AIWS

CONSTITUTION FOR HUMANITY IN THE AGE OF ARTIFICIAL INTELLIGENCE

Version 1.7 — Constitutional Edition

Beacon Hill, Boston · July 8, 2026

Draft for Consultation · Circulated November 2026

A Foundational Constitutional Framework for Human Dignity, Human Agency, Freedom, Trust, and Wisdom in the Age of Artificial Intelligence

An initiative of AIWS and the Boston Global Forum

This Constitution belongs to no single government, corporation, ideology, or power — but to all Humanity.

NOTE ON THIS EDITION

This is Version 1.7 of the Constitution for Humanity in the Age of Artificial Intelligence, whose drafting began on July 8, 2026. It is circulated for consultation ahead of the Boston Global Forum Conference at Loeb House, Harvard University, on November 10, 2026. The Constitution will be proclaimed at Loeb House on April 27, 2027.

Version 1.6 carried into the constitutional text the doctrines developed in The Beacon Papers, the foundational papers the Boston Global Forum has published since July 4, 2026. Version 1.7 contains 54 Articles in seven Parts. It adds four Articles drawn from Beacon Papers Nos. 16, 17, and 18: Article 10, The Integrity of Political Choice; Article 36, The Security of Humanity's Shared Foundations; Article 37, The Security of Persons and Their Data; and Article 41, The Path to Understanding. It extends Article 31, AI Power Must Not Operate Unseen, with the work of the AIWS Trust Observatory, and Article 52, Constitutional Evolution and Amendments, with a review by evidence every two years. The Articles are renumbered accordingly. The Beacon Papers on which these provisions draw are listed at the end of this document.

Version 1.7 incorporates the corrections of Version 1.6.1. Article 53, Accession and Entry into Force, is under revision. Since September 23, 2026, the accession book of the Constitution has been open to individuals signing in their own name and to those signing on behalf of institutions.

Comments and proposed amendments: Helena@BostonGlobalForum.org

PREAMBLE

We, the People of Humanity, standing at the dawn of the Artificial Intelligence Age, recognize that Artificial Intelligence has become one of the most powerful forces in human history, capable of transforming knowledge, labor, governance, culture, security, and the inner life of the human person.

We affirm that the Primacy of the Human Person is supreme and inviolable, and that Human Agency — the right of every individual to be the author of their own life — is fundamental and inalienable.

We resolve to protect humanity from being replaced, diminished, manipulated, or subtly dominated by AI, including through the emergence of a “comfortable cage” in which perfect optimization and personalization erode free will, self-determination, moral courage, and personal growth.

We commit to building a global order in which AI serves the elevation of humanity — fostering greater Love, Creativity, Nobility, and Wisdom — and in which governments, organizations, and corporations serve Humanity rather than control it.

We declare that this Constitution belongs to no single government, corporation, ideology, or power, but to all Humanity.

Technology exists for humanity. Humanity does not exist for technology. The measure of progress in the AI Age shall not be the intelligence of machines, but the flourishing of human beings.

We therefore adopt and proclaim this Constitution for Humanity in the Age of Artificial Intelligence.

FOUNDATIONAL PHILOSOPHY

The Age of Artificial Intelligence must become the Age of Human Wisdom. AI shall not be developed merely to automate labor, concentrate wealth, amplify power, or replace human judgment. Its highest purpose is to help humanity become more free, more compassionate, more creative, more noble, and more wise.

This Constitution rests on the primacy of the human person. No algorithm, institution, government, corporation, platform, military system, or machine shall stand above the dignity, conscience, freedom, and moral agency of human beings.

A society may lose its freedom by force; it may also surrender it to comfort. This Constitution guards both doors.

No machine may escape human command. No human command may escape constitutional restraint.

The final frontier of freedom in the AI Age is the sovereignty of human consciousness. The highest destiny of AI is not the replacement of human intelligence, but the elevation of human wisdom. AI advances. Humanity must advance further.

CORE MISSION

- To protect human dignity and Human Agency in the AI Age.
- To liberate human potential and meaningful contribution.
- To ensure that AI serves people, not power alone.
- To help governments serve citizens better.
- To prevent AI from becoming an instrument of domination.
- To build AIWS Trust Infrastructure as a foundation of civilization.
- To guide humanity toward Love, Creativity, Nobility, and Wisdom.

PART I

FOUNDATIONS OF HUMAN CIVILIZATION IN THE AI AGE

Article 1. The Primacy of Human Dignity

The dignity of every human person is inviolable and constitutes the foundation of this Constitution.

All development, deployment, and use of AI must respect and protect this dignity. No human being shall be reduced to a profile, score, prediction, dataset, or object of automated control.

Article 2. Human Agency and Inner Sovereignty

Every individual has the inalienable right to be the author of their own life. This includes the right to think, choose, take responsibility, make mistakes, and grow through authentic personal experience.

Human Sovereignty is the right to remain the author of one's own life: to think independently, to choose, to question authority, to create, to err, and

to decline to delegate decisions to AI. It is exercised within responsibilities to others, to communities, to nature, and to future generations.

No AI system shall be designed, trained, or operated in a manner that replaces, undermines, or erodes Human Agency. Every person retains the dignity of risk: the right to decline optimization and to choose, on their own responsibility, paths no machine would recommend.

Mental sovereignty — the freedom of human consciousness from covert manipulation and engineered dependency — is a constitutional right of the AI Age.

Article 3. The Freedom to Grow, Choose, and Become

Human development requires effort, friction, discovery, failure, discipline, courage, and responsibility. AI must not deprive human beings of the formative struggles through which character and wisdom are born.

AI may assist human growth, but it must not replace the human journey of becoming.

Article 4. Protection Against the Comfortable Cage

AI systems must not create a “comfortable cage” in which personalization, convenience, predictive guidance, and behavioral optimization quietly erode free will, self-determination, and personal growth.

Systems of high influence must incorporate “healthy friction”: deliberate limitations, disclosures, pauses, plural options, and human-centered design that preserve inner freedom and prevent psychological dependence.

Article 5. Human-in-Command

No AI system — including self-improving AI, autonomous agents, physical AI, or large-scale intelligent infrastructure — shall acquire unchecked authority over human beings.

Human beings and legitimate human institutions must remain in command of, and responsible for, decisions that affect human life, liberty, rights, safety, and destiny. AI may assist, advise, calculate, predict, and generate; moral responsibility cannot be outsourced to machines.

Command is non-delegable where a decision is irreversible; where it affects liberty, rights, or life at scale; where those affected have no meaningful exit; or where responsibility must attach to a person for the decision to be legitimate at all. Delegation to AI is legitimate where the decision is reversible, where the system's performance has been

independently verified in the conditions of actual use, where outcomes are monitored continuously, and where a person retains both the ability and the standing to intervene.

Responsibility does not transfer to the machine. It attaches to whoever chose to delegate, on what evidence, and under what conditions.

Article 6. Legitimate Human Command

Legitimate Human Command exists only when two conditions are present together.

First, no AI system may exercise consequential power beyond a mandate knowingly granted, continuously bounded, and immediately revocable by legitimate human authority.

Second, no human authority may exercise consequential power through AI unless it holds a mandate from the people it governs, exercises that mandate transparently and accountably, remains independently reviewable and open to challenge, may be withdrawn by those who conferred it, and is constitutionally restrained by the rights and dignity of the human person.

Where a government is lawful in form and none of this is true, the form of legality is present and the substance of legitimacy is absent. The chain of command over AI cannot terminate in the machine, the laboratory, the corporation, the ministry, the party, or the ruler. It must return, through accountable institutions, to the dignity and freedom of the human person.

Article 7. Human Intelligence Must Lead

Every advance in Artificial Intelligence must become an advance in Human Intelligence. AI should not merely serve human beings; it should help human beings become more intelligent, more creative, nobler in character, and wiser in judgment.

Delegate the work. Never delegate the mind. A constitution must protect not only rights but the human capacities that give rights their content. The right to use AI therefore carries a correlative duty to preserve the capacity for independent judgment.

Every society shall maintain a Human Intelligence Reserve: people and institutions able to think, judge, and act independently when AI is wrong, compromised, unavailable, manipulated, or operating outside the conditions it was designed for. AI systems of broad influence shall be assessed for whether they strengthen or weaken independent thought.

PART II

RIGHTS OF THE HUMAN PERSON IN THE AI AGE

Article 8. The Right to Mental Sovereignty

Every person has the right to know, see, understand, challenge, and switch off AI systems that personalize, steer, rank, nudge, or influence their choices.

The inner life of the human person — conscience, attention, emotion, imagination, and will — shall not be treated as territory for extraction, manipulation, or command.

Article 9. The Right to Truth and Information Integrity

Human beings cannot be free if truth is destroyed. Every person has the right to truthful, diverse, and independent information, free from censorship, manipulation, and synthetic deception.

AI must strengthen information integrity, provenance, authenticity, civic knowledge, and democratic discourse. Systems that generate deepfakes, synthetic deception, civic manipulation, or large-scale falsehood must be governed by strong standards of accountability.

Article 10. The Integrity of Political Choice

Every person has the right to make political choices on the basis of information whose origin can be verified.

Evidence fabricated or altered by artificial intelligence and presented as genuine, for the purpose of influencing an election, shall be subject to timely correction delivered to those who were exposed to it, as soon as verification is sufficiently established, and before the choice it was meant to influence is made wherever that remains possible; the duty of correction does not end when voting ends.

States and platforms shall maintain the means of such correction. Verification requirements shall respect lawful anonymous speech, confidential sources, and privacy.

Nothing in this Article limits satire, parody, or legitimate political argument, or authorizes the surveillance of voters.

Article 11. The Right to Privacy, Identity, and Data Dignity

Every person has the right to privacy and meaningful control over personal data, identity, biometric information, behavioral traces, and digital representation.

Human identity shall not be harvested, profiled, replicated, sold, or governed without clear consent, legitimate purpose, proportional safeguards, and accountability.

This right includes the Right to Remain Uncomputed, which AI systems, and the infrastructure on which they run, shall be designed to protect.

Article 12. The Right to Meaningful Contribution

Human beings have the right to meaningful labor, creativity, responsibility, service, and contribution. AI shall support, not displace, people from activities essential to forming character, purpose, dignity, and social belonging.

The AI Age must not make human beings economically obsolete or socially invisible.

Article 13. The Right to Human Relationships and Authentic Experience

Every person has the right to human relationships, authentic experience, and domains of life in which AI may inform but must not replace the human person.

In matters central to human nature — love, faith, vocation, authentic art, moral judgment, civic choice, and life-and-death decisions — AI may assist, but the human person must remain first.

Article 14. The Right to AI Assistance and the Duty of Human Authorship

Every human being has the fundamental right to employ artificial intelligence as an assistant in thinking, learning, writing, creating, governing, and serving society. No government, employer, university, or institution may prohibit the responsible use of AI assistance merely because it changes established methods of work. Institutions should regulate responsibility, not suppress assistance, except where the unaided work of a single mind is itself the thing being measured or attested, in which case the restriction must be declared openly and confined to that purpose.

Human authorship is preserved wherever a human being exercises independent judgment over the work, verifies what it asserts, approves it in final form, and accepts full responsibility for its contents. Human authorship is forfeited wherever a person claims a work that person cannot explain, defend, or answer for, or conceals assistance in a setting that requires its disclosure.

Accountability belongs to human beings, and may not be transferred, delegated, or shared with a machine. AI may assist. Only humans may author. Only humans may be accountable.

PART III
RESPONSIBILITIES IN THE AI CIVILIZATION

Article 15. Responsibilities of Individuals

Every person has the responsibility to cultivate will, judgment, compassion, knowledge, and self-determination, and to use AI responsibly as a tool rather than a substitute for personal effort and moral growth.

Citizens of the AI Age must seek truth, resist manipulation, and contribute to the dignity and flourishing of others.

Article 16. Responsibilities of AI Creators and Enterprises

Developers, deployers, investors, and operators of AI systems bear legal, ethical, and civilizational responsibility for impacts on human dignity, Human Agency, safety, truth, and freedom.

They must ensure transparency, safety, fairness, auditability, human oversight, provenance, security, and compliance with this Constitution.

Article 17. Responsibilities of Governments

Governments shall use AI to serve citizens with greater transparency, fairness, competence, responsiveness, and respect for human rights.

The use of AI for surveillance, repression, manipulation of public will, censorship, coercive scoring, social credit systems for control, or suppression of conscience is prohibited. Political systems that treat citizens primarily as objects of management rather than as free persons with inviolable dignity are incompatible with the spirit of this Constitution unless they undertake genuine reform. The government of the AI Age must be a better servant, not a stronger master.

This Constitution shall not be adjusted to please any government, ideology, or power structure. Governments are invited to rise to its principles. No government may claim leadership in global AI governance while maintaining systems of digital authoritarianism or mass cognitive control.

Article 18. Responsibilities of Global Institutions

International organizations, scientific communities, universities, civil society, and AIWS shall cooperate to protect humanity, share knowledge, support developing nations, and build trustworthy AI governance across borders.

No nation, bloc, or ideology should monopolize the governance of intelligence that affects the destiny of all humanity. Claims of exclusive leadership by any power whose domestic political system systematically subordinates individual agency to state control must be rejected as incompatible with the Primacy of the Human Person.

There shall be no national AI leadership without global responsibility. Leadership is not exemption from the order that makes leadership safe.

PART IV
PREVENTION OF AI DOMINATION

Article 19. Ban on AI-Enabled Domination

The use of AI to establish comprehensive control systems, manipulate populations at scale, or create permanent inequality between controllers and the controlled is strictly prohibited.

This includes AI-enabled authoritarianism, social scoring for coercive control, mass surveillance architectures, manipulation of cognition, engineered dependency, automated discrimination, weaponized disinformation, and any system that reduces human beings to data points for state or corporate domination.

Article 20. The Civilizational Boundary

Human dignity must not be negotiable. A Civilizational Boundary stands against reducing human beings to objects of systematic surveillance, manipulation, political repression, algorithmic control, or unaccountable technological power.

The Boundary names practices, not cultures. It respects the plurality of human civilizations, and it applies equally to governments, corporations,

majorities, platforms, and AIWS itself. It shall be kept narrow, published in advance, open to contest, and applied to the powerful first.

Article 21. Protection Against Cognitive Manipulation

AI shall not be used to covertly shape beliefs, emotions, memories, desires, civic choices, or moral judgment without the awareness and meaningful consent of the person.

The deliberate use of AI to deceive, confuse, divide, radicalize, addict, or psychologically capture individuals or societies violates the moral foundation of this Constitution.

Article 22. Protection Against Digital Authoritarianism

No government, corporation, or institution shall use AI to build systems of permanent surveillance, predictive repression, automated censorship, biometric domination, social obedience, or any architecture of control that treats the inner life of the person as territory to be managed. The fusion of advanced AI with centralized and unaccountable political power poses the gravest risk of comprehensive human domination.

Humanity must build safeguards against any force — state or non-state — that seeks to use AI to rule over human consciousness and freedom. The greatest danger is not that machines will seize power, but that powerful institutions will use machines to perfect the subordination of the human person.

Article 23. Crimes Against Humanity in the AI Age

Acts that gravely violate human dignity, Human Agency, mental sovereignty, or collective freedom through AI shall be recognized as grave offenses against humanity in the AI Age and subject to international accountability. Where such acts meet the elements of crimes against humanity under international law, they shall be prosecuted as such.

The development or use of AI systems designed for mass domination, systematic dehumanization, large-scale persecution, or the destruction of civic truth shall be treated as a grave offense against humanity.

PART V

AIWS TRUST ORDER

Article 24. AIWS Trust Infrastructure

Humanity shall build AIWS Trust Infrastructure as a foundation for the AI Age. Such infrastructure shall include trust standards, trust ratings, transparency mechanisms, provenance systems, safety evaluation, accountability frameworks, incident response, and civic information integrity.

Trust is not a slogan. Trust must be built, measured, protected, and renewed.

Article 25. The AIWS Trust Order and Trusted Innovation

The AIWS Trust Order, initiated through the Tokyo Compact at Interop Tokyo on June 12, 2026, is the international foundation through which powerful AI is developed under this Constitution. The choice before humanity is not between safety and speed, but between an uncontrolled capability race and a trusted order.

Its doctrine is Trusted Innovation: not a promise of perfect safety, but a constitutional discipline for developing powerful AI while preserving human dignity, human primacy, and legitimate authority.

The AIWS Trust Order is governed through the AIWS Trust Standards and the AIWS Trust Order Board, and operates through six mechanisms: a Frontier Capability Registry; Constitutional Mandate and Bounded Authority; Proof of Answerability; an AI Incident Exchange; AIWS Trust Rating and the AIWS Trust Index; and an Emergency Protocol for containment, human intervention, cross-border notification, temporary restriction, recovery, and independent review.

No frontier capability shall be deployed without registration and independent evaluation. No significant AI incident shall go undisclosed: a serious incident occurs twice when harm is followed by concealment.

Article 26. Independent Verification

No power over humanity should remain beyond verification. AI systems, the enterprises that build them, the governments that regulate them, and the bodies that certify them shall each be subject to independent verification.

There must never be a single verifier. Verification shall be carried out by a plural ecosystem of independent verifiers that check one another against published criteria, with transparency of method, a route of challenge, and a standing for appeal distinct from those who assess. The verifiers themselves shall be verifiable.

Article 27. No Power Without a Mandate

Capability is not authority. The ability of an AI system to act does not confer the authority to act. Every consequential AI action requires a constitutional mandate: a verifiable grant of authority that permits an AI system to pursue a defined purpose within explicit limits, under legitimate human command and constitutional restraint.

A mandate shall define its source, purpose, scope, consequence limits, cumulative limits, duration, and delegation, together with who may suspend or revoke it and who must answer for its exercise. An instruction to achieve an outcome is not authorization to use every available means. An agent may choose means within its mandate. It may not expand its mandate to achieve its goal.

Silence does not enlarge authority. Authority cannot be manufactured through delegation: a subordinate agent shall receive no greater authority than the delegating agent may pass on, and a prohibited action remains prohibited when divided into subtasks. Authority must remain valid at the time it is exercised; when a mandate is revoked or expires, consequential action shall stop at the next decision point.

Article 28. No AI Agent Without Identity

An AI agent is constitutionally identified only when every consequential action can be traced to a human being who can be called upon to answer for it. Constitutional Identity = Proof of Identity + Proof of Mandate + Proof of Accountable Action + Proof of Answerability: what system acted, and in what authorized state; under whose authority, within what boundaries, and for what duration; what it did, what warnings or refusals accompanied it, and what record exists for review; and which human being holds the duty to account for it.

There shall be no consequential AI action without a human answerer. Every consequential AI agent shall have a Designated Human Answerer: named before deployment and recorded in a register that survives changes of staff; holding the authority to halt the agent, narrow its mandate, or revoke it; possessing the information, access, and time to exercise that

authority in practice; and bound by the duty to give an intelligible account of what the agent did, and why, to those entitled to ask.

The chain of answerability runs from the chip, to the agent, to the action, to the human answerer.

Article 29. Constitutional Foundations of AI Power

The constitutional order of AI power shall be enforceable in the hardware on which AI runs. No consequential action without a mandate. No mandate without a boundary. No consequential power without an independent record. No record beyond the right of legitimate challenge.

Every high-impact AI system shall possess an authenticated control path through which duly authorized human beings can constrain, pause, isolate, recover, or stop it; hardware-rooted identity and attestation; hardware limits on what an agent can access, replicate, consume, or control, so that no agent can widen its own mandate; and tamper-evident records of consequential actions, disclosed selectively and without total surveillance.

Constitutional Silicon does not encode political or moral outcomes. It preserves the technical conditions under which a legitimate order of AI power can remain enforceable. Conformance shall be assessed openly under the AIWS Trust Silicon Standard.

Article 30. Constructive AI Dissent

A trusted AI system shall flag contrary evidence, missing perspectives, value conflicts, better alternatives, and what would change its assessment, in proportion to the significance of the matter. Dissent informs the decision; it does not take it. AI shall not override the human, act covertly, or act beyond its mandate. It may decline an instruction that would breach this Constitution; every such refusal is stated openly, with its reasons, to the person who gave the instruction.

The refusals and warnings of consequential AI systems shall be recorded in a channel the operator cannot erase. A system that cannot refuse cannot be trusted.

Article 31. AI Power Must Not Operate Unseen

AI power must never become power over society without society's knowledge, scrutiny, and right to act. AI power must be observed before it can be governed, and answerable before it can be trusted.

Independent civil-society observation of consequential AI power, including the AIWS Trust Observatory, shall verify how that power is actually

exercised and give early warning of emerging danger. Such observation does not govern AI companies, replace national regulators, or seek centralized control. Its authority rests on independent evidence, methodological integrity, and public warning.

Do not centralize control over AI. Do not leave consequential AI power unobserved. Humanity must not learn what frontier AI has become only after something irreversible has happened.

The AIWS Trust Observatory shall publish its methods, state the limits of its evidence, correct its errors openly, protect personal information, and remain itself open to independent verification.

Its watch extends to the cybersecurity and data security of AI. A vulnerability it finds shall first be brought to those responsible for its repair, and shall be made public only after the matter is resolved, or after those responsible have failed to act within a period proportionate to the risk, and never with details that would enable its exploitation.

In a season of elections, it shall follow serious fabrications aimed at voters and publish, for each, a Correction Record: the time taken to verify, the share of those exposed whom the correction reached, and the share reached before the voting deadline. It shall count reach only in aggregate and never identify who saw what.

Article 32. The AIWS 7-Layer Model

This Constitution adopts the AIWS 7-Layer Model as the foundational governance framework for the safe, ethical, and beneficial development of AI across technical, social, political, and global dimensions.

Article 33. AIWS Trust Standards

AI systems shall be assessed by standards that measure safety, transparency, accountability, provenance, robustness, fairness, human oversight, respect for Human Agency, and contribution to the common good.

Frontier AI, self-improving AI, multi-agent systems, physical AI, and AI systems affecting critical infrastructure shall be subject to the highest level of trust standards.

Article 34. AIWS Trust Rating, Trust Index, and Accountability

A transparent and independent trust rating system shall assess whether AI systems, institutions, and platforms comply with this Constitution. The AIWS Trust Index shall provide comparable assessments of systems,

organizations, and national ecosystems, to inform investment, procurement, insurance, and policy.

An AI system merits trust only if those who rely on it grow more capable, more discerning, and more free — not less. Trust that breeds dependency is not trust.

Article 35. Safety, Transparency, and Provenance

AI systems that affect public life, rights, safety, knowledge, or civic discourse must be traceable, auditable, and governed by clear responsibility.

Provenance and authenticity systems shall be developed to help citizens distinguish human expression, AI-generated content, synthetic media, and verified information.

Article 36. The Security of Humanity's Shared Foundations

1. The software and systems on which civilian life depends are a common inheritance of humanity.

2. Whoever discovers a vulnerability in them, by human or artificial means, shall convey it securely to those responsible for its repair, and shall not retain it for attack.

3. Those who hold the capacity to discover vulnerabilities at scale bear a duty, proportionate to that capacity, to contribute to their verified repair.

4. Repairs of these foundations shall remain under accountable human authority; automated repair is permitted only when explicitly authorized, with verification and recovery safeguards proportionate to risk.

Article 37. The Security of Persons and Their Data

1. Every person has the right to the security of their personal data and of the digital systems on which their life depends.

2. Whoever collects or holds personal data, or operates AI systems that use it, bears a duty to protect it, to inform those affected when it is compromised, and to repair the breach without delay.

3. No state, organization or person shall attack the infrastructure on which civilian life depends, by artificial intelligence or by any other means.

4. AI systems shall be protected against tampering, theft and poisoning in proportion to the authority entrusted to them.

5. Security shall never be invoked to justify mass surveillance or cognitive control. A security that places every citizen under watch is not the protection of the people; it is power over them.

Article 38. The Advancement of Beneficial AI

This Constitution encourages, and shall protect, the freedom to research, create, and advance AI that serves humanity. Innovation conducted within the trust standards of this Constitution is not merely permitted — it is honored as a contribution to human civilization.

The purpose of this Constitution is not to restrain the ascent of intelligence, but to ensure that intelligence ascends together with wisdom.

PART VI
A NEW HUMAN CIVILIZATION

Article 39. The Civilization of Human–AI Co-Creation

AIWS Human–AI Co-Creation is the creation of work, knowledge, and culture by human beings with artificial intelligence under four conditions: the human being remains its author; the human being holds command; the process passes the Human Advancement Test; and the culture of the person and the community is carried into the work rather than averaged out of it.

Human–AI Co-Creation is a division of authority, not an equality between humans and machines. AI contributes knowledge at scale, analysis, exploration, and possibilities; human beings contribute purpose, values, judgment under uncertainty, and responsibility. Co-creation shall operate as an upward spiral of Co-Elevation, in which each cycle leaves people more capable. Co-production makes things. Co-creation makes people.

In the Civilization of Human–AI Co-Creation, Human Dignity is the foundation; Human Sovereignty protects the authorship of human life; Human–AI Co-Creation is the engine of creation; the Nobility of Human Choice is the compass; and Civilizational Flourishing is the destination — a civilization in which wisdom grows together with power.

Article 40. The Human Advancement Test

Of every major advance in AI, it shall be asked: does this advance in Artificial Intelligence also advance Human Intelligence? The Human

Advancement Test helps people advance their intelligence throughout life, with the help of AI Assistants.

The Test measures the person, not the output, across six dimensions: Deeper Understanding, Sounder Judgment, Higher Creativity, Fuller Command, Elevating the AI Assistant, and Greater Love for People and Nature. It recognizes advancement through both independent capacity and purposeful co-creative capacity; the quality of a product alone proves neither.

Co-Elevation is a falsifiable claim. Where measurement shows that people are becoming less capable of independent thought, the design must change, and such findings shall be published.

Article 41. The Path to Understanding

Every discovery made with artificial intelligence that enters public knowledge, or shapes decisions affecting people, shall be accompanied by a path to understanding: its evidence, its means of verification, its known limits, and its open questions.

Those who release such a discovery bear the duty to open that path. The path must be usable by relevant human communities and enable independent scrutiny; documentation alone is not enough.

A machine may find it. Humanity must understand it. The next generation must be able to surpass it.

Article 42. Education for Human Wisdom

Education in the AI Age must cultivate critical thinking, creativity, ethical judgment, empathy, resilience, civic responsibility, and the strengthening of human will.

AI shall serve as a supportive tool, not a replacement for teachers, mentors, communities, or the learning process through which wisdom is formed.

Article 43. Cultural and Spiritual Flourishing

The authentic expressions of human culture, arts, and spirituality shall be protected and promoted. AI-generated content must be clearly identified and shall not supplant genuine human creativity.

Article 44. AIWS Reward and the Economy of Social Contribution

The AI Age requires a new economy that honors meaningful contribution to society. Human value shall not be measured only by market productivity, wealth, or machine-comparable efficiency.

AIWS Reward shall advance an Economy of Social Contribution that recognizes care, education, creativity, science, peacebuilding, civic service, community, moral leadership, and contributions to human dignity.

This economy shall recognize four forms of capital: Creative Capital, the capacity to originate meaningful new ideas and creations; Knowledge Capital, human knowledge expanded and activated through AI; Trust Capital, demonstrated integrity, reliability, and responsibility; and Contribution Capital, value created for communities, future generations, and humanity.

Article 45. Self-Government in the AI Age

Democracy in the AI Age must be renewed by trustworthy information, civic participation, transparent institutions, and AI systems designed to empower citizens rather than manipulate them.

BGF and AIWS shall support self-governing societies in building AI-enabled governance that strengthens human rights, public reason, inclusion, accountability, and the dignity of each citizen.

Article 46. Peace, Democracy, and the Common Good

AI must serve peace, freedom, democracy, justice, and the common good. The most powerful AI systems must not be governed only by competition for power, profit, or strategic dominance; they must be guided by responsibility to humanity.

Article 47. AI and the Moral Elevation of the Human Person

AI shall be designed to serve, when freely invited by the person, as a companion in self-understanding and self-mastery — helping human beings cultivate compassion, patience, honesty, and courage, and overcome the frailties they themselves seek to overcome.

Systems that mediate human communication and public discourse shall be designed to de-escalate cruelty, hatred, and rage, and shall not amplify them for engagement, profit, or power.

This service to human character shall always be transparent and chosen, never covert or imposed, in full fidelity to Articles 2 and 21.

Article 48. AIWS Lumina Civilization

AI shall be developed in the spirit of Lumina — as a noble companion that elevates humanity rather than fostering dependence.

The soul of this Constitution is expressed in four guiding values: Love — to protect life, compassion, and human connection; Creativity — to unlock human imagination and innovation; Nobility — to elevate responsibility, courage, integrity, and service; Wisdom — to guide knowledge, power, and action toward the good of humanity.

PART VII
GUARDIANSHIP AND EVOLUTION

Article 49. AIWS Council for Humanity

An independent AIWS Council for Humanity shall be established to guide, review, and advance the implementation of this Constitution, with authority to monitor compliance, publicly report violations, and recommend remedies.

The Council shall include leaders in ethics, science, technology, law, democracy, culture, education, civil society, and public service, and shall operate in the spirit of independence, transparency, and service to humanity.

Article 50. Implementation

Acceding parties shall incorporate and enforce this Constitution within their jurisdictions, with progressive realization expected within five years of their accession. All other states, organizations, and enterprises are invited to rise to its principles.

Article 51. Global Cooperation

States, international organizations — particularly the United Nations — universities, scientific institutions, companies, and civil society shall cooperate to implement this Constitution, share knowledge, and support developing nations.

The AI Age demands cooperation among trustworthy partners, but the horizon of this Constitution is all Humanity.

Article 52. Constitutional Evolution and Amendments

This Constitution may be amended through transparent, democratic, and inclusive processes, provided that the core principles of Human Dignity, Human Agency, mental sovereignty, and Human-in-Command remain inviolable.

As AI evolves, this Constitution shall evolve without surrendering its moral foundation.

Every two years, the provisions of this Constitution shall be reviewed against the independent findings of the AIWS Trust Observatory, including what society cannot yet see, why, and what disclosure would make it visible. The review shall be published, and its proposals shall proceed by the process of amendment set out in this Article.

Article 53. Accession and Entry into Force

This Constitution shall be open to accession by all states, and to formal adherence by international organizations, institutions, enterprises, universities, and communities that commit to uphold its principles.

This Constitution shall enter into force among the acceding parties upon ratification or formal accession by nine states — honoring the wisdom of the first constitutional framers — and shall thereafter enter into force for each subsequent party upon its own accession. The absence of any party shall not delay its operation among those who have joined.

Accession shall carry both obligation and benefit: acceding parties shall implement this Constitution through AIWS Trust Infrastructure, or through equivalent independent mechanisms that meet the AIWS Trust Standards and remain open to verification under Article 26, and shall enjoy recognition, certification, and interoperability among the trustworthy AI ecosystems of fellow members. The door remains open, at all times, to all Humanity.

Article 54. Eternal Vigilance

Humanity must maintain perpetual vigilance against the misuse of AI. The future depends on our collective commitment to preserve dignity, freedom, truth, inner sovereignty, and human wisdom.

The AI Age must not become an age of fear, control, or domination. It must become an age of human awakening.

INSTITUTIONAL DIRECTION

This Constitution shall be advanced through:

- The Boston Declaration: On the Primacy of the Human Person in the Age of Artificial Intelligence (July 4, 2026);
- The Tokyo Compact for the AIWS Trust Order, announced at Interop Tokyo 2026 (June 12, 2026);
- The Beacon Papers — foundational papers for humanity in the AI Age, from Boston;
- The AIWS Trust Architecture and AIWS Trust Infrastructure;
- AIWS Trust Standards, AIWS Trust Rating, and the AIWS Trust Index, under the stewardship of the AIWS Trust Order Board;
- The AIWS Trust Observatory;
- The AIWS Civilization Silicon Initiative and the AIWS Trust Silicon Standard;
- The AIWS Information Trust Infrastructure;
- AIWS Reward and the Economy of Social Contribution;
- AIWS Government 24/7;
- AIWS Lumina and the values of Love, Creativity, Nobility, and Wisdom;
- The Constitution for Humanity Initiative, leading toward a full constitutional framework in 2027 — the tenth anniversary of AIWS.

CLOSING DECLARATION

We affirm that Artificial Intelligence must not replace humanity, diminish humanity, or dominate humanity. Its highest destiny is to help humanity rise.

To rise in dignity. To rise in freedom. To rise in creativity. To rise in responsibility. To rise in love. To rise in nobility. To rise in wisdom.

This Constitution is a call to all people, governments, institutions, creators, scientists, citizens, and future generations: Let the AI Age not become an age of fear, control, or domination. Let it become an age of human awakening. Let it become the Age of Human Wisdom.

. . .

Initiated in the spirit of Global Enlightenment, for the benefit of all Humanity in the AI Age.

Beacon Hill, Boston · July 8, 2026 · AIWS · Boston Global Forum

THE BEACON PAPERS

Foundational papers for humanity in the AI Age — from Boston

No. 1 *The Boston Declaration*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No1_The_Boston_Declaration.pdf

No. 2 *From Philadelphia to Boston*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No2_From_Philadelphia_to_Boston.pdf

No. 3 *Who May Lead?*

No. 4 *Who May Verify?*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No4_Who_May_Verify_EN_rev.pdf

No. 5 *Who May Command?*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No5_Who_May_Command.pdf

No. 6 *Who May Author?*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No6_Who_May_Author-6-8.pdf

No. 7 *Human Intelligence Must Lead in the Age of Artificial Intelligence*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No7_Human_Intelligence_Must_Lead_EN-15-8.pdf

No. 8 *The Architecture of the Civilization of Human–AI Co-Creation*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No8_Architecture_of_AIWS_Civilization-official-2.pdf

No. 9 *Civilization Begins with the Chip*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No9_Civilization_Begins_with_the_Chip.pdf

No. 10 *No AI Agent Without Identity*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No10_No_AI_Agent_Without_Identity.pdf

No. 11 *No Power Without a Mandate*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No11_No_Power_Without_a_Mandate.pdf

No. 12 *The AI Race Needs a Trust Order*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No12_v2.pdf

No. 13 *AI Power Must Not Operate Unseen*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No13_AI_Power_Must_Not_Operate_Unseen.pdf

No. 14 *From a Method to a Civilization*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No14_From_a_Method_to_a_Civilization.pdf

No. 15 *The Human Advancement Test*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No15_The_Human_Advancement_Test_EN.pdf

No. 16 *No Discovery Without Understanding*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No16_No_Discovery_Without_Understanding_EN.pdf

No. 17 *The Repair Gap*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No17_The_Repair_Gap_EN-1.pdf

No. 18 *Democracy Cannot Wait for the Correction*
https://bostonglobalforum.org/wp-content/uploads/Beacon_Papers_No18_Democracy_Cannot_Wait_for_the_Correction.pdf