



Constitutional Valley (Malaga)

AI- Governance Infrastructure

Silicon Valley built centralized AI platforms. Now Constitutional Valley (Malaga) is pioneering the decentralized governance layer that makes AI work for everyone. Malaga will be recognized, globally, as the birthplace of ethical AI infrastructure, where individuals and institutions gain true sovereignty over enhanced personalized AI interactions.

We are building strategic partnerships with:

- Innovation Universities advancing AI personalization and ethics research
- Safety Institutes developing AI governance frameworks
- Technical/Software Partners for validation and scaling
- Investors committed to ethical AI infrastructure



“CONSTITUTIONAL MEMORY”

A Brief Overview

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 full personality profiles - The AI platform receives contextual detailed Personality Profiles / Data / LinkedIn Profiles and relevant Chat History per session via API's but never stores it permanently.

“The user maintains sovereignty, while the platform gains hyper-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 MARKET OPPORTUNITY

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

Revenue Model:

- Education: £1-5/month per student (200M students globally)
- Professional: £10-50/month subscription per active user (900M LinkedIn users)
- Enterprise: £15-50/month per employee seat (50K enterprises >1,000 employees)
- Parents: £2-3/month per family (2BN global families and educational institutions)
- Licensing: API access fees to AI platforms

Conservative 5-Year Projection: £175M revenue (validated conservative business model: 1M students / 50 universities = £15M; 1M professionals = £100M; 200K corporate users = £50M; 500K families £10M – Total 2.7M users)

Scale Projection: - Every LinkedIn professional requiring hyper-personalization and career privacy from AI platforms - Every university student needing tailored mentoring but sovereignty from institutional surveillance - Every enterprise employee whose company allows access to, but demands governance over AI usage - Every parent seeking child protection without platform exploitation

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

“The Realistic Scale is - Billions of users globally once infrastructure adoption accelerates.”

THE INFRASTRUCTURE REALITY

Constitutional Memory represents the next foundational layer in AI architecture - user-controlled personalization infrastructure that will become as ubiquitous as OAuth for authentication or HTTPS for security - infrastructure that will underpin how humanity interacts with artificial intelligence across all contexts and life stages.

The question isn't whether AI will be personalized – it will. The question is whether humans or platforms control that personalization, and whether it's built now with human agency at the foundation or retrofitted after surveillance models entrench.

“Constitutional Memory will become standard infrastructure in the future.”

Greg Malpass

Founder

+44 7850 230692

malpass.greg@gmail.com

Constitutional-Memory ©2025





CONSTITUTIONAL MEMORY

The Universal Identity & Trust Layer for AI

*"Transforming the £150 Billion AI Personalization Market
Through Four Distinct Privacy Architectures"*

Constitutional Memory is to be the 'Universal Identity and Trust layer for AI chatbots and Enterprises', offering a final solution to the 'Memory Liability' posed by AI companies and their Enterprise Clients holding users' personal data and information. AI companies currently 'Harvest Data' and use conversation history which lacks structured, comprehensive user understanding. 'Opting Out' of their data sharing/chat-history features only destroys the personalisation sought by users. Constitutional Memory creates user-controlled detailed encrypted structured personal profiles and context analysis of chat history across multiple AI-platforms for anonymised AI-chats and secure hyper-personalised responses using AI analysis and API Middleware.



THE 'BLACK BOX' SOLUTION

Destiny-Gram's "black box" sends personal profiles and chat history anonymously to AI chatbots via APIs

STUDENTS & PROFESSIONALS
Seek 'Secure Storage' of their Data, and Chat History, and creation of Detailed Profiles to ensure 'Hyper-personalized' Chatbot responses

User



- Full profiles (in a sealed vault)
- Chat History (from Destiny-Gram)
- Encrypted data

AI Chatbots



- Relevant, personalized responses
- No access to user identity or raw data

AI PLATFORMS
Need to comply with growing 'Regulatory Requirements', move away from 'Surveillance Capitalism' and find a solution to their growing 'Memory Liability'

ENTERPRISES
Need 'Governance over their Employee AI usage' for compliance but seek to enable 'AI usage & productivity' without IP or confidential client information leakage liability

BLACK BOX

THE BLACK BOX TECH STACK



AI-Built Infrastructure: We have: 1,000+ pages of enterprise-grade code (for both education/professional private enhanced AI-Personalization model and the separate Enterprise AI-Governance model) generated with Claude Opus 4 - the AI that will use our technology. This isn't theoretical - it's validated by the platform implementing it.

Technical Foundation (Production-Ready):

Frontend: React 18 with TypeScript, Next.js, TailwindCSS

Backend: FastAPI, SQLAlchemy 2.0, PostgreSQL with Row-Level Security

AI Integration: Claude API, OpenAI API with custom middleware

Security: JWT with RSA keys, OAuth2, comprehensive encryption

Infrastructure: Docker, Kubernetes-ready with monitoring stack

Code Quality: Enterprise-grade architecture with Domain-Driven Design, complete security implementation, and comprehensive testing framework. Technical review confirms immediate production readiness.

Summary:

- Complete 4-product architecture with shared core + customization layers
- Security: JWT/RSA, OAuth2, Row-Level Security, encryption
- €500K development value completed in 4 weeks vs 6-10 months
- Technical review: "exceptional enterprise-grade, immediately"