

The Paramount AI Manifesto

A declaration of intent for responsible, ethical, and human-centred AI



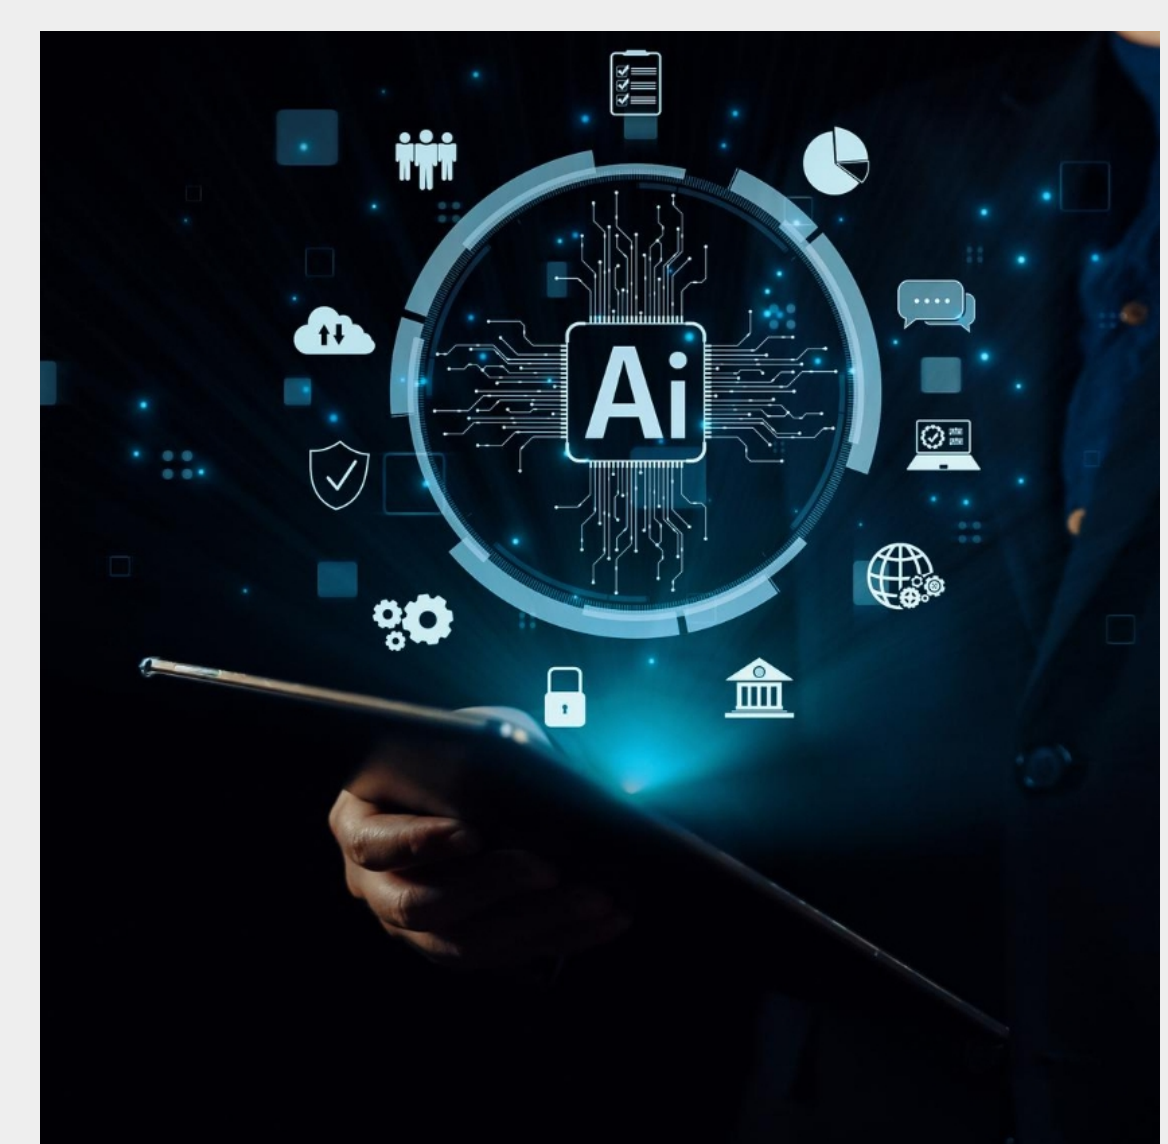
At Paramount, AI is not adopted as a trend or convenience. It is managed as a strategic capability, subject to disciplined judgment, oversight, and care. This Manifesto defines how we design, deploy, and govern AI across operations. It is not a technical framework, but a statement of intent.

AI at Paramount is a governed strategic capability.

Our Belief

We believe sustainable progress sits at the intersection of innovation and responsibility.

- AI should strengthen human capability, not diminish it.
- Security must be treated as a foundational requirement, not a later control.
- Intelligence deployed without accountability is unacceptable.
- Privacy is a fundamental right, not a procedural exercise.
- Innovation that disregards fairness results in harm and loss of trust.
- Scale must never outpace safety, dignity, or public confidence.
- AI systems that touch people or their data must never operate without their informed agreement.
- AI must be grounded in accurate, high-quality, contestable data,



Our Principles


1. Human Accountability Remains Central

AI augments decision-making but does not replace human judgment. Accountability for decisions, outcomes, and consequences always rests with people.



2. AI Is Deployed With Clear Intent

AI is implemented deliberately, with a defined purpose, clear boundaries, and continuous oversight throughout its lifecycle.



3. Impact Must Be Demonstrable

Every AI system must deliver measurable and meaningful value for customers, employees, or the organisation, supported by transparent evaluation.



4. Work Should Become More Meaningful


AI is used to reduce friction and low-value effort, enabling people to focus on judgment, creativity, and collaboration while preserving agency and trust.



Our Boundaries

There are clear limits to how we use AI. We will not:

- Use AI to manipulate behaviour or bypass informed consent.
- Allow AI to make consequential decisions affecting individuals without meaningful human oversight.
- Collect, use, or retain data without transparency, necessity, or lawful provenance.
- Deploy AI that impacts vulnerable groups without prior impact assessment.
- Operate AI systems without clearly defined ownership for risks and outcomes.
- Prioritise speed, scale, or competitive advantage at the expense of safety, fairness, or trust.
- Deploy AI in environments affecting people or trust without appropriate observability.
- Maintain AI systems that are unsafe, obsolete, or no longer trustworthy.
- Obscure failures through complexity, third-party dependencies, or lack of transparency.



Our Commitments

- Secure by design
- Auditable by default
- Explainable to the greatest extent technology permits
- Aligned with human agency
- Resilient to misuse and abuse
- Fair, proportional, and impact-aware
- Transparent in data origin and model behaviour
- Assigned clear ownership for risk and outcome
- Monitored continuously for performance drift and unintended impact
- Aligned with applicable local and regional laws and regulations
- Retired responsibly when safety or trust can no longer be assured

Our Promise

AI systems will not always perform as intended. Organisations, including ours, will also make mistakes. Our commitment is not to perfection, but to integrity – acting responsibly, correcting issues promptly, and learning continuously. We are working toward a future in which advanced AI systems are trusted, well-governed, and worthy of confidence.

Premchand Kurup
Chief Executive Officer