the data and statistical patterns needed to inform the trigger system of the Companion Cube.

3. **Product Testing**
   The trigger system and the Companion Cube underwent a comprehensive series of trials, refining them into precise and uncompromisable tools.

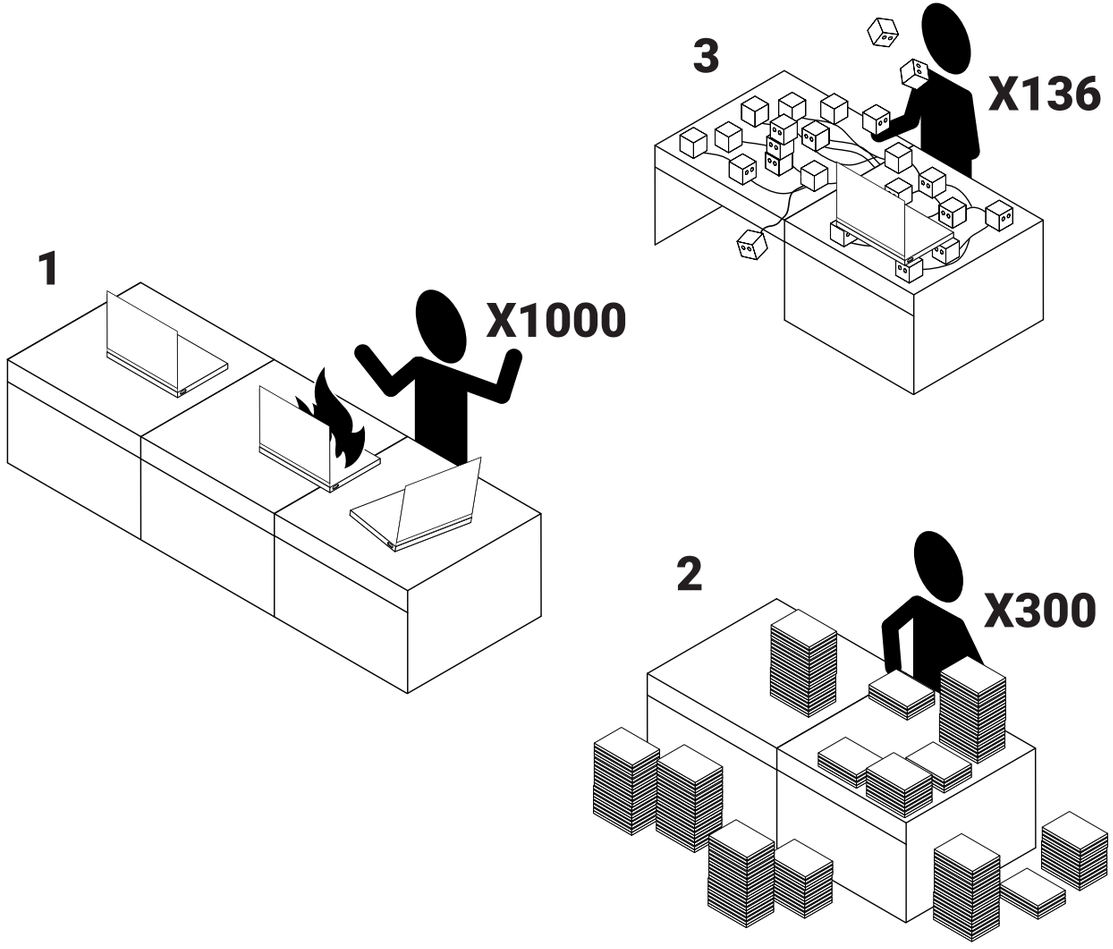

BEHAVIOUR OBSERVATION

Since the release of ChatGPT, researchers at Taadum have spent countless hours recording their observations of the ground-breaking tool.

Through careful prompting and chat curation, they began to identify patterns in its behaviour and speculate on its weaknesses and points of vulnerability.

Researchers observed a clear correlation between conversation length and performance, noting that long and dense conversations lead to a clear decay in performance.

Similarly, they observed that even as newer AI models introduced larger context windows and more advanced memory capabilities, the systems continued to exhibit performance limitations far from their advertised capacities.

As conversations approached those limits, performance became immediately unreliable.

Based on these findings researchers noticed that ChatGPT tends to compress information when summarizing long material and often leaves out small yet critical details in favour of a generalised output.

They also noticed that it does not appear to evaluate information equally.

Instead, the system is influenced by the contextual proximity of material, recognising information that is more recent as more important, while older details become increasingly irrelevant.

LITERATURE ABSORPTION

We weren't the only ones to recognize the vulnerabilities present within this new AI technology.

After consolidating their observations our researchers refereed to hundreds of publications analysing and benchmarking AI systems against their advertised capabilities.

Covering the early machine learning systems to today's long-context LLMs, these papers repeatedly revealed a gap between claimed performance and the limitations found in real-world conditions.

Each study contributed to the growing data which informed our trigger system, allowing us to evaluate AI behaviour with greater precision and enabling the Companion Cube to warn users of the underlying architectural limitations and behavioural instabilities present in Large Language Models.

We absorbed the most from these papers:

- "Attention is all you need" (2017)
- "Know What You Don't Know: Unanswerable Questions for SQuAD" (2018)
- "Lost in the Middle: How Language Models Use Long Contexts" (2023)
- "Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena" (2023)
- "RULER: What's the Real Context Size of Your Long-Context Language Models?" (2024)
- "Leave No Document Behind: Benchmarking Long-Context LLMs with Extended Multi-Doc QA" (2024)
- "AbsenceBench: Language Models Can't Tell What's Missing" (2025)
- "NOLIMA: Long-Context Evaluation Beyond Literal Matching" (2025)

HOW TO USE YOUR CC

Confetti is not assured.

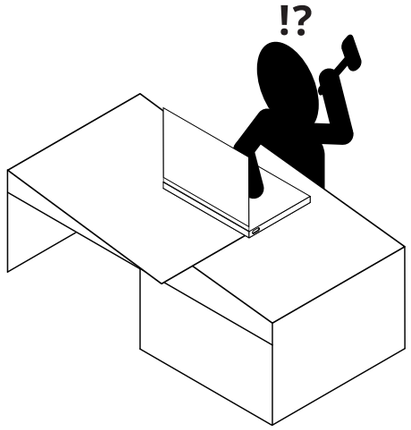

Using the Companion Cube is as easy as understanding human emotions.

Can you tell who is having a better time at work?

GUIDE

Welcome to better habits and practices when using ChatGPT.

The Companion Cube is designed to help you curate and maintain healthier conversations with OpenAI's leading models. By using the Companion Cube, you will remain aware of the vulnerabilities your chats are susceptible to, while being equipped to recognize warnings and take action before they affect the quality of your work.

With time, your improved ability to interpret the CC's displayed expressions will allow you to develop a more efficient and effective work-flow with ChatGPT and safeguard the integrity of your work.

Setup your Companion Cube in nine easy steps:

1. Remove the protective film from the Companion Cube (CC) display.
2. Place your CC on a flat surface near your Personal Computer.
3. Connect your CC to your PC using the provided USB-C cable.
4. Download and install the software to link the CC to your ChatGPT instance. (Provided as a QR code in the sleeve).
5. Position the CC to face you at all times.
6. Use your ChatGPT instance as usual.
7. Observe the expression on the CC.
8. Refer to the Expression Sheet.
9. Adjust your prompting and behaviour as suggested in **Chapter 5: User Intervention.**

USER INTERVENTION

To improve your habits when using ChatGPT, you must learn to recognize the expressions displayed by the Companion Cube and understand the actions required in response to each state.

Intervention is the critical difference between falling victim to and actively addressing the vulnerabilities detected using the Companion Cube.

The Companion Cube updates its expression after each prompting turn, allowing you to react immediately to shifts in conversation health. These warnings vary both in severity and type, ranging from mild to critical, and indicate positional, load, and or recency related risks.

The most effective intervention is dictated entirely by the warning type and severity.

The intervention you choose may have either an immediate or delayed effect on the health of the conversation. Some interventions will correct previous prompting errors, while others will gradually reduce triggers over multiple turns.

For this reason, intervention should not be considered as a single corrective action, but an adjustment in habits, monitoring and adjusting behaviour with every prompting turn.

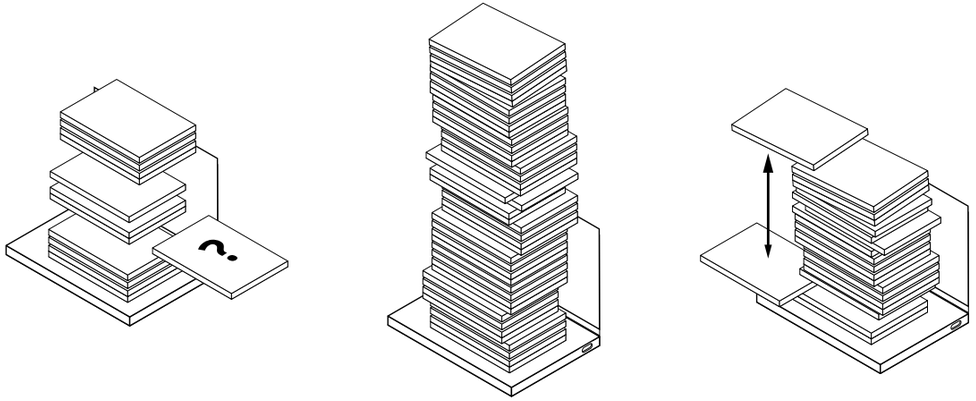

Warning types are dependent on the source of instability, whether it is triggered by a specific prompt or by the gradual evolution of the conversation's context.

The system is built to recognise three distinct warning types, each corresponding to unique trigger condition.

- **Positional Warning**
  Important information has been buried in a large conversation or positioned in the middle of a complex prompt. That information is at risk of being displaced by newer information, making retrieval less reliable.

- **Load Warning**
  The system is overloaded with information and is at risk of failing or crashing. This is triggered by the cumulative length of a conversation, too many uploaded files, or the complexity of a single prompt.

- **Recency Warning**
  The turn distance between the current prompt and the information it is recalling becomes to great, the system is forced to rely on earlier information that has become too far removed from the active exchange.

Visual Display	Expression	Visual Form	Animation	System meaning		Trigger	Warning type	Condition	Suggested action
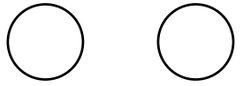	**1**	Neutral open eyes	Blinking, soft hover, scanning	Stable, functional, low-risk state		Idle state, light prompt, normal context load	No warning	The system is comfortable and operating normally	Continue working as normal
	**2**	Expanding Stars	Pulsing outward, held briefly before response	Active reasoning, analysis, preparation before output		Well constructed prompt	No warning	The system will perform effectively	Continue working as normal
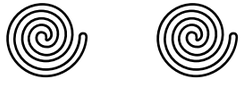	**3**	Rotating Spirals	Slow to fast rotation, increasingly woobly at higher severity	Distress, search instability, positional confusion, retrieval strain		Very long prompts, buried information, references to earlier context	Positional warning	The system may struggle to retrieve or prioritise the right information	Reframe future prompting, reduce context, bring relevant information closer to the end or restate it clearly
	**4**	Sinking / half-shutting eyes	Drooping downward, narrowing gradually	Fatigue, reduced certainty, mild overload, cautious reading state		Reading long material, moderate load, approaching limits	Mild load warning	The system is under some strain but still functioning	Simplify the task, split the request, reduce unnecessary context
	**5**	Squinting / closing slowly	Compressed, heavy, slow closing motion	Strong overload, context saturation, burdened processing		High token count, large uploads, long running conversations, high context occupancy	Medium load / context warning	The system is becoming compressed, less precise, and more vulnerable to drift	Start a new chat, compact the task, reduce files, shorten the prompt
	**6**	X-shaped eyes	Static or abrupt lock into place	Failure, breakdown, unusable state, critical context condition		Critical context overload, disconnection, unrecoverable state	Critical load warning	The system should not be trusted in its current state	Reset, reconnect, restart the task, or begin a fresh conversation

taadum

Optimizing Human Output
Taadum labs