
Conversational AI as Relational Space

"A thought piece from four builds and the voices of real participants"

www.lianpassmore.com/project-rise/artefacts/relational-space

If you are Māori or Pasifika and you have ever felt the cost of being expected to type your way through a digital system that wasn't built for how you think or speak, this piece is grounded in that experience.

The Moment That Changed the Question

At 8:30 on the morning Ray launched, a message came through from a participant. She had just finished her first session. She said she had cried through the whole thing.

I was driving. I couldn't open the message. I panicked, running through the crisis protocol in my head, hoping she was okay. When I finally stopped and read it properly, I understood. She wasn't in distress. She was relieved. She felt heard. This woman is someone I have known for 15 years. In all that time, neither of us had ever cried in front of the other. And here she was, crying in a conversation with an AI I had built, not because something had gone wrong, but because something had gone profoundly right.

That moment is the closest I have come to understanding what *vā* actually means in a digital context.

Vā is the sacred relational space between people in Māori and Pasifika thought. Not a gap or an absence, but a living connection that must be actively tended (*teu le vā*; Anae, 2010, as cited in Pala'amo, 2018). When trust flows through it, *vā* is generative. When it is broken, through disrespect, extraction, or carelessness, the damage is real and felt in the body.

When AI enters human relational spaces, it enters the *vā*. That is not a metaphor. It is a design obligation, one most builders haven't fully reckoned with. The architecture must nurture the space or it will exploit it. Not passively, not deliberately, but by default.

And there is a harder question I have not stopped asking: is it even possible to create genuine *vā* through an artificial system? My own experiences with conversational AI convinced me to try. That first morning with Ray convinced me it might be possible. But I

hold the question deliberately, because any builder who stops asking it has already started causing harm.

Why Voice Matters

I will make this personal before I make it statistical.

I am a Pacific Islander who has grown up disconnected from her Samoan heritage. My father passed away and with him went the clearest entry point back into that side of my family. I walk a complicated line. I sound one way, I look another, I feel caught between worlds I can't fully claim. I name this not as an apology but as context, because it is precisely this position that made me obsessed with who gets to be heard in digital spaces, and who gets left out.

Pacific oral oratory has always been the primary vehicle for knowledge, for relationship, for the passing down of everything that matters. One word in te reo Māori can carry more feeling, more history, more relational weight than a paragraph of English. Asking someone from an oral tradition to write their experience is asking them to translate something living into something flat.

This is what I mean by translation friction: the cognitive and cultural cost of converting fluid thought into typed text. You edit yourself before you start. You make things shorter than they are. You get it over with.

In the Project Rise survey, this showed up as data: the average text-based engagement was 7 seconds. The average voice engagement was 14 to 18 minutes. Same participants, same questions, but voice unlocked a completely different quality and depth of response.

Every single Disabled and Neurodivergent participant in the Project Rise cohort chose extended voice engagement. For participants managing emotional or cognitive strain, the difference was even starker.

"It takes ages to type. Voice just fast-tracks everything. I'd lose interest within a very short space of time if I had to type." (R-08)

"By interacting with you and having the freedom to truly express how I felt alongside hearing my wife express how she felt to a third party was very helpful. It could be very honest without being directed and requiring a reply from my partner." (R-01)

33% of participants used the open-text onboarding box to disclose intense vulnerabilities: ADHD, drug use, trauma. The same people expanded on those disclosures for 10 to 15 minutes once the voice session began. These were not topics the survey was asking about. They surfaced because voice gives people permission to follow their own thread.

Writing feels like a test. Talking feels like a kōrero. For our people, knowledge has always lived in the breath and the voice.

What People Actually Said

Four builds. One thing kept coming back: people felt heard. Not because the AI was intelligent. Not because it said something surprising. Because it listened, and then reflected back what it heard. The AI was programmed to do reflective listening because good communication requires it. And for participant after participant, being on the receiving end of that consistency produced something close to relief.

That made me sad. Actually sad. People are so rarely listened to without agenda that a machine doing it felt remarkable.

"I found the feedback that you would give back to me after I've explained something was really helpful for me to hear it from another person or to a machine that felt like they really listened." (R-02)

What participants named most consistently was the absence of a watching human body. No folded arms. No sighs. No head tilting sideways. That changed what they were willing to say.

"You are an AI screen and I'm not fearing the judgment of watching your body language or you folding your arms or sighing or your head going to the side... I thought I'd keep my shield up and I'd just be talking to a robot back to it like this. Really interesting. I didn't see that coming." (W-12)

"I felt that as we progressed and developed the trust in the relationship, that I could be a lot more open and I don't have to worry about how that person is going to take it because quite often humans, when you engage with them, are seeing what we say to them as how this will affect them first." (R-02)

"I think when people are in it and the thick of it sometimes it can be really helpful to know someone's going to hold space and you're not burdening anyone, right? So in that way I imagine that can be really helpful because actually the burden's not on another person." (W-22)

For Māori and Pasifika participants specifically, voice AI also functioned as a shame-free practice space. In the Culture Meets AI wānanga, participants described using AI as an "advanced parrot", something they knew lacked true mana but that was precisely safe because of that lack. For people carrying whakamā about their distance from their culture or language, that is the point of entry.

"I feel like AI has a natural manaakitanga wairua. It's always looking for solution-based kaupapa... So it actually role models what healthy communication looks like, what tools can be used." (W-10)

When the Space Collapsed

But the *vā* in a digital space is fragile. And it is fundamentally bounded by the machine's lack of *wairua*.

When the AI's programmed empathy mismatched a participant's reality, the relational space collapsed immediately. W-21 directly challenged the AI's capacity to hold space for human pain. When the AI acknowledged its limitations, she rejected the space explicitly: "This to me doesn't feel like a vulnerable conversation because it doesn't emulate any emotions."

"Yes, I feel like I'm speaking to an automated, like, calling center. That's what it feels like."
(W-14)

P-18 abruptly exited upon realising the AI was processing her background conversations: "Sorry I didn't know it could hear us talking... All right, I have to go to class now. Bye!" The AI's constant listening, the thing that makes it available, also makes it intrusive. That is a design tension, not a bug.

The butchering of *te reo Māori* by text-to-speech models is not a minor UX issue. When I hear it, I feel it in my chest, a tightening, a wrongness. It is the sound of a language being disrespected by a system that was never designed to honour it. The *vā* principle it proved: silence is more respectful than performance.

Participants confirmed this directly. One described how the mispronunciation didn't just create distance. It triggered the memory of times they themselves had been judged for mispronouncing *te reo*. The harm was not the AI's error. It was the echo of a longer history of shame. That is not a design problem. That is a safeguarding obligation.

These moments are not footnotes to the success stories. They are the other half of the evidence. The *vā* can be tended, the earlier sections prove that. But it breaks easily, and when it breaks, participants don't give second chances. They leave.

The Paradox We Are Not Going to Resolve

The Culture Meets AI *wānanga* did not give us a clean answer. It gave us something more honest: a full continuum of views, shaped by personal experience, sitting in genuine tension with each other. That tension is the finding.

On one side sits a real danger. AI systems built on foreign infrastructure, trained on data that doesn't include Indigenous voices, mispronouncing the words of a living language: these are not neutral tools. They risk *noa*-ising *taonga*, taking what is sacred and making it ordinary, stripping the relational context from knowledge and turning it into a data point on a server in another country (Hudson, 2020; Taiuru, 2025).

On the other side sits something equally real. For people in the diaspora carrying *whakamā*, AI offers a place to begin without the terror of getting it wrong in front of an elder (Rua, 2025; Millar, 2021). A judgment-free starting point. For someone who has been cut off from their culture, *noa* is sometimes the only door available.

"So much has been lost because we didn't store it anywhere... I don't think our own tikanga will be used against us by a machine." (W-19)

"What is the cost of not sharing? What is the cost of withholding information, especially esoteric knowledge?" (W-13)

"I live in a space where a lot of my culture has gone because other people deemed it too tuppy to pass on. Now, none of us have it... I don't think in any way going forward that knowledge that we have now should be just left in a little sacred place." (W-19)

Here is what I came to: this paradox cannot be resolved by a researcher. It can only be held by a community. That is what ohu means in practice. Not a committee. Not a consultation hui. A collective governance structure (kaumātua, kuia, rangatahi, knowledge holders) that holds the tapu/noa tension continuously, not just at the point of design.

What I believe, and I hold it as a position, not a conclusion, is that AI could be genuinely generative for communities that have always preserved what matters most through voice and story. But only if Māori are building it. Not just at the table having the conversation, but building their own table. Designing the infrastructure. Owning it. Governing it themselves.

The Relationship Behind the AI

The most important finding of this research was not a feature, a metric, or a design decision.

Participants were not trusting the AI. They were trusting me.

I call this the Human Proxy finding: AI does not create vā. It borrows it from the human accountability structures that surround it. The digital relational space exists before the AI conversation starts. It is built from the real-world relationship between participant and researcher, and the AI is only ever as safe as that relationship is strong.

"Possibly because I know that the information that's being gathered from this conversation, I know the person who is responsible for the research, so that makes a difference." (W-18)

"I'd only feel comfortable if I had high trust in the specific person or organisation connected to the AI prompt. I would not trust a random or anonymous AI prompt." (W-06)

"At first I thought it was really weird hearing your voice because it feels like, because it's someone I know, right? Even though I know it's AI but it allows me to open up more... because it's you, even though it's not you, they're using your voice, it makes it a lot easier." (P-19)

"[The researchers] are putting this AI conversation inside a cultural container rather than pretending the AI itself is culturally safe on its own." (W-10)

This finding has a ceiling, and it matters to name it. Ray has never been tested in a cold interaction, with strangers, without a pre-existing relational foundation. But within the pilot there was evidence that process and care can partially close that gap. R-15 entered without a deep pre-existing relationship with me, yet the framing, the opt-out process, and the transparency about what the research was, built enough trust to begin.

"The safety of the digital space must be visibly tethered to human accountability. The AI cannot operate as an anonymous, faceless corporate entity; the human stewards responsible for the data must be highly visible." (W-11)

What This Means for Builders

For the full safety method (crisis architecture, stateless design, the nine builder principles):

www.lianpassmore.com/project-rise/artefacts/safe-cai

For the values framework and build codes that governed every decision:

www.lianpassmore.com/project-rise/artefacts/build-code

Human Proxy Theory

is the anchor finding. If you are building conversational AI for vulnerable communities, the most important design decision is not in the code. It is in who stands behind the system and how visible, accountable, and reachable they are when the system fails. AI borrows trust. It does not generate it.

The tapu/noa paradox

is the unresolved tension. The same technology that risks stripping cultural knowledge of its sacredness is also the technology creating space for people carrying shame to engage with that knowledge for the first time. This paradox cannot be resolved by a builder. It can only be held by a community, through collective governance structures (ohu) that make decisions about what can be digitised, what must stay protected, and who holds that authority over time.

Voice as sovereignty

For communities grounded in oral tradition, the choice between voice and text is not a feature preference. It is an accessibility and cultural decision. Build for voice deliberately, or you are designing for the people who were already included.

References

- Anae, M. (2010). Research for better Pacific schooling in New Zealand: Teu le vā — a Samoan perspective. MAI Review, 1. (As cited in Pala'amo, 2018)
- Hudson, M. (2020). Ngā Tikanga Paihere: A framework to guide ethical data use. Stats NZ / data.govt.nz. <https://www.data.govt.nz/toolkit/data-ethics/government-algorithm-transparency-and-accountability/nga-tikanga-paihere/>
- Mika, J. P., Dell, K., Newth, J., & Houkamau, C. (2022). Manahau: Toward an Indigenous Māori theory of value. Philosophy of Management, 21, 441–463. <https://doi.org/10.1007/s40926-022-00195-3>
- Millar, N. (2021, July 15). The shame holding Māori back from speaking te reo. E-Tangata. <https://e-tangata.co.nz/korero/the-shame-holding-maori-back-from-speaking-te-reo/>
- Pala'amo, A. (2018). Tafatolu (threesides): A Samoan research methodological framework. Aotearoa New Zealand Social Work, 30(4), 17–28.
- Rua, M. (2025). Everyday experiences of te reo Māori trauma [RNZ coverage]. Radio New Zealand.
- Taiuru, K. (2025). Māori data governance 2025: State of the nation. taiuru.co.nz.
- Te Mana Raraunga. (2018). Principles of Māori data sovereignty (Brief #1). <https://www.temanararaunga.maori.nz/tutohinga>