CONSTITUTIONAL MEMORY

The Infrastructure Layer for Human Agency in AI Interactions

Prepared for: Lex Fridman
From: Gregory Malpass MBA (LBS/NYU)
Date: November 2025

THE FUNDAMENTAL PROBLEM

Current AI systems present a false binary choice that undermines human agency:

Option A: Generic AI with no personalization → Limited utility, frustrating interactions
Option B: Platform-controlled personalization → Surrender data sovereignty, enable surveillance

This isn't a feature problem. It's an infrastructure gap. We built the internet without user-controlled identity (leading to password chaos), then retrofitted OAuth. We built the web without encryption (leading to mass surveillance), then retrofitted HTTPS.

We're building AI without user-controlled personalization. History suggests we'll regret this.

THE TECHNICAL SOLUTION

Constitutional Memory = User-controlled AI personality profiles that are:

- **Transparent:** Users see exactly what AI knows about them
- **Editable:** Users control their own data, not platforms
- **Portable:** Works across Claude, ChatGPT, Gemini - not locked to one platform
- **Private:** Sealed vault architecture - platforms never access raw profile data
- **Contextual:** Different profiles for student/professional/parent contexts

Technical Architecture:

```
User Profile Storage (User's Device/Cloud)
         ↓
API Integration Layer (OAuth-style authorization)
         ↓
AI Platform (Claude/ChatGPT/Gemini)
         ↓
Personalized Response (without platform data retention)
```

The AI platform receives contextual profile data per session but never stores it permanently. User maintains sovereignty. Platform gains personalization capability without liability.

VALIDATION RESULTS

Empirical Testing (31 comparative examples): - Generic AI responses vs. Constitutional Memory-enhanced responses - **62% improvement** in response relevance, specificity, and usefulness - Demonstrated across professional advice, technical questions, personal development contexts

Example: - Generic query: “How should I approach my career?”
- Generic response: “Consider your skills, interests, and market demand…”
- Constitutional Memory response: “Given your 25 years in international business development, MBA background, and pivot into AI governance research, focus on positioning yourself as a bridge between technical AI development and institutional policy implementation…”

The difference is transformative. And it scales.

THE REGULATORY WINDOW

EU AI Act (Articles 5 & 52): Mandates transparency and user control in high-risk AI systems
GDPR Article 17: Right to erasure creates tension with AI training models
UK AI Safety Institute: Prioritizing data sovereignty frameworks

Timeline: - **Now → 18 months:** Voluntary adoption phase, first-mover advantage - **18-24 months:** Regulatory mandates likely begin in EU/UK - **24+ months:** Late-movers face compliance retrofitting

First-mover captures the standard. OAuth wasn't first to solve authentication - it was first to solve it *well* and became the standard everyone adopted.

THE MARKET OPPORTUNITY

Immediate Addressable Market: - **Education:** Universities deploying AI tutors without student data exploitation (FERPA compliance) - **Enterprise:** Companies wanting employee AI productivity without liability (GDPR compliance)
- **Professionals:** LinkedIn-scale user base wanting career privacy from AI platforms - **Parents:** Child protection from AI platforms without surrendering family data

Revenue Model: - Education (EDU): £1-5/month per student - Consumer/Professional (PRO): £10-50/month subscription per active user - Enterprise: £15-50/month per employee seat
- Families (SHIELD): £2-3/month per family (child protection) - Licensing: API access fees to AI platforms

Conservative 5-Year Projection: £175M revenue (validated business model)
Realistic Scale: Billions of users globally once infrastructure adoption accelerates

Exit Strategy: Expected acquisition interest from Microsoft/ Anthropic/Google tier seeking compliant personalization infrastructure

---

WHY THIS MATTERS PHILOSOPHICALLY

Your conversations with Yoshua Bengio, Stuart Russell, and Max Tegmark explore AI safety and human agency. Constitutional Memory addresses this at the infrastructure level.

The question isn't whether AI will be personalized - it will. The question is whether humans or platforms control that personalization.

Current trajectory: Platforms control personalization → Users surrender agency → Surveillance capitalism extends into AI age

Constitutional Memory trajectory: Users control personalization → Platforms compete on AI quality, not data extraction → Human agency preserved

This is the authentication layer for AI personalization. Someone will build it. The question is whether it's built with human agency as the foundation or retrofitted later after surveillance models entrench.

THE PARTNERSHIP PROPOSITION

What I've Built: - 900+ pages technical specifications (API architecture, security protocols, UX frameworks) - Validation framework with measurable results (62% improvement) - Constitutional Memory methodology (personality profiling, context management) - Regulatory compliance strategy (EU AI Act, GDPR, FERPA alignment)

What I'm Proposing: - Your team runs commercialization (product, engineering, go-to-market) - I remain minority partner providing technical architecture, research validation, strategic guidance - You leverage your platform, technical credibility, and audience alignment with this mission

Why You: Your podcast explores the philosophical depth of human-AI interaction. Your audience cares about agency, transparency, and technology serving humanity. You understand both the technical implementation and the civilizational stakes.

This needs someone who sees infrastructure, not just another app.

CURRENT STATUS & NEXT STEPS

Development Phase: - Technical architecture: Complete - Validation framework: Complete and empirically tested - Business model: Validated with financial projections - Regulatory positioning: Aligned with emerging frameworks

Funding Status: - Bootstrap development phase (self-funded) - Seeking commercialization partner with platform and technical execution capability - University research partnership in parallel (separate track)

Immediate Next Step: 60-minute conversation to discuss technical architecture, market positioning, partnership structure.

Contact:

Gregory Malpass MBA (LBS/NYU)
25+ years international business (£10B+ deal structuring, 60+ countries)
Constitutional Memory / Destiny-Gram
+44 7850 230692
malpass.greg@gmail.com

The Bottom Line:

OAuth solved identity. HTTPS solved privacy. Constitutional Memory solves personalization.

This will become standard infrastructure. The question is whether it’s built with human agency at the foundation or retrofitted after surveillance models entrench.

Interested?