



MICKAI ENGINEERING CORPUS · SOVEREIGN FUTURES

AI Babel

Sovereign tongues, cultural fractures, and the new global mind. Linguistic nuance preserved through user-owned local models.

VOLUME IV · EDITION ONE · MAY 2026

AUTHOR

Micky Irons

Founder and named inventor, Mickai LTD. Cumbria.

MICKAI ENGINEERING CORPUS · SOVEREIGN FUTURES

Mickai LTD · Companies House 17166618 · press@mickai.co.uk

UK IPO public register GB2607309.8 to GB2610422.4 · Trade mark UK00004373277

ABSTRACT

AI Babel

Sovereign tongues, cultural fractures, and the new global mind. Linguistic nuance preserved through user-owned local models.

Cloud translation systems optimise for the languages with the largest training corpora, which means minority and indigenous languages are systematically rounded off. Under a sovereign substrate, the per-language model lives on operator hardware, the corpus is signed by its custodian, and every translation is recorded with cryptographic provenance in the OAR chain. The ebook examines the engineering, the political economy of corpus dominance, and the practical case for AI-mediated diplomacy under signed transcripts.

CONTENTS

Sovereign Futures, Volume IV

Part I Foundational Essay

- The homogenisation problem
- What a signed translation enables
- The minority-language case

Part II Worked Examples

- A treaty drafted under signed transcript
- A reconstructive corpus for a dormant language
- Hybrid art forms with cryptographic provenance

PART I · FOUNDATIONAL ESSAY

AI Babel

The homogenisation problem

A cloud translator optimises for the languages with the largest training corpora. The political stakes are not theoretical. Welsh, Cornish, Manx, Gaelic, isiZulu, Quechua, Aymara, the eight hundred living languages of Papua New Guinea: each gets a shrinking share of vendor attention as the leading model providers consolidate. The substrate question is whether the only way to read a minority language in 2035 is through a model that is being kept on life support by a vendor that no longer finds it commercially worthwhile.

The sovereign architecture inverts the question. The per-language model lives on operator hardware. The training corpus is signed by its custodian. Every translation is recorded with cryptographic provenance in the OAR chain. The vendor cannot starve a language out of existence by withdrawing cloud capacity, because there is no cloud capacity to withdraw.

What a signed translation enables

Diplomacy is the obvious application. Simultaneous translation across a negotiation is currently a vendor leasehold on the transcript. Under sovereign architecture, the transcript is signed at commit by each participant's substrate, and any later disputed phrase can be replayed, verified, and adjudicated offline. Treaty drafting under cryptographic transcript is qualitatively different from treaty drafting under whatever audit log the cloud vendor chose to retain.

Less obvious is the hybrid art form. When a poet writes in a language with a signed corpus, and a translator reads in another, the chain records the lineage. The reader sees what the model did with the original phrasing. The poet's estate sees the same chain. The translator's signature is recorded. Cultural production survives the vendor that distributed it.

The minority-language case

Cumbria has a relevant test case in the form of the Cumbric language, extinct since the twelfth century, partially reconstructed from place names. A sovereign architecture cannot resurrect Cumbric. It can hold the reconstructive corpus, the toponymic record, and the scholarly attributions in a chain that survives the academic department that built them. That is the unglamorous engineering case for sovereign multilingual AI. Not preservation as marketing, preservation as substrate.

PART II · WORKED EXAMPLES AND EXPANSIONS

Practitioner notes against the foundational essay

A treaty drafted under signed transcript

A worked example: a bilateral cyber-defence agreement between the UK and a Five-Eyes partner. The negotiation runs over six weeks. Each session is simultaneously translated under both parties' sovereign substrates. The transcripts are signed at commit by each participant's hardware-bound key. Any later disputed clause can be traced to the moment it entered the draft, the speaker who proposed it, the model that translated it, and the parallel translations in the other party's record. The treaty itself is signed once both records reconcile. Subsequent enforcement runs against the same chain.

A reconstructive corpus for a dormant language

A reconstructive project for, say, Manx, runs across the Isle of Man heritage office, the University of Liverpool linguistics department, and a network of voluntary native speakers. Each contributor signs their contribution against a shared sovereign substrate. The corpus accumulates under cryptographic provenance: who recorded what, when, under whose attribution. A model trained on the corpus carries the corpus signature in the chain. Every output is traceable to source. The language survives the institutional sponsor.

Hybrid art forms with cryptographic provenance

The most interesting application is the hybrid art form. A Welsh-language poet writes a sequence; a Quechua-language translator carries it across; a third reader hears both in their own tongue through a personal sovereign translator. Each step is signed. The chain records the original phrasing, the translator's reading, the third reader's locale. Cultural exchange becomes auditable cultural exchange. The translator's craft is recorded as their own intellectual contribution, not absorbed by a vendor's training set.

AUTHOR

Micky Irons

Author. Micky Irons is the founder and named inventor of Mickai, the Sovereign Intelligence Operating System. Based in Cumbria. UK IPO public register GB2607309.8 to GB2610422.4 plus the four new May 2026 filings on cross-implementation OAR verification, pluggable post-quantum signing, federated voice cloning, and audit-by-default sovereign CLI command trace. Companies House 17166618. Trade mark UK00004373277 (classes 9 and 42). The agentic marketing runtime documented at mickai.co.uk/articles/amt-crunchbase-40k-to-500-in-seven-days moved the founder Crunchbase profile from approximately 40,000 to 500 in seven days.

COLOPHON

Mickai Engineering Corpus, Sovereign Futures, Volume IV

Mickai LTD. Companies House 17166618. 20 Wenlock Road, London, N1 7GU. press@mickai.co.uk. mickai.co.uk. UK IPO public register GB2607309.8 to GB2610422.4. Trade mark UK00004373277 (classes 9 and 42). Brand voice audited under the Mickai AMT preflight gate. Sovereign Futures Edition One, May 2026.