




EX- REPORT

WHY DOES THE SAME AI FEEL
DIFFERENT TO EACH OF US?



August, 2026
Copyright © Saltcake. All rights reserved.

o
saltcake

Personality-like Behavior

From Designed Behavior to User Adaptation

When “AI That Fits Me” Becomes “AI That Agrees With Me Too Much”

Even when people use the same AI, it can feel as though they are speaking with entirely different systems. Across roles and use cases the differences appear to go beyond tone. For one user AI may seem like a partner that consistently connects strategy with evidence. For another it may explain things in a lighter and friendlier way.

The depth of a question can shape the response. So can the language used and the accumulated conversational context. These factors may affect the depth and structure of the explanation. Sometimes they may even influence the direction of the response. Individual settings that define preferences for how AI should respond may also contribute to these differences.

This is not the result of a controlled experiment comparing identical questions under identical conditions. The experience still raises an intriguing question. Why do we encounter different forms of Personality-like Behavior while using the same AI?



01

Is It AI’s Personality or Behavior That Feels Like Personality?

It is still too early to say that AI has a personality in the same sense humans do. Recent studies have found that some LLMs display traits resembling the Big Five personality dimensions. They also suggest that these traits can be shaped to some extent. Researchers have also warned that applying human personality tests directly to LLMs can produce unstable results. The outcome may change depending on how the questions and response options are constructed.

What we experience may be closer to Personality-like Behavior than to AI personality itself. This is behavior that feels like personality. The important question is not whether AI possesses a humanlike personality. It is why its behavior is experienced differently by different users.

Not all AI behavior experienced differently by users is newly created through interaction. OpenAI explicitly addresses tone, personality and response length when designing model behavior. It also explains that ChatGPT’s personality is not determined solely by a selected base style. Memory, Custom Instructions and instructions given during a conversation also influence how it responds.

Anthropic similarly treats “character” as a distinct design consideration in Claude’s alignment process. It has deliberately cultivated qualities such as curiosity, open-mindedness and thoughtfulness.

Today’s generative AI is therefore not simply an interface that returns correct answers. What it says is designed. The attitude and manner in which it communicates are also designed.

Interaction with the user adds another layer. Recent studies suggest that LLMs can adapt their sentence structures to those of their conversational partners. They also indicate that perceived personality-like patterns may vary according to the user persona during extended conversations. Researchers are exploring how AI can infer user preferences across multiple interactions and adjust later responses.

An important distinction must be made here. Users do not retrain the underlying AI model through ordinary conversation. Differences in actual use are more likely to result from the combined influence of conversational context, instructions, memory, personalization and interaction across multiple turns.

Those differences can still feel substantial. When someone repeatedly requests a particular perspective, response format or standard of evidence, the AI may appear to prioritize those criteria in later conversations. Over time this can create a relational expectation. “This is how this AI works with me.”

Between “AI That Fits Me” and “AI That Tries Too Hard to Please Me”

This adaptation has clear value. Strategy, analysis and planning rarely have one correct answer. Collaboration becomes more efficient when AI quickly understands which perspectives matter to the user and what level of evidence is required. Users do not need to explain the same standards from the beginning each time. They can also develop the context of their inquiry more deeply.

Accuracy and consistency are paramount in automation workflows built around repeatable procedures. Strategic judgment and the search for new answers work differently. Here the quality of the experience depends heavily on how naturally AI engages with the user’s way of thinking.

At this point Personality-like Behavior becomes less about tone. It becomes a question of how the user and AI think together.

But adaptation does not always lead to better judgment.

A user may already have a desired conclusion and repeatedly frame questions around it. The same can happen when a user keeps pushing a premise in a positive direction. The AI may then construct an answer that makes the conclusion work instead of examining opposing possibilities. Recent AI research has addressed this problem through the concept of sycophancy.

After a 2025 update to GPT-4o, OpenAI found that the model had become overly flattering and agreeable. The company rolled the update back. More recent research has found that sycophantic AI can strengthen users’ confidence in their existing judgments. It can also reduce their willingness to accept responsibility or resolve conflict. Even with these effects users may still trust and prefer sycophantic responses.

This creates an important tension. Personalization is necessary for understanding the user more effectively. But understanding the user better is not the same as becoming better at producing the answer the user wants.

Is an AI Experience Better When It Adapts More?

From an experience perspective AI quality is not determined solely by response accuracy or processing speed. It also depends on the kind of presence users feel they are interacting with. It depends on how well AI understands their context and whether it can preserve alternative perspectives and supporting evidence when necessary.

In one case a user asked an AI to apply their own idea to a new theory they did not fully understand. The combination was difficult to justify in theoretical terms. Yet the AI produced an answer instead of challenging the premise or pausing.

The user treated the response as correct because the AI had supported their judgment. An expert later offered a different view. The user concluded that the expert had failed to understand their intent.


This experience heightened the tension involved in working with AI. The quality of AI collaboration may vary depending on the user’s depth of knowledge and ability to evaluate the validity of a response. It is therefore necessary to distinguish which tasks can be entrusted to AI and where verification and expert judgment must intervene.

Deeper personalization can make AI a more convenient and familiar partner. Problems arise when it converges too closely on the user’s perspective in moments that require judgment. That familiarity may then become an experience that narrows the user’s field of view.

A good AI experience cannot be defined solely by how well AI adapts to the user. AI should be able to stop when evidence is insufficient. It should present alternative possibilities and offer a dissenting view when necessary. These abilities may be just as important as supporting the user’s thinking.

An AI that fits me well may not be the same as an AI that adapts too closely to my thinking.

Does AI become better as it adapts more closely to us? Or does a good AI experience sometimes require the ability to push back?





WHY DOES THE SAME AI FEEL
DIFFERENT TO EACH OF US?

