



WEWOKU.COM

The Constitutional Memory Problem

"AI governance isn't a regulatory burden—it's the infrastructure opportunity of the decade. After three years developing what I call "Constitutional Memory," I've written about why AI personalization requires governance infrastructure, not just policy. The \$150B question: How do we enable AI hyper-personalization without surveillance? Full article below. Would love to hear from others working on AI governance, infrastructure, or the intersection of tech and public policy." – Greg Malpass Founder & CEO, Constitutional Memory

Why AI Governance Needs Infrastructure, Not Just Policy

As AI systems evolve from ephemeral tools into persistent decision-making entities, we face a governance challenge that neither technology companies nor policymakers are adequately addressing: How do AI systems gather relevant data, remember, update, analyse, advise, AND remain bound by normative commitments over time?

This isn't a theoretical question. It's the infrastructure challenge that will define whether AI serves humanity or surveils it.

The Problem We're Not Discussing

Current AI governance debates focus heavily on immediate concerns: bias in training data, transparency in decision-making, safety in deployment. These matter. But they miss a deeper structural issue.

When an AI system "learns" deeply about you across multiple interactions—understanding your personality, experience, preferences, anticipating your needs, shaping your choices—it develops what amounts to institutional memory. Yet unlike human institutions, which have constitutions, oversight bodies, and amendment processes, AI memory operates in a governance vacuum.

Consider:

- Who decides what an AI system is allowed to remember about you?
- How do those rules adapt as AI capabilities evolve?
- What prevents "constitutional drift" as systems scale and adapt?
- Who has authority to contest or amend these commitments?

These aren't features to be added later. They're infrastructure requirements.

Why This Matters Now

The European Union's AI Act, the US Executive Orders on AI, and emerging regulations globally all recognize that AI governance requires more than voluntary compliance. But regulation without infrastructure is wishful thinking.

We regulate banks through capital requirements and reserve systems—infrastructure that makes oversight possible. We regulate food safety through supply chain traceability—infrastructure that enables accountability. Yet we're attempting to regulate AI personalization without building the governance infrastructure that would make such regulation enforceable.

The result? A dangerous paradox:

- **78% of users want AI personalization** (Accenture, 2024)
- **82% fear how their data is used** (Pew Research, 2024)
- **\$150 billion market opportunity** sits in the gap between these numbers

This paradox exists because current AI systems require users to surrender data control/chat history to gain personalization benefits. That's not a feature of AI technology—it's a failure of AI infrastructure.

Constitutional Memory as Infrastructure

Through three years of systematic research and collaboration with leading AI systems, we have been developing what I call "constitutional memory"—infrastructure that addresses this governance gap.

The concept draws from constitutional design in human institutions. Just as democratic constitutions embed values, establish accountability mechanisms, and create processes for legitimate amendment, constitutional memory provides the technical and governance infrastructure for AI systems to:

1. **Embed commitments transparently** (what rules govern this AI's memory?)
2. **Version and audit changes** (how have those rules evolved, and why?)
3. **Enable contestability** (who can challenge or amend these commitments?)
4. **Detect drift** (how do we know if systems deviate from their constraints?)
5. **Maintain user sovereignty** (who ultimately controls personal AI memory?)

This isn't just philosophy. It's architecture. And architecture is infrastructure.

The Research-to-Infrastructure Pipeline

We've been proactive in early engagement with institutions like the Ada Lovelace Institute in London, which focuses on ensuring data and AI work for people and society. Their interest in the governance dimensions of "memory-bearing AI systems" validates that this challenge sits at the intersection of technical architecture, policy development, and public interest research.

But research alone won't solve this. We need to build the infrastructure layer that sits between foundation models (GPT, Claude, Co-Pilot, Gemini) and end-user applications—enabling personalization without surveillance, compliance without surveillance capitalism, and accountability without centralized control.

This is what ‘Constitutional Memory’ is currently building: AI governance infrastructure for the era of persistent, long-term, learning systems.

Why Infrastructure, Not Just Policy

Policy without infrastructure fails. Consider:

- GDPR gave users data rights but didn't provide infrastructure to exercise them at scale
- "Explainable AI" requirements exist without systems architecture to make explanations meaningful
- "Right to be forgotten" laws collide with AI training data pipelines that lack deletion mechanisms

Constitutional Memory provides the missing layer:

- **For users:** Sovereign control over how AI systems remember and personalize
- **For enterprises:** Compliance infrastructure for AI governance regulations
- **For developers:** Architectural patterns for building accountable AI systems
- **For policymakers:** Enforceable mechanisms for the regulations they're writing

The Opportunity Ahead

AI governance is often framed as a cost centre—regulatory compliance that slows innovation. This is backwards.

The company that solves the \$150 billion personalization paradox by providing governance infrastructure doesn't capture a market—it creates one. Just as cloud infrastructure (AWS, Azure) enabled the application layer to flourish, AI governance infrastructure will enable the next generation of AI applications to scale responsibly.

This infrastructure layer needs to be built by people who understand:

- **The technical architecture** of AI systems at scale
- **The policy landscape** being shaped by regulators globally
- **The user experience** of AI personalization
- **The commercial realities** of building sustainable infrastructure businesses

It also needs to be built by people who care deeply about the outcome—who see this not as a technical problem to solve but as a societal challenge to address with urgency and integrity.

What Comes Next

Over the coming months, we will be exploring partnerships with:

- **Strategic co-founders** who bring US (and EU) market expertise, policy connections, and infrastructure-building experience, and CARE about governance/ethics
- **University research partners** (eg. MIT, NYU, Stanford, Carnegie Mellon) for academic validation of the governance frameworks
- **Enterprise early adopters** who recognize that AI governance is a competitive advantage, not just a compliance burden

If you're working on AI governance, building AI infrastructure, or investing in the intersection of technology and public policy, I'd welcome the conversation.

The Constitutional Memory problem isn't going away. The question is whether we build the infrastructure to solve it now—when we still have agency over the architecture—or later, when surveillance models have become entrenched and constitutional rights over our AI memory are harder to claim.

I believe we build it now. With the right partners, the right architecture, and the right mission.

Greg Malpass is Founder & CEO of ‘Constitutional Memory’, developing AI governance infrastructure for privacy-first personalization. With 25+ years in international business and infrastructure, an MBA from London Business School/NYU Stern, and deep engagement with both AI systems and policy institutions, he's building the missing layer between foundation models and applications.

Constitutional Memory is currently exploring strategic co-founding partnerships. Full technical specifications and research documentation available to qualified partners under NDA terms. Connect on LinkedIn or reach out at destinyinvestors@btinternet.com

AIGovernance # AIInfrastructure # ResponsibleAI # TechPolicy # AIEthics # DataSovereignty # Innovation # TechForGood