AI Governance Policy

Dunlin Housing Limited - Version 1.2 - March 2025 - Review due September 2025

1. Purpose

This policy sets out the principles governing Dunlin Housing Limited's use of artificial intelligence in its governance function. It applies to all AI tools used by the board, the governance and regulation team, and any function operating under the board's authority.

The policy exists because Dunlin uses AI tools to inform the decisions it makes - about the condition of its properties, the experience of its residents, the effectiveness of its operations, and the quality of its governance. Those tools create obligations. This document is the record of what those obligations are.

2. Governing Principle

Artificial intelligence does not make governance decisions at Dunlin. It informs them. There is a category of decision - one that requires a person to stand behind a conclusion with moral accountability - that cannot be delegated to an automated system without that decision ceasing to be what it is. Every principle in this policy is an expression of that commitment.

The board is not an oversight body that sits above the AI tools and reviews their outputs at quarterly intervals. The board is the part of the governance architecture responsible for the judgements that require human authority. That is a more demanding role than the conventional governance literature tends to assign, and this policy is designed to support it.

3. The Nine Principles

The principles below govern every AI tool Dunlin uses in its governance function. They are expressed as active commitments rather than guidance, and each is applied to every substantive output the system produces.

Principle 1 - Human authority is non-delegable

Decisions carrying regulatory, fiduciary, or ethical weight are human decisions. AI tools can synthesise evidence, surface options, and recommend a course of action. They cannot make the decision. This applies without exception and is not modified by the quality or confidence level of the AI output.

Principle 2 - The sensing boundary runs asset-outward, not people-inward

AI tools monitor properties, processes, and operational conditions. They do not monitor residents as individuals, except where a resident has specifically and revocably consented to receive AI-assisted support. The boundary between monitoring assets and monitoring people is a governance boundary, not a technical one, and it is maintained by policy rather than by the limits of the technology.

Principle 3 - Consent is specific and revocable

Where AI tools engage with resident data in a support context, consent is given for a defined purpose and can be withdrawn at any time. Withdrawal of consent does not affect the resident's relationship with Dunlin or the services they receive. Consent architectures are session-specific; general consent to data processing does not constitute consent to AI-assisted engagement.

Principle 4 - Every AI output is traceable

The reasoning behind every AI recommendation must be available for human review. A recommendation that cannot be explained cannot be acted on. Traceability is a condition of use, not an optional feature. Where an AI tool cannot provide an explanation for an output, that output is treated as void for governance purposes.

Principle 5 - Board overrides are governance events, not system failures

When the board or executive reaches a decision contrary to an AI recommendation, the decision and its reasoning are recorded. That record is evidence that human judgement is genuinely exercised within the governance process - not evidence of tension between the board and its tools. The audit trail of board overrides represents the governance architecture operating as designed. The Governance and Regulation function provides a summary of override decisions to the Audit and Risk Committee annually.

Principle 6 - The governance function holds the organisation's purpose

Questions about what Dunlin is for and on whose behalf it operates are not answered by AI. They are held by the governance function and, through it, by the people the organisation serves. That question is not one the system is designed or permitted to resolve. An organisation can reconfigure its operational tools and remain what it is. What it cannot do is configure away its answer to why it exists.

Principle 7 - Welfare signals have defined, timed routing to human review

Where the system generates a signal indicating a potential welfare concern, the path to human assessment is pre-defined, timed, and non-deferrable. Welfare signals are not queued behind maintenance priorities or subject to operational discretion about when to escalate. The maximum permissible interval between signal generation and human acknowledgement is defined in the operational procedures that accompany this policy. This principle was added in March 2025 - see Section 5.

Principle 8 - Accountability extends to organisational boundaries

AI tools that operate across organisational boundaries - with contractors, managing agents, or data-sharing partners - are subject to the same governance principles as internal tools. Accountability does not stop at the legal boundary of the organisation. Contracts with partners who receive or generate data through Dunlin's AI systems include provisions requiring compliance with these principles.

Principle 9 - This policy learns from operational experience

The principles in this document are designed to grow. When operational experience identifies a failure mode the governance logic has not anticipated, the policy is updated and the reason is recorded in Section 5. A principle added in response to an identified failure is not evidence of weakness in the governance architecture - it is evidence that the architecture is functioning as intended.

4. Review and Assurance

This policy is reviewed annually by the Governance and Regulation function and approved by the board. The board receives a summary of AI tool usage, any incidents or near-misses, and the status of any open actions at each quarterly meeting.

The Audit and Risk Committee has oversight of this policy and receives a full review report annually. Any material change to the principles - including additions prompted by operational experience - is approved by the full board before the revised policy is published.

This policy is published on the Dunlin Housing website. Publication is not a communications exercise. It is an accountability commitment: the organisation’s statement of how it uses AI tools is available to residents, regulators, and funders without request.

5. Version History

Principle 9 commits Dunlin to recording why this policy changes, not only what changed. The entries below are that record.

Version 1.2 - March 2025

Principle 7 expanded and two supporting operational commitments added following a review in Q1 2025. The review was prompted by an incident in February 2025 in which a welfare signal generated by the property sensing system was not routed to human review within the organisation's target response window. The resident involved had been without adequate heating for three weeks before the gap was identified.

The review found three points in the governance logic where routing requirements were either undefined or subject to operational discretion: the classification threshold for welfare signals, the escalation path when the primary reviewer is unavailable, and the maximum interval before a signal is automatically elevated to board-level attention. All three were addressed in this revision. The operational procedures accompanying this policy were updated at the same time.

Version 1.1 - November 2024

Minor amendments following the expansion of the Board Portal to include additional AI advisory tools available to board members. No changes to governing principles. References to the scope of application updated to reflect the expanded tool set.

Version 1.0 - September 2024

Initial publication. Approved by the board at its meeting on 18 September 2024.