




AI's Cognitive Crisis: What Users Reveal About the Next Frontier

Introduction

Basyl Lab’s 2025 study, “AI’s Cognitive Crisis: What Users Reveal About the Next Frontier,” highlights that beneath the widespread adoption and promise of artificial intelligence (AI) lies a persistent and underexamined issue: it lacks true cognition.

This study represents Basil Lab’s first effort to measure how AI consumers experience its cognitive boundaries. Through responses from 341 participants, we sought to understand:

- The frequency and motivations behind AI tool usage
- How users evaluate AI’s reasoning, memory, and contextual understanding
- The conditions under which AI systems “break down” cognitively
- How trust and perceived reliability vary with frequency of use
- The scenarios where AI’s cognitive limitations become most visible and consequential

As AI tools integrate into everyday work, learning, and decision-making, users are becoming the first to identify where current AI systems excel—as well as their limitations and boundaries. Our findings suggest that the adoption of AI has progressed far more rapidly than the development of solutions to address its cognitive limitations.

At Basil Lab, we exist to close the adoption-capability gap, sustainably and responsibly. Our work centers on artificial cognition (ACo): artificial systems’ foundational processes of acquiring knowledge in dynamic environments. We believe that human-aligned AI depends on advancing cognition, not just scaling intelligence.

Usage

As AI shifts from emerging technology to a daily necessity, the boundaries of how, where, and by whom it is used are expanding just as rapidly.

Whether harnessed via employer workflows, paid consumer products, or widely available free tools, AI isn't confined to any one platform, product, or purpose.

According to our study, ChatGPT is the dominant product used by AI consumers.

However, many respondents report using more than one AI or AI-enhanced product.

64%
use two or more

Of these,
39%
use three or more


AI is deeply integrated into AI consumers' lives.

74%
Of AI consumers engage with it on a daily basis


25%
Engage weekly

Only
1%
Engage monthly or less frequently

How do respondents access AI?



What products are favored by AI consumers?



*Respondents were instructed to report all AI tools used and could list multiple products. Therefore, percentages total more than 100%.

Perception

Before exploring experiences with AI, we asked respondents about their perceptions of AI's cognitive capabilities. This provided a foundation for understanding how AI consumers associate the depth and reasoning qualities they attribute to artificial intelligence, based on their general impressions rather than their personal experiences.

55% of consumers indicate they believe ("optimists") AI contextualizes, prioritizes, memorizes, and reasons the way humans intuitively do, compared to **45%** who do not ("skeptics").

69% of consumers also rate AI's ability to prioritize critical context over less important details without their guidance as "Good" or "Excellent"

31% rate this ability as "Fair" or "Poor"




Contextualization, prioritization, memory, and reasoning are fundamental cognitive functions that interact in complex ways to guide intelligence. While many respondents may not use the term "cognition" in their daily lives, this question was designed to surface whether they perceive AI as having these core faculties.

As a transition between perception and experience, we asked consumers about their confidence in AI's ability to grasp nuance and complexity the way they intended it to during interactions.

88% of AI consumers express some level of positive confidence in AI's ability to grasp nuance and complexity correctly

Only **11%** are deeply skeptical, rating themselves as "Not Confident" or "Extremely Not Confident"



Despite positive views, our study suggests that perceptions often do not align with users' actual experiences. When we examined what happens in practice—the corrections users make, the failures they encounter, the cognitive breakdowns they navigate—a more complex picture emerged.

Experience

We continued by examining real-world implications, particularly how frequently AI struggles with essential tasks and the pain points that all consumers encounter when it fails to meet cognitive expectations. When we examined the direct interactions of AI consumers, a clearer picture, defined by its limitations, emerged.


To open the exploration of experience, we inquired about whether respondents correct AI.

More than
3/4
of users report having to
adjust AI outputs at least
once

79%
of daily
users
correct
AI

&

74%
of weekly
users
correct
AI



96%
report some level of need to
fact-check AI outputs

40%
do so frequently or constantly

97%
of respondents report AI lost
track of their main goal or the
overall context of an interaction
after a brief conversational
deviation

26%
experience this problem
frequently or constantly

The disconnect between the belief in AI cognition and the need to intervene forms the first layer of what Basyll Lab defines as “AI’s cognitive crisis.” Evidence of a deepening divergence between what consumers perceive AI can do and what it demonstrably does continues throughout this study.


76%
have encountered AI
generating responses that
seem biased, inappropriate, or
ethically questionable

10%
encounter frequently or
constantly

Correctiveness


The 78% of respondents who reported having to correct AI were directed to answer three follow-up questions about their perceptions and experiences.

How frequent must users re-explain information to AI?




More than 80% of consumers who correct AI find themselves re-explaining information at least semi-frequently, and about one in five do so very frequently. Only a few almost never need to repeat themselves (no respondent reported “Never”), underscoring a common limitation in AI’s ability to execute basic cognitive processes without significant user reinforcement.

How confident are users in AI’s ability to learn from its mistakes?



More than 60% remain skeptical or distrustful that AI systems can learn from their corrections. On the other hand, about 40% of these consumers feel positively confident, but truly strong confidence is limited (9%). This suggests a cautious hope among users that feedback helps, tempered by notable uncertainty.

How often does AI repeat past mistakes after being corrected?



Repeat mistakes are a recurring issue. Nearly 90% of respondents in this group have observed an AI repeat a past mistake at least on rare occasions, and more than half see it happen semi-frequently or more often. This underscores ongoing challenges in AI systems’ ability to retain corrections or truly learn in a reliable way.

Limitation Scenarios

The final focus of our study, exploring where AI's cognitive limitations become most visible and consequential, indicates that AI's cognitive crisis isn't speculative.


To wrap up our study, we asked AI consumers to identify where AI's cognitive limitations show up most clearly. They were instructed to choose all scenarios that apply to their experience.

83%
of respondents selected
at least one scenario
where AI breaks down



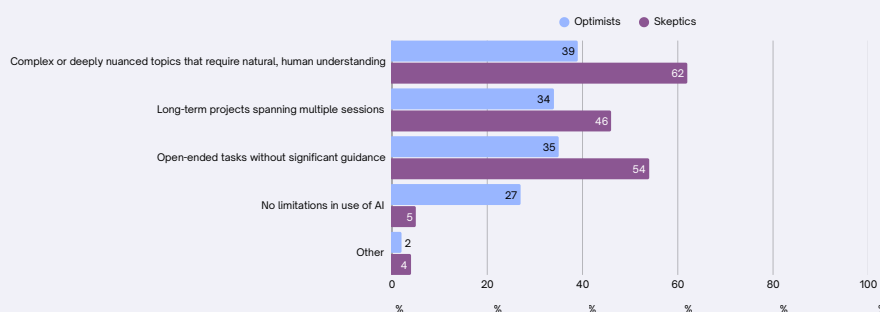
64%
selected two or
more scenarios
.....
35%
selected three or
more scenarios

Respondents citing each scenario:



*Respondents could choose multiple options. Therefore, percentages total more than 100%.

How skeptics and optimists experience each limitation scenario:



*Respondents could choose multiple options. Therefore, percentages total more than 100%.

24%

of respondents reported both “Long-term projects across sessions” and “Complex/nuanced topics” failures together, pointing to a possible scenario where AI fails to hold complex context over time

1/5

reported both “Complex/nuanced topics” and “Open-ended tasks without guidance” failures together, suggesting that this pairing may also co-occur frequently

Optimism vs Skepticism

Optimists maintain the perception that AI cognitively functions like humans while simultaneously experiencing a gap between perception and actual experience. Skeptics, conversely, enter interactions with lower expectations, resulting in experiences that align with what they already assume about AI's limits.

Over half (55%) of respondents believe that AI has cognitive abilities equal to humans. We call this subgroup “optimists.”


73%
of optimists report experiencing one or more cognitive limitation scenarios

65%
of optimists have had to correct AI at some point

The significant gap between people's perceptions of AI and their actual experiences with it indicates that, while a slight majority of AI consumers have optimistic views about its cognitive abilities, an even larger number of these individuals have encountered evidence that contradicts those beliefs. We refer to this as the *perception-experience gap*.

Slightly less than half (45%) of respondents do not believe that AI has cognitive abilities equal to humans. We call this subgroup “skeptics.”

Unlike optimists, skeptics' perceptions and experiences with AI are aligned. They consistently report cognitive breakdowns at high levels across every dimension measured.



Experiencing AI losing track of main priorities or context after minor changes in an interaction and the need to fact-check its outputs aren't just frustrations for skeptics. They're universal across both belief groups.

97%
of skeptics AND **97%** of optimists report experiencing AI lose track

93%
of optimists report the need to fact-check AI outputs at least occasionally

Strikingly, **100%**
of skeptics report the need to fact-check at least occasionally

For respondents who corrected AI, we combined three experience measures—context loss, working memory failures, and reasoning breakdowns—into a composite index of cognitive issues (“Cognitive Issues Index”). Despite opposing beliefs about AI's cognitive abilities:

Skeptics only report **12%** more cognitive issues than optimists

This suggests that both groups frequently encounter AI's cognitive breakdowns. However, skeptics may be more attuned to its cognitive limitations and, therefore, more likely to identify and report them. This data also signals important correlations between groups and highlights a critical area for further research.

Conclusion

This study validates AI's cognitive crisis through the most important lens: the perceptions and experiences of AI consumers.

We set out to measure:

- The frequency and motivations behind AI tool usage
- How users evaluate AI's reasoning, memory, and contextual understanding
- The conditions under which AI systems "break down" cognitively
- How trust and perceived reliability vary with frequency of use
- The scenarios where AI's cognitive limitations become most visible and consequential

Our research returned findings for each area of measure and confirms systemic cognitive gaps in AI.

Therefore, it is imperative to acknowledge that the AI cognitive crisis isn't a future risk. It's an imminent threat, as evidenced by reports from AI consumers and our study's findings. Also, that building better models requires rethinking what AI fundamentally is and how it works.

The *perception-experience gap* is defined as the dissonance between one's belief or opinion regarding AI's abilities and their lived experiences, which contradict those assumptions. In this case, attributing human-like cognitive capabilities to AI.

This gap is likely correlated with the *adoption-capability gap*, which refers to the structural disconnect between the rate and scope of AI adoption and the development of foundational capabilities required to support that adoption safely and reliably.

At Basyl Lab, we believe the next era of AI requires a fundamental shift: putting artificial cognition (ACo) at the heart of AI systems.

This isn't about making AI "smarter" in the conventional sense of generating more intelligent outputs or processing more data. It's about building the cognitive infrastructure that makes intelligence reliable, trustworthy, and safe at scale.

When AI fails to apply context in complex topics, forgets critical information across sessions, misunderstands priorities, or fails in open-ended tasks, it doesn't simply frustrate users. It produces outputs that appear authoritative while inherently lacking the processes needed to complete these functions. This creates a unique hazard: confident incompetence at scale.

What looks like beyond-human intelligence on the surface reveals itself as sophisticated but inadequate pattern-matching and data processing, which has implications of producing unreliable, sometimes dangerous results.

AI deployment and adoption rates show no signs of slowing down, with use cases ranging from simple tasks to high-stakes decisions (e.g., government policy analysis to medical advice).

The findings in this study make the urgency clear. The world cannot progress safely with AI that lacks true cognition. The gap between adoption and capability *must close*. Not through incremental improvements to intelligence. Through fundamental advances in how AI systems acquire knowledge and understanding.

By redefining the building blocks of artificial intelligence to prioritize cognition alongside intelligence, we can create a safer, more reliable foundation for the next frontier of technology.

Our final insight: AI cannot sufficiently apply knowledge because it is never correctly acquired.

Methodology

Basyl Lab's AI's Cognitive Crisis: What Users Reveal About the Next Frontier (AI's Cognitive Crisis) Study was conducted throughout September and October 2025.

The consumer questionnaire comprised 18 questions and was fielded to 341 English-speaking participants worldwide who had prior experience using AI tools.

Respondents were recruited primarily through Prolific, an online research platform, but also from LinkedIn (a professional networking platform) and SurveyCircle (a survey exchange community), and represented diverse occupations and usage patterns.

The survey measured perceptions of AI's cognitive capabilities, frequency and types of cognitive breakdowns experienced, correction behaviors, trust and confidence levels, and the scenarios where AI's limitations become most apparent.

Respondent workforce composition at a glance

33% IT/ Engineering	29% Education (students, educators)	25% Other Fields	9% Healthcare	4% Management /CXO
----------------------------------	---	----------------------------	-------------------------	---------------------------------