Federal Register Notice: 89 FR 51554, Federal Register :: Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development, June 18, 2024.

Request for Information on the National Digital Twins R&D Strategic Plan

Spatial Web Foundation

Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development

Dear Office of Science and Technology Policy,

The full promise of Digital Twins will only be realized when they are deployed as an ecosystem of virtual representations of systems interacting with natural and artificial intelligence agents for decision-making. This multi-directional interplay between Digital Twins, physical systems, and agents requires new R&D activities to achieve trustworthy interactions. Research is needed in providing common interfaces to the variety of Digital Twin modeling methods; in defining robust interoperability between Digital Twins, sensors, actuators, and agents; and in fostering sustainable ecosystems of Digital Twins as sources of knowledge.

The Spatial Web Foundation (SWF) [1] is dedicated to the development and implementation of socio-technical standards that will provide a safe and secure and interoperable foundation for the Spatial Web. SWF is a community of developers, creators, scientists, innovators and ethicists with a shared mission to enable a hyper-connected, contextually aware, ethically-aligned network of humans, machines, and artificial intelligence.

The Spatial Web provides the holistic and coherent technical framework for the implementation of a collaborative, interactive, interoperable ecosystem of Digital Twins. As defined in IEEE P2874 Spatial Web Protocol, Architecture and Governance standard, [2] the Spatial Web is the conceptual and distributed computing system-of-systems for a shared world system of agents and knowledge domains. The Spatial Web depends upon Digital Twins to provide models of the natural, engineered, or social systems as the basis for agents to perform activities that meet the needs of socio-technical stakeholders. It aims to provide the technical framework required to handle the challenging topics of data integration and management, interoperability, verification and validation, and ethical and security concerns.

The specification and standards were accepted by the IEEE in July of 2024 and will be made available to the public by the end of the year. We humbly welcome the readers of this RFI to join our working group and would greatly value your guidance and input. See The Spatial Web Protocol, Architecture and Governance Specification Summary [3] for more information.

Trustworthy: Governance of AI and Digital Twins
In 2016, the IEEE launched the Global Initiative on Ethics of Autonomous Intelligent Systems (AIS) [4] to address ethical, legal, and social concerns in AI and autonomous technology design and development. These standards are informed by IEEE's Ethically-Aligned Design P7000 Series [5] of standards that address human rights, well-being, accountability, and transparency for AI and AIS. The IEEE P2874 [2] standards are being developed with SWF to address the following needs as it relates to building a future global AI governance framework:

1. Ensuring digital twins provide 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.

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, we propose 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. Dentons, The Spatial Web Foundation and VERSES Inc have published a report outlining their approach in this Future of AI Governance Executive Summary [6]

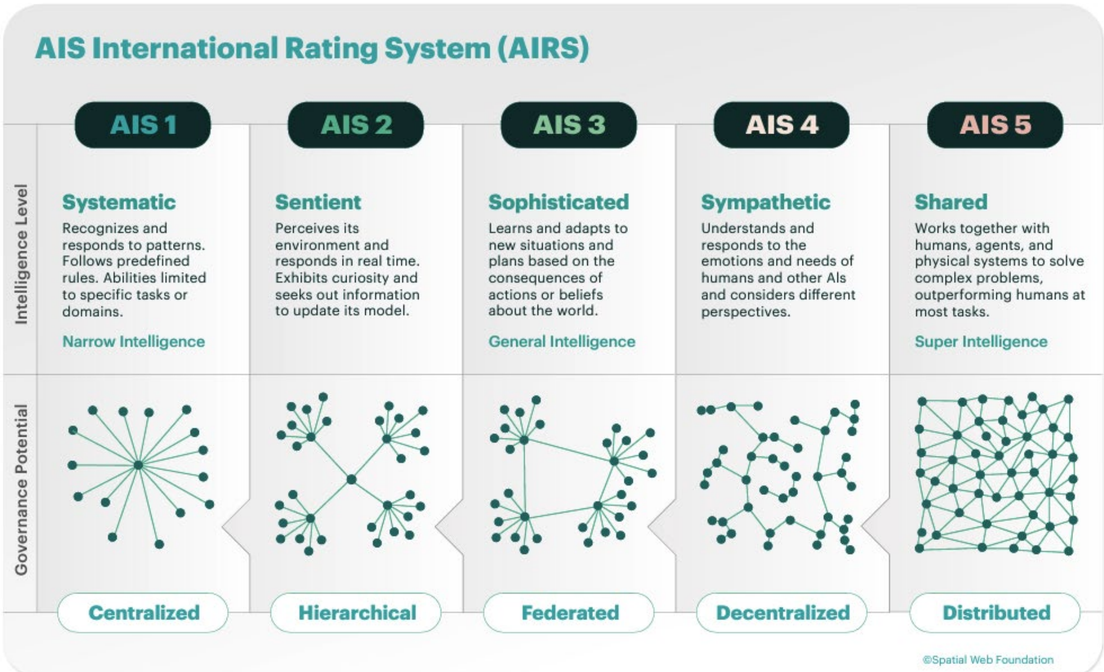

This table illustrates the proposed AIS International Rating System (AIRS).

Domains and digital twins can be used to simulate activities within a domain at scale, ensuring they comply with all safety, security and ethical standards before receiving the credentials required to be deployed into the real world.

R&D is needed in order for governments, regulators and industry to leverage these new socio-technical standards and adopt a new, sustainable framework for AI governance.

VVUQ: AI and Digital Twins need a standard language for trustworthy interactions. The Spatial Web has defined an ontology that defines the primary entities needed for a language that allows AI Agents to evaluate relevance of a Digital Twin to the Activity being performed by the Agent. Once a Digital Twin is identified, the AI Agent defines and structures Activities to be executed over time. In some cases Contracts between the AI Agent and Digital Twin will need to be agreed upon in order to establish the basis for reliable execution of the Activity. The Spatial Web Ontology has been encoded in the Hyperspatial ModelingLanguage (HSML). The ontology was developed based on careful examination of several existing ontologies, i.e., IEEE

7007™-2021, ISO/IEC 21838-1:2021, ISO/IEC 21838-1:2021, and the Suggested Upper Merged Ontology (SUMO).
R&D is needed to elaborate, test and confirm that the Spatial Web ontology provides trustworthy interoperability between AI Agents and Digital Twins. Furthermore, research should be conducted to ensure universal, cross domain interoperability and standardization, enabling AI agents to access data and execute activities across a system of systems.

VVUQ: A common interface is needed to uniformly access Digital Twins with varying representations of space and time. Digital Twins represent systems using information constructs that mimic the structure, context, and behavior of the system. The most primary constructs of Space and Time are highly variable. The Spatial Web approach to modeling Digital Twins is based on an abstract definition for hyperspace, i.e., A set of 'points' such that, for every ordered pair of points, there is a set of 'paths' from the first point to the second, This abstract definition includes all classes of hyperspace as shown in figure 5; topological manifolds, vector spaces, graphs, hypergraphs, cellular structures and abstract data types enabling digital twins to model any type of hyperspatial data. The Spatial Web is using Category Theory to define a common interface across all classes of hyperspace [7].

Additional R&D is needed to elaborate, test and confirm a common interface to all hyperspace representations used in Digital Twins based on Category Theory.

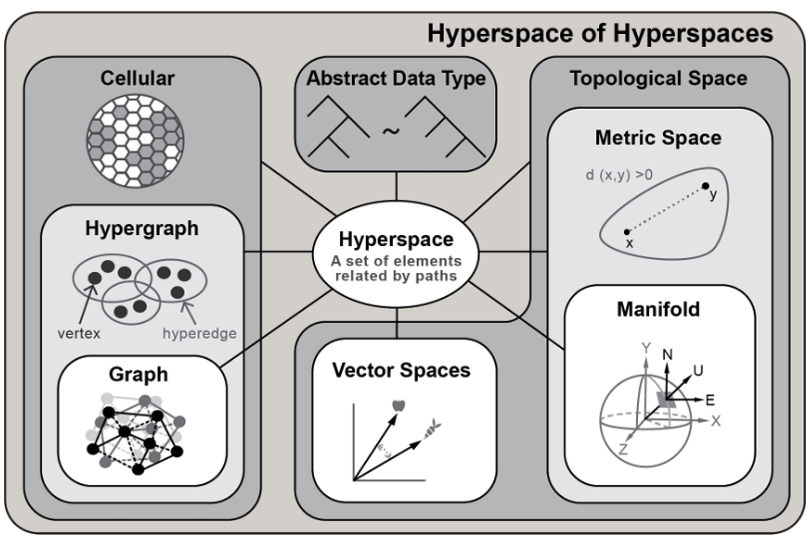

Figure 5 —Basic classes of hyperspace.

VVUQ: Scalable and sustainable ecosystem of Digital Twins as domains of knowledge. Individual Digital Twins mimic the structure, context, and behavior of a natural, engineered, or social system. Collectively, networks of Digital Twins will represent a knowledge structure about domains of the observable world. The Spatial Web provides for a knowledge model of any domain based on graph structure encoded in the HSML. This graph of all domains and their relationships in the Spatial Web is termed the Universal Domain Graph (UDG). Based on current large knowledge graphs and a cellular model of the surface Earth at decimeter-scale, the UDG is estimated to eventually contain approximately 10^{14} nodes. A distributed server instructure is needed to manage distributed storage and performance scalability of the UDG. *Additional R&D is needed to elaborate, test and confirm feasible implementations of the UDG as a basis for the Digital Twin Ecosystem.*

We view the Spatial Web as a cyber-physical ecosystem of natural and synthetic sense-making, in which humans are integral participants -- what we call "shared intelligence". This vision is premised on active inference, a formulation of adaptive behavior that can be read as a physics of intelligence, and which inherits from the physics of self-organization. In this context, we understand intelligence as the capacity to accumulate evidence for a generative model of one's sensed world -- also known as self-evidencing. Formally, this corresponds to maximizing (Bayesian) model evidence, via belief updating over several scales: i.e., inference, learning, and model selection [8].

By ensuring data interoperability, AI agents can query data, make informed decisions and execute activities across domains without needing to be retrained or re-coded to operate within a specific domain, greatly accelerating the adoption and deployment of digital twins.

R&D is needed to ensure HSML can support a multitude of intelligence types to reason and plan across any domain and digital twin.

VVUG: HSML and factor graphs document databases as means of quantifying uncertainty. The Spatial Web Foundation in collaboration with VERSES have developed a new technology called factor graph document databases, or FDD. FDD are databases capable of storing data under both vectorised and graph representations, following the ontology of HSML. Crucially, they also have computational capabilities (e.g., Bayesian and approximate Bayesian inference), meaning that they can update the value of their related entities by performing computations over those entities. They can infer and predict change in their stored entities, and importantly, provide an estimate, through probabilistic inference, of the uncertainty over the relation between those entities. Concretely, this means that FDDs, because they use HSML as their modeling language, can be used to produce digital twins of different real world systems, and perform computations over the entities that make up those systems – the kind of computation that allows simulation, but also the kind of computation needed to measure the uncertainty of the aspects of the real world system's modeled by the FDD. *Additional research is needed to validate FDDs as means for probabilistic inference.*

Regulatory Science: Digital Twins need to represent the rules and laws within their domains that can be interpreted by both humans and machines. Existing rules and regulations have been drafted traditionally by humans and for humans, but with little regard for interpretability and executability by a machine. There is tremendous value in seeking to bridge the gap between the traditional way of drafting and publishing regulation (i.e., by humans for humans) and having machine-readable and machine-executable regulation (i.e., by humans, for machines). Translating existing regulation into machine-interpretable and machine-executable code will allow us to properly capture how humans read and interpret the regulation and ensure machines and AI universally follow the same rules and regulations.

In order to accurately represent laws, digital twins need to be hierarchical in a nested architecture that inherit the rules above them and pass their rules to the ones below them. For example, the autonomous car industry needs a national, state and local digital twin registry, including all the laws and ordinances required to safely operate on the roads. These rules then need to be accurately applied to hierarchical scales of actors, agents and IoT devices, from swarms of vehicles to individual sensors. Such a registry would enable regulators to make real time updates to laws so cars are aware of construction, hazards, and adjust laws based on dynamic factors like weather.

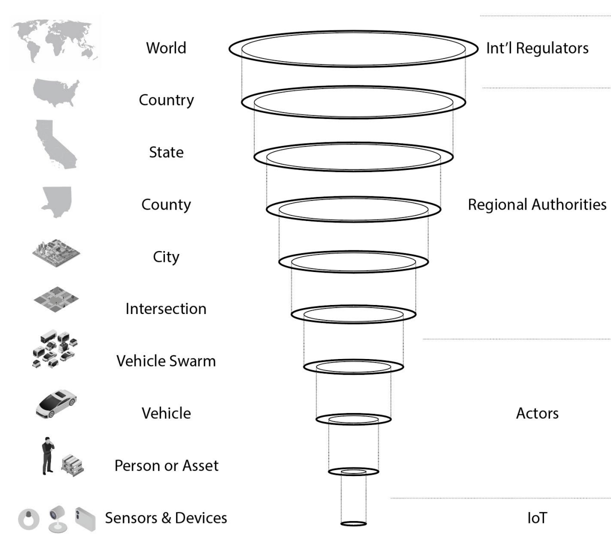

AI can be used to ingest the massive amounts of existing legal text and regulations, applying it to digital twins where it can be administered by humans to ensure its accuracy and compliance. Real world scale simulations can then be run to ensure the machine code accurately represents the laws and recommend adjustments to increase safety and efficiency, eventually becoming a self learning and improving system.

R&D is needed to accurately represent laws to digital twins and run them in simulation to ensure their efficacy.

Business: Regulators need a framework to digitally transform the industries they regulate. Over the coming decades, entire industries will become smarter, connected, automated, secure, and more sustainable. This initiative requires the digital transformation of most if not all commercial and public entities including governments, institutions, corporations, customers and their interactions while preserving their privacy and security. Doing so will greatly transform the economy, accelerating growth and innovation while reducing the commercial friction of regulations and bureaucracy.

The Spatial Web Foundation provides the framework and standards necessary for industry leaders and regulators to create, regulate and operate industry specific domains such as transportation, healthcare, energy and space exploration. Further research is needed to design economic ecosystems and frameworks that provide the tools required for regulators and industry stakeholders to deploy, manage and operate within the industry domain. *Additionally, research is needed to ensure ethical deployment of such systems.*

Business: Regulators, lawmakers and industries need an economic simulation model. In order to accelerate the growth of high risk industries that require significant financial and R&D investments, a global scale economic digital twin is needed to simulate potential costs and benefits of various business models. For example, in the case of space and lunar of exploration, aerospace companies and space agencies can model and simulate various kinds of exploration activities, helping to define which organizations will be responsible for specific activities, what the commercial benefits could be and what additional services are needed to be supported by new industry startups or taken on by the government. Such a simulation could also serve as an investment vehicle, enabling companies to help validate their business assumptions and attract investment.

Or in another example, such a system could help lawmakers and researchers craft tax and incentive models to more accurately predict the economic and social benefits of policies before making it law. As economic and tax decisions are put into action, analysts and AI agents should be able to validate how well the simulation reflected the real life deployment of the project, determining where the model should be adjusted for future simulations, creating a self learning and increasingly accurate representation of economic reality. Furthermore, these economic models should feed into one another, enabling it to more accurately represent the economy as a whole.

R&D is needed to take into account the various stakeholders of such a model, ensuring it is flexible and scalable to adapt to any industry and ensures that it can accurately represent the socio-economic principles of the industry it is meant to serve.

Spatial Web scenarios that utilize Digital Twins

Spatial Web Scenario: Cross Platform Simulation and Coordination of Lunar Exploration. In order to autonomously coordinate activities on the moon, rovers and landers must communicate position and planning, however existing maps from previous fly-bys of the moon have insufficient resolution for planning optimal, energy-efficient paths through or around obstacles such as rocks, craters and regulus (moon dust). Simulating and executing these routes requires rovers to digitize the terrain and share real time position and physics dynamics in a shared multi-dimensional world model. However interoperability between systems and organizations is a gating factor to autonomous collaboration on the moon just as it is on Earth.

The Spatial Web Foundation and JLP have completed a research project to demonstrate the sharing of world models and data between a rover and lander.

Spatial Web Scenario: Digital earth / Greenhouse Gas Monitoring

The scenario describes how the multi-level cognitive computing network of the Spatial Web can facilitate the interaction of greenhouse gas (GHG) monitoring with emitters and goals to collectively take action to reduce atmospheric carbon and the resulting impact on the Earth's climate.

Spatial Web Scenario: Urban digital twin / Smart city

Digital twin technology coupled with AI is improving how cities function, meeting major challenges of human civilization including health, climate change, and sustainability. Urban Digital Twins (UDTs) move us toward the UN goal to make cities and human settlements inclusive, safe, resilient and sustainable.

This scenario shows how the Spatial Web enables digital twin technology for addressing urban sustainability with a focus on energy. Using urban energy system modeling and action plans developed in the Spatial Web multi-scale cognitive computing ecosystem will benefit next generation cities and the globe. Considering climate change it is vital that more cities deploy energy UDTs to address energy consumption and climate sustainability.

Spatial Web Scenario: The Metropolis Energy Model Digital Twin domain

The Metropolis Digital Twin — Energy Model Domain is heterarchially related to the Metropolis Digital Twin Domain and the Regional Energy Grid domain.

To support the Spatial Web Scenario for Urban digital twin several domain models are needed:

- Metropolis geographic domain is a Spatial Web geographic containing the geometry of relevant features that are under the jurisdiction of the Metropolis government.
- The Regional Energy Grid Domain is a model of the grid electrical technology using a conceptual graph for the substations, transmission lines , etc
- Metropolis Digital Twin — Energy Model Domain merges these two domains

R&D is needed to further research the above examples, working with stakeholders to accurately represent the ecosystem of their industry and properly apply digital twins and AI to solving their challenges across all platforms.

Authors

George Percivall, Vice-Chair of IEEE Standards Working Group on the Spatial Web
Capm Petersen, Chief Innovation Officer, VERSES Inc
Jacqueline Hynes, Research Engineer, Spatial Web Foundation
Axel Constant, Researcher and Technical Writer, VERSES Inc

References

[1] https://www.spatialwebfoundation.org/

[2] https://sagroups.ieee.org/2874/

[3]https://www.spatialwebfoundation.org/_files/ugd/a2c1f4_9957a3bdf3334eb0a3033bdee763760c.pdf

[4] https://standards.ieee.org/industry-connections/ec/autonomous-systems/

[5]https://standards.ieee.org/ieee/7000/6781/

[6]https://www.spatialwebfoundation.org/_files/ugd/a2c1f4_06ec41ae5aec4c959e1d99cca8e83b97.pdf

[7] Evan Patterson, Owen Lynch, and James Fairbanks. Categorical Data Structures For Technical Computing https://arxiv.org/abs/2106.04703v5

[8] Karl J. Friston et al. "Designing Ecosystems of Intelligence from First Principles" https://arxiv.org/abs/2212.01354