NATIONAL GEOSPATIAL-INTELLIGENCE AGENCY

Know the Earth... Show the Way... Understand the World

Trillion-Pixel GeoAI Challenges 2025

GeoAI Workshop Report

NGA NATIONAL GEOSPATIAL-INTELLIGENCE AGENCY | OAK RIDGE National Laboratory

Toward GEOINT Operational Efficiency:
Foundation Digital Twins and Geospatial AI Leads the Way

Trillion-Pixel GeoAI Challenges Workshop Report

THE NATIONAL GEOSPATIAL-INTELLIGENCE AGENCY
and
OAK RIDGE NATIONAL LABORATORY

September 2025

Prepared by
THE NATIONAL GEOSPATIAL-INTELLIGENCE AGENCY
and
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-BATTELLE LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

Contents

Acknowledgements 4
Session Speakers 4
Committee 4
Notetakers and Editors 4
Executive Summary 5
Abstract 5
Key Themes and Discussions 5
Key Challenges identified 5
Recommendations for Advancement 6
Session 1: Executive Roundtable – Geospatial AI Strategy 7
Session 2: Data Infrastructure for Foundational Digital Twin 8
Session 3: World Foundation Models for GeoAI Agents 9
Session 4: Scalable Computing Infrastructure for Enterprise HPC 10
Session 5: Trustworthiness in GeoAI Systems 12
Session 6: New Public-Private Partnerships for Accelerated Operational Impact 13
Conclusion 15

Acknowledgements

Session Speakers

Mr. Keith Allen, *National Geospatial-Intelligence Agency*
Dr. Swatantran Anu, *Corteva Agriscience*
Dr. Robert Ashcraft, *Natl Geospatial-Intelligence Agency*
Dr. Orhun Aydin, *Saint Louis University*
Mr. Greg Brunner, *Esri*
Mr. Bill Caniano, *National Geospatial-Intelligence Agency*
Mr. Robert Cardillo, *Taylor Geospatial Institute*
Dr. Carter Christopher, *Oak Ridge National Laboratory*
Dr. Philipe Dias, *Oak Ridge National Laboratory*
Mr. Gary Dunow, *US Geospatial Intelligence Foundation*
Dr. Rafael Ferreira Da Silva, *Oak Ridge Natl Laboratory*
Mr. David Gleason, *Amazon Web Services*
Mr. James Griffith, *Natl Geospatial-Intelligence Agency*
Dr. Nathan Jacobs, *Washington University*
Mr. Geoffrey King, *Greater Saint Louis Inc.*
Dr. Bill Kramer, *University of Illinois Urbana-Champaign*
Mr. Tom Lash, *Vibrant*
Dr. Zongyi Li, *MIT*
Ms. Jen Marcus, *Taylor Geospatial Engine*
Mr. Matt McClure, *NASA*
Mr. James McDaniel, *Natl Geospatial-Intelligence Agency*
Dr. Michael Mohoney, *UC Berkeley*
Dr. Albert Reuther, *MIT*
Mr. Rohan Sharma, *University of Missouri Saint Louis*
Mr. Howard Small, *Natl Geospatial-Intelligence Agency*
Dr. Kristin Sobolik, *University of Missouri Saint Louis*
Ms. Kari Stone, *National Geospatial-Intelligence Agency*
Ms. Ashley Suiter, *National Geospatial-Intelligence Agency*
Mr. Kyle Tucker-Davis, *Scale AI*
Dr. May Yuan, *University of Texas at Dallas*

Committee

Dr. Sonya Bahar, *University of Missouri Saint Louis*
Dr. Budhu Bhaduri, *Oak Ridge National Laboratory*
Dr. Sanjiv Bhatia, *University of Missouri Saint Louis*
Mr. Noah Butler, *Technology Entrepreneur Center T-REX*
Ms. Mike Fitzpatrick, *Natl Geospatial-Intelligence Agency*
Ms. Jaimeson Fredericks, *NGA*
Ms. Bree Futrell, *Oak Ridge National Laboratory*
Dr. Trevor Garner, *Natl Geospatial-Intelligence Agency*
Ms. Kathleen Klein, *Technology Entrepreneur Center T-REX*
Dr. Dalton Lunga, *National Geospatial-Intelligence Agency*
Ms. Marlie Mollett, *Natl Geospatial-Intelligence Agency*
Ms. Jamie Roberts, *Natl Geospatial-Intelligence Agency*
Ms. Christine Woodard, *NGA*
Dr. Lexie Yang, *Oak Ridge National Laboratory*

Notetakers and Editors

Dr. Robert Ashcraft, *Natl Geospatial-Intelligence Agency*
Ms. Kel Gruber, *University of Missouri Saint Louis*
Ms. Anushka Kalavagunta, *University of Missouri STL*
Mr. James McDaniel, *Natl Geospatial-Intelligence Agency*
Dr. Cory Reed, *National Geospatial-Intelligence Agency*
Mr. Rohan Sharma, *University of Missouri Saint Louis*

Executive Summary

Abstract

The 2025 Trillion-Pixel GeoAI Challenges Workshop convened experts across government, industry, and academia to address the critical transition of geospatial artificial intelligence from isolated pilot projects to robust, enterprise-scale operational systems. One focus of the workshop was the strategic convergence of foundational AI models and digital twins to enable dynamic, predictive representations of the physical world. Participants examined significant operational hurdles, highlighting the urgent need for engineered trustworthiness in AI, seamless data interoperability across legacy systems, and resilient infrastructure. To advance the GeoAI enterprise, the workshop emphasized fostering innovative public-private partnerships to build a unified geodata ecosystem, establishing objective standards for AI explainability and security, and adopting a hybrid computing strategy that optimally balances on-premises High-Performance Computing (HPC), commercial cloud platforms, and edge capabilities to support mission-critical operations.

Key Themes and Discussions

- Convergence of Digital Twins and Foundation Models: A primary theme was the strategic integration of foundational AI models with digital twins to create dynamic, predictive representations of the physical world. Discussions centered on moving beyond static data products toward integrated, model-level data that enables real-time updates and "what-if" simulations.
- The Critical Role of Public-Private Partnerships (PPPs): The workshop emphasized that accelerating GeoAI development and operational impact requires new, innovative partnership models. Sessions highlighted the need to break down cultural and bureaucratic barriers between government, industry, and academia to foster collaboration, share resources, and co-develop solutions.
- Shift to Enterprise-Scale and Edge Computing: Conversations revolved around building the next generation of scalable infrastructure required for GeoAI. This included a focus on hybrid approaches that leverage both specialized, on-premises High-Performance Computing (HPC) and the flexibility of commercial cloud platforms, while also pushing computational capabilities to the edge to improve efficiency and support tactical operations.

A central topic of the 2025 workshop was the strategic imperative to create a unified, enterprise-wide GeoAI ecosystem. Discussions moved beyond theoretical concepts to address the practical challenges of implementation, emphasizing the convergence of foundation models and digital twins to enable predictive analysis and dynamic global change detection. The sessions highlighted a significant shift from siloed, product-centric workflows toward a more integrated, model-based approach where data from diverse sources is fused into a single, authoritative representation. This vision is supported by a push for new public-private partnerships to overcome cultural divides and accelerate innovation, along with improving data interoperability and a strategic re-evaluation of computing infrastructure to balance the power of centralized HPC with the agility of edge computing.

Key Challenges identified

- Establishing Trust and Ensuring Explainability: A recurring challenge is the "black box" nature of complex AI models, which creates significant hurdles for user trust and adoption. The difficulty in verifying model reasoning, quantifying uncertainty, and defending against

adversarial attacks was identified as a major barrier to using GeoAI in high-stakes operational environments.

- Data Interoperability and Foundational Standards: The workshop identified the lack of universal standards as a critical bottleneck. Challenges include integrating data from disparate sources, establishing a common geodetic reference frame (especially for the vertical dimension), and moving legacy systems from product-centric data duplication to a single, authoritative virtual reference model.
- Bridging the Gap Between Pilot Projects and Production Systems:
  Panelists frequently noted the significant chasm between successful small-scale AI pilots and robust, enterprise-level production systems. This gap is created by the immense difficulty and cost of data curation, model monitoring, and building the resilient infrastructure required to operationalize GeoAI capabilities at scale.

The workshop identified several fundamental challenges that must be addressed to realize the full potential of GeoAI. A primary concern is establishing trustworthiness in AI systems; the inherent opacity of deep learning models makes it difficult for users to understand their outputs, quantify uncertainty, and defend against sophisticated adversarial attacks. Another significant hurdle is the persistent lack of data interoperability. Experts detailed the immense difficulty of aligning diverse geospatial datasets to a common standard and transitioning from inefficient, product-centric data systems to a unified digital twin architecture. Finally, a major theme was the struggle to move beyond successful but limited pilot projects to create robust, scalable, and operationally resilient enterprise-level GeoAI systems, a gap that stems from the immense underlying challenges in data preparation, infrastructure, and workforce development.

Recommendations for Advancement

- Engineer Trustworthiness into System Design: Instead of treating trust as a subjective user perception, it must be engineered into GeoAI systems from the ground up. This involves developing objective, measurable standards for validation, integrating explainability (XAI) techniques throughout the development lifecycle, and creating robust defenses against adversarial manipulation.
- Foster a Collaborative Geodata Ecosystem: To overcome data silos, the community should prioritize the creation of a "geodata cooperative." This involves leveraging existing standards bodies like the OGC, promoting the co-production of data on common standards (e.g., "Real 3D"), and incentivizing public-private partnerships to share data and infrastructure resources.
- Adopt a Hybrid and Balanced Computing Strategy: Organizations should move away from a "one-size-fits-all" infrastructure model. The recommended approach is to develop a pragmatic, hybrid strategy that balances investments between on-premises HPC for specialized tasks, pay-as-you-go commercial clouds for scalable model training, and a robust network of edge devices for efficient, low-latency operational support.

To advance the GeoAI enterprise, the workshop put forward several key recommendations. First and foremost is the need to proactively engineer trustworthiness into the core design of AI systems by establishing objective validation standards, prioritizing explainability, and building resilient security measures. Secondly, it is imperative to foster a collaborative geodata ecosystem by building on existing standards, promoting public-private data sharing agreements, and moving toward a unified virtual data model to break down legacy silos. Finally, the community must adopt a balanced and mission-focused computing strategy, creating a hybrid infrastructure that

strategically allocates resources between specialized on-premises systems, scalable cloud platforms, and efficient edge computing to meet the diverse demands of the modern geospatial environment.

Session 1: Executive Roundtable – Geospatial AI Strategy

Presentations:

The session began with presentations from four panelists, each offering a unique perspective on the future of geospatial AI. Jen Marcus highlighted the Taylor Geospatial Engine's focus on accelerating the transition of academic research to commercial capability through collaboration between researchers, industry practitioners, and other stakeholders. She emphasized the importance of piloting projects, such as their work on defining field boundaries, and conducting market assessments to guide long-term strategy. Kari Stone of the NGA outlined the agency's strategic priorities, which include advancing computer vision, promoting responsible AI through training and accreditation, and exploring the potential of Large Language Models (LLMs) and multimodal AI. She acknowledged that while these technologies offer great promise, challenges related to cost, trust, and global data coverage persist. Additionally, the session was framed by an opening emphasis on the growing strategic importance of AI and Digital Twins within NGA's priorities, reinforcing that these efforts are part of a broader institutional shift. Marcus also underscored that AI development requires significant infrastructure, including data centers, energy, cooling systems, and a large multidisciplinary workforce, highlighting that innovation depends on a full ecosystem, not just models.

Gary Dunow from the USGIF spoke on the importance of creating a trusted environment for AI to foster industry participation and maintain a competitive edge globally. He stressed the role of USGIF in building a collaborative ecosystem and the community's reliance on NGA for leadership and standards. Finally, Jim Griffith, also from NGA, detailed three key priorities for the development of foundational digital twins: establishing a comprehensive 3D representation of the globe, securing high-performance computing and storage, and developing a global change detection capability. He advocated for a geodata cooperative to ensure data accuracy and accessibility and expressed a preference for procuring AI capabilities from commercial vendors to optimize government resources. He further specified that the 3D global model would require approximately 2–3 meter resolution and noted that much of the necessary data already exists but must be better organized and structured. He also introduced the concept of an AI-driven "daily take" of the Earth to enable continuous global change detection, emphasizing that this capability should integrate multiple data sources beyond traditional remote sensing.

Question/Answer Session:

The discussion that followed the presentations delved into several critical areas concerning the implementation and strategy of geospatial AI and digital twins. A recurring theme was the importance of collaboration and breaking down organizational silos. Panelists agreed that fostering partnerships between government, industry, and academia is essential for innovation. Strategies discussed included implementing pilot projects to demonstrate value, engaging in cross-disciplinary collaboration, and creating forums for open dialogue to ensure all stakeholders, including end-users, are part of the conversation. The consensus was that initiatives should be built from the grassroots level to ensure they meet real-world needs. In addition, panelists emphasized the need for pragmatic investment strategies, noting that not all use cases require full 3D Digital Twins; "sometimes 2D is enough." They stressed avoiding overengineering solutions,

using high-performance computing selectively, and shifting away from inefficient hub-and-spoke architectures toward edge computing approaches to improve efficiency and reduce costs.
Another key topic was the role of AI in developing and maintaining standards. Panelists acknowledged the difficulty in enforcing standards but emphasized their importance for interoperability and data consistency. The discussion highlighted the need for leadership in setting standards and the potential for AI to help modernize data and bridge gaps between old and new data formats. The conversation also touched upon investment strategies, with panelists advocating for a more efficient use of high-performance computing resources and an edge computing approach to improve efficiency. The importance of understanding user needs to avoid creating overly complex or expensive solutions was also stressed, with the general agreement that organizations should "buy what they can buy" and focus on developing a skilled workforce to manage and integrate these purchased models.
The panelists also addressed the distinction between Digital Earth and Digital Twins, explaining that the latter is typically used for more localized and detailed applications. They discussed the challenges of real-time updates in digital twin systems, noting that while global real-time updates are not currently feasible, they are critical for regional applications, especially in disaster response scenarios. The discussion concluded with a look at the future, with panelists expressing optimism about the potential of emerging technologies like quantum computing and the innovative capabilities of the next generation of data scientists and analysts.

Session 2: Data Infrastructure for Foundational Digital Twin

Presentations:

The session featured three main presentations that outlined the conceptual framework and practical challenges of building a foundational digital twin. Dr. Carter Christopher of Oak Ridge National Laboratory began by contrasting traditional, linear geospatial workflows with the paradigm of a digital twin. He introduced the concepts of "digital threads," which connect the physical world to a virtual model, and the "twinning rate," the frequency at which the virtual model is updated to reflect real-world changes. He argued that the focus must shift from creating siloed data products to integrating various data sources at the model level, using high-performance computing and container orchestration to create a dynamic, holistic representation from which tailored products can be "peeled off" as needed.
Following this, Howard Small, a prominent geodesist from NGA, detailed the critical importance of a common geodetic reference frame. He explained that for a digital twin to be accurate, all data must be anchored to a consistent surface, such as the WGS 84 ellipsoid. He highlighted the complexity of this, noting that different models and datums exist and that establishing a universal vertical reference for elevation (the "Z" value) remains a significant global challenge. The key takeaway was that data interoperability and accuracy depend on meticulous management of these foundational reference systems.
Lastly, Keith Allen from NGA provided insight into the agency's modernization efforts. He described the massive undertaking of creating a high-resolution global elevation model, which has already produced over 21.7 trillion pixels of data. He announced a partnership with the Army on the "One World Terrain" program and the development of a new product specification, "Real 3D," based on Cesium 3D tiles. This initiative aims to create synergy between government and industry partners, allowing them to co-produce data on a common standard and move away from the static, incompatible datasets of the past.

Question/Answer Session:

The panel discussion delved deeper into the practicalities and implications of creating a foundational digital twin. A central theme was the importance of interoperability, which the panelists agreed was essential for effective operations, from military decision-making to disaster response. The ability to rapidly share and align data from different sources without lengthy manual conversions could shorten decision cycles from days to hours or even minutes. This requires not only technological solutions but also adherence to common standards, such as WGS 84, to ensure all data aligns correctly. The panel stressed leveraging existing standards bodies like the Open Geospatial Consortium (OGC) rather than creating new ones.

The discussion also addressed the challenge of moving from legacy, product-centric data systems to a model where a single feature is stored once but can be visualized in multiple ways. The panelists acknowledged that authoritative databases currently contain duplicate entries for different products (e.g., a lighthouse represented differently on maritime and topographic charts). The vision for the digital twin is to have a single virtual reference for that object, with various threads and attributes that support different user needs. This leads to questions of customer adoption; currently, most users are not asking for a digital twin, so a key task is educating them on the benefits of receiving dynamic, on-demand data versus traditional static map sheets. Success for the digital twin initiative will not be measured by creating a photorealistic virtual globe, but by its operational impact. Key metrics include reducing mission timelines, the rate of technology adoption by users, and the ability to meet a wider range of mission requirements more efficiently. In the longer term (5+ years), the ultimate promise of the digital twin lies in enabling advanced simulations and predictive "what-if" analysis.

Finally, the panel touched upon the critical issue of data provenance and security. In an era of adversarial attacks, ensuring the integrity and traceability of data from its source to its application is paramount. This is especially true when using "black box" AI models. The strategy is to maintain control over the entire data processing pipeline in a modular architecture, ensuring a trusted government agent can vouch for the data's quality and accuracy, rather than simply relying on an end product from an external source.

Session 3: World Foundation Models for GeoAI Agents

Presentations:

The third session shifted the focus to the next generation of AI, exploring the concept of "World Foundation Models" for GeoAI agents. The presentations centered on the idea of creating AI systems that don't just passively analyze data but can actively understand, reason about, and interact with the physical world. Nathan Jacobs from Washington University detailed his research on using diverse open data sources to understand geospatial dynamics, such as combining social media data with imagery to map transient attributes like "lushness" or "gloominess," synthesizing new satellite imagery from text prompts, and developing interactive GeoAI agents that learn from human search patterns.

Zongyi Li from MIT presented a new architectural design for GeoAI, focusing on creating a general and flexible model that can be applied to different problems. He discussed using "neuro orbiters" and embedding physics into the models to achieve domain consistency across various scales and to handle complex geometries, from coastlines to city streets. Coming from a military background, Kyle Tucker-Davis of Scale AI offered a practical perspective, arguing that the focus of AI development is often misplaced on the model itself. He stressed that the real challenge lies in

the data "foundry," as GeoAI inherits the hardest problems of both Big Data (volume, velocity, variety) and AI (data preparation, curation, and monitoring), creating a significant gap between successful pilots and robust production systems.
Anu Swatantran from Corteva Agriscience grounded the discussion in the philosophy of geography, noting that while we can monitor the "what" with thousands of satellites, the data remains in silos, preventing us from understanding the "how" and "why." She echoed the challenge of bridging the chasm between pilot projects and production-scale systems. Finally, Michael Mohoney from UC Berkeley provided a counterpoint from his work on physical systems, emphasizing that a "foundation model" is a base upon which to build. He explained that while building the initial foundation model requires immense resources, adapting it for a specific use case is far less intensive, presenting a significant opportunity for the community to leverage these powerful, large-scale models across a wide range of scientific and geospatial challenges.

Question/Answer Session:

The discussion following the presentations was a deep dive into the ambitious and complex challenge of creating true World Foundation Models. A major theme was the fundamental difficulty of integrating the laws of physics and real-world dynamics into AI systems. Panelists discussed that current machine learning models are excellent at pattern recognition but struggle with causality and structured reasoning. The consensus was that a hybrid approach is needed, one that combines deep learning with symbolic AI and physics-based simulations to create models that can generalize beyond their training data and understand the underlying principles governing the environment.
Another key area of discussion was the role of simulation and virtual environments. The panelists agreed that creating high-fidelity virtual replicas of the world would be essential for training and testing GeoAI agents safely and effectively. These "digital sandboxes" would allow developers to expose AI agents to a vast range of scenarios, including rare and dangerous events, without real-world consequences. This would be critical for validating agent behavior and ensuring they act reliably and predictably before being given any level of autonomy in operational environments. The conversation highlighted the need for these simulations to be not just visually accurate, but also physically realistic.
Finally, the panel grappled with the concept of "agency" itself and the trustworthiness of these future AI systems. The discussion touched upon the significant ethical and safety implications of creating AI that can interact with and modify the world. The panelists emphasized the need for robust explainability (XAI) and verification techniques to ensure that human operators can understand and trust the reasoning behind an agent's decisions and actions. The conversation concluded that while the vision of World Foundation Models is powerful, the path to achieving it requires solving fundamental research challenges in AI, simulation, and trustworthiness to ensure these powerful new tools are used responsibly.

Session 4: Scalable Computing Infrastructure for Enterprise HPC

Presentations:

Session four focused on the critical challenge of building a scalable computing infrastructure capable of supporting the immense demands of GeoAI and foundational digital twins at an enterprise scale. Albert Reuther from the MIT Supercomputing Center opened by describing his team's approach: deeply understanding the underlying mathematics of a problem to build the most efficient hardware architecture to solve it. He highlighted their latest deployment, a massive

system with 636 GPUs, emphasizing the need to balance network and storage to create an "economically efficient ecopod" that serves diverse user needs.
Rafael Ferreira Da Silva from the ORNL Leadership Computing Facility expanded the view to distributed AI-driven workflows across multiple facilities. His R&D focus is on autonomous science networks that utilize not just High-Performance Computing (HPC), but also cloud, accelerators, and emerging quantum technologies. He outlined the grand challenges facing the community, including managing massive data growth, orchestrating complex workflows across distributed systems, and the critical need to enable trust and robustness in Foundation Models by keeping human oversight "in the loop." Bill Kramer, a Research Professor at the University of Illinois and director of the New Frontiers Initiative, brought decades of experience to the panel. He discussed the deep synergy between data and computing, drawing on his history of analyzing massive datasets from fusion reactors, the Human Genome Project, and managing the Blue Waters supercomputer, which was instrumental in completing the Earth Digital Elevation Model (DEM).
Matt McClure presented NASA's approach to these challenges, given their vast array of missions and partners across scientific disciplines. He explained that NASA is building specific science foundation models for each division to bridge the gap between scientific discovery and operational use by astronauts, noting that all of this is fundamentally powered by compute. Finally, David Gleason from AWS provided a commercial cloud perspective, observing that while customers have different end goals, they share common underlying problems when trying to operationalize Generative AI. He traced the natural evolution of customer solutions: starting with basic models, encountering performance and flexibility challenges, moving to Retrieval-Augmented Generation (RAG) to index unstructured text for specific contexts, and ultimately evolving toward the use of autonomous AI agents.

Question/Answer Session:

The discussion following the presentations provided a deeper dive into the strategic trade-offs of designing and funding the future of geospatial computing. An interesting component was the "convergence vs. specialization" debate. The panelists generally agreed that while general-purpose cloud platforms offer incredible scale and flexibility, the unique demands of GeoAI, like processing massive, multi-modal sensor data and the need for energy-efficient processing at the edge, suggest that specialized hardware and software will continue to play a crucial role. The consensus was that a "one-size-fits-all" solution is unlikely to be optimal; instead, a hybrid approach that leverages the best of both worlds will be necessary.
Another significant area of discussion was the practical challenge of balancing investments. The panelists explored how an organization should allocate its budget between maintaining its own on-premises HPC capabilities, leveraging the pay-as-you-go model of commercial clouds, and building out a network of edge devices. The conversation highlighted that the ideal balance depends heavily on the specific mission requirements, data security constraints, and latency needs. For large-scale model training and data archiving, the cloud is often the most cost-effective solution. However, for real-time operational needs in tactical environments, on-premises and edge computing remain indispensable.
Finally, the panel touched upon the readiness of these systems for enterprise-wide adoption. They discussed the feasibility of having a single, unified platform that can handle both the immense task of training foundational models and the more agile requirements of running production inference across an organization. The panelists concluded that while this is a desirable goal, the current reality is that these two functions often require different architectural setups. The discussion underscored the need for a well-defined infrastructure strategy and a clear-eyed assessment of the

costs, risks, and performance trade-offs inherent in building a scalable, resilient, and secure computing backbone for the future of GEOINT.

Session 5: Trustworthiness in GeoAI Systems

Presentations:

Session five tackled the critical issue of trustworthiness in GeoAI systems, defining it as the degree to which users, stakeholders and decision-makers can fundamentally rely on a system's outputs and behavior. The presentations established that this is especially challenging in the geospatial domain due to inherent issues like spatial uncertainty, data gaps, and potential biases in collection. The speakers broke down the problem of building trust into three core areas: first, the challenge of defining and measuring trustworthiness itself through clear, objective standards; second, the necessity of ensuring explainability (XAI) so that users can understand the reasoning behind a system's conclusions; and third, the critical need to defend against adversarial threats that could manipulate a system's output.

This critical issue was tackled with perspectives that spanned foundational research, academic theory, corporate strategy, and government policy. Phillipe Dias, a researcher with the GeoAI group at ORNL, opened by discussing the dual nature of generative AI and foundational models by highlighting their massive improvements alongside new trustworthiness challenges. He detailed ORNL's efforts to leverage these models for synthesizing high-resolution satellite imagery and forecasting future spatial patterns using diffusion models and historical datasets. However, Dias warned that the billions of parameters in these models create massive interpretability hurdles. He noted that profit-driven entities often prioritize speed and cherry-picked results over transparency, making neutral evaluation, uncertainty quantification, and deep fake detection essential. Building on the theme of model opacity, Orhun Aydin from Saint Louis University focused on uncertainty-aware GeoAI algorithms for digital twins. He explored the "cone of uncertainty," noting that while we understand the mechanics of smaller models, larger models increasingly become opaque black boxes. Citing the historical example of the ELIZA chatbot, Aydin cautioned against the unintended consequences of AI, particularly the danger of unintentionally removing human experts from the decision-making loop.

The discussion then shifted toward defining and operationalizing trust. Dr. May Yuan from the University of Texas at Dallas challenged the group with a foundational question: Can GeoAI truly be trustworthy? She argued that trustworthiness must be engineered into the model itself, rather than relying on a user's subjective perception, as humans psychologically build trust simply by being happy with an output. Dr. Yuan highlighted the dangers of geographic variance in input data, the lack of generalizability in over-tuned GeoAI models, and the critical need to address biases in LLMs before they are deployed as autonomous interfaces. Providing an industry perspective, Greg Brunner from Esri explained how AI is infusing all aspects of GIS across GeoAI, AI assistants, and AI agents. He emphasized that trust is a foundational requirement built on principles like security, privacy, and fairness. To foster this, Esri utilizes "transparency cards" that provide users with detailed statistical reports on their models. However, he acknowledged that trust is a two-way street, and users ultimately must evaluate the tools hands-on to determine trustworthiness for their specific operational needs.

Finally, Ashley Suiter, the acting Chief of Responsible AI at the NGA, detailed how the agency manages the deluge of both data and new AI capabilities. She outlined NGA's comprehensive four-pillar approach to responsible AI: People (users and developers), Programs (alignment), Policy (governance), and Platforms (scalable standards). Suiter stressed that trustworthiness validation

cannot be a separate step; it must be deeply integrated into existing cyclical requirements and risk management processes. Highlighting that the NGA manages AI risk at the specific use-case level, she underscored that "People" represent the most critical pillar. To ensure human responsibility and oversight at every step, from the developer to the end-user, she noted that the NGA has implemented rigorous accreditation and GEOINT responsible AI training programs.
Together, the panelists presented a framework for addressing these challenges, emphasizing that trust is not a feature to be added at the end but must be a foundational component of the entire system design. The presentations set the stage for a deeper discussion on how to create systems that are not only accurate but also reliable, secure, and transparent enough to be confidently used in high-stakes national security environments.

Question/Answer Session:

The discussion following the presentations delved into the practical and philosophical difficulties of establishing trust in AI. A major theme was the challenge of creating measurable, objective standards for trustworthiness. Panelists agreed that a single metric for "trust" is impossible, as the required level of reliability and explainability is highly dependent on the specific operational context. For example, the trust requirements for an AI system recommending a target in a real-time military operation are far more stringent than for one analyzing long-term climate trends. The conversation highlighted the need for a nuanced approach, developing specific standards and validation techniques for different types of GeoAI applications.
Another key area of debate was the inherent tension between model performance and explainability. The panelists discussed the "black box" problem, where the most accurate and powerful AI models are often the most difficult to interpret. They explored various XAI techniques designed to peer inside these models but acknowledged that a perfect solution does not yet exist. The consensus was that organizations must carefully consider this trade-off, sometimes opting for a slightly less performant but more interpretable model to ensure that human operators can understand, verify, and ultimately trust the AI's recommendations, especially when lives are on the line.
Finally, the discussion addressed the very real and growing threat of adversarial attacks. Panelists provided examples of how easily GeoAI systems, particularly those based on computer vision, can be deceived by carefully crafted, often imperceptible, changes to input data. The conversation stressed the need for a proactive security posture, moving beyond traditional cybersecurity to include robust testing, validation, and continuous monitoring to detect and defend against these novel threats. The overarching conclusion was that building trustworthy GeoAI requires a holistic approach that integrates technical solutions for security and explainability with a deep understanding of the human users and the high-stakes decisions they are required to make.

Session 6: New Public-Private Partnerships for Accelerated Operational Impact

Presentations:

The sixth session provided a comprehensive overview of how innovative Public-Private Partnership (PPP) models can be leveraged to accelerate the development and operational impact of GeoAI technologies, particularly within the St. Louis region. Tom Lash, CEO of Vibrant, drew upon his long, extensive background in geospatial intelligence to help frame the discussion around the intersection of federal technology needs and commercial industry capabilities. Geoffrey King from Greater St. Louis Inc. expanded on this regional focus by discussing his role in managing the

"GeoFutures" initiative. He highlighted the organization's strategic aim to intentionally strengthen and deepen the local geospatial economy by fostering a collaborative, multi-sector ecosystem. The panel also featured strong perspectives from the government-turned-industry and academic sectors. Robert Cardillo, representing Taylor Geospatial and bringing his deep experience as the former Director of the National Geospatial-Intelligence Agency (NGA), offered a forward-looking vision for the community. He warned that the industry's future could be severely inhibited if it remains overly focused on past methods and traditional procurement models. Instead, he stressed that the future of the geospatial enterprise will be fundamentally defined by its ability to build and leverage dynamic external relationships. Finally, Dr. Kristin Sobolik, Chancellor at the University of Missouri–St. Louis (UMSL), provided the crucial workforce development perspective. Recognizing that UMSL serves as the primary workforce driver for the local economy, she emphasized that it is imperative for the university to actively disrupt the traditional model of higher education. By doing so, UMSL ensures it continues to be a vital, responsive asset for the St. Louis community and the growing geospatial ecosystem.

The panelists presented a portfolio of modern PPP strategies designed to harness the respective strengths of government, industry, and universities. These models included research consortia for pre-competitive R&D, shared infrastructure partnerships to democratize access to high-performance computing and large datasets, and innovation sandboxes that provide controlled environments for experimentation in sensitive sectors like defense. A key example of this approach was the St. Louis regional GeoAI ecosystem, anchored by NGA, which supports over 7,500 geospatial jobs and integrates sectors such as agriculture, life sciences, finance, and advanced manufacturing. Panelists emphasized that this regional clustering model, combined with projected GeoAI market growth from $25B to $36B, demonstrates how PPPs can drive both national security outcomes and regional economic competitiveness. Panelists further emphasized that trust in these partnerships depends on shared data, clear rules, and consistent practices, highlighting NGA's evolving approach of integrating classified and unclassified environments, such as the Next NGA West campus design with approximately 20% unclassified and 80% classified space, as a practical model for enabling collaboration while maintaining security.

Question/Answer Session:

The discussion following the presentations delved into the practical challenges and success factors for implementing these new Public-Private Partnership models. A major theme was the inherent "culture clash" between the fast-paced, profit-driven private sector and the more bureaucratic, risk-averse culture of government. Panelists agreed that overcoming this is a primary hurdle. A key strategy discussed was the importance of building personal relationships and establishing clear, shared goals from the outset. The consensus was that successful partnerships require dedicated "champions" on both sides who can navigate their respective organizational structures and maintain momentum.

Another significant area of conversation was the management of intellectual property (IP). The panelists acknowledged that IP rights are often a major sticking point, with industry needing to protect its commercial interests while the government requires sufficient rights to use the technology for its missions. The discussion highlighted the need for flexible and creative agreements that move beyond traditional government contracts. Models like pre-competitive R&D, where IP is shared among consortium members, and innovation sandboxes, where IP rights can be negotiated on a case-by-case basis, were presented as effective ways to balance these competing interests and prevent IP concerns from derailing collaboration.

Finally, the panel explored how to scale these partnerships from regional successes to a national ecosystem. The conversation emphasized the importance of creating replicable, standardized frameworks for PPPs that can be adapted to different needs and contexts. The panelists also stressed the role of intermediary organizations, such as university-affiliated research centers and non-profits, which can act as neutral third parties to facilitate communication, manage administrative overhead, and build trust between government and industry partners. The discussion concluded with an optimistic outlook, suggesting that while challenges remain, these new partnership models are essential for ensuring the nation remains at the cutting edge of geospatial and AI innovation.

Conclusion

The 2025 Trillion-Pixel GeoAI Challenges workshop underscored the critical need to transition geospatial artificial intelligence from isolated pilot projects to robust, enterprise-scale operational systems. The presentations and discussions highlighted the strategic convergence of foundation models and digital twins to build predictive, dynamic representations of the physical world for GEOINT. Key challenges identified included establishing genuine trustworthiness in opaque AI systems, overcoming the immense difficulty of ensuring data interoperability across legacy platforms, and building the scalable, hybrid computing infrastructure required for enterprise-wide adoption. The workshop emphasized the importance of fostering a unified GeoAI ecosystem to support these advancements, particularly within the context of the National Geospatial-Intelligence Agency's strategic priorities.

The workshop also explored the broader implications of operational GeoAI in national security, disaster response, and economic competitiveness. By leveraging foundation models and digital twins, the GEOINT community can enhance traditional analysis, shorten decision cycles, and unlock predictive capabilities. However, issues related to the significant gap between pilot projects and production systems, the complexities of managing public-private partnerships, and the persistent "black box" nature of advanced AI models remain pressing concerns. These must be addressed through continued research and deliberate collaboration between government agencies, industry partners, and academic institutions.

The workshop made several recommendations for making progress on these challenges, including:

- Focus on pragmatic investment by prioritizing the procurement of proven commercial AI capabilities and developing a skilled workforce to manage and integrate them, rather than over-investing in building all solutions from the ground up.
- Prioritize viewing and implementing data interoperability, based on common standards, as a foundational requirement for the entire enterprise.
- Invest in creating high-fidelity, physically realistic virtual "sandbox" environments to safely train and validate GeoAI agents, allowing for testing against a wide range of scenarios before they are given autonomy in operational settings.
- Adopt pragmatic, hybrid computing strategies that balance the use of on-premises, cloud, and edge resources to meet specific mission needs.
- Developing objective, measurable standards to engineer trustworthiness into system design from the beginning of the development lifecycle, requiring objective validation, integrated explainability (XAI) techniques, and robust security against adversarial attacks, rather than treating trust as a feature to be added later.

- Establish a standardized, replicable framework for Public-Private Partnerships that includes flexible agreements and utilizes neutral intermediary organizations to help overcome the cultural and bureaucratic gaps between government and industry.

Finally, the Trillion-Pixel GeoAI Challenges workshop emphasized the need for a cohesive, community-wide effort to accelerate innovation and achieve operational impact. Participants called for a more integrated geodata ecosystem, championing new partnership models and a collective commitment to overcoming shared technical and cultural hurdles. Moving forward, the fusion of foundation models and digital twins holds immense potential to revolutionize the GEOINT enterprise, provided that the community works collaboratively to address the technical, ethical, and infrastructural challenges required to build a trusted, scalable, and truly operational GeoAI future.