

Written Testimony on
Governance of Artificial Intelligence:
Governing AI for Shared Prosperity in a Democracy

Presented at

“Safeguarding Data and Innovation: Building the Foundation for the Use of Artificial Intelligence”

Submitted to

**House Energy and Commerce Committee
Subcommittee on Innovation, Data, and Commerce**

Submitted by

**Laura MacCleery, Senior Policy Director
Claudia Ruiz, Senior Civil Rights Analyst
Policy & Advocacy
UnidosUS**

Raul Yzaguirre Building
1126 16th Street, NW, Suite 600
Washington, DC 20036-4845

October 18, 2023

On behalf of UnidosUS, we respectfully submit this testimony on the pressing issue of governing artificial intelligence and data. UnidosUS is a nonprofit, nonpartisan organization that serves as the nation’s largest Hispanic civil rights and advocacy organization. Since 1968, we have challenged the social, economic, and political barriers that affect Latinos through our unique combination of expert research, advocacy, programs, and an Affiliate Network* of nearly 300 community-based organizations across the United States and Puerto Rico.

As we [described](#) at the AI Insight Forum in September, artificial intelligence (AI) holds enormous promise, but also poses potential threats to civil rights and our democracy absent ethical, constitutionally designed governance. America stands at a crossroads as AI capabilities race ahead, requiring a new social compact between technology and democracy.

A forthcoming series of papers from UnidosUS will propose three pillars for inclusive AI governance—first, integrated oversight for new and existing systems based on constitutional and democratic standards that operate as a baseline for human freedoms; second, participatory engagement by impacted groups; and third, human infrastructure investments to build capacity and ensure the benefits of technology are shared. Below, we suggest some models and ideas for implementing each of these pillars.

Policymakers have proposed principles and standards to govern AI systems, including the [National Institute of Standards and Technology](#) (NIST), [White House, Office of Science and Technology Policy](#) (OSTP), [National AI Advisory Committee](#) (NAAIC), and members of Congress through legislative proposals. However, principles alone are insufficient—our task now is to use these to develop a pragmatic roadmap for implementing ethical AI in practice. This moment demands converting principles into pragmatic action.

Given its reach and power, to deploy AI responsibly and ethically will require new and innovative forms of governance. Systems should ensure their accountability to the people they impact the most—including workers, creators, communities of color and lower-income people, and others who have been left behind and left out by traditional research or the digital divide (and are thus invisible to the models). For too many, the last decades of technological innovation failed to allow them to reap the benefits of a global economy.

As AI systems grow more capable, we must assure that these technologies empower rather than undermine human potential, autonomy, and dignity. Widening economic disparities, information silos on social media, and a scarcity of public spaces already strain social cohesion and imperil democratic principles. If improperly governed, AI could exacerbate polarization, anti-democratic impulses, tribalism, and racial tensions. We must proactively ensure these systems do not inherit and amplify existing biases that can fan the flames of social division.

If legislative approaches advance without an attempt to account for what we already know to be deeply troubling about AI uses, including biased outcomes and rights-infringing practices

* UnidosUS Affiliate Network, <https://www.unidosus.org/about/affiliates/>.

that run counter to constitutional rights around due process, freedom of expression, and the right to be left alone, Congress will have missed the most compelling opportunity it will likely ever have to comprehensively align new technologies with the opportunities they present.

The Promise and Perils of AI

Inclusive, ethical AI systems could expand opportunities for Latinos in areas like education, language access, and employment. Personalized learning software, translation tools that break down language barriers, and job matching platforms custom-built for diverse candidates all represent ways technology could uplift rather than marginalize, disadvantaged groups. Thoughtfully crafted oversight can steer these transformational tools toward expanding freedom rather than undermining it.

Yet the millions of Latinos we represent face double-edged risks from proliferation of AI systems. On the one side, gaps in digital access and tech education exclude many Latinos from emerging economic opportunities in AI fields. About 35% of Latinos lack home broadband, limiting development of the skills required for tech roles where Latinos are already severely underrepresented.

Coupled with lower educational pathways into tech fields, these dynamics widen racial wealth divides. Although the technological advances of the past decade have revolutionized consumption, the racial wealth gap has remained stubbornly unchanged for the past forty years. A shocking lack of diversity in the tech sector persists, with Latinos accounting for only [8% of the STEM](#) workforce, despite also accounting for almost [78% of all new workers](#) by 2030. What is more, the tools to alleviate worker displacement and address equity are underdeveloped.

At the same time, Latinos and other communities of color are subjected to expansive governmental surveillance technologies like predictive policing algorithms and biometric screening tools used against immigrant groups. Opaque AI models used in high-stakes decisions around criminal justice, lending, and benefits entrench historical biases and discrimination, yet recourse is limited or nonexistent. While principles like due process, equal protection, and privacy underpin our laws, outdated regulations fail to provide any accountability for rights-infringing uses of AI.

No democratic government should have the ability to track the activities and thoughts, the whereabouts and networks, the psychological vulnerabilities and personal correspondence, of all of its people. Democracy's lifeblood is the accommodation of diversity and dissent, enabled by associational freedom and zones of autonomy for citizens to develop beliefs away from state interference or coercion. Mass surveillance infrastructures enabled by unfettered AI can devastate the foundations on which democratic self-rule relies.

We must be vigilant against normalizing governmental uses of AI that cross constitutional red lines. Mass [biometric monitoring](#), [predictive policing](#) absent accountability, and [immigrant](#)

[tracking](#) tools raise present threats that demand oversight. America's democratic vitality hinges on governing AI innovations with liberty at the forefront. Ultimately, the choice that is often posed between data security and privacy, on the one hand, and effective law enforcement, on the other, is a false one—[good design](#) can make both a reality, once appropriate incentives and [protections](#) are in place.

For example, predictive policing tools trained on [flawed crime statistics](#) have been found to [disproportionately target](#) low-income neighborhoods of color by falsely correlating race with criminality. Similarly, [sentencing algorithms](#) drawing on racially skewed conviction data likewise entrench harsher outcomes for minorities. [Lending](#) and [credit access](#) algorithmic discrimination also persists, despite legal authorities that require lenders to eliminate it. And medical algorithms and tool design, such as we saw with [pulse oximeters](#) during the pandemic, [can perpetuate health](#) disparities if trained on unrepresentative datasets.

We therefore support the urgent need for data safeguards that are aligned with global approaches and reflect the likely near-future state of changes in the European Union (EU) around both privacy and AI. Although it is tempting to continue avoiding the hard questions, key baselines around data collection practices and permissible uses, consumer rights and product liabilities, and ethical data design and defaults, are now essential. Without guardrails, we will continue to see the types of harms outlined above. We could also face tailored forms of manipulation and control, which can now be easily designed at scale given these new tools, and therefore threaten fragile democratic systems.

Thankfully, the notion that we cannot have both privacy and our cherished freedoms is outmoded and a false choice. As the EU has made clear—security and liberty are essential values to inform ethical design. For example, systems should implement data minimization aligned with protections against unreasonable search and seizure.

Segmenting data use by purpose, rather than allowing unfettered collection, is also important. Transparency and rights that allow impacted people to contest automated decisions reflect due process principles, and enabling user controls over data can help to safeguard freedom of association and expression. Additionally, technical solutions like federated learning and on-device processing can reinforce data minimization and user controls that avoid centralized data storage. With careful design, we can craft data safeguards that honor both privacy and constitutional liberties.

Direction from Congress Is Needed to Address Harms and Establish a Nimble but Effective System of Oversight Aligned with Constitutional and Democratic Systems

Much of the discussion thus far in Congress has largely focused on the exciting frontier models—powerful natural language chatbots, also called Large Language Models, or LLMs—sometimes disregarding the ways that existing (even long-deployed) AI-driven technologies, such as decision models that control access to credit, housing, or public benefits, or furnish predictions for law enforcement, immigration, and judicial decision-making, today inflict harm

on millions, often without redress or attention to the constitutional and legal rights that should apply.

This intense attention on the newest forms of AI and its risks and promise is understandable, but regrettable. Congress should pay close attention to the novel risks from foundational models. However, to address the legitimate harms and create meaningful AI governance, Congress will also have to create oversight of AI where its uses have been neglected and permitted to inflict harm without accountability over the past decade. We must not get distracted by the “shiny new toy” and discount the ways AI already impacts communities, [financial opportunity](#), [workers](#), [health](#), [education](#), and more.

Importantly, a scatter-shot approach to governance will not yield the right level of consistency to guide AI development, now or in the future. Merely regulating frontier models risks enabling the continued development of super-tools that inflict these and other harms, unchecked, as their capacities and reach proliferate across domains. On the other hand, focusing only on specific use cases will miss the forest for the trees.

For this reason, governing AI will require a hybrid approach with coordinated oversight spanning both foundational models and specific high-impact applications in sectors from healthcare to lending where algorithms already operate. Comprehensive reforms must address current and emerging technologies, combining centralized expertise with sector-specific regulations and align them with democratic and constitutional principles.

Congress should develop a specialized regulator for foundation systems, as well as separate dedicated incubators or departments for each major domain regulator with specific authorities and obligations to oversee AI’s implications for research, products and services within their domain. A cross-governmental working group led by the Office of Science and Technology Policy (OSTP) should coordinate activities and approaches, and essential lessons and developments with cross-domain implications should be shared.

Activities on both the foundational regulatory and domain levels should be grounded in legal, constitutional, and ethical principles that support a healthy democracy and humane uses of technology, and address emergent and current harms and risks. Congress should set flexible directional goals, provide ample and clear authority for new standards and evidence-based regulatory processes, and create a strong legislative record to support iterative problem-solving by agencies to address evolving needs and drive innovation that improves the technology’s fit to our values over time.

Because conventional regulatory processes often struggle to keep pace with the speed of technological change, Congress must empower regulators with the flexibility and resources to iterate governance in line with AI’s rapid evolutions. Rulemaking cycles measured in years will fail to address novel challenges in time. We need more nimble oversight mechanisms that will be capable of timely course corrections based on emerging evidence and lessons from a learning community of practice.

Well-designed standards and goals would guide engineers at leading companies to move beyond narrow technical objectives to embrace constitutional and legal standards as primary design criteria. Rules should make clear the relationship between, for example, notions of due process and data transparency, as described above. AI alignment methodologies currently in use by industry that focus on human values like fairness, transparency and accountability should also inform technical development of standards.

Practical new requirements should reflect these goals and be calibrated to reflect risks. For example, Congress should empower the relevant regulators in each domain to regularly audit algorithms and set directional goals that allow flexibility for AI's rapid iterations and experimentation with models, while assuring core safety and transparency remains intact and incentivizing fairer outcomes over time. Immediate disclosures of emergent risks and options to mitigate them should be required, and new deployments should be well managed for safety, security, and other relevant regulatory parameters through licensing, ongoing oversight, and a culture of accountability, candor, and collaboration.

In sum, key principles that are already in use by industry for AI alignment—including making systems helpful, harmless, honest, accountable, transparent, and fair—should be guiding lights for engineers and policymakers alike. Where current real-world systems lag these goals and use AI algorithms, they should be brought into alignment over time with these best practices and standards, in a practice of leveling up. Embedding ethical objectives in the technical architecture itself is vital. Comprehensive governance can reinforce those aims and drive productive innovation.

Predictable oversight frameworks are also imperative to incentivize technology firms to invest in developing rights-preserving innovations aligned with democratic values. Ethically focused companies need market signals that building technologies that respect principles like privacy and accountability will be competitively advantageous. For this reason, it is important that statutory language provide clear standards that drive improvement over time—requiring explainability, accountability, and transparency to be maximized consistent with evolving technological capacities.

Enacting prudent governance now will provide critical assurance to technology leaders that innovations designed responsibly from the outset will be rewarded over more reckless products that monetize surveillance and manipulation. America has an opportunity to lead in demonstrating that values-driven innovation has a viable path to market. Establishing ethical AI as standard practice rather than an upstart niche is essential for an empowering technological future.

Reimagining Inclusive Governance of AI Development for Impacted Communities

Ethical engineering is only part of the equation. Robust public oversight and enforceable rules grounding technology in democratic values are equally vital. America should lead in demonstrating how societies thrive when innovations align with, rather than undermine, Constitutional rights. Prudent governance of transformative technologies is particularly crucial for marginalized groups like Latinos and other people of color whose experiences are often invisible in both policy and technology development.

A core challenge in ethical AI development is ensuring that biased inputs do not lead to discriminatory outputs that violate democratic principles. Historical data powering algorithms often reflects past prejudice against minorities and other groups, reproducing injustice through new technologies. To address the problem and the disruptive effects of technologies, we need new ways of collaborating across sectors to leverage existing domain expertise and to go beyond check-the-box forms of “consultation” with impacted communities to full partnership.

Merely debiasing algorithms is insufficient—ongoing community participation is essential to assess real-world impacts on disadvantaged groups and inform iterative improvements. While frameworks like [NIST's](#) highlight the need for community consultation, specific processes and decision-making power must be defined to make this engagement meaningful, rather than symbolic.

For these reasons, it is crucial for Congress to create a powerful seat at the table for communities impacted by this transformation, so that they may claim their role in determining its uses and our shared future. Putting impacted communities at the heart of our new system of AI governance is essential to uplift ethical considerations, generate evidence on real-world impacts and make it relevant to oversight, and inform the development of better and fairer AI uses over time.

Building technical fluency across participants and establishing clear procedures for “consultation with teeth” are key to inclusive AI governance. AI governance frameworks should include clear obligations for both developers and regulators to conduct ongoing impact reviews. For their part, oversight boards with community representation should have the authority to modify or block high-risk model deployments based on rights, ethical and equity criteria. These should evaluate risks not just to individuals but to communities and society, guided by metrics that track effects on key goals like inclusion, equity, privacy, transparency, and democratic accountability.

Higher-risk AI systems—such as credit and lending algorithms, healthcare diagnostics, and applicant screening tools—are those with significant decision-making authority over crucial aspects of living, including financial opportunity, employment, worker oversight, health access, or education. For such impactful uses, domain-specific advisory boards comprised of impacted community members, subject matter experts and professionals, and legal and ethics experts should be codified into law. These boards would have defined authority to assess AI systems for

disparate impacts using audits and criteria, which should include accountability, transparency, and other ethical dimensions. Boards could have a determinative role in approving, modifying, or rejecting the use of higher-risk systems prior to deployment based on their robust review processes.

Proactive governance measures like mandating diverse training data, independent algorithmic audits, human oversight requirements for critical decisions and functions, and rights-based technical constraints are also imperative, but are insufficient alone. Community participation is essential for surfacing insights technologists alone may miss regarding accessibility, bias, and other impacts. Preventing discriminatory outputs requires continuous vigilance— we must remain cognizant that historical data biases can propagate injustice and be pushing towards fairness and accountability for outcomes in our regulatory designs.

Ambitious yet achievable standards for inclusive evidence-gathering and iteration will help ensure AI systems align with and strengthen shared values over time. Methodologies for ongoing impact reviews should leverage diverse expertise and established best practices, such as:

- Algorithmic audits evaluating systems for biases based on testing inputs and outputs.
- User studies surveying diverse groups on their experiences with AI systems through interviews and observed interactions.
- Consultations with civil society groups focused on relevant issues of equity, privacy, and digital rights.
- Multidisciplinary expert review boards assessing systems and policies from technical, ethical, social, and legal perspectives.

Additional methods like community data audits, participatory design sessions, and “red teaming” to probe vulnerabilities, can provide further evidence. Mediated learning exchanges between developers and community leaders, while complaint redress processes are overseen by regulators, and participatory pilots could help to translate insights into actionable improvements.

An oversight body, such as an AI ethics and safety council, should be tasked with coordinating reviews, tracking compliance, and enforcing practices. Metrics on dimensions like input diversity, harm reduction, and user empowerment should set ambitious yet practical standards. Insights from regularly timed annual reviews should directly inform policy updates, system design changes, adjustments to training data, and be publicly communicated to stakeholders given the evolving nature of the technologies. Review requirements should also be tied to deployment authorizations to ensure that there is accountability. In summary, inclusive evidence-gathering and timely iterations will help ensure AI systems are developed safely, ethically, and in line with public needs and values.

For lower-risk AI applications, a voluntary certification program could help consumers identify technologies with development that included such ethical assessments. Certification levels like

bronze, silver and gold could indicate the rigor of evaluations completed and logos displayed at point-of-use. International alignment on certification standards, where feasible, would maximize global interoperability.

To support essential research and development, policymakers could also establish and fund an AI Safety Institute modeled after the National Institutes of Health. This Institute would partner with academia, industry, and civil society to address gaps in AI ethics understanding and solutions where market incentives fall short. Potential areas of focus could include bias detection and mitigation, algorithmic transparency, privacy-enhancing technologies, and safeguards against manipulation.

The Institute would translate promising research into governance guidance and help to inform technical standards. It could provide grants, host conferences, develop training programs, and undertake participatory pilot deployments to bridge theory and practice. Taking inspiration from biomedical research institutions, an AI Safety Institute could grow the diverse talent and evidence base needed to fairly and broadly distribute societal benefits.

We imagine a new social compact between government, technology companies, and the people served by technology. With intentional design, impacted communities can inform how we steer these powerful technologies toward equitable progress. This would, in turn, expand the legitimacy, reach, accuracy and relevance of these technologies. Multi-perspective input is one key to fostering innovation that serves society broadly, spots gaps in data, develops accurate assessments of impacts, and drives responsible innovation. Most crucially, impacted communities deserve transparency from the AI systems affecting their lives.

Supporting Access to Shared Benefits and to Address Displacement and Impacts

All Americans deserve to reap the fruits of technological progress while retaining an ability to redress its harms. To address the need for these transformational technologies to produce widely shared benefits, Congress should establish a dedicated fund modeled on the [CDC Foundation](#) that could support digital skills education and training, community-based AI auditors, participatory technology workshops, and other capacity building to close knowledge and equity gaps. Grants to local organizations would enable national assessments of AI's impacts on disadvantaged groups and workforce needs and could build expertise and knowledge of technological tools within impacted groups.

It should include funding for community organizations to build AI-specific expertise and support participation in the types of community engagement outlined above. Building technical fluency will allow impacted communities to help steer these powerful technologies toward sustainable and equitable progress and improve their accountability.

To democratize the benefits, there is also a need for human infrastructure capable of leveraging AI, including building capacity in and among communities that are often invisible in the development and deployment of technologies. We must close the skills gap and talent

gap—and a robust program to do so will address the needs of a burgeoning industry while helping to mitigate persistent sources of economic inequality and the lack of shared opportunity.

Substantial funding for creating [a more diverse pool of](#) digital economy workers, through partnerships with culturally competent community-based organizations, as well as [skills-based hiring](#) that connects workers of all skills and backgrounds to an AI economy, would help workers who [are currently excluded](#) from upwardly mobile, future-focused employment find career pathways into the AI economy. Such investments would allow the benefits of AI to be more fully shared, and could help develop expertise and spark interest.

An AI human infrastructure investment program could be modeled on the [CDC Foundation](#)—an independent nonprofit and the “sole entity created by Congress” to mobilize private sector and philanthropic resources to support the CDC’s work on public health. It could be funded by a mix of donations from the technology sector, alongside a federal endowment that grows through licensing fees or other federal supports, to provide grants for community organizations initiatives such as:

- Equipping organizations and their members to use the tools and providing feedback (helping to ensure the digital divide does not render communities or individuals invisible);
- Working to solve the talent gap and develop the evidence base for upskilling and skills-based employment through partnerships and innovative approaches to recruitment of an economically and racially diverse workforce;
- Increasing capacity for diverse voices to engage as peers in regulatory and oversight processes around AI tools, including funding nonprofits to train community auditors on AI systems and conduct audits of algorithms, data collection and uses.

The fund could be overseen by an independent governing board with representatives from government, industry, academia, civil rights groups, and community organizations and be housed within an existing agency or as a nonprofit organization to manage the grantmaking process. Its priorities and processes could be informed by an annual national assessment of AI skills gaps, workforce needs, and barriers to digital inclusion. Nonprofit community organizations could apply to receive grants for purposes such as:

- Providing AI and technology job training programs and internships in impacted areas;
- Developing curricula and certifications for community members to become AI auditors;
- Funding data and policy experts to enable meaningful technical input into AI systems;
- Supporting participatory design workshops and exchanges between developers and community residents;
- Working with recipients to document community partnerships, inclusion plans, and to measure and reflect on impact;

- Funding programs to enable collaborations with schools, employers, and industry partners to create pathways to equitable AI workforce participation.

The fund could also maintain a publicly accessible database of grants awarded, results achieved, and best practices for building AI expertise in local communities to build field knowledge and insights.

This kind of holistic, community-centered approach would build shared capacity to help govern AI responsibly. It would also help to unlock the full potential of America's diversity as a competitive advantage in AI development, maximizing the power of our institutions in the service of improvements to technologies. Sufficient investment at scale in inclusive participation to build and support an AI economy is therefore an imperative for shared prosperity and innovation. The dividend for our national economy could be astonishing.

From Principles to Implementation: Creating an Ethical AI Ecosystem

The table below illustrates our vision for connected and values-driven AI governance. Realizing the full promise of AI while upholding our shared values requires connecting these concepts. It is essential that the rule of law govern these powerful technological tools. Yet ethical principles and technical standards alone are insufficient without inclusive governance processes that give impacted communities an empowered voice in designing and deploying these technologies. And participatory oversight mechanisms need a prepared public, which demands investments to build AI expertise and access within local communities.

We will know we are innovating in the right way only when our tools achieve and advance transparency, access, safety, and equity, among other values, and when the process is participatory and inclusive. But we cannot address biased outcomes, develop inclusive technology, manage the disruptive downsides of AI and its destabilizing potential, or effectively balance law enforcement needs against our valued liberties, unless we center our vision on near-term solutions (while creating stable guardrails around emerging uses that may pose a future risk).

Without open and full participation in an ethical and regulatory framework, AI risks dehumanizing effects like ubiquitous surveillance, manipulation for profit, replicating bias that reinforces injustice, and automated systems that displace human discretion in impactful decisions. Power imbalances mean that AI will primarily benefit the already powerful, exacerbating inequality.

Technical progress must be paired with social progress. Only a holistic approach will enable technology and society to advance in equal measure. Lawmakers must enact comprehensive reforms that move beyond principles to functional guardrails, meaningful participation, and expanded capacity. By linking ethical guidelines, participatory structures, and human infrastructure, we can build an AI future that reflects our democratic ideals.

Most importantly, inclusive democratic practices should infuse AI governance itself. Impacted communities deserve structured involvement in shaping these powerful technologies. With broad collaboration and human-centered design, AI can be governed ethically at the outset, rather than regulated after the fact and once harms are already entrenched.

American leadership in developing AI that expands opportunity while respecting Constitutional values is essential for an inclusive future. Progress will demand openness to creative partnerships and evolving best practices. But we have never shied away from big challenges when ideals hang in the balance. The future of AI governance ultimately relies on rediscovering this pragmatic idealism. If we approach this challenge with courage, wisdom, and democratic faith, we have an opportunity to write a new chapter on technology and progress. The stories that future generations tell about AI may someday inspire the world—if we dare shape them. The task is monumental, but so is the opportunity before us.

##

We applaud the Committee's commitment to exploring effective ways to safeguard innovation and data in the era of AI. We also would welcome the opportunity to collaborate and provide input. Should you have any questions, please contact Laura MacCleery, Senior Director of Policy, at lmaccleery@unidosus.org, and Claudia Ruiz, Senior Civil Rights Analyst, at cruiz@unidosus.org.

CONSTITUTIONAL PRINCIPLES	APPLICATION TO AI FOUNDATION MODELS
Rule of law	AI is a tool bounded by human ethics and legal accountability, and meaningful human oversight is required for all critical decisions and functions
Checks and balances	Democratic governance and regulatory systems provide a check on marketplace harms and abuses; empower users and impacted communities, and standards and processes require timely and nimble adaptation of rules and standards, keeping pace with technological advances
Democratic legitimacy	Undue concentrations of power are addressed by participatory mechanisms engaging diverse representatives from impacted communities, including in design, development, implementation, and monitoring. Additional forms of accountability to users, impacted communities, and co-creators are embedded in legal standards, including design defaults, credit mechanisms, data safeguards and stewardship obligations, and legal redressability. Transparency, fairness, accountability and explainability are core principles that facilitate public knowledge and broaden participation in AI development and shared benefits.
Preservation of zones of personal liberty and limits on overreach	Rules balance need for law and order with the right to be left alone and to move freely in the world; privacy by design and guardrails on both governments and private actors that preserve individual freedoms and minimize the risk of coercion, manipulation, and domination over individuals and groups
Accountability for harms	Bans on specific high-risk uses accompanied by stricter review of uses in matters essential to human functioning and freedoms; legal liability for AI-driven infliction of specific harms (for example, bias, fraud or predation) and a general duty of care to safeguard against intrusions on human liberties, democratic participation, safety, and security
Protection of freedom of speech and expression	A system of AI-assisted monitoring, provenance and watermarking ensures that truthful and political speech are supported and protected while false and misleading or fraudulent forms of communication are discernable and subject to enforcement review as appropriate; creative and satirical uses are defined and legally protected
<p style="text-align: center;">Complementary Domain-Specific Approaches in Health, Financial Services, Employment, Education and More</p> <p>Coordinated with cross-departmental leadership, domain-tailored rules cover specific risks and promote access to benefits of AI in areas like financial wellbeing, health care, research, voting and elections, education, labor and jobs, national security, and other common applications, developed by extending agency’s legal and practical authorities and expertise, through collaborative governance and multi-stakeholder participation, alongside funding to address human impacts and cover rapid response and longer-term needs. Impacted communities should have a defined participatory role and be supported through disruptions and positioned to share in the benefits of advances. Approaches should balance support for responsible innovation with security, safety, and reliability, and track timely and evidence-based assessments of impacts that spur positive iterations and evolve acceptable practices to maximize human agency and shared benefits, and to minimize potential harms. They should be aligned with the constitutional framework and developed in coordination with democratic allies globally.</p>	