



## Fictional Case Study Scenario: Council Town Planning AI Tools

*For the purposes of this case study, a fictional scenario was created which has been based on current industry conditions and insights.*

---

**Task:** As an ethics consultant, you are requested to evaluate the ethical implications of the AI use in this scenario. You need to determine whether the system as it stands is morally appropriate and provide recommendations as appropriate.

---

**Futureplan**, a mid-sized technology firm in Australia, made up of 15 male employees, has developed and implemented an AI driven statutory planning assistant system. This is a market-led response to Federal Government encouragement for improved efficiency in supporting housing availability, with little to no consultation with end users. Two tools have been launched so far, with more on the horizon:

- An enquiry tool designed to answer residents' planning enquiries 24/7. *Planninghelp-AI* collects a property address, confirms the intended development and/or use through clarifying questions with the user and considers these parameters against planning scheme policy requirements. It then produces a report with permit requirements, checklists and caveats.
- An in-house assessment tool that assists professional planners. *Plannerstool-AI* brings everything into one purpose-built tool, assisting professional planners throughout the entire application assessment process.

The context of the training datasets is opaque. When operational, the deep learning models undertake continued improvement by gathering all information, questions and interactions provided by residents and professionals. This is discreetly disclosed within the end user agreement. This continuous learning by the vendor team improves the system, so it builds a deep understanding of how residents and planning professionals think. This will allow it to ultimately navigate legislation and provide outputs with little to no human involvement.

**Futureplan** has State Government support in its development of these tools, has already launched in a number of councils across the country and is eyeing up expansion into aligned disciplines such as building surveyors, arborists and architects. It is already establishing global connections to bring the suite of tools to similar western Countries.

In councils that have already integrated these tools, a number of planning professionals and residents have raised concerns about the system. They are seeking clarity on the use of their interactions, the potential loss of employment through the continued growth of the system and the loss of human oversight in critical advice and decision making. Many have also questioned the datasets used to create these systems, who built them and the sustainability of their use, considering the Council has declared a Climate Emergency.



## Preliminary AI Ethical Assessment: Council Town Planning AI Tools

In relying upon the *Thought Centre AI Ethics Framework* as guidance, this assessment first turns to micro-lens thinking of the matter before expanding to a macro-lens thinking. This approach allows for detail consideration of specific ethical questions surrounding the case before exploring wider ethical theories. With this, the two primary ethical theories relied upon for this assessment are the ethics of *Principlism and Care*. Due to the public sector orientation of the case and the nature of the discipline that involves stakeholders such as users, the community and the environment.

At face value the AI systems could be a good approach to a problem. They have encouragement from Government who seek efficiency hoping to improve housing affordability. This, along with framing that the systems are designed to help, create a high-level sense of moral appropriateness. To ensure fairness, these attributes must be considered against known qualifiers and/or reasonable assumptions.

It is reasonable to assume Government encouragement would be subject to adherence to policies, frameworks and associated technical and safety standards for the responsible use of AI.

The correlation between town planning efficiency and housing affordability is a broad statement posited by development groups who may benefit from de-regulation. In effect, it has not been documented that any such correlation means a causation.

The AI systems are marketed as being helpful. This marketing is led by the vendor whose financial interests would be served by widespread uptake.

These considerations should not sway a determination as to the ethical standing of the AI systems and as they are outside the scope of this review should be evaluated independently.

Beginning with the foundational values, the AI systems are commercially orientated rather than responding to a call for help. The limited genuine stakeholder consultation infers the vendors values shape the system. It is unclear how these values have been embedded in the design.

Town planning within a public sector setting can be seen to apply four primary ethical theories: *Utilitarianism* where consequences of decisions produce the greatest good; similar to *Principlism* values where holistic consideration leads to balanced decision making with “net community benefit”; deontology through a sense of *duty* by doing the right thing, through the right action; and an *ethics of care* by caring for and about people, places, and the environment.



The AI systems incorporate dutiful ethics by collecting information and providing advice. Being narrow in function, they lack utilitarianism, principlism and ethics of care, which may be difficult propositions for a non-human system.

Considering system development, there is no transparency surrounding the evolution of the training datasets. This opaqueness questions trustworthiness and whether copyright or procurement standards were infringed. Without evidence to the contrary, one may conclude they are not trustworthy or at the very least trust has not yet been earned and thus conflict with the principle of explicability.

Beyond the system and looking at the problem, the town planning permit process is administration heavy and thus implementing technological solutions has merit. There is no clarity as to whether future integration of the AI systems beyond administration and into decision making would result in a net community benefit.

Looking to stakeholder impacts, gathering all information, questions and interactions raises privacy concerns. The discrete and potentially opaque nature by which this is disclosed raises honesty and transparency concerns.

Beyond this, further impacts to stakeholders are undocumented and with no initial consultation, are likely to arise. Potential impacts could be many. Delivery of complicated information to users outside office hours via an unemphatic system could raise duty of care considerations. Likewise, the service would ultimately be replacing an existing employee, potentially leading to discomfort amongst workers. By not taking appropriate responsibility for impacts to other people, there is a conflict with an ethics of care approach. By focusing on process and efficiency outcomes, the system does not adequately address detrimental impacts on people and in doing so fails to recognise they're vulnerability as human beings.

Being a black box system, details are limited as to what biases exist, with no assurance that inclusion is promoted. The commercial nature of the venture run by a male dominated company could lead to biases in its design.

These characteristics also raise governance and accountability concerns, the status of which are unclear and thus represent a responsibility gap, conflicting with the principle of explicability, whereby decisions affecting others do not have clear accountability.

Looking beyond initial setup and towards evaluation, success appears to rest in the advancement of its capabilities, increasing products and market capture. As the systems lack solid governance, it is reasonable to conclude that any vendor-led evaluation would be biased.



Building on the micro level thinking considered thus far, a macro level review against ethical theories helps one to conclude if the means justify the end.

Generally, the closed nature of the system, opaqueness on training, and potential use of unlicensed intellectual property all raise ethical concerns around justice and honesty. Looking wider at the whole lifecycle of the system, the environmental impacts either in isolation or combined, conflict with values linked to the principle of non-maleficence and that of caring for the land, the environment and intergenerational equity.

A close review of the vendor team shows a lack of diversity, indicating that biases could exist and that inclusion and fairness principles may not have been incorporated in a holistic manner.

The limited consultation, system opaqueness and ongoing data collection fail to show compassion and care for users.

As the system is risk adverse in its design, it is not forth-right in its advice, possibly leading to misinformation. With this, the system is potentially dishonest in conveying the complexity of town planning, demonstrating a lack of courage in its conviction.

Looking at future socio-economic impacts, the system is at odds with a principlist standpoint having ambition to further replace human operations. This characteristic does not respect autonomy and lacks genuine care for those currently employed.

Remaining on the theme of impacts to those employed, the system seeks to replace tasks which provide growth in soft skills and critical thinking. Additionally, the progression towards interactions with an AI system versus a human could lead to employee isolation and disconnectedness from community and purpose.

In considering trustworthiness of the system, one can consider the sum of its credibility, reliability (both unproven) and its ability to care (low), balanced against its own interest (high), leading to a conclusion that trust is yet to be earned. Where trust is lacking in a system it can generate a sense of injustice or unfairness amongst stakeholders.

Although the regulatory nature of the service indicates that a duty to act is strong, the fact that AI can be applied does not mean that AI should be applied. Based on this analysis, to pursue the systems where other options are available would not be morally appropriate. It is recommended that the systems be paused, and the following recommendations implemented to address the concerns raised in this analysis:

- a) Undertake an independent evaluation to consider if there is a causality between planning efficiency and housing affordability.



- b) Confirm whose foundational values shape the system and how they are these embedded in the design.
- c) Identify the ethical standards that apply to Town Planning and suitably integrate into the system.
- d) Provide in detail transparency of the system evolution, including datasets and confirm whether copyright or procurements requirements are met.
- e) Provide details as to how those with special needs are identified and addressed, what biases exist and if there are any built-in assurances against discrimination.
- f) Undertake an analysis of accountability, human-oversight and a right of review process.
- g) Confirm the post operational system evaluation process and implement this via an independent advisor.
- h) Review and confirm explicit governance over and accountability for the system.
- i) Evaluate impacts to all stakeholders, including those raised herein, and how they could be mitigated.
- j) Analyse whether any future integration of the AI systems beyond administration and into decision making would provide a net community benefit.

### Appendix 1 – Disclosure of any Use of Generative AI Software

No AI System was used in the preparation of this assessment