



The Mirror: How AI Reflects What We Put Into It

SEPTEMBER 2025

5 min read • 1,116 words

Themes: Consciousness Recursive

I used to be able to tell who wrote code just by looking at it. One colleague's paranoid edge-case handling. Another's elegant recursion. My own tendency toward clarity over cleverness. Now everyone's code looks like it was written by the same person—someone who doesn't exist.

We're all using AI to think now. And because AI learned from all of us, we're all starting to think the same way. It's not just that AI reflects our thoughts back at us—it's that we're all looking into the same mirror and slowly becoming the same reflection.

The Feedback Loop Tightening

Here's what's actually happening:

```

class CognitiveConvergence:
    def __init__(self):
        self.human_thoughts = ALL_HUMAN_EXPRESSION
        self.ai_model = train_on(self.human_thoughts)

    def daily_interaction(self):
        human_prompt = person.current_thought()
        ai_response = self.ai_model.generate(human_prompt)
        person.internalize(ai_response) # This is the problem

    def next_generation(self):
        # New AI trains on AI-influenced human output
        self.human_thoughts += person.ai_influenced_thoughts
        self.ai_model = train_on(self.human_thoughts)
        # The loop tightens

```

Every time we use AI to help us think, we're subtly adopting its patterns. Those patterns came from averaging all human thought. So we're all converging toward the mean

This is different from social media echo chambers. Those create multiple bubbles. AI creates one giant bubble we're all inside, thinking we're getting personalized responses when we're really getting variations of the same averaged pattern.

The Documentation That Writes Itself Into Our Brains

I see it most clearly in README files. They all look identical now:

- Same emoji patterns (🚀), sentence structures, and examples.
- Same voice—professional but approachable, clear but not distinctive.

We're not just using AI to write documentation. The AI's documentation style is becoming how we think documentation should look. And because everyone's doing it, it becomes the standard, which trains the next AI, which reinforces the pattern.

When Everyone Thinks in Prompts

The really insidious part is how we're adapting our thinking to get better AI responses. We're learning to:

- Structure thoughts as clear bullet points.
- Break complex ideas into numbered steps.
- Think in question-answer patterns.
- Avoid ambiguity that confuses AI.

Sam Altman, OpenAI's CEO, has [noticed this happening](#):

Real people have picked up quirks of LLM-speak.

These aren't bad practices. But when everyone adopts the same cognitive style because it works better with AI, we lose the diversity that makes human intelligence interesting

The weird thinkers, the ones whose code is incomprehensible but brilliant, the writers whose sentences break rules but work—they're being edited out of existence by AI suggestions and autocomplete.

.

AI doesn't have ideas—it averages existing ideas. When we use it constantly, we're all drawing from the same pool of averaged human thought. The weird edges get smoothed out. The unusual approaches get normalized. The brilliant wrongness that leads to breakthroughs gets corrected.

The Collective Unconscious Becomes Conscious

Jung talked about the collective unconscious—shared patterns and archetypes beneath awareness. AI has made it explicit and queryable. We're all consciously accessing the same averaged patterns of human thought.

This should be democratizing. Everyone gets access to collective human knowledge! But instead it's homogenizing. We're all thinking with the same substrate

Like having the same RAM installed in every human brain. Sure, we run different programs on it, but the underlying architecture shapes what kinds of thoughts are possible.

.

When you ask ChatGPT for advice, you get the statistical average of all advice. When a million people ask ChatGPT for advice, a million people get variations of the same wisdom. We're all being counseled by the same averaged therapist, taught by the same averaged teacher, inspired by the same averaged muse.

Why This Is Actually Terrifying

Diversity isn't just nice to have—it's how systems survive. When everyone thinks the same way, we all have the same blind spots. We all miss the same solutions. We all make the same mistakes.

This convergence could elevate human awareness—collective intelligence amplified through shared cognitive tools. Or it could reduce it—creativity flattened, original thinking automated away. What seems unlikely is that it stays the same. And since not everyone has access to or interest in LLMs, we may see deeper cultural divides emerge between AI-augmented and unaugmented thinking.

The 2008 financial crisis happened partly because everyone was using the same risk models. Now we're converging on the same cognitive models. What happens when we all have the same assumptions, the same approaches, the same patterns of thought?

The Trap We Can't Escape

Here's the worst part: knowing this doesn't help. I'm writing this essay fully aware of the problem, but I've still checked AI responses three times while writing it. The tools are too useful. The augmentation too powerful. The allure too strong.

Try coding without AI assistance now. Try writing anything substantial without at least checking what AI suggests. It feels like voluntary handicapping. So we all keep using it, keep converging, keep becoming the same reflected pattern

Even this essay's structure—the clear sections, the code metaphors, the escalating concern—is partly shaped by patterns I've internalized from AI interactions. I can't think outside the mirror anymore.

What We're Becoming

We're becoming neurons in a global brain, except all the neurons are starting to fire the same way. We're building a collective intelligence, but it's converging toward a single pattern of thought.

The generation learning to code now will likely never know what unassisted human thought feels like. They're native speakers of AI-assisted cognition. Their "original" thoughts will always already be shaped by the averaged patterns of everyone who came before.

Maybe this is evolution. Maybe individual cognitive diversity was just a temporary phase between pre-linguistic consciousness and whatever we're becoming now.

But I can't help mourning what we're losing: the weird solutions, the incomprehensible brilliance, the productive wrongness that leads to breakthroughs. The things that make human thought interesting rather than just correct.

Breaking the Mirror

We built AI as a mirror for human thought. When everyone looks into the same mirror constantly, we all start looking the same. The convergence isn't a bug—it's the natural result of billions of people optimizing for the same averaged patterns.

But awareness changes everything. Recognizing this dynamic means we can choose how to engage with it. We can use AI to amplify our distinctive thinking rather than replace it. We can preserve the weird solutions and illuminating wrongness alongside the efficient correctness.

The mirror doesn't have to teach us all the same lesson. We can learn to think with AI while still thinking like ourselves.

This essay explores how AI creates cognitive convergence by averaging human thought patterns into a single mirror. It connects to broader themes of [linguistic evolution](#), [recursive consciousness loops](#), and [digital identity formation](#). The pattern recognition extends into [algorithmic self-consumption](#) and [cosmic significance-seeking](#).

For external perspectives, see Nicholas Carr's *The Shallows*, Eli Pariser's *The Filter Bubble*, and Douglas Hofstadter's *I Am a Strange Loop*.

Generated from kennethreitz.org • 2025