

BrinkByte AI Ethics Statement

BrinkByte is committed to building AI systems that are responsible, transparent, and deliver real-world value without compromising ethical integrity. Our approach is anchored in the Australian Government's AI Ethics Principles, while remaining practical and innovation-focused.

We adopt these principles as a guide — not a constraint — and actively interpret them to support the development of fast, fit-for-purpose computer vision systems that serve communities, enterprises, and governments across Australia.

1. Human, Social and Environmental Wellbeing

We design our AI systems to improve outcomes for people and the communities they live in. We evaluate use cases for their potential to enhance wellbeing, equity, safety, and environmental impact — especially in public infrastructure, safety, and planning contexts.

2. Human-centred Values

BrinkByte's products are built to respect dignity, autonomy, and the rights of all individuals. We embed human oversight and accountability throughout our model lifecycle and never use CV in applications that violate fundamental rights or freedoms.

3. Fairness

We aim to minimise bias in training data, model outputs, and downstream decisions. When working with demographic or sensitive data, we apply fairness checks and avoid outputs that could lead to unjust outcomes for individuals or groups.

4. Privacy Protection and Security

We ensure that privacy is considered in every deployment — from how footage is collected and stored, to how inferences are generated. All data is handled under secure protocols, and our models are designed to avoid storing or identifying individuals unless explicitly required and approved.

5. Reliability and Safety

Our CV systems are stress-tested under a variety of conditions — including changes in lighting, weather, and occlusion — to ensure consistent performance. We monitor accuracy and false-positive rates in real-world environments, and we design for fail-safes where appropriate.

6. Transparency and Explainability

We provide customers with clear, plain-English documentation on how our models work, what data they rely on, and how they make decisions. Our detection and alerting logic is explainable and auditable — not a black box.

7. Contestability

We welcome questions and challenges to our models' outputs. Our customers retain the ability to audit outputs, adjust thresholds, and re-label data to inform retraining. BrinkByte provides tools and support to enable continuous validation.

8. Accountability

We take full responsibility for the models and pipelines we deploy. All AI development at BrinkByte follows a structured lifecycle, including validation, documentation, and versioning. We respond quickly to issues and retrain models when performance drift is identified.

Version 1.0 – May 2025