

Executive Summary:

The Future of Global AI Governance

July 2023

大成 DENTONS

SPATIAL WEB
FOUNDATION

VERSES

This report introduces a groundbreaking generation of socio-technical standards that will help drive the global discourse on artificial intelligence (AI) governance. These standards, if adopted, have the potential to address what current regulatory and governance efforts cannot: the critical issues of interoperability, explainability, and AI's exponential advance toward greater intelligence and autonomy. In addition, the report introduces a rating system for AI that proposes new governance frameworks to accommodate increasingly intelligent and autonomous AI systems.

This first-of-its-kind perspective on global AI governance combines the legal expertise of the world's largest law firm, Dentons, the AI acumen of VERSES, and guidance on socio-technical standards from the Spatial Web Foundation.

In 1950, Alan Turing posed a fundamental question, "Can machines think?" Today, we are on the verge of answering that question with the creation of AI systems that can imitate our cognitive abilities, interact with us naturally, and even appear to demonstrate human-like thinking. The recent and rapid ascent of generative AI in particular has ushered in a new era for artificial intelligence, introducing systems that showcase remarkable human-level abilities in generating content such as text, media, and software code.

These new AI systems hold great promise. They also raise serious concerns about safety, accountability, ethics, intellectual property rights, fairness, accuracy, bias, explainability, and privacy. As AI continues to evolve and increasingly operates seamlessly in our everyday lives and in the physical world, these risks will undoubtedly multiply.

The rapid development and widespread adoption of generative AI has fueled a global debate about the most effective path forward on AI governance. Although there have been calls for a new era of law and regulation, and a growing global consensus about the need for universal AI governance frameworks, the best path forward remains uncertain.

For example, although traditional market-driven approaches to AI regulation may encourage innovation, they also raise concerns about a "winner-takes-all" approach that could result in a

poverty of options for consumers, an unhealthy concentration of power, and even geopolitical unrest. Traditional governmental approaches to regulation and legislation likewise face limitations. Regulation and legislation are often slowly adopted, and could result in conflicting frameworks. This approach, if not measured, has the potential to stifle growth, prevent interoperability, and create a hard-to-navigate patchwork of AI guidelines throughout the globe.

Current approaches to developing AI regulation and governance frameworks are only addressing the tip of the iceberg. They have largely focused on imposing **social standards** on the organizations and individuals who develop and deploy AI. For example, many jurisdictions focus on imposing requirements around transparency, explainability, accountability, safety, and other societal or human-centered values.

Although these approaches are logical, they generally lack corresponding technical guidance, creating implementation and compliance challenges. They will also hit a ceiling. By primarily focusing on the developers and deployers of AI rather than the AI systems themselves, regulators run the risk of neglecting the trajectory of AI systems as they become increasingly intelligent and autonomous. It is therefore crucial for the future of AI governance that regulatory and governance frameworks include the potential of AI

systems to self-improve, self-adapt, and establish networks with other AIs, sensors, and machines.

Critically, future regulatory and governance efforts must take into account the inevitable convergence of AI with robotic and Internet of Things (IoT) systems, giving rise to Cyber-Physical Systems (CPS). The convergence of CPS with AI could lead to the emergence of **Autonomous Intelligent Systems** (AIS) that operate autonomously across digital and physical domains. AIS will usher in a new generation of the web, powering various intelligent applications, from smart assistants to smart cities to smart supply chains. Similar to the autonomic nervous system's intelligent regulation of the body, AIS could seamlessly orchestrate countless activities in the background of our lives, with increasing autonomy. The word autonomy means "self-regulating" or self-governing. Therefore, effective AI regulation and governance must account for future AI systems that can govern themselves.

Technical standards, the forgotten foundations of technological societies, play a critical role in future AI regulation and governance—but they may not be enough.

In May 2023, the G7 leaders issued a joint communiqué calling for the development and adoption of international technical standards to govern the development and deployment of AI. Similar calls for **technical standards** have recently come from private and public stakeholders. These calls are logical, as technical standards have played a pivotal role in shaping the modern world, providing significant benefits for governments, developers, manufacturers, and consumers alike.

"Cars run on gas from any gas station, the plugs for electrical devices fit into any socket, and smartphones connect to anything equipped with Bluetooth. All of these conveniences depend on technical standards, the silent and often forgotten foundations of technological societies."
– The New York Times

But this time, technical standards alone may not be enough. AI is poised to impact both technology and society in unprecedented ways. The G7

communiqué recognized this tension by stressing the importance of "social implementation" of technical AI standards. The US Department of Commerce's National Institute of Standards and Technology (NIST) likewise has warned against an over-reliance on purely technical standards when it comes to addressing AI. NIST has specifically suggested a hybrid approach of developing socio-technical standards. Going a step further than purely technical standards, **socio-technical standards** aim to bridge the gap between technology and society.

Emerging technologies guided by socio-technical standards could enable AI and AIS to be technically sound, socially beneficial, safe, compliant with laws, and able to be aligned with societal norms and values.

Developing socio-technical standards for AI governance could provide countless benefits. Socio-technical standards could enable governments and regulators to act in a coordinated fashion to proactively ensure compatibility and compliance. Such standards could provide a customizable framework that would allow regulators across various regions to implement their own regulations for AI while maintaining global interoperability. Perhaps most importantly, socio-technical standards would account for the potential trajectory of AI from the "narrow" AI of today (i.e., designed for specific tasks within limited domains), to the "general" AI on the horizon (i.e., those able to learn, understand, and apply knowledge across different domains), and ultimately to tomorrow's "super intelligent" AI (i.e., that which surpasses human-level intelligence). In the end, socio-technical standards may provide a path to AI governance that keeps pace with the exponential evolution and increasing autonomy of intelligent machines.

Introducing a new generation of socio-technical standards to govern AI.

In 2016, the Institute of Electrical and Electronics Engineers (IEEE)—the world's largest association of technical professionals, whose mission is to advance technology for the benefit of humanity—launched the Global Initiative on Ethics of Autonomous Intelligent Systems to address ethical, legal, and social concerns in AI and autonomous technology design and development.

In 2020, The IEEE P2874 Spatial Web Working Group was formed to lead the development of socio-technical standards for AI and AIS alignment, interoperability, and governance. These standards, which are currently being drafted, are informed by IEEE's Ethically-Aligned Design P7000 Series of standards that address human rights, well-being, accountability, and transparency for AI and AIS.

The IEEE P2874 standards are being developed to address the following needs as it relates to building a future global AI governance framework:

- 1. Ensuring a shared understanding of meaning and context between humans and AIs.**
- 2. Ensuring explainability of AI systems, enabled by the explicit modeling of their decision-making processes.**
- 3. Ensuring interoperability of data and models that enable universal interaction and collaboration across organizations, networks, and borders.**
- 4. Ensuring compliance with diverse local, regional, national, and international regulatory demands, cultural norms, and ethics.**
- 5. Ensuring authentication and credentialing, driving compliance and control over critical activities, with privacy, security, and explainability built-in by design.**

This new generation of socio-technical standards are being developed to scale at the speed of AI evolution. If adopted globally, these standards

could enable us to steer AI systems, even those that exceed human-level intelligence. These standards lay the foundations for the efficient integration and adoption of AI technologies while minimizing the risk inherent in AI.

In order to leverage these new socio-technical standards, governments and regulators will need to adopt new frameworks for AI governance.


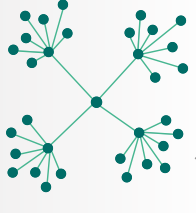



As AI and AIS evolve toward greater levels of intelligence and autonomy, new forms of governance will necessarily emerge, taking into account factors such as the intelligence capabilities, autonomy levels, and trustworthiness of AI systems.

This report outlines a proposal for various AIS governance frameworks to be adopted, ranging from centralized to federated to distributed, each one differing in the degree of control and decision-making authority they afford. Such frameworks, if adopted, would enable stakeholders to select the desired level of autonomy they are comfortable granting to an AI system to operate within a particular domain.

In addition, to assist regulators and stakeholders in assessing which governance frameworks are appropriate, and to comprehensively evaluate the capabilities and limitations of AIS under these frameworks, this report proposes the adoption of a multilevel AIS International Rating System (AIRS). Akin to the Society of Automotive Engineers (SAE) levels used for self-driving vehicles, AIRS would apply to any AI-powered system, ranking the levels of intelligence and autonomy exhibited and providing potential corresponding governance frameworks at each level.

These levels and governance frameworks could replace the hierarchical ranking system being adopted in many jurisdictions around the globe.

AIS International Rating System (AIRS)

	AIS 1	AIS 2	AIS 3	AIS 4	AIS 5
Intelligence Level	<p>Systematic</p> <p>Recognizes and responds to patterns. Follows predefined rules. Abilities limited to specific tasks or domains.</p> <p>Narrow Intelligence</p>	<p>Sentient</p> <p>Perceives its environment and responds in real time. Exhibits curiosity and seeks out information to update its model.</p>	<p>Sophisticated</p> <p>Learns and adapts to new situations and plans based on the consequences of actions or beliefs about the world.</p> <p>General Intelligence</p>	<p>Sympathetic</p> <p>Understands and responds to the emotions and needs of humans and other AIs and considers different perspectives.</p>	<p>Shared</p> <p>Works together with humans, agents, and physical systems to solve complex problems, outperforming humans at most tasks.</p> <p>Super Intelligence</p>
Governance Potential	 <p>Centralized</p>	 <p>Hierarchical</p>	 <p>Federated</p>	 <p>Decentralized</p>	 <p>Distributed</p>

©Spatial Web Foundation

This table illustrates the proposed AIS International Rating System (AIRS).

While current regulation is focused on individual AI models, this report argues for the need to govern the network upon which AI and AIS will operate.

Today's AI systems are generally regarded as siloed models that do not interact with one another. The future of AI, however, lies in interoperable networks of autonomous AI agents working together to solve complex, dynamic problems. Socio-technical standards like those being developed within the IEEE P2874 WG will provide a framework for a holistic and adaptive approach to managing these increasingly complex autonomous AIS networks.

Networks of AIS may include different levels and types of intelligent machines, requiring an AI governance framework that can accommodate dynamic and increasing levels of complexity, intelligence, and autonomy.

Socio-technical standards would grant stakeholders the flexibility to govern as they deem appropriate, while also ensuring seamless interoperability across the network, regardless of the diversity of systems involved.

When socio-technical standards are coupled with the various governance models this report proposes, such standards enable the governance of the activities of AI systems and networks directly—all while maintaining the flexibility to accommodate a plurality of local, national, and international regulatory demands, social values, and human rights.

Socio-technical standards facilitate harmonious interactions that align with human values, in part by enabling law as code.

We cannot govern AI systems and the machines they will power in the same way we govern humans. Machines do not respond to punishment, nor are they currently bound by empathy or ethical concerns. To create AI systems that will remain aligned with our values and interests, we will need a means of encoding our principles and guidelines directly into the AI software and systems.

The IEEE P2874 Standards can convert regulations into a machine-readable, machine-sharable, and machine-executable format.

This “law as code” breakthrough offers a means to close the gap between societal values, regulatory policies, legal processes, and technological mechanisms, leading to more equitable and accessible benefits and services of AI.

Early versions of the socio-technical standards in development at the IEEE have already been successfully field tested. The European Commission's [Flying Forward 2020](#) (FF2020) research project is revolutionizing urban air mobility by enabling AI-powered drones to comply with regional laws as they autonomously navigate urban airspace across five EU countries. Through a geospatial digital infrastructure and new socio-technical standards, FF2020 has demonstrated that autonomous compliance to law—regionally and across the EU—is not only possible but proven.

Imagine the impact of AIS built on interoperable, global socio-technical standards that operate across every sector of the economy. These systems

could align with our values and foster symbiotic relationships with humans, amplifying our strengths and assisting us in overcoming our weaknesses. They could be a trusted guide through an ever more complex, interconnected, and dynamic world, transcending the very notion of “smart” by enabling us and the systems they operate in to become ever smarter.

Interoperable socio-technical standards will enable AI and AIS to seamlessly adapt to real-time changes, enhancing our safety and enjoyment. These standards become the bedrock upon which smart devices and systems collectively evolve to become smarter through their participation. The benefits of global, interoperable standards include more efficient transportation, more effective and affordable healthcare, more personalized and accessible education, better coordination in responding to humanitarian and environmental crises, and a global supply chain that operates at an unprecedented level of resilience and optimization, reducing waste and energy consumption.

AI that harnesses the power of natural intelligence could pave the way to smarter governance and eventually a smarter world.

In the not too distant future, AI systems may evolve beyond large language models to become small, agile, adaptable, explainable, and increasingly autonomous “Intelligent Agents.” Similar to how intelligence manifests in nature, agents would share knowledge, ask questions, and demonstrate curiosity about the world.

Unlike today’s complex, data-hungry AI systems, these agents would require minimal training, relying instead on smaller amounts of highly contextualized data—what one might call “smart data” as opposed to the currently ubiquitous “big data” approaches.

While agents are specialized to accomplish specific tasks at an expert level, they would have the ability to communicate with one another, continuously exchanging knowledge to tackle complex, dynamic challenges, all while retaining the capacity to acquire new information. Essentially, these agents would operate in a manner akin to human

intelligence but at a scale and speed far beyond our capabilities.

A proven framework known as "Active Inference," developed by the renowned neuroscientist and theorist Karl Friston, has the potential to bridge the gap between neuroscience and AI. Active Inference has demonstrated practical applications across multiple disciplines, including neuroscience, psychology, theory of mind, and robotics. Currently, efforts are underway to utilize Active Inference in the development of genuinely intelligent AI systems.

By harnessing insights from the inner workings of the human brain, Active Inference-based AI systems can "introspect," offering detailed explanations for their actions while continuously improving their intelligence and understanding of the world around them. Active Inference can enable smarter governance by empowering AI systems to become explainable, accountable, and adaptive.

These systems can then better align with our goals and values, cultivating trust and transparency even as they become self-governing.

In this report, we explore how smarter AI governance goes beyond restrictions and controls, enabling us to steer AI toward its most promising capabilities. By leveraging socio-technical standards and new approaches to AI, smarter AI governance results in safe, explainable, compliant systems. These systems could enable a smarter world where Autonomous Intelligent Systems free us from mundane physical and mental tasks, allowing us to focus on innovation, self-actualization, creativity, and new horizons.

Download the full report:

www.verses.ai/aigovernance

Authors

Gabriel René, Philippe Sayegh, Capm Petersen, Chase Pletts, Dan Richardson, Peter Stockburger, Eric Tanenblatt

For more information please contact:

The Spatial Web Foundation

inquiries@thespatialweb.org

VERSES

inquiries@verses.ai

Dentons

peter.stockburger@dentons.com

The Spatial Web Foundation is a non-profit organization dedicated to the development and ethical use of technology, specifically in the creation and implementation of Spatial Web standards. The foundation's core initiatives include developing open standards and protocols, promoting interoperability, and educating the public and policymakers. We support transparency and accountability in the development and use of technology, promote responsible innovation, engage in dialogue and collaboration with various stakeholders, and adopt a proactive approach to addressing emerging ethical challenges. Ultimately, the Spatial Web Foundation seeks to create a more inclusive, accessible, and equitable internet that empowers individuals and communities to connect, create, and thrive. For more information: spatialwebfoundation.org

VERSES is a cognitive computing company specializing in next-generation Artificial Intelligence. Modeled after natural systems and the design principles of the human brain and the human experience. Built on open Spatial Web standards (IEEE P2874) our platform transforms disparate data into knowledge models that foster trustworthy collaboration between humans, machines, and AI, across digital and physical domains. Imagine a smarter world that elevates human potential through innovations inspired by nature. Learn more at [VERSES](#), [LinkedIn](#), and [Twitter](#).

Dentons is designed to be different. As the world's largest global law firm with 21,000 professionals in over 200 locations in more than 80 countries, Denton's can help you grow, protect, operate and finance your business. Denton's polycentric and purpose-driven approach, together with their commitment to inclusion, diversity, equity and ESG, ensures they challenge the status quo to stay focused on what matters most to you. For more information visit: www.dentons.com