

ETHICS & AI TRANSPARENCY

Code of Ethics

The commitments that govern our use of AI in community and stakeholder engagement.

YUZU LABS PTY LTD, TRADING AS YUZU COMMUNICATIONS · VERSION 1.0 · AUGUST 2026

Consultation data isn't ordinary data. Behind every response is a person – often speaking about decisions that affect their home, their community or their livelihood. Used carelessly, AI can misrepresent what a community said: sentiment misread across cultural and linguistic lines, quieter voices collapsed into an 'other' category, the loudest themes mistaken for the most important ones. When that analysis informs a real decision, the error isn't statistical – it's ethical.

This Code sets out how we use AI in our work, where it falls short, and the standards we hold ourselves to. It has guided how we work since Yuzu was founded. Here it is written down – so that we can be held to it.

OUR COMMITMENT

Technology that amplifies human expertise. AI supports professional judgement – it never replaces it. That order never reverses.

Our commitments

- 01 Human oversight, always.** Every AI-assisted analysis is reviewed, refined and validated by experienced practitioners. AI supports professional judgement; it does not replace it. That order never reverses.
- 02 Transparent and reviewable.** No black boxes. Findings can be checked against the responses they came from, and corrected where the context demands it. If we can't show you where a finding came from, it doesn't go in the report.
- 03 Accountable.** The reports we deliver carry our name. When something is wrong, we say so and we fix it.
- 04 Your data stays yours.** You own your data and the outputs – always. Your data isn't retained on our servers, we never use client data to train AI, and on request we delete your data and the record of analysis entirely.
- 05 Privacy by default.** Personal details – names, addresses, contact information – are removed on the local machine before analysis, so only de-identified text is ever sent onward for processing. You can do this before sending us anything, or we'll do it for you.
- 06 No hype.** We describe what our technology can and can't do – honestly. If AI isn't the right tool for your problem, we'll tell you.

What AI can't do

- AI models carry bias from the data they were trained on.
- Sentiment analysis can misread cultural context, community language and nuance.
- Theme detection can over-represent the loudest voices and lose the quietest ones.
- No algorithm understands the history behind a community's concerns, or what's at stake locally.

These limitations are why our model exists: practitioners design the inputs, and practitioners review every output. The technology is powerful. It is never left unsupervised.

Where we use AI – and where we don't

Where we use it

- ✓ Identifying themes and sentiment in large volumes of feedback
- ✓ Categorising responses
- ✓ Removing repetitive analysis work
- ✓ Helping track issues over time

Where we don't

- ✗ Making decisions
- ✗ Producing findings without practitioner review
- ✗ Replacing the professional judgement our clients rely on

How we hold ourselves to this

Every AI-assisted analysis carries a practitioner's name and a practitioner's review. If we can't show you where a finding came from, it doesn't go in the report. And if AI isn't the right tool for your problem, we'll tell you – plainly.

This is a living commitment. We review this Code as our tools and our practice evolve, and we welcome being held to it.