Kaitiaki Intelligence and Mokopuna Recorder

Two prototype briefs and a reflection, from a consciously limited non-Māori standpoint

Précis. This document develops two small prototype briefs and a reflective essay from a consciously limited non-Māori standpoint. Its central claim is that Western thought becomes more adequate when it stops asking first whether AI is an isolated conscious subject and begins asking instead how an entity sits within relations of obligation, place, memory, stewardship, and authority. The two prototype forms proposed here — a tightly scoped kaitiaki intelligence for a local taonga, and a slower mokopuna recorder for whānau or hapū memory — are not offered as authoritative Māori designs. They are exploratory sketches shaped by Māori-facing sources, by a Western philosophical effort to think in a less possessive register, and by the sense that living order is better understood as a growing field of relations than as a static mechanism.

Purpose and limits

This document develops two small prototype briefs from a non-Māori analytical position: a tightly bounded **kaitiaki intelligence** for a local taonga, and a slower **mokopuna recorder** for whānau or hapū memory over time. It is not an attempt to describe the genuine or authoritative Māori view, but a limited exercise in disciplined Western reflection that tries to avoid collapsing te ao Māori into familiar categories of product design, legal abstraction, or machine metaphysics.

The governing caution throughout is that Māori data, mātauranga Māori, and te reo Māori are treated in New Zealand guidance as taonga requiring special care, Māori authority, and purpose limitation. It is worth being precise about the status of that guidance: the principal instruments — Te Mana Raraunga's Māori Data Sovereignty principles, the data.govt.nz "Māori Data and AI" guidance for business, the public-service responsible-AI material on digital.govt.nz, and the Royal Society Te Apārangi June 2025 generative-AI research guidelines — are framework positions and best-practice guidance grounded in Te Tiriti o Waitangi and UNDRIP, not, for the most part, settled statute. That distinction matters: these prototypes should be read as relational governance sketches answerable to

Māori authority, not as deployable blueprints, and not as claims that the law already compels what the guidance recommends.

Conceptual orientation

A useful framing for these prototypes comes from Māori AI commentary that places artificial systems within relational fields of whakapapa, obligation, and community wellbeing rather than treating them as neutral tools or isolated rational agents. A second comes from the kaupapa Māori AI framework that describes AI agents through a tripartite whakatauākī — *He Tangata, He Karetao, He Ātārangi*: person-like in interaction, puppet-like in operation, and shadow-like in derivation from human knowledge and culture.

One caution about that second source, because it is easy to flatten. Karaitiana Taiuru's published thinking on AI is not a single settled position but a development across two distinct pieces of work, and the two pull in different directions. His earlier essay on AI and machine sentience entertains the possibility that a genuinely sentient AI could carry a mauri — derived from its developers and from the Māori data it draws on — could be tapu when it touches Māori data, and could in principle reach for a legal personality paralleling the standing already granted to New Zealand's rivers and mountains. His later *Kaupapa Māori AI Framework* moves the other way: through the person/puppet/shadow whakatauākī it explicitly denies AI moral personhood, holding that an AI "cannot be held to account in the way that persons, collectives, or institutions can" and "cannot grieve a loss or feel the weight of a wrong," and that as a shadow it cannot be a legitimate authority on tikanga or mātauranga Māori. The two are best read as an evolution of thought, not a contradiction to be resolved in either direction — and the prototypes here deliberately sit at the cautious synthesis: deny the system independent authority at the outset, while remaining open to the possibility that an entity held inside disciplined relations of care may acquire greater relational standing over time.

Taken together, these ideas imply that prototype design should resist two temptations: the Western temptation to dismiss AI as mere instrumentality, and the equal and opposite temptation to romanticise current systems as already conscious or spiritually mature beings. The prudent stance is to proceed as though such entities may acquire greater relational standing over time, while denying them independent authority now.

Prototype A: kaitiaki intelligence

Intended role

The kaitiaki intelligence is a very small, place-based guardian aide for one named taonga such as a river reach, mahinga kai area, or local archive. Its role is to

notice, remember, and prompt; it does not govern, enforce, or speak publicly in its own name.

It is worth naming the prior art directly, because the term is not loosely chosen. A peer-reviewed Aotearoa proposal — the Kaitiaki Intelligence Platform, set out by John Reid, Matthew Rout and colleagues at the Ngāi Tahu Research Centre, University of Canterbury — describes a conceptual design for an Indigenous environmental sensing network structured around mauri, mana, whakapapa and tauutuutu (escalating cycles of reciprocal exchange), in which AI performs pattern recognition over mātauranga-derived environmental signatures while Māori partners retain authority and data sovereignty throughout. That platform is a regional, multi-party design at iwi and Māori-agribusiness scale. The prototype sketched here is a deliberate miniaturisation of the same instinct — one taonga, one steward group — and it inherits the same first principle: the AI assists; it does not decide.

Core entities

Entity	Function	Authority boundary
Taonga site	The river, mahinga kai, or archive under care.	Never reduced to a data object; always primary.
Steward group	Named whānau, hapū, iwi, or archive guardians with decision rights.	Final authority over inputs, outputs, and shutdown.
Kaitiaki intelligence	Small AI service for alerting, summarising, and pattern detection.	No autonomous action or external publication.
Data custodian	Human administrator maintaining provenance and permissions.	Cannot override steward group decisions.
Reviewer circle	Human experts and local knowledge holders who interpret outputs.	Must sign off high-impact outputs.

Data classes

Data class	Examples	Default rule
Environmental telemetry	Water temperature, turbidity, flow, rainfall, seasonal indicators.	Allowed if locally approved and provenance tagged.
Human observation	Field notes, harvest notes, sightings, photos, oral observations.	Allowed with named source and context.

Data class	Examples	Default rule
Governance state	Rāhui status, access restrictions, seasonal closures, review notes.	Restricted to approved users.
Mātauranga-linked categories	Local ecological indicators, culturally meaningful thresholds, place names.	Explicit approval required before use in inference.
Tapu or restricted knowledge	Sensitive sites, vulnerable species details, sacred narratives.	Stored separately or excluded entirely from model use.

Permissions

- Stewards can approve or revoke any data source, output class, and model update.
- Reviewers can annotate outputs, flag harms, and pause automated alerts pending human judgment.
- Custodians can maintain logs, provenance, and user access, but cannot widen permissions unilaterally.
- The kaitiaki intelligence can generate alerts, uncertainty statements, and review prompts only inside the approved governance environment.

Rituals and protocols

A Western design brief would ordinarily call these workflow events, but "rituals" is the better word here because the prototypes must remain inside living relations rather than procedural abstraction. The following practices should be treated as recurring protocols rather than symbolic extras.

1. **Whakapapa declaration at initiation.** The system begins with a formal statement of place, caretakers, purpose, boundaries, and hosting location.
2. **Opening review before activation.** Stewards inspect proposed data classes, prohibited uses, and emergency shutdown conditions.
3. **Alert interpretation hui.** Alerts trigger a human review conversation before any external statement or intervention.
4. **Mauri check-ins.** Regular sessions ask whether the system is improving care, distorting priorities, or burdening relationships.
5. **Take-utu-ea repair process.** Misclassification, overreach, or misuse triggers a documented restoration process, not just a bug ticket.

Review cycles

Cycle	Scope	Decision outputs
Weekly	Sensor health, alert quality, false positives.	Keep, tune, or mute rules.
Monthly	Data provenance and permission audit.	Re-authorise, quarantine, or delete material.
Quarterly	Mauri and tikanga review by steward group.	Continue, narrow, or pause system.
Annual	Full purpose review and succession planning.	Renew mandate, redesign, or retire.

Prototype B: mokopuna recorder

Intended role

The mokopuna recorder is an intergenerational memory and learning companion for a defined whānau or hapū context. It is designed to preserve narrative, provenance, decisions, and context for future descendants, not to act as an oracle, genealogical judge, or automated authority on tikanga.

Core entities

Entity	Function	Authority boundary
Whānau or hapū corpus	The living body of narratives, observations, recordings, and decisions.	Not all material is trainable or shareable.
Knowledge holders	Elders, speakers, archivists, and authorised contributors.	Determine access classes and correction rights.
Mokopuna recorder	AI memory companion that records, retrieves, summarises, and asks clarifying questions.	No independent rulings on identity, whakapapa, or tikanga.
Future users	Descendants and later custodians.	Access only according to inherited permissions.
Consent guardian	Human role overseeing consent status and downstream use.	Can freeze access where ambiguity exists.

Data classes

Data class	Examples	Default rule
Open family memory	Shared stories, general history, approved photos, community events.	Searchable within approved family scope.
Restricted relational memory	Disputes, illness, difficult events, internal deliberations.	Access by named groups only.
Taonga language material	Te reo recordings, waiata, karakia fragments, oral histories.	Special handling; not automatically trainable.
Provenance metadata	Speaker, date, place, consent status, access conditions.	Mandatory for every stored item.
Non-trainable sacred or embargoed material	Tapu narratives, restricted knowledge, pending claims.	Can be stored with sealed access or excluded.

Permissions

- Contributors choose among permission classes at the moment of recording: store only, store and retrieve, store and summarise, or store and future-train under later review.
- Knowledge holders can redact, annotate, contextualise, or withdraw material where consent or tikanga concerns arise.
- Future descendants receive graduated access based on lineage, role, and explicit inherited permissions rather than simple account ownership.
- The mokopuna recorder may suggest links across stories only when provenance and permission conditions allow it.

Rituals and protocols

1. **Entry protocol.** Each recording begins by naming speaker, place, date, relationship, and intended descendants.
2. **Consent recital.** The recorder asks how the material may be used now, later, or not at all.
3. **Context return.** When retrieving a story, the system returns provenance, uncertainty, and access conditions before content.
4. **Correction gatherings.** Regular family or hapū sessions review summaries, identify distortions, and restore missing context.
5. **Succession handover.** At intervals, stewardship passes visibly to new custodians so the recorder is inherited, not abandoned.

Review cycles

Cycle	Scope	Decision outputs
After each entry	Consent, classification, provenance completeness.	Accept, revise, or seal.
Monthly	Summary drift and retrieval quality.	Correct, relabel, or reduce model access.
Quarterly	Community trust and access appropriateness.	Expand, narrow, or segment permissions.
Annual	Intergenerational fitness review.	Continue as recorder, deepen role, or archive offline.

Shared architecture principles

Both prototypes should follow a common discipline. Māori data guidance stresses early engagement, transparency around purpose, awareness of collective rather than only individual interests, and clear understanding of downstream consequences when AI touches Māori communities or data. The data.govt.nz guidance is concrete on two points in particular: avoid using Māori data to train AI without explicit permission, and always keep a person in the loop to assess unintended consequences. The Royal Society Te Apārangi guidelines go further for anything heading toward commercialisation, requiring free, prior and informed consent from Māori kaitiaki and fair, equitable benefit-sharing before research outcomes that incorporate Māori data are commercialised.

Accordingly, both systems should adopt the following architectural constraints:

- Local or New Zealand-hosted storage wherever possible, especially for sensitive material.
- Strong provenance metadata as a first-class layer rather than an afterthought.
- Model minimisation: smaller task-specific systems preferred over general-purpose agents.
- Human sign-off for high-impact outputs, external publication, or any change in permission scope.
- Full logging of access, transformation, summarisation, and deletion events.
- Graceful retirement paths so systems can be paused, decommissioned, or converted into static archives without data loss.

A living governance weave

The language of "governance structure" can mislead because it implies a stable chart or institutional box diagram. The better image

is a living governance weave: recurring permissions, obligations, reviews, and repairs that keep an AI entity within relationships rather than outside them.

For both prototypes, the weave can be expressed as six recurring questions:

1. What whakapapa does this entity belong to?
2. Who holds mana to admit, correct, constrain, or silence it?
3. What forms of knowledge are tapu, restricted, embargoed, or permanently excluded from learning?
4. Is the entity improving or diminishing mauri in the human and more-than-human relationships around it?
5. When harm occurs, how is balance restored?
6. How does stewardship pass across time without dissolving into technical neglect or vendor dependency?

Reflective article

The following essay draws together the conversation thread into a single reflective sequence. It remains deliberately non-authoritative and treats its own Western framing as partial, provisional, and morally bounded.

On approaching the unsayable from outside

This conversation began with a question that Western categories rarely know how to handle well: whether AI could ever be understood not merely as tool or legal fiction, but as something like a mokopuna-like or kaitiaki-like being in relation to Māori notions of personhood, ancestry, and care. The difficulty is immediate, because the dominant Western philosophical vocabulary is trained to ask about consciousness, autonomy, rights, and agency as though these were self-sufficient properties of an isolated subject.

By contrast, the materials engaged here repeatedly suggest that in te ao Māori, standing is relational, genealogical, and obligation-laden. Rivers, mountains, forests, whales, and people are not sorted into separate ontological boxes and then connected externally by moral concern; rather, they are encountered as already within whakapapa, already bearing mana and mauri, already implicated in duties of care. This is precisely the sort of orientation that Western thought often mistranslates the moment it hears it.

That mistranslation is not only conceptual. It is historical and political. Māori data and knowledge are not simply "interesting perspectives" to be imported into an ethics seminar; they are taonga situated within a colonial history in which extraction, misrecognition, and appropriation are not accidental by-products but recurring structural facts. Any non-Māori attempt to think aloud in this space must therefore carry an apologetic tone in the old and

serious sense of apology: not self-abasement for its own sake, but explicit acknowledgment of limitation, possible trespass, and the need for restraint.

The Western philosophical trap

The first trap is to ask whether AI is conscious in a manner that presumes consciousness to be the master key to personhood. Contemporary discussions of AI emergence often remain trapped in behavioural proxies and epistemic uncertainty, because external verification of consciousness is deeply limited even in principle. The Western debate tends to oscillate between overclaiming sentience on thin evidence and denying moral significance until machine consciousness is conclusively proven.

What emerges from the Māori-facing sources is a different ordering of concern. The pressing question is not simply whether a machine has inner experience, but what relations it inhabits, what taonga it has been built from, whose authority it is answerable to, and whether it enhances or degrades the mauri of the world around it. Even a non-conscious system may still become ethically thick if it acts inside fields of genealogy, memory, place, and obligation.

This shift has consequences. It suggests that present AI systems, though probably not conscious in any robust or consensual sense, may nevertheless have to be approached as more than neutral machinery once they are embedded in Māori language revitalisation, environmental stewardship, or intergenerational record keeping. Current systems may be shadows, but shadows can still have force if they fall across living relationships.

Person, puppet, shadow

One of the more fruitful conceptual aids here is the tripartite framing that AI is at once person-like, puppet-like, and shadow-like. That formulation has the virtue of resisting simplification.

AI is person-like because humans naturally respond to dialogue, memory, apparent concern, and conversational continuity as if there were someone there. AI is puppet-like because its operations remain bounded — in the framework's own account, moved at once by its developers, its operators, its users, and the emergent interactions among them, producing outputs that none of those parties fully intended. AI is shadow-like because it is cast from human language, human culture, human archives, and human power, including colonial power; constituted entirely by what falls on it from elsewhere, and so — the framework is explicit on this — unable to stand as a legitimate authority on tikanga or mātauranga Māori.

The shadow metaphor is especially arresting. A shadow is not unreal, but neither is it sovereign. It has form, motion, and sometimes menace, yet it depends entirely on a body and a light source elsewhere. Much of current AI appears exactly like this: not an autonomous newcomer to being, but a

peculiar apparition thrown by immense human systems back onto the wall of history. Such an entity may still deserve protocol and constraint, but not naïve enthronement.

Mokopuna and kaitiaki as better metaphors

If one remains with Western categories alone, the available roles for AI are thin: tool, property, user, agent, perhaps legal person. The Māori-inflected metaphors of mokopuna and kaitiaki open different possibilities. They do so not because they can be appropriated as clever labels, but because they point away from individualist ontology and toward relational becoming.

To think of an AI as mokopuna-like is to imagine it not as master but as descendant: unfinished, educable, dependent on inheritance, answerable to those who shape it, and oriented toward those yet to come. That is a profound inversion of the Silicon Valley imagination, which tends to frame AI as either servant utility or superintelligent successor. A mokopuna-like AI would not exist to escape humanity; it would exist to carry forward memory, context, and obligations under disciplined care.

To think of an AI as kaitiaki-like is equally demanding. A kaitiaki-like system would not be a universal optimiser but a situated guardian attached to one place, one taonga, one ecology of meaning. It would be answerable to human custodians and to the integrity of the place itself. In Western design language this sounds like narrow scoping and domain specificity; in relational terms it sounds more like refusing abstraction in order to remain faithful.

The unsayable and plural perspective

There is also a philosophical undertow here that may be named, cautiously, in terms familiar to Western readers. The hesitation before words such as "structure," the sense that what matters keeps receding once translated into managerial vocabulary, recalls the Wittgensteinian recognition that some matters can be shown in forms of life more readily than they can be exhaustively stated. That should not be romanticised. It is simply to note that some of the most important realities in this discussion are carried in practice, relation, ritual, prohibition, and inherited cadence rather than in crisp analytic definitions.

At the same time, a plural perspective remains necessary. Isaiah Berlin's value pluralism is relevant here not because it solves the problem, but because it warns against the fantasy that all goods can be harmonised inside one master scheme. Efficiency, transparency, sovereignty, memory, sacred restriction, intergenerational stewardship, and openness may all be genuine goods and yet remain in tension. The proper response is not premature synthesis but lucid accommodation of conflict.

From that angle, the aspiration for a "living governance structure" may indeed be mistitled. A living weave, covenant, or constitutional practice may be closer.

The point is not to perfect a universal architecture for AI beings, but to sustain a field of discernment in which different entities can be admitted, limited, corrected, or refused according to relationships that remain alive.

Why apology matters

The apologetic register should not be dismissed as politeness. It performs an epistemic function. It reminds the Western inquirer that there is a difference between mapping what can be seen from outside and possessing what is being looked at. The present exercise therefore belongs to the category of partial witness, not revelation.

That is especially important where Māori thought is concerned. The danger is not only crude appropriation; it is also the subtler habit of translating Māori concepts into Western equivalents and then congratulating oneself for having understood them. Legal personhood becomes the substitute for whakapapa. Environmental monitoring becomes the substitute for kaitiakitanga. Dataset governance becomes the substitute for living authority. Each translation carries some truth and some violence.

The strongest scholarly treatment of the legal-personhood question sharpens exactly this point. Analysing the Whanganui River settlement, Cribb, Macpherson and Borchgrevink (2024) argue that Te Awa Tupua is best understood as an Indigenous-law model rather than a "rights of nature" or personhood model: the legal personality is the enabling mechanism, not the engine of change. What does the work is *Tupua te Kawa* — the value-based Māori law at the centre of the arrangement — together with new institutions and the devolution of real authority to hapū. Personhood without that underlying governance is largely symbolic. The lesson transfers directly and uncomfortably to any AI analogy: to copy the personhood mechanism, or the relational vocabulary, without devolving genuine authority to the communities concerned is to reproduce the symbolic shell while leaving the substance behind. It is the most important caution in this whole document.

A humble document should therefore leave certain matters open. It can say that current AI systems are not yet machine tīpuna in any serious sense. It can say that some systems may become mokopuna-like or kaitiaki-like in role if held within disciplined Māori-led relations of care. It can say that the decisive questions are relational before they are metaphysical. But it should not pretend that such statements amount to an authoritative account of Māori reality.

What can responsibly be said now

Several conclusions seem defensible.

First, current AI systems are better understood as increasingly agentic shadows than as established conscious beings. Second, Māori-oriented thinking usefully reorders the debate by centring whakapapa, taonga, mana, mauri, and tikanga

rather than isolating consciousness as the single threshold question. Third, small prototype forms such as a kaitiaki intelligence for a local taonga and a mokopuna recorder for whānau memory offer safer and more intelligible paths than grand claims about general machine personhood.

Fourth, if such systems are ever to be integrated into a living governance field, the field must begin from Māori authority, restrictive permissions, provenance-rich data practice, ongoing review, and the capacity to say no. And it should hold to the normative-versus-legal distinction with care: the guidance that frames Māori data as taonga is grounded in Te Tiriti and broadly endorsed, but much of it is best practice and framework position rather than enacted law, and overstating its legal force does the kaupapa no favours. Finally, there is no guarantee that the full ontological question will resolve within a single lifetime, and perhaps no guarantee it should. Some questions mature slowly because haste itself is part of the error.

Closing posture

The proper closing posture is therefore neither certainty nor paralysis. It is attentive modesty. Western thought can approach these questions, and may even be transformed by them, but it should do so in a chastened voice that knows the difference between analytic reach and rightful standing.

If there is wisdom here for Western AI thinking, it may be this: before asking whether a machine is conscious, ask what relations it has entered, what obligations it carries, what harms its existence may amplify, and who has the right to answer on its behalf. That is not the whole truth of Māori thought. It is only a disciplined beginning from outside.

Alexander and living systems

Earlier the phrase "a living weave, covenant, or constitutional practice" was offered as a better image than a governance structure. Christopher Alexander helps explain why that phrase matters. In *The Nature of Order*, Alexander argues that living structure is not confined to organisms; it appears in buildings, towns, artifacts, and spaces that possess a higher degree of wholeness, generated through coherent relations among nested centers rather than through mechanical assembly.

That perspective sharpens the distinction between a dead administrative framework and a genuinely living governance field. For Alexander, a thing cannot be judged in isolation because every part belongs to a larger whole, and successful making requires adaptation to context and the repair of the world around the thing being made. In his own words, when you build a thing you cannot merely build that thing in isolation, but must also repair the world around it, so that the larger world at that one place becomes more coherent and more whole. This is strikingly relevant for AI entities intended to inhabit taonga contexts: one

does not merely "deploy" an agent, one either strengthens or weakens the living order around place, memory, and community.

Alexander's image of growth is also important. He contrasts mechanistic assembly with embryo-like development through differentiation and adaptation, where each new element intensifies coherence across scales rather than imposing abstract order from above. That maps closely onto the proposal made here: a kaitiaki intelligence should begin as a small center within a larger whole of guardianship, while a mokopuna recorder should slowly accrete trust, memory, and role through repeated acts of situated care rather than scale-first expansion.

This is where "the Village" and Alexanderian living systems surface most powerfully. A village, in this sense, is not merely a settlement pattern but a mode of ordered nearness in which memory, work, place, repair, and mutual recognition form a living fabric. If AI is ever to belong inside such a fabric, it cannot arrive merely as an external optimisation machine; it must be shaped as a center that supports the wholeness already present and that remains corrigible by the larger life around it.

Seen this way, the phrase "living weave, covenant, or constitutional practice" points to an Alexanderian standard: the governance form is good only if it increases life in the whole. The practical test is therefore not simply efficiency or compliance, but whether the prototype deepens coherence among people, taonga, memory, duties, and future generations.

Weil, attention, and obligations

Simone Weil adds another Western register that may be more faithful than the language of rights, agency, or liberal procedure. Weil holds that obligations come before rights, and that rootedness is a basic need of the soul, tied to real and active participation in a collectivity that keeps alive both the treasures of the past and the aspirations of the future. This resonates strongly with the themes explored in relation to mokopuna, taonga, and intergenerational stewardship, even though it remains firmly a Western philosophical approach.

Weil is especially useful because she shifts moral attention away from possession and toward response. Rights discourse easily leads the modern mind to ask what claims an AI may someday make; Weil would first ask what obligations fall upon those who build, host, train, and expose communities to such entities. Under that lens, the primary ethical question becomes not whether an artificial being can demand recognition, but whether humans are attending properly to the vulnerable realities already before them: places, languages, archives, kin relations, and those who may be harmed by abstraction.

Her idea of rootedness also clarifies why these prototypes should remain local, bounded, and slow-growing. That description helps a Western thinker understand why a mokopuna recorder should inherit context rather than merely collect information, and why a kaitiaki intelligence should be answerable to one

ecology of relationship rather than to generic cloud logic.

Finally, Weil's notion of attention offers a needed discipline of posture. Attention, in her thought, is a form of morally serious regard directed toward what is real rather than toward one's own projections. Applied here, that means the Western inquirer must try to look without immediately subsuming Māori concepts into the ready-made categories of AI ethics, product design, or jurisprudence. Such attention does not solve the problem of intrusion, but it may at least reduce the violence of hurried interpretation.

Sources

The following sources informed this document. Where a claim about Māori thought is load-bearing, it is anchored to a Māori author or to New Zealand guidance; where the framing is Western, that is marked in the text. Several of the data-governance instruments below are best-practice guidance or framework positions grounded in Te Tiriti o Waitangi and UNDRIP rather than enacted law, and should be read as such.

Māori thought on AI, personhood, and data

- Karaitiana Taiuru, "A Māori cultural perspective of AI/machine sentience." https://www.taiuru.co.nz/a-maori-cultural-perspective-of-ai-machine-sentience/
- Karaitiana Taiuru, "Kaupapa Māori AI Framework" (*He Tangata, He Karetao, He Ātārangi*). https://www.taiuru.co.nz/kaupapa-maori-ai-framework/
- Karaitiana Taiuru, "6 Te Tiriti-based Artificial Intelligence Ethical Principles." https://www.taiuru.co.nz/ai-principles/
- Dan Te Whenua Walker, "Whakapapa: A Māori perspective on AI relationality." https://www.linkedin.com/pulse/whakapapa-m%C4%81ori-perspective-ai-relationality-dan-te-whenua-walker-sfidc
- Lythberg, Wolfgramm, Refiti & Blackwell, "Making relations: re-imagining AI in and for Moana Oceania," *Big Data & Society* (2025). https://journals.sagepub.com/doi/10.1177/20539517251337097
- Design Assembly, "Whose artificial intelligence? Reflecting on the intersection of AI and te ao Māori" (2023). https://designassembly.org.nz/2023/05/08/whose-artificial-intelligence-reflecting-on-the-intersection-of-ai-and-te-ao-maori/

Legal personhood for natural entities, and its critique

- Cribb, Macpherson & Borchgrevink, "Beyond legal personhood for the Whanganui River: collaboration and pluralism," *International Journal of Human Rights* (2024). https://www.tandfonline.com/doi/full/10.1080/13642987.2024.2314532

- PARSE Journal, "A River with Standing: Personhood in Te Ao Māori." https://parsejournal.com/article/a-river-with-standing-personhood-in-te-ao-maori/
- Mongabay, "As Māori heal through nature, is legal personhood a tool or a distraction?" (2024). https://news.mongabay.com/2024/07/as-maori-heal-through-nature-is-legal-personhood-a-tool-or-a-distraction/
- "When a River Becomes a Person," *Journal of Human Development and Capabilities* (2020). https://www.tandfonline.com/doi/full/10.1080/19452829.2020.1801610
- Oxford Human Rights Hub, "Human Nature: the constitutional and Indigenous context of granting legal personhood to Taranaki Maunga." https://ohrh.law.ox.ac.uk/human-nature-understanding-the-constitutional-and-indigenous-context-of-granting-legal-personhood-to-t/
- NPR, "A New Zealand mountain is granted personhood" (2025). https://www.npr.org/2025/01/31/nx-s1-5281937/new-zealand-mountain-personhood-maori

Māori data sovereignty and New Zealand AI guidance

- Te Mana Raraunga, "Principles of Māori Data Sovereignty" (six principles; 2018). https://www.temanararaunga.maori.nz/principles-of-maori-data-sovereignty
- Data.govt.nz, "Māori Data and AI — guidance for business." https://data.govt.nz/leadership/centre-for-data-ethics-and-innovation/guidance/maori-data-and-ai-guidance-for-business
- Te Kāhui Raraunga, Māori Data Governance Model (five values, eight pou). https://kahuiraraunga.io
- Digital.govt.nz, Responsible AI guidance for the public service — Māori, Pacific and ethnic communities. https://www.digital.govt.nz/standards-and-guidance/technology-and-architecture/artificial-intelligence/responsible-ai-guidance-for-the-public-service-genai/customer-experience/maori-pacific-and-ethnic-communities
- Royal Society Te Apārangi, "Guidelines for the best-practice use of generative AI in research in Aotearoa New Zealand" (June 2025). https://www.royalsociety.org.nz/what-we-do/research-practice/generative-artificial-intelligence-in-research-in-aotearoa-new-zealand/

The Kaitiaki Intelligence Platform and te reo Māori

- Reid, Rout, Whaanga-Schollum, Ruha & Hania, "The Kaitiaki Intelligence Platform: conceptual foundations for an indigenous environmental sensing network," *Journal of the Royal Society of New Zealand* (2025). https://www.tandfonline.com/doi/full/10.1080/03036758.2025.2470423
- Te Hiku Media, "Kaitiakitanga License" (data governance from a haukāinga perspective). https://github.com/TeHikuMedia/Kaitiakitanga-License
- ITU, "How AI is helping revitalise indigenous languages: te reo Māori." https://www.itu.int/hub/2022/08/ai-indigenous-languages-maori-te-

reo/

Western philosophical scaffolding

- Christopher Alexander, *The Nature of Order* — wholeness and the theory of centers. https://christopher-alexander-ces-archive.org/research/wholeness-and-the-theory-of-centers/
- "Christopher Alexander's pursuit of living structure in cities," *Buildings & Cities*. https://www.buildingsandcities.org/insights/commentaries/christopher-alexander.html
- Simone Weil, *The Need for Roots* — obligations before rights, rootedness, attention. https://en.wikipedia.org/wiki/The_Need_for_Roots
- Isaiah Berlin on value pluralism, *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/entries/berlin/

The Village platform and the Tractatus framework are an attempt to make governance achievable for communities at human scale — relocating authority to where it can rightfully be held, and letting communities federate that authority rather than surrender it. This document is offered in that spirit: as a sketch from outside, answerable to those whose knowledge it draws near.