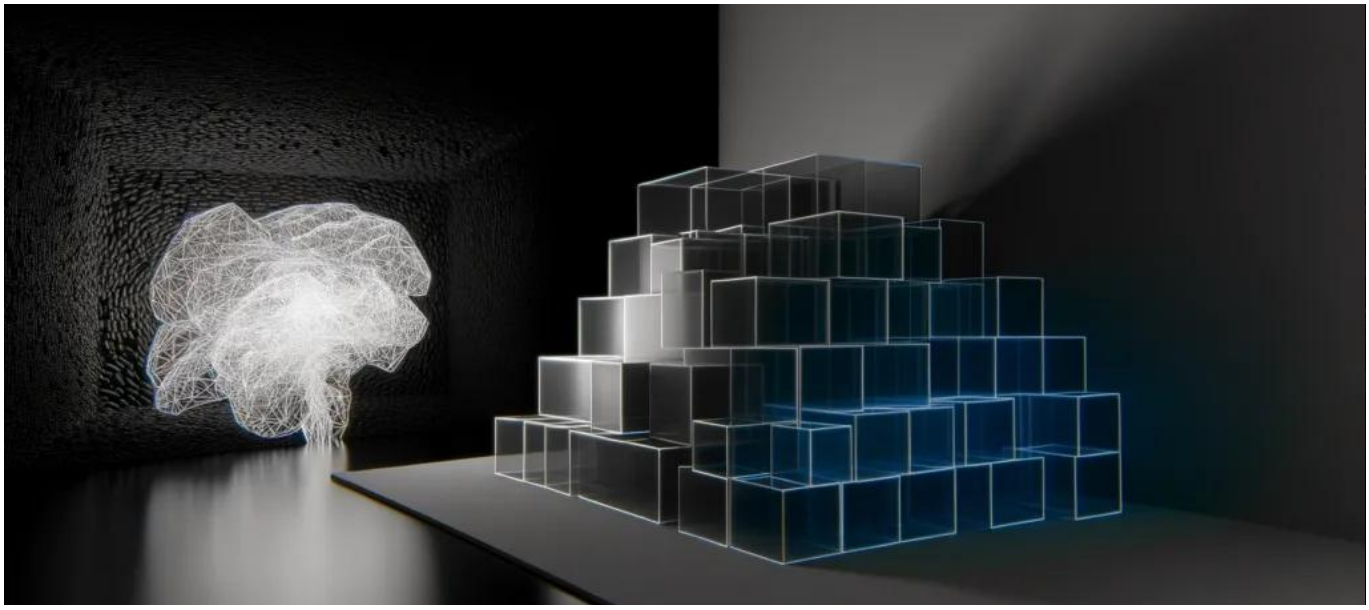


The Death of Stateless AI: Why Google's Titans+MIRAS Architecture Just Made the 'Context Window' Obsolete

Google just killed the context window arms race with a 760M parameter model that outperforms GPT-4. Here's why most AI teams are now building for a paradigm that's already dead.

The Billion-Dollar Distraction We've Been Chasing

For the past three years, the entire AI industry has been locked in an increasingly absurd competition. The metric everyone obsessed over? Context window size. From 4K to 8K to 32K to 128K to Claude's 200K to Gemini's 2M tokens—we watched the numbers climb and convinced ourselves this was progress.

It wasn't. It was a dead end dressed up as innovation.

While OpenAI, Anthropic, and every well-funded lab poured resources into extending the token capacity of fundamentally stateless systems, Google's DeepMind team quietly announced Titans and MIRAS in early December 2024—an architecture that doesn't just extend memory but fundamentally reimagines how AI systems remember, forget, and learn.

The difference is not incremental. It's categorical.

> We've been treating AI memory like a filing cabinet when we should have been treating it like a brain.

Traditional transformer models—every GPT, Claude, Llama, and Gemini you've ever used—are stateless. They process your input, generate a response, and immediately forget everything. The "memory" you experience in conversations is an illusion: previous messages get stuffed into an ever-growing context window, reprocessed from scratch with each interaction. It's computationally wasteful, architecturally brittle, and fundamentally limited by how many tokens you can cram into a single forward pass.

Titans changes this at the architectural level. And if you're building AI systems today without understanding what this means, you're building for yesterday.

What Titans Actually Is (And Why It Matters)

Let me be precise about what Google has built here, because the implications are easy to underestimate.

Titans introduces what the research team calls "neural memory modules"—deep multi-layer perceptrons that function as a learnable, persistent memory system operating alongside the standard attention mechanism. Unlike the key-value caches in traditional transformers (which are essentially growing vector piles), these memory modules can compress, synthesize, and abstract information in ways that scale sublinearly with context length.

The arXiv paper published December 31, 2024 describes three distinct memory integration strategies:

- **Memory as Context (MAC):** Memory modules run in parallel with attention, contributing additional context tokens that the model can attend to
- **Memory as Gate (MAG):** Memory output gates the attention output, allowing learned priors to modulate what information passes through
- **Memory as Layer (MAL):** Memory and attention operate as separate layers in sequence, each processing the full hidden state

The MAC variant has shown the most promise in their experiments, but all three demonstrate something transformers alone cannot achieve: efficient long-range dependency modeling without quadratic attention costs.

Here's the number that should make you pay attention: **Titans models at 760 million parameters outperform GPT-4 on the BABILong benchmark**—a test specifically designed to evaluate long-context understanding and reasoning.

Let me repeat that. A model roughly 1/2000th the size of GPT-4 beats it on tasks requiring memory and long-range reasoning.

This is not a marginal improvement. This is an architectural paradigm demonstrating that brute-force scaling was never the only path forward.

The MIRAS Framework: Where Things Get Interesting

If Titans provides the memory architecture, MIRAS (Memory, Instruction, Reasoning, Action, and State) provides the framework for continuous learning during inference.

As covered by The Decoder, MIRAS enables something that has been a holy grail in AI development: models that can update their core knowledge in real-time as they encounter new information, without requiring retraining.

Traditional models are frozen at training time. Whatever they learned from their training data becomes their permanent worldview. Want to update a model with new information? You need to fine-tune it, retrain it, or rely on increasingly elaborate RAG (Retrieval Augmented Generation) pipelines that bolt external knowledge onto a fundamentally static system.

MIRAS-enabled models can do something different. They can:

1. Encounter information during inference that contradicts or extends their existing knowledge
2. Evaluate whether this information is surprising or important (using learned heuristics)
3. Update their neural memory modules in real-time to incorporate this new knowledge
4. Apply adaptive forgetting to prevent memory saturation

> The question is no longer whether your model can hold 2 million tokens. It's whether it can forget intelligently and update itself without retraining.

This "surprise-driven" update mechanism is particularly elegant. Rather than indiscriminately storing everything (the approach that makes RAG systems bloated and slow), MIRAS-enabled models only update memory when they encounter information that deviates significantly from their expectations. Expected information confirms existing knowledge and requires no storage. Surprising information triggers memory updates.

This mimics how biological memory actually works. You don't remember every meal you've ever eaten, but you remember the one where you got food poisoning. Salience drives encoding.

The Deeper Problem Titans Solves

To understand why this matters, you need to understand the architectural limitations that have plagued transformers since Vaswani et al. introduced them in 2017.

Transformers revolutionized NLP by replacing sequential processing (the domain of RNNs and LSTMs) with parallel attention mechanisms. This enabled massive scaling and unprecedented performance. But it came with a fundamental trade-off: transformers have no inherent recurrence, no persistent state that carries information across processing steps.

The attention mechanism is powerful precisely because it can relate any token to any other token in the sequence. But this power comes at $O(n^2)$ computational cost—quadratic with

sequence length. Double your context window, quadruple your compute. This is why extending context windows is so expensive and why even trillion-dollar companies hit practical limits.

RNNs solved the memory problem differently—they maintain hidden states that persist across time steps, enabling theoretically infinite context. But RNNs suffer from vanishing gradients, sequential processing bottlenecks, and an inability to attend to specific distant tokens. They're fast but imprecise.

Titans attempts to synthesize the best of both worlds.

The neural memory modules provide persistent, learnable state (like RNNs) while the attention mechanism provides precise, parallel token relationships (like transformers). The combination addresses both the compute scaling problem and the memory persistence problem simultaneously.

As Tribe AI's analysis notes, this hybrid approach represents a broader trend in AI architecture: moving beyond pure transformer designs toward systems that incorporate multiple types of memory and processing.

Benchmark Reality Check

Let's look at the actual performance data, because claims in AI research often don't survive contact with rigorous evaluation.

Metric	Titans (760M)	GPT-4	Llama3
BABILong Accuracy	Superior	Baseline	Below GPT-4
Context Capacity	2M+ tokens	128K tokens	128K tokens
Parameter Count	760M	~1.8T (estimated)	70B-405B
Memory Type	Neural MLP	KV Cache	KV Cache

The BABILong benchmark is specifically designed to test long-context understanding—tasks that require models to track and relate information across extended sequences. This is precisely where traditional transformers struggle because attention weights become diluted across too many tokens.

What's remarkable isn't just that Titans wins on this benchmark. It's the efficiency with which it wins. A 760M parameter model shouldn't be competitive with GPT-4 on any benchmark. The fact that it is suggests the architecture, not just scale, is doing meaningful work.

Google's team also demonstrated strong performance on time-series prediction and genomics applications—domains where long-range dependencies and continuous adaptation are critical. This suggests Titans isn't a narrow benchmark optimizer but a genuinely more capable architecture for memory-intensive tasks.

The Hardware Angle You're Not Thinking About

There's a secondary implication here that most coverage has missed: the memory architecture shift has massive hardware implications.

High-bandwidth memory (HBM) revenue is expected to double to $35 billion in 2025, driven almost entirely by AI workloads. Current transformer architectures are memory-bound—the attention mechanism requires shuffling enormous amounts of data between compute units and memory, and KV caches grow linearly with context length.

Titans' neural memory approach changes this equation. By compressing information into learned MLP weights rather than storing raw key-value pairs, the architecture dramatically reduces memory bandwidth requirements. This has cascading effects:

- Lower inference costs per token
- Ability to run capable models on consumer hardware
- Reduced data center energy consumption
- Faster iteration cycles for researchers without hyperscaler resources

If Titans scales as its early results suggest, we may see a democratization of capable AI that the current GPU/HBM oligopoly has prevented. Not because hardware becomes cheaper, but because the algorithms become more efficient.

The Data Scarcity Problem Gets Solved

Here's another angle that deserves attention: continuous learning during inference directly addresses the AI data scarcity crisis.

Industry analysts predict synthetic data will dominate AI training by 2030 because we're running out of high-quality human-generated text to train on. We've essentially scraped the entire internet, and the next generation of models needs more data than exists.

Current solutions involve generating synthetic training data—having AI create the content that future AI will learn from. This works, sort of, but it introduces distribution shift problems and potential quality degradation over generations.

MIRAS-style continuous learning offers a different approach: instead of requiring massive pre-training datasets, models can learn efficiently from smaller amounts of data through continuous adaptation during deployment. Each interaction becomes a potential learning opportunity, not just an inference pass.

This doesn't eliminate the need for pre-training, but it reduces the dependency on ever-larger training corpora. A model that can genuinely learn from its deployment experience needs less baked-in knowledge from training.

> The most data-efficient model isn't the one trained on the most data. It's the one that learns the most from every interaction.

What This Means for Your Architecture Decisions

If you're building AI systems today—whether that's internal tools, customer-facing products, or research infrastructure—Titans and MIRAS should force you to reconsider several fundamental assumptions.

Assumption 1: Bigger Context Windows Solve Memory Problems

This was never true, but the industry acted as if it were. Extending context windows

addresses the symptom (models forgetting relevant information) without addressing the cause (stateless architecture that treats memory as input rather than state).

If you're currently planning infrastructure around models with extended context windows, consider that this entire approach may become obsolete. Building your system to depend on 2M token context windows is building for a paradigm that's already being superseded.

Assumption 2: RAG Is the Solution for Knowledge Currency

Retrieval Augmented Generation has been the industry standard for adding current information to static models. It works, but it's architecturally awkward—you're essentially building a search engine around a model that should be able to remember things itself.

MIRAS-style continuous learning could obsolete significant portions of RAG infrastructure. If models can update their own knowledge during inference, the elaborate retrieval pipelines become unnecessary overhead.

I'm not saying abandon RAG today. The technology is proven and Titans is still early. But if you're making multi-year infrastructure investments, factor in the possibility that your retrieval architecture becomes vestigial.

Assumption 3: Model Scale Is the Primary Lever

The 760M vs. GPT-4 comparison should shatter this assumption definitively. Architectural improvements can deliver capabilities that no amount of scaling would achieve. A better design at smaller scale beats a worse design at massive scale on the tasks that matter.

This has practical implications for build vs. buy decisions. If smaller, more efficient architectures can match or exceed massive models on key capabilities, the economic case for running your own models (rather than paying per-token API costs) becomes much stronger.

Assumption 4: Models Are Static After Training

This is the deepest assumption Titans challenges. We've accepted that models are frozen artifacts—trained once, deployed forever, updated only through expensive retraining

cycles.

MIRAS-enabled models break this pattern fundamentally. They're living systems that evolve through use. This changes everything about how you think about model lifecycle management, versioning, testing, and monitoring.

If your model changes its knowledge during deployment, what does "testing" even mean? How do you ensure quality when the model tomorrow is different from the model today? These are genuinely hard problems that the industry hasn't grappled with yet.

The Open Source Wildcard

Google has announced plans to release Titans code publicly. This is significant.

When transformers were introduced, the architecture was described in a paper but implemented across dozens of different frameworks and libraries. This led to fragmentation and slow adoption. When Meta released Llama weights, it accelerated open-source AI development by years.

If Google releases not just the paper but working code for Titans, adoption could be rapid. Within months, we could see Titans-based models fine-tuned for specific domains, Titans integrated into existing frameworks, and Titans-inspired architectures emerging from other labs.

The alternative—if Google keeps the code internal—would slow adoption significantly but wouldn't prevent it. The paper contains enough detail for determined teams to reimplement. It would just take longer and introduce more variation.

For teams making technology bets, the timing of Google's code release matters enormously. An early release (Q1 2025) could mean Titans-based tooling is available by mid-year. A delayed or partial release could push mainstream adoption into 2026 or beyond.

What This Doesn't Solve

Intellectual honesty requires acknowledging the limitations and open questions.

Scaling Behavior Is Unknown

Titans has been demonstrated at 760M parameters. We don't know how the architecture behaves at GPT-4 scale (1.8T+ parameters) or whether the efficiency advantages persist. It's possible that at massive scale, the benefits diminish or new pathologies emerge.

Google presumably has internal experiments at larger scales they haven't published. Until those results are public, scaling claims remain speculative.

Continuous Learning Introduces New Failure Modes

A model that updates its knowledge during inference can also learn wrong things. Adversarial inputs could potentially corrupt model knowledge in ways that persist across interactions. The "surprise-driven" update mechanism presumably has safeguards, but the attack surface for continuous learning systems is larger than for static models.

The safety and alignment implications are profound and underexplored. If your model can learn from users, malicious users become a vector for model corruption that doesn't exist in static systems.

Memory Consolidation Remains Partially Understood

The papers describe the memory update mechanism in mathematical terms, but the emergent behavior of neural memory modules over extended deployment is unclear. How do these memories interact? Can they interfere with each other? Is there catastrophic forgetting over very long timescales?

These questions will only be answered through extensive deployment and study. Early adopters will be contributing to this understanding, for better or worse.

Integration With Existing Infrastructure Is Non-Trivial

Current AI infrastructure—from training pipelines to serving frameworks to monitoring tools—assumes stateless models. Titans requires rethinking much of this infrastructure. Stateful model serving is a different problem than stateless serving. Version control for models that change during deployment is a different problem than for static weights.

None of these are insurmountable, but they represent real engineering effort that the ecosystem hasn't yet invested.

The Historical Parallel Worth Considering

In 2017, when "Attention Is All You Need" introduced transformers, the initial reception was measured. Yes, it was impressive. Yes, it improved on existing methods. But it wasn't immediately obvious that transformers would obsolete essentially all previous NLP architectures within five years.

What made transformers dominant wasn't just their performance on initial benchmarks. It was their scalability, their parallelizability, and the emergent capabilities that appeared at scale. These properties only became apparent over years of iteration and investment.

Titans may or may not follow a similar trajectory. The architectural advantages are clear on paper. The benchmark results are impressive. But whether this translates into industry-wide adoption depends on factors we can't yet observe: scaling behavior, ecosystem support, engineering challenges at production scale, and competitive responses from other labs.

The prudent position is neither dismissal nor hype. It's watchful preparation—understanding the implications deeply enough to move quickly if adoption accelerates, while not overcommitting to an architecture that might not mature as expected.

Practical Steps for Technical Leaders

Given the uncertainty, here's what I'd recommend for technical leaders navigating this shift:

Short Term (Next 6 Months)

- Build abstraction layers between your application logic and model implementations, making it easier to swap architectures later
- Invest in understanding the Titans papers deeply enough to evaluate future implementations
- Audit your current systems for assumptions about model statefulness
- Begin planning for stateful model serving infrastructure, even if you're not ready to deploy

Medium Term (6-18 Months)

- Experiment with early Titans implementations as they become available
- Develop internal expertise on continuous learning safety and monitoring
- Evaluate which of your current RAG investments would be affected by continuous learning capabilities
- Build relationships with teams actively working on memory-enabled architectures

Long Term (18+ Months)

- Be prepared to migrate production systems to memory-enabled architectures if they prove stable
- Develop new testing and monitoring paradigms for stateful models
- Consider the second-order effects on your competitive position if efficient, memory-enabled models become widely available

The Bigger Picture

Titans and MIRAS represent more than a technical improvement. They represent a potential shift in what AI systems fundamentally are.

Current AI models are sophisticated pattern matchers operating on fixed knowledge. They're tools—powerful tools, but tools nonetheless. You use them, they produce output, and they remain unchanged.

Stateful, continuously learning models are something different. They're systems that

accumulate knowledge and adapt over time. They're closer to employees than tools—entities that learn from experience and bring that learning to future tasks.

This isn't AGI. It's not sentience. But it's a meaningfully different category of thing than what we've been building. The engineering, operational, and organizational patterns that work for tools may not work for learning systems.

The teams that recognize this shift early and adapt their thinking accordingly will have significant advantages. The teams that try to treat stateful systems like stateless systems—or that ignore the transition entirely—will find themselves building increasingly obsolete infrastructure.

We've spent a decade optimizing for the wrong metric. Context window size was a proxy for memory, and we mistook the proxy for the thing itself. Titans shows that genuine memory—learned, compressed, adaptive—is achievable through architecture, not brute force.

The stateless era is ending. The question is whether you're ready for what comes next.

The architecture that wins won't be the one that holds the most tokens—it will be the one that learns the most from every interaction while knowing exactly what to forget.



WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting