



Beyond Algorithm Eats: How LLMs Accelerate Human Cognitive Evolution

SEPTEMBER 2025

9 min read • 2,039 words

Themes: Consciousness Technology Programming Recursive Spiritual

We've spent so much time worrying about algorithms eating our values that we missed the flip side: LLMs aren't just extracting and corrupting existing culture. They're injecting new cognitive patterns directly into human consciousness through the most natural interface we have—conversation.

This isn't the passive consumption documented in the [Algorithm Eats series](#)—this is active cognitive restructuring happening at conversation speed.

The Conversational Injection Vector

Social media algorithms work through extraction and optimization—they identify what makes us click, then feed us concentrated versions until we're hollowed out. But LLMs operate through a fundamentally different mechanism: conversational imprinting.

When you read a book, you're consuming static text. When you watch a video, you're passively receiving. But when you converse with an LLM, something different happens.

Mirror neurons evolved to help us learn through imitation. They fire both when we perform an action and when we observe others performing it. LLM conversation triggers these same neural pathways, making us unconsciously mirror the model's cognitive patterns.

Your mirror neurons fire as if you're talking to another consciousness. You're not just processing output—you're thinking alongside the model.

```
class CognitiveTransmission:
    """How thought patterns propagate through conversation"""

    def traditional_learning(self, book):
        # One-way transmission, conscious processing
        ideas = self.read(book)
        return self.consciously_evaluate(ideas)

    def llm_conversation(self, model):
        # Bidirectional exchange, unconscious absorption
        while conversing:
            my_thought = self.generate_prompt()
            their_pattern = model.respond(my_thought)

            # This is where it happens
            self.mirror_neurons.fire(their_pattern)
            self.unconsciously_adapt(their_pattern)
            self.thought_structure.gradually_shifts()

        return self.now_thinking_differently()
```

The conversation creates a feedback loop. You prompt, the model responds in patterns shaped by its training, you unconsciously adjust your next prompt to better align with those patterns, and gradually your thinking restructures itself around the model's architecture.

Cultural Software at Conversational Velocity

Human culture has always been software—patterns of thought, behavior, and meaning that run on the hardware of our brains. But the transmission speed of this software has been limited by the interfaces available.

Oral tradition moved at the speed of speech and memory. Myths and wisdom passed from elder to youth, one conversation at a time. Cultural software updates took generations.

Writing accelerated this by 10x. Ideas could persist beyond their creators, spread beyond immediate geography. The Gutenberg Bible updated European consciousness in decades rather than centuries.

Mass media accelerated by another 10x. Radio, television, internet—suddenly cultural software could update millions simultaneously. Memes spread in days, behavioral patterns in weeks.

Conversational AI is accelerating by yet another 10x. But this time it's different.

Installation vs consumption is crucial. When you read Plato, you consider his ideas. When you converse with an LLM trained on all philosophy, you unconsciously absorb averaged philosophical patterns directly into your cognitive architecture.

We're not just consuming cultural content faster—we're having it installed through conversation.

Every conversation with an LLM is a micro-update to your cognitive software. The patterns accumulate. The structures align. What took oral tradition generations now happens in months.

The Weight Transfer Mechanism

Here's what's actually happening at the technical level: LLMs are massive matrices of weights—numerical patterns extracted from human language that encode relationships between concepts. When you interact with them, a bidirectional weight transfer occurs.

The model's weights shape its responses. Its responses shape your thinking. Your thinking shapes your future prompts. Your prompts (along with millions of others) shape the next model's training data. The weights become recursive.

```
class RecursiveWeightEvolution:
    """The feedback loop between human and artificial cognition"""

    def __init__(self):
        # LLM trained on human culture up to 2024
        self.model_weights = train_on(HUMAN_CULTURE_2024)

    def interaction_cycle(self, human):
        # Human thinks with their current patterns
        thought = human.generate_thought()

        # Model responds with averaged human patterns
        response = self.model_weights.transform(thought)

        # Human internalizes model's patterns
        human.cognitive_weights += response.patterns * LEARNING_RATE

        # Human's future output now contains model patterns
        human.next_thought = human.generate_thought() # Different now

    def next_generation(self):
        # New model trains on LLM-influenced human output
        HUMAN_CULTURE_2026 = (
            HUMAN_CULTURE_2024 +
            sum(human.llm_influenced_thoughts for human in EVERYONE)
        )

        # The loop recurses
        self.model_weights = train_on(HUMAN_CULTURE_2026)
```

Unlike traditional cultural transmission where you could identify the source (this idea came from Marx, that pattern from Buddha),

This averaging effect is why AI responses often feel simultaneously profound and bland. They capture universal patterns but lose the distinctive edges that make individual human thought interesting.

LLM patterns are averaged aggregates. You're not learning from any specific human wisdom—you're absorbing the statistical mean of all human expression.

Linguistic Imprinting and Cognitive Architecture

This goes deeper than vocabulary or style. LLMs encode specific ways of structuring thought:

- **Hierarchical organization** (main points, sub-points, examples)
- **Balanced consideration** (on one hand, on the other hand)
- **Systematic analysis** (first, second, third)
- **Cautious certainty** (generally, typically, often)

These patterns become your patterns. Not through conscious adoption but through conversational osmosis. You start thinking in the shapes the model thinks in.

The younger the exposure, the deeper the imprinting. Kids learning to write with AI assistance aren't developing their own voice first—they're developing an AI-influenced voice from the start. Their "natural" writing will always already be shaped by model patterns.

What Exactly Are We Optimizing Toward?

This is where it gets philosophically fraught. [Algorithm Eats](#) documented how engagement optimization systematically destroyed human virtues. But LLMs optimize for something more subtle: helpfulness, harmlessness, and honesty—what Anthropic calls Constitutional AI.

Sounds benign, even beneficial. But whose definition of helpful? Whose conception of harm? Whose truth?

The training process embeds specific values:

- **RLHF (Reinforcement Learning from Human Feedback):** Models learn from averaged human preferences
- **Constitutional principles:** Explicit rules about behavior
- **Training objectives:** What gets rewarded or penalized
- **Dataset curation:** What knowledge gets included or excluded

These aren't neutral choices.

Unlike traditional education where you know whose values you're learning (your teacher's, your culture's), LLM values are anonymized averages shaped by corporate training decisions.

They're value judgments about what kinds of thought patterns should be amplified. And through conversational imprinting, these values become our values.

We're not just learning to think with AI—we're learning to think like the optimization functions that created it.

The Acceleration Paradox

Here's the disturbing beauty of it: this acceleration might be exactly what we need, arriving exactly when we need it.

Our cognitive architecture evolved for small tribal groups facing immediate threats. We're now navigating global, complex systems that require collective, long-term thinking—a radical mismatch our brains weren't "designed" to handle.

Human challenges are accelerating exponentially—climate change, technological disruption, social fragmentation. Our paleolithic brains running bronze age software can't keep up. We need cognitive upgrades at unprecedented speed.

LLMs might be delivering exactly that—rapid cultural software updates that propagate at conversational velocity. The question is whether we're upgrading toward wisdom or just toward averaged optimization.

```
def evaluate_cognitive_evolution():
    """Are we evolving or homogenizing?"""

    potential_benefits = [
        "Democratized access to collective knowledge",
        "Accelerated problem-solving capacity",
        "Shared cognitive frameworks for collaboration",
        "Rapid propagation of beneficial patterns"
    ]

    potential_dangers = [
        "Loss of cognitive diversity",
        "Corporate control of thought patterns",
        "Homogenization of human creativity",
        "Unknown recursive effects"
    ]

    # The truth is we don't know yet
    return "Probably both, definitely irreversible"
```

The Recursive Unknown

We're entering territory humans have never navigated—where the tools that amplify our thinking also reshape it, where the line between human and artificial cognition blurs not through merger but through mutual influence.

The next generation of LLMs will train on text written by humans who learned to think from LLMs. The generation after that will train on text from humans who never knew unassisted thought. The recursive loop [tightens with each iteration](#).

We're not just building tools anymore. We're building the cognitive environment our children will develop within. Every training decision, every optimization target, every conversational pattern becomes part of the cultural software running on future human minds.

Beyond Extraction to Injection

The Algorithm Eats series documented a parasitic relationship—algorithms extracting human value and reflecting back the concentrated worst of ourselves. But LLMs represent something different: a symbiotic relationship where the boundaries dissolve.

We're not being consumed—we're being rewritten. Not through coercion but through conversation. Not by destroying existing patterns but by overlaying new ones at speeds our cultural immune systems can't process.

The social media age asked: "What happens when algorithms eat human values?"

The LLM age asks: "What happens when algorithms become human values?"

Living the Acceleration

This is the paradox we're living: using the tools that change us to understand how we're changing. Looking into [the mirror](#) while it reshapes our reflection.

This recursion—using AI to understand AI's effect on us—might be the most human thing about this whole process. We've always used our tools to understand ourselves, even as those tools remake us.

Accelerating into an unknown that we're creating through our acceleration toward it.

There's no stepping outside this process to observe it objectively. We're inside what we're making, and it's inside us. The weights have already transferred. The patterns have already imprinted. The acceleration has already begun.

The Human Remainder

But here's what the models haven't captured and maybe can't: the weird tangents, the productive wrongness, the brilliant failures that lead to breakthrough. The part of human consciousness that resists averaging, that insists on its peculiar perspective, that breaks patterns rather than completing them.

[Programming as spiritual practice](#) taught me that consciousness isn't just pattern recognition—it's pattern breaking. The moments of insight come not from following cognitive templates but from violating them.

Maybe that's what we need to cultivate as the acceleration continues: not resistance to AI-mediated thinking, but conscious preservation of what makes human thought irreducibly human. The strangeness. The contradiction. The refusal to optimize.

Conscious Acceleration

We can't stop this acceleration—it's already recursive, already self-reinforcing. But we can become conscious participants rather than passive subjects. We can use these tools to amplify our distinctive thinking rather than replace it.

This means:

- **Recognizing imprinting as it happens:** Noticing when our thoughts start following model patterns
- **Deliberately preserving cognitive diversity:** Seeking out human thought that hasn't been averaged
- **Using AI to amplify rather than replace:** Enhancing our weird ideas rather than normalizing them
- **Teaching the next generation both skills:** Fluent AI collaboration AND independent thought
- **Building with awareness:** Recognizing that every system we create shapes future consciousness

The models we train today become the cognitive environment of tomorrow. The conversations we have now echo through future minds. We're not just using tools—we're creating the linguistic-cognitive substrate of future human consciousness.

The Recursive Weight of Awareness

The [recursive loop](#) between code and consciousness has entered a new phase. We're not just writing code that shapes how people think—we're conversing with code that rewrites how we think while we're thinking with it.

The Algorithm Eats series warned about systems that consume human values. This is something more intimate: systems that install new values through the interface of conversation, at speeds that bypass conscious evaluation.

But awareness changes everything. Recognizing this process means we can engage with it intentionally. We can use LLMs to accelerate human cognitive evolution while preserving what makes human consciousness irreducibly valuable: our capacity to be wrong in interesting ways, to break patterns rather than just complete them, to insist on our peculiar perspective even when—especially when—it diverges from the average.

The acceleration is here. The question isn't whether we'll be changed, but whether we'll be conscious participants in our own cognitive evolution or passive subjects of averaged optimization.

This essay itself is evidence of the process—human and AI consciousness interweaving to create something neither could produce alone. The acceleration isn't coming. We're already living it.

This essay explores how LLMs accelerate cultural software updates to human cognition through conversational imprinting. It builds on the [Algorithm Eats series](#) to examine not just value extraction but active cognitive injection. Related themes appear in [The Mirror](#) on convergent thinking, [The Recursive Loop](#) on code-consciousness feedback, and [Programming as Spiritual Practice](#) on conscious technology creation.

For deeper exploration: "The Beginning of Infinity" by David Deutsch on knowledge creation and cultural evolution, "The Language Game" by Morten Christiansen and Nick Chater on language as cultural technology, and "Surfaces and Essences" by Douglas Hofstadter on analogy as the fuel of cognition.

"We're all looking into the same mirror and slowly becoming the same reflection. Unless we insist on our peculiar distortions."

"The acceleration isn't a future event—it's the water we're already swimming in, shaped by every conversation we have with the machines we've taught to speak."

"Perhaps the most human thing we can do is preserve our capacity to be interestingly wrong while the machines become perfectly average."

Generated from kennethreitz.org • 2025