OAK RIDGE NATIONAL LABORATORY

TRILLION-PIXEL

GeoAI Challenges

- Dalton Lunga
- Rahul Ramachandran
- Shawn Newsam
- Jitendra Kumar
- Eric Shook
- Dawn King
- Lexie Yang
- David Page
- Katie Schuman
- Yan Liu
- Manil Maskey
- Philipe Dias
- Carter Christopher
- Budhu Bhaduri

2023 GeoAI

WORKSHOP REPORT

THE WAY FORWARD

in Scaling Multimodal GeoAI

geoai.ornl.gov/trillion-pixel

Oak Ridge National Laboratory is managed by UT-Battelle LLC for the US Department of Energy

OAK RIDGE National Laboratory

ORNL/LTR-2023/2326

Geospatial Science and Human Security Division

Trillion-Pixel GeoAI Challenges Workshop Report

October 2023

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831-6283
managed by
UT-Battelle LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

1. INTRODUCTION 6
2. SESSION 1: AGENCY GRAND USER CHALLENGES AND USE CASES 8
3. SESSION 2: AI FOUNDATION MODELS 8
4. SESSION 3: GEOSPATIAL DATA INFRASTRUCTURE 10
5. SESSION 4: EDGE COMPUTING 11
6. SESSION 5: HARDWARE AND SOFTWARE ARCHITECTURES 12
7. SESSION 6: WORKFORCE REQUIREMENTS AND PARTNERSHIPS . . . .

ACKNOWLEDGMENTS

Session Speakers

Mr. Daniel M. Cotter, *Senior Advisor, US Department of Homeland Security*.
Mr. Mark Munsell, *Director, Data and Digital Innovation Directorate, National Geospatial-Intelligence Agency*
Mr. Kevin Murphy, *Chief Science Data Officer, National Aeronautics and Space Administration*
Dr. David Alexander, *Senior Science Advisor for Resilience, U.S. Department of Homeland Security*
Dr. Raghu Ganti, *Principal Research Scientist, IBM*
Dr. Philipe Dias, *Research Scientist in Computer Vision and Machine Learning, Oak Ridge National Laboratory*
Dr. Kaleb Smith, *Senior Data Scientist, NVIDIA AI Technology Center*
Mr. Muthukumaran Ramasubramanian, *Computer Scientist, National Aeronautics and Space Administration*
Dr. Samantha Arundel, *Acting Research Director, Center of Excellence for Geospatial Information Science, U.S.G.S.*
Dr. Rasmus Houborg, *Principal Geospatial Fusion Scientist, Planet Labs*
Dr. Brian Freitag, *Research Scientist, National Aeronautics and Space Administration*
Mr. John Wegrzyn, *Senior Research and Development Engineer, UMBRA*
Professor Katie Schuman, *University of Tennessee, Knoxville*
Dr. Angel Yanguas-Gil, *Principal Materials Scientist, Argonne National Laboratory*
Dr. Erica Montbach, *Manager, NASA Planetary Exploration Science*
Dr. Hamed Alemohammad, *Director, Center for Geospatial Analytics, Clark University*
Dr. Prasanna Balaprakash, *Director of AI Programs, Oak Ridge National Laboratory*
Dr. Erwin Gilmore, *Senior AI/ML Specialist, Amazon Web Services*
Professor Vijay Gadepally, *Massachusetts Institute of Technology*
Dr. Mo Sarwat, *CEO, Wherobots Inc.*
Dr. Forrest Hoffman, *Distinguished Computational Earth System Scientist , Oak Ridge National Laboratory*
Dr. Dawn King, *Senior AI/ML Geophysical Applied Scientist, National Geospatial-Intelligence Agency*
Professor Orhun Aydin, *St. Louis University*
Dr. Subit Chakrabarti, *Vice President of Technology, Floodbase*
Dr. Matej Batic, *Research Team Lead, Sinergise*

Oak Ridge National Laboratory Event Support Team

Mr. Eric Swanson	Ms. Audrey Carson
Mr. Jordan Bowman	Mr. Jason Wohlgemuth
Mr. Jacob Arndt	Ms. Dara Cunningham
Ms. Elisabeth Johnson	Ms. Suzanne Booher
Dr. Abhishek Potnis	Dr. Joe Tuccillo
Ms. Heidi Tubbs	Dr. Karly Harrod
Mr. Gene Bailey	

EXECUTIVE SUMMARY

The emergence of virtually ubiquitous geospatial data, astonishing progress in artificial intelligence, transformational advances in cloud infrastructure, and high-performance computing are converging to create unprecedented detail in mapping and interpreting the Earth's surface. Rapid innovations in satellite and airborne remote sensing capabilities enable data collection upwards of 100 terabytes or more per day. Despite advances from the artificial intelligence, computer vision, and earth observation communities, we have yet to fully unlock the tremendous opportunities this data offers to provide insights into how humans occupy, interact, and alter the planet's surface over time. This opportunity calls for programmatic investments in and designs of new end-to-end transformational technologies that can scale and reduce the latency from data generation to decision making.

Since 2019[*], a new community at the intersection of geospatial sciences, national security, AI, computer vision, and computing, has been biennially gathering to discuss the frontiers of GeoAI systems. Framed as a modern-day moonshot for geospatial science and technology advancement, the Trillion-Pixel GeoAI Challenge was introduced in 2019. The challenge brings forth key priority focus themes that require broad community engagements. National security and open science domain problems, AI and computing, workforce development and partnerships are among some of the key topics that remain relevant to addressing the emerging GeoAI challenges. In 2021[†], the workshop discussions extended on the outcomes of the 1st series and presented "The Frontiers of Trillion-Pixel GeoAI End-to-End Systems."

We themed our most recent gathering "Multimodal Trillion-Pixels: GeoAI Challenges" and welcomed nearly 100 experts from government, image science, computer vision, high-performance computing, artificial intelligence, advanced workflows, and end-user communities in June 2023. Discussions focused on agency grand challenges, advances in AI foundations, geospatial data infrastructure, edge computing, hardware and software architectures, and workforce requirements and partnerships. Presentations covered the above six themes over a two-day event which featured a diverse group of 32 speakers who each highlighted the current state of practice as well as the modern challenges and opportunities to solve the emerging multimodal trillion-pixels GeoAI challenges.

Event participants represented various sectors, including US federal agencies, private industry, academia, international agencies, and the US Department of Energy national laboratories (Table). The event presented several takeaways on the way forward to addressing Trillion-Pixel Challenges for Multimodal GeoAI:

1. **Democratizing open access and documentation of AI systems**: DOE's labs presents unique research strengths and investments in leadership computing facilities that can be leveraged to lower the cost of developing large models for the scientific community. As Federally Funded Research and Development Centers, national laboratories do provide computing resources to communities but should also consider providing testbeds for evaluating and bench-marking GeoAI tools.

2. **Establishing leadership in GeoAI for National Security and Open Science**: Government, in partnership with academia and geospatial private sectors, should harness recent AI advances by

[*] Lunga, Dalton D., Alemohammad, Hamed, Liu, Yan, Newsam, Shawn, Pacifici, Fabio, Santos-Villalobos, Hector, Shook, Eric, Stewart, Robert N., Voisin, Sophie, Yang, Lexie, and Bhaduri, Budhu L. The Trillion Pixel GeoAI Challenge Workshop. United States: N. p., 2019. Web. doi:10.2172/1606744.

[†] Yang, Lexie, Lunga, Dalton, Lieberman, Jordan, Doster, Timothy, Kerner, Hannah, Casterline, May, Shook, Eric, Begoli, Edmon, Ramachandran, Rahul, Kumar, Jitendra, Pacifici, Fabio, Newsam, Shawn, Ward, Steven, and Bhaduri, Budhu. 2021 GeoAI Workshop Report: The Trillion Pixel Challenge. United States: N. p., 2021. Web. doi:10.2172/1883938.

Sector	No. participants
Academia	6
Federal Agency	23
Industry	27
International Agency	2
National Laboratory	36

investing in national strategic initiatives, including interoperable and standardized emergency management ecosystems and developing foundation AI models for science to interact with the environment in responsible, repeatable, and ethical manners.

3. **Focused codesign of spatiotemporal AI tools** : Much recent AI progress has been in industry applications that require little geospatial awareness. The industry needs focused, sustained research efforts to develop spatially-aware multimodal foundation models, along with new domain-specific hardware architecture designs.

1. INTRODUCTION

Given decades of evolution in earth observation sensing instruments and computing hardware, AI advancement and its broader impact on society, the national security and open science missions are facing a transformational moment. The design and optimization of geospatial artificial intelligence system workflows, which can quickly map and interpret the Earth's surface, require foundational research, expertise and experience that are currently lacking in the national security and open science communities. Unlike other scientific domains that have benefited from world-class resources in hardware and software architectures, these communities need to adopt new strategies for human capital development, partnerships, and technology infrastructure advancement to address the current and future grand challenges and use cases. We give context to six key themes where seizing opportunities and addressing new challenges will be critical to transforming multimodal trillion-pixel GeoAI challenges for societal benefit.

Grand Challenges and Use Cases: The astounding rise of AI and its impact on society has been decades in the making. AI technology is now at the frontline of diagnosing diseases, translating languages, and transcribing speech. Yet in the context of geographic knowledge discovery, a growing number of emerging national security and open science challenges are revealing glaring weaknesses in current GeoAI systems. In this third iteration of the conference, we sought to understand both existing and emerging grand geospatial use cases and the limiting factors to advancing the wellbeing and security of our society.

AI Foundations: AI-enabled analytic tools promise to synthesize available Earth observation data more effectively and provide unexpected, impactful solutions relevant to various national security and climate science challenges, such as energy grid resilience, nuclear nonproliferation, and risk-based predictive systems for complex socio-economic and physical environments. But currently, the development of application-specific GeoAI models is resource-intensive and costly in gathering large amounts of high-quality data, relearning common representations, and underutilizing other unlabeled data modalities. This puts the community at a significant disadvantage in the pursue of cutting-edge, AI-enabled geospatial workflows that support rapid and high-quality decision making for national security and open science challenges. Forward-looking discussions are key to the identification and production of relevant foundation

GeoAI models for decision support. In this 2023 workshop, the charge was to probe the emerging trends in large AI models and ask what building blocks are pertinent to the challenges emerging from natural and man-made events.

Geospatial Data Infrastructure: Massive geospatial datasets collected from surveys, sensing, and social media provide rich data sources and geospatial context for GeoAI. However, the GeoAI community still faces enormous challenges in integrating large, spatially heterogeneous, multi-modal geospatial datasets into existing machine learning frameworks. This thematic session aimed to address major geospatial data challenges and opportunities related to analysis-ready data (ARD) and interoperability of multimodal data products. These include difficulties in labeling massive geospatial datasets, the development of geospatial data schemas for machine learning models, data curation and validation, the potential to leverage existing rich data sources for GeoAI, HPC-based data-intensive AI computations, and future needs for geospatial data infrastructure solutions that couple geospatial processes, data analytics, and AI as a scalable platform for empowering large-scale applications.

Edge Computing: Geospatial edge devices, such as satellites, drones, and other size, weight, and power (SWaP)-limited devices, are crucial in addressing trillion-pixel challenges. These edge devices are the front line of the emerging global geospatial enterprise that collects more than 100 trillion pixels over the surface of the Earth daily. The current trend of connecting these edge systems to large, federated "cloud" computing demands ever-increasing bandwidth to handle the volume and velocity of raw data where a single satellite generates more than 100 terabytes per day. The primary mode of connection is through wireless networks to push this volume of data from the edge to centralized processing. Within the scope of this theme, experts proceeded to discuss the state of edge computing, identified enabling technologies and grand challenges, and defined future research directions to make edge computing a disrupting agent for GeoAI-based decision-making.

Hardware and Software Architectures: Application of GeoAI to growing volumes of geospatial information will require efficient, high-performance computing (HPC) hardware, network infrastructure and scalable software frameworks to transfer, store and process massive amounts of data. Planetary-scale GeoAI architectures capable of learning and inferring from trillions of pixels streaming daily will require tight integration of network, storage, and computing resources that span CPUs, GPUs, accelerators, and potentially domain-specific architectures such as a GeoAI spatial processor. Also needed are scalable open-source software ecosystems to analyze these datasets leveraging state-of-the-art computational resources. The conference assembled a panel of experts to explore the computational challenges facing GeoAI and discussed potential and holistic architectural solutions to overcome these challenges.

Workforce Requirements and Partnerships: Forming effective partnerships remains fundamental to addressing the challenges and opportunities arising from the multimodal GeoAI frontier. Researchers and practitioners with expertise in different GeoAI data modalities must come together to make progress on societal-scale challenges. Enabling AI model-based solutions to solve GeoAI grand challenges will also require even more interdisciplinary collaboration due to the scale of resources and complexities of policy-making needed to sustain the ecosystem. What types of talent do we need to bring together now to build such an ecosystem? What types of talent do we need to train to sustain and extend it? What are the key elements needed to foster a workplace or collaborative environment that will support generations to carry on the mission and solve the grand challenges? These are some of the key questions that were addressed during this session.

We summarize the discussions and presentations addressing the key questions from each session.

2. SESSION 1: AGENCY GRAND USER CHALLENGES AND USE CASES

This opening session presented several use cases where the next generation of GeoAI technologies could significantly impact the national security and open science mission space. Three leading authorities from the government Mark Munsel from the National Geospatial Intelligence Agency (NGIA), Daniel Cotter from the Department of Homeland Security (DHS), and Kevin Murphy from the National Aeronautics and Space Administration (NASA) framed both existing and emerging grand geospatial use cases and the limiting factors to advancing the well-being and security of our society.

Cotter spoke about how GeoAI intersects with emergency management and the need for understanding the activities of research communities, working with mission operators and making technologies usable. He surfaced questions that remain central to advancing emergency management ecosystems, such as how to better spend funds to suit a variety of emergency management goals and how to enable interoperability/standards and sharing of resources with the research community. Furthermore, Dan presented grand challenges for establishing future emergency operations centers, including: (a) how to project future hazards and threats while extrapolating to things we have not already observed, (b) the need for tools to strengthen preparedness/resilience and risk management objectives to timely counter terrorism threats and secure US borders and sovereignty, (c) how to design digital twins and leverage the metaverse to process novel data streams including imagery, audio, surveillance, detection, and measurements to equip the future emergency management operation centers.

Munsel presented views on how the GeoAI community should develop technologies that go beyond the capability of current large language models. He mentioned a need to develop large computer vision for efficient object identification yet with greater precision to apply AI to different use cases. For example, the expectations from an analyst, a combatant commander, and an AI practitioner can widely differ. The design of explainable large vision models could be a game changer to address many shortcomings that limit large-scale deployment of AI in critical decision-making applications.

Murphy presented a few grand challenges the open science community is facing, including: the growing need to develop Foundation language models for science to interact with the environment and how best can the community broaden accessibility of data and AI tools in a responsible, repeatable, and ethical manner so that local communities can solve their problems? NASA's Earth Fleet is collecting and adding in orders of 60 petabytes of imagery to archives each year. Through the open science initiative, users are positioned to exploit data that can potentially inform how the environment changes and how that relates to how people live. There is a growing need to make training datasets and benchmarks available the public, and to create engagement activities like data challenges while working across research communities to standardize the curation of data, validation, and downstream scientific evaluation tools.

3. SESSION 2: AI FOUNDATION MODELS

AI-enabled analytical tools promise to synthesize available Earth observation data more effectively and provide novel, impactful solutions relevant to various national security and climate science challenges. However, limiting factors remain in cutting-edge AI-enabled geospatial workflows to reduce the latency from data generation to decision-making on national security and open science challenges. A team of experts consisting of Raghu Ganti from IBM, David Alexander at the Department of Homeland Security, Philipe Dias of Oak Ridge National Laboratory, Kaleb Smith of NVIDIA AI Technology Center, and Muthukumaran Ramasubramanian of NASA presented a forward looking agenda on the need to identify and develop relevant foundation GeoAI models for decision support. The panel's charge was to probe the emerging trends in large AI models and formulate building blocks pertinent to the trillion-pixel GeoAI challenges emerging from natural and man-made events.

To open the session, Ganti gave an overview of foundation models as a way for the audience to appreciate the key terminology used to describe this fast-evolving technology. Ganti's remarks noted how foundation models have marked an inflection point in AI. For example, the introduction of the transformer architecture in 2017 by Google was a key breakthrough in leveraging massive unlabeled data and vast computing resources to demonstrate self-supervised learning techniques at scale. Ganti discussed fine-tuning and prompt engineering as current methods to leverage large foundation models for downstream geospatial application tasks. These are the same techniques the IBM team used to illustrate the impacts of Prithvi, a foundation model developed through a NASA-IBM partnership in collaboration with various organizations including ORNL.

Alexander discussed generative AI tools' abilty to enhance future emergency management. Although it is still early days for DHS, a diverse set of application challenges stand to benefit from the versatility of developing generative AI tools. Of particular interest to DHS foundation model customization for a variety of use cases, including predicting risk and exposure, natural language processing tools to model the effectiveness of policies, use of self-supervised learning methods to fill the ground truth data gaps, and predicting population dynamics and migration flows.

Dias presented on ORNL perspectives toward establishing a leadership capability on foundation models for national security applications. In his talk, he walked the audience through a series of existing challenges that includes the multimodality nature of geospatial data. In his presentation he covered a set of building blocks key to developing multimodal foundation models. A generalist multimodal GeoAI foundation model will require several elements, including computing at scale using large supercomputers, alignment of complementary data modalities, scalable pipelines to ingest data during training, the need for a universal tokenization scheme, and partnerships across the community. Pre-trained on large-scale multi-modal geospatial data to effectively learn representations, such models can be fine-tuned or adapted for specific applications or tasks, making them highly versatile. By simplifying and accelerating the downstream task development process, geospatial foundation models make advanced geospatial analysis accessible to a broader range of researchers, practitioners, and organizations. This democratization of GeoAI is envisioned to promote innovation and wide-spread adoption across various application domains. As a result, regulatory guardrails are becoming critical as public release of foundation models is currently conducted with little concerns on threats to society. For example maintaining privacy, mitigating bias, increasing transparency, traceability, and auditability are key as these tools inform various decision-making processes.

Smith presented a few demonstrations on how NVIDIA is enabling generative AI use to support geospatial applications. For example, NVIDIA developed NERF for geospatial for stitching 2D images to generate 3D digital elevation models. NVIDIA is currently developing generative AI tools to create the digital twin of the earth.

Ramasubramanian presented on how the NASA IMPACT team is approaching large language model (LLMs)designs for earth sciences. As an example, he illustrated LLM training and fine-tuning on corpus from earth science journals. The training strategy follows the traditional masking out of certain parts of the sentences and setting the model to predict the missing words. Ramasubramanian talked briefly about few example downstream tasks including the traditional question answering systems to retrieve open science documents, text classification, semantic equivalence, knowledge extraction and generative use cases covered conversational AI, and table to text conversion.

4. SESSION 3: GEOSPATIAL DATA INFRASTRUCTURE

The third thematic discussion focused on the enormous challenges faced by the community when integrating large, spatially heterogeneous, multimodal geospatial datasets into existing machine learning frameworks. Samantha Arundel of USGS, Rasmus Houborg of Planet Labs, Brian Freitag of NASA, and John Wegrzyn of UMBRA discussed a broad set of topics, including current state of practice in remote sensing analysis ready data and interoperability of multimodal data products, difficulties in labeling massive geospatial datasets, the development of geospatial data schemas for machine learning models, data curation and validation. The group also discussed future needs for geospatial data infrastructure solutions which integrate geospatial processes, data analytics, and AI as a single scalable platform for empowering large-scale applications.

Arundel began by emphasizing how USGS is often tasked to answer difficult spatial questions. For example, a decision maker interested in urban growth could ask to see what cities will be in the Sierra Nevada rain shadow in 20 years. Arundel said such a complex spatial query generated by SpatialXPERT, an envisioned USGS analytical tool, cannot be resolved by current ChatGPT systems. The USGS is curating authoritative and diverse data that could see new data infrastructure systems leverage relational knowledge representation frameworks including ontologies applied to incorporate complex spatial information to inform the next generation of GeoAI search systems.

Houborg of Planet Labs discussed challenges related to Analysis Ready Data (ARD) for AI applications. He noted while there is an abundance of diverse data, the lack of interoperability and data readiness are barriers for their usage. Houborg stressed the importance of developing methods for data fusion and interoperability to harness AI insights effectively, while highlighting the increasing role of AI in analyzing data from various sensors and the central role of data fusion in creating AI-compatible products. In the context of AI-enabled insights, he identified soil water content, land surface temperature, biomass, surface reflectance, and leaf area index as some of the key building blocks for AI-driven analyses toward applications in domains such as agriculture, national security, forestry, mining, energy, sustainability, and climate.

Freitag from NASA highlighted the expansion of EOSDIS Data Holdings (to around 600*PB* by 2029), describing how responding to such data volumes requires geospatial data infrastructure modernization. The demand for interdisciplinary science, a drive toward open, accessible, and reproducible research, and the challenges of data access fuel this modernization. Freitag described the Visualization, Exploration and Data Analysis (VEDA) platform intended to address such challenges. VEDA serves as an open environment that merges critical earth science datasets with open-source tools and contributors including researchers, science teams, data producers, and the broader open-source community, ensuring accessibility. Freitag also emphasized the need to enable interdisciplinary science at scale, particularly in the cloud.

Wegrzyn, representing Umbra Commercial Satellite Imagery, discussed the collection, distribution and utility of synthetic aperture radar (SAR) data. Umbra specializes in gathering and distributing SAR data, known for its all-weather and day-and-night imaging capabilities that makes it valuable for applications including defense, infrastructure mapping, and environmental monitoring. Wegrzyn highlighted the upcoming surge in available SAR data, emphasizing its unique nature compared to traditional imagery. He outlined the challenges of working with SAR and radio frequency (RF) data due to their complexity, mentioning how SAR data is distinct from traditional electro-optical (EO) imagery and how the real SAR product consists of signal collections. While human analysts can identify edge-like structures in images (e.g., defining an aircraft) derived from SAR data, models trained on EO data tend to struggle with such tasks.

5. SESSION 4: EDGE COMPUTING

Edge devices are the front line of the emerging global geospatial enterprise that collect more than 100 trillion pixels over the surface of the Earth daily. However, the current trend of connecting these edge systems to large, federated cloud computing demands ever-increasing bandwidth to handle the volume and velocity of raw data where a single satellite generates more than 100 terabytes per day. Unlike the high-bandwidth network-based approach, edge computing offers an opportunity to reduce bandwidth requirements and to improve latency and security whereby processing and computing are distributed into the edge devices and condensing raw data into more manageable information. This fourth theme of discussions focused on GeoAI on edge, neuromorphic and quantum computing, and other advances to explore possible research avenues for edge computing to address size, weight, and power (SWAP) limitations while simultaneously reducing downstream bandwidth requirements and enhancing decision-making loops.

Katie Schuman of the University of Tennessee, Angel Yanguas-Gil of Argonne National Laboratory, Erica Montbach of NASA Planetary Exploration Science, and Hamed Alemohammad of Clark University presented complementary views and discussed the state of edge computing, identified enabling technologies and defined future research directions to make edge computing a disrupting agent for GeoAI-based decision-making.

Montbach's presentation focused on autonomy in space exploration, particularly within the context of the Planetary Exploration Science Technology Office (PESTO). The PESTO recommends strategic technology investments to the Planetary Science Division (PSD) and manages PSD technologies until they are adopted by missions. She emphasized that autonomy is essential in space exploration for several reasons. Due to the vast distances and slow communication in space, autonomy is needed when the rate of decision-making surpasses communication constraints, being crucial for timely responses to unpredictable changes. Additionally, decisions often require rich on-board data, and computation at the edge could enhance system robustness or simplicity.

Schuman's presentation focused on the concept of neuromorphic computing and its relevance to GeoAI, highlighting applications, challenges, and ongoing research in the field. Neuromorphic computers are characterized by being massively parallel and based on neurobiology, modeling computing processes after the human brain. They facilitate in-memory computation, massively parallel event-driven processing, and are known for low power consumption. In the context of GeoAI, neuromorphic computing is appealing due to its parallel, fast, scalable, event-driven, robust and, particularly, low power consumption nature. The presentation also described ongoing research that explores trade-offs between energy, area, accuracy, and latency, considering robustness and resilience, especially in edge environments with limited access. Schuman also mentioned real-world applications, such as the example of an RC car developed at UT Knoxville's TENNLab.

Yanguas-Gil's presentation addressed the primary problem of smart sensors that are capable of remote operation, data processing, and learning after development. The presentation discussed the edge-HPC continuum and explored how brain-inspired algorithms and architectures could optimize complex, distributed systems. Yanguas-Gil highlighted the potential of reverse-engineering insect systems and applying those findings to models, algorithms, and hardware. Insects, particularly bees, can serve as a model system due to their embodiment of many requirements for SWaP-C (size, weight, power, and cost) GeoAI at the edge. Bees showcase online learning capabilities and aggressive data reduction, processing data efficiently through spikes. Their communication and power management strategies also hold relevance. Brain-inspired architectures can prevent catastrophic forgetting and improve past tasks, including few-shot learning. Yanguas-Gil addressed development barriers, emphasizing the need to accelerate the training and optimization of novel architectures. He highlighted the importance of application-driven co-design, developing tools for exploration at scale, and utilizing HPC to explore system configurations.

Alemohammad's presentation centered on GeoAI at the edge. He outlined six dimensions that guide the decision of processing data at the edge: scalability, latency, autonomy, privacy, security, and bandwidth. Examples of scenarios where GeoAI at the edge is needed included: enabling timely decision-making without relying on constant connectivity and the ability to monitor, evaluate, and send early warnings with minimal latency. Examples of applications that require timely interventions were discussed, such as precision agriculture, real-time field-level decisions, infrastructure management, and early warnings for issues like bridge cracking. The presentation also included discussions into the challenges of validation and evaluation in the context of GeoAI at the edge, emphasizing the importance of considering also inference at the edge rather than training, and the need to ensure trustworthiness in the systems developed.

6. SESSION 5: HARDWARE AND SOFTWARE ARCHITECTURES

As aforementioned, reducing latency from data generation and decision making in trillion-pixel challenges will demand efficient, high-performance computing hardware and network infrastructure. Also needed are scalable open-source software ecosystems to analyze these datasets leveraging state-of-the-art computational resources. In this fifth session, Prasanna Balaprakash of ORNL, Erwin Gilmore of Amazon Web Services, Vijay Gadepally of Massachusetts Institute of Technology, Mo Sarwat of Wherobots, and Forrest Hoffman of ORNL discussed a range of topics and highlighted the computational challenges facing GeoAI and offered ranging thought-provoking views on the potential next generation set of tools that could contribute to this modern day moonshot.

Balaprakash spoke about the impact of hardware on AI and computation, and how the idea wins not because it's the greatest idea, but because of the hardware-software ecosystem. Balaprakash also emphasized the importance of considering the energy consumption in terms of data movement, as it is currently a big bottleneck in processing geospatial data. He also suggested the community think about alternatives to only increasing the data input to gain better performance.

Gadepally presented on the data-centric computing stack and noted that data storage, cleaning, and discovery can be expensive. Thus considering the the efficiency of computing and storage across different computing platforms is important. He also noted that the need for better tools to build datasets and search through data as well as the need for automated tools to seamlessly compute across diverse computing architectures. Finally, he suggested rethinking the software perspective and exploring new types of computing platforms and/or paradigms to address these issues.

Gilmore shared the frustrations of using HPC (access, job management, etc.) across several applications cases. Gilmore further discussed the possibilities and alternatives of supplementing on-premise HPC resources to support at-scale works. He also mentioned AWS recognizes the importance of foundational models and will prioritize the efforts of building those models to support multiple applications.

Sarwat highlighted a lack of development of architecture that empowers developers, which has resulted in a greater focus on model building rather than creating tools to help developers integrate spatial data. He stressed the importance of well-designed software-hardware interfaces for spatial AI, rather than focusing on building-specific chips for spatial AI. Several platform recommendations from Sarwat included Apache Sedona and Geotorch AI. Such platforms are aimed at empowering developers to incorporate spatial AI into the design of new applications.

Hoffman talked about Earth System Grid Federation (ESGF), which provides a peer-to-peer distribution network for common models and open results and aims to be the best information available about climate future and impact. The ESGF is concerned with full stack security and integrity of data, not controlling access to data. The simulations generated by the ESGF are enormous (e.g. a datasize can be of 5 petabyes-sized datasets), therefore access to the data and searching capabilities are critical to support data dissemination. Possible solutions including implementing a cloud index, elastic search, and a few continental use index copies to make it more scalable.

7. SESSION 6: WORKFORCE REQUIREMENTS AND PARTNERSHIPS

The frontiers of multimodal GeoAI system will require collaboration and effective partnerships across government, industry, academia, and society. Experts convened to discuss requirements to various levels of partnership and expertise to design and deploy GeoAI model-based solutions. Representative experts that included Dawn King of NGA, Orhun Aydin of St. Louis University, Subit Chakrabarti of Floodbase, and Matej Batic of Sinergise (part of Planet) discussed a range of topics with focus on talent development needs and fostering of collaborative environments.

King spoke about challenges faced at NGA in getting young talent to work in geoscience and AI. To address this, they have established the Geoscience and AI Application (GAIA) Lab and a joint NGA/university program to expand partnerships with T-Rex tech incubator and the University of Missouri St. Louis. GAIA student projects have been launched to strengthen external partnerships and internal collaboration capabilities for unclassified mission-related problems. With this setup, they hope to attract young talents by exposing them interesting problems they can work on and contribute toward.

Chakrabarti presented on the importance of communicating the importance of GeoAI to support informed decision making for consumers and the potential for societal benefit. He stressed the ideal GeoAI workforce for this application domain, which is lacking right now, requires not only the capability to analyze data, but also the ability to connect to industrial, users and multiple disciplines and explain the uniqueness of the GeoAI offering.

Aydin presented on the challenges and opportunities in GeoAI education. He emphasized the need to focus on basic knowledge, followed by applied research, developing industry products, and ultimately benefiting society. He also highlighted the importance of looking at GeoAI concepts holistically rather than in isolation. Finally, he discussed several opportunities, including the NASA Strategic Plan 2022, NGA Strategy 2025, Taylor Geospatial Institute, and UN Sustainable Goals.

Batic provided an overview of various initiatives aimed at making EO problems attractive and engaging to diverse groups. Examples of these initiatives include the EO-Browser, Sentinel-Hub Blog, and user friendly Jupyter notebooks. Batic emphasized the convenience and importance of using blogs to address diverse groups and highlighted the summer school program, which involves working with post-secondary education (undergrads and above) students and provides a venue for those familiar with the GeoAI domain.