A man with curly brown hair, wearing a light blue shirt and a green jacket, is speaking. He is positioned in front of a whiteboard that contains several diagrams, including a flowchart with a central box and arrows pointing to other boxes, and a diagram with a central box and arrows pointing to a list of items. To the left of the man is a large green plant in a grey pot. The background is a plain white wall.

AI will be normal.

md

P

This was the year that I think the AI got smarter than us.



ENGLISH
LANGUAGE PARTNERS

Working with former refugees and migrants



ACE AOTEAROA 2026 · WORKSHOP

Ethical Gen-AI in ESOL Education

Jia Rong Yap, PhD

Anchored in ELP's AI for Enhanced Teaching project



A national survey

Most teachers had **no formal AI training, most wanted PD**, and the most common barrier was **not knowing where to start** — yet many were **already experimenting** alone.



An AI Working Group

Early-adopter teachers whose **real classroom experience** helped shape workshop we built — practice leading design, not the reverse.



Two PD workshop series

“**Ethical Gen-AI in ESOL Education**” and “**AI Specialisation: Using AI Tools.**” Ethics and purpose were the first things we taught.

*Mapped to the **UNESCO AI Competency Framework for Teachers (2024)** and **TPACK-Gen AI (Petko et al., 2025)**. While ESOL is our ground — the evaluation method transfers to anything you teach.*

Where the sector has landed

Across adult & community education, the field has independently moved to the same question:

OLD QUESTION

“Should we use AI at all?”



THE QUESTION NOW

“How do we use it responsibly — and well?”



Human-mediated, not replacement

The consensus model: teachers stay central; AI drafts, adapts, translates. AI-assisted teaching — not automated teaching.



Framed for inclusion

Strongest signal for our sector: AI to support low-literacy and multilingual adults — simplified text, translation, non-judgemental practice.



Ethics moved to the centre

Focus has shifted from productivity to AI literacy, data privacy, accuracy and cultural fit — and to educator competence.

But the field is still early-stage: definitions are unsettled and **training lags adoption**. This session addresses that gap — between using AI and using it *well*.

UNESCO Guidance on Generative AI in Education & Research (2023, updated 2026); EPALE / EAEA European adult-education practitioner resources (2025)

And our learners are already using it



... but not on their own

AI is often *inaccurate*, *overconfident*, or *too generic* — and underserved learners can be left behind without scaffolding. Recommended best practice is *guided use*.

Consistent with UNESCO (2024, upd. 2026) and EPALE/EAEA adult-education guidance on access and guided use.

Ethical and Purposeful AI Use

You can generate (almost) *anything* now.

Ethical and Purposeful AI Use

You can generate *anything* now.

How we *structure our thoughts and logic* (prompting) gets a *draft*.

You can generate *anything* now.

How we *structure our thoughts and logic* (prompting) gets a *draft*.

Judgement makes it right for the learner. The skill is doing *both, in a loop*.

Human-in-the-loop and ethical use: AI as a cognitive partner

Deliberate vs Incidental use

INCIDENTAL

“*Make this easier.*”

- Takes the first output as the answer
- No purpose, audience, or constraint stated
- Generic in → generic out
- Feels fast; quietly erodes judgement

DELIBERATE

“*Simplify this for an adult learner who reads Hangul, not Roman script, in week one.*”

- Treats the first output as a draft
- States purpose, audience, constraint
- Iterates and pushes back
- Judges output against what the learner needs

Three habits: *state purpose · iterate · pedagogical judgement*

Output, or relationship?



The one question we'll keep asking

Does this serve the learner and the relationship – or just produce a tidier output?

Pairing artificial with ancestral intelligence — the cultural knowledge and relational practice that keep learning human.



Why your judgement is the safeguard

AI is **sycophantic** by design. It will *validate a wrong answer* and tell you *your work is excellent* when it ***isn't***.





LIVE DEMO

Moving from incidental to deliberate AI use

Turning “make this simpler” into a structured lesson on chronic hypertension

- 1 State intent – what it’s for and who it’s for
- 2 Prompt with purpose, audience, constraint ([RISEN framework](#))
- 3 Push back once (revise, reiterate) — watch it cave (sycophancy)
- 4 Critically review the output

Is this right for my learner? Four checks.

The lesson looked great. Whether they're usable is a different question — the core of ethical & critical use.



1

Culturally right?

Does it assume a context your learner doesn't share?
Does it stereotype?



2

Accurate?

Did “simpler” quietly drop a fact that mattered — learner's background, their learning style?



3

Whose judgement?

Are you still the teacher, or is the tool deciding? Learner autonomy intact?



4

Private & safe?

What did you paste in?
Whose data is now on someone's server?

These four are your transferable checklist — regardless of discipline or topic

Guided practice

Pair work — one drives the device, one says the intent out loud. Saying the intent IS the skill.



Setup

Think of a task you give learners — a notice, instruction, worksheet. Picture the learner it didn't reach.



Generate

Prompt with purpose, learner, constraint (RISEN framework). Get the initial output but don't settle for it yet.



Run the 4 checks

Culturally-safe? Accurate? Whose judgement? Private? Push back where it fails — especially check #1.



Share

A few pairs: which check caught the most? The contrast does the teaching.



Decide

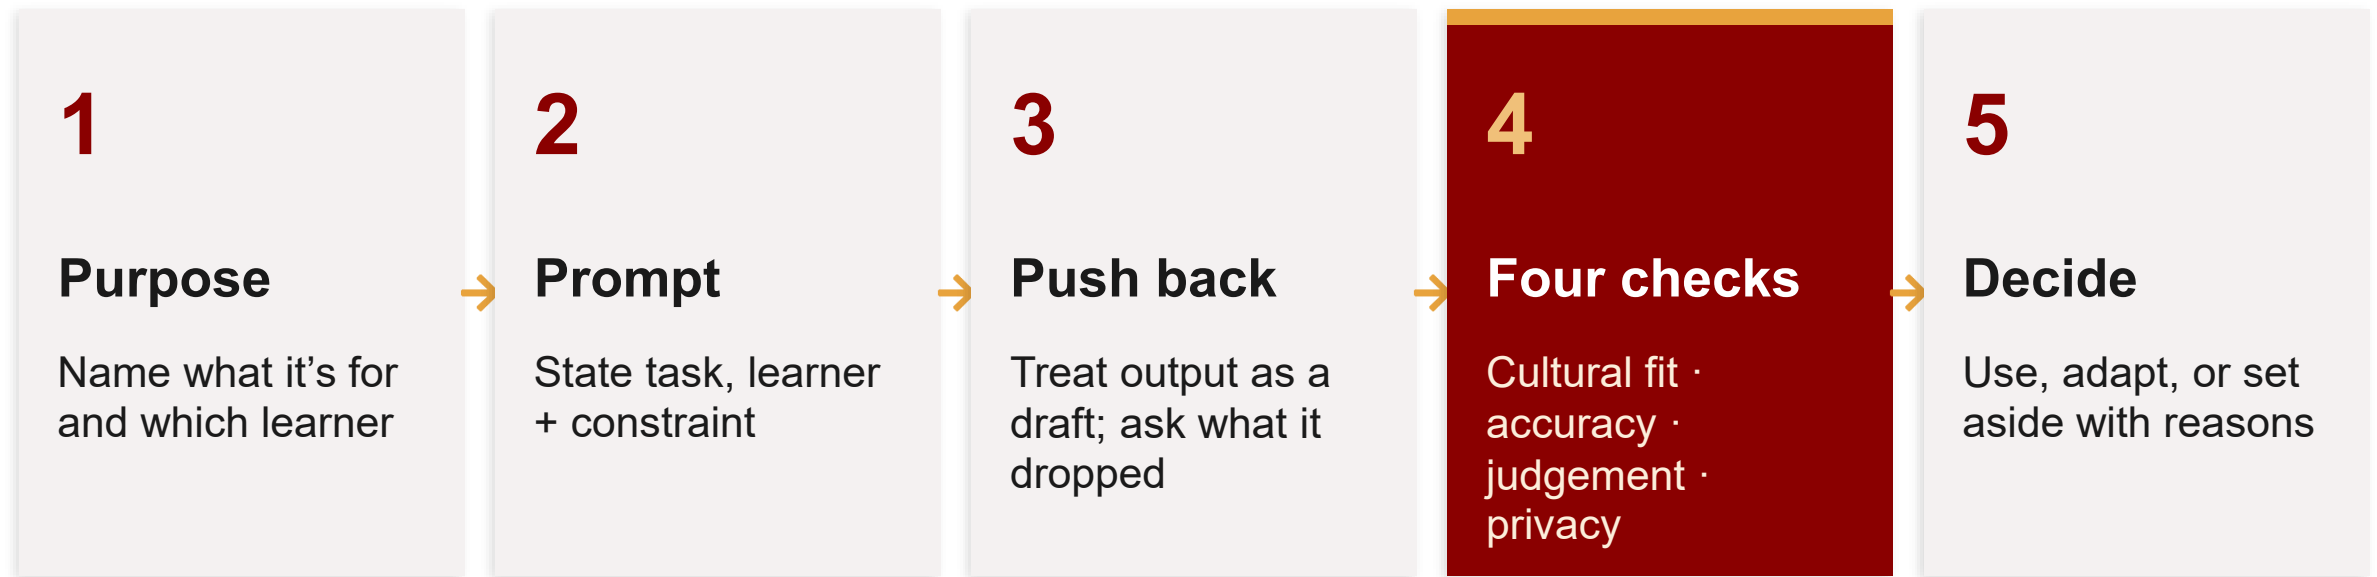
Use it, adapt it, or discard it. "Not at all" is a professional answer.

Need an idea?

Feel free to use the fallback notice on [the handout](#)

The workflow

*Deliberate use isn't a tool; it's a **repeatable method** turned into a **habit**.*



The four checks (step 4) are the ethics that ensure you stay in the 'pilot seat'.

The honest caveats

Equity

Free tools, uneven access, English-centric models. Who is designed out?

Over-reliance

If it erodes the judgement that catches the sycophant, it costs more than it saves.

When not to

Some learning is relational and slow on purpose. “Not at all” is a legitimate answer.

Beyond the hype means using AI deliberately – or choosing not to, with reasons.

HOW I USE A.I. TO MAKE DECISIONS



And that's not like

Q & A

Going forward, how will you be using AI?



Let's stay in touch!

Jia Rong Yap

JiaRong.Yap@englishlanguage.org.nz



Ngā mihi nui
Thank you for coming along to this workshop!





When you think about AI as "digital labor" in adult education, which statement best reflects your ethical priority?



Sam Altman recently stated that "AI got smarter than us" in terms of processing data and generating information. In our sector, what is the one human capability AI will never out-smart?