

Taonga in the Latent Space

A community-scale world model for Indigenous data sovereignty

John G. Stroh · My Digital Sovereignty Ltd. · 13 slides with presenter notes

The bet against the trench

- Nov 2025: Yann LeCun leaves Meta to found AMI Labs — seed round ~\$1.03B at a \$3.5B valuation, reported as Europe's largest
- “The AI industry is completely LLM-pilled.” “Everybody is ... digging the same trench.”
- His charge: LLMs “can't truly reason or plan, because they lack a model of the world. They can't predict the consequences of their actions.”
- “...they're not going to be built on LLMs.”

What a world model is

- Perception → world model → cost → actor (LeCun, 2022)
- World model = “a kind of ‘simulator’ of the relevant aspects of [the] world” — predicts future states as a function of actions
- Mode-1: reactive, no model consulted. Mode-2: “reasoning and planning ... akin to model-predictive control”
- A text-only LLM is stuck in Mode-1 — no model to run a consequence through

JEPA — predict meaning, not pixels

- Joint Embedding Predictive Architecture: predict future representations in latent space, not raw pixels or tokens
- “It is not generative.” “We advocate against the use of generative architectures.”
- Learn to ignore what isn't worth predicting; predict what matters
- V-JEPA 2 (2025): a world model doing zero-shot robot planning — proof it ships

The world worth modelling first is not the planet

- Nothing in LeCun's definition requires the world to be pixels — only a state, actions, and dynamics
- An institution is a discrete world whose state is already a record — the expensive half (learning state from sensors) doesn't arise
- Community-scale is tractable; a planet-model that dethrones the frontier is not
- Rome was not built in a day — lay one true stone

The Village already governs that world

- Sovereign records = the community's state, signed and replayable — the ground truth a world model needs
- The gate already exists: checks “before action execution ... the process cannot bypass the check” — allow / deny / escalate-to-human
- The Situated Language Model made language cognition sovereign
- Our Sovereign-Record Architecture paper (§15) names the gap: sovereignty over data — “not by itself ... sovereignty over cognition”

The missing middle: a community-state model

- Predict what a proposed action would do to the community's state — consent, taonga, rights, mana, relationships
- Maps onto LeCun: perception=encoder, world model=community-state model, cost=kaitiaki critic stack, actor=the agent
- Turns an institutional agent from Mode-1 to Mode-2 — foresee the effect on the people, then act
- Steered, not retrained; built inside tenant isolation (Adversary A7's refusal property)

Start filling it with what the community sees

- Past text, like adding documents to a situated model: a farmers' market clip, footage of the breached stop banks, a graduation
- LeCun's V-JEPA video models are open and runnable today — they turn video into representations, not pixels
- And live feeds — river gauges, tide and swell, rainfall, water-quality probes: perception that updates continuously and runs the next hours forward
- A community needs only its OWN world — a small corpus it already holds — not a million hours of internet video

Who is kaitiaki of the world model?

- Plural Values named the gap: “no kaitiaki on the artifact”
- LeCun's cost has an intrinsic part that is “immutable (not trainable)” — in the Village that is the constitutional floor (Te Tiriti, tapu, human-in-the-loop)
- The trainable critic sits under the community's signed steering pack — “the benchmark is downstream of the kaitiaki”
- The unit of deliberation is the latent diff: a legible, signed, revertible change to the model of your world

Federation: how it takes root

- Federate, Don't Align — taken one layer down, to world models
- Each Village holds the model of its own world; none holds a model of everyone's
- Cross-community questions federate bilaterally — scope-limited, revocable, provenance intact
- H-JEPA turned on its side: a polycentric mesh, no centre to capture

The one word, used once

- I'll call where this leads the first step toward “sentient AI” — my coinage, not LeCun's
- Not consciousness, not feeling — accountability: a machine that has a world and can be governed in it
- Tractatus 12.6: “Agency cannot be simulated, only respected”
- Model the world the community lives in — never the community's agency

The limits, stated plainly

- Not shipped — proposed. No community-state model runs today; no aspirational building, no aspirational writing
- Closes governance sovereignty, not substrate sovereignty (open weights, foreign hardware remain)
- Federation mechanism is live and tested; no inter-community network yet
- Written from a limited, non-Māori standpoint; appropriate use in te ao Māori is for tangata whenua to determine

A model can answer. An agent can act.

- But only an agent that can imagine what its act will do to a community — and let the community decide — has earned the right to act at all
- The world model is how it imagines. The kaitiaki is who decides.
- The machine models the world; the people remain its authors

Taonga in the Latent Space — John G. Stroh, My Digital Sovereignty Ltd., June 2026. agenticgovernance.digital · CC BY 4.0