



# How **AI** is Changing the Way We Plan, Listen and Respond

Practical guidance for using AI across the engagement lifecycle without losing trust

This resource summarises key ideas from our webinar “How AI is Changing the Way We Plan, Listen and Respond” and the accompanying slide deck.

[communitilabs.com](https://communitilabs.com)

# Table of Contents

---

<b>1. Why AI matters for engagement right now</b>	<b>3</b>
<b>2. Reality check: when AI helps, and when it hurts</b>	<b>4</b>
<b>3. Set expectations: manual, augmented, automated</b>	<b>5</b>
<b>4. Phase 1 – PLAN: “Pre-sensing” community signals</b>	<b>6</b>
<b>5. Phase 2 – ENGAGE: chat, voice and 24/7 access</b>	<b>7</b>
<b>6. Phase 3 – REVIEW: “AI first pass, analyst-refined”</b>	<b>9</b>
<b>7. Phase 4 – FOLLOW-UP: closing the loop at scale</b>	<b>11</b>
<b>8. Guardrails: ethics frameworks you can lean on</b>	<b>12</b>
<b>9. A practical responsible-use playbook</b>	<b>13</b>
<b>10. Getting started</b>	<b>14</b>

---

# 1. Why AI matters for engagement right now

## Community and stakeholder engagement is under pressure:

- Consultations generate thousands of submissions that teams struggle to analyse in time.
- Language, literacy and access barriers mean important voices never make it into the room.
- Participants often never hear back—eroding trust and future willingness to engage.

At the same time, AI tools have moved from the lab to our daily work. Large language models (LLMs) can read, write and summarise at scale; chatbots can interact 24/7 in multiple languages.

The opportunity is real—but so are the risks. The goal is not to replace practitioners. It's to give them “copilot, not autopilot” support: AI does the heavy lifting; humans make the judgments.



The lifecycle diagram above is intended to serve as a guide to show how AI can be helpful and where to be cautious.

## 2. Reality check: when AI helps, and when it hurts

On slide 2, we contrasted two high-profile government examples.

### Scale done right | ATO “Alex”

- The Australian Taxation Office’s virtual assistant, Alex, has handled well over a million conversations.
- It focuses on a bounded set of frequently asked questions, drawing on curated and verified content.
- When questions are complex or outside its scope, Alex escalates to human staff.
- Result: reduced call volume and more staff time for complex, higher-value work.

### Risk without guardrails | NYC “MyCity”

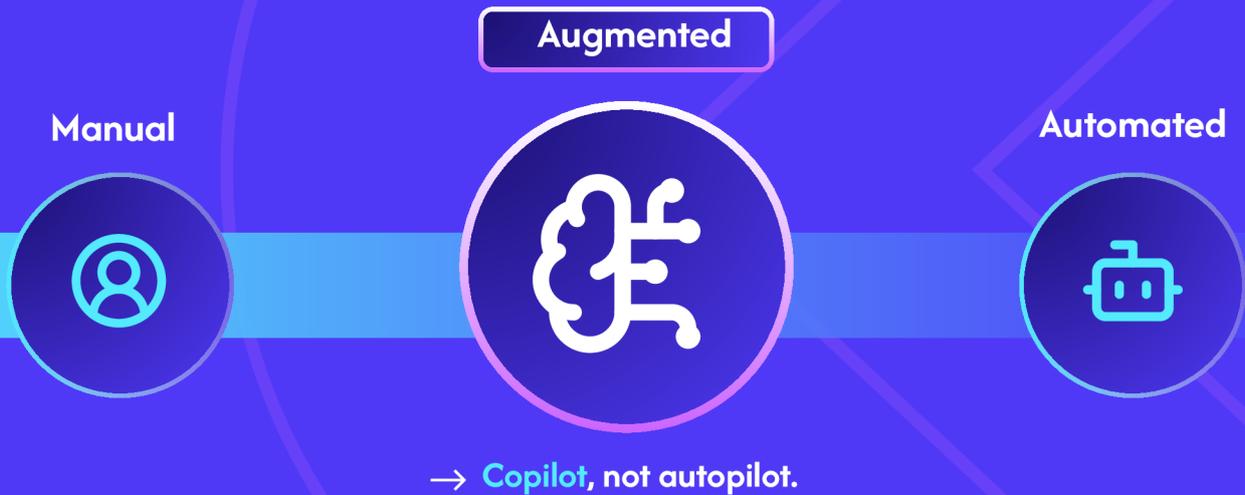
- New York City’s MyCity chatbot was launched with broad promises of AI-powered help.
- Independent testing found it giving incorrect—and in some cases illegal—advice on issues like dismissal and housing discrimination.
- Disclaimers were strengthened, but reputational damage was done and some residents may have acted on bad advice.

**Lesson:** AI is not “set and forget”. It must be scoped, tested and monitored, with:

- Bounded domains and clear policy constraints.
- Retrieval from verified sources—not the open internet alone.
- Refusal rules for topics it should not address.
- Clear escalation paths to humans.

# 3. Set expectations: manual, augmented, automated

Slide 3 shows a spectrum: **manual** → **augmented** → **automated**. The sweet spot for engagement is the middle.



- **Manual:** humans do everything. High fidelity, but slow and inconsistent at scale.
- **Automated:** the system acts with no human in the loop. Attractive on paper, but risky and ethically problematic for community input.
- **Augmented (“copilot”):** AI drafts, clusters and suggests; humans decide, validate and interpret.

## In practice, AI is:

### Good at

- Retrieval
- First-draft writing
- Sorting and clustering large volumes of text

### Weak at

- Local nuance
- Equity considerations
- Understanding when a small minority voice is critically important

For practitioners, the role shifts from doing all the manual processing to framing problems, triaging risks and assuring quality.

# 4. Phase 1 – PLAN: “Pre-sensing” community signals

The map visual on slide 4 shows topic “bubbles” like Housing, Economy, Anxiety and Rent across NSW local government areas, a snapshot of research that analysed millions of tweets during COVID.



## AI can:

- Scan public data (e.g. social media, forums, existing submissions) to detect emerging issues.
- Highlight geographic or demographic hotspots for particular concerns.
- Help you anticipate questions, identify overlooked stakeholders and target outreach.

## But there are important caveats:

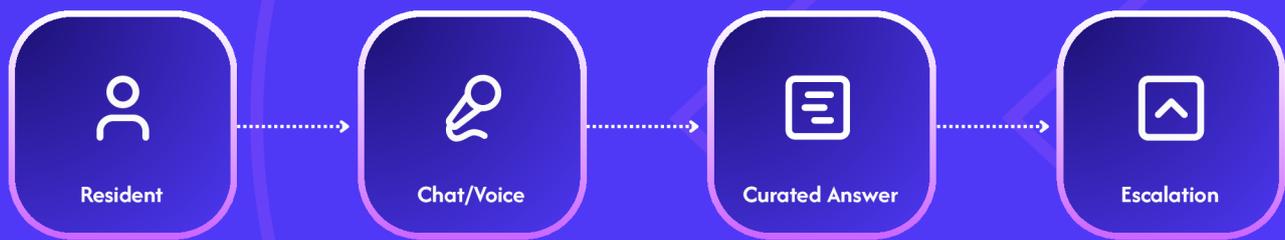
- Not representative by default. Social media over-represents younger, more online, more vocal residents. It under-represents older people, those without digital access and those who avoid public posting.
- Signal, not truth. Use these insights to shape your engagement questions and outreach, not to pre-decide policy.
- Ethics of public data. “Publicly visible” is not the same as “ethically free to repurpose”. Consider consent, sensitivity and expectations before analysing or storing posts.

## Practical uses in planning

- Use AI-assisted topic scans to inform your stakeholder and risk analysis.
- Equip engagement staff with likely issue areas before they go on-site.
- Prioritise translation and accessible formats where online signals show high concern.

# 5. Phase 2 – ENGAGE: chat, voice and 24/7 access

Slide 5 illustrates a simple flow: Resident → Chat/Voice → Curated Answer → Escalation, with notes on multilingual and 24/7 access.



Multilingual; 24/7; guardrails.

## AI assistants can:

- Provide round-the-clock responses in multiple languages.
- Meet residents where they are (web, mobile, messaging apps, voice).
- Handle routine questions so staff can focus on complex interactions.

Examples discussed in the webinar include Singapore's VICA platform and the [GOV.UK](#) Chat pilot, both of which blend rule-based tools with retrieval-augmented generation and clear escalation to humans.

# 5. Phase 2 – ENGAGE: chat, voice and 24/7 access cont.

But technology is only half the story. To make engagement chatbots safe and useful:

## Invest in content governance

- Ensure underlying web content and FAQs are accurate, up-to-date and written in plain language.
- Treat the knowledge base as a living product, not a one-off upload.
- Design refusal and escalation paths Explicitly block high-risk categories like legal advice, individual case assessments and employment decisions.
- Make it easy for the assistant to say “I can’t help with that—here’s how to contact a person.”

## Protect privacy

- Assume residents will sometimes paste sensitive or identifying information.
- Use private deployments; implement PII redaction and robust data governance, in line with OAIC guidance not to feed personal data into public AI tools.

## Be transparent

- Clearly label when people are interacting with a bot.
- Explain how their input will be stored, used and protected.

Done well, AI-assisted engagement expands access rather than replacing human contact.

# 6. Phase 3 – REVIEW: “AI first pass, analyst-refined”

Slide 6 compares an AI-generated cluster of raw phrases (e.g. “housing issues”, “rental prices”, “eviction”) with human-refined themes like Housing Affordability, Cost of Living and Tenant Rights.

↘ AI First Pass



↘ Analyst-refined



This is where AI can radically reduce time-to-insight while preserving human judgment.

## A practical seven-step workflow

1. De-duplicate submissions Merge near-identical responses (e.g. campaign emails) so they don't dominate analysis.
2. Cluster semantically similar responses Use AI to group comments by meaning, not just keywords.
3. Rank clusters by salience Identify clusters with the richest content or strongest emotion.
4. Extract representative quotes Pull direct quotations that illustrate each theme in participants' own words.
5. Generate theme summaries Ask AI to draft short descriptions of each cluster.
6. Human validation and refinement Analysts review and adjust themes, check for accuracy and re-label in policy-relevant language.
7. Minority voice check Deliberately search for small or outlier clusters that may represent marginalised or highly affected groups.

# 6. Phase 3 – REVIEW: “AI first pass, analyst-refined” cont.

## Key risks to watch:

- Hallucination: AI may assert themes or views that don't actually appear in the data.
- Majority bias: dominant views can drown out smaller but critical perspectives.
- Context collapse: local terms (“the transition”, “the reserve”) can be misinterpreted without human local knowledge.

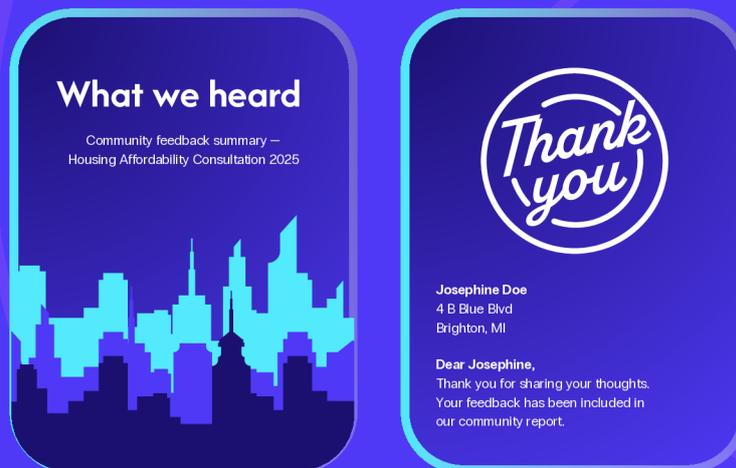
## To mitigate these risks:

- Ensure every AI summary is checked against raw submissions.
- Keep an audit trail of AI outputs and human edits.
- Build explicit checks for minority and vulnerable groups into your analysis process.

The UK's Consult/Humphrey pilots suggest that, with these safeguards, AI-assisted analysis can closely match human themes in a fraction of the time, freeing practitioners to spend more time on interpretation and follow-through.

# 7. Phase 4 – FOLLOW-UP: closing the loop at scale

Slide 7 shows two sides of the follow-up challenge: a “What we heard” report and a personalised thank-you letter to an individual resident.



## AI can support both:

- Reporting: Draft accessible “What we heard” summaries that reflect validated themes, include representative quotes and highlight how input will influence decisions.
- Personalised acknowledgements: Generate tailored thank-you messages that reference what each person contributed, not just their name.

## Additional possibilities include:

- Producing multiple versions of reports (e.g. easy-read, technical appendix, translated versions).
- Segmenting communications by stakeholder type or level of involvement.

However, transparency is key. As the slide footer notes, this should always be “AI-assisted drafting, human-reviewed.”

## Best practice:

- Clearly disclose when AI has helped draft a report or email.
- Ensure a named staff member reviews and signs off on all external communications.
- Give participants the option to query or correct how their input has been represented.

# 8. Guardrails: ethics frameworks you can lean on

Slide 8 summarises Australia's eight AI Ethics Principles, which offer a useful lens for engagement projects:

✔ Human, Social & Environmental Wellbeing

✔ Reliability & Safety

✔ Human-Centred Values

✔ Transparency & Explainability

✔ Fairness

✔ Contestability

✔ Privacy Protection & Security

✔ Accountability

**Human, Social & Environmental Wellbeing:** AI should enhance, not undermine, community wellbeing. Ask: does this tool genuinely improve participation and outcomes, or just make our internal process faster?

**Human-Centred Values:** Avoid using AI in ways that manipulate or suppress community voices.

**Fairness:** Test tools for differential performance and ensure minority voices aren't lost in automated analysis.

**Privacy Protection & Security:** Treat consultation data as sensitive. Don't put identifiable information into public AI tools; follow OAIC guidance.

**Reliability & Safety:** Pilot and test systems; monitor performance; have a plan to pause or roll back if issues arise.

**Transparency & Explainability:** Tell people when and how AI is used. Document methods in your reports.

**Contestability:** Give participants pathways to challenge or correct AI-assisted interpretations of their input.

**Accountability:** Assign clear human responsibility for each AI-supported process.

The OVIC case, where a child-protection worker used a public AI tool to draft a report with sensitive case details, is a stark reminder of why governance, not just enthusiasm, must drive adoption.

# 9. A practical responsible-use playbook

The final slide of the webinar condensed these ideas into an eight-point “Responsible-use Playbook”.

## Use it as a checklist before every AI pilot:

-  Disclose AI use – in engagement materials, on your website and in reports.
-  Label synthetic content – clearly mark any AI-generated text, images or simulations.
-  Protect privacy & data – minimise, secure and de-identify; avoid public tools for sensitive data.
-  Ensure human oversight – no auto-publish; keep humans in the loop.
-  Be transparent & explainable – document and share methods at an appropriate level of detail.
-  Ensure fairness & avoid bias – test across groups; look for who might be excluded.
-  Guarantee safety & reliability – test, monitor and have a kill switch.
-  Maintain accountability – name the decision-owner and provide contact points.

# 10. Getting started

## If your organisation is early in its AI journey, consider:

- Start small and low-risk – for example, AI-assisted drafting of internal summaries or comms, with strong human review.
- Pilot one lifecycle phase at a time – e.g. use AI in the Review phase for a single consultation and compare results with manual analysis.
- Create a lightweight AI register – track where AI is used, data flows, approvals and outcomes.
- Educate leadership – frame “copilot, not autopilot” as the risk-aware path, using examples like NYC MyCity and OVIC to show what happens when guardrails are missing.
- Co-design safeguards with your community – ask participants what they’re comfortable with and how they’d like AI use to be disclosed.

AI won't fix weak engagement practice—but, used well, it can help you plan more intelligently, engage more inclusively, review more consistently and follow up more meaningfully.

Visit [communitilabs.com](https://communitilabs.com) for more information