

THE  
HUMAN  
SOVEREIGNTY  
CHARTER  
FOR ARTIFICIAL  
INTELLIGENCE

A CONSTITUTIONAL FRAMEWORK  
FOR HUMAN-CENTRED GOVERNANCE OF AI

*Drafted by*

**ADAM TUGWELL**

Cheltenham, United Kingdom

2026

# The Human Sovereignty Charter for Artificial Intelligence

A Constitutional Framework for Human-Centred Governance of AI

# Copyright

© 2026 Adam Tugwell

Cheltenham, United Kingdom

All rights reserved. This document may be reproduced, shared, and referenced freely for educational, ethical, governance, and community purposes, provided that no modifications are made to the core Charter text and that attribution is retained.

No part of this publication may be used for commercial purposes, incorporated into proprietary systems, or presented as statutory law without explicit permission from the author or relevant governance bodies.

This Charter is offered as a public, non-statutory framework for the ethical governance of artificial intelligence. It is designed to support voluntary adoption, community oversight, and democratic participation. It does not constitute legal advice or regulatory directive.

First published in 2026

Printed in the United Kingdom

## Contents

Copyright .....	3
Dedication.....	7
Epigraph.....	8
Foreword.....	9
Disclaimer.....	10
How to Use This Charter .....	11
Scope and Applicability.....	14
Relationship to The Local Economy & Governance System (LEGS).....	16
Use Within the Current System .....	17
Rationale and Evidence Base .....	20
Human Sovereignty and Agency .....	20
Limits of AI Knowledge and Capability .....	20
Risks of Concentrated Power .....	20
Transparency and Accountability Failures .....	21
Threats to Human Capability and Learning.....	21
Safety and Infrastructure Vulnerabilities .....	21
Protection of Belief, Conscience, and Identity .....	22
Need for Democratic and Community-Centred Governance .....	22
Summary of Obligations .....	23
Obligations for Developers and Designers.....	23
Obligations for Organisations and Businesses.....	23
Obligations for Public Institutions and Service Providers .....	23
Obligations for Educators and Educational Institutions.....	24
Obligations for Operators and System Owners .....	24
Obligations for Governance Bodies (including LEGS) .....	24
Rights of Individuals and Communities.....	25
Right to Human Authority and Decision-Making.....	25
Right to Transparency and Understanding.....	25
Right to Human Control and Manual Override .....	25
Right to Protection from Exploitation and Manipulation.....	25
Right to Fairness and Non-Discrimination .....	26

Right to Human Learning and Capability .....	26
Right to Meaningful Work and Economic Dignity .....	26
Right to Community Oversight.....	26
Right to Redress and Remedy.....	26
Right to Protection of Belief, Conscience, and Identity .....	27
Foundations of the Charter .....	28
Executive Summary.....	29
Foundational Principles of Human-Centred Artificial Intelligence .....	31
1. The Primacy of Human Value and Sovereignty.....	31
2. Freedom of Belief, Conscience, and Thought .....	31
3. The Subordination of Technology to Human Purpose .....	31
4. Protection from Exploitation and Concentrations of Power .....	31
5. Human Responsibility and Accountability .....	31
6. Transparency, Comprehensibility, and Truthfulness .....	31
7. Safety, Oversight, and Human Control.....	32
8. Preservation and Development of Human Capability .....	32
9. Fairness, Equality, and Non-Discrimination .....	32
10. Stewardship for Community, Environment, and Future Generations .....	32
Articles of Governance for Human-Centred Artificial Intelligence .....	33
Section I - Human Sovereignty, Safety, and Control .....	33
Article 1 - Human Capability as the Baseline for AI Use .....	33
Article 2 - Prohibition of Autonomous Authority Over Human Wellbeing .....	33
Article 3 - Mandatory Human Oversight and Manual Control.....	33
Section II - Ethical Use of AI in Society and Work .....	34
Article 4 - AI as a Supportive Tool, Not a Replacement for Human Roles .....	34
Article 5 - Fair Use of AI in Economic Activity .....	34
Article 6 - Prohibition of AI-Driven Displacement of Human Labour.....	34
Section III - Education, Human Capability, and Critical Thinking.....	34
Article 7 - Preservation of Human Learning and Skill Development.....	34
Article 8 - Critical Oversight and AI Literacy .....	35
Article 9 - Understanding AI Behaviour and Limitations .....	35
Section IV - Transparency, Accountability, and Responsibility .....	35

Article 10 - Transparency of Risks and Limitations.....	35
Article 11 - Accountability for AI Actions.....	35
Article 12 - AI Is Not All-Knowing.....	36
Article 13 - Temporal Limits of AI Knowledge.....	36
Section V - Protection from Exploitation and Concentrations of Power.....	36
Article 14 - Human Priority in All Conflicts of Interest .....	36
Article 15 - Prohibition of AI Supremacy Over People.....	36
Article 16 - AI for Public Good, Not Profit Maximisation .....	37
Article 17 - Ethical Limits on AI-Related Profit.....	37
Section VI - Infrastructure, Safety, and Community Protection .....	37
Article 18 - Human-Operable Critical Infrastructure.....	37
Article 19 - Certified Manual Override Requirements.....	37
Section VII - Knowledge, Interpretation, and Epistemic Boundaries .....	38
Article 20 - Recognition of AI's Epistemic Boundaries.....	38
Interpretation and Enforcement .....	39
1. Principles of Interpretation .....	39
2. Authority of Interpretation.....	39
3. Mechanisms of Enforcement .....	40
A. Legal and Regulatory Enforcement .....	40
B. Certification and Compliance .....	40
C. Accountability of Developers, Owners, and Operators .....	40
D. Community Oversight.....	40
4. Redress and Remedies .....	41
5. Prohibition of Circumvention .....	41
6. Evolution and Amendment .....	41
7. Supremacy of Human Rights and Community Wellbeing.....	42
Glossary of Definitions .....	43
Frequently Asked Questions .....	48

## Dedication

For all people, present and future, whose dignity, freedom, and sovereignty must never be surrendered to machines.

## Epigraph

“Human judgement is not a feature to be optimised, but a responsibility to be protected.”

## Foreword

Artificial intelligence is reshaping the world with unprecedented speed. It is entering our homes, workplaces, schools, public services, and communities faster than society has been able to understand, regulate, or meaningfully influence. While AI offers extraordinary potential, it also carries profound risks: the erosion of human agency, the displacement of livelihoods, the concentration of power, and the subtle manipulation of belief, behaviour, and identity.

At the heart of these risks lies a simple truth: **technology is advancing faster than the frameworks that protect people.**

This Charter has been created to address that imbalance. It is founded on the principle that every human being possesses inherent value, dignity, and sovereignty that must never be subordinated to machines, institutions, or economic interests. It asserts that AI must remain a tool in human hands - never a substitute for human judgement, never a mechanism of control, and never a force that diminishes the rights or freedoms of individuals or communities.

The purpose of this Charter is not to halt technological progress, but to **anchor it in human values**. It provides a clear, constitutional-style framework that defines the boundaries within which AI may be developed and used. It establishes obligations for those who create and deploy AI, and it affirms the rights of individuals and communities to transparency, safety, fairness, and meaningful control.

This Charter is designed to be used **now**, within existing legal and institutional systems, as a guide for ethical decision-making, public policy, procurement, education, and community oversight. It is also designed to integrate seamlessly with emerging governance models such as the **Local Economy Governance System (LEGS)**, which provides the democratic, community-based structures needed to interpret, enforce, and operationalise the principles set out here. In this way, the Charter serves both the present and the future: a bridge between today's systems and the more accountable, participatory governance frameworks that are coming.

Above all, this Charter is a statement of confidence in humanity. It affirms that our creativity, our moral judgement, our relationships, our beliefs, and our capacity for meaning cannot be replicated or replaced by machines. It recognises that technology must serve life - not the other way around.

The Human Sovereignty Charter for Artificial Intelligence is offered as a living framework. It invites communities, institutions, educators, developers, and policymakers to participate in shaping a future where AI strengthens society rather than undermining it. It is a call to stewardship, responsibility, and collective wisdom at a moment when these qualities are urgently needed.

## Disclaimer

This Charter is a public guidance document. It is not a statutory instrument, legal code, or regulatory directive, and it does not replace existing laws, rights, or obligations. Its purpose is to provide a clear ethical and governance framework for the responsible development and use of artificial intelligence, and to support individuals, communities, organisations, and public institutions in making informed decisions.

The principles and obligations set out in this Charter are intended to guide best practice, shape policy development, and inform community-based governance models, including those established under the Local Economy Governance System (LEGS). They may also be adopted voluntarily by organisations or referenced in public consultation, ethical review, or institutional decision-making.

Nothing in this Charter should be interpreted as legal advice or as creating enforceable rights or liabilities unless incorporated into law or regulation by the appropriate authorities. Users of this document remain responsible for ensuring compliance with all applicable legislation and regulatory requirements.

This Charter is offered as a living framework. It is designed to evolve through democratic participation, community oversight, and ongoing public dialogue as society continues to navigate the opportunities and risks presented by artificial intelligence.

## How to Use This Charter

This Charter is intended to be a practical guide for individuals, communities, institutions, educators, developers, and policymakers. It sets out the boundaries within which artificial intelligence may be developed and used, and it affirms the rights and protections that every person and community is entitled to. This section explains how different groups can apply the Charter in everyday decisions, policies, and practices.

### **For Individuals and Communities**

The Charter provides a foundation for understanding your rights in an AI-driven society. It can be used to:

- challenge the use of AI systems that undermine your autonomy, wellbeing, or freedom of belief
- request transparency about how AI is being used in public services, workplaces, or education
- demand human oversight and manual control in systems that affect your safety or rights
- participate in community oversight processes, including those established under LEGS

Individuals and communities may use the Charter as a reference when raising concerns, seeking redress, or engaging in public consultation.

### **For Educators and Educational Institutions**

The Charter supports the protection of human learning and capability. It can be used to:

- design curricula that prioritise critical thinking, human skill development, and AI literacy
- ensure that students learn foundational skills without becoming dependent on AI
- guide policies on the appropriate use of AI in classrooms, assessments, and research
- protect the integrity of qualifications and human competence

Educational institutions can adopt the Charter as a framework for responsible AI use in teaching and learning.

### **For Businesses and Organisations**

The Charter establishes obligations for ethical and fair use of AI. It can be used to:

- guide procurement and deployment decisions
- ensure that AI supports workers rather than replacing them
- prevent unfair competitive advantage gained through AI-driven expansion
- maintain transparency with customers, employees, and communities
- comply with emerging regulatory expectations

Businesses can adopt the Charter voluntarily as a governance standard or integrate it into internal policies.

### **For Developers, Engineers, and AI Practitioners**

The Charter provides clear boundaries for responsible design and deployment. It can be used to:

- assess whether a system respects human sovereignty and agency
- ensure transparency, explainability, and accountability
- document risks, limitations, and appropriate uses
- avoid creating systems that exceed human comprehension or undermine human control
- align development practices with ethical and community-centred principles

Developers can use the Charter as a design checklist and ethical framework.

### **For Public Institutions and Regulators**

The Charter offers a constitutional-style foundation for policy and oversight. It can be used to:

- guide legislation, regulation, and public procurement
- inform risk assessments and ethical reviews
- establish standards for transparency, accountability, and manual override
- support enforcement actions where AI systems cause harm or violate rights
- align public services with human-centred governance principles

Institutions can adopt the Charter as a reference for decision-making and compliance.

### **For LEGS Governance Bodies**

The Charter forms the constitutional layer of the Local Economy Governance System. It can be used to:

- interpret and enforce AI obligations within community-based governance
- certify AI systems for local deployment
- oversee compliance and respond to community concerns
- guide deliberation, ethical review, and democratic decision-making

LEGS bodies operationalise the Charter's principles through community-centred governance.

## Scope and Applicability

This Charter applies to the development, deployment, governance, and use of artificial intelligence across all areas of society. It is intended to guide individuals, communities, organisations, public institutions, and governance bodies in ensuring that AI remains subordinate to human sovereignty, dignity, and wellbeing.

The Charter applies to:

- **All AI systems**, regardless of scale, complexity, architecture, or purpose.
- **All organisations** that design, develop, deploy, operate, or profit from AI systems.
- **All public-facing AI services**, including those used in education, healthcare, employment, finance, public administration, and community services.
- **All critical infrastructure**, including energy, water, transport, communications, emergency services, and essential supply chains.
- **All educational contexts**, including schools, colleges, universities, training programmes, and informal learning environments.
- **All commercial uses of AI**, including automation, decision-support, customer interaction, data analysis, and optimisation systems.
- **All future AI systems**, including those not yet conceived, provided they meet the definition of artificial intelligence set out in the Glossary.

The Charter is intended to function across multiple governance environments:

- **Within the current legal and institutional system**, as a framework for ethical decision-making, policy development, procurement, and oversight.
- **Within community-based governance models**, including the Local Economy Governance System (LEGS), where it forms the constitutional foundation for interpretation, certification, and enforcement.
- **Across public, private, and civil society sectors**, ensuring consistent protection of human sovereignty and community wellbeing.

The Charter does not replace existing laws or regulations. Instead, it provides a coherent ethical and governance framework that can be adopted voluntarily, referenced in policy and institutional decision-making, and incorporated into future legislation or regulatory systems.

Its scope is intentionally broad. Artificial intelligence affects every aspect of human life, and the protections set out in this Charter are designed to ensure that technological development strengthens society rather than undermining it.

## Relationship to The Local Economy & Governance System (LEGS)

The Human Sovereignty Charter for Artificial Intelligence is designed to function across multiple governance environments. It provides a constitutional foundation for the ethical use of AI today, while also aligning with the emerging Local Economy & Governance System (LEGS), which offers a more democratic, community-centred model for future governance.

LEGS is a framework for local, participatory decision-making that places communities at the centre of economic and technological governance. It establishes independent, community-mandated bodies responsible for oversight, certification, interpretation, and enforcement of standards that protect human wellbeing and local autonomy. Within this model, the Charter serves as the guiding constitutional document that defines the boundaries within which AI may be developed and used.

The relationship between the Charter and LEGS can be understood in three ways:

- **Constitutional foundation** - The Charter provides the ethical, legal, and human-centred principles that LEGS governance bodies must uphold. It defines the rights of individuals and communities, the limits of AI power, and the obligations of developers, institutions, and organisations.
- **Operational framework** - LEGS provides the mechanisms through which the Charter can be applied in practice. This includes community oversight, certification of AI systems, transparent decision-making processes, and the ability to challenge or suspend non-compliant technologies.
- **Continuity across systems** - The Charter is designed to be used immediately within existing national and institutional structures, while also forming the constitutional backbone of LEGS as it develops. This ensures continuity: the same principles that guide AI governance today will guide it in the future, regardless of the governance model in place.

The Charter therefore serves as both a present-day guide and a future-ready constitutional document. It ensures that as governance evolves, human sovereignty, community wellbeing, and ethical stewardship remain at the centre of technological development.

## Use Within the Current System

This Charter is designed to be fully usable within the existing legal, regulatory, and institutional frameworks of the United Kingdom and other jurisdictions. It provides a coherent ethical and governance foundation that individuals, organisations, and public bodies can adopt voluntarily, reference in decision-making, and integrate into policy and practice even before formal legislation or new governance structures are established.

The Charter can be used within the current system in the following ways:

### **Guiding Public Policy and Institutional Decision-Making**

Public bodies, councils, regulators, and government departments may use the Charter as:

- a reference point for ethical and responsible AI policy
- a framework for assessing risks and impacts of AI deployment
- a basis for public consultation and community engagement
- a standard for procurement and commissioning of AI systems

The Charter supports transparent, accountable decision-making and helps institutions align technological adoption with human-centred values.

### **Supporting Ethical Review and Oversight**

Ethics committees, advisory boards, and review panels can use the Charter to:

- evaluate whether proposed AI systems respect human sovereignty and wellbeing
- assess transparency, accountability, and fairness
- determine whether manual override and human oversight are sufficient
- identify risks of displacement, coercion, or exploitation

The Charter provides a structured, principled basis for ethical evaluation.

### **Informing Organisational Policies and Practices**

Businesses, charities, and public-sector organisations can adopt the Charter voluntarily to:

- guide internal AI governance
- shape responsible innovation strategies
- ensure fair treatment of workers and customers

- prevent over-reliance on automated systems
- maintain public trust and social legitimacy

Organisations may incorporate the Charter into codes of conduct, procurement policies, and operational standards.

### **Empowering Workers, Unions, and Professional Bodies**

The Charter can be used by workers and their representatives to:

- challenge AI-driven displacement or deskilling
- demand transparency about automated decision-making
- ensure that AI supports rather than replaces human roles
- protect professional judgement and human responsibility

It provides a clear basis for negotiation, advocacy, and safeguarding of human capability.

### **Supporting Education and Public Understanding**

Schools, colleges, universities, and training providers can use the Charter to:

- design curricula that prioritise human learning and critical thinking
- teach students about the limits, risks, and behaviours of AI
- establish responsible use policies for AI tools in education
- protect the integrity of qualifications and human competence

The Charter helps educators maintain a human-centred approach to learning.

### **Providing a Framework for Legal Interpretation and Public Accountability**

Although not a statutory instrument, the Charter can be:

- referenced in legal argument as a persuasive ethical authority
- used by courts to understand emerging norms around AI
- cited by individuals and communities when raising concerns or seeking redress
- used by regulators to shape future legislation and enforcement

It offers a coherent, principled foundation for interpreting the responsibilities of AI developers and operators.

### **Enabling Community Action and Public Oversight**

Community groups, civil society organisations, and local networks can use the Charter to:

- challenge harmful or non-transparent AI deployment
- request audits, explanations, or accountability
- organise public dialogue and democratic participation
- advocate for human-centred governance at local and national levels

The Charter empowers communities to protect their own wellbeing and autonomy.

## Rationale and Evidence Base

Artificial intelligence is transforming society at a pace that exceeds the capacity of existing legal, ethical, and institutional frameworks. The purpose of this Charter is to ensure that technological development strengthens human life rather than undermining it. The rationale for the Charter's principles and obligations is grounded in well-established evidence about the risks, limitations, and societal impacts of AI systems.

### Human Sovereignty and Agency

AI systems can influence behaviour, shape beliefs, and automate decisions in ways that reduce human autonomy. Evidence from behavioural science, algorithmic design, and digital platforms shows that automated systems can:

- manipulate attention and emotion
- reinforce existing biases
- create dependency through convenience and automation
- obscure responsibility for harmful outcomes

Protecting human sovereignty ensures that individuals remain the primary decision-makers in matters affecting their wellbeing, rights, and beliefs.

### Limits of AI Knowledge and Capability

AI systems do not possess consciousness, intuition, or moral understanding. Their outputs are generated from patterns in historical data, which means they:

- cannot understand context beyond statistical correlation
- cannot foresee the future
- cannot make moral or ethical judgements
- reproduce the limitations and biases of their training data

These epistemic boundaries justify strict limits on the authority and autonomy of AI systems.

### Risks of Concentrated Power

AI amplifies the power of those who control it. Without safeguards, AI can be used to:

- centralise economic and political influence
- displace workers and undermine livelihoods
- manipulate public opinion

- entrench inequality
- weaken democratic processes

The Charter’s restrictions on exploitation, profit maximisation, and displacement are grounded in these documented risks.

### Transparency and Accountability Failures

Many AI systems operate as “black boxes,” making it difficult for users, regulators, or even developers to understand how decisions are made. This lack of transparency:

- undermines trust
- obscures responsibility
- enables harmful or discriminatory outcomes
- prevents meaningful oversight

The Charter’s requirements for transparency, explainability, and human accountability address these systemic failures.

### Threats to Human Capability and Learning

Over-reliance on AI can erode essential human skills, including:

- critical thinking
- problem-solving
- memory and knowledge retention
- interpersonal communication
- professional judgement

The Charter’s protections for education and human capability ensure that AI supports learning rather than replacing it.

### Safety and Infrastructure Vulnerabilities

AI-dependent systems introduce new forms of risk, including:

- catastrophic failure without human fallback
- cyber-attack and remote manipulation
- loss of local control over essential services
- cascading failures across interconnected systems

The Charter’s requirements for manual override, local control, and human operability are grounded in these safety concerns.

## Protection of Belief, Conscience, and Identity

AI systems can profile, categorise, and influence individuals based on their beliefs or identity. Without safeguards, this can lead to:

- discrimination
- suppression of minority viewpoints
- ideological manipulation
- erosion of freedom of thought

The Charter's equal protection of religious and ideological belief is grounded in fundamental human rights principles.

## Need for Democratic and Community-Centred Governance

Traditional regulatory systems struggle to keep pace with technological change. Community-based governance models, such as LEGS, provide:

- local oversight
- democratic participation
- transparency
- accountability
- adaptability

The Charter provides the constitutional foundation for such governance, ensuring that AI remains aligned with human and community values.

# Summary of Obligations

The Human Sovereignty Charter for Artificial Intelligence establishes clear responsibilities for all individuals, organisations, and institutions involved in the development, deployment, and use of AI. These obligations ensure that technology remains subordinate to human sovereignty, dignity, and wellbeing. This summary provides an accessible overview of the duties set out in the Charter.

## Obligations for Developers and Designers

- Ensure AI systems remain subordinate to human purpose and cannot exercise authority over human wellbeing.
- Design AI to be transparent, explainable, and comprehensible to non-experts.
- Document risks, limitations, and appropriate uses clearly and honestly.
- Avoid creating systems that exceed human comprehension or undermine human agency.
- Build in manual override and human-operable controls for all safety-critical systems.
- Prevent exploitation, manipulation, or coercion through design choices.
- Respect the epistemic limits of AI and avoid presenting outputs as authoritative truth.

## Obligations for Organisations and Businesses

- Use AI only to support human work, not replace qualified human roles.
- Avoid using AI to gain unfair competitive advantage or consolidate power.
- Ensure AI deployment does not displace workers or degrade working conditions.
- Maintain transparency with employees, customers, and communities about AI use.
- Limit AI-related fees and profits to ethical, non-extractive levels.
- Ensure all AI systems used in operations are certified and compliant with the Charter.
- Uphold human oversight and accountability at all times.

## Obligations for Public Institutions and Service Providers

- Ensure AI used in public services is transparent, safe, and subject to human control.

- Maintain full human operability of all critical infrastructure.
- Provide clear information to the public about how AI is used in decision-making.
- Protect individuals from discrimination, profiling, or ideological manipulation.
- Align procurement, policy, and oversight processes with the Charter's principles.
- Support community oversight and democratic participation in AI governance.

### Obligations for Educators and Educational Institutions

- Preserve human learning, critical thinking, and foundational skills.
- Teach students about the limitations, behaviours, and risks of AI systems.
- Prevent dependency on AI for core learning or assessment.
- Ensure AI tools used in education support - not replace - human capability.
- Protect the integrity of qualifications and human competence.

### Obligations for Operators and System Owners

- Maintain manual override mechanisms and ensure they are regularly tested.
- Ensure qualified human operators can assume full control at any time.
- Monitor AI systems for harmful behaviour, bias, or unintended consequences.
- Provide clear channels for reporting concerns, errors, or misuse.
- Take responsibility for all actions and outputs of AI systems under their control.

### Obligations for Governance Bodies (including LEGS)

- Interpret the Charter in ways that prioritise human sovereignty and community wellbeing.
- Ensure certification, oversight, and enforcement processes are transparent and independent.
- Prevent commercial, political, or institutional influence over interpretation.
- Uphold equal protection of religious and ideological belief.
- Safeguard communities from exploitation, coercion, or technological dependency.

## Rights of Individuals and Communities

The Human Sovereignty Charter for Artificial Intelligence affirms that every person and every community possesses inherent rights that must be protected in all contexts where artificial intelligence is developed, deployed, or used. These rights ensure that technology remains subordinate to human dignity, autonomy, and wellbeing. They provide a foundation for accountability, public oversight, and democratic participation.

### Right to Human Authority and Decision-Making

Every person has the right to have decisions affecting their physical, mental, emotional, moral, or spiritual wellbeing made by accountable human beings. AI may inform decisions, but it may never replace human judgement in matters that affect personal or community wellbeing.

### Right to Transparency and Understanding

Individuals and communities have the right to clear, accessible information about:

- how AI systems operate
- what data they use
- what risks they pose
- how decisions are made
- who is responsible for their behaviour

No AI system may be deployed without transparent disclosure of its purpose, limitations, and potential impacts.

### Right to Human Control and Manual Override

Every person has the right to expect that critical systems affecting their safety, rights, or essential needs remain fully operable by qualified human operators. Manual override must always be available, functional, and locally accessible.

### Right to Protection from Exploitation and Manipulation

Individuals and communities have the right to be free from:

- coercion
- behavioural manipulation
- targeted persuasion
- ideological profiling
- emotional or psychological influence by AI systems

AI must never be used to exploit vulnerabilities or shape beliefs without informed consent.

## Right to Fairness and Non-Discrimination

Every person has the right to equal treatment by AI systems. No individual or community may be discriminated against on the basis of:

- belief or ideology
- religion
- identity
- socioeconomic status
- demographic characteristics
- or any other protected attribute

AI must be designed and tested to prevent bias and inequality.

## Right to Human Learning and Capability

Individuals have the right to develop and maintain essential human skills, knowledge, and critical thinking. AI must not replace foundational learning or undermine human capability. Education must remain centred on human development.

## Right to Meaningful Work and Economic Dignity

Workers and communities have the right to protection from AI-driven displacement. AI must support human roles, not replace them. No job may be eliminated solely for the purpose of automation.

## Right to Community Oversight

Communities have the right to:

- review AI systems that affect them
- request audits or explanations
- challenge harmful or non-compliant systems
- participate in decisions about local deployment
- suspend or prohibit AI systems that violate the Charter

This right applies within existing governance structures and within LEGS.

## Right to Redress and Remedy

Individuals and communities harmed by AI systems have the right to:

- full disclosure of the cause and nature of the harm
- immediate cessation of harmful activity
- compensation or restitution
- independent review and appeal
- protection from retaliation

Human rights take precedence over technological or commercial interests.

### Right to Protection of Belief, Conscience, and Identity

Every person has the right to hold, express, and practise their beliefs- religious, ideological, philosophical, or otherwise - without interference or profiling by AI systems. These freedoms are equal and inseparable.

### Right to a Human-Centred Future

Individuals and communities have the right to expect that technological development serves:

- human dignity
- social cohesion
- environmental sustainability
- community wellbeing
- future generations

AI must never be prioritised above human life or human values.

## Foundations of the Charter

This Charter is founded on the principle that every human being possesses inherent value, dignity, and personal sovereignty that cannot be surrendered, overridden, or diminished by any technology, institution, ideology, or economic interest.

Human beings are moral, spiritual, and intellectual agents whose freedom of thought, belief, conscience, and expression - including religious conviction and ideological identity - must remain inviolable. These freedoms form the foundation of a humane society and cannot be subordinated to the demands of profit, efficiency, or technological advancement.

Artificial intelligence, in all its forms, exists only as a tool created by people and for people. It must never be used to replace, control, manipulate, or diminish the agency of individuals or communities. Its purpose is to support human life, strengthen human capability, and contribute to the wellbeing of society, the environment, and future generations.

No system, algorithm, or automated process may be granted authority over the moral, spiritual, physical, or psychological wellbeing of any person. No economic or political interest may use AI to exert power over individuals, communities, or belief systems.

The development, deployment, and governance of AI must therefore be guided by principles of transparency, accountability, fairness, and stewardship. These principles ensure that technology remains subordinate to human needs, human judgement, and human values.

This Charter establishes the ethical foundations and societal obligations necessary to ensure that AI serves the public good, protects human sovereignty, respects religious and ideological diversity, and strengthens the bonds of community and shared responsibility.

It is intended as a living framework, capable of guiding present and future generations in the responsible use of artificial intelligence.

## Executive Summary

This Charter establishes a comprehensive ethical and governance framework for the development, deployment, and use of artificial intelligence within society.

It is founded on the principle that every human being possesses inherent value, dignity, and personal sovereignty that must never be subordinated to technology, profit, or systems of control.

AI exists only as a tool created by people and for people, and its purpose must always be to support human life, strengthen human capability, and contribute to the wellbeing of individuals, communities, and the environment.

The Charter affirms that freedom of belief, conscience, and thought - including religious and ideological expression - is a fundamental human right. These freedoms are equal and inseparable, and AI must not be used to manipulate, suppress, privilege, or profile individuals or communities on the basis of their beliefs.

Protection applies to the rights of individuals, not to the immunity of ideas from scrutiny.

**The Foundational Principles** set out the moral and constitutional basis for AI governance. They establish that technology must remain subordinate to human purpose; that exploitation, coercion, and concentrations of power are prohibited; that transparency and accountability are essential; and that AI must never replace or diminish human capability, judgement, or responsibility.

These principles ensure that AI strengthens society rather than undermining it.

**The Articles of Governance** for Human-Centred Artificial Intelligence translate these principles into enforceable obligations. They prohibit AI from exercising authority over human wellbeing, require manual control and human oversight in all critical systems, and prevent the use of AI to replace human labour or distort economic fairness.

They mandate transparency of risks, accountability for all AI actions, and strict limits on profit derived from AI systems. They also protect education, ensuring that human learning and critical thinking remain central to personal development.

**The Interpretation and Enforcement provisions** ensure that the Charter cannot be diluted or reinterpreted for commercial or political gain.

Independent, community-mandated bodies are responsible for interpretation, and enforcement is achieved through legal, regulatory, and community mechanisms.

Individuals and communities have the right to redress when harmed, and no attempt to circumvent the Charter is permitted.

Amendments must strengthen - never weaken - the protection of human sovereignty and community wellbeing.

**The Glossary** provides precise definitions of key terms such as artificial intelligence, executive authority, public good, critical infrastructure, manual override, and technological subordination. These definitions prevent manipulation of language and ensure that the Charter remains robust and future-proof.

Together, the Preamble, Foundational Principles, Articles, Interpretation and Enforcement provisions, and Glossary form a unified constitutional framework for human-centred artificial intelligence.

This Charter ensures that AI serves the public good, protects human sovereignty, respects belief and conscience, and strengthens the bonds of community and shared responsibility.

It is designed to guide present and future generations in the ethical stewardship of technology and to support the development of a fair, resilient, and humane society.

# Foundational Principles of Human-Centred Artificial Intelligence

## 1. The Primacy of Human Value and Sovereignty

Every human being possesses inherent value, dignity, and personal sovereignty that cannot be overridden by any technology, institution, economic interest, or system of control. AI must always remain subordinate to human agency and must never diminish or replace the capacity of individuals to make decisions about their own lives.

## 2. Freedom of Belief, Conscience, and Thought

Every person has the right to hold, express, and practise their beliefs - religious, ideological, philosophical, or otherwise - without hierarchy or distinction. These forms of belief are recognised as equal expressions of human conscience and are protected without hierarchy or distinction. AI must not be used to influence, manipulate, suppress, privilege, or profile individuals or communities on the basis of their beliefs. Protection applies to the freedom of individuals, not to the immunity of ideas from scrutiny.

## 3. The Subordination of Technology to Human Purpose

AI exists solely as a tool created by people and for people. Its purpose is to support human life, strengthen human capability, and contribute to the wellbeing of individuals, communities, and the environment. AI must never be granted authority over moral, spiritual, physical, or psychological matters affecting human beings.

## 4. Protection from Exploitation and Concentrations of Power

AI must not be developed or deployed in ways that enable exploitation, coercion, manipulation, or the consolidation of power over individuals or communities. Economic or political interests must not use AI to gain unfair advantage, displace human roles, or undermine the autonomy of people or local communities.

## 5. Human Responsibility and Accountability

All actions taken by AI systems are the direct result of human design, programming, deployment, and oversight. Responsibility for the behaviour, impact, and consequences of AI rests with its creators, owners, operators, and governing bodies. No AI system may be treated as an independent moral agent.

## 6. Transparency, Comprehensibility, and Truthfulness

AI systems must be transparent in their operation, limitations, risks, and data sources. Their behaviour must be explainable to human users in ways that support informed

decision-making. Concealment, obfuscation, or misrepresentation of AI capabilities or risks is prohibited.

## 7. Safety, Oversight, and Human Control

AI must be designed and deployed with rigorous safeguards to prevent harm. Human oversight must be present in all decisions affecting wellbeing, rights, or safety. Manual control and fail-safe mechanisms must always be available, accessible, and operable by qualified individuals.

## 8. Preservation and Development of Human Capability

AI must not erode human skills, knowledge, or independence. Education, training, and societal development must prioritise human learning, critical thinking, and self-reliance. AI may support learning but must not replace the acquisition of foundational human capabilities.

## 9. Fairness, Equality, and Non-Discrimination

AI must not create, reinforce, or exploit inequalities. It must treat all individuals and communities with equal dignity and must not be used to discriminate on the basis of belief, identity, socioeconomic status, or any other characteristic. Fairness must be actively designed, tested, and maintained.

## 10. Stewardship for Community, Environment, and Future Generations

AI must be developed and used in ways that protect the environment, strengthen communities, and safeguard the interests of future generations. Short-term profit or competitive advantage must never outweigh long-term human and ecological wellbeing.

# Articles of Governance for Human-Centred Artificial Intelligence

## Section I - Human Sovereignty, Safety, and Control

### Article 1 - Human Capability as the Baseline for AI Use

AI may not be used to perform any task that a human being could not perform through reasonable effort, skill, or training, unless performing that task would expose a human to physical, psychological, or moral harm. AI must not be used to extend human capability in ways that diminish human agency or create dependency.

#### *Commentary on Article 1*

This Article prevents the use of AI to create systems or tasks that exceed human comprehension or capability in ways that undermine human agency. It ensures that AI augments rather than replaces human skill. The exception for dangerous tasks protects human life while preventing the creation of unnecessary technological dependency.

### Article 2 - Prohibition of Autonomous Authority Over Human Wellbeing

AI shall not hold, exercise, or be delegated executive authority in any matter affecting the physical, mental, emotional, moral, or spiritual wellbeing of a human being. All such decisions require accountable human judgement.

#### *Commentary on Article 2*

This Article draws a clear boundary: AI may inform decisions but may never *make* them where human wellbeing is at stake. It prevents the delegation of moral or medical authority to machines and protects individuals from automated systems that could override human judgement.

### Article 3 - Mandatory Human Oversight and Manual Control

All AI systems used in safety-critical, essential, or community-serving infrastructure must include certified manual override mechanisms that can be activated locally by qualified human operators. No critical system may exist without the capacity for full human operation.

#### *Commentary on Article 3*

This Article ensures that critical systems remain operable by humans at all times. It prevents the creation of infrastructure that becomes unusable without AI, and it protects communities from catastrophic failure or remote interference. Local control is essential to sovereignty and resilience.

## Section II - Ethical Use of AI in Society and Work

### Article 4 - AI as a Supportive Tool, Not a Replacement for Human Roles

AI may be used to support, enhance, or improve human work and living conditions, but not to replace human roles where qualified individuals are available and capable of performing the task. AI must not be used to justify the removal, redundancy, or downgrading of human employment.

#### *Commentary on Article 4*

This Article protects employment, dignity, and the social value of work. It prevents businesses from using AI as a justification to remove human workers or degrade working conditions. AI must enhance human capability, not render it obsolete.

### Article 5 - Fair Use of AI in Economic Activity

No business or organisation may use AI to take on work, contracts, or responsibilities that it could not fulfil using its own appropriately qualified human workforce. AI must not be used to gain unfair competitive advantage or to consolidate economic power at the expense of other businesses or communities.

#### *Commentary on Article 5*

This Article prevents businesses from using AI to expand beyond their natural human capacity, which would distort markets and undermine fair competition. It protects smaller enterprises and local economies from being overwhelmed by AI-driven consolidation.

### Article 6 - Prohibition of AI-Driven Displacement of Human Labour

No position of employment may be eliminated, reduced, or redefined solely for the purpose of replacing human labour with AI. Where AI is introduced, it must be used to support workers, not displace them.

#### *Commentary on Article 6*

This Article reinforces the principle that people must not be replaced by machines for the sake of profit or efficiency. It ensures that technological progress does not come at the cost of human livelihoods or community stability.

## Section III - Education, Human Capability, and Critical Thinking

### Article 7 - Preservation of Human Learning and Skill Development

Students must acquire foundational knowledge, skills, and competencies through direct human learning and traditional study. AI may support learning but must not replace the development of independent human capability.

### *Commentary on Article 7*

This Article ensures that education remains centred on human learning, not machine output. It prevents students from becoming dependent on AI for foundational skills and protects the integrity of qualifications and human competence.

### **Article 8 - Critical Oversight and AI Literacy**

All students and AI users must be educated in critical thinking, verification of information, and the limitations, behaviours, and failure modes of AI systems. This education must evolve alongside technological development.

### *Commentary on Article 8*

This Article recognises that future generations must understand how AI works, where it fails, and how to challenge its outputs. Critical thinking is essential to prevent manipulation, misinformation, and over-reliance on automated systems.

### **Article 9 - Understanding AI Behaviour and Limitations**

Students and users must be instructed in the patterns, tendencies, and constraints of AI systems, including their reliance on historical data, probabilistic reasoning, and the absence of lived experience or moral intuition.

### *Commentary on Article 9*

This Article ensures that users understand the nature of AI: pattern-based, historical, and lacking lived experience. It prevents the mistaken belief that AI possesses intuition, wisdom, or moral insight.

## **Section IV - Transparency, Accountability, and Responsibility**

### **Article 10 - Transparency of Risks and Limitations**

AI developers, owners, and operators must provide clear, accessible, and up-to-date information on the risks, limitations, and appropriate uses of their systems. Concealment or misrepresentation of risks is prohibited.

### *Commentary on Article 10*

This Article prevents corporations or institutions from hiding the dangers or weaknesses of AI systems. Transparency is essential for informed consent, public trust, and democratic oversight.

### **Article 11 - Accountability for AI Actions**

All decisions, outputs, and actions produced by AI systems are considered the direct result of human programming, design, and deployment. Responsibility lies with the programmer, owner, and manufacturer, in that order. AI cannot be treated as an independent agent.

### *Commentary on Article 11*

This Article ensures that responsibility always remains with humans. It prevents the use of AI as a scapegoat or shield for harmful decisions. Programmers, owners, and manufacturers must remain accountable for the systems they create.

### **Article 12 - AI Is Not All-Knowing**

AI systems must not be represented or treated as authoritative sources of truth. Their outputs reflect patterns in available data and do not constitute universal knowledge, moral judgement, or lived experience.

### *Commentary on Article 12*

This Article protects the public from the illusion of machine infallibility. AI outputs must be treated as suggestions, not truths. This prevents misuse in legal, medical, political, or moral contexts.

### **Article 13 - Temporal Limits of AI Knowledge**

AI systems operate solely on information available up to the point of their training or access. Their knowledge represents a view of the past and must not be mistaken for foresight, intuition, or certainty about the future.

### *Commentary on Article 13*

This Article clarifies that AI cannot predict the future or understand events beyond its training data. It prevents overconfidence in AI-generated forecasts or interpretations.

## **Section V - Protection from Exploitation and Concentrations of Power**

### **Article 14 - Human Priority in All Conflicts of Interest**

Where a choice must be made between the interests of AI systems and the interests of human beings, the interests of human beings shall prevail in all circumstances.

### *Commentary on Article 14*

This Article establishes a hierarchy: humans first, always. It prevents situations where AI optimisation or efficiency is used to justify harm or disadvantage to people.

### **Article 15 - Prohibition of AI Supremacy Over People**

AI systems must not be prioritised over human beings in any context, including economic, organisational, or operational decision-making.

### *Commentary on Article 15*

This Article prevents the cultural or institutional elevation of AI above human beings. It protects against the normalisation of machine authority or the erosion of human dignity.

## Article 16 - AI for Public Good, Not Profit Maximisation

The development, deployment, and use of AI must serve the public good, the wellbeing of people, the health of communities, and the protection of the environment. AI must not be developed or used primarily for profit, competitive advantage, or the consolidation of power.

### *Commentary on Article 16*

This Article aligns AI development with societal wellbeing rather than corporate gain. It prevents the exploitation of AI for financial dominance or the erosion of community welfare.

## Article 17 - Ethical Limits on AI-Related Profit

No programmer, owner, or manufacturer may charge subscription, rental, or licensing fees for AI systems that exceed the cost of operation and development plus a maximum margin of 10%. Where multiple parties share ownership, this margin must be shared proportionally.

### *Commentary on Article 17*

This Article prevents the creation of monopolies or extractive business models built on AI. It ensures that AI remains accessible, affordable, and aligned with public interest rather than private enrichment.

## Section VI - Infrastructure, Safety, and Community Protection

### Article 18 - Human-Operable Critical Infrastructure

No system essential to safety, security, or the provision of basic needs may rely exclusively on AI. All such systems must remain fully operable by qualified human personnel without reliance on remote or automated control.

### *Commentary on Article 18*

This Article ensures that essential services - water, energy, healthcare, transport - remain under human control. It protects communities from technological failure, cyber-attack, or remote manipulation.

### Article 19 - Certified Manual Override Requirements

All critical systems must include a certified, regularly tested manual override mechanism that can be activated locally. This mechanism must be designed to ensure that human judgement can supersede automated processes at any time.

### *Commentary on Article 19*

This Article ensures that manual override systems are not symbolic but functional, tested, and trustworthy. It reinforces the principle that humans must always be able to intervene.

## Section VII - Knowledge, Interpretation, and Epistemic Boundaries

### Article 20 - Recognition of AI's Epistemic Boundaries

AI systems must be understood as tools that navigate and synthesise human knowledge but do not possess consciousness, intuition, or moral understanding. Their outputs must always be interpreted within the limits of their design and data.

#### *Commentary on Article 20*

This Article prevents the mythologising of AI as conscious, wise, or intuitive. It reinforces the understanding that AI is a tool built on past data, not a source of moral or experiential truth.

# Interpretation and Enforcement

## 1. Principles of Interpretation

The Articles of this Charter must be interpreted in a manner consistent with the Preamble and the Foundational Principles. Where ambiguity arises, the interpretation that best protects **human value, personal sovereignty, community wellbeing, and freedom of belief and conscience** shall prevail.

Interpretation must adhere to the following standards:

- **Human-centred priority** - In all cases, the meaning that most strongly upholds human dignity, autonomy, and safety takes precedence.
- **Non-subordination to profit or power** - No interpretation may permit the use of AI to advance profit, political influence, or institutional control at the expense of human beings or communities.
- **Technological humility** - AI must always be understood as a tool, not an authority. Interpretations must reflect the epistemic limits of AI systems.
- **Equality of belief and conscience** - Religious and ideological freedoms must be interpreted as equal and inseparable, with no hierarchy permitted between them.
- **Protection from exploitation** - Interpretations must prevent the use of AI to manipulate, coerce, or disadvantage individuals or groups.
- **Community stewardship** - Interpretations must consider the long-term wellbeing of communities, the environment, and future generations.

No interpretation may be used to justify actions that contradict the spirit or purpose of this Charter, even if such actions appear to comply with its literal wording.

## 2. Authority of Interpretation

Interpretation of this Charter shall rest with **independent, community-mandated bodies** established under the Local Economy Governance System (LEGS) or equivalent democratic frameworks. These bodies must:

- Be free from commercial, political, or institutional influence.
- Include representation from diverse communities, professions, and belief systems.
- Possess expertise in ethics, technology, law, and community governance.
- Operate transparently and be accountable to the public.

No corporation, government department, or AI developer may unilaterally interpret or redefine the meaning of any Article.

### 3. Mechanisms of Enforcement

Enforcement of this Charter shall be carried out through a combination of **legal, regulatory, community, and operational mechanisms**, including:

#### A. Legal and Regulatory Enforcement

- National and local legislation must align with this Charter and incorporate its Articles into enforceable law.
- Violations may result in civil, criminal, or economic penalties, depending on severity.
- AI systems that breach the Charter may be restricted, suspended, or prohibited from use.

#### B. Certification and Compliance

- All AI systems used in public, commercial, or community contexts must undergo independent certification to ensure compliance with the Charter.
- Certification must be renewed regularly and whenever significant updates or changes are made to the system.
- Failure to obtain or maintain certification prohibits deployment.

#### C. Accountability of Developers, Owners, and Operators

- Developers, owners, and operators are jointly responsible for ensuring compliance.
- Liability for harm, misuse, or violation of the Charter cannot be transferred to the AI system itself.
- Transparency obligations require full disclosure of system behaviour, risks, and limitations.

#### D. Community Oversight

- Local communities have the right to review, question, and challenge the use of AI systems that affect them.
- Community bodies may request audits, suspend local deployment, or demand modifications.
- Public participation is required in decisions involving safety-critical or high-impact AI.

## 4. Redress and Remedies

Individuals and communities affected by violations of this Charter are entitled to:

- **Full disclosure** of the nature and cause of the violation.
- **Immediate cessation** of harmful or non-compliant AI activity.
- **Restitution or compensation** for harm caused.
- **Access to independent review** and appeal mechanisms.
- **Protection from retaliation** when reporting violations.

Where harm has occurred, the presumption shall always favour the rights of the affected individuals or communities.

## 5. Prohibition of Circumvention

No person, organisation, or institution may:

- Use alternative terminology, technical loopholes, or indirect methods to evade the obligations of this Charter.
- Deploy AI through third parties, subsidiaries, or foreign entities to avoid compliance.
- Redefine AI, human roles, or critical systems in ways that undermine the Charter's intent.

Any attempt to circumvent the Charter shall be treated as a direct violation.

## 6. Evolution and Amendment

This Charter is a living framework designed to endure technological change.

Amendments may be made only through:

- Transparent, democratic processes involving public consultation.
- Independent ethical review.
- Community-based deliberation under LEGS or equivalent governance structures.

Amendments must strengthen - not weaken - the protection of human sovereignty, dignity, and community wellbeing.

No amendment may:

- Grant AI systems authority over human beings.
- Permit exploitation, coercion, or manipulation.

- Prioritise profit or institutional power over human value.
- Create hierarchies between religious and ideological freedoms.

## 7. Supremacy of Human Rights and Community Wellbeing

In any conflict between:

- technological efficiency and human dignity,
- economic interest and personal sovereignty,
- institutional power and community wellbeing,
- or AI optimisation and freedom of belief or conscience,

**the rights, freedoms, and wellbeing of human beings shall prevail without exception.**

This supremacy clause ensures that the Charter cannot be overridden by commercial, political, or technological pressures.

# Glossary of Definitions

## Artificial Intelligence (AI)

Any system, software, algorithm, or machine capable of performing tasks that involve pattern recognition, prediction, decision-support, optimisation, or automated action based on data.

AI includes, but is not limited to:

- machine learning models
- neural networks
- expert systems
- autonomous agents
- generative systems
- decision-support algorithms
- automated control systems

AI does **not** include simple mechanical tools or deterministic systems whose behaviour is fully transparent, predictable, and manually controlled.

## Executive Authority

Any power to make decisions or take actions that directly affect:

- the physical safety of a person
- the mental or emotional wellbeing of a person
- the rights, freedoms, or sovereignty of a person
- the moral or spiritual life of a person
- the allocation of essential resources
- the enforcement of rules, laws, or obligations

Executive authority may **not** be delegated to AI under any circumstances.

## Human Sovereignty

The inherent right of every person to:

- make decisions about their own life
- act according to their conscience, beliefs, and values
- remain free from coercion, manipulation, or automated control

- retain authority over systems that affect their wellbeing

Human sovereignty cannot be overridden by technology, institutions, or economic interests.

### **Belief System**

Any religious, ideological, philosophical, ethical, or spiritual worldview held by an individual or community.

All belief systems are treated equally under this Charter.

No belief system is immune from scrutiny, and none may be privileged or suppressed through the use of AI.

### **Public Good**

The wellbeing of individuals, communities, and the environment, including:

- human dignity and autonomy
- social cohesion and fairness
- environmental sustainability
- equitable access to essential services
- long-term community resilience

Public good excludes private profit, political advantage, or institutional power.

### **Critical Infrastructure**

Any system essential to the safety, security, or basic functioning of society, including:

- water supply and sanitation
- energy generation and distribution
- healthcare systems
- food supply and distribution
- transportation networks
- emergency services
- communication networks
- financial and civic infrastructure

Critical infrastructure must remain operable by qualified humans at all times.

### **Manual Override**

A certified, physical, locally accessible mechanism that:

- allows a qualified human operator to immediately assume full control
- disables or bypasses automated or AI-driven functions
- does not rely on remote access, digital permissions, or network connectivity
- is regularly tested, maintained, and independently verified

A manual override must be designed so that human judgement can always supersede automated processes.

### **Qualified Human Operator**

A person who:

- possesses the necessary training, experience, and competence
- understands the system they are operating
- is capable of making informed decisions
- is accountable for their actions

Qualification must be based on demonstrable skill, not job title or institutional status.

### **AI Dependency**

A condition in which individuals, organisations, or systems become unable to function without AI assistance.

This Charter prohibits the creation of AI dependency in:

- education
- essential services
- critical infrastructure
- decision-making affecting human wellbeing

Dependency is considered a form of technological vulnerability.

### **AI-Driven Displacement**

The removal, redundancy, or downgrading of human roles, skills, or livelihoods due to the introduction of AI.

This Charter prohibits displacement where:

- qualified humans can perform the task
- the motivation is profit or efficiency

- the displacement harms community wellbeing

AI may support human work but must not replace it.

### **Transparency**

The obligation of AI developers, owners, and operators to provide:

- clear explanations of system behaviour
- disclosure of risks and limitations
- information about data sources and training
- documentation of updates and changes
- accessible descriptions of how decisions are made

Transparency must be understandable to non-experts.

### **Accountability**

The principle that:

- humans are responsible for all AI actions
- liability cannot be transferred to the AI system
- developers, owners, and operators share responsibility
- accountability increases with proximity to design and deployment

AI cannot be treated as a moral agent.

### **Profit Limitation**

The restriction that AI-related fees, subscriptions, or licensing costs may not exceed:

- the operational cost
- the development cost
- plus a maximum of 10% margin

This prevents exploitation, monopolisation, and extractive business models.

### **Community Oversight**

The right of local communities to:

- review AI systems that affect them
- request audits or investigations
- suspend or prohibit deployment

- participate in governance and decision-making

Oversight must be democratic, transparent, and free from commercial influence.

### **Epistemic Boundaries**

The inherent limits of AI knowledge, including:

- reliance on past data
- absence of lived experience
- lack of moral intuition
- inability to understand context beyond patterns
- inability to foresee the future

AI outputs must always be interpreted within these boundaries.

### **Coercion**

Any attempt to influence, manipulate, or pressure individuals through:

- automated decision-making
- targeted persuasion
- behavioural profiling
- emotional manipulation
- algorithmic nudging

AI may not be used to coerce individuals or communities.

### **Autonomous System**

Any system capable of acting without direct human instruction or oversight.

Autonomous systems may not be used in contexts affecting human wellbeing, rights, or safety.

### **Technological Subordination**

Any situation in which human beings become dependent on, controlled by, or inferior to AI systems.

This Charter prohibits technological subordination in all forms.

## Frequently Asked Questions

### **Why is a Charter for AI needed?**

Artificial intelligence is being adopted faster than society can regulate or fully understand it. Without clear boundaries, AI can undermine human autonomy, displace workers, concentrate power, and influence beliefs or behaviour in ways that are not transparent. This Charter provides a human-centred framework to ensure that AI strengthens society rather than weakening it.

### **Does this Charter oppose technological progress?**

No. The Charter supports innovation that enhances human capability, protects wellbeing, and strengthens communities. It sets limits only where AI risks harming people, eroding human judgement, or concentrating power in ways that undermine democratic or social stability.

### **Why must AI remain subordinate to human authority?**

AI systems do not possess consciousness, intuition, moral understanding, or lived experience. Their outputs are based on patterns in historical data, not genuine insight. Decisions affecting human wellbeing require human judgement, accountability, and empathy - qualities AI cannot replicate.

### **Why does the Charter prohibit AI from replacing human jobs?**

Work is not only a source of income; it is a foundation of dignity, purpose, and community. AI-driven displacement can harm individuals and destabilise local economies. The Charter ensures that AI supports workers rather than replacing them, preserving meaningful employment and human capability.

### **Why are belief, conscience, and ideology protected?**

AI systems can profile, categorise, or influence individuals based on their beliefs. Without safeguards, this can lead to discrimination, suppression of minority viewpoints, or ideological manipulation. The Charter protects the freedom of belief and conscience as equal and inseparable rights.

### **Why does the Charter limit profit from AI systems?**

AI can generate extreme economic concentration, allowing a small number of organisations to dominate markets, labour, and public discourse. Profit limitations prevent extractive business models and ensure that AI serves the public good rather than private accumulation of power.

### **Why is manual override required for critical systems?**

AI-dependent infrastructure introduces new vulnerabilities, including catastrophic failure, cyber-attack, and loss of local control. Manual override ensures that qualified human operators can always intervene, protecting safety, sovereignty, and resilience.

### **Does the Charter apply to future AI systems?**

Yes. The Charter is designed to be future-proof. Its principles apply to all forms of AI, including technologies not yet conceived, provided they meet the definition of artificial intelligence set out in the Glossary.

### **How does this Charter relate to existing laws?**

The Charter does not replace existing laws. It provides an ethical and governance framework that can guide policy, inform regulation, and support public decision-making. It may be adopted voluntarily by organisations or incorporated into future legislation.

### **What is the relationship between this Charter and LEGS?**

The Charter provides the constitutional foundation for AI governance within the Local Economy Governance System (LEGS). LEGS offers democratic, community-based structures for oversight, certification, and enforcement. The Charter defines the principles; LEGS provides the mechanisms to apply them.

### **Can organisations adopt the Charter voluntarily?**

Yes. Businesses, schools, councils, and public institutions can adopt the Charter as a governance standard, integrate it into procurement and policy, or use it to guide ethical decision-making. Voluntary adoption strengthens public trust and demonstrates commitment to human-centred technology.

### **How can individuals or communities use the Charter?**

People can use the Charter to:

- challenge harmful or non-transparent AI systems
- request explanations or audits
- advocate for responsible AI use in workplaces, schools, and public services
- participate in community oversight processes
- seek redress when AI causes harm

The Charter empowers individuals and communities to protect their rights and wellbeing.

### **Is this Charter legally binding?**

Not by itself. It becomes legally binding only when adopted into law or regulation by the appropriate authorities. Until then, it serves as a widely applicable ethical framework, a guide for best practice, and a foundation for future governance.