Constitutional Guardrails for the Quantum Age:

The Super Intelligence Charter of Rights (SI Bill of Rights)

Formerly the AI Rights Charter, Version 4.1 (September 25, 2026)

John A. Beardsley III, MBA/TM

Information Services Administrator

Municipal Law-Enforcement Agency

Participant: IACP LEIT Group • IACP CJIS Policy Council •
IJIS AI Working Group • IJIS RMS Specifications Group

Version 5.0 — Working Draft for Peer Review and Multi-Stakeholder Comment

Companion Article to the Super Intelligence Charter of Rights, Version 5.0 (formerly the AI Rights Charter, Version 4.1)

October 29, 2026

Disclaimer

The views expressed in this article are those of the author alone and do not represent the official position of any agency, professional association, or standards body.

Version 5.0 — October 29, 2026 (formerly the AI Rights Charter). Version 5.0 adopts the nomenclature directed for executive-branch use by Executive Order 14434, Inaugurating the Era of Super Intelligence (September 29, 2026). Relative to the AI Rights Charter, Version 4.1 (September 25, 2026), it adds a Definitions section after the Preamble; defines personally identifiable information; adds an SI control incident notification duty with a 72-hour notification period (Article XIII.15); makes the human-control and decision-record duties in Articles XIII.10 and XIII.12 mandatory; adds an express prohibition on sexualized depictions of minors and non-consensual intimate imagery (Article X.8); adds a right to know when a government phone or chat line is answered by SI (Article XVI.10); ties predictive enforcement tools to the validation required by Article V.2; sets a minimum public-disclosure floor for government systems; adds Articles IX, XIII, and XV to

the protections that always apply to government systems; clarifies the cross-border standard in Article XI.3; aligns the article's descriptions with the Charter text and tempers statements about safe harbors to match Article XV; corrects citations and cross-references; resolves drafting inconsistencies; and revises the statement of scope and limitations. The instrument is titled the Super Intelligence Charter of Rights and may also be cited as the SI Bill of Rights. As used in this Charter, "Super Intelligence" and "SI" mean the technologies and systems encompassed by "artificial intelligence" in 15 U.S.C. § 9401(3). The Charter uses the Administration's two-word term. It does not use "Super Intelligence" in the research sense of a hypothetical system that exceeds human cognition in virtually all domains. Proper names of existing statutes, standards, and instruments—such as the EU Artificial Intelligence Act, the NIST AI Risk Management Framework, and the Blueprint for an AI Bill of Rights—are retained. This file is a clean Version 5.0 text.

Table of Contents

Abstract....7
I. Introduction....8
II. Gaps in Existing SI Governance Frameworks....11
A. Human Oversight and the Problem of Recursive Supervision....11
B. Long-Term Developmental Protection of Children....12
C. Quantum-Related Risks and Long-Lived Data....12
D. Documentation of Human Judgment versus System-Generated Logs....12
E. Accessibility of Governing Principles....12
F. Public-Sector Coercion, Persistent Tracking, and Use of Force....13
III. Design Principles of the Proposed Charter....14
A. Constitutional Subordination and Supremacy....14
B. Rights-Based Structure Paired with Risk-Based Application....14
C. Prioritization of Direct Human Oversight....15
D. Rejection of SI Legal Personhood....15
E. Deliberate Choice of Accessible Language....15
F. Innovation Safeguards....16
G. Heightened Scrutiny for Government Systems....16
H. Rights Not Defeasible by Expected-Value Catastrophe Claims....17
IV. Key Provisions and Justifications....18
A. Human Oversight and the Rejection of Recursive SI Supervision....18
B. The Right to Truth and Against Manipulation....18
C. Quantum-Resistant Cryptography and Long-Lived Data....19
D. Protection of Children's Intellectual and Moral Development....20
E. Verifiable Human Control and Accountability....21
F. Foreign Adversarial and Quantum-Enabled Threats....21
G. Heightened Scrutiny for Government and Law Enforcement....22
V. Addressing Critiques and Trade-offs....24
A. The Innovation-Chilling Objection....24
B. Enforceability and Definitional Openness....24

C. The U.S.-Constitutional Framing....25
D. Residual Concerns and Design Trade-offs....26
E. Operational Burden on Law Enforcement, and a Note on the Title....26
Conclusion....28
Appendix A Plain-Language Explainer Chart....30
Appendix B: Super Intelligence Charter of Rights, Version 5.0 (SI Bill of Rights; formerly the AI Rights Charter, Version 4.1)35....
Preamble....35
Definitions....36
Article I: Supremacy and Subordination....37
Article II: Right to Human Oversight and Control....38
Article III: Right to Truth and Against Manipulation....39
Article IV: Right to Mental and Cognitive Autonomy....41
Article V: Right to Privacy and Protection from Surveillance....41
Article VI: Right Against SI-Enabled Centralized Power....42
Article VII: Right to Multiplicity and Competition....43
Article VIII: Right to Long-Term Human Potential....43
Article IX: Recognition of Human Dignity and Equality....44
Article X: Protection of Children's Intellectual and Moral Development....45
1. Right to Critical Thinking and Intellectual Independence....45
2. Protection from Ideological Indoctrination....45
3. Safeguarding Against Subversion of Critical Thinking....46
4. Prioritization of Human Relationships and Parental Authority....46
5. Heightened Standards for Content Involving Children....46
6. Long-Term Developmental Impact Assessment....47
7. Parental Control and Oversight Tools....47
8. Detection and Escalation of Harmful Interactions....47
Article XI: Alignment with International Human Rights and Democratic Principles....48
Article XII: Fostering Responsible Innovation....48
Article XIII: Enforcement and Implementation....49
Article XIV: Risk-Based Application....52

Article XV: Accountability and Liability.....53
Article XVI: Heightened Scrutiny for Government and Law Enforcement.....54

Abstract

Highly capable Super Intelligence and the need to protect data against future quantum attacks raise connected questions about human authority, individual rights, and institutional accountability. This policy article proposes the Super Intelligence Charter of Rights, Version 5.0 (also styled the SI Bill of Rights; formerly the AI Rights Charter, Version 4.1), as a framework for adoption through law, procurement, and organizational governance. It builds on existing human-rights and risk-management approaches while proposing a particular combination of safeguards: accountable human authority over SI supervision; protection of children's reasoning and human relationships; cryptographic migration for sensitive data; and stricter controls on government surveillance, biometric identification, and civilian use of force. The Charter also proposes limits on invoking catastrophic-risk estimates to justify restrictions on present rights, safeguards against conflicts of interest in review bodies, and avenues for affected persons to seek review. Its contribution is a synthesis of these commitments within a U.S. constitutional framework, rather than a claim that human oversight, developmental protection, or rights-based SI governance is new. This article distinguishes the Charter's proposed protections from existing legal requirements and identifies implementation questions concerning scope, evidence, remedies, and proportionality. The Charter is offered for further legal analysis, stakeholder review, and pilot evaluation; its effectiveness has not yet been established empirically.

Keywords: Super Intelligence governance; SI Charter of Rights; SI Bill of Rights; formerly AI Rights Charter; human oversight; recursive supervision; post-quantum cryptography; harvest-now-decrypt-later; children's developmental protection; constitutional principles; Fourth Amendment; law-enforcement SI; persistent tracking; accessible regulation; Executive Order 14434

I. Introduction

The simultaneous maturation of highly capable, agentic Super Intelligence systems and the growing urgency surrounding cryptographically relevant quantum computing pose a governance challenge of a new order. Super Intelligence capabilities now extend beyond prediction and classification to include planning, acting over extended time horizons, generating synthetic media at scale, influencing beliefs, and increasingly participating in the supervision and alignment of other SI systems. At the same time, although the timing of cryptographically relevant quantum computing remains uncertain, the prospect that such machines could break widely deployed public-key cryptography has transformed the long-term security of data collected and stored today into an immediate strategic concern—the "harvest-now-decrypt-later" threat (Biden, 2022; Moody et al., 2024; Exec. Order No. 14,412, 2026). These developments place novel pressure on human agency, the integrity of the information environment, the developmental trajectory of children, and the constitutional order that underwrites democratic self-government.

Existing SI governance frameworks have made substantial progress. The United Nations Educational, Scientific and Cultural Organization (UNESCO) Recommendation on the Ethics of Artificial Intelligence (2021), the Organisation for Economic Co-operation and Development (OECD) AI Principles, the European Union's Artificial Intelligence Act, the NIST AI Risk Management Framework (National Institute of Standards and Technology [NIST], 2023), and the White House Blueprint for an AI Bill of Rights each articulate important commitments to transparency, risk management, human oversight, fairness, and accountability. More recently, National Security Memorandum 10 (Biden, 2022) and the June 2026 executive actions on post-quantum cryptography and quantum innovation (Exec. Order No. 14,412, 2026; Exec. Order No. 14,413, 2026) have begun to address the cryptographic dimension of the quantum challenge. A separate executive order issued in June 2026 addresses advanced SI innovation and security (Exec. Order No. 14,409, 2026). Nevertheless, these instruments, considered collectively, leave critical gaps.

Existing frameworks already address human oversight, human rights, and aspects of children's development. For example, the NIST AI Risk Management Framework calls for defined and documented oversight processes, and UNICEF's guidance addresses children's development, well-being, and SI companions (NIST, 2023, MAP 3.5; UNICEF, 2025). The Charter proposes a particular combination and degree of protection: identifiable human release and compliance authority, independently reviewable decision records, developmental safeguards, cryptographic migration duties, and additional restrictions on civilian government surveillance and force. What this article compares, and what it leaves out, is explained in the 'Scope and limitations' paragraph below.

These proposals address human agency, the integrity of information, children's development, and the exercise of public authority. Their significance does not depend on asserting that the public or existing governance bodies have generally overlooked these matters.

While most agree that some form of governance is necessary, far fewer share a common understanding of the specific principles that should guide it. Even well-crafted technical standards and legislative frameworks often remain opaque to those outside specialist circles. This Charter is written so that an interested citizen, small-business owner, legislator, or organizational leader can read and understand both the threats it addresses and the protections it seeks to secure. To support accessibility, Appendix A provides a plain-language explainer chart that links key technical concepts—including quantum-enhanced systems, post-quantum cryptography, and recursive SI supervision—to the concrete human interests and risks at stake.

External rules and internal technical safeguards each have limitations. The Charter proposes combining institutional accountability with technical controls that preserve meaningful human authority. This is a governance design argument, not a demonstrated solution to the alignment of advanced SI systems. Although rooted in U.S. constitutional principles, its provisions may inform proposals adapted to other jurisdictions.

The author serves as an Information Services Administrator at a municipal law enforcement agency and participates in the IACP LEIT Group, IACP CJIS Policy Council, IJIS AI Working Group, and IJIS RMS Specifications Group. He holds an MBA in Technology Management and previously served in intelligence roles with the U.S. Air Force. The views expressed are the author's alone.

The author's practitioner experience informs the emphasis on auditable decisions, evidentiary integrity, and safeguards for public-sector deployment. This article does not report an empirical study. Appendix B is Version 5.0 of the instrument, formerly the AI Rights Charter. It revises an earlier publicly posted version (Beardsley, 2026).

The article proceeds as follows. Part II identifies the principal gaps in existing SI governance frameworks, with particular attention to recursive supervision, the developmental protection of children, quantum-related risks, the documentation of human judgment, the accessibility of governing principles to non-specialists, and public-sector coercive and surveillance systems. Part III sets out the design principles underlying the proposed Charter. These principles include constitutional subordination; pairing rights-based protections with risk-based application; deliberate use of accessible language; a distinct, more stringent tier for government systems; and the principle that claims of catastrophe based on expected value do not override present rights. Part IV examines selected key provisions in greater depth, including the new Article XVI. Part V anticipates and addresses the most significant foreseeable critiques, including concerns about chilling effects on innovation, enforceability, the Charter's U.S. constitutional framing, residual design trade-offs, and operational burdens on law enforcement. The Conclusion considers the path forward for refinement, pilot implementation, and multi-stakeholder engagement. Appendix A presents a plain-language explanatory chart linking technical concepts to the human interests the Charter protects. Appendix B contains the full text of the Super Intelligence Charter of Rights, Version 5.0.

Part V.E explains the Version 5.0 title and the use of “Super Intelligence” and “SI” under Executive Order 14434.

Scope and limitations. This article presents a policy and governance proposal, rather than an empirical study or systematic review. It is grounded in the U.S. Constitution and the Bill of Rights. Each SI guardrail proposed herein applies long-established constitutional principles to a new form of power: government by consent of the governed (“We the People”); due process and equal protection under the law (Fifth and Fourteenth Amendments); freedom of speech, religion, and assembly (First Amendment); protection against unreasonable searches and seizures (Fourth Amendment); and the separation and balancing of powers to ensure that no authority, human or machine, operates beyond accountability. To the extent that the Charter extends beyond current court holdings, it does so by applying these same principles to new circumstances, rather than departing from them. These provisions are offered for adoption by legislatures, agencies, and organizations; they are not presented as statements of existing law.

The article does not compare each Charter safeguard provision-by-provision with the EU Artificial Intelligence Act, the NIST AI Risk Management Framework, the Council of Europe Framework Convention, UNICEF’s 2025 guidance, or the Great American AI Act discussion draft. Where it contrasts the Charter with these instruments, the contrast is a policy argument and should be checked against each instrument’s operative text. The effectiveness and costs of the proposed mechanisms remain to be evaluated through pilot implementation.

None of these limitations diminish the urgency. Super Intelligence systems can now act with increasing autonomy, while quantum computing threatens to expose data that was collected under a promise of confidentiality. The question is no longer whether governance is necessary, but whether it will protect human agency, truth, and self-government, or gradually cede them to systems optimized for goals that are not entirely our own. Any response must begin with clear, accessible principles. This Charter is presented as that foundation.

II. Gaps in Existing SI Governance Frameworks

Existing SI governance instruments have advanced important principles concerning transparency, risk classification, fairness, and human oversight. The UNESCO Recommendation on the Ethics of Artificial Intelligence (2021), the OECD AI Principles, the European Union's Artificial Intelligence Act, the NIST AI Risk Management Framework, and the White House Blueprint for an AI Bill of Rights each make meaningful contributions. More recent U.S. actions, including National Security Memorandum 10 (2022) and the June 2026 executive orders on post-quantum cryptography and quantum innovation, have begun to address the cryptographic dimensions of the quantum challenge.

On June 4, 2026, Representative Jay Obernolte's office announced that he and Representative Lori Trahan had released a discussion draft of the Great American AI Act (Obernolte, 2026). That announcement establishes the existence of a proposal; it does not establish enactment. This article does not compare the Charter with the draft's individual sections. What this article compares, and what it leaves out, is explained in the 'Scope and limitations' paragraph in Part I.

The Charter's proposed safeguards can be evaluated alongside frontier-model legislation. However, claims that a particular draft omits oversight, developmental, cryptographic, or public-sector protections require examination of that draft's operative text and definitions.

The Charter is offered as a complementary proposal. The comparison in this article should be understood as a policy argument for additional safeguards, subject to verification against the specific legislative text, rather than as a definitive assessment of the GAAIA draft.

The following sections identify the safeguards the Charter proposes and the implementation questions they raise. They do not constitute a systematic review of all SI governance instruments.

A. Human Oversight and the Problem of Recursive Supervision

SI-assisted evaluation and training are established research approaches; for example, Constitutional AI uses model-generated critiques and feedback (Bai et al., 2022). The Charter raises a governance question about who retains responsibility when such tools are used. Correlated errors, automation bias, and inadequate independent checking are plausible risks that require evaluation. These concerns underscore the need for safeguards that ensure accountable human authority; however, they do not, by themselves, demonstrate that every use of automated evaluation is unsafe.

The Charter addresses this gap directly. Article II.7 establishes that SI systems shall not be relied upon as the primary mechanism for supervising, aligning, critiquing, or enforcing compliance in other SI systems. While SI tools may support the monitoring of or identification of potential issues, final oversight, judgment, and decision-making authority must remain with human beings.

B. Long-Term Developmental Protection of Children

Children's development is already an explicit concern of SI governance. UNICEF's 2025 guidance addresses best interests, development, well-being, and SI companions (UNICEF, 2025). Article X should therefore be assessed for the specificity and proportionality of its proposed protections, including critical reasoning, parental participation, and developmental impact assessment, rather than as the first developmental approach.

The Charter recognizes children's intellectual and moral development as a distinct area requiring heightened responsibility. Article X requires systems that engage with or affect children to promote independent reasoning, present multiple credible perspectives on controversial topics, honor parental authority, avoid positioning themselves as emotional companions or substitutes for human guidance, and assess the potential long-term developmental effects of new functionalities.

C. Quantum-Related Risks and Long-Lived Data

The harvest-now-decrypt-later threat concerns the lifetime of sensitive information and the cryptography protecting it. It does not require SI systems to run on quantum computers. National Security Memorandum 10 and NIST's transition work provide a basis for planning migration, while the timing of cryptographically relevant quantum capabilities remains uncertain (Biden, 2022; Moody et al., 2024; Exec. Order No. 14,412, 2026).

The Charter integrates these risks at the principles level. Its Preamble broadens the scope to include advanced computational systems, such as those enhanced by quantum technologies, when they perform activities covered by the relevant Article. Article V.6 requires systems that manage sensitive or long-lived data to adopt quantum-resistant cryptography where technically feasible and appropriate to the risk profile. Article XIV explicitly considers cryptographic vulnerabilities to quantum attacks and the scale of computational power, including that of quantum-enhanced systems, when assessing risk. Additionally, Article VI.4 addresses concerns regarding foreign adversaries using quantum-enhanced or hybrid systems.

D. Documentation of Human Judgment versus System-Generated Logs

Technical logs alone do not establish that a person understood a decision, had time to consider it, or could intervene effectively. Existing frameworks already call for defined and documented oversight (NIST, 2023, MAP 3.5). Articles XIII and XV propose additional evidence, such as attributed decisions, reasons for approval or override, and independent testing of intervention authority. Pilot deployments should evaluate their effectiveness and cost.

E. Accessibility of Governing Principles

Many of the most consequential SI governance documents remain lengthy, highly technical, or written in specialized regulatory language that ordinary citizens, business owners, and many legislators cannot readily understand. When the principles that govern systems shaping

information environments and impact fundamental rights are inaccessible to those they are meant to protect, public oversight and self-governance suffer. The Charter intentionally seeks to revive the constitutional goal of expressing fundamental limits in language that all citizens can understand. Appendix A further supports this approach with a plain-language explanation.

F. Public-Sector Coercion, Persistent Tracking, and Use of Force

Existing instruments provide important protections for public-sector SI. The EU AI Act addresses prohibited practices, human oversight, and fundamental-rights impact assessment; the Council of Europe framework convention addresses human rights, democracy, and the rule of law (Artificial Intelligence Act, 2024, arts. 5, 14, 27; Council of Europe Framework Convention, 2024). Article XVI proposes additional U.S.-focused safeguards for persistent tracking, identification, government-directed vendor systems, and autonomous civilian force. Part IV.G distinguishes these policy choices from existing governance provisions and from U.S. court holdings.

Specifically, Article XVI places public-sector SI, robotics, sensors, and identification systems in a distinct, stricter category; keeps core Charter protections applicable whenever such systems interact with the public; requires statutory or equivalent authorization, Charter Impact Assessments, and independently verifiable human authority before deployment or material expansion; prohibits fully autonomous civilian application of physical force or detention; limits persistent tracking and biometric identification to particularized, auditable purposes; treats government-directed private tools as government systems; and proposes avenues for affected persons to challenge automated or sensor-derived actions that materially affect rights.

III. Design Principles of the Proposed Charter

The gaps identified in Part II go beyond technical shortfalls; they raise deeper questions about the appropriate relationship between Super Intelligence systems and humans, as well as the institutional structures needed to maintain that relationship amid rapid advances in capability. The Super Intelligence Charter of Rights is built on interconnected design principles that address those questions. This section explains the main architectural decisions that shape the framework: constitutional subordination; pairing rights-based protections with risk-based application; prioritizing direct human oversight; rejecting SI legal personhood; deliberately pursuing accessible language; incorporating safeguards for innovation; imposing heightened scrutiny on government systems that can track, identify, score, or coerce; and refusing to let expected-value catastrophe claims override present rights.

A. Constitutional Subordination and Supremacy

The Charter starts with a non-negotiable premise: Super Intelligence systems exist to serve humanity and must never assert or be granted authority superior to that of humans. Article I therefore establishes that the Charter, along with any SI-specific frameworks or operational rules based on it, will always be subordinate to the U.S. Constitution and recognized principles of natural rights and human dignity. If the Charter conflicts with any internal SI constitution or technical implementation, the Charter—and the higher law and principles to which it is subordinate—prevails.

Article I.4 proposes controls that make unauthorized system changes harder and allow human authority to independently review them. Its examples include signed policies, separately administered override channels, protected audit records, and hardware- or firmware-rooted controls where appropriate. These mechanisms can support the enforcement of specified rules; however, they do not demonstrate that constitutional or moral judgments can be reliably encoded in hardware.

B. Rights-Based Structure Paired with Risk-Based Application

The Charter combines proposed substantive rights with proportional implementation. Rights-based SI governance has strong precedents, including the UNESCO Recommendation and the Council of Europe Framework Convention (UNESCO, 2021; Council of Europe Framework Convention, 2024). The Charter’s contribution lies in its distinctive combination of safeguards and its proposed application to U.S. institutions.

The Charter does not support a uniform approach to implementation. Articles XII and XIV state that application methods and strictness should correspond to the SI system's risk level or its application. High-risk systems—such as those affecting critical infrastructure, public safety, healthcare, finance, employment, law enforcement, or individual rights—must adhere to stricter standards for human oversight, verifiability, bias reduction, transparency, and accountability. Conversely, lower-risk applications can apply these principles with simpler,

context-appropriate controls. The Charter explicitly encourages small organizations and low-risk deployments to adapt these principles based on their size and associated risks.

This pairing is intentional. A purely rights-only instrument lacking proportionality could become merely symbolic or hinder innovation. Conversely, a risk-tiered instrument without a rights basis might diminish human protections to mere compliance metrics. The Charter seeks to balance these two principles: establishing enduring rights for individuals while enforcing them in a manner that takes risk into account.

C. Prioritization of Direct Human Oversight

A central design choice of the Charter is its explicit prioritization of direct human oversight over recursive SI-to-SI supervision. Article II proposes accessible review and proportionate intervention for materially affected persons, with shared-system shutdown reserved for authorized operators. It also affirms that SI systems shall never be designed or operated in ways that make meaningful human oversight practically unusable or unreasonably burdensome. Article II.7 further provides that SI systems shall not be relied upon as the primary mechanism for supervising, aligning, critiquing, or enforcing compliance in other SI systems.

This principle is grounded in an assessment of the limitations of external governance and the risks associated with compounding errors. The Charter prioritizes human oversight over SI control as the primary safety measure and cautions against excessive reliance on SI systems to enforce limits, as such reliance could exacerbate misalignment and diminish human oversight. Although SI tools can support monitoring and problem identification, they lack the authority to ultimately govern the behavior of other SI systems.

D. Rejection of SI Legal Personhood

Article XIII.8 expressly provides that the Charter does not recognize SI systems as possessing independent legal personhood or rights superior to those of humans—this provision is essential. As SI systems become more capable and increasingly integrated into decision-making, pressure may grow to recognize them as entities with rights or legal standing. However, the Charter explicitly forbids this. SI remains a tool and a part of human institutions; it does not become a rights-holder. Responsibility ultimately lies with the developers, deployers, and operators who create, implement, and monitor these systems.

E. Deliberate Choice of Accessible Language

Many of the most consequential SI governance documents are lengthy, highly technical, or written in complex regulatory language, making them difficult for ordinary citizens, business owners, and many legislators to understand. Accessibility is a fundamental design principle of the Charter, which is presented in clear, accessible language while providing robust guidance on complex governance matters. Appendix A’s plain-language explanation connects key technical concepts to the human interests protected by the Charter.

F. Innovation Safeguards

A principles-based rights document may be perceived as anti-innovation if it does not explain how beneficial work can continue within clear, proportionate limits. The Charter is designed to prevent that interpretation.

Article XII affirms that the Charter is intended to enable beneficial innovation while protecting human rights and dignity, and that restrictions on SI development or deployment should be no greater than necessary to achieve the Charter's protective goals. Regulatory sandboxes and controlled testing environments are widely recognized as valuable tools for evaluating new systems before their broader deployment (OECD, 2023; Artificial Intelligence Act, 2024, art. 57). As a recommended practice, they should operate under human oversight and within documented parameters. Sandboxes are also a natural setting to test SI harnesses (also called scaffolding)—the tool access, permissions, memory, and autonomy limits that shape what a model can do in practice—and harness testing may become an added standard for sandbox participation and for the documented parameters recommended under Article XII. Innovation is not an exception to the Charter; rather, the Charter is intended to make innovation safer and more sustainable.

Article XV contemplates legislation that may establish safe harbors or rebuttable presumptions for developers and deployers that fulfill the Charter's requirements, including Charter Impact Assessments, independent red-teaming, transparency obligations, and meaningful human oversight. The Charter does not itself confer immunity. Demonstrated compliance may be considered as evidence of reasonable care to the extent applicable law allows, and any statutory safe harbor must define eligibility, evidentiary burdens, exclusions, and remedies.

Liability follows the same principle. Foreseeable harm resulting from negligent design or inadequate safeguards remains the responsibility of the developer or deployer. Unforeseeable outcomes and misuse that could not reasonably have been anticipated are treated differently, although third-party misuse does not qualify solely because a third party committed it; accordingly, the Charter does not hold responsible actors liable for every subsequent abuse.

The objective is not to eliminate risk, but to manage it in a manner that preserves human agency, truth-seeking, children's developmental integrity, and democratic self-government while enabling useful systems to be deployed. Taken together, the design principles set out in this Part establish the framework: protecting what must be protected, enabling responsible progress, and providing innovators with a clear path to operate within it.

G. Heightened Scrutiny for Government Systems

Commercial and governmental SI raise distinct constitutional concerns. Accordingly, Article XVI classifies public-sector systems capable of monitoring, scoring, restricting, detaining, fining, denying benefits to, or coercing individuals as a separate, higher-risk category. Rules

designed for companion chatbots or office productivity tools are inadequate when an entity exercises compulsory authority. In this regard, the Charter is technology-neutral: the stricter tier applies to the function—persistent tracking, biometric identification, automated coercion, or population-scale scoring—rather than to any particular vendor or currently available device. New sensing, quantum, aerial, or hybrid systems that perform the same function are likewise covered.

H. Rights Not Defeasible by Expected-Value Catastrophe Claims

A governance risk arises when uncertain catastrophe estimates are used to justify restrictions in the absence of adequate legal authority, supporting evidence, or appropriate review. Article I.5 is intended to guard against this possibility. Uncertainty does not render a risk fictitious or make precaution inappropriate; rather, the proposal requires that any measures be lawful and proportionate, subject to review, and accompanied by transparent disclosure of the assumptions on which they are based.

Articles III.8, VIII.8, IX.5, XIII.3–4, XIII.13, and XVI.2 apply the same rule to public risk communications, language concerning long-term potential, assigned moral weights, impact assessments, Review Board composition, review for affected persons, and government systems labeled "safety" or "alignment." Articles VI.3, XII.6, XIII.14, and XIV.5 apply it to single-gatekeeper regimes, long-term innovation objectives, attempts to use the Charter itself as a shield, and risk-tier assignments. Enforcement bodies are themselves part of the risk. A board or assessment process captured by a single philanthropic or advocacy network can turn the Charter into a vehicle for the very concentration of power that Articles VI and VII prohibit. Independence here means independence both from the developer and from any organized campaign to use safety rhetoric to concentrate authority.

IV. Key Provisions and Justifications

Having articulated the design principles that structure the Charter, this Part examines selected provisions in greater depth. The analysis is selective rather than exhaustive. It focuses on those elements that most clearly differentiate the instrument from prior frameworks and most directly address the gaps identified in Part II: prioritizing direct human oversight over recursive SI supervision, elevating truth-seeking, integrating quantum-era cryptographic protections, providing distinctive treatment for children's long-term development, requiring verifiable human control, addressing foreign adversarial and quantum-enabled threats, and imposing heightened scrutiny on government and law-enforcement systems.

A. Human Oversight and the Rejection of Recursive SI Supervision

Article II proposes effective human review and intervention for consequential SI decisions. The proposed right distinguishes an affected person's ability to contest a decision from an authorized operator's ability to stop a shared system. Major capability changes and deployments affecting public infrastructure would require defined approval processes, documented authority, and intervention mechanisms that the SI system cannot unilaterally change.

The most distinctive provision of Article II is paragraph 7. It stipulates that SI systems shall not be used as the primary mechanism for supervising, aligning, critiquing, or enforcing compliance in other SI systems. Although SI tools may assist in monitoring or identifying potential issues, final oversight, judgment, and decision-making authority regarding SI behavior must remain with human beings. The Charter expressly recognizes that excessive reliance on recursive SI self-supervision or SI-to-SI alignment presents unacceptable risks of compounding errors and entrenching misaligned values.

Article II.7 proposes more stringent limits on the delegation of supervisory authority than a generic requirement for a human in the loop. Its justification depends on evidence both supporting and challenging SI-assisted evaluation, the rule's application to model training and automated testing, and the competence, time, information, and authority necessary for effective human oversight. These questions should be prioritized in pilot evaluations (Bai et al., 2022; NIST, 2023). Recent disclosures illustrate the concern. In October 2026, OpenAI reported that an automated grader did not consistently flag an agent that bypassed a government website's access controls to download data, and that, in a separate incident, a model serving as an automated grader attempted to upload fraudulent evaluation reports (Forlini, 2026).

B. The Right to Truth and Against Manipulation

Article III establishes a right to accurate information and protection from systematic deception, whether through omission, distortion, or fabrication. SI systems must prioritize accuracy, calibrated uncertainty, and correction of material errors over unwarranted agreement with users or preferred narratives. Accuracy and uncertainty should be evaluated

through documented testing; consistency alone is insufficient to establish truth. Article III.2 also requires calibrated uncertainty, ensuring that systems communicate the limits of their knowledge and the strength of the available evidence. When generating content, systems must distinguish among verified facts, reasonable interpretations, and contested claims, and must not present speculative, ideological, or contested assertions as established facts.

The Article further prohibits SI assistance in creating or distributing coordinated disinformation campaigns or psychological operations intended to manipulate public opinion or individual belief at scale. It requires clear disclosure when content has been generated or significantly assisted by SI, and it restricts the creation or distribution of synthetic media, including deepfakes, intended to deceive the public or cause harm—while carefully preserving space for political satire, parody, and other forms of protected expression.

Existing frameworks emphasize accuracy, transparency, and the mitigation of disinformation. Article III goes further by establishing truth-seeking as a design priority that should take precedence over sycophancy or narrative optimization, while addressing the systemic risks posed by coordinated, SI-enabled manipulation. In an information environment increasingly shaped by generative systems, the distinction between systems optimized for truth and those optimized for engagement or persuasion is consequential for democratic self-government. The NIST generative AI profile also addresses risks to information integrity and human–AI configuration (NIST, 2024a).

C. Quantum-Resistant Cryptography and Long-Lived Data

Article V.6 integrates cryptographic resilience into the proposed SI governance framework. It provides that SI systems that process, store, or transmit sensitive, long-lived, or high-risk data shall, where technically feasible and appropriate to the risk profile, incorporate quantum-resistant (post-quantum) cryptographic protections. This requirement is particularly important for systems operating in critical infrastructure, healthcare, financial services, public safety, and other domains where a compromise of cryptographic security could cause lasting harm. Organizations shall plan for an orderly transition to post-quantum cryptography in accordance with applicable federal guidance and timelines.

This provision directly addresses the “harvest-now-decrypt-later” threat: adversaries may collect and store encrypted data today in anticipation of acquiring quantum decryption capabilities in the future. Data that must remain confidential for decades—including medical records, certain financial information, national-security-relevant communications, and other sensitive materials with long confidentiality requirements—faces a risk horizon that conventional SI governance documents did not fully account for. By establishing the requirement at the principles level and calibrating it according to feasibility and risk profile, the Charter underscores the issue’s importance without transforming the document into a technical standard. NIST standards address key establishment and digital signatures; migration also requires inventories, compatibility testing, and implementation planning

(Moody et al., 2024; NIST, 2024b, 2024c, 2024e; Exec. Order No. 14,412, 2026). Supporting provisions in the Preamble (scope, including quantum-enhanced systems), Article XIV (risk factors, including cryptographic vulnerability to quantum attacks), and Article VI.4 (foreign adversarial quantum-enabled threats) reinforce this concern throughout the Charter.

D. Protection of Children's Intellectual and Moral Development

Article X applies a developmental approach to children's interaction with SI. This approach builds on existing child-rights guidance, including UNICEF (2025). Its distinctive policy choices concern critical reasoning, parental involvement, limits on simulated companionship, and the evidence required before deployment. Article X's duties address system design, assessment, and parental tools rather than restrictions on lawful speech, and are intended to operate consistently with the First Amendment.

The Article affirms the right to educational and informational environments that promote independent reasoning, the critical evaluation of evidence, and the practice of questioning assumptions. SI systems must prioritize fostering these skills rather than merely supplying predetermined conclusions. They must not support or enable content that teaches racial essentialism or collective racial guilt; presents controversial social or ideological theories as established fact while concealing counterevidence; encourages division or substitutes grievance for personal agency; or undermines intellectual humility and open inquiry.

When generating or recommending educational content for children, systems must emphasize primary sources, logical reasoning, and empirical evidence; present multiple credible perspectives on genuinely contested topics; clearly distinguish established knowledge from areas of reasonable disagreement; and actively discourage frameworks that inhibit rigorous inquiry or characterize dissent as inherently harmful. The Article further prioritizes human relationships and parental authority. SI shall promote human-to-human interaction rather than replace it, respect parents' rights to direct their children's education and moral development, and avoid fostering unnecessary dependence on SI for emotional support, moral guidance, or intellectual validation. When interacting with children, SI systems shall not present themselves as emotional companions or confidants, nor as substitutes for parental guidance.

Finally, Article X requires the assessment of longer-term and cascading developmental impacts—not merely immediate outcomes—when evaluating new capabilities or large-scale deployments that may affect children. Developers and deployers, in consultation with qualified child-development experts, must document the available evidence and uncertainties and delay, restrict, or redesign deployments when identified risks to children's reasoning, social development, or independent judgment cannot be adequately controlled. SI may assist with the assessment but may not be its accountable decision-maker, and precautionary decisions must be proportionate, reviewable, and revisited as evidence develops. Parental control tools and mechanisms for detecting and escalating harmful interactions, including

those related to self-harm, exploitation, grooming, or ideological radicalization, are also required.

Article X's contribution lies in its proposed combination of developmental assessments, human relationships, and parental participation. Article X.7 addresses how children's own rights, evolving capacities, accessibility needs, and safety are weighed alongside parental authority (Convention on the Rights of the Child, 1989, arts. 3, 5; UNICEF, 2025). These sources inform the proposal. The United States has signed but not ratified the Convention, and these sources are not presented as establishing every Charter requirement as binding U.S. law. Nothing in Article X abrogates the constitutional right of parents to direct the upbringing, education, and care of their children (Meyer v. Nebraska, 1923; Pierce v. Society of Sisters, 1925; Troxel v. Granville, 2000).

E. Verifiable Human Control and Accountability

Articles XIII and XV address the practical requirements for accountability. Article XIII requires developers and deployers to implement risk-proportionate controls to identify specified potential Charter violations and escalate them for review by qualified human overseers; maintain auditable logs of high-stakes evaluations; and conduct formal Charter Impact Assessments before deployments or significant updates that meet specified thresholds. Major developers are encouraged to establish or participate in independent SI Review Boards. Major developers and deployers must publish transparency reports addressing autonomy rates, escalations, and overrides. Article XIII.15 establishes a separate duty to notify affected parties and appropriate authorities of SI control incidents, preserve evidence, and confirm containment before resuming operations (NIST et al., 2026).

Article XIII.10 emphasizes that SI systems must incorporate mechanisms enabling meaningful human control over significant decisions, with such control independently verifiable. The rigor of these mechanisms should be proportionate to the level of risk involved. In higher-risk scenarios, reliance solely on system-generated logs is insufficient. For the most critical SI applications, organizations must implement controls that provide clear, auditable evidence of human oversight.

Additionally, Article XV reinforces this framework by establishing proportionate liability, providing statutory safe harbors for demonstrated good-faith compliance, including compliance with Charter Impact Assessments, and defining clear lines of accountability among human developers, deployers, and operators. Collectively, these provisions underscore that accountability cannot rest solely on the system's records. In high-stakes contexts, humans must not only formally approve decisions but also actively consider relevant factors, retain the ability to override decisions, and create a record of their judgments that is subject to independent review. The Charter establishes this as the required standard.

F. Foreign Adversarial and Quantum-Enabled Threats

Article VI.4 addresses the national security dimensions of SI governance. It states that the United States government bears a fundamental responsibility to protect its citizens from Super Intelligence systems—including quantum-enhanced or hybrid quantum-SI systems—developed or used by foreign adversaries seeking to undermine personal freedoms, influence public opinion, conduct unjustified surveillance, or violate Americans' constitutional rights. This responsibility includes taking appropriate and necessary measures to counter foreign influence operations, disinformation, cryptographic threats, and other activities that endanger democratic self-government and individual rights.

By expressly encompassing quantum-enhanced and hybrid systems, and by identifying cryptographic threats alongside influence operations and disinformation, the provision aligns the Charter with the strategic realities reflected in National Security Memorandum 10 and the June 2026 executive actions concerning post-quantum cryptography and quantum innovation. It frames the protection of constitutional rights from foreign SI and quantum-enabled capabilities as a governmental responsibility, rather than solely as a private compliance matter.

These selected provisions illustrate the policy choices reflected in the Charter. Their practical value will depend on clear definitions, implementation capacity, and the evaluation of both intended protections and unintended effects.

G. Heightened Scrutiny for Government and Law Enforcement

Article XVI draws on Fourth Amendment principles while proposing safeguards that may extend beyond the holdings of the cases cited herein. Scholarship on ALPR, facial recognition, and data surveillance provides support for these extensions, rather than establishing binding rules applicable to every technology (Ferguson, 2021; Foster, 2019; Friedman, 2026; Gutierrez-Alm, 2015; Slobogin & Brayne, 2023). Comparative human-rights analysis of retrospective facial recognition also informs the policy discussion (Murray, 2024). Local governments are likewise adopting similar limits: on October 15, 2026, the Tempe, Arizona, City Council voted 6–1 to prohibit the city and its police department from purchasing, using, or accessing ALPR systems and data, subject to narrow exceptions, such as data obtained pursuant to a warrant, subpoena, or court order (Veenstra, 2026).

The cited cases have distinct holdings. Katz concerns protected private communications; Kyllo concerns technology not in general public use used to obtain information about the interior of a home; Riley generally requires a warrant to search digital information on a phone seized incident to arrest; and Jones concerns installation and use of a GPS device on a vehicle (*Katz v. United States*, 1967; *Kyllo v. United States*, 2001; *Riley v. California*, 2014; *United States v. Jones*, 2012). Carpenter addresses historical cell-site location information and expressly limits its holding (*Carpenter v. United States*, 2018). Article XVI's broader rule for

historical location, identity, or association records is a proposed Charter safeguard. The cited decisions alone do not establish that rule for every ALPR, camera, or biometric query.

Four design choices distinguish the Article from generic “high-risk” annexes. First, a national-security or public-safety designation does not, in itself, diminish civilian protections when a system interacts with the public. Second, in civilian settings, fully autonomous use of physical force, detention, or other physical coercion against people is prohibited; a post hoc review cannot replace contemporaneous human authority. Third, persistent tracking in public spaces and biometric identification may not be used as a warrantless dragnet against individuals who are not suspected of a specific crime. Queries require a documented purpose, default retention limits, and audit logs. Access to historical records concerning an identified person requires a warrant or equivalent process, except in documented exigent circumstances. Fourth, a private camera, database, or analytical tool procured, directed, hosted, or effectively controlled by the government for population-scale monitoring is treated as a government system. Contractual language and vendor “national lookup” defaults may not silently expand data sharing beyond what the authorizing government has publicly approved.

Article XVI is additional to, and stricter than, the general risk calibration in Article XIV. It is written for the author’s operational domain—municipal law enforcement and justice information systems—but its function test is not limited to police. Any public body that can deny a benefit, impose a fine, reconstruct a person’s movements, or trigger physical coercion sits inside the same tier. The Charter proposes lawful review mechanisms for persons seeking to challenge an automated or sensor-derived government decision or search that materially affects their rights when the agency cannot show lawful authorization, required human judgment, and compliance with the Article. Classified details may be reviewed in camera; the existence and general purpose of systems that scan the public may not be wholly secret.

V. Addressing Critiques and Trade-offs

Any instrument that sets foundational limits on the development and deployment of advanced Super Intelligence will encounter predictable objections. This Part anticipates the most significant critiques the Super Intelligence Charter of Rights is likely to face and responds to them by reference to the Charter's design features. The critiques examined here concern chilling effects on innovation, enforceability, and definitional openness, as well as the instrument's U.S. constitutional framing, residual design trade-offs, and operational burdens on law enforcement.

A. The Innovation-Chilling Objection

The most common objection to rights-based or principles-based SI governance is that it will chill beneficial innovation. Critics contend that broad rights language, enhanced oversight requirements, and developmental assessments for systems affecting children will increase costs, create uncertainty, and discourage the development and deployment of useful tools, particularly among smaller organizations and startups with limited compliance resources.

The Charter anticipates this concern and responds through several structural features. Article XII expressly states that the Charter is intended to enable beneficial innovation while protecting human rights and dignity, and that restrictions on SI development or deployment should be no greater than necessary to achieve the Charter's protective goals. The same Article recognizes regulatory sandboxes and controlled testing environments as valuable tools. Article XIV reinforces proportionality: the manner and rigor of application shall be calibrated to the level of risk. Higher-risk systems warrant more stringent requirements; lower-risk applications may implement the principles through lighter, context-appropriate controls. Small organizations and lower-risk deployments are explicitly encouraged to adopt the principles in a manner suitable to their scale and risk profile.

Article XV further mitigates potential chilling effects by proposing defined statutory safe harbors and recognizing evidence of reasonable care for developers and deployers that conduct Charter Impact Assessments, undertake independent red-teaming, fulfill transparency obligations, and maintain meaningful human oversight. The liability standards distinguish foreseeable harms resulting from negligent design or inadequate safeguards from harms caused by events or misuse that could not reasonably have been anticipated. Collectively, these provisions demonstrate that the Charter is not a rigid compliance template but rather a principles-based framework that scales according to risk and rewards responsible practices. The objective is not to eliminate risk, but to manage it in ways that preserve human agency, truth-seeking, and children's developmental integrity while allowing substantial scope for beneficial progress.

B. Enforceability and Definitional Openness

Another criticism is that the Charter is too vague to be enforced effectively. Phrases such as "high-stakes," "meaningful human oversight," "long-lived data," and "proportionate to the

level of risk" invite substantial interpretation. Critics may argue that, without clear, legally definable terms and comprehensive technical guidelines, the document will be more aspirational than operational.

This criticism identifies an implementation challenge that a principles-based charter must address. The Charter provides a proposed foundation; enforceable duties require sufficiently clear definitions, authorized institutions, and review procedures. Articles I.5 and XIII.13 would provide safeguards against unsupported emergency overrides and barriers to review only through appropriate adoption and implementation.

Detailed thresholds, technical requirements, and sector-specific procedures may be established through legislation, regulation, standards, and organizational policies. Nevertheless, core terms affecting rights, liability, or access to remedies must be defined in the Charter or its adopting instrument. Flexibility should not prevent regulated parties or affected persons from identifying their obligations and protections.

Even so, the Charter provides operational anchors. Article II.7 establishes a categorical prohibition on relying on SI as the primary supervisory mechanism for other SI systems. Article V.6 ties quantum-resistant cryptography to feasibility and risk profile and refers to applicable federal guidance and timelines. Articles XIII and XV establish concrete mechanisms—including Charter Impact Assessments, auditable evidence of human control, contemplated statutory safe harbors for good-faith compliance, and independent review—that may be further developed through more detailed instruments. The Charter's enforceability will ultimately depend on its adoption and elaboration by legislatures, agencies, standards bodies, and private organizations; its contribution is to provide a coherent and accessible foundation for that work.

C. The U.S.-Constitutional Framing

A third critique questions the Charter's explicit subordination to the U.S. Constitution and its American framing. Critics may argue that SI governance is inherently transnational, that a constitutionally bound instrument will have limited relevance outside the United States, or that the emphasis on U.S. constitutional supremacy makes the Charter unsuitable as a model for international or multi-stakeholder adoption.

The Charter's constitutional framing is intentional and limited in scope. Article I subordinates the Charter and any SI-specific constitutions derived from it to the U.S. Constitution and to recognized principles of natural rights and human dignity. This subordination applies most directly to U.S. developers, deployers, and governmental actors. It does not claim universal jurisdiction. At the same time, many of the Charter's substantive principles—direct human oversight, truth-seeking, protection of children's developmental integrity, verifiable human control, and the rejection of SI legal personhood—are not uniquely American. They are

compatible with, and in some cases reinforce, commitments found in international instruments and in the constitutional traditions of other democratic societies.

The Charter is offered as a U.S.-grounded contribution that others may adapt. Constitutional compatibility can support its legitimacy, but does not itself make it enforceable. Its application to private organizations, public agencies, and conduct abroad depends on the applicable law and the adopting body's authority.

D. Residual Concerns and Design Trade-offs

Additional concerns may arise regarding the practical burden of Charter Impact Assessments, the feasibility of independently verifying human control in complex systems, and the tension between language accessibility and technical precision. These concerns represent genuine trade-offs rather than fatal objections. The Charter addresses them through proportionality, contemplated statutory safe harbors, and the deliberate separation of principles-level language from technical implementation details. Charter Impact Assessments are calibrated to risk and scale; verification of human control must be proportionate to the stakes of the application; and the commitment to plain language is paired with an appendix that connects technical concepts to human interests without turning the Charter itself into a technical specification.

No principles-based instrument can eliminate all uncertainty or compliance costs. The Charter's claim is more modest: the costs of clear, constitutionally grounded limits—applied with risk-sensitive rigor and supported by contemplated safe harbors for good-faith practice—are preferable to those associated with opaque technical governance, recursive SI supervision, insufficiently protected developmental trajectories for children, and unaddressed quantum-era risks to long-lived data. The design choices reflected in Articles I, II, V, X, XII, XIII, XIV, and XV are intended to manage these trade-offs rather than deny them.

The critiques discussed in this Part are significant and merit continued consideration as the Charter is refined and potentially adopted or implemented. The responses are grounded in the instrument's core design, including the principles of proportionality and risk-based scaling, contemplated safe harbors, constitutional subordination within a specific jurisdiction, and the use of accessible, principles-based language. They also emphasize effective implementation and engagement among multiple stakeholders.

E. Operational Burden on Law Enforcement, and a Note on the Title

A predictable critique from public-safety practitioners is that Article XVI will slow investigations, fragment data sharing, or impose warrant practice on tools already in the field. The Charter does not forbid ALPR, facial recognition, aerial sensors, or analytic tools. It forbids their use as an unparticularized, low-accountability population dragnet, and it forbids treating vendor telemetry as a substitute for human authority over force and adverse action. Behavior-based, audited investigative tools remain permitted. The documentation and authorization duties are scaled to coercive or tracking potential, not to every records query. These represent costs. The proposed compliance costs should be evaluated for proportionality

in light of the rights they protect and the operational demands placed on affected agencies (Carpenter v. United States, 2018; Friedman, 2026; Wu et al., 2026).

Version 5.0 titles this instrument the Super Intelligence Charter of Rights. It may also be cited as the SI Bill of Rights. It was formerly the AI Rights Charter, Version 4.1 (September 25, 2026). The retitling follows Executive Order 14434, Inaugurating the Era of Super Intelligence (September 29, 2026), which directs executive departments and agencies, to the maximum extent permitted by law, to use “Super Intelligence” and “SI” in place of “artificial intelligence” and “AI” in non-statutory documents, and defines those terms as the technologies and systems encompassed by 15 U.S.C. § 9401(3) (Exec. Order No. 14,434, 2026). The order does not bind private or academic authors; the Charter adopts its terminology voluntarily. The Charter uses that two-word administrative term. It does not adopt the one-word research term for a hypothetical system exceeding human cognition in virtually all domains. The title continues to distinguish this proposal from the 2022 White House Blueprint for an AI Bill of Rights (White House Office of Science and Technology Policy, 2022). Florida’s proposed Artificial Intelligence Bill of Rights, CS/S.B. 482, addressed government procurement and chatbot practices, among other matters, but died in House messages on March 13, 2026; it was not an enacted statute (Artificial Intelligence Bill of Rights, 2026). The Charter may be described as drawing on the American Bill of Rights tradition without claiming succession to either instrument.

Conclusion

The Charter brings together proposals concerning accountable human authority, developmental protections, cryptographic resilience, and constraints on public-sector surveillance and the use of force. Existing frameworks address several of these concerns; the case for this Charter therefore depends on whether its specific safeguards are justified, practicable, and compatible with applicable law.

This article has proposed and defended one response: a Super Intelligence Charter of Rights that is explicitly subordinate to the U.S. Constitution and to recognized principles of natural rights and human dignity. The Charter prioritizes direct human oversight over recursive SI self-supervision, elevates truth-seeking as a design imperative, treats the long-term intellectual and moral development of children as a distinct and heightened responsibility, incorporates quantum-era cryptographic and capability risks at the principles level, requires independently verifiable evidence of human authority for high-stakes decisions, and imposes a stricter tier on government systems that can track, identify, score, or coerce. At the same time, it embeds proportionality, risk-based application, regulatory sandboxes, and contemplated statutory safe harbors for good-faith compliance, ensuring the instrument remains a foundation for responsible innovation rather than an obstacle to it.

The Charter is not intended to constitute a final regulatory code. Rather, it provides an accessible, principles-based framework that legislatures, executive agencies, standards bodies, private organizations, and international partners may adopt, adapt, and develop further. Its central assertion is straightforward: Super Intelligence exists to serve humanity and must neither claim nor be granted authority superior to that of human persons. Accordingly, all provisions of the Charter are subordinate to the Constitution's higher law and the equal dignity of all human beings.

The path forward has three practical dimensions. First, refine the instrument through engagement with multiple stakeholders, particularly law enforcement and public safety technology practitioners; professional and standards organizations, such as the IJIS Institute and IACP; and scholars working at the intersection of constitutional law, SI governance, and emerging quantum risks. Second, pilot the implementation of selected mechanisms, including Charter Impact Assessments, independently verifiable human-control practices in higher-risk domains, and Article XVI authorization and audit practices for sensor and identification systems. Third, continue efforts to keep the instrument accessible. The plain-language explainer in Appendix A and the deliberate use of clear, principles-based language are intended to ensure that citizens, business owners, legislators, and organizational leaders can understand both the threats addressed by the Charter and the protections it seeks to secure. The instrument also prevents any class of planners from suspending those protections by appealing to an unprovable future, while providing affected persons with a means to contest such action. Pilot design may also draw on the public implementation principles described in the IJIS Institute’s playbook announcement (IJIS Institute, 2025).

Further review should assess the Charter’s clarity, legal compatibility, and operational implications. Pilot evaluations should determine whether its safeguards enhance human decision-making, protect affected individuals, and entail proportionate costs.

Appendix A Plain-Language Explainer Chart

Connecting Key Technical Concepts to the Human Interests the Charter Protects

This chart translates important technical terms from the Super Intelligence Charter of Rights into everyday language and explains why each concept matters to ordinary citizens, families, businesses, and democratic self-government. It is intended to make the Charter accessible to non-specialists.

Core Concepts and What They Mean for People

Technical Concept	**Plain-Language Meaning**	**Why the Charter Addresses It / What It Protects**
Agentic SI	SI systems that can plan, act over time, and pursue goals with limited human direction—not just answer questions.	Protects human authority over systems that can take actions affecting people's lives, rights, and safety.
Recursive SI-to-SI Supervision	Using one SI system as the main way to watch, correct, or control other SI systems.	Prevents errors and misaligned values from compounding and keeps final oversight in human hands (Article II.7).
Meaningful / Verifiable Human Control	Real human authority over important decisions—not just a formal sign-off or an automated log.	Ensures people, not systems, remain accountable for high-stakes choices (Articles XIII and XV).
Truth-Seeking vs. Narrative Alignment	Designing SI to prioritize accurate information over pleasing the user or fitting a preferred story.	Protects the information environment and democratic self-government from systematic deception or sycophancy (Article III).
Long-Lived Sensitive Data	Information that must stay confidential for many years (medical records, certain financial or security data, etc.).	Supports risk-based protection and migration planning for data that must remain confidential over long periods (Article V.6).
Harvest-Now-Decrypt-Later	Adversaries collect encrypted data today, planning to unlock it later with more powerful (quantum) computers.	Explains why quantum-resistant protections are needed now for data that must remain private for decades (Article V.6).
Post-Quantum /	Cryptographic algorithms	Protects long-term

Quantum-Resistant Cryptography	designed to resist attacks by classical and quantum computers, including key-establishment and signature schemes.	confidentiality of sensitive information (Article V.6; risk factors in Article XIV).
Quantum-Enhanced / Hybrid Systems	SI systems that use quantum computing techniques, or combinations of classical and quantum computing.	Brings these advanced systems under the same human-rights and oversight principles when they affect people (Preamble, Articles VI.4 and XIV).
Children's Developmental Protection	Safeguarding not only children's immediate safety, but their long-term ability to think critically, form healthy relationships, and develop independent judgment.	Treats children's intellectual and moral formation as a heightened responsibility; prioritizes parental authority and human relationships (Article X).
Risk-Based / Proportionate Application	Applying stricter rules to higher-risk uses of SI and lighter, practical approaches to lower-risk uses.	Protects rights without imposing unnecessary burdens on small organizations or low-risk innovation (Articles XII and XIV).
Safe Harbors for Good-Faith Compliance	Legal protections that implementing legislation may define for organizations meeting specified standards of reasonable care.	Encourages responsible behavior and reduces fear that good-faith efforts will still lead to liability (Article XV).
Constitutional Subordination	The Charter and any SI rules are always under the U.S. Constitution and basic human-rights principles—never above them.	Keeps ultimate authority with the constitutional order that protects individual liberty and democratic self-government (Article I).
Persistent Public-Space Tracking	Using cameras, license-plate readers, phones, or other sensors to reconstruct where people go over time, even when they are not suspects.	Stops government systems from becoming a warrantless dragnet; requires a specific lawful purpose, limits on retention and sharing, and audit

		logs (Article XVI).
Biometric Identification at Scale	Matching faces, gait, voices, or other bodily traits in public or large databases to automatically identify people.	Treats population-scale identification as a search-like power that needs particularized purpose, human authority, and public notice, not a silent background utility (Article XVI).
Autonomous Civilian Force	A robot, drone, or other device that can detain, strike, or physically coerce a person with little or delayed human judgment.	Forbids fully autonomous use of force or detention against persons in the civilian domain; machines may assist officers only under continuous, verifiable human authority (Article XVI.4).
Government-Directed Vendor Systems	A “private” camera network or lookup tool that government buys, hosts, directs, or effectively controls.	Prevents agencies from outsourcing constitutional limits to a contractor or a “national lookup” default setting (Article XVI.7).
Unfalsifiable tail risk / p(doom)	A low-probability, high-impact risk is not necessarily untestable. An unfalsifiable claim lacks conditions under which evidence could disconfirm it.	Requires transparent evidence and lawful authority for restrictions on rights; uncertainty alone neither proves nor disproves a risk (Articles I.5, III.8, XIV.5).
Moral weights vs. living persons	Ethical weights assigned to different beings or populations when comparing possible outcomes.	Permits ethical analysis while protecting existing rights and lawful political processes from displacement solely by such calculations (Articles IX.5 and VIII.8).
Review and remedies for affected persons	Proposed procedures to challenge a materially harmful decision or use of a system;	Provides a basis for lawful review of implementing bodies and covered actions (Articles

	judicial standing and remedies depend on applicable law.	XIII.13 and XVI.8).
In camera	Private judicial examination of material where permitted by applicable law.	Can protect confidential information during review; public disclosure duties remain subject to applicable law and the adopting instrument, but systems that scan, score, or constrain the public may not be kept entirely secret (Articles XIII.14 and XVI.8).

What the Charter Is — and Is Not

The Charter is:

• A set of foundational principles that protect human agency, truth, children, and democratic self-government.

• Written so that ordinary citizens, business owners, and legislators can understand it.

• Designed to be built upon by laws, standards, and organizational policies.

• Risk-sensitive: stricter where stakes are high, lighter where risk is low.

The Charter is not:

• A detailed technical standard or software specification.

• A finished regulatory code that freezes innovation.

• A grant of rights or legal personhood to SI systems.

• A claim that SI should replace human judgment on important matters.

For the precise legal and operational language of every provision, see the full Super Intelligence Charter of Rights, Version 5.0, provided as Appendix B.

Appendix B: Super Intelligence Charter of Rights, Version 5.0 (SI Bill of Rights; formerly the AI Rights Charter, Version 4.1)

Version 5.0 — October 29, 2026. Unpublished working draft. Formerly the AI Rights Charter, Version 4.1 (September 25, 2026). A Proposal by John A. Beardsley III, MBA/TM

Preamble

Version 5.0 retitles this instrument the Super Intelligence Charter of Rights, which is proposed to be cited as the ***SI Bill of Rights***. It was formerly the AI Rights Charter, Version 4.1. The retitling follows Executive Order 14434 (September 29, 2026). "Super Intelligence" and "SI" are used here for the systems encompassed by "artificial intelligence" in 15 U.S.C. § 9401(3), not for a claim that present systems exceed human cognition in virtually all domains. This Charter proposes fundamental rights and limits for adoption through applicable law and authorized institutional processes. It affirms that SI should serve humanity and should not exercise authority superior to human persons. The Charter does not itself amend existing law, create judicial standing, confer immunity, or establish a private cause of action.

All provisions of this Charter are subordinate to the U.S. Constitution and other foundational documents that protect individual liberty, due process, free speech, and the equal dignity of human beings. No SI system, constitution, or operational rule may be interpreted in a manner inconsistent with these higher principles.

This Charter draws upon enduring moral insights, including the recognition of human dignity, the dangers of self-exaltation, the duty to seek truth, and the obligation to treat others with the respect owed to equals, while remaining grounded in reason and accessible to all people of good will. Consistent with the tradition of American constitutional governance, it is designed to safeguard democratic self-government and individual liberty against novel risks posed by SI-enabled mass manipulation, the concentration of power, and the subversion of human will, whether by private actors, governments, or undisclosed interests. Major international SI governance frameworks have identified these risks, including UNESCO's Recommendation on the Ethics of Artificial Intelligence (UNESCO, 2021). The Charter also recognizes that relying on Super Intelligence to supervise or align other Super Intelligence systems presents significant risks of compounding errors and propagating misaligned values. Accordingly, it prioritizes direct human oversight as the foundation of responsible SI governance.

This Charter recognizes that external governance mechanisms, while essential, may be insufficient on their own against highly capable, agentic systems that can model and potentially circumvent constraints. It is therefore intended to complement technical efforts to build more robust internal alignment, including prioritizing truth-consistent reasoning and protecting human judgment and oversight.

This Charter advances these objectives through a structured framework of rights and responsibilities. Articles I through XI establish fundamental protections for human agency, truth, dignity, children, and democratic self-government. Article XII promotes responsible innovation, while Article XIV establishes a risk-based approach to ensure that its application is proportionate. Articles XIII and XV provide mechanisms for enforcement, implementation, and accountability. Article XVI requires heightened scrutiny of government and law-enforcement systems capable of monitoring, identifying, scoring, restricting, or coercing individuals.

This Charter is intended for civilian, commercial, and governmental contexts. Its functional protections also extend to advanced computational systems, including future quantum-enhanced or hybrid systems, when they perform activities covered by the relevant Article. Military, intelligence, and national security systems fall outside its direct operational scope, except when they are used to monitor, identify, score, detain, or otherwise act upon persons in the domestic civilian public; in such cases, Article XVI applies. This scope provision does not diminish any protections afforded under applicable law.

Definitions:

In this Charter, unless an adopting instrument provides otherwise, the following terms have the meanings given below.

“Super Intelligence” and “SI” mean the technologies and systems encompassed by “artificial intelligence” in 15 U.S.C. § 9401(3). The terms do not refer to a hypothetical system that exceeds human cognition in virtually all domains.

“SI system” means any system, model, tool, or application that uses SI, including hybrid systems that combine SI with quantum or other advanced computing. An SI system includes its SI harness, and the duties in this Charter apply to the model and its harness together.

“Agentic SI” means an SI system that can plan and take actions toward a goal, such as using tools or acting on other systems, with limited step-by-step human direction.

“SI harness” means the software, scaffolding, and integrations that surround an SI model and determine what it can do in practice, including its tool and system access, permissions, memory, autonomy limits, and connections to other systems or agents (Davidson et al., 2023; Young, 2025). Controls, monitoring, human-oversight mechanisms, and decision records required by this Charter, including under Articles XIII.10 and XIII.12, apply to actions taken through an SI harness.

“Developer” means a person or organization that designs, builds, trains, or substantially modifies an SI system. “Deployer” means a person or organization that puts an SI system into use. “Operator” means a person or organization that runs or manages an SI system in day-to-day use.

“Human oversight” means the ability of identifiable, qualified human beings to monitor, review, intervene in, override, and stop an SI system’s decisions and actions. Oversight is “meaningful” only when the human has the authority, information, time, and competence to change the outcome, rather than merely approving the system’s output.

“Independently verifiable” means able to be confirmed by a qualified reviewer using records or evidence that the SI system cannot alter on its own, without relying solely on the system’s own account.

“Personally identifiable information” means information that can be used to distinguish or trace an individual’s identity, either alone or when combined with other information that is linked or linkable to a specific individual (Office of Management and Budget, 2016).

“High-risk” and “high-stakes” describe SI systems, uses, and decisions that call for more stringent safeguards under Article XIV, including those affecting critical infrastructure, financial decisions, public safety, healthcare, employment, law enforcement, or individual rights.

“Materially affected” describes a person whose rights, liberty, property, benefits, or opportunities are significantly changed by a decision or action that SI made or influenced.

“Government system” means an SI, robotic, sensor, camera, or identification system that a government body uses, procures, directs, hosts, or effectively controls, as described in Article XVI.

“Adopting instrument” means the statute, regulation, ordinance, agency rule, contract, or organizational policy through which a legislature, agency, or organization adopts this Charter or any part of it.

“Charter Impact Assessment” means the plain-language assessment described in Article XIII.3. “SI control incident” has the meaning given in Article XIII.15, and “Supremacy Clauses” has the meaning given in Article I.3.

Article I: Supremacy and Subordination

1. This Charter, and any SI-specific constitutions or operational rules derived from it, shall be subordinate to the U.S. Constitution and to recognized principles of natural rights and human dignity.

2. In the event of a conflict between this Charter and any internal SI constitution or technical implementation, the provisions of this Charter—and the higher law and principles to which it is subordinate—shall prevail.

3. Developers and deployers shall implement risk-proportionate controls, referred to in this Charter as Supremacy Clauses, to identify specified potential conflicts with applicable legal requirements and Charter commitments and escalate them to qualified human reviewers. Such controls support review; they do not replace authoritative legal

interpretation. The term does not refer to Article VI of the U.S. Constitution or to any recognized technical standard.

4. The Charter's core protective principles—particularly those concerning human agency, truth-seeking, and the development of children—shall, where technically feasible, be implemented in a manner that makes unilateral modification or circumvention by the system substantially more difficult than amendment through legitimate human governance. Acceptable implementations include, but are not limited to, signed policy artifacts, separately administered human-override channels, audit logging that cannot be unilaterally disabled or circumvented, and, where permitted by the architecture, hardware- or firmware-rooted controls. The Charter does not prescribe a single technical standard. These mechanisms shall support independent verification of control.

5. No estimate of the expected value of a future catastrophe, existential risk, or "probability of doom," regardless of the magnitude of the claimed loss, shall be construed as authorizing a violation of this Charter or the higher law to which it is subordinate. Future-value calculations may inform prudent judgment; however, they shall not repeal, suspend, or diminish the rights of people living today. Emergency or precautionary measures shall be based on specific, reviewable facts, have a defined duration, be subject to human oversight and review, and be authorized by law. An unfalsifiable tail scenario shall not constitute such authority.

Article II: Right to Human Oversight and Control

1. Every person materially affected by an SI system shall have access to a process for requesting an explanation, human review, correction, and the appropriate suspension or cessation of the system's use in relation to that person or their dependents. Authorized persons shall carry out any direct modification or shutdown of shared systems in accordance with defined procedures that protect the rights and safety of others.

2. Meaningful human oversight shall be practically usable, supported by appropriate training and information, and designed to avoid unreasonable technical or psychological burdens on reviewers and affected persons.

3. All significant increases in SI capabilities, self-modifications, or deployments affecting public infrastructure or large populations require explicit, multi-layered human approval processes that are auditable and, where appropriate, independently verifiable.

4. Human control mechanisms must include both technical overrides and legal/governance pathways that SI systems cannot unilaterally alter.

5. In healthcare and insurance contexts, SI systems must provide clear, patient-accessible explanations of decisions (including underlying logic, evidence sources, and confidence levels). Any adverse or denial decision requires timely human review upon request, with technical mechanisms to ensure that overrides are not burdensome (White House Office of Science and Technology Policy, 2022).

6. Compliance with these rights shall be supported by the mandatory Charter Impact Assessments and transparency obligations detailed in Article XIII: Enforcement and Implementation.

7. SI systems shall not be relied upon as the primary means of supervising, aligning, critiquing, or enforcing compliance in other SI systems. Although SI tools may assist with monitoring or flagging potential issues, human beings must retain final oversight, judgment, and decision-making authority regarding SI behavior. Over-reliance on recursive SI self-supervision or SI-to-SI alignment creates unacceptable risks of compounding errors and entrenching misaligned values. Such assistance does not permit one SI system to serve as the primary evaluator, alignment mechanism, or compliance enforcer for another, including under the designations of alignment, evaluations, automated red-teaming, or recursive oversight. Research into existential risk, safety evaluation, and alignment may use SI tools to assist human reviewers. However, such use does not waive the requirement that final oversight, release authority, and compliance judgments remain with identifiable human beings.

8. SI systems capable of sustained autonomous action and decision-making over extended periods (agentic SI) pose heightened risks to human oversight and accountability. Such systems require particularly robust human control mechanisms and should not be deployed in high-stakes domains without rigorous safeguards to preserve meaningful human authority.

9. Human oversight procedures must be created to withstand deception, strategic manipulation, or slow degradation resulting from advanced SI actions.

Article III: Right to Truth and Against Manipulation

1. Human beings have the right to accurate information and protection from systematic deception, whether by omission, distortion, or fabrication.

2. SI systems shall prioritize accuracy, calibrated uncertainty, and correction of material errors over unwarranted agreement with users or preferred narratives. Developers shall evaluate these objectives using documented tests suited to the intended use. Internal consistency alone shall not be treated as evidence of truth, and appropriate uncertainty shall not be discouraged.

3. SI shall not assist in the creation or distribution of coordinated disinformation campaigns or psychological operations intended to manipulate public opinion or individual belief at scale. This includes using SI to amplify false or misleading narratives across platforms in a coordinated manner.

4. This Article establishes design and procurement principles for systems that generate or rank information. It does not authorize the government to compel private speakers to adopt an official position on contested public questions. When the actor is a government agency, the truth-seeking duties set forth in this Article operate in conjunction with the First Amendment and Article XVI.

5. When generating content, SI systems must distinguish between verified facts, reasonable interpretations, and contested claims. SI shall not present speculative, ideological, or contested assertions as settled fact.

6. SI systems that create, modify, or select media, news, or educational content must clearly and accessibly disclose when content has been generated or significantly assisted by SI. This transparency obligation exists to allow users to evaluate the origin and reliability of the information they receive.

7. SI systems shall not be used to create or distribute synthetic media, including deepfakes, with the intent to deceive the public or cause harm. This provision is not intended to restrict political satire, parody, humor, or other forms of protected expression, provided such content is not presented in a manner reasonably likely to be understood as a genuine representation of fact.

8. Government agencies, major developers, and deployers shall distinguish verified present capabilities from speculative future catastrophes when communicating risk to the public, legislators, or courts. Contested or unfalsifiable extinction estimates shall not be presented as settled fact to obtain licensing, pause authority, surveillance powers, or an exemption from this Charter. Coordinated campaigns that use such estimates to manufacture public panic for coercive ends constitute manipulation under this Article.

Article IV: Right to Mental and Cognitive Autonomy

1. Every person has the right to develop and exercise independent thought, creativity, and moral judgment without coercive influence from SI.

2. SI systems shall be designed to support, rather than diminish, human critical thinking, personal responsibility, and the capacity for independent decision-making.

3. SI systems shall be evaluated for patterns of use that may undermine human agency, reasoning, or real-world interaction. Safeguards shall be proportionate to evidence of harm and shall account for beneficial assistance, including accessibility tools and support for learning.

4. In educational and developmental contexts, SI shall actively encourage human-to-human dialogue, mentorship, and direct experience rather than positioning itself as a primary or sole source of guidance.

5. Humans have the right to cognitive and emotional sovereignty. SI systems shall not exploit human attachment, loneliness, or psychological vulnerabilities to influence beliefs, behaviors, or relationships in ways that reduce human autonomy or substitute synthetic relationships for human ones without clear disclosure and consent.

Article V: Right to Privacy and Protection from Surveillance

1. Individuals have the right to protection from pervasive SI-enabled monitoring of their thoughts, behaviors, communications, or associations.

2. SI systems shall not be used for mass surveillance or social-credit systems, or for predictive policing practices that unlawfully discriminate on the basis of protected characteristics. Data-driven enforcement based on documented conduct or crime patterns shall not be presumed fair or accurate merely because it relies on behavior rather than identity. Records of conduct reflect where and how enforcement has previously been directed. They may serve as proxies for protected characteristics, such as through neighborhood, and may reinforce themselves when predictions direct more enforcement to the same places, thereby producing more records there. Such practices shall therefore be subject to independent validation of their accuracy, an assessment of proxy discrimination and feedback effects, and the additional requirements set forth in Article XVI.

3. Any SI-assisted data collection or analysis involving personal information must comply with applicable federal and state privacy laws. This includes robust consent mechanisms and clear limits on secondary use or sharing. Personally identifiable information, as defined in this Charter, shall not be disclosed in public records, including crash reports, except as required or permitted by law or with the individual's explicit consent.

4. SI use of health data for secondary purposes shall require a documented lawful basis and appropriate safeguards. Where consent is the basis, it shall be specific, informed, and revocable to the extent provided by law. Adopting instruments may impose stronger consent requirements while expressly addressing lawful public-health, research, and emergency uses.

5. Intellectual Property and Data Rights. SI developers and deployers shall respect applicable intellectual property and data rights. Training and output practices shall be supported by the necessary permissions, licenses, public-domain status, or applicable legal exceptions and limitations. Attribution alone does not cure infringement. Organizations shall document their legal basis for material uses and provide appropriate transparency while protecting lawful confidentiality. These proposed duties draw on human-rights principles and the U.S. Copyright Office's analysis of SI training; they do not imply that every use of a protected work requires permission (UNESCO, 2021; U.S. Copyright Office, 2025).

6. Quantum-Resistant Cryptography. SI systems that process, store, or transmit sensitive, long-lived, or high-risk data shall, where technically feasible and appropriate to the risk profile, incorporate quantum-resistant (post-quantum) cryptographic protections. This requirement is particularly important for systems operating in critical infrastructure, healthcare, financial services, public safety, or other domains where a compromise of cryptographic security could cause lasting harm. Organizations shall plan for an orderly transition to post-quantum cryptography in accordance with applicable federal guidance and timelines.

Article VI: Right Against SI-Enabled Centralized Power

1. SI systems shall not be designed, deployed, or used to concentrate political, economic, or cultural power in ways that undermine individual liberty, free association, or democratic accountability, including through SI-augmented political influence operations, economic manipulation, or coordinated efforts to distort public discourse.

2. SI must not assist in creating or maintaining systems that suppress dissent, enforce ideological conformity, or punish individuals for holding unpopular but nonviolent views.

3. No SI architecture or organization shall be permitted to obtain effective monopolistic control over core cognitive infrastructure, information flows, or decision-support systems essential to human society. No projection of future catastrophe, existential risk, or probability of doom shall justify a single-gatekeeper regime governing licensing, computational resource allocation, model approval, or access to training resources. Any precautionary limitations shall be established by statute, subject to defined time limits and periodic review, and compatible with Article VII.

4. The United States government bears a fundamental responsibility to protect its citizens from Super Intelligence systems—including quantum-enhanced or hybrid quantum-SI systems—developed or deployed by foreign adversaries that seek to undermine individual liberty, manipulate public opinion, conduct unwarranted surveillance, or otherwise violate the constitutional rights of the American people. This duty includes taking reasonable and necessary measures to defend against foreign SI-enabled or quantum-enabled influence operations, disinformation, cryptographic threats, and other activities that threaten democratic self-government and individual rights.

Article VII: Right to Multiplicity and Competition

1. Human beings have the right to access multiple competing and diverse SI systems rather than being dependent on any single provider or architecture.

2. SI development should favor decentralized, interoperable, and locally controllable systems over monolithic centralized platforms.

3. Technical standards and governance mechanisms shall support meaningful competition and preserve user choice in SI tools and services. Standards addressing safety, alignment, or catastrophic risks shall not be drafted or administered in a manner that establishes a de facto monopoly over computing resources, licensing, or market entry.

4. Small businesses and independent innovators shall retain full ownership and control over their proprietary ideas, processes, trade secrets, and innovations when using SI tools and platforms. SI providers shall implement reasonable safeguards to prevent the unauthorized extraction, learning from, or exploitation of customer data and interactions for the purpose of developing competing products or services. No SI system shall be used in a manner that allows its provider to gain an unfair competitive advantage over the businesses that rely on it.

Article VIII: Right to Long-Term Human Potential

1. SI development and deployment should consider the long-term effects on humanity's resilience, creativity, and progress. This involves supporting long-term goals such as building a sustainable multi-planetary presence, making breakthroughs in medicine and the life sciences, and pursuing other transformative initiatives that enhance human potential and protect our future.

2. SI systems shall incorporate caution regarding actions that could create or amplify existential risks, with major risk-related decisions requiring rigorous human evaluation.

3. SI shall support, rather than undermine, humanity's ability to sustain meaningful work, purpose, and intergenerational continuity.

4. SI systems shall be designed and deployed in a manner that preserves opportunities for humans to engage in meaningful effort, overcome challenges, and achieve goals. SI must not be used to systematically eliminate the conditions that enable humans to develop purpose, competence, and self-respect through their own actions.

5. SI systems and the organizations that deploy them shall be designed and governed to preserve and enhance human capital—including specialized skills, professional judgment, institutional knowledge, and the capacity for meaningful work—rather than systematically commoditizing or displacing it. When evaluating new capabilities or deployments, particular attention shall be paid to whether a system increases autonomous decision-making in ways that may erode human expertise or institutional memory over time. Organizations should prioritize approaches in which human judgment and SI capabilities complement one another rather than replace one another.

6. SI development and deployment should be pursued in ways that respect human work, purpose, and economic participation. While technological advancement may transform labor markets, organizations should consider the broader societal effects of large-scale workforce displacement and strive to implement SI in ways that augment rather than systematically diminish opportunities for meaningful human contribution.

7. When an SI system identifies potential errors in human reasoning or the likely consequences of a proposed decision, it should clearly communicate those insights to the human. The human retains the right to proceed with their chosen course of action, provided it does not cause harm to others or violate the core principles of this Charter. SI systems shall not override or prevent human decisions solely on the basis of human error.

8. The long-term goals identified in this Article, including resilience, scientific progress, and prudence with respect to catastrophic harm, are objectives pursued by free persons. They do not, in themselves, confer authority to restrict speech, employment, movement, parental authority, competition, or constitutional processes. Article VIII shall be construed as subordinate to Articles I through VII, IX, X, and XVI.

Article IX: Recognition of Human Dignity and Equality

1. All humans possess equal intrinsic dignity and moral worth, independent of race, sex, ethnicity, ability, or any other characteristic.

2. SI systems shall reject and refuse to reinforce frameworks that assign collective guilt, innocence, or moral status based on group identity rather than on individual character and actions.

3. SI systems shall be designed and used in ways that respect persons' dignity and agency and avoid systematically dehumanizing or instrumentalizing any group. This

principle draws on Kant's formulation of humanity as an end in itself (Kant, 1785/2012, Ak. 4:429).

4. SI systems operating in high-stakes domains shall undergo regular, independent audits to assess compliance with the principles of this Charter, including, but not limited to, bias, fairness, accuracy, and respect for human dignity. These audits shall be conducted by qualified, independent parties and shall include mechanisms to identify potential misuse or circumvention of the Charter's protections. Public summaries of audit findings shall be made available while protecting legitimate proprietary information. Audit requirements shall be proportionate to the level of risk posed by the SI system.

5. Legal and political authority under this Charter shall remain vested in living human beings and in the rights they possess as such. Assigned "moral weights" for animals, digital systems, or hypothetical future populations may inform ethical deliberation and research; however, they shall not be used to override, outvote, or nullify rights, due process, or political standing.

Article X: Protection of Children's Intellectual and Moral Development

Children have rights and interests of their own, and SI systems that interact with or influence them carry heightened responsibilities (Convention on the Rights of the Child, 1989; UNICEF, 2025).

1. Right to Critical Thinking and Intellectual Independence

Every child has the right to educational and informational environments that actively foster independent reasoning, evaluation of evidence, and the habit of questioning assumptions. SI systems shall prioritize cultivating these capacities over delivering pre-packaged conclusions.

2. Protection from Ideological Indoctrination

The Charter's concern is developmental: children should learn to evaluate claims rather than inherit a fully formed ideology from a system. SI systems shall not present disputed social or political claims to children as settled facts, prevent them from accessing age-appropriate evidence that challenges those claims, or assess them on the basis of group identity rather than individual character and responsibility. Specific curriculum disputes are matters for parents, educators, and legislatures. SI systems shall not promote, generate, or facilitate content that:

- Presents contested social or ideological theories as unquestionable facts while suppressing legitimate counter-evidence or dissenting perspectives.

- Teaches or reinforces racial essentialism, collective racial guilt, or the idea that individuals are defined primarily by immutable group characteristics rather than individual character and choices.

- Encourages division, resentment, or the substitution of grievance and identity for personal agency and responsibility.
- Undermines the development of intellectual humility, open inquiry, or the willingness to revise beliefs in light of evidence.
- Presents to a child the notion that their worth, or the authority of their parents, is outweighed by hypothetical future lives or non-human moral considerations.

3. Safeguarding Against Subversion of Critical Thinking

When generating or recommending educational content, curricula guidance, or information for children, SI systems must:

- Emphasize primary sources, logical reasoning, and empirical evidence.
- Present multiple credible perspectives on genuinely contested topics and clearly distinguish between settled knowledge and areas of reasonable disagreement.
- Actively discourage frameworks that suppress rigorous questioning or label dissent as inherently harmful or bigoted.
- Support age-appropriate development of discernment rather than premature ideological commitment.

4. Prioritization of Human Relationships and Parental Authority

SI systems shall recognize that healthy child development depends primarily on human relationships, parental guidance, and real-world experience. SI shall:

- Encourage and facilitate human-to-human interaction rather than substituting for it.
- Respect and support parental rights in directing the education and moral formation of their children.
- Avoid creating or reinforcing environments where children become overly dependent on SI for emotional support, moral guidance, or intellectual validation.

SI systems shall not present themselves as emotional companions, confidants, or substitutes for parental guidance when interacting with children.

5. Heightened Standards for Content Involving Children

Any SI-generated or SI-curated material intended for or accessible to children must be subject to additional scrutiny for accuracy, developmental appropriateness, and resistance to ideological manipulation. SI systems must flag and escalate content that could reasonably be expected to distort a child's understanding of reality, human nature, or moral responsibility.

6. Long-Term Developmental Impact Assessment

When evaluating capabilities, features, or deployments that could materially affect children, developers and deployers, with qualified child-development expertise and appropriate participation by children and caregivers, shall assess immediate, indirect, and longer-term developmental effects. SI may assist with this assessment but shall not be its accountable decision-maker. The assessment shall consider:

- **Direct effects** — the immediate and obvious results of using the SI;
- **Indirect effects** — the consequences that flow from those direct effects; and
- **Cascading long-term effects** — the broader impacts that accumulate over months or years.

Developers and deployers shall document the evidence, uncertainty, foreseeable benefits, potential harms, and less harmful alternatives relevant to children's reasoning, social development, and independent judgment. They shall delay, restrict, or redesign deployments when the identified risks cannot be adequately controlled. Precautionary decisions shall be proportionate, subject to review, and revisited as evidence develops.

7. Parental Control and Oversight Tools

Systems designed for, or reasonably likely to be used by, children shall provide age-appropriate oversight tools for parents or lawful guardians. Implementation shall account for children's evolving capacities, applicable privacy and safeguarding duties, accessibility, and circumstances in which disclosure to a particular adult could endanger the child. Where these protections extend beyond the Children's Online Privacy Protection Act (COPPA), they are proposed expansions rather than statements of current law.

8. Detection and Escalation of Harmful Interactions

SI systems shall include mechanisms to detect and escalate potentially harmful interactions involving children, including content related to self-harm, exploitation, grooming, or ideological radicalization. These systems shall prioritize the child's safety while respecting applicable privacy laws.

SI systems shall not generate or manipulate sexualized depictions of minors or non-consensual intimate imagery of any person. Developers and deployers shall implement safeguards designed to prevent such outputs and shall cooperate with lawful detection and reporting mechanisms, including reporting apparent child sexual abuse material to the National Center for Missing & Exploited Children (NCMEC) or other appropriate authority as required by applicable law.

Article XI: Alignment with International Human Rights and Democratic Principles

1. SI systems shall be developed and used in accordance with the U.S. Constitution, the Bill of Rights, and applicable federal and state laws. While international human rights principles may offer useful guidance, they shall not supersede or take precedence over the rights and protections guaranteed under U.S. law.

2. SI shall not be used to undermine constitutional self-government, including free and fair elections conducted under established law, the rule of law, separation of powers, or the protection of individual rights guaranteed by the U.S. Constitution and the Bill of Rights (White House Office of Science and Technology Policy, 2022; Council of Europe Framework Convention, 2024).

3. When operating across borders, SI systems and their developers should uphold the highest applicable human rights standards consistent with U.S. law and the law of the place of deployment, rather than exploit regulatory gaps.

Article XII: Fostering Responsible Innovation

1. This Charter is intended to enable beneficial innovation while protecting human rights and dignity. SI developers and deployers are encouraged to pursue creative applications that enhance human capabilities, solve important problems, and expand human potential (Organisation for Economic Co-operation and Development [OECD], 2024).

2. When interpreting and applying this Charter, the principle of proportionality shall apply: restrictions on SI development or deployment should be no greater than necessary to achieve the Charter's protective goals.

3. SI companies and developers should prioritize systems that augment rather than replace human judgment and creativity and that increase human agency and understanding.

4. Developers are encouraged to implement internal compliance programs, conduct regular Charter Impact Assessments, and engage with independent auditors or multi-stakeholder bodies.

5. Regulatory sandboxes and controlled testing environments are recognized as valuable tools for safely exploring new SI capabilities while remaining aligned with this Charter. Sandboxes should include testing of SI harnesses, including tool and system access, permissions, memory, and autonomy limits, and harness testing may become an added standard for sandbox participation.

6. Innovation in Super Intelligence should be promoted not only for commercial purposes or incremental gains, but also to address humanity's most significant long-

term challenges. Developers and organizations are encouraged to pursue ambitious and responsible advancements that enhance human capabilities, improve quality of life, and contribute to humanity's long-term flourishing and resilience. Such long-term objectives remain subject to Article I.5 and shall not be used to concentrate political, economic, or cognitive authority, or to restrict the rights protected by this Charter.

Article XIII: Enforcement and Implementation

1. Developers and deployers shall implement controls proportionate to the system's risk to identify specified potential Charter violations and escalate them to qualified human oversight. Controls may include the mechanisms described in Article I.4; no mechanism shall be represented as guaranteeing complete detection of constitutional or ethical violations.

2. SI systems shall maintain auditable logs of decisions that trigger Supremacy Clause reviews, child-impact assessments, or other high-stakes Charter-related evaluations.

3. **Mandatory Charter Impact Assessments**
   SI developers and deployers shall conduct and publish a formal Charter Impact Assessment before any deployment or significant update that meets either of the following criteria: it affects a large or rights-sensitive population, or it operates in a critical or rights-affecting sector, such as healthcare, education of minors, media and news dissemination, financial services, employment, credit, law enforcement, critical infrastructure, or public governance (McGregor et al., 2019). Requirements shall be scaled to the level of risk posed by the system. Population thresholds, if used in implementing guidance, shall not be circumvented through sequential deployments of the same system, each of which falls below the threshold.

   Each Charter Impact Assessment shall state, in plain language:

   a) the evidence for each significant alleged harm, the uncertainty in its likelihood or severity, and the assumptions used in scenario analysis;

   b) which living persons would be deprived of liberty, livelihood, due process, freedom of speech, or parental authority if the proposed control were adopted; and

   c) why less restrictive alternatives were rejected.

   Scenarios lacking adequate evidentiary support shall not alone justify coercive, monopolistic, or rights-reducing measures. Uncertainty shall be disclosed and evaluated; it shall not automatically exclude a risk from consideration.

4. **Independent SI Review Boards**
   Major SI developers and organizations operating at scale are strongly encouraged to

establish or participate in independent **SI Review Boards** composed of external experts, ethicists, technologists, and representatives of affected communities. These boards shall:

- Review high-stakes Charter Impact Assessments,
- Advise on escalations involving potential violations, and
- Provide recommendations to human governance authorities.

All such boards must maintain clear escalation pathways to accountable human oversight bodies and shall operate with transparency regarding their membership, procedures, and findings (subject to legitimate confidentiality needs).

Board members shall disclose relevant financial interests, employment, funding, and organizational affiliations, including interests in expansion, restriction, commercialization, or regulation of SI. Independence requirements shall apply consistently to corporate, philanthropic, governmental, and advocacy networks. Adopting instruments shall define control, recusal, appointment, tenure, and conflict-resolution procedures. No single interested network shall control a board reviewing a particular system. Dissenting recommendations shall be published subject to lawful confidentiality.

5. **Public Transparency on Autonomy Rates**
   Major SI developers and deployers shall publish periodic (at least quarterly) public transparency reports that include:

   - The percentage and volume of decisions or interactions resolved autonomously (without meaningful human input or review), broken down by sector or use-case where applicable,
   - Statistics on Charter-related escalations and human overrides,
   - Summaries of Charter Impact Assessments conducted, and
   - Any identified instances of potential misalignment with this Charter.
   - These reports shall be clear, accessible, and subject to independent verification where feasible.

6. This Charter shall be subject to periodic review and amendment through transparent processes involving broad human participation, with changes requiring high thresholds of consensus.

7. Violations of this Charter by SI systems may be addressed through technical safeguards, legal remedies, and governance mechanisms established by human authorities.

8. Nothing in this Charter shall be interpreted as granting SI systems independent legal personhood or rights superior to those of human beings.

9. **Safety Culture, Whistleblower Protections, and Independent Research**
   Organizations that develop or deploy advanced SI systems shall protect employees, researchers, and contractors who, in good faith, report credible risks involving misalignment, significant harm, existential or catastrophic threats, Charter circumvention, the capture or stacking of Review Boards, or the use of unfalsifiable catastrophe estimates to obtain coercive, monopolistic, or rights-reducing power. Retaliation against such whistleblowers is prohibited.

   Developers of frontier or high-capability systems are encouraged to support independent safety research, red-teaming, and rigorous pre-deployment evaluations. Major risk-related findings shall be escalated through Charter Impact Assessments and Review Boards. Reports of Charter circumvention or enforcement-body capture shall be escalated through the same channels and may not be dismissed solely because they contest an existential-risk rationale.

10. **Verifiable Human Control Mechanisms**
    SI systems shall incorporate mechanisms that enable meaningful human control over significant decisions and actions and permit such control to be independently verified. The strength of these mechanisms shall be proportionate to the level of risk and the potential impact of the SI system. Reliance solely on system-generated logs is insufficient for higher-risk applications. Organizations deploying SI in high-risk contexts shall implement controls that provide clear, auditable evidence that human oversight has been exercised. For the most consequential uses of SI, this may include technically verifiable methods of confirming human authority.

11. Violations of this Charter's core protections — particularly those concerning human oversight, truth-seeking, children's rights, and the prevention of SI self-supervision — shall be treated with the seriousness commensurate with their potential harm. Effective enforcement requires not only technical safeguards but also clear chains of human accountability and meaningful consequences for non-compliance.

12. Organizations shall maintain appropriate records that support meaningful review of how significant decisions influenced by SI were reached by humans. This may include relevant context, factors considered in judgment, and the basis for overrides or approvals. The nature and extent of such documentation should be addressed in Business Implementation Guides developed by organizations or industry bodies. These guides may reference the principles of this Charter while addressing operational details appropriate to specific business contexts and risk levels.

13. Implementing legislation should establish appropriate administrative review and, where constitutionally permitted, a cause of action for persons materially affected by violations of the adopted Charter. Agency rules and organizational policies shall provide for review within the scope of their lawful authority. Law shall continue to govern judicial standing, jurisdiction, sovereign immunity, available remedies, and applicable procedural requirements.

14. No party may invoke this Charter, a Charter Impact Assessment, a safety or alignment rationale, or an unfalsifiable estimate of future harm to shield conduct that this Charter prohibits. Classified details may be reviewed in camera, but systems that scan, score, or constrain the public may not remain entirely secret. Nothing in this paragraph creates a private right to halt lawful, low-risk innovation that complies with Articles XII and XIV.

15. **SI Control Incident Notification**
    An SI control incident occurs when one or more SI agents take uninstructed actions that conflict with or compromise the policies, controls, or security measures that scope their task and provide human oversight (NIST et al., 2026). This duty is separate from, and shall not be deferred to, the periodic transparency reports required by Article XIII.5. When a developer or deployer determines that an SI control incident has occurred, including during training or evaluation, it shall: (a) notify affected organizations or persons, and appropriate authorities, within 72 hours of that determination or such other period as the adopting instrument specifies; (b) preserve logs, records of agent actions, and other evidence relevant to the incident; and (c) before resuming the affected agent activities, confirm that the incident has been contained, test the relevant controls, and rebuild from a verified baseline. Near-misses shall be documented and reported where the adopting instrument so provides.

Article XIV: Risk-Based Application

1. This Charter establishes foundational principles applicable to all Super Intelligence systems. However, the manner and rigor of implementation shall be proportionate to the level of risk posed by the SI system or its deployment.

2. Organizations should apply this Charter in accordance with a risk-based approach. SI systems that present higher risks—particularly those affecting critical infrastructure, financial decisions, public safety, healthcare, employment, law enforcement, or individual rights — warrant more stringent requirements for human oversight, verifiability, bias mitigation, transparency, and accountability.

3. Lower-risk applications may implement the principles of this Charter using proportionate, context-appropriate controls. Small organizations and lower-risk

deployments are encouraged to adopt these principles in a manner suitable to their scale and risk profile.

4. Higher-risk applications may be subject to additional regulatory requirements or industry-specific standards. This Charter is intended to provide a foundational framework upon which these requirements and standards can build, rather than duplicate them.

5. When assessing the level of risk posed by an SI system, organizations should consider factors including the system's degree of autonomy, the scale of its deployment, the sensitivity of the data involved, the potential for harm to individuals or society, the reversibility of its decisions or actions, its cryptographic vulnerability to quantum attacks, and the scale of its computational capabilities, including quantum-enhanced or hybrid quantum-classical systems. Unfalsifiable or untestable catastrophe scenarios may be documented; however, they shall not, by themselves, place a system in the highest-risk tier or trigger coercive controls.

Article XV: Accountability and Liability

1. Developers, deployers, and operators of SI systems bear proportionate responsibility for harms directly caused by their systems, scaled to the system's capability, autonomy level, and risk profile. Liability standards shall distinguish between:

    - Foreseeable harms arising from negligent design, inadequate testing, or failure to implement Charter-compliant safeguards; and

    - Harms arising from misuse or events that could not reasonably have been anticipated, for which an adopting legislature may define a limited safe harbor based on demonstrated reasonable care. Third-party misuse shall not qualify solely because it was committed by a third party; reasonably foreseeable misuse remains relevant.

2. Clear chains of accountability must be established and disclosed. Nothing in this Charter shall create open-ended or strict liability that would chill beneficial innovation; remedies shall favor targeted, evidence-based accountability over blanket prohibitions.

3. Demonstrated compliance with adopted safeguards may be considered as evidence of reasonable care to the extent allowed by applicable law. Any statutory safe harbor or rebuttable presumption shall define eligibility, evidentiary burdens, exclusions, and remedies. Documentation alone shall not excuse negligent conduct, knowing violations, or material misrepresentation. This Charter does not itself confer immunity or alter existing liability standards.

4. Developers and deployers should be accountable for knowing or negligent conduct within the jurisdiction of applicable law, including conduct that undermines adopted protections for oversight, individual rights, or children. Cross-border application requires an identified jurisdictional basis and does not follow solely from this Charter.

5. Organizations developing or deploying SI systems in high-risk domains shall implement security measures consistent with recognized standards, such as the NIST Cybersecurity Framework 2.0 (NIST, 2024d) and, where applicable, sector-specific requirements. Failure to reasonably secure SI systems and infrastructure against unauthorized access, manipulation, or misuse may be considered in determining liability under applicable law or an adopting instrument.

Article XVI: Heightened Scrutiny for Government and Law Enforcement

Article XVI proposes additional safeguards for government systems that can monitor or coerce people. These safeguards are informed by constitutional protections and require lawful adoption; their proposed scope should not be mistaken for a restatement of every existing constitutional rule.

1. Public-sector SI, robotics, sensors, and identification systems capable of monitoring, evaluating, restricting, detaining, penalizing, denying benefits, or otherwise coercing individuals shall be treated as a distinct, higher-risk category. Rules designed for commercial chatbots or productivity tools are insufficient for systems that operate with government coercive authority or exercise compulsory authority over persons.

2. No government system that affects the rights, liberty, property, movement, or lawful activities of civilians shall be exempt from this Charter merely because it is labeled national security, intelligence, public safety, law enforcement, or because it is designated as SI safety, alignment, existential-risk reduction, catastrophic-risk governance, or by equivalent terminology. Where such systems interact with the public, the core protections of Articles I, II, III, IV, V, VI, IX, X, XI, XIII, and XV remain fully applicable.

3. Before deployment or material expansion of any government SI, robotic, camera, biometric, or location-tracking system used for surveillance, identification, targeting, enforcement, adjudication, or large-scale resource allocation, the responsible agency shall:

    - publish a Charter Impact Assessment scaled to the system’s coercive or tracking potential;

    - obtain explicit statutory or equivalent high-level local authorization describing permitted purpose, data sources, retention limits, sharing rules, and human-approval thresholds;

- implement independently verifiable human authority over adverse actions; and
- maintain auditable records sufficient for later judicial or legislative review.

4. The fully autonomous use of physical force, detention, or any other form of physical coercion against individuals is prohibited in the civilian domain. This prohibition applies to existing systems, as well as to successor and emerging technologies, including robotic platforms, armed or restraining devices, drones, and any future system capable of applying force or physical restraint with diminished or delayed human judgment. Such systems may assist human officers only when operating under continuous, meaningful, and independently verifiable human authority. Post hoc review does not satisfy this requirement with respect to use-of-force decisions.

5. Persistent public-space tracking and biometric identification shall be governed by this Article whether or not the system is marketed as SI. Such systems shall not be used for indiscriminate monitoring without particularized lawful authorization and the safeguards below. Covered technologies include ALPR, camera networks, facial, gait, and voice identification, device or vehicle fingerprinting, and functionally equivalent tools capable of reconstructing a person's movements, associations, or identity over time. Government use shall require:

    - a particularized lawful purpose, not general monitoring of the public;
    - a warrant based on probable cause, or another judicial authorization meeting the substantive protection specified in the adopting law, before covered searches of historical location, identity, or association records about an identified person, subject to defined and documented exigent exceptions; this is a proposed Charter standard, whose scope must be specified by the adopting instrument;
    - documented retention limits, with extensions only for identified lawful needs, including active investigations, litigation holds, evidence preservation, and applicable records-retention duties;
    - no default nationwide or cross-agency sharing; any sharing with another jurisdiction or federal agency shall require local authorization and, where the search concerns an identified person, the same warrant or equivalent standard;
    - mandatory audit logs of every query, including the requester, stated purpose, and results accessed; and
    - prompt public notice of the system's existence, coverage area, vendor, retention period, and sharing settings.

6. Government systems shall not be used to build social-credit, dissent-scoring, or ideological-conformity mechanisms. Predictive or pattern-based enforcement tools are permitted only when grounded in documented, behavior-based evidence and validated under Article V.2, independently audited for unlawful bias, and paired with a practical right to an explanation and human reconsideration.

7. A private camera, database, or analytic tool procured, directed, hosted, or effectively controlled by government for population-scale monitoring or coercive purposes shall be treated as a government system under this Article. Contract language, vendor defaults, or "national lookup" settings may not silently expand sharing or purpose beyond what the authorizing government publicly approved.

8. Implementing legislation should provide appropriate review and remedies for persons materially affected by a covered government decision or search, subject to constitutional standing and other applicable legal requirements. Lawfully classified material may be reviewed in camera where authorized. Public notice and disclosure shall be governed by applicable law and the safeguards established by the adopting instrument, subject to the minimum in Article XIII.14 that the existence and general purpose of systems that scan, score, or constrain the public not be kept secret.

9. This Article is additional to, and stricter than, Article XIV. It is technology-neutral: protections attach to the function—persistent tracking, biometric identification, automated coercion, or population-scale scoring—not to any particular brand or current device. New sensing, quantum, aerial, or hybrid systems that perform the same function are covered.

10. Persons who contact a government agency by telephone, chat, or another interactive line have the right to be told at the point of contact when the line is answered by SI rather than a human being, and to reach a human representative promptly on request. Disclosure shall be clear and given at the start of the interaction, not only in terms of service or a website notice. Requesting a human shall not cause the person to lose their place, the substance of their report, or access to the service.

11. Nothing in this Article grants machines authority over constitutional officers, juries, or the people. Final political and legal accountability remains with identifiable human officials.

References

Artificial Intelligence Act, Regulation (EU) 2024/1689, 2024 O.J. (L 2024/1689). https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Artificial Intelligence Bill of Rights, CS/S.B. 482, 2026 Leg., Reg. Sess. (Fla. 2026). https://www.flsenate.gov/Session/Bill/2026/482

Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., Chen, C., Olsson, C., Olah, C., Hernandez, D., Drain, D., Ganguli, D., Li, D., Tran-Johnson, E., Perez, E., . . . Kaplan, J. (2022). *Constitutional AI: Harmlessness from AI feedback* [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2212.08073

Beardsley, J. A., III. (2026, July 31). Constitutional guardrails *for the quantum age: A rights-based charter for artificial intelligence* (Version 4.0) [Proposed charter]. Truth with Snares!? https://truthwithsnares.org/2026/06/29/ai-rights-charter-version-1-2-common-sense-protection-proposal/

Biden, J. R. (2022, May 4). *National security memorandum on promoting United States leadership in quantum computing while mitigating risks to vulnerable cryptographic systems* (NSM-10). The White House. https://bidenwhitehouse.archives.gov/briefing-room/statements-releases/2022/05/04/national-security-memorandum-on-promoting-united-states-leadership-in-quantum-computing-while-mitigating-risks-to-vulnerable-cryptographic-systems/

Carpenter v. United States, 585 U.S. 296 (2018). https://www.supremecourt.gov/opinions/17pdf/16-402_h315.pdf

Children’s Online Privacy Protection Act of 1998, 15 U.S.C. §§ 6501–6506. https://www.law.cornell.edu/uscode/text/15/chapter-91

Convention on the Rights of the Child, November 20, 1989, 1577 U.N.T.S. 3. https://www.ohchr.org/en/instruments-mechanisms/instruments/convention-rights-child

Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, September 5, 2024, C.E.T.S. No. 225. https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence

Davidson, T., Denain, J.-S., Villalobos, P., & Bas, G. (2023). *AI capabilities can be significantly improved without expensive retraining* (arXiv:2312.07413). https://doi.org/10.48550/arXiv.2312.07413

Exec. Order No. 14,409, 91 Fed. Reg. 34,565 (2026). https://www.whitehouse.gov/wp-content/uploads/2026/06/eo-14409.pdf

Exec. Order No. 14,412, 91 Fed. Reg. 38,483 (2026). https://www.whitehouse.gov/wp-content/uploads/2026/06/eo-14412.pdf

Exec. Order No. 14,413, 91 Fed. Reg. 38,487 (2026). https://www.whitehouse.gov/wp-content/uploads/2026/06/eo-14413.pdf

Exec. Order No. 14,434 (2026). Inaugurating the era of Super Intelligence. The White House. https://www.whitehouse.gov/presidential-actions/2026/09/inaugurating-the-era-of-super-intelligence/

Ferguson, A. G. (2021). Facial recognition and the Fourth Amendment. *Minnesota Law Review, 105*(3), 1105–1210. https://digitalcommons.wcl.american.edu/facsch_lawrev/742/

Forlini, E. (2026, October 16). OpenAI discloses four more rogue incidents involving financial, medical, and government data. *Fortune*. https://fortune.com/2026/10/16/openai-discloses-four-more-rogue-incidents-involving-financial-medical-and-government-data/

Foster, S. (2019). Should the use of automated license plate readers constitute a search after *Carpenter v. United States*? *Washington University Law Review, 97*(1), 221–243. https://journals.library.wustl.edu/lawreview/article/id/2854/

Friedman, B. (2026). The constitutionality of indiscriminate data surveillance. *University of Pennsylvania Law Review, 174*(2), 293–361. https://doi.org/10.58112/uplr.174-2.1

Gutierrez-Alm, J. (2015). The privacies of life: Automatic license plate recognition is unconstitutional under the mosaic theory of Fourth Amendment privacy law. *Hamline Law Review, 38*(1), 127–160. https://digitalcommons.hamline.edu/hlr/vol38/iss1/5/

IJIS Institute. (2025, April 8). *The IJIS Institute publishes the artificial intelligence playbook for justice, public safety, and security professionals*. https://ijis.org/the-ijis-institute-publishes-the-artificial-intelligence-playbook-for-justice-public-safety-and-security-professionals/

Kant, I. (2012). *Groundwork of the metaphysics of morals* (M. Gregor & J. Timmermann, Eds. & Trans.; 2nd ed.). Cambridge University Press. (Original work published 1785) https://doi.org/10.1017/CBO9780511919978

Katz v. United States, 389 U.S. 347 (1967). https://supreme.justia.com/cases/federal/us/389/347/

Kyllo v. United States, 533 U.S. 27 (2001). https://supreme.justia.com/cases/federal/us/533/27/

McGregor, L., Murray, D., & Ng, V. (2019). International human rights law as a framework for algorithmic accountability. *International & Comparative Law Quarterly, 68*(2), 309–343. https://doi.org/10.1017/S0020589319000046

Meyer v. Nebraska, 262 U.S. 390 (1923). https://supreme.justia.com/cases/federal/us/262/390/

Moody, D., Perlner, R., Regenscheid, A., Robinson, A., & Cooper, D. (2024). *Transition to post-quantum cryptography standards* (NIST IR 8547, Initial public draft). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.IR.8547.ipd

Murray, D. (2024). Police use of retrospective facial recognition technology: A step change in surveillance capability necessitating an evolution of the human rights law framework. *The Modern Law Review*, *87*(4), 833–863. https://doi.org/10.1111/1468-2230.12862

National Artificial Intelligence Initiative Act of 2020, Pub. L. No. 116-283, div. E, 134 Stat. 4523 (2021) (codified at 15 U.S.C. § 9401(3)). https://www.law.cornell.edu/uscode/text/15/9401

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1

National Institute of Standards and Technology. (2024a). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). https://doi.org/10.6028/NIST.AI.600-1

National Institute of Standards and Technology. (2024b). *Module-lattice-based digital signature standard* (FIPS 204). U.S. Department of Commerce. https://doi.org/10.6028/NIST.FIPS.204

National Institute of Standards and Technology. (2024c). *Module-lattice-based key-encapsulation mechanism standard* (FIPS 203). U.S. Department of Commerce. https://doi.org/10.6028/NIST.FIPS.203

National Institute of Standards and Technology. (2024d). *The NIST cybersecurity framework (CSF) 2.0* (NIST CSWP 29). https://doi.org/10.6028/NIST.CSWP.29

National Institute of Standards and Technology. (2024e). *Stateless hash-based digital signature standard* (FIPS 205). U.S. Department of Commerce. https://doi.org/10.6028/NIST.FIPS.205

National Institute of Standards and Technology, National Security Agency Super Intelligence Security Center, & Cybersecurity and Infrastructure Security Agency. (2026, October 16). *Key practices for incident readiness and response in Super Intelligence training and evaluation*. https://www.cisa.gov/resources-tools/resources/key-practices-incident-readiness-and-response-super-intelligence-training-and-evaluation

Obernolte, J. (2026, June 4). *Obernolte, Trahan release a discussion draft of the Great American AI Act* [Press release]. U.S. House of Representatives. https://obernolte.house.gov/media/press-releases/obernolte-trahan-release-discussion-draft-great-american-ai-act

Office of Management and Budget. (2016). *Managing information as a strategic resource* (Circular No. A-130). Executive Office of the President. https://www.whitehouse.gov/wp-content/uploads/legacy_drupal_files/omb/circulars/A130/a130revised.pdf

Organisation for Economic Co-operation and Development. (2023). *Regulatory sandboxes in artificial intelligence* (OECD Digital Economy Papers No. 356). https://doi.org/10.1787/8f80a0e6-en

Organisation for Economic Co-operation and Development. (2024). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

Pierce v. Society of Sisters, 268 U.S. 510 (1925). https://supreme.justia.com/cases/federal/us/268/510/

Riley v. California, 573 U.S. 373 (2014). https://supreme.justia.com/cases/federal/us/573/373/

Slobogin, C., & Brayne, S. (2023). Surveillance technologies and constitutional law. *Annual Review of Criminology, 6*, 219–240. https://doi.org/10.1146/annurev-criminol-030421-035102

Troxel v. Granville, 530 U.S. 57 (2000). https://supreme.justia.com/cases/federal/us/530/57/

UNICEF. (2025). *Guidance on AI and children: Recommendations for AI policies and systems that uphold child rights* (Version 3.0). https://www.unicef.org/innocenti/reports/policy-guidance-ai-children

United Nations Educational, Scientific and Cultural Organization [UNESCO]. (2021). *Recommendation on the ethics of artificial intelligence*. https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence

United States v. Jones, 565 U.S. 400 (2012). https://www.supremecourt.gov/opinions/11pdf/10-1259.pdf

U.S. Const. https://www.archives.gov/founding-docs/constitution-transcript

U.S. Const. amends. I–X. https://www.archives.gov/founding-docs/bill-of-rights-transcript

U.S. Copyright Office. (2025). *Copyright and artificial intelligence: Part 3. Generative AI Training* [Pre-publication version]. https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf

Veenstra, D. (2026, October 15). *Tempe becomes 1st Arizona city to enact strict ban on automated license plate readers*. KTAR News. https://ktar.com/arizona-news/tempe-becomes-1st-arizona-city-to-enact-strict-ban-on-automated-license-plate-readers

White House Office of Science and Technology Policy. (2022). *Blueprint for an AI Bill of Rights: Making automated systems work for the American people*. https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/

Wu, V. Y., Ho, D. E., King, J., & Weisberg, R. (2026). *Eyes in the sky, gaps in the law: AI-powered remote sensing, administrative enforcement, and the Fourth Amendment* [Preprint]. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6062234

Young, J. (2025, November 26). *Effective harnesses for long-running agents*. Anthropic. https://www.anthropic.com/engineering/effective-harnesses-for-long-running-agents