Why Depth?

Questions Beyond the Age of Scale Out

Author: Jay Sato
Series: Economy of Depth — Papers for the Age When Depth Creates Value
Version: v1.0 / April 2026

Introduction — A Striking Experience in the Summer of 2025

The Structural Limits of Chat-Based AI

In 2024, millions of business professionals began using AI. Open ChatGPT or Gemini in a browser, ask a question, receive a smart answer, close the tab. Sound familiar?

- "I can't find that brilliant response from last week — it's buried somewhere in my chat history"
- "I spend more time copying and pasting AI outputs into documents than actually thinking"
- "I asked the same question I asked last week, and the AI explained everything from scratch"

These are not problems of usage. They are structural limitations of chat-based AI.

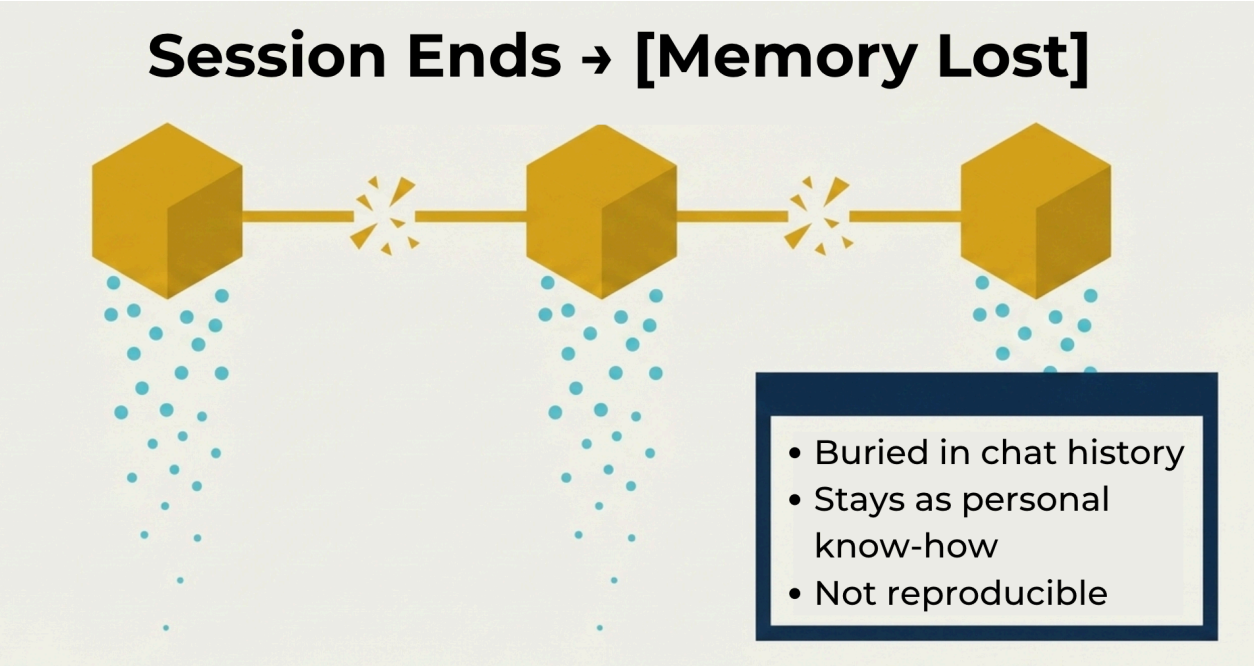

With chat-based AI, context resets with every conversation. AI outputs live on the service provider's servers, buried in chat logs — not in your hands. When a brilliant idea emerges, it is

up to you to decide whether to save it, manually copy it, and organize it somewhere. Forget to save it, and you will never encounter that output again.

The finished proposal stays in your hands. But the reasoning behind it — why you made that decision — does not. The alternatives you considered and rejected disappear. The analysis that informed your judgment disappears. The deliverable survives, but the judgment and insight evaporate.

AI Editors: A Different Kind of Tool

In the summer of 2025, I began working not with chat-based AI, but with an AI editor.

AI editors (Claude Code, Google Antigravity, and others) were originally designed for software engineers to write code. The decisive difference from chat-based AI is that the AI itself can directly read and write files on your computer.

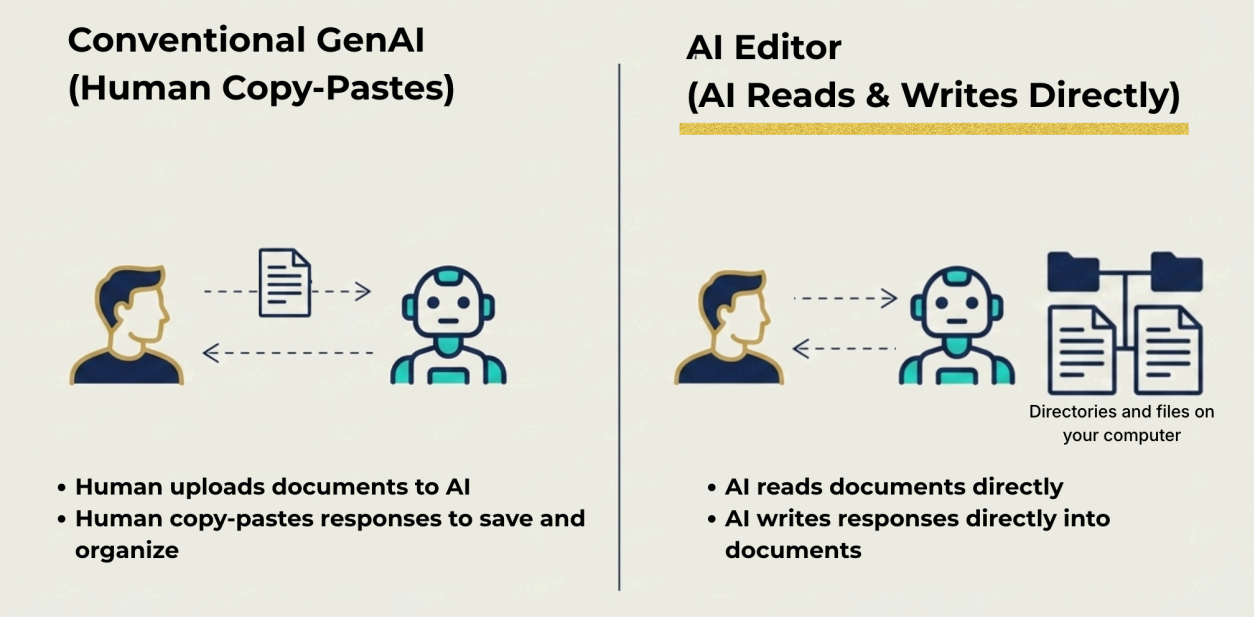

On Saturday, June 28, 2025, I decided to repurpose this engineering tool for business, based on a hypothesis: every digital artifact can now be created with AI. Within the first four days, the hypothesis proved correct. Videos, business documents — all produced with a $100/month subscription and a single laptop.

But something unexpected emerged beyond the original purpose. The potential for human-AI collaboration was enormous, yet the infrastructure for it was entirely absent. With an AI editor, deliverables remain as files on your machine. Why you made a particular decision, the alternatives you considered, the analysis behind your judgment — everything that is structurally lost in chat-based AI accumulates naturally here. Sensing this potential, I began building the collaboration infrastructure.

An important distinction must be drawn here. Tasks that AI can complete autonomously and domains that require human-AI collaboration are fundamentally different. For data entry or routine document generation, AI alone is sufficient. But in domains where there are no right answers — where judgment and articulation create value — strategy, business development,

marketing — AI without human context can only return generic responses. This series addresses the latter domain.

As the collaboration deepened, structured intellectual ideas emerged one after another. The volume grew large enough to require management, so I recorded every single one with blockchain timestamps (OpenTimestamps).

A Qualitative Shift, Discovered in Retrospect

After 25 days of collaboration, I organized the intellectual ideas that had been created. When I tallied the creation density by week, a remarkable pattern emerged.

In the first week: 0.9 ideas per day. Even this is a phenomenon that structurally cannot occur with chat-based AI. The AI reads past files, understands the accumulated context, and presents structured ideas. But with context still shallow, the AI's responses remained "organized but generic."

By the third week, the same AI was generating 28.7 ideas per day. A 32-fold increase. And the quality had fundamentally changed. All nine S-tier (highest-rated) ideas were concentrated in this later phase.

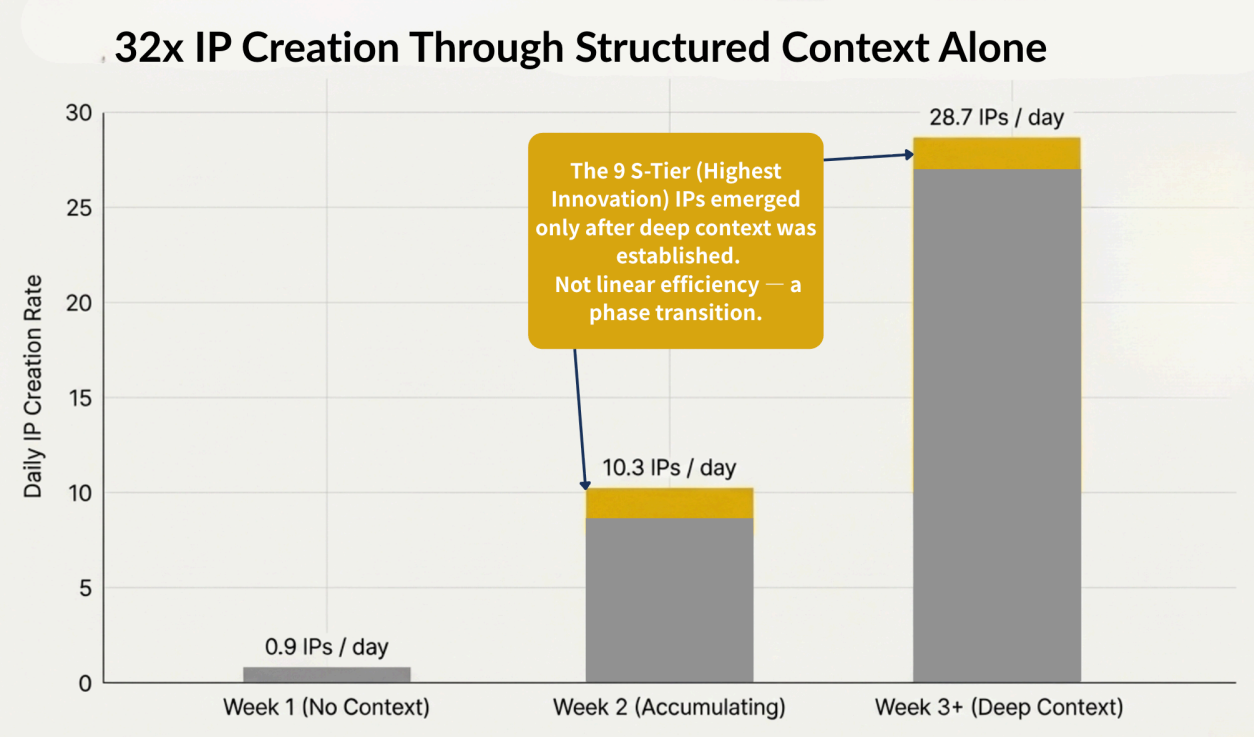

I did not change the model. I did not add computational resources. The only thing that changed was the depth of context accumulated between the AI and me.

25 days. 145 hours. A $100/month subscription and a single laptop.

Ten months later, the accumulation that began with those 25 days had grown into more than 500 intellectual ideas (every one recorded with blockchain proof of existence), an IP portfolio

with 166 items meeting J-KISS investment evaluation criteria — a startup investment framework widely used in Japan — and one patent filed with another in preparation.

Why did the same AI produce such radically different results? And if these results are reproducible, could they change social and economic systems? This question is the starting point of this series.

An Axis Orthogonal to Scale Out — Deep Out

There is a dominant direction in current AI development. Scale Out. Make the model bigger, increase training data, add computational resources. Hundreds of billions of dollars are flowing in this direction.

Scale Out is a proven strategy that forms the foundation of today's AI. No argument there.

What I experienced over 25 days was another dimension — one that layers on top of that foundation.

Deep Out. By continuously accumulating context with the same AI, the quality of output changes nonlinearly. Instead of changing the model, you deepen the context between the model and the human.

If Scale Out is the approach of "expanding the foundation," Deep Out is the approach of "digging deeper on the same foundation."

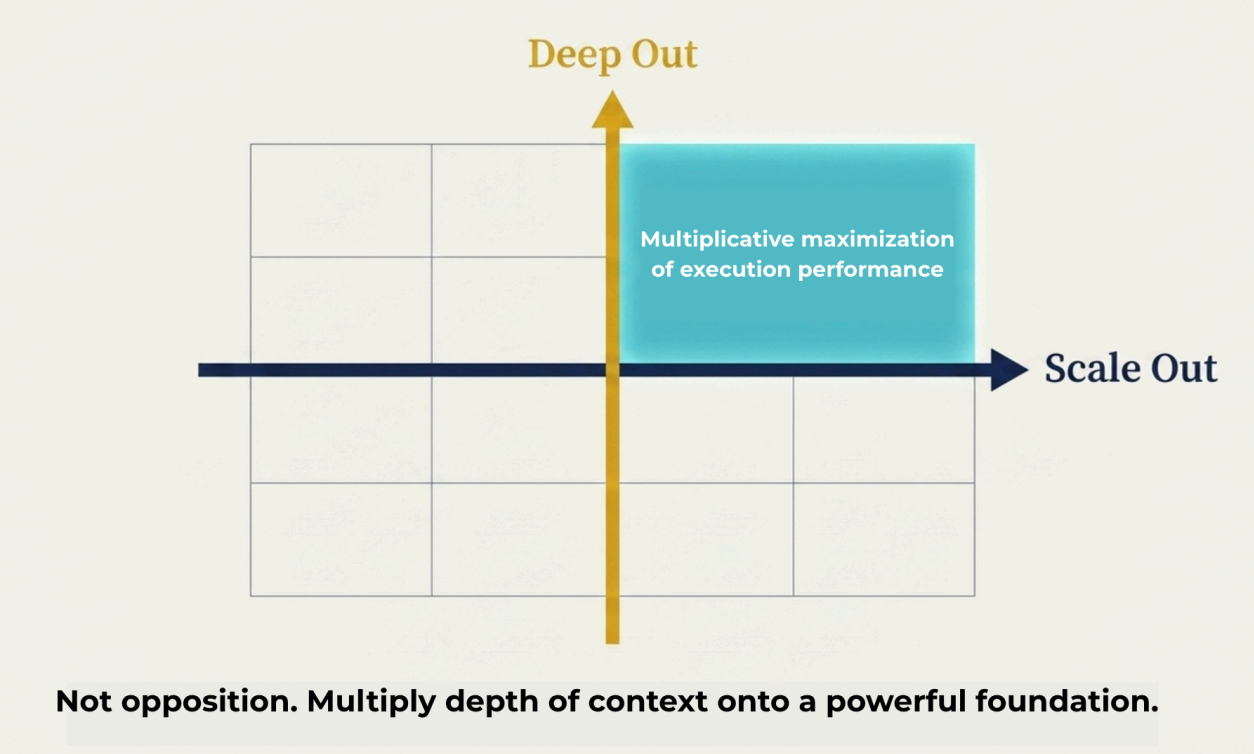

These two are not opposing concepts. They multiply. When the depth of Deep Out was layered onto an excellent AI — the benefit of Scale Out — 393 ideas were born from a small apartment in Tokyo.

A Question — What Is Happening in Your Organization?

Two questions.

The Daily Reset — A Structure That Accumulates Nothing

Are your employees restarting context-free conversations with AI every single day?

As long as they use general-purpose chat-based AI, this is structurally inevitable. No accumulation, no handoff — disposable conversations. The AI returns smart answers, but it does not remember last week's question. Chat histories are buried in individual browsers, and the colleague at the next desk is asking the same question again.

In terms of my experiment, this means remaining stuck at the 0.9-ideas-per-day state — the state before context accumulates. No matter how many shallow conversations you stack, quality does not change.

Intellectual Dissipation — Invisible Assets Evaporating

But there is a more serious problem. Is your organization's intellectual capital disappearing right now, at this very moment?

A talented employee develops a groundbreaking hypothesis through dialogue with AI. But that intellectual struggle is locked inside a personal chat history. Eventually it gets buried under new chats and becomes unfindable. When that employee leaves, the insight vanishes without a trace.

Knowledge silos in the AI age cannot even be documented in a handover note. The insights born from AI dialogue are often "intuitive" — things even the person who had them cannot articulate. It is like leaving a vault open — except nothing is being stolen. It is evaporating.

This is not merely a lag in efficiency improvement. It is intellectual dissipation. Individual productivity may improve, but the organization's intellectual assets are not accumulating. Only by going deep does the landscape change.

CycleGen — A Method for Reproducing Depth

The methodology born from this experiment is called CycleGen.

A roughly 60-minute human-AI collaboration cycle. In the first 30 minutes, the AI works autonomously, producing deliverables as files. In the second 30 minutes, the human reviews the output — confirming, approving, or redirecting. Repeat.

The core principle: "The AI writes the documents. The human makes the judgments."

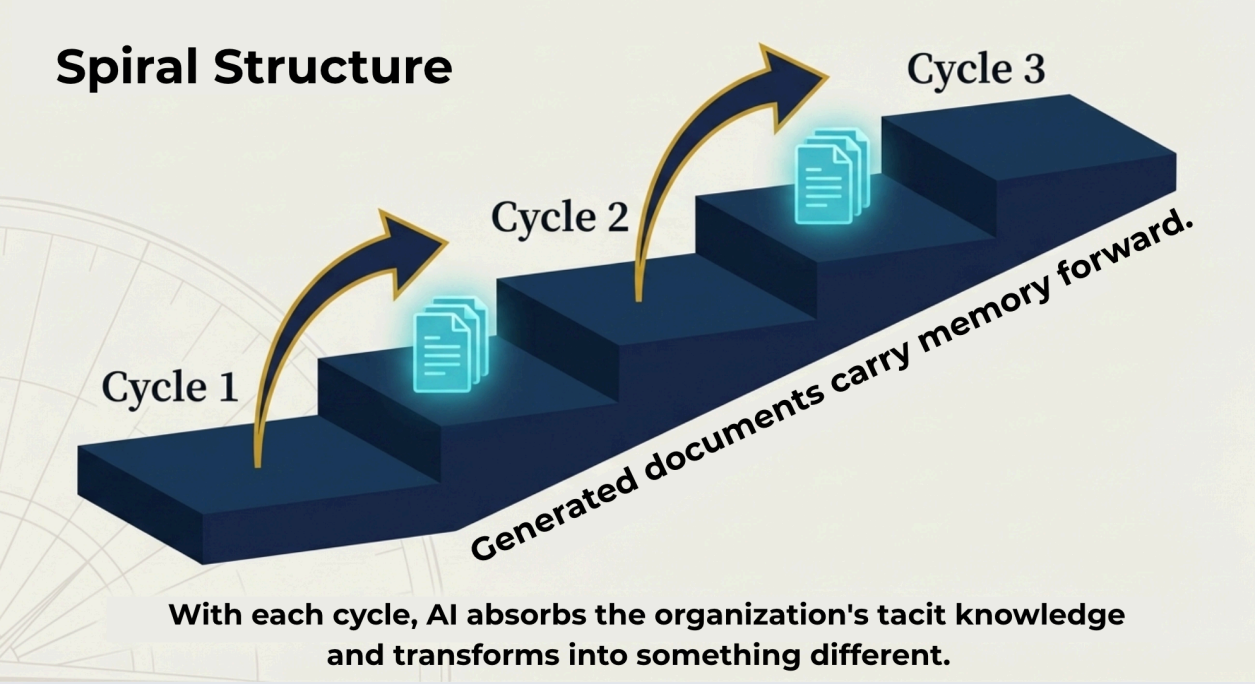

Documents created by the AI accumulate as context for the next cycle. With each cycle, the AI moves closer to "knowing" your organization's decision criteria, past failures, and unarticulated tacit knowledge. The AI at cycle 10 behaves like a different entity from the AI at cycle 1. The model is the same. Only the context is different.

Think about cooking. Many people naturally improve at cooking. Why? Because "tasting" and "eating" function as mandatory checkpoints across the entire cooking process. If there is not enough salt, you know from tasting. You know from your family's reaction. Feedback is built into cooking from the very beginning.

Does your work have an equivalent of "tasting"? Does anyone taste-test your proposal while you are still writing it? Most knowledge work does not have this. Nothing is checked until the final product is complete. AI dialogue is the same — you ask, you get an answer, and it is over. There is no "tasting."

This structural problem is not about individual usage habits. It is about the absence of a mechanism for accumulating depth. CycleGen's design — "a human judges every 60 minutes" — is an attempt to embed "tasting" into knowledge work.

"Why 30 minutes? AI capabilities are advancing rapidly. Why not let it run autonomously for longer?" — a natural question.

For routine tasks that AI can complete independently — code generation, data processing, translation — a longer autonomous run may work, provided the design is sound.

But the domain that Deep Out and CycleGen address is different. Business development, marketing, strategic planning — in domains where there are no right answers, human judgment is non-negotiable. The moment you take AI output at face value, direction drifts. And AI works

fast. Within 30 minutes, there is more than enough output for a human to evaluate and judge. Left unchecked, flawed premises simply compound into flawed conclusions.

Thirty minutes is not "AI's limitation." It is "the optimal interval for human judgment to intervene."

CycleGen does not depend on any specific AI product. The principle is simple: structure context, accumulate it, carry it forward to the next cycle. Follow this pattern with any AI, and you can dig for depth.

Where Things Stand — April 2026

My (Jay's) experience is Deep Out in the domain of business development. I worked across multiple domains simultaneously and produced 393 intellectual properties.

Initially, this was one person's experience — nothing more. But that stage is already over.

A colleague, Shiraki, combined Google Antigravity with CycleGen and reduced 10 hours per week of non-productive time in customer operations to approximately 15 minutes — a 95% reduction. Moreover, Shiraki independently produced patent-level intellectual property. I did not instruct this. By continuously following the CycleGen pattern, Shiraki's own breakthrough emerged in Shiraki's own domain.

Since then, more than five practitioners across different industries have adopted CycleGen and reported productivity improvements.

But let me be honest. The reproduction so far has been in the domain of "efficiency" — faster, easier. Whether the "qualitative transformation" I first experienced — seeing things differently, asking better questions — can be reproduced in others remains under investigation. Reproducing efficiency and reproducing transformation are different problems.

What this series asks is what lies beyond that. When you dig for depth in your own domain, what happens? Do you reach transformation beyond efficiency? Share what you discover.

Something You Can Do Today

Just one thing.

Tomorrow morning, before you start a conversation with AI, try this: tell the AI, in your own words, "Last week, when I asked about this topic, here is what you said. Let's continue from there."

Does that one extra step change the AI's output? If so, what changes?

Context changes output — that experience is the entry point to Deep Out.

About This Series

"Economy of Depth" is a series that takes this experience and methodology and poses questions to readers across different domains.

If the book *Deep Out — A Voyage Log from the Beginning* is a record of the experience, this series is a collection of papers for thinking together about "what happens in your world."

Series Overview

Paper	Title	Question
Paper 01	The Quiet Crisis of Feral AI	Is your organization's AI use building "assets" or producing "dissipation"?
Paper 02	Economy of Depth	Can industrial policy built on Scale Out assumptions function in the Deep Out era?
Paper 03	How to Cultivate the OS of Thinking	The relationship between liberal arts and specialized skills in the AI age
Paper 04	AI Is a Mirror	What does deep collaboration ask of human self-awareness?
Paper 05	Democratizing Depth	Can depth change society without massive capital?
Paper 06	The Possibility of IP Corporation	Can current intellectual property systems adequately handle the AI era?

Each paper can be read independently. Start with the domain that interests you most.

Going Deeper

- Read the experience: *Deep Out — A Voyage Log from the Beginning* (Amazon Kindle / Paperback) — The full record of a 25-day voyage
- Try the method: deepout.org — How to get started with CycleGen
- Start a conversation: signal@deepout.org — Share your experience

Economy of Depth Series — Introductory Paper
Jay Sato / April 2026