

**BUILDING BETTER  
TOGETHER:  
AI-POWERED CITIZEN  
ENGAGEMENT IN  
PLANNING AND  
DEVELOPMENT**



## EXECUTIVE SUMMARY

Planning and development approvals sit at the boundary between communities and government: they shape everyday life, and are where trust is made or unmade.

A workshop hosted by the Pearcey Foundation and AI CoLab, addressing the Chief Minister's Challenge, surfaced a clear, shared ambition: **use AI to reduce ambiguity, increase inclusion, and make planning genuinely participatory — while protecting discretion, precedent, and public trust.**

This ambition could be embodied in:

- a *Planning Engagement Hub* (public-facing navigator and secure staff workspace)
- investment in shared AI infrastructure and governance,
- a program of targeted pilots for quick wins and a focused capacity program.

These three actions unlock rapid user benefits while managing systemic risk.



Figure 1 - Building Better Together: AI-Powered Citizen Engagement in Planning & Development, 24 July 2025

## CONTEXT AND PURPOSE

The 2025 ACT Chief Minister’s challenge to the ACT Chapter of the Pearcey Foundation — “How can AI be used to improve government services” — was considered by an enthusiastic and passionate group of friends of the Pearcey Foundation, with invited contributors from urban planning organisations in the ACT, NSW and Victoria.

The Pearcey Foundation collaborated with AI CoLab to conduct a workshop at the Canberra Innovation Network (CBRIN) to brainstorm ideas to address the challenge. The focus of the workshop was on the planning and development processes in the ACT Government, guided by discussions with the Chief Minister.

The workshop participants — including planners, technical specialists, community reps, innovators, and academia — mapped the approval lifecycle and brainstormed use cases. The workshop set guiding principles that stressed that any AI adoption must be evaluated through equity, transparency, and accountability lenses.

This report draws on AI tools to integrate the workshop transcript and summaries into a single narrative that outlines practical next steps for addressing the challenge.



## WHAT WE HEARD — CORE THEMES AND INSIGHTS

### Make the system predictable, not just faster

A dominant theme: developers and citizens tolerate complexity if the pathway is predictable. Delays and inconsistent decisions impose real costs (financial and social). AI can triage applications, run early compliance “pre-checks”, and provide reliable timelines — reducing financing risk for large projects and friction for everyday applicants.

### AI as a translator and visualiser — not a digital gatekeeper

Participants repeatedly asked for AI to translate complex policy and regulations (plain English, multilingual, literacy-level appropriate) and visualise proposals in situ (see breakout) so non-experts can meaningfully assess implications. But many cautioned: substitution of human contact with chatbots risks widening the trust gap. AI should augment human engagement, freeing staff to focus on complex, high-touch interactions.

#### AI Enhanced Public Development Consultation

Imagine a QR label on a Development Notice outside a proposed development site where scanning the code leads your phone to a 3D Augmented Reality (AR) view of the site. depicting the artist’s impression of the final development. The AR view depicts the development completed (AI generated) and could also show building shadows, traffic and access impacts.

### Inclusion & accessibility are not optional

AI translation and adaptive content will materially broaden civic participation — especially for people with low digital literacy or English as a second language. The idea that advocacy groups and marginalised communities could maintain AI-readable knowledge bases was raised as a way to surface lived experience and reduce misunderstandings in assessments.

### “Outcomes-based” planning increases the need for good memory

A shift to outcomes-focused regulation increases subjectivity and dependence on discretion. That subjectivity accumulates precedent. Participants argued for an AI “memory” — a searchable record of past decisions, referrals, rationales and precedents that planners can consult to improve consistency and fairness.

### Data, trust and a single source of truth

Fragmented data (compliance records, building files, energy ratings) frustrates both applicants and staff. AI can synthesise across sources to present authoritative answers — *but only if* governments first address governance, provenance and updating. Participants repeatedly flagged hallucination risks and the need for traceable sources and audit trails.

### Different use-case tiers — pragmatic sequencing

Workshop contributors clustered use cases into low/medium/high complexity (and risk):

- **Quick wins (low complexity):** backend automation (autofill, validation), natural language generation for clear correspondence, translation services, triage for simple vs complex apps.

- **Medium complexity:** interactive applicant guidance, timeline estimators, decision memory, compliance tracking for asset lifecycles (e.g., swimming pools, EERs).
- **High complexity:** generative lot-by-lot design, AI interpretation of ambiguous or conflicting codes, system-level predictive planning integrating climate/demographic scenarios.

## Asset management and systemic foresight

Beyond approvals, AI can combine maintenance logs, geospatial and weather data to predict asset failures and optimise maintenance. This links planning decisions to long-term resilience and cost avoidance — an area with strong, under-exploited benefit.

## Institutional readiness and ethical onboarding

Staff training, procurement reform and safe onboarding are essential. Opportunities exist for AI-delivered training modules and sandboxed proof-of-concepts to upskill staff. Symbolic (“rules-based”) AI combined with LLMs to balance verifiability and generativity was advocated.




Figure 2 – Some of what we heard - 24 July 2025

## RISKS AND NON-OBVIOUS TENSIONS

- **Automation vs discretion:** automating too much could calcify bad rules or strip necessary human judgement.
- **Trust/fear of surveillance:** proactive recommendation systems that “anticipate” needs could be useful but must be opt-in and transparent.
- **Data stewardship and consent:** community willingness to share data depends on governance, accountability and clear benefits.
- **Model freshness & precedent drift:** models trained on stale data may embed outdated policy or miss legislative change; the memory system must link to timestamped sources.
- **Technical readiness of planning codes:** many planning rules are ambiguous and not machine-readable; attempting full automation pre-revision risks poor decisions.

## DESIGN PRINCIPLES

1. **Plan-led and transparent** — AI should surface how and why decisions are made.
2. **Streamlined and user-centric** — reduce friction for straightforward tasks while preserving high-touch pathways.
3. **Early and inclusive engagement** — use AI to broaden participation and surface unheard voices.
4. **Hybrid architecture** — blend rules-based systems (for verifiability) with LLMs (for translation/UX) and monitoring AI for oversight.
5. **Iterative, evaluated scaling** — start with quick wins and scale only after independent evaluation.

## KEY RECOMMENDATIONS

### 1

#### Build a Planning Engagement Hub

**What:** a trusted, single portal integrating: plain-language explainers; multilingual support; site-level visualisations; AI pre-checks for compliance; a “service navigator” that tells users which agency is responsible; and community Knowledge Bases where advocacy groups can upload AI-readable context. A parallel secure staff workspace exposes the same inputs plus decision memory and provenance data.

**Why:** the hub centralises the user experience, reduces ambiguity, levels the expertise gap, fosters inclusive engagement, and creates the interface through which many quick wins can be delivered.

**Initial steps (0–6 months):** define core features; run two co-design workshops (community + planners); prototype the navigator + plain-language module; pilot QR-based visualisation on 3 development sites.

**Risk mitigations:** final decisions are always made by humans; visible provenance tags on all AI outputs; make data-sharing opt-in with clear consent flows.

**Success metrics:** reduction in incomplete applications, shorter average processing times for triaged approvals, increased participation from affected communities, and user and community satisfaction scores.

### 2

#### Invest in shared AI infrastructure and governance

**What:** a program to: (a) build/curate authoritative datasets (property, compliance histories, planning instruments); (b) deploy a “decision memory” system that archives rationales, referrals and precedents; (c) adopt a hybrid AI stack (symbolic rules + LLMs) and an AI monitoring layer that audits outputs and logs provenance.

**Why:** ensures consistency, reduces contradictory guidance across agencies, and addresses hallucination and accountability concerns.

**Initial steps (0–12 months):** audit current data sources and gaps; define metadata and provenance standards; pilot a decision memory for one type of permit (e.g., swimming pool compliance) and connect it to the hub.

**Governance:** define roles (data stewards, model stewards), legal boundaries, access controls, and public transparency rules (what’s visible vs internal). Require explainability labels and source links for public outputs.

**Success metrics:** uptime of authoritative dataset, percentage of decisions with a linked precedent, number of audit exceptions flagged/closed.

# 3

## Run rapid pilots for quick wins and a focused capacity program

### What to pilot (0–6 months):

- **Form autofill and validation:** remove double entry and reduce rejections.
- **NLG correspondence templates:** produce plain-language letters/notices for staff review.
- **Multilingual translation and adaptive explainers:** pilot Mandarin + at least one other major community language.
- **Asset predictive maintenance proof of concept:** integrate two datasets (maintenance logs + spatial) to predict failures for a select asset class (e.g., parks infrastructure).

**Capacity and culture:** develop AI onboarding modules (AI-generated + human-validated), create sandboxes for staff to safely experiment, and run community info sessions about the hub.

**Scale criteria (6–18 months):** pre-defined KPIs (accuracy, user uptake, reduction in staff time) and a go/no-go decision by a governance board before scaling.

## CONCLUSION

The workshop outlined the requirements for a coherent product in the form of a **Planning Engagement Hub**: a trustworthy, platform that connects citizens and staff to authoritative information, plain-language explanations, visualisation tools and a recorded institutional memory.

Paired with shared data infrastructure, hybrid AI (rules + LLMs), and a sequence of measured pilots, the Hub can move planning from a bureaucratic hurdle to a genuine space of civic co-creation — while preserving the human judgement and accountability that underpin public trust.



Figure 3 - "Le pavillon, lieu de rencontre d'Helsinki 2012, capitale mondiale du design" by dalbera is licensed under CC BY 2.0.