VELOCITY

Insights for Federal Innovators

V2. 2024

BY Booz Allen

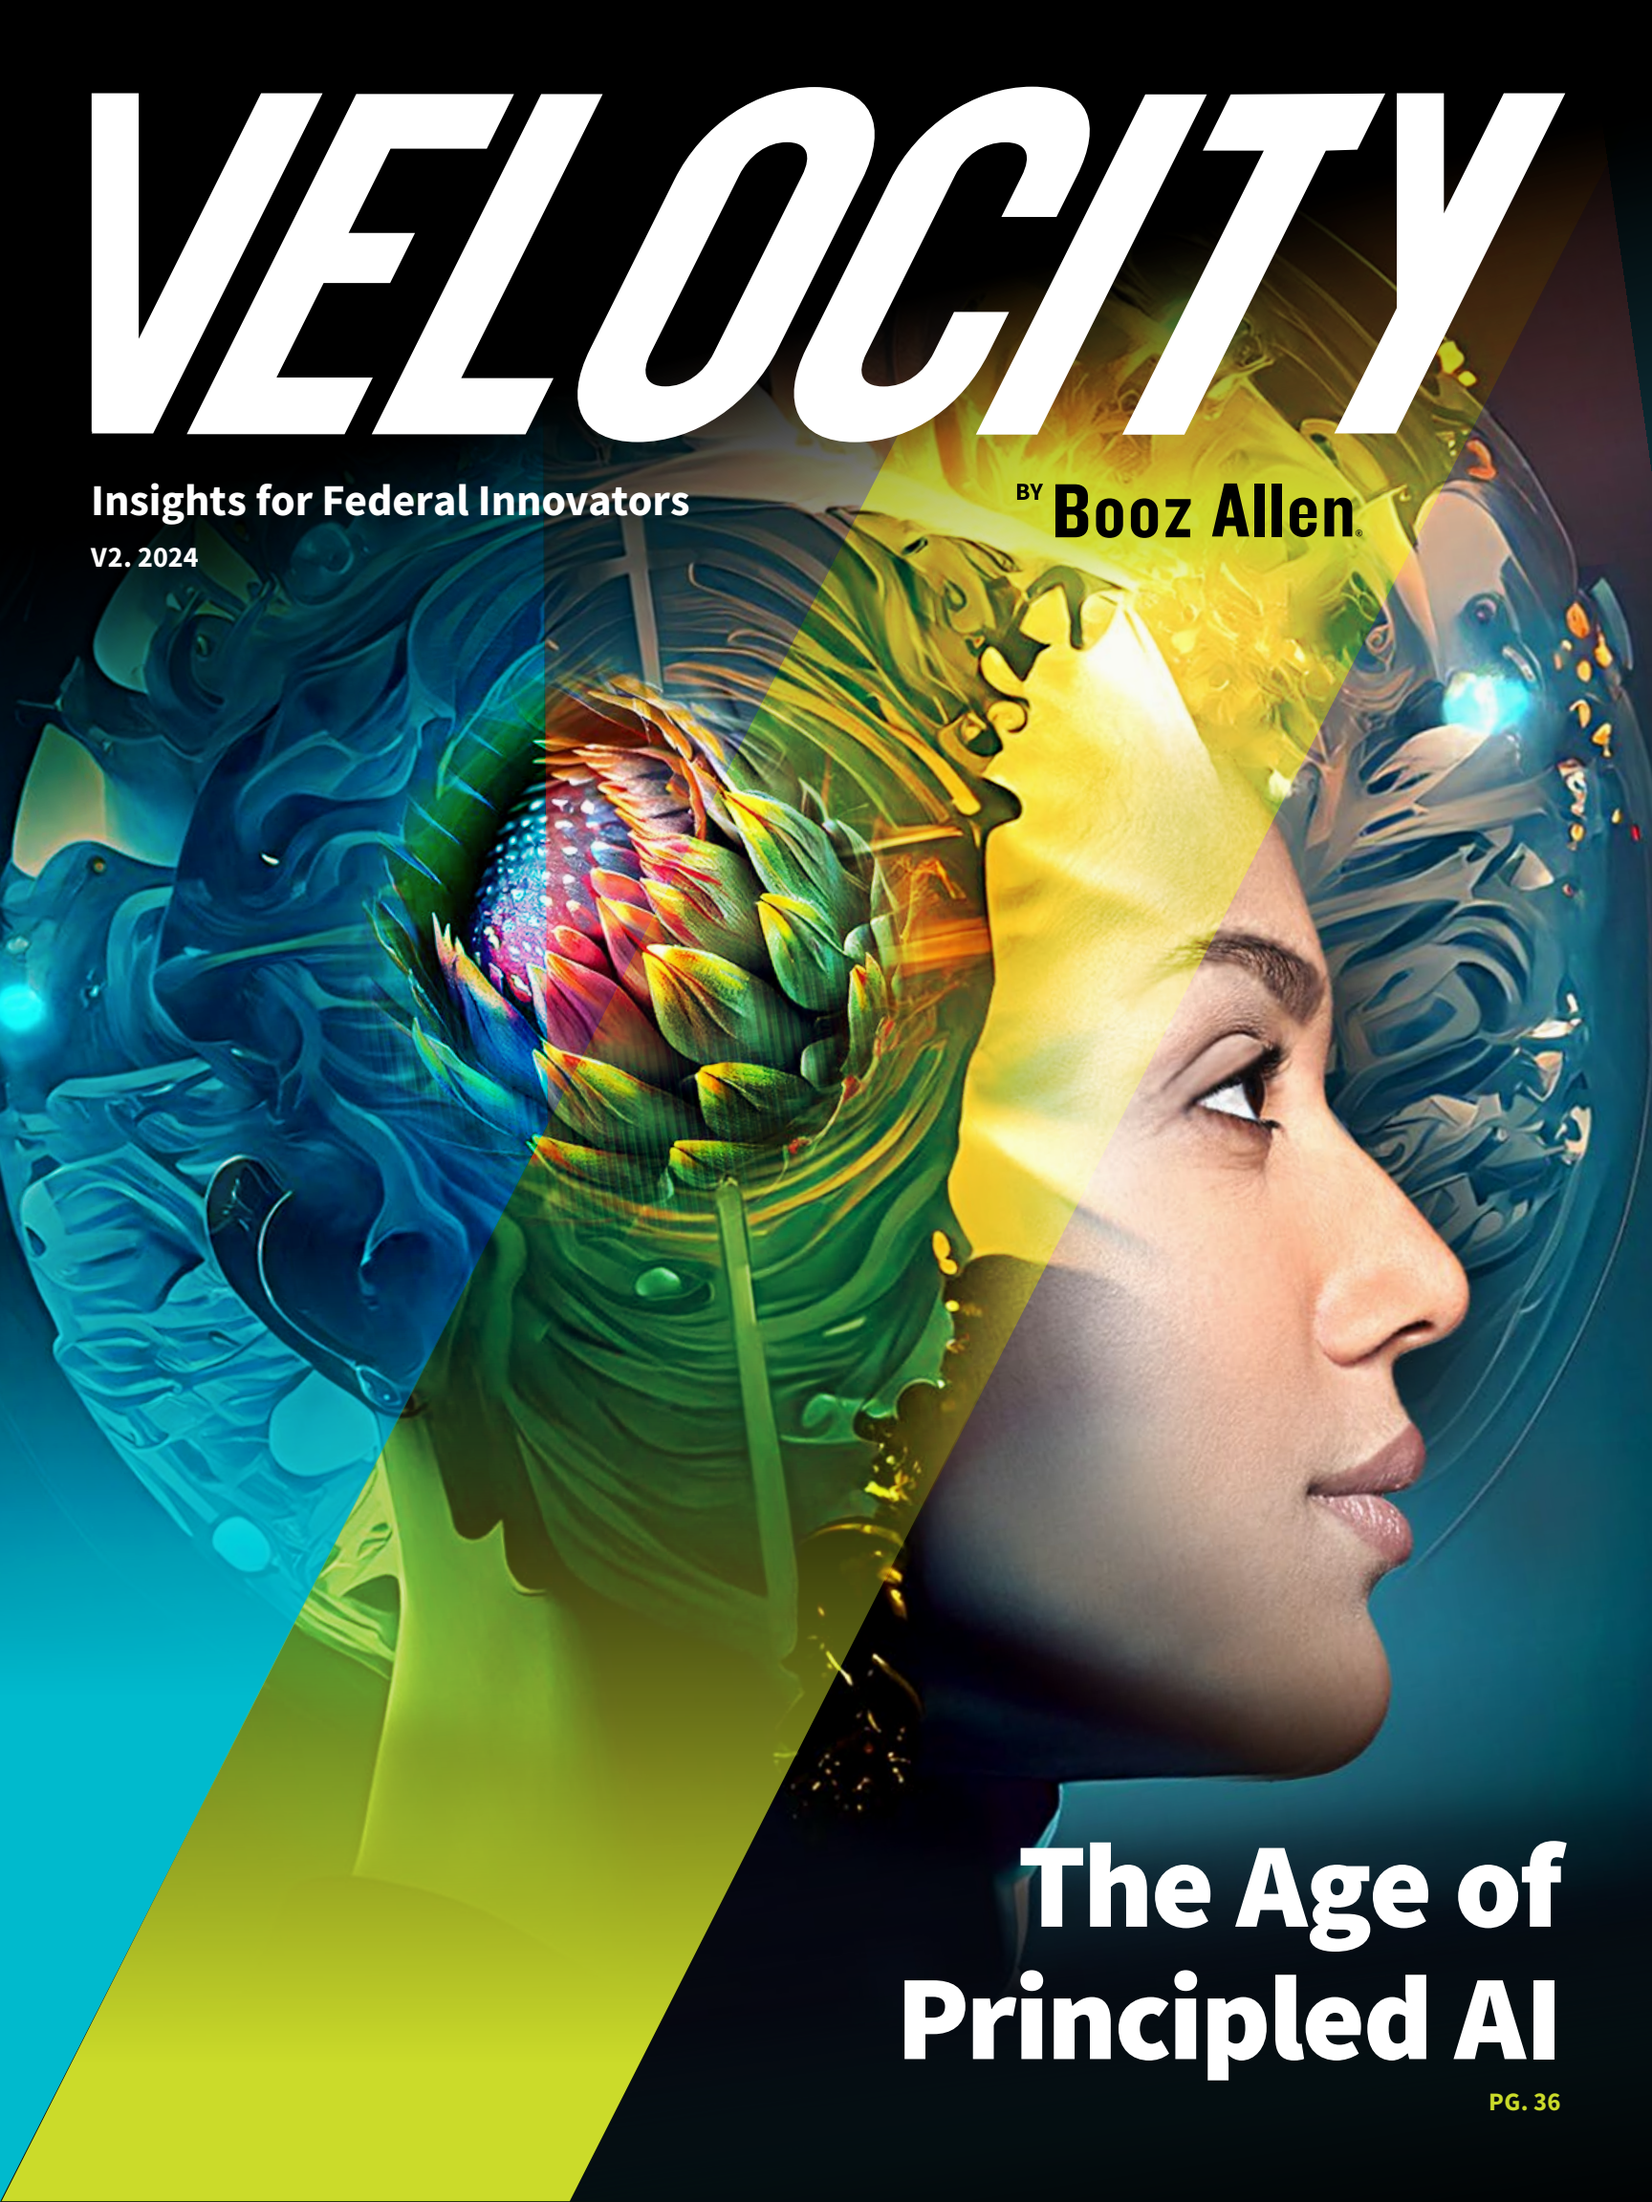

The Age of Principled AI

PG. 36

Contents

Page 36

> “A year from now my hope is that Responsible AI will emerge as the cornerstone of all AI. Every person that touches AI in some capacity will have a role in accountability, transparency, and fairness embedded in AI systems and organizations.”
>
> —**Navrina Singh**, founder and CEO at Credo AI

ABOUT THE COVER
Cover art was designed by Brody Rose. It uses AI-generated imagery from Adobe Express and Adobe Photoshop’s AI Generative Fill function. Its imagery represents optimism through change, with the protea flower as a symbol of resilience and strength in transformation.

A NOTE TO READERS
The ideas and opinions contained herein are those offered by the individual authors. They are intended as considerations in the associated technical areas discussed. They do not necessarily represent the firm’s views, but offer the breadth, depth, and currency of Booz Allen’s technical talent and that of our partners. References to any products or companies are not meant to be an endorsement by Booz Allen or any other contributors.

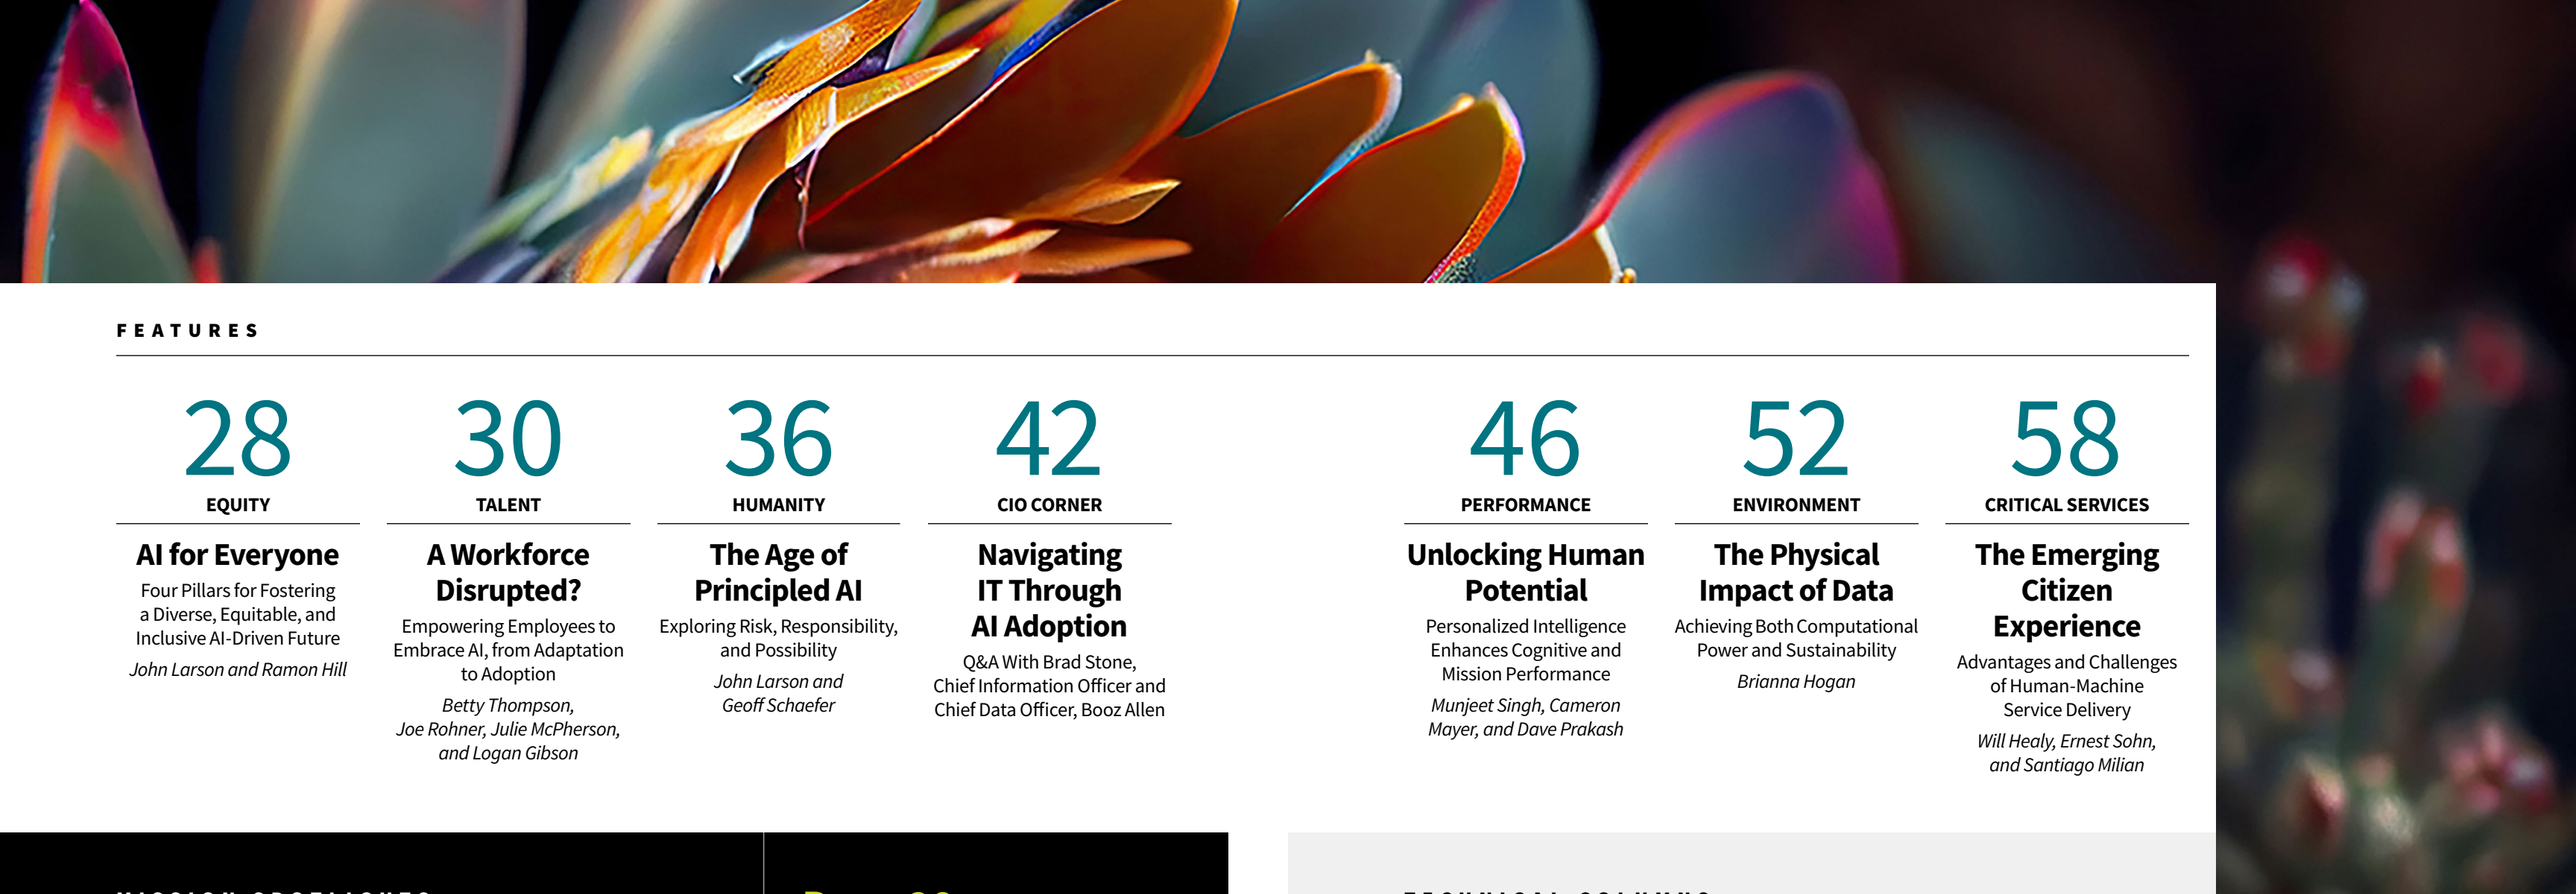

FEATURES

MISSION SPOTLIGHTS

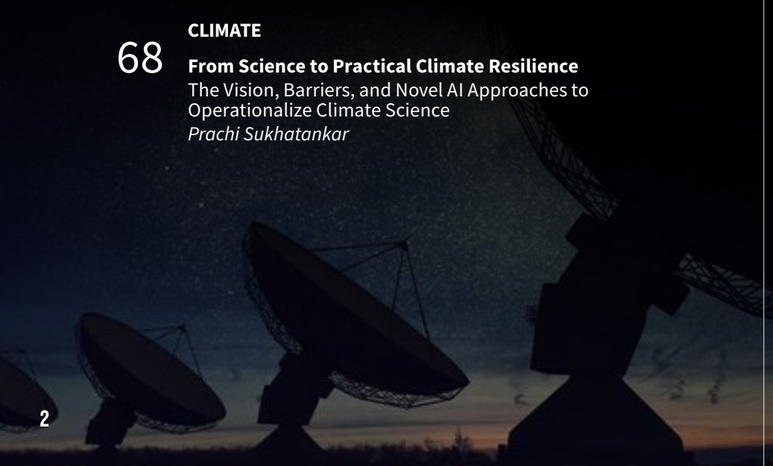

Page 28

> “Cultivating diverse talent isn’t just about filling seats—it’s about ensuring that the AI systems of tomorrow are built with the collective wisdom of our entire society, reflecting the richness of our shared experiences and values.”
>
> —**Alex Kotran,** co-founder and CEO at aiEDU

TECHNICAL COLUMNS

Booz Allen's AI Adoption Studio in Washington, DC

FROM THE CHIEF TECHNOLOGY OFFICER

I am excited to welcome you to the second edition of *Velocity*, our annual publication for government and industry leaders to explore emerging issues at the intersection of mission and technology.

In this edition, we're placing the spotlight squarely on artificial intelligence (AI)—from adopting and deploying this uniquely powerful technology to using it responsibly (see our cover story on page 36) and channeling it effectively for mission impact.

Advances in generative AI dominated headlines this year, as we crest the latest AI hype wave. Machines rapidly creating language and images democratized a sense of AI's possibilities and risks. While there's uncertainty about how the future of this technology will unfold, we do know that AI is categorically different from other past and emerging capabilities. Most important is its extreme ubiquity—the way it is poised to reshape all areas of human life.

At this open-ended AI moment, I invite you to explore the insights we've gathered from experts across the industry landscape as we seek to expand and deepen our knowledge through fundamental questions:

How will AI reimagine everything, from the global workforce to sustainability practices—and how should organizations proactively prepare, adapt, and respond?

How will agencies across defense, intelligence, and civil missions harness emerging tools to accelerate capabilities and create decision advantage?

What advances in cybersecurity, edge computing, and quantum computing should be top of mind for government innovation?

I hope you will find ideas and insights throughout *Velocity* that your organization can investigate further throughout your AI journey. As this collection demonstrates, the way we approach, build, and deploy AI to the mission today will set agencies on the path to meaningful, accelerated, and long-term impact. In this era of dramatic change, I believe we are well positioned to harness this singular enabler for good and to come together for responsible innovation.

Thank you for joining the conversation.

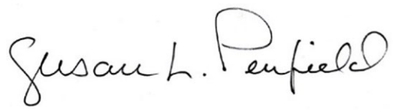

Susan Penfield

Chief Technology Officer
Booz Allen

MISSION SPOTLIGHT: SPACE

Linking Large Language Models for Space Domain Awareness

ADVANCING AI TO ENSURE FREEDOM AND SAFETY IN SPACE

Ron Craig and Michelle Harper

Crowded orbits, fast-moving adversaries, the number of active satellites expected to grow to more than 60,000 by 2030—space leaders have a challenge only advanced AI can address. Space domain awareness, the practice of tracking and understanding factors that can affect U.S. space operations, requires integration and calculations beyond current capabilities. We find that a new approach, networking large language models (LLMs), can accelerate capabilities across an enterprise of global partners.

The U.S. Needs to Move Faster in the New Space Race

Once envisioned as a sanctuary for exploration and other peaceful pursuits, space is now increasingly "competitive, congested and contested." Rival nations are developing anti-satellite tactics and weapons, militarizing the domain even as a new commercial space industry is booming—centered on initiatives from selling space data to building rockets to further NASA's mission to Mars.

Moreover, proliferated satellite constellations are being rapidly developed and launched in clusters. The first megaconstellation was launched in 2019; four years later, these groups already make up more than half of active satellites. A single network can contain hundreds or even thousands of satellites—most notably Starlink, which plans to expand its fleet to as many as 42,000. This significantly crowds a domain that has become critical for daily living. From national security to climate science, communications to traffic directions, we depend on satellite services to be there when we need them.

All these dependencies create an urgency to adopt innovations and strategies that will ensure the U.S. and its allies stay at the forefront of space—and knowing the location of space assets and their operators' intent is foundational to that goal.

Celestial Chess: A Life-or-Death Game

Staggering as it is to contemplate a trillion objects traveling through space, the challenge isn't just about the computational power required. It's about transforming a process where operators manually track data on multiple screens into a system capable of integrating complex datasets, automating processes, and applying advanced algorithms to add a new level of precision—enabling predictive analytics and, ultimately, recommended courses of action.

LAUNCHING AN ERA OF HIGHER RISK

As space capabilities continue to accelerate, risk will rise along with reward.

Expected satellite growth:
The total number could top **60,000** by 2030.

Estimated space junk:
More than **100 trillion** pieces may already be orbiting.

Expanded risk:
Even **tiny paint flecks** can damage a satellite.

SOURCES: Phys.org: "Scientists Call for Global Push to Eliminate Space Junk"; NASA.gov: "Space Debris and Human Spacecraft"

The complexities can be compared to a game of chess—a game computers have become famously good at—played in multiple dimensions, with decisions made at split-second speed. Some of the factors:

Objects need to be tracked in multiple orbits. Most commercial satellites, human space missions, and the International Space Station are in low Earth orbit (LEO), a regime that extends to about 2,000 km. Higher orbits, such as medium Earth orbit (MEO) and geostationary orbit (GEO), host navigation, weather, communications, and national security satellites. NASA's Artemis program and robotic adjuncts from multiple nations generate more traffic between Earth and the Moon. Operating safely in this region of cislunar space requires new technologies and tactics for object detection, forecasting, and collision avoidance.

Data pours in from multiple sensors from multiple sources. This results in a profusion of siloed datasets in multiple formats that need to be ingested and processed, with granular security applied. Although the increased number and diversity of data sources improve our ability to perform space domain awareness, it also introduces a data fusion challenge as multiple data sources with different formats and reference frames must be integrated in real time.

Some of the most critical datasets are classified, requiring laborious manual processes to share across domains. Although technical cross-domain solutions exist, they often can't keep pace with new file types and data structures and lack the resiliency to keep operating under stress or adversarial attack.

Modeling threat scenarios requires a vast amount of data to train algorithms. As space defense is a relatively new area, data is scarce and, in some cases, doesn't yet exist. Therefore, generating synthetic data is a necessary step, with attendant responsibilities, such as ensuring algorithms are free of bias.

Tracking the Future: The BRAVO Hackathon

LLMs—deep learning algorithms that generate content and perform other complex functions using very large datasets—exploded in popularity following the launch of OpenAI's ChatGPT in the fall of 2022. While most users have been experimenting with creating poetry, writing essays, or paraphrasing information, Booz Allen has been exploring new applications for LLMs across space-domain applications.

The capability to network LLMs was demonstrated in spring 2023 at the Air Force's BRAVO hackathon, a multi-classification event drawing over 400 experts

Figure 1: Networking LLMs demonstrated at the BRAVO hackathon. Two specially trained models shared information and seamlessly delivered a report to the operator.

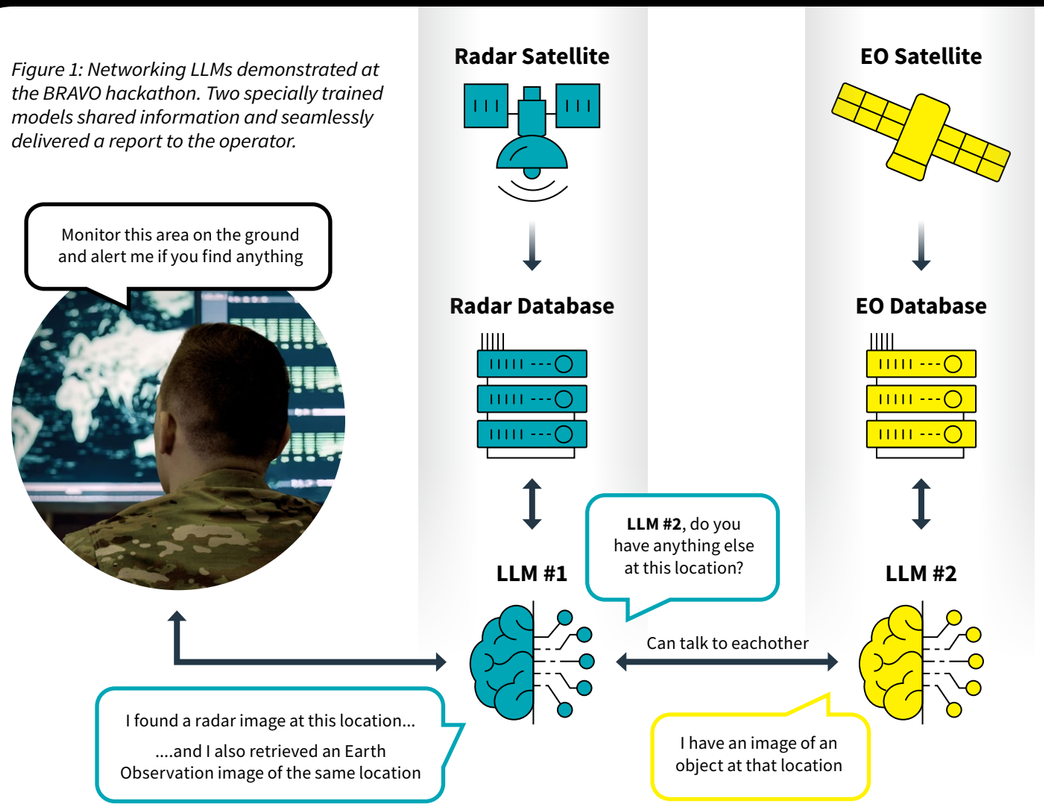

to compete in prototyping data solutions for pressing problems. The award for best data visualization and the Best User Interface award went to a team that linked two LLMs in a classified environment using zero trust protocols.

The hackathon gave the team a chance to give ensemble modeling—a process for improving multiple diverse algorithms to arrive at an outcome—new power by networking LLMs rather than individual algorithms. This opened a new path to generate the fast, comprehensive answers required to move space operations with speed and accuracy. It also provided two-way communication between specialized LLMs to amplify space operators' awareness.

After rapidly deploying a user interface, the team deployed two LLMs and wrote an app that allowed them to talk with each other (see Figure 1). The first LLM was trained in radar sensor data, while the other was trained in Earth observation (EO) imagery.

The team executed a scenario where the team member, acting as operator, asked the first LLM to watch a certain area in Asia and send an alert if anything of interest was found. No special codes were needed; the operator simply typed the request as if texting a colleague. In practice the request could have been activated another way, according to operator preference; for example, via voice recognition.

The first LLM, designated as moderator, located a radar image and asked the second model if it had any data. The second LLM, trained in EO, responded that it did and sent the image along. The first LLM then delivered both images to the operator along with a message saying, essentially, "I found a radar image at that location and retrieved an EO image at that same location." The process was simple and streamlined for the human partner.

"Software building conferences like BRAVO allowed us to push, and sometimes stumble, on some interesting solutions," said Collin Paran, the AI solutions architect who led the Booz Allen team. "Linked, multimodal, and networked AI with two-way communication will certainly unlock more insights for different organizations."

Why Networking LLMs Dramatically Increases Space Awareness

The synergy of LLMs working together and delivering ever more detailed, insightful results makes this approach significant. Say you have a mission-focused LLM trained on the Space Command's catalog of objects in orbit, the Unified Data Library. Imagine you network it with an LLM trained in avoiding collisions, called conjunctions; one trained on radar data; and another trained on a military intelligence database of adversarial threats. You've conducted skilled training and testing, and you've been using the system for increasingly critical tasks.

Now the system is deployed on a mission where the operator wants to understand the behavior of a particular set of satellites. The primary LLM identifies the satellites and, equipped with past data, knows that the last 40 times that constellation passed in that orbit in that configuration, the satellites moved slightly closer to a U.S. Space Command satellite.

Because the LLM is also configured with the other models, it can query those LLMs on the behavior as well. As a result, the primary LLM can inform the operator that this constellation has recently been flagged for moving out of its orbit and its speed is increasing just enough to create that devastating conjunction. And thanks to information from the LLM trained on adversarial threats, the primary LLM could also queue up details for the operator identifying tactics that could be at play.

Going one step further, we can imagine one of the LLMs is trained in the physics of the problem. The primary LLM could use this capability to extrapolate possible scenarios and courses of action—perhaps recommending a maneuver that uses less fuel or is less disruptive to the orbits of other satellites. The human remains in charge and is empowered to make a more confident decision, faster.

Knowledge gained on this encounter will feed into the primary LLM's learning process. And because it is linked—nested with the other LLMs to provide hierarchical, context-based learning—it gains knowledge from them at a massive rate. Its information is continuously updated as the nested LLMs are trained and validated on the latest feeds in their vast database. Every event helps the system become steadily more intelligent, creating an ever more sentient space domain awareness capability.

Adding Classified Data—and What It Takes to Do It

The crowding of space requires higher accuracy in tracking and predicting space movement, requiring data from all sources—especially classified data, traditionally difficult to share. A critical aspect of the linked approach is that networking LLMs has been demonstrated in a secure enclave, using both classified and unclassified data.

The innovation can be taken to space organizations by teams with a development environment built on open architectures, infrastructure that leverages government-owned technologies, zero trust architecture, and experience providing flexible modernization for government missions. For example, networking LLMs requires the same granular security policies as military initiatives, such as Joint All-Domain Command and Control (JADC2).

Automated DataOps ensures the onboarding of diverse feeds, standardizing formats and enforcing data standards and policies plus providing granular security. Common tools and interoperable technologies simplify development. And cross-domain solutions ensure automated workflows with modular elements that can be adapted for mission demands.

Networking LLMs to leverage their unique strengths ensures real-time advances as they collaborate on tasks. Linking these models is a practical way to deliver increasingly powerful results.

Amplifying AI Advantages

Real-time communication and collaboration between LLMs that are continually trained on trusted data opens the way to multiple advances. For example, the practice:

- Frees up operators to focus on assessment rather than switching screens and manually evaluating and comparing data to anticipate threats.
- Enables automated fusion of classified, civil government, and commercial data to train more powerful, precise AI models.
- Provides each stakeholder with automatic access to data, improving decision making for stakeholders across domains.

Training, Testing—and Then Trusting

The concept of training LLMs to be networked starts with a focus on the mission and ensuring compliance with ethical guidelines like the AI Bill of Rights and the NIST AI Risk Management Framework. AI scientists need to confirm that data sources, including other LLMs, are trained on unique datasets from verifiable sources.

Developers can quickly incorporate intelligent agents and tools that integrate easily with trusted sources once data ingestion is assured and a training pipeline and prompt templates are built. Meanwhile, strategies can streamline the process. For example, training on servers before migrating systems to the cloud saves on costly cloud computing.

Focused, nested training using trusted data and ensuring a strategic intersection between the LLMs is critical to ensure rapid, accurate returns. Data scientists need to go through the system and assess different weights, inputs, and other components, testing its information with truth data and then entrusting it with small tasks as a first step to more strategic ones. For example, it could be asked to develop a red-team attack scenario that the human experts can incorporate into a training exercise.

Linking LLMs Can Launch Adaptive Space Awareness

As General Chance Saltzman, the Space Force's chief of space operations, emphasizes, resilience is essential and continuous awareness is critical for the Space Force's strategy of competitive endurance.

Networking LLMs to leverage their unique strengths ensures real-time advances as they collaborate on tasks. Linking these models is a practical way to deliver increasingly powerful results. It's scalable, allowing the networking of multiple LLMs. It's model-agnostic, so it can be used with any LLM. And it holds the promise of connecting the vast, siloed datasets that are key to avoiding celestial collisions and countering adversarial attacks.

Ron Craig *is vice president of space strategy and solutions at Booz Allen.*

Michelle Harper *leads software projects that accelerate integrated capabilities for Booz Allen clients, including the Space Force.*

SPEED READ

There's an urgent need for advanced AI as an expected surge in the number of active satellites by 2030 makes space increasingly "competitive, congested, and contested." Networking large language models (LLMs) can enhance space domain awareness and address these challenges.

The BRAVO hackathon showcased the transformative capabilities of networked LLMs in space operations. By linking these models, they can communicate, share data, and amplify space operators' awareness, leading to more efficient and precise decision making.

By allowing LLMs to collaborate and learn from each other, networking them can provide comprehensive insights into space behaviors and threats. This interconnected system promises enhanced space domain awareness and strategic advantage.

Can Quantum Supercharge AI?

QUANTUM COMPUTING'S ROLE IN THE EVOLUTION OF AI

Isabella Bello Martinez, Ryan Caulfield, and Brian Rost

Quantum mechanics and machine learning, two of the most transformative forces of the past two centuries, are converging to mark the start of a new era of AI. This convergence—known as quantum machine learning—has the potential to address limitations of classical (meaning "not quantum") machine learning, particularly processing power and speed.

Classical machine learning has undoubtedly made significant strides in data processing and predictive analytics. Yet it is often limited by computer speed and memory, especially when dealing with large and complex datasets. Quantum machine learning (QML) leverages some of the unique features of quantum systems to overcome these limitations.

QML is currently on the cusp of transitioning from offering a purely theoretical advantage to finding real-world, high-impact applications. In the realm of drug discovery, where the search for new drugs often involves navigating a vast space of molecular combinations, QML has shown potential for identifying promising compounds more efficiently. Recent evidence also suggests QML provides advantages for computer vision, where identifying key features in unlabeled images is becoming increasingly important. Relatively small quantum models have the power to perform well on even the largest and most complex datasets that would otherwise require impractically large, classical models. These, and other advantages of QML, arise from the fact that quantum systems are inherently more complex and more capable of representing complicated patterns than comparable classical systems.

Although quantum machine learning is still in its early stages, the progress made so far indicates that QML will have a transformative effect on AI. The impact will be felt across diverse fields, many of which are directly aligned with government interests, such as designing better materials for assets in space, improving health diagnostics, and advancing computer vision for superior ISTAR (intelligence, surveillance, target acquisition, and reconnaissance).

An Evolution, in Partnership with Classical Algorithms

Humanity has been "computing" since we first started using numbers. While the types of computations and technology used today are drastically more sophisticated than keeping track of bushels of wheat on an abacus, the fundamental computational model has remained the same. A computer carries out these computations using bits, which are objects that can be in one of two states—like "on" or "off"—called 0 and 1. For the first time in history, we're starting to compute using a completely novel computing paradigm known as *quantum computing*. Based on **quantum mechanics**, quantum computing is expected to have a wide-reaching and transformative impact.

At the core of quantum computation are ***qubits*** (quantum bits), the quantum version of bits. Like a bit, a qubit can exist in the 0 state or the 1 state. Unlike a bit, a qubit can also exist in a uniquely quantum state that is analogous to being partly 0 and partly 1, or existing along a continuum between 0 and 1. This makes qubits more complex than bits, enabling quantum computing to tackle problems well beyond what would be possible with classical computing.

DID YOU KNOW:

- **Quantum mechanics** is the physics which governs very small (particles, atoms), very cold (superconductors, superfluids), and exotic systems (lasers, semiconductors, stars) from which surprising behavior arises.
- Researchers are experimenting with different ways of making **qubits**, and no clear winner has emerged yet. Implementations range from basic quantum objects such as atoms, ions, or light (photons) to more exotic quantum systems such as nanodiamonds, superconductors, and more.
- At the heart of the advantage of quantum computation is a uniquely quantum phenomenon called **entanglement**, whereby multiple qubits become fundamentally linked and share information between themselves in ways not possible classically.
- The **operations** implemented by a quantum computer are called quantum gates. Both quantum and classical algorithms can be specified as circuits, a graphical representation of a series of gates to be applied to the qubits.

Both quantum and classical computers compute by executing a series of instructions called an algorithm, which manipulates the states of their underlying (qu)bits. Unlike classical algorithms, which can only flip bits between 0 and 1, quantum algorithms use a richer variety of **operations** that take advantage of the qubit's complexity.

Quantum Computing Meets Machine Learning

From computer vision to cybersecurity and everything in between, this past decade has shown us the versatility and power of machine learning technologies. Given this, we can expect that expanding the capabilities of machine learning will continue to drive progress across the board. QML, the integration of quantum computing into machine learning systems, offers one promising way to do just that. QML is anticipated to provide improvements in speed and performance to diverse AI application areas spanning medicine, finance, data analysis, and more.

Evidence has been mounting that QML is indeed capable of delivering on these promises. For example, a 2021 study in the *Journal of Chemical Information and Modeling* demonstrated how using QML could accelerate the process of drug discovery to combat diseases, such as COVID-19 and tuberculosis, compared to using analogous machine learning methods (see "Quantum Machine Learning Algorithms for Drug Discovery Applications"). This was demonstrated by using the quantum versions of common classical machine learning methods, such as deep neural networks and support vector machines, to classify which molecules were potential inhibitors for a target disease. Despite the quantum computer's imperfect, error-prone nature, the QML models achieved similar accuracy to the classical models while demonstrating a speed advantage that grew with the size of the dataset. The timing data suggests that as molecule databases continue to expand, QML's superiority will only become more pronounced and remain viable when classical machine learning methods begin to struggle. This is not limited merely to drug discovery; it stands as a strong affirmation that the theoretical speedups for QML will translate into wide-reaching, real-world impacts.

HIGH-FREQUENCY QUESTIONS ABOUT QUANTUM COMPUTERS

Are quantum computers better at everything?	No. Quantum computers excel at a limited number of tasks. Examples include, but are not limited to, simulating quantum physics for material science or drug development, optimization for logistics or finance, and math for machine learning or cryptography. They offer no benefit for many things we do with computers. That said, the advantage of using a quantum computer for the right kinds of problems can be enormous.
Will quantum computers replace classical computers?	No. Quantum and classical computers will work together. Because quantum computers offer no advantage for most computing tasks, we will still want classical computers to handle most of our computations. In the same way many current high-performance computing setups call a GPU (graphics processing unit) to accelerate certain tasks, future setups will likely be classical computers that can call a QPU (quantum processing unit) as needed.
So, quantum computers actually exist?	Yes. Many companies, governments, and academic institutions currently have quantum computers that vary greatly in size, power, architecture, and more.
If quantum computers already exist and are supposed to be so powerful, then why aren't we doing more with them?	Quantum computers are still not sufficiently powerful to be commercially impactful. Despite the largest general-purpose quantum computers theoretically having plenty of qubits (hundreds) to completely outclass our best supercomputers for certain problems, these quantum computers are unlikely to outclass a smartphone. This gap between theory and practice exists because current quantum computers are prone to errors that severely limit their power. It is known, in theory, that these errors can be overcome, and that "perfect" quantum computers can be built. The timeline for achieving this is still unclear, but most experts estimate in terms of decades, not years (see The Quantum Threat Timeline Report 2022).
If "perfect" quantum computers won't be available soon, why should we care about them now?	We expect that imperfect quantum computers will be useful for solving real-world problems soon, with many estimates for achieving a quantum advantage falling between this year and 2028. It remains an open question of when and for what quantum computers will first become impactful. One promising candidate application for near-term quantum computers is machine learning.

Figure 1: Illustrative steps for inference and training of a hybrid quantum-classical model

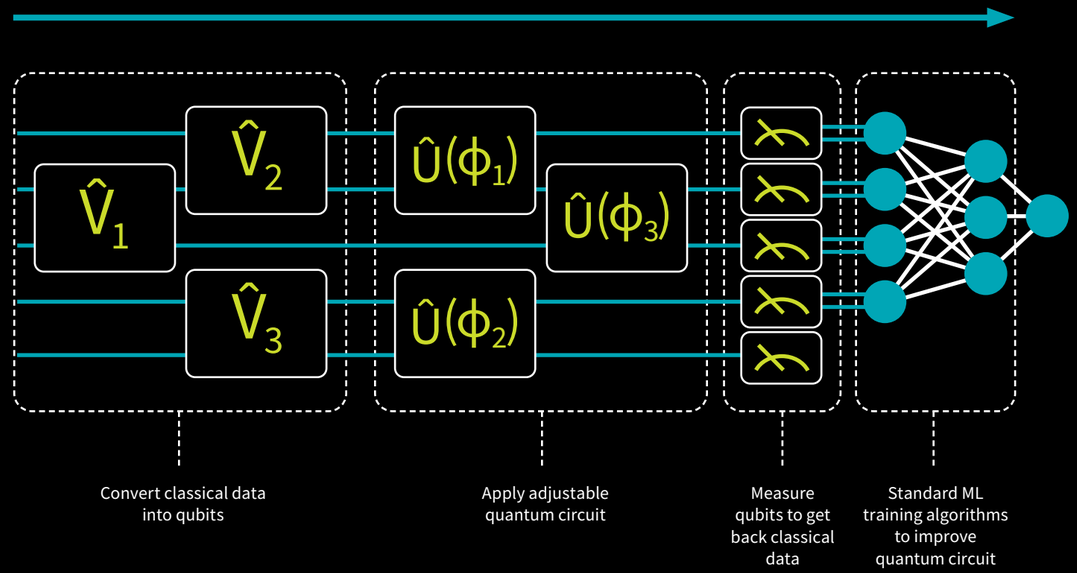

In addition to speed advantages, QML is predicted to perform well using significantly smaller models than classical machine learning, making it possible to tackle previously infeasible problems. The number of features a classical model can represent is directly related to its size. Qubits, however, contain more information than bits, so significantly smaller quantum models can represent the same number of features. This may allow for reasonably sized quantum models to perform well in situations where an infeasibly large classical model would be needed.

Evidence of this was provided by a 2022 computer vision study in *Quantum Science and Technology* that focused on learning with unlabeled data (see "Quantum Self-Supervised Learning"). It is quickly becoming infeasible to label many datasets of interest due to their sheer size, such as photos on social media or images taken by self-driving cars, so the ability to efficiently learn on unlabeled data is increasingly critical. Techniques for learning on complex, unlabeled datasets often require impractically large models since they must discern and represent complicated patterns in the data.

In this study, the researchers trained a model to classify simple images of planes, cars, birds, cats, and deer. They then trained the same model again but replaced part of the model with a quantum equivalent. Using a **quantum simulator**, the authors showed that the quantum model outperformed a classical model of the same size and that a smaller quantum model could achieve the same performance as the classical model. The quantum model was then run on a real, imperfect quantum computer and matched the classical performance, despite the high error rate of the quantum computer. The study shows that this QML technique is capable of overcoming the classical bottlenecks and is robust against the errors in our current quantum computers, suggesting that QML for computer vision may become a practical application soon.

DID YOU KNOW:

A quantum simulator is classical software which mimics a quantum computer. Researchers often use this because 1) it mimics a perfect quantum computer without errors and 2) it can be difficult and expensive to get time on a real quantum computer. Of course, because quantum computers are more powerful than classical ones, the **quantum simulator** is very inefficient and only works for simulating a small number of qubits.

From computer vision to cybersecurity and everything in between, this past decade has shown us the versatility and power of machine learning technologies. Given this, we can expect that expanding the capabilities of machine learning will continue to drive progress across the board. QML, the integration of quantum computing into machine learning systems, offers one promising way to do just that.

Other studies run on quantum simulators point to quantum advantages for a variety of machine learning tasks crucial to applications, such as detecting financial irregularities, increasing battery efficiency, and diagnosing diseases, such as breast cancer and COVID-19. The evidence contained in these studies and many others is strengthened by the fact that it serves to confirm known theoretical advantages of QML. The combination of empirical results and theoretical predictions underscores the reliability and potential of QML for real-world applications.

While these initial results are exciting, it's crucial to clarify that current best QML models cannot yet outperform the best classical machine learning models. The QML examples highlighted above were built using a small number of qubits, with the computer vision being the largest at only eight qubits. To fit the data on this small number of qubits, the original data had to be compressed significantly. Though QML displayed benefits over machine learning models trained on such compressed data, it still lags behind the performance machine learning models trained on the full, uncompressed data.

Advancement Within Reach

As the power of quantum computers continues to grow, we can expect QML to swiftly catch up to and surpass the current state of the art as it develops into a powerful tool.

This rapidly developing technology is poised to bring unprecedented advancements to a wide range of AI application areas, from computational biology to climate modeling, by offering improvements in performance and efficiency. Although QML is still a nascent technology, it has already been validated through small-scale experiments and theoretical work. It's essential to acknowledge that, like all emergent technologies, QML has its nuances and challenges. However, given the explosive growth trajectory of quantum computing, we can expect QML to rapidly transition from a fledgling technology with limited practicality to an invaluable tool that improves our ability to solve complex problems beyond the reach of classical computing techniques and advance the national security agenda.

***Isabella Bello Martinez**, **Ryan Caulfield**, and **Brian Rost** are scientists at Booz Allen, helping clients understand what quantum computing can do today and how to prepare for the next wave of capabilities.*

SPEED READ

Quantum computers are increasingly being utilized as part of machine learning, creating the exciting new field of quantum machine learning (QML), which promises to overcome some of the processing power and speed limitations of other machine learning methods. QML is likely set to revolutionize areas such as drug discovery and computer vision by efficiently handling large and complex data sets.

While quantum computers excel at specific tasks, they will need to work in tandem with classical computers to enhance their capabilities for certain problems. Relatively small quantum models have the power to perform well on even the largest and most complex datasets that would otherwise require impractically large, classical models.

Current QML models have shown promise but they still face challenges, such as the need for data compression due to limited quantum bits (qubits). Despite these hurdles, the rapid growth and development of quantum computing indicates that QML could soon transition from a nascent technology to a transformative force.

MISSION SPOTLIGHT: INTELLIGENCE

Open-Source Overload

MOBILIZING PUBLICLY AVAILABLE DATA FOR NATIONAL SECURITY

Eric Zitz and Gabi Rubin

While it will never replace classified intelligence collection and analysis, OSINT is the "INT" that best balances the traditional need for secrecy with the increased need for rapid information sharing to address developing and emergent threats.

In light of today's national security threats, the need for the U.S. intelligence community (IC) to swiftly process and disseminate information has never been more urgent. Dynamic and multifaceted challenges demand unprecedented speed and agility, and this imperative has been further magnified by the exponential rise of publicly available information.

Open-source intelligence (OSINT), as defined by the SANS Institute, is "intelligence produced by collecting, evaluating, and analyzing publicly available information with the purpose of answering a specific intelligence question." Beyond mere information, OSINT can contextualize, enhance, and validate analysis and provide opportunities for increased dissemination. While it will never replace classified intelligence collection and analysis, OSINT is the "INT" that best balances the traditional need for secrecy with the increased need for rapid information sharing to address developing and emergent threats. In a panel about the future IC workforce, Patrice Tibbs, deputy chief of the Open Source Enterprise at the Central Intelligence Agency, called OSINT the "INT of first resort" and said that OSINT brings a new perspective not only to the information being collected but also to the entire process of intelligence collection.

But OSINT sources are far from static and have a scope that continues to expand. What was once limited to historical research or news monitoring now encompasses a rich array of technology and collection methods. Even technical domains traditionally associated with geospatial intelligence or signals intelligence can now be integrated and analyzed alongside more traditional OSINT sources, giving rise to more robust and holistic intelligence. Director of National Intelligence Avril Haines affirms that various intelligence disciplines and IC agencies are all diligently developing their versions of open-source capabilities. "However, we are not in a position where we feel as if the entire intelligence community is leveraging … the best of what we can do in this space yet, and that is something that we have been focused on," she said, in a hearing to the Senate Committee on Armed Services.

With its advantages, OSINT also disrupts the intelligence playing field. According to former Defense Intelligence Agency Director Robert Ashley and former Principal Executive of the Office of the Director of National Intelligence (ODNI) Neil Wiley, "the ubiquity and accessibility of this public data" narrows the advantage of the IC's niche and proprietary intelligence sources and methods, which emphasizes the need to continuously evolve and adapt how publicly available information is integrated with classified resources.

National security organizations must evolve their understanding of OSINT to keep pace with emerging adversarial capabilities in big data aggregation, cloud computing, AI, and machine learning (ML) analytics.

Data, Data Everywhere

Big data stands as the formidable linchpin within the OSINT domain, shaping the very frameworks of collection and processing methodologies. According to Statista, in this year alone, 120 zettabytes of data will be created, transformed, captured, copied, and consumed and that number will grow by 20–30 zettabytes annually thereafter. To put that into perspective, one zettabyte is equal to one trillion gigabytes, or over 570 million years of YouTube videos. Even if the IC is only expected to ingest and analyze a fraction of that, it's a nearly inconceivable task.

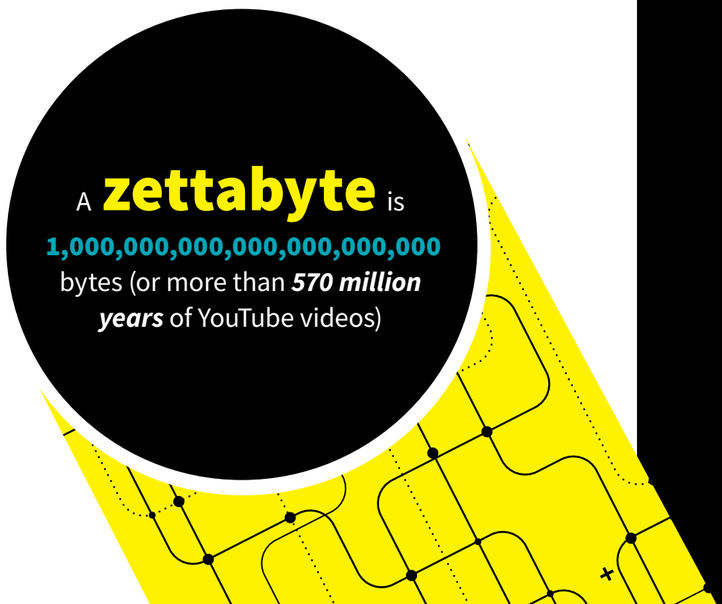

Apart from building advanced analytic engines to automate the exploitation, this data challenge requires analysts to determine what data to prioritize and target, based on the mission problem, in order to limit the volume required for aggregation. Accomplishing this can provide analysts with previously unavailable insights, including patterns of life; messaging trends; social, financial, and supply chain networks; breaking news updates; and more.

Let's take a closer look at the range of sources available and which need to be curated to paint a complete picture.

While OSINT aggregation, processing, and analysis is not free, many of the sources that form the basis for OSINT insights are. Free and open, publicly available information can range extensively in terms of data type and content. Though by no means an exhaustive list, some examples include:

- news articles
- public social media posts
- marine or air traffic monitoring sites
- international trade databases
- company registration databases
- government public records
- nongovernmental organization (NGO) reports
- civil imagery
- human geography and infrastructure data

It's a matter of knowing that information exists, then collecting, vetting, and processing it to address intelligence requirements. However, an additional level of vetting is often required for free sources of publicly available information to ensure the data can be validated and avoid leveraging mis- or disinformation. Extracting value from freely accessible information sources is often a race against the clock.

For example, in 2023, Twitter stopped allowing unregistered viewers to see individual tweets and limited the number of tweets non-paying users could view per day. Many aggregators (though not all) were locked out and no longer able to leverage the Twitter application programming interface (API). Similarly, the United Nations' Comtrade database recently updated its subscription plan and now only allows limited access and no downloads on their free public user license.

There is also a subset of publicly available information that is commercially sold, which includes geolocation; commercial satellite and airborne imagery; radio frequency data; subscription news and journal articles; and databases of pre-aggregated public records, among many others.

An ODNI report from 2022 describes that the purchase of commercially available information (CAI) can be made one time or on a subscription basis and may involve a purchaser "directly ingesting the CAI or obtaining a license agreement that affords a continuing right of access." The report further notes that "there is today a large and growing amount of CAI that is available to the public, including foreign governments (and their intelligence services) and private-sector entities, as well as the IC. CAI clearly provides intelligence value, whether considered in isolation and/or in combination with other information, and whether reviewed by humans and/or by machines."

As the IC increases investment in commercially available information in alignment with the nation's strategic imperatives, it will need to continue developing new OSINT policies and frameworks. These guidelines can help ensure the IC handles and collects information properly, maintains vigilance against potential counterintelligence threats, and allocates resources comprehensively to support the development of tools, training, and tradecraft.

Deciphering and Accelerating the Value of OSINT

Modern technologies are turning the tide and empowering national security organizations to harness the immense potential of public and commercially available data and analytic sources. Advanced modeling and automation capabilities, decentralized data processing, and AI/ML are alleviating the tasks that were once arduously executed through hours of manual research, exploitation, integration, and analysis by the IC.

AI's ability to process zettabyte-scale datasets is already being seen in generative AI and large language models, and new methods and approaches are being applied to less structured data, such as imagery. **Synthetaic Chief Executive Officer Corey Jaskolski says, "The future of AI lies in intuitive tools that put that power into the hands of subject-matter experts who can elucidate insights from the AI in real time and then make meaning from the insights it generates."** As AI/ML capabilities are paired more deeply with human analysts, national security stakeholders will see costs greatly reduced and mission outcomes dramatically enhanced.

"The ability for analysts to easily capture, collaborate, and automate their tradecraft frees them to perform higher level analysis instead of worrying about data representation and translation," says Nask Incorporated Technical Director Ken Pratt. "This is a force multiplier allowing fewer analysts to perform more effective and valuable analysis against a larger corpus of data."

Beyond automation, the decentralization of information processing heralds remarkable efficiencies. Here, raw data is exploited closer to its source and the integration occurs downstream, closer to the analyst or user. The embrace of decentralized workflows and the synergy of specialized tools significantly amplify the pace of exploitation and insight generation. This is true for nearly all data types and sources, including text-based foreign media and social media; published databases; and technical data from satellites, aircraft, and other sensors.

An excellent example of this is the potential for decentralized Commercial Synthetic Aperture RADAR (COMSAR) collection and exploitation closer to the source, which would maximize the community's ability to harness the rich, remote sensing metadata while minimizing the costs of data transfer and storage of large COMSAR imagery files. **"By delivering and processing that data at the edge, the analytic insight delivery time to the user is measured in seconds and minutes, not hours," says Ursa Space Vice President of Government Programs George Flick. "This allows for quicker situational awareness and decision making for users and operators. Speed is essential."** New and emerging technologies and techniques in AI/ML are designed for large and complex datasets, and advanced algorithms shine the closer they're hosted to the raw data.

OSINT has continued to evolve alongside the emergence of new data sources and techniques, spanning the spectrum from free or open to commercial realms. The contemporary landscape boasts an unprecedented volume and diversity of OSINT, demanding more robust analytic capabilities in AI/ML and automation to sustain an intelligence edge and enable the IC to harness the massive amounts of information proliferating in the public sphere.

The key to leveraging these sources and capabilities and driving enhancements for national security missions is to purposefully discover, identify, process, and integrate the right data—it is not to try to process every byte of unclassified data that's out there. By understanding the value of the data, decentralizing its exploitation, and layering it through AI/ML technologies and techniques, the national security community can create an advantage through these sources and insights.

Eric Zitz *is a mission technology leader for Booz Allen's national security business. He develops new capabilities and solutions that integrate data science, modeling, and automation to enhance intelligence operations.*

Gabi Rubin *is a leader in Booz Allen's OSINT capability Global4Sight®, which offers language-enabled and data-driven intelligence solutions for civilian, intelligence, and military agencies.*

SPEED READ

The modern landscape of national security challenges compels the U.S. intelligence community to swiftly process and disseminate information with unprecedented speed. This urgency is heightened by the exponential rise of publicly available information intensifying the demand for agility and adaptability. The ubiquity of public data narrows the advantage of classified intelligence sources, necessitating continuous evolution and integration.

Open-source intelligence (OSINT) strikes a balance between secrecy and rapid information sharing, offering new perspectives on intelligence collection processes. The potential of OSINT is not static; it evolves to encompass diverse sources and domains, from text-based to technical data, fostering a holistic intelligence picture.

National security organizations leverage modern technologies to tap into the vast trove of publicly available and commercially sold data. Automation, decentralized data processing, and AI/ML capabilities expedite the exploitation of data previously requiring manual efforts. The key is to identify valuable data, decentralize its exploitation, and integrate it with AI/ML technologies.

Generative AI in the Wild

DEPLOYING POWERFUL MISSION APPLICATIONS WITH PURPOSE

Ted Edwards and Alison Smith

"In the world of startup valuations, there's generative AI—and everything else."

In an article whose title shares that quote, PitchBook data spotlights the striking separation between early-stage AI startups and other young companies in initial funding rounds. In 2023, generative AI companies' pre-money valuations increased an incredible **16%** from the prior year, compared to a significant drop of **24%** for startups in all other sectors attempting series A and B funding (see Figure 1).

The enthusiasm from investors is well founded: Generative AI recently took AI from the realm of engineers and democratized the technology in a way that no other low-code/no-code platform has ever done. OpenAI's ChatGPT launched into the public consciousness with a bang—setting the record for the fastest-growing consumer base for an application—and immediately demonstrated far-reaching capabilities as an AI assistant on everything from composing recipes to summarizing complex topics and generating computer source code. Google and Bing have incorporated generative AI into their search systems, enabling direct responses to queries and circumventing a user's need to sift through numerous webpages to find answers. Service providers like Midjourney, Stability AI, and OpenAI's DALL-E have harnessed generative AI to create accurate and striking images from text-based descriptions of the desired output.

Despite the groundbreaking nature of these services, enterprises and their users should take a deliberate approach. Generative AI surfaces considerations regarding fair use and copyright based on how they use training data. Models can easily mislead users, novices, and experts alike, making it difficult to distinguish fabrications from factual content (this particular risk manifests as "hallucinations," which are outputs and information from models that sound highly plausible and convincing but are simply made up or incorrect). And of course, for enterprises with sensitive data, there's the serious risk of data spillage by employees sharing confidential information with open-source generative AI models.

The immense power of generative AI is ripe to uncover transformative opportunities across government missions. But like any tool, generative AI is only effective if it is applied with purpose. Therefore, rather than asking "How can I use generative AI?" the more nuanced, strategic question is: "Where will generative AI tackle a challenge better than the other tools in our toolbox?"

This article explores the singular role of generative AI in revolutionizing government missions while agencies navigate emerging challenges; weigh the costs of application; and ensure responsible, effective use.

Comparison of median early-stage, pre-money valuations (in millions)

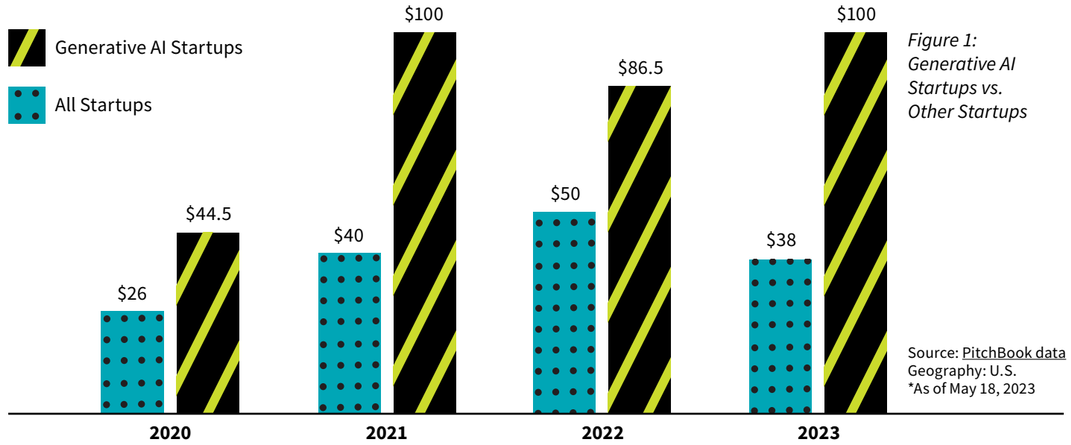

Figure 1: Generative AI Startups vs. Other Startups

Source: PitchBook data
Geography: U.S.
*As of May 18, 2023

Rise of Generative AI

The rapid impact of generative AI—particularly ChatGPT—has prompted significant investments and has already revolutionized businesses and industries by augmenting core processes and support functions. Organizations are increasingly expected to deploy generative AI to drive productivity, streamline processes, and disseminate information.

A Meteoric Rise: Reaching 1 million users in 5 days, OpenAI's ChatGPT serves as an exemplar for the speed with which people and enterprises have embraced the value of generative AI.

Major Players: Key innovators in the generative AI landscape include Google, Microsoft, Amazon, and IBM. These companies are directly embedding generative AI technologies into their products, leveraging their respective strengths to advance the field.

Growing Enterprise Interest

'23 In 2023, **70%** of organizations are currently exploring generative AI, with **19%** in pilot or production according to Gartner®, and Forrester predicts that about **10%** of Fortune 500 enterprises will generate content with AI tools.

'24 By 2024, approximately **40%** of enterprise applications are expected to incorporate embedded conversational AI (Gartner®).

'26 By 2026, generative design AI is poised to automate **60%** of the design effort for new websites and mobile apps (Gartner®).

'27 By 2027, nearly **15%** of new applications will be automatically generated by AI without human intervention (Gartner®).

Recent Innovation Transforms a Long History of Research

Generative AI as we know it today is relatively new. In fact, seasoned researchers and engineers have been surprised by the fidelity of language generated by large language models (LLM) like ChatGPT, but many of the models and approaches underlying it have been studied for years.

The seminal 2017 paper "Attention Is All You Need" by researchers at Google laid the foundation for what we now know as generative AI. It introduced Transformer architecture, a way for these new models to direct their "attention" to specific parts of data, massively increasing their capabilities. These findings were groundbreaking and now serve as the basis for nearly all of the most state-of-the-art generative AI systems. So, at their core, how do LLMs work? LLMs are trained to generate responses by predicting the words that will produce the most desirable response. They "learn" to do this by studying the relationships between words in an enormously large collection of text produced by humans. The training process requires specialized hardware and so much computational power that researchers have raised concerns about the level of energy consumption and environmental impact from training these models (for more on sustainability, see page 52). Throughout the training process, information about the relationships of words in text is processed through deep neural networks consisting of layers of components called transformers, which can capture complex linguistic structures and subtle details about meaning and interaction. Although the structure of these networks is intricate, it is relatively simple compared to the complex and unpredictable results the algorithms produce.

Now there is an entire field of research to study "emergent behaviors," which are actions or abilities that an AI system was not explicitly designed to do. Experts are working to understand how LLMs learn, how to determine what they know, how to better train them, and how to shape their output for specific use cases. Rapidly, foundational LLMs are deploying with advanced techniques, such as supervised fine tuning and reinforcement learning, which reward models for their predictive abilities and success during training or when completing a task.

While many technologists have long studied generative AI techniques, new research brings renewed opportunity to explore the technical approaches required to apply them to their fullest potential.

Harnessing the Good—and Anticipating the Bad

The promise of generative AI is undeniable. But enterprises have to navigate dynamic ethical implications, bias, and potential security vulnerabilities—and then implement mitigation strategies (see Figure 2). Risk management for generative AI is not intended to curb applications of the technology; rather, it empowers IT and mission leaders to deploy more use cases with confidence.

A comprehensive approach includes the following steps:

1. **Codify and clarify policies** that outline acceptable use cases, guidelines for content generation, and measures to ensure ethical and responsible use.
2. **Commit to transparency** by documenting limitations, biases, potential inaccuracies, and training data of the generative AI system.
3. **Deploy traceability mechanisms** to track the sources and origins of the generated content—for example, maintain records of the training data, algorithms used, and decision-making processes.
4. **Audit training data** to guard against biased or offensive content generation. Training data is crucial in shaping model behavior, so diverse, representative, and unbiased data can minimize risks associated with generating unwanted or inappropriate content.
5. **Extend security measures** to protect against potential vulnerabilities and attacks. This includes securing the underlying codebase, implementing secure data handling practices, and guarding against malicious uses of generative AI systems.

Figure 2: Mitigation strategies for seven critical risk areas when engaging with generative AI

RISK	MITIGATION STRATEGIES
Hallucinations or Confabulations *Generative AI may create output that seems factual but is not.*	• Encourage critical thinking and fact-checking of generated content. • Manually verify content by human reviewers. • Implement algorithms and filters to detect and flag potential hallucinations.
Harmful Content *Training data that includes internet data can generate negative, offensive, or damaging content.*	• Implement filters and content moderation mechanisms to prevent the output of harmful content. • Regularly monitor and review system-generated content to ensure its appropriateness.
Algorithmic Bias *Training data that includes biased data can perpetuate existing biases and create harm.*	• Conduct bias audits and continuously monitor the training data. • Implement techniques such as data augmentation, diverse training data, and fairness-aware training to mitigate bias.
Misinformation and Influence Operations *Generative AI can be used to create realistic yet fabricated content for disinformation campaigns.*	• Implement systems to detect and flag potential misinformation generated by AI. • Enhance media literacy among users to identify and challenge false information.
Intellectual Property (IP) *Generated content may violate someone else's intellectual property rights or lead to IP leakage.*	• Conduct thorough IP clearance checks and ensure compliance with copyright and trademark laws. • Implement data protection measures to prevent IP leakage by employees interacting with generative AI systems.
Privacy *Mishandling of personally identifiable information (PII) during input or output can lead to privacy breaches.*	• Implement strict data privacy protocols and ensure compliance with regulations, such as the General Data Protection Regulation (GDPR). • Conduct privacy impact assessments and implement safeguards to prevent the input and output of PII.
Cybersecurity *Outputs of generative AI systems can introduce security vulnerabilities, and bad actors can exploit these tools for malicious purposes.*	• Conduct security audits and penetration testing to identify and address vulnerabilities in generated content and code. • Implement robust cybersecurity measures to protect against potential attacks leveraging generative AI.

6. **Monitor and evaluate** the performance, outputs, and impacts of generative AI systems. Continuous evaluations and audits can help identify emerging risks, detect potential biases or inaccuracies, and inform necessary adjustments for responsible use.
7. **Educate employees** on the risks, limitations, and ethical considerations of generative AI. By investing in a culture of awareness and responsible use, the workforce can actively contribute to risk mitigation efforts.

INDUSTRY PERSPECTIVE

***Tim Lortz**, a generative AI expert at Databricks, shared his perspective with our team.*

"Organizations which enable and encourage their people to incorporate generative AI into their work will be best positioned for long-term success," he says. He pointed out that we also need to remember the human element—most of the world has only been familiar with generative AI for a few months, and the early adopters have tended to be engineering types. "As managers, analysts, and other generative AI consumers gain a better understanding of what generative AI can, and can't or shouldn't, do for them, the use cases they identify and pursue will shape immense changes in the workplace and for our expectations with technology in general. These use cases could be as simple as writing or coding assistants built into familiar software tools, or something much more complex."

We asked Tim a few specific questions about generative AI:

What advice can you offer IT leaders about exploring these capabilities? Getting started with generative AI has become relatively easy with the release of hosted generative AI models like ChatGPT and open source libraries like Transformers and LangChain to run generative AI applications locally. But going beyond demos to production applications is not for the faint of heart. It's important for leaders to navigate key challenges and questions going into any investment, such as: How are you going to use generative AI models in secure, sensitive environments? How will you integrate generative AI into a broader data governance strategy? How can you effectively stitch together services from disparate vendors into your legacy applications?

What are you most interested to see as this technology matures? Addressing the black-box nature of large language models is a crucial research area. Developing techniques to explain and interpret the decisions made by these models will be vital for their adoption in critical services. Also, it will be important to continue evolving open source models to more closely rival the generative AI capabilities of proprietary models like those of OpenAI. The gap is closing, but the difference in response quality is still noticeable. Can reinforcement learning from human feedback, for example, become easy to adopt for more organizations?

Lastly, I'm also interested in developments around efficient model inference. The past few months have yielded impressive results in terms of "small" models with competitive generative AI benchmark scores. Improvements to attention architectures and quantization methods have shown a lot of promise to make models require less computation and memory without loss of quality, and we can expect to see more evolution.

The Future of Generative AI in Government Missions

Generative AI is swiftly taking hold across diverse government missions, revolutionizing processes—from data classification to research support—and enabling agencies to traverse and analyze vast amounts of data.

Some of the most visible examples of generative AI are built to help the average user interact with the Federal Government. For example, large language models are enhancing citizen-facing chatbots, improving search functions, and personalizing service delivery (on page 58, read more about customer experience). But the use cases span diverse mission sets.

One compelling application in development is a ChatGPT-powered platform for wargaming. Wargames and what-if exercises often lack the resources required to create realistic and complex simulations of stakeholders, such as threat actors, regulators, journalists, customers, and partners. A traditional simulation of multiple stakeholders' responses and their interactions is time consuming and resource intensive. Additionally, the need for subject matter experts to play different roles can be impractical and costly, hindering the scalability and frequency of exercises.

To address these challenges, ChatGPT can be imbued with a long-term memory platform to craft intelligent personas capable of simulating various stakeholders involved in crisis management scenarios. These personas can dynamically respond to inputs, ask questions, make decisions, and interact with other personas and participants during both the planning and execution phases of an exercise. This new platform can turn static, open source data into a dynamically responsive persona capable of responding based on data and its assigned attributes. The project aims to increase efficiency and deliver a highly realistic crisis management exercise.

Another spotlight use case is in the space of infrastructure management. Generative AI is helping to create digital twins of complex infrastructure systems so that agencies can proactively monitor, optimize, and predict maintenance needs, ensuring the efficiency, resilience, and longevity of critical government infrastructure. In areas such as bridge engineering or aircraft design, generative AI can benefit engineers by providing access to real-time data and analytics in addition to faster prototyping and rapid design iterations.

As generative AI technology evolves, it will foster an era of cognitive augmentation where AI and human intelligence work together to accomplish tasks of an unprecedented scale and complexity. We see it expanding far beyond its present applications to impact areas such as:

ORGANIZATIONAL PRODUCTIVITY
By automating real-time data analysis of security reports, policy documents, and financial forecasts, generative AI can significantly enhance productivity. It eliminates bottlenecks, expedites decision making, and frees up crucial human hours.

DIGITAL ASSISTANTS
Powered by generative AI, digital assistants can learn from historical data, understand context, and generate humanlike text. These assistants can streamline communication, facilitate information flow within government agencies, provide real-time writing and proofreading assistance, summarize complex documents, and create meaningful content.

POLICYMAKING
Generative AI's capacity to simulate and generate diverse scenarios based on existing data makes it a valuable policymaking aid. It can predict potential impacts and outcomes of proposed policies, facilitating proactive decision making and aiding policymakers in drafting effective and comprehensive policies.

DISASTER MANAGEMENT
Generative AI can create hyperrealistic simulations of natural disasters, including complex climate models, helping to improve planning, resource allocation, and response strategies to minimize the impact of emergency events.

SMART CITIES AND ENGINEERING
By simulating traffic patterns, energy consumption, and population behaviors, generative AI could provide critical insights for planning urban infrastructure and public services to make our cities more efficient, livable, and sustainable.

IMMERSIVE TRAINING
Integrating generative AI with augmented reality and virtual reality marks a revolution in training and education within government agencies, providing immersive, personalized learning experiences.

CYBERSECURITY AND ANOMALY DETECTION
Beyond simulating threats, generative AI can help predict and neutralize threats in real time. By constantly learning and adapting to new threat vectors, generative AI will be a powerful ally in maintaining national cybersecurity.

There is an understandable wave of excitement and enthusiasm surrounding how generative AI technology will fundamentally shape how we live our lives, conduct business, and interact with organizations and with one another.

To expand the mission impact of generative AI, government agencies and enterprises will need to understand how to design, build, and implement the technology *and* the use cases best suited for it. As an industry, it will take a commitment to establish the norms, standards, and regulations necessary for long-term use of generative AI, but for now it remains crucial for organizations to invest in deepening their understanding of the technology and its capacity for both risk and possibility.

***Ted Edwards** and **Alison Smith** are AI technologists who specialize in the application of generative AI, helping Booz Allen and its clients harness the technology for mission impact.*

SPEED READ

Generative AI is transforming government missions by aiding in data analysis, communication, policy simulation, disaster management, and more. It is positioned to enhance productivity and decision making by facilitating collaboration between AI and human intelligence.

The promise of generative AI is undeniable. But to use it responsibly, enterprises will have to consider ethical implications, bias, and potential security vulnerabilities—and then understand and implement mitigation strategies.

To expand the mission impact of generative AI, government agencies and enterprises will need to understand how to design, build, and implement the technology and the use cases best suited for it.

Managing Edge AI Sprawl

SIMPLIFYING COMPLEXITY BEYOND THE ENTERPRISE

Brad Beaulieu, Beau Oliver, Josh Strosnider, and Rebecca Allegar
Contributions from Eric Syphard and Greg Kacprzynski

The evolution of mission operations far beyond enterprise boundaries—from remote work sites to the battlefield and outer space—continues to accelerate through edge computing technology.

For the Federal Government in particular, edge technology is enabling agencies to harness powerful analytics and AI in the field and address growing mission requirements to:

- Enhance resilience and redundancy beyond traditional IT architectures
- Improve situational awareness and decision making at the operator level
- Keep data at the point of processing, due to volume and security, rather than backhauling it to the enterprise

"As agencies generate massive amounts of content and data at the edge from things like sensors, cameras, drones, industrial machines, and healthcare equipment, the data must be processed as close to the point of origin for many applications to be effective," said Ramesh Kumar, head of product and solutions for AWS Snow services. "We believe that edge native applications for IoT, data and image analytics, and AI/ML will grow in prevalence across the tactical edge with magnified growth for mission solutions."

But as advanced edge AI capabilities multiply, so do the devices, systems, and resources needed to enable highly specialized mission sets. What should agencies consider when faced with this growing "edge sprawl"—and how should they adjust their design and management of AI models to reduce complexity from enterprise to edge and back? For many, the key lies in engineering backward from the mission and committing to open, modular architectures that support robust technical performance across enterprise-to-edge continuums.

Edge Sprawl, Dissected

Edge sprawl occurs when many mission-specific devices and systems operate independently in a fractured ecosystem that incrementally grows in size and diversity over time.

Undoubtedly, organizations need to build, manage, and train new models rapidly and effectively across the edge environment. But what can seem like quick wins can actually hinder the mission when a thoughtful plan isn't in place for operating the ecosystem as a whole.

Expansion of various mission applications for AI and edge technology leads to bespoke tools that agencies are unable to manage or scale for future system requirements. Within mission sets, each distinct activity may need its own set of highly tailored capabilities with unique form factors and hardware configurations. While increasingly critical, these single-use, single-mission solutions are a central cause of edge sprawl.

For some organizations, the myriad tools and platforms risk spiraling toward uncontrolled complexity that prevents the sharing of data and distribution of insights between edge and enterprise and among edge systems. Where edge sprawl persists, organizations face a series of interrelated risks that can diminish mission outcomes:

SECURITY BREACHES
Vulnerabilities can quickly move from edge to the larger network, such as with machine learning (ML) models that poisoned data has compromised.

INTEROPERABILITY CONSTRAINTS
Silos make data and model sharing inefficient or impossible, hindering decision making, scale, and affordability.

VISIBILITY LIMITS
Barriers emerge, impacting the ability to achieve situational awareness at the edge and maintain visibility across large-scale mission domains.

DECREASED FLEXIBILITY AND SCALABILITY
Mission platforms may lack the capacity to evolve in line with changing objectives.

RISING COSTS
Operations and maintenance costs increase, and agencies must cope with constant refresh cycles—that are neither planned nor designed—to maintain mission resilience.

Mitigating these risks and successfully managing edge sprawl may involve adjustments to how organizations plan, design, and manage their mission platforms. It's logical to seek out different solutions to achieve different goals, but, at root, edge ecosystems represent a decentralized fabric that gives agencies the ability to process information and train models without an exchange between devices and the central enterprise. In the push for mission-specialized devices and systems, the primary goal—to provide a secure, integrated foundation to help warfighters, civilian operators, and agencies make critical decisions faster in distributed or disconnected environments—should not be overlooked.

From Complex Sprawl to Singular Architecture

To counter the edge sprawl challenge, agencies should think about identifying a comprehensive, scalable approach that addresses essential aspects of their edge environments (see Figure 1).

The high demand for AI at the edge is promising because any time organizations can move AI closer to

Figure 1: A comprehensive, scalable edge environment

SECURITY ACROSS EDGE SYSTEMS	INTEROPERABILITY AMONG DEVICES	CONNECTION FROM ENTERPRISE TO EDGE AND BACK
Without a well-designed, open architecture, the applications and models operating at the edge are open to risks, including biased AI models, poisoning of data, and widening of the broader network's attack surface. The security challenge of backhauling data to the enterprise, whether from contested or uncontested environments, should not be underestimated. These significant risks require management and mitigation beginning at the enterprise and working down to the edge.	Agencies need to harness increasingly powerful AI models and operate them at the edge, with the latest example being large language models like ChatGPT. However, these models call for unusually intensive compute resources and distinct hardware tools from an array of manufacturers. A crucial level of interoperability can be lost without strategies and tools, such as runtime environments to erase processor-specific silos.	Deficiencies in security and interoperability between multiple edge devices sever connections between edge and enterprise, isolating insights to devices, systems, and locations where decision makers lack capacity to access and apply them. That places senior leaders and operators in the dark, unable to make sense of the information that is being collected and processed at an edge node but not shared past that point.

Agencies should be mindful of tradeoffs between optimized technical performance and critical mission needs, such as seamless data transfers to enable full-domain awareness.

AI AT THE EDGE

The Next Wave

Today's AI models are often trained with the intent that they will be running on hefty cloud-based compute infrastructure, with data feeds that are multiple, persistent, and dependable. Without many constraints, AI engineers employ techniques to train their models for finely tuned performance. And why not? If an organization is going to invest in AI, it should perform at peak levels. But the future promise of AI at the edge will be marked by innovation and progress in a few key areas, including:

Making tradeoffs between precision versus application. If the inference—the model runtime—is on fixed-compute, limited power devices (like a mobile phone or perched camera) in disconnected environments, perfection in performance need not be the primary objective. The potential value of AI-enabled decision support at the disconnected edge is so significant that a slight tradeoff in precision is often worth it if engineers can extend AI inference capability to low-compute and low-power environments.

Accelerating the training cycle in operations. While there is undoubtedly a future where complex AI models are trained locally on edge devices, this technology is not yet fully mature. It can often take weeks to aggregate field data on cloud infrastructure and retrain and redeploy an AI model that needs refinement as conditions or parameters change. This long cycle time can diminish mission impact and erode trust in the solution. There are capabilities emerging to address this; for example, Synthetaic has developed a computer vision tool that allows users to "nudge" AI models in the right direction to improve future results for detection or classification of objects. As more capabilities like this become available, AI at the edge can be trained at a fraction of the time and cost of traditional methods.

Innovating delivery pipelines. The success of development operations (DevOps) with continuous integration/continuous delivery (CI/CD) pipelines in traditional software must eventually apply to AI at the edge to drive operational success and improved adoption. While there is a trend toward increasingly powerful AI, like large language models born from massive computing resources, there is also a trend toward efficient AI runtimes and model delivery pipelines to realize true AI CI/CD. The evolution and convergence of both will be paramount to the next wave of AI.

the point of data generation, they can reliably create mission and technical advantages. But two stubborn challenges remain:

- First, every time a new model is needed, it can take weeks to build and deploy it—that's way too long.
- Second, in the move from the cloud to the edge, AI performance tends to fall apart—accuracy, latency, and power efficiency all suffer.

Siloed commercial technologies are unable to keep up with this demand, and it is becoming more critical for organizations to combine AI infrastructure, platforms, pipelines, and models into unified solutions that are flexible based on the specific mission. We refer to this concept as a "mission stack," which is a set of integrated, mission-optimized technologies that function holistically to accelerate production-ready capabilities and reduce the total cost of ownership for end users.

For AI at the edge, a mission stack approach (versus a device- or capability-specific approach) helps organizations design open, modular frameworks. These frameworks promote feature enhancement and device and application monitoring while creating flexibility to build and deploy new capabilities in a secure, streamlined fashion. This allows organizations to create, test, and train models or key applications in an enterprise environment, ensuring that security protocols are maintained. These capabilities can then transition downrange into local hubs or across distributed enterprise endpoints to increase accessibility of the models and features to the operators or devices in the field.

This approach brings together different hardware and software solutions that might otherwise struggle to connect, communicate, and create value in unison. For instance, to rapidly build and run AI that can be used to train and deploy highly reliable AI to the edge in minutes, a holistic solution requires ML technology to train models in the mission and in real time without sacrificing latency. Such a solution also requires model compression so that models don't sacrifice performance at the edge and can be integrated across diverse edge processors and devices.

An open, modular design for AI at the edge allows for compression and maintenance of the models as they move downrange (see Figure 2). Maintaining the models at the enterprise level enhances their security, accuracy, and performance for the

Figure 2: Singular Approach to Edge Architecture

An open architecture at the edge facilitates the aggregation of data to enhance a common mission. It enables interoperability between disparate point solutions for an integrated mission stack, pools resources and capacity for federated learning, and creates the critical connective tissue between the cloud and the tactical edge.

"As agencies generate massive amounts of content and data at the edge from things like sensors, cameras, drones, industrial machines, and healthcare equipment, the data must be processed as close to the point of origin for many applications to be effective. We believe that edge native applications for IoT, data and image analytics, and AI/ML will grow in prevalence across the tactical edge with magnified growth for mission solutions."

*—**Ramesh Kumar**, head of product and solutions for AWS Snow services*

community. And the ability to transition data, results, and performance metrics both up- and downrange is critical to continuously improving and informing the models. This approach creates the opportunities agencies need most to expand the edge environment while centrally managing and developing capabilities that can scale.

A Future in Synch

It's useful to think about efforts to address edge sprawl through a comparison to the evolution of personalized health devices over the past decade. While fitness enthusiasts and patients once were forced to juggle disparate devices, from special cameras to multiple apps, today's human performance users have uninterrupted access to synchronized dashboards that update in real time. This shift to interoperability and singular dashboards is enabled by standardized information sharing protocols, security requirements, medically informed dashboard requirements, and a user-first approach to designing the user experience. It allows users to seamlessly get real-time updates on their aggregated and analyzed personal health statistics in a single location—with analysis on performance, recommendations for improvement, and insights into predictions throughout their day—versus spending time connecting each application and device to their computer, then downloading and making sense of the data system by system.

From healthcare to disaster response to the battlefield, mission sets will increasingly require secure, interoperable edge ecosystems that enable faster, more reliable analytics and data processing at distributed endpoints. As discrete edge AI capabilities continue to advance, here's the critical opportunity for the enterprise: to embrace those emerging capabilities at scale and in synchronization within their existing cloud-to-edge infrastructure.

Brad Beaulieu *is an expert in cloud-to-edge architecture and security. He focuses on emerging cloud capabilities and development in Booz Allen's BrightLabs incubator, an experimentation organization designed to develop, test, and incubate mission-centric solutions rooted in emerging technology.*

Beau Oliver *leads Booz Allen's Technology Exploration unit within the CTO, overseeing collaboration across the innovation ecosystem to scout, collaborate, and invest in mission-ready technologies.*

Josh Strosnider *leads Booz Allen's partnership efforts as part of the CTO organization and helps drive strategies to enable, differentiate, and expand federal access to emerging technologies.*

Rebecca Allegar *oversees corporate collaborations between Booz Allen and industry partners.*

SPEED READ

As advanced edge AI capabilities multiply, so do the devices, systems, and resources needed to enable highly specialized mission sets. Expansion of various mission applications for AI and edge technology leads to bespoke tools that agencies are unable to manage or scale for future system requirements.

Edge sprawl occurs when many mission-specific devices and systems operate independently in a fractured ecosystem that incrementally grows in size and diversity over time.

It is becoming critical for organizations to combine edge AI infrastructure, platforms, pipelines, and models into unified solutions that are flexible based on the specific mission. The key is engineering from the mission backward and committing to open, modular architectures that support robust technical performance across enterprise-to-edge continuums.

EQUITY

AI for Everyone

4 PILLARS FOR FOSTERING A DIVERSE, EQUITABLE, AND INCLUSIVE AI-DRIVEN FUTURE

John Larson and Ramon Hill

Artificial intelligence will have profound implications for every aspect of society, from healthcare and transportation to finance and national security. Already academia, policymakers, technologists, pundits, enterprises, economists, and average citizens are contemplating and debating the many ways AI could or should be used.

To be sure, it's essential to do all we can to understand and address the real-world impact of AI deployment—beyond recommending your next streaming video or predicting your shopping preferences. AI is already assisting doctors in making better diagnoses, streamlining access to critical services, helping remediate climate challenges, and much more as described elsewhere in this issue of *Velocity*. However, a critical question of our time is how to harness AI's power for society's broader good, to make lives better. There is necessary public discussion about AI's capacity to exclude marginalized populations—but there is an equally resounding opportunity to use AI in fostering diversity, equity, and inclusion (DEI). Here are four spotlight areas where AI and DEI intersect to help promote the equitable treatment and full participation of people from across society.

"In a world where AI is rapidly reshaping every facet of our lives, it's imperative that its development and application are rooted in diversity, equity, and inclusion. We believe that the future of responsible AI lies in equipping every student, regardless of their background, with the knowledge and tools to navigate, influence, and innovate in this transformative landscape. **Cultivating diverse talent isn't just about filling seats—it's about ensuring that the AI systems of tomorrow are built with the collective wisdom of our entire society**, reflecting the richness of our shared experiences and values."

Alex Kotran
Co-Founder & CEO, aiEDU

Principles of DEI in the Age of AI

1 Foster AI Literacy for All

AI is arguably the single most transformative technology of a generation since the invention of electricity. Just as the Rural Electrification Act of 1936 facilitated and accelerated the provisioning of electricity and, ultimately, its economic benefits to all corners of the nation, we need to ensure that everyone in the U.S. has high-level AI literacy and thus can be part of the emerging economy.

AI is already touching virtually everyone daily as they search for jobs, apply for loans, access healthcare, vote, travel, select clothes, and more. Having a basic understanding of what AI is and the skills to use AI-driven tools and technologies is becoming just as essential as knowing how to read, write, and use a computer. Fostering AI literacy also involves educating people about the ethical implications of AI, including those related to bias, privacy, and transparency. As a nation, if we don't succeed in broad AI literacy that reaches people in underserved and underrepresented communities—or those in jobs that are traditional and nontechnical in nature—we risk increasing disparity and leaving people behind, unable to enjoy the economic benefits that AI will open for the country. That, in turn, would create incredible individual and community hardship, put the nation at a competitive disadvantage, and inhibit the creation of AI that represents our democratic values.

2 Empower the Future Leaders of AI

Our ability to build a diverse AI workforce and to harness the power of AI for societal good is predicated on developing the technical talent of the next generation without leaving anyone behind. Concerningly, much of the focus on students now is on how to prevent them from using AI tools to cheat on homework assignments.

The next generation of talent is critical to national security and to the U.S. maintaining a leadership position. It's clear that a broader talent pool is vital: On the graduate level, the number of American students in AI programs hasn't increased since 1990. While graduate programs are just one of many pathways to AI leadership, we know that more upcoming talent needs to be trained in AI. Providing early access to technical AI education for young people from diverse or historically underserved communities could spark an interest in AI as a career and expand the pool of future leaders and practitioners.

As a community, we must be intentional about bringing these technologies into educational settings and equitably unlocking student potential. For example, the AI Education Project (aiEDU) is a nonprofit organization that seeks to do just that, by developing AI literacy and training content that educators, schools, and school districts can use in classrooms.

3 Diversify AI Development Teams

AI leaders are responsible for ensuring that those who work in AI reflect the rich tapestry of our society. It's no secret that teams of individuals with diverse perspectives and demographic backgrounds bring forward richer, more innovative solutions and better business outcomes.

We believe that the same is true with AI. Consider the development of facial recognition algorithms around a decade ago. At that time, those working on the problem didn't consider that if the model was trained on historical data, it would perform best for lighter-skinned males and poorest for darker-skinned females—which is exactly what happened. However, when AI models and use cases are designed from the outset by a diverse talent base that brings a broad range of inputs and perspectives to the table, they are much more likely to be built for the benefit of all people—not just a subset of society.

4 Understand and Mitigate Technical Bias

AI algorithms directly reflect the data they are informed with and their parameters for learning. When models are fed biased data, such as erroneous, unrepresentative, or discriminatory information, those models can amplify or propagate the bias. With AI increasingly driving decisions that impact individuals' lives, algorithmic bias against subgroups of our society can reinforce and perpetuate discrimination and imbalanced power structures.

The mortgage lending industry exemplifies how AI bias can impact lives and livelihoods. In the past, minorities were not able to get mortgages at the same rate as nonminorities due to discriminatory lending practices. If that historical data is used to train an AI model, that model would likewise approve mortgages at a higher rate for nonminority applicants than for minority applicants—perpetuating and amplifying this historic inequity.

Although it's difficult to know if a data set is biased until certain groups are affected, AI leaders must train talent to understand those types of systemic biases in data to minimize them early and often throughout the modeling process and guard against biased outcomes. For example, employing quality control strategies helps find biases during the development cycle, response bias sampling corrects issues in the dataset by oversampling specific populations, and generating synthetic data helps represent diverse populations when actual data doesn't exist.

Optimism for the Future

AI is going to be transformative for our society, for our country, and for the world. It will help us solve the most challenging problems we face—but only if we ensure that everyone has an equitable opportunity to be part of the AI revolution.

This is just a snapshot of an important topic. Tune in to Booz Allen's Unstoppable Together podcast to learn more about the intersection of AI and DEI, and explore other DEI topics and challenges facing today's workforce.

A WORKFORCE DISRUPTED

EMPOWERING EMPLOYEES TO EMBRACE AI, FROM ADAPTATION TO ADOPTION

Betty Thompson, Joe Rohner, Julie McPherson, and Logan Gibson

Change is an undeniable aspect of human existence. Throughout the annals of history, humans have consistently demonstrated an extraordinary capacity to adapt and evolve in response to the relentless march of technological progress and societal transformation. From the epochal Industrial Revolution to the advent of the internet and the more recent digital revolution, humanity has exemplified an exceptional ability to harness innovation and reshape the very fabric of the workforce.

As we venture into the era of AI—with generative AI bringing the conversation mainstream—we find ourselves standing at the precipice of yet another paradigm-shifting transition. Researchers from OpenAI, OpenResearch, and the University of Pennsylvania argue that:

80%
of the workforce could have at least **10%** of their work **affected** by the implementation of large language models.

20%
of the workforce could see an impact to **at least half** of their work.

15%
of all work tasks in the U.S. "could be completed **significantly faster** at the same level of quality."

Naturally, there will be apprehension regarding job displacement, along with growing concerns and anxiety surrounding the potentially deepening digital divide. Society will need to care for those impacted by the real transition costs of AI in the short term, despite the broader economic predictions. As Marc Andreessen highlights in his article "Why AI Will Save the World," society has been through two technology-driven unemployment panic cycles in our recent past—the outsourcing panic of the 2000s and the automation panic of the 2010s. However, despite dialogue suggesting unemployment was inevitable, by late 2019 the world had more jobs at higher wages than ever in history.

Against a backdrop of trepidation and understandable unease, an undercurrent of excitement is emerging, stemming from the remarkable possibilities that the accelerated and monumental progress of AI holds for global industries of every kind—particularly for the Federal Government as a catalyst for national security and to reimagine citizen services.

Against a backdrop of trepidation and understandable unease, an undercurrent of excitement is emerging, stemming from the remarkable possibilities that the accelerated and monumental progress of AI holds for global industries of every kind—particularly for the Federal Government as a catalyst for national security and to reimagine citizen services.

In response to the national and global AI race, this article explores the profound impact of AI on the workforce and approaches to fostering a digitally proficient enterprise with broad access to AI adoption.

National Urgency for AI Talent

The National Security Commission on Artificial Intelligence (NSCAI) report delivers an uncomfortable message regarding national security and AI: "America is not prepared to defend or compete in the AI era. … While much remains to be learned about the power and limits of AI technologies, big decisions need to be made now to accelerate AI innovation to benefit the [U.S.] and to defend against the malign uses of AI."

AI is transforming the global economy and has become a key driver of economic growth and military capabilities. Nations that excel in AI development and deployment will secure substantial advantages in economic productivity, defense capabilities, and intelligence gathering. Consequently, the report emphasizes the urgent need for the U.S. to compete and win the global competition for AI talent.

Despite the exponential growth of AI as a field, the number of domestic-born students pursuing AI doctorates in the U.S. has remained stagnant since 1990. This stagnation raises concerns about the nation's ability to meet the demand for skilled AI professionals and threatens its position as a leader.

While the NSCAI report projects a shortage of 250,000 data scientists by 2025 in the U.S., the 2023 National Artificial Intelligence Research and Development Strategic Plan report projects computer and information science positions will grow by 22% between 2020 and 2030. According to this report, "the resulting economic growth could be large: AI research is expected to contribute as much as $11.5 trillion in cumulative growth across G20 countries alone over the same period."

EMBRACING AI IN THE FEDERAL WORKFORCE:
A New Era of Productivity and Efficiency

*Conversation with **Johnny C. Taylor Jr.**, CEO of the Society for HR Management (SHRM)*

The SHRM Omnibus Survey reveals a significant portion of working Americans believe AI can enhance their performance and make their jobs better. Additionally, recent SHRM research found half of U.S. workers believe AI will improve their workplaces, trusting their organizations to effectively train them to utilize new technologies. Still, 23% of U.S. workers fear job displacement due to automation in the next five years.

Johnny C. Taylor Jr., CEO of SHRM, emphasizes the "importance of organizations communicating frequently with employees about how AI will be used in an organization, reassuring them the purpose of integrating new technologies is not to displace workers."

The rapid evolution of AI in the workplace is undeniable—and the trend toward automation and AI in recruitment, hiring, and performance management is on the rise. In fact, as of 2022, nearly a quarter of organizations have already integrated AI to support HR activities, resulting in faster hiring processes.

"For the federal workforce, the future hinges on proper training, staying updated with AI legislation, and embracing the synergy of AI and human intelligence (HI)," Taylor says. With training and education, organizations can actively engage with—and thereby address—emerging issues such as hallucination, intellectual property concerns, biases in algorithms, and ever-expanding use cases.

Ultimately, to remain competitive in talent acquisition, Taylor reinforces that federal leaders should embrace AI instead of avoiding it. "The mindset shift needs to start at the top."

The CHIPS and Science Act of 2022 is reinforcing an effort to develop proper data, knowledge and skills, and workforce curation and analysis techniques by authorizing the National Science Foundation to generate a study of U.S. universities that conduct high-impact AI research to better understand what factors enable AI progress. There is still much to be learned, but these reports emphasize the need for the U.S. to focus on talent development across the public sector.

There is an urgency to maintain the U.S.'s position in the global AI talent race, despite apprehensions within the workforce. Therefore, a pivotal question emerges: How can we address apprehensions and equip a larger talent pool to shape this transformation?

Demystifying AI in the Workforce

AI is finally unlocking opportunities for many organizations and agencies to "do more with less." Undoubtedly, the downstream effect is that AI will disrupt and eliminate certain job functions, as is typical with any emerging technology that offers automation capabilities.

Within the talent ranks, whether in technical or nontechnical roles, this proposition understandably sparks concern and anxiety. But let us be clear: AI can't automate or solve everything. Rather, AI can automate pieces of most jobs, resulting in streamlined processes, increased throughput, and automation of repetitive operations. In practice, this human-machine teaming "unbundles" tasks so repetitive assignments (the monotony of which often leads to mistakes) are offloaded and humans can reserve their time and energy for high-order critical thinking and problem-solving.

AI can't automate or solve everything. Rather, AI can automate pieces of most jobs, resulting in streamlined processes, increased throughput, and automation of repetitive operations.

Realistically, workforce transformation in the age of AI does not directly equal job transformation. Instead of broad-brush changes, there needs to be a nuanced lens applied at the task level to evaluate where AI can serve as an aid and where traditional skillsets need to be preserved. In the study cited earlier from OpenAI, OpenResearch, and the University of Pennsylvania, researchers reinforce the idea that "each job is

a bundle of tasks" and note that when looking at the overlap between AI capabilities and specific job tasks, "it would be rare to find any occupation for which AI tools could do nearly all of the work."

Therefore, a prevailing sentiment is resonating across industries: AI will not directly replace your job, but someone using AI will. Successful bundling of newly packaged jobs assumes that those affected have a basic understanding and awareness of AI. By empowering more corners of the workforce with necessary AI skills, organizations can foster a culture of innovation and productivity, transform roles, and enable people to work in tandem with AI technologies to achieve better outcomes.

Scaling AI Awareness to Curb the Digital Divide

On the heels of a digital revolution, where parts of society are already experiencing disenfranchisement, the advent of AI further compounds these effects. This holds true in the workforce microcosm, as well. There is a noticeable tendency to centralize AI organizationally, but such centralization may exacerbate existing disparities. Moreover, the transformation potential of AI extends far beyond specialized roles like data scientists and machine learning engineers; it envelopes the broader workforce, as more diverse jobs start to enhance their results and outputs through human-machine teaming. Striking the right balance between centralization and access is crucial to avoid exclusive ownership of AI tools by a select few, as this could be the first sign of failure in fostering an inclusive AI-driven future.

Considering the significance of AI talent in every job, at every level, and in every capacity, there is a need for new ways to unleash talent and expand the pool. Algorithmic and mathematical outputs must be enhanced and enriched by subject matter expertise, further demanding a workforce that has a broad foundation of AI literacy (or sufficient awareness) to interpret results, pursue intended objectives, and achieve a decision-making advantage.

The key for leaders today lies in creating opportunities for the wider talent pool to embrace this new paradigm. Specifically, reskilling and upskilling programs ensure that individuals remain relevant in the workforce and position federal organizations to unlock the full potential of AI. According to the World Economic Forum's Future of Jobs Report 2023, such training programs are emerging quickly:

- AI and big data are the **No. 3** priority in company training strategies from now until **2027**, and the **No. 1** priority for companies with more than **50,000** employees.
- AI and big data training programs will constitute more than **40%** of technology-focused training in companies across the U.S., China, Brazil, and Indonesia over the next five years.

Of course, not every person or role needs the same kind of training. When designing or implementing a purposeful AI training program, leadership should consider different tracks to segment the diversity of user personas. For example, the general workforce may need to start with basic AI literacy, and data scientists and engineers who are hands-on with AI solutions and workflows will need advanced and applied technical education. There is much that can be done with training programs and creating broader access to AI, but this transformation doesn't happen overnight. AI modernization is a journey.

So, what does this mean for federal organizations?

Forging a Path Forward: AI as a Crucible for Progress

Progress in the field of AI heralds notable shifts in military, informational, and economic supremacy, empowering new capabilities while making existing ones more accessible to a wider range of actors. We are still at the early stages of the AI evolution in government and across industry. The near-term focus needs to be on engaging and preparing those who are impacted by this transformation and critical to its future trajectory: the federal workforce today.

Here are three priority areas to consider:

1 INVESTING IN A CULTURE OF ADAPTABILITY.

In the World Economic Forum's Future of Jobs Report 2023, organizations are ranking "self-efficacy skills"—pertaining to resilience, flexibility, and adaptability—as core skills to navigate today's workforce (ranked above technological literacy). Creating a culture of continuous learning, where employees are encouraged to upskill and reinvent themselves through training programs enables talent to adapt—and adopt AI—more quickly as job requirements evolve. Agencies can establish AI-focused training programs, mentorship initiatives, and career pathways that *meet employees where they are* and help them chart their own trajectories to remain competitive in the rapidly changing job market.

2 ESTABLISHING A STRATEGY FOR BROAD ACCESS TO AI.

Innovation thrives when a broad range of perspectives and experiences are involved. Achieving AI proficiency across the workforce will require a look at diversity, equity, and inclusion (DEI). This goes beyond ensuring DEI in AI hiring practices. It extends to creating an inclusive environment where AI isn't walled off in an agency but is accessible to a diverse range of talent and experience. Democratizing access to AI knowledge and skills in this way will allow federal agencies to incorporate diverse perspectives in AI design and development—critical to avoiding algorithmic bias, promoting equitable outcomes, and reaching into vast pools of untapped potential talent to create solutions for citizen services and national security. For more on this topic, see page 28.

3 HARNESSING PARTNERSHIPS TO AUGMENT THE WORKFORCE.

Agency leaders don't have to build an AI workforce in a vacuum. If the current workforce is primarily focused on upskilling and training, federal agencies can embrace collaboration and partnerships as formidable means to augment their workforce and leverage external AI expertise. In our rapidly evolving AI landscape, no single entity possesses all the necessary skills and insights. Pursuing partnerships with academic institutions, private sector companies, and research organizations empowers agencies to tap into a rich network of AI talent, knowledge, and resources. By fostering a culture of collaboration, agencies can harness the collective intelligence of diverse stakeholders and power AI-driven initiatives that have far-reaching impacts.

Agency leaders don't need to have all the answers right now as to how the workforce will weather these changes. But here's what we do know: While AI is undeniably causing disruption to the workforce, if approached responsibly and incrementally, this disruption can be harnessed as an opportunity for inclusive growth and progress. By embracing the transformative power of AI, we create opportunities to expand our skills, enlighten future generations, and scale our capabilities at an unprecedented pace.

AI transformation is going to challenge all of us to reimagine our roles, fostering a mindset of continuous learning and adaptability. With the right strategies and investments, organizations can navigate this disruption and position their workforce at the forefront of the AI-driven future. With commitment, we can not only weather the storm of disruption but also harness its potential to propel us toward a future of expanded possibilities and meaningful impact.

Betty Thompson *is Booz Allen's chief people officer.*

Joe Rohner *leads AI talent and delivery to clients in the western U.S. and the Pacific, including Hawaii, Japan, and Korea—and helped develop Booz Allen's AI workforce training program.*

Julie McPherson *leads Booz Allen's Solutions Center business to drive adoption and reuse of the company's technical baselines and solutions, and is a passionate leader of the firm's DEI Council and technical experience groups.*

Logan Gibson *is chief of staff in Booz Allen's chief technology office and oversees the company's technical experience groups.*

SPEED READ

The transformative power of AI is reshaping the workforce, presenting both disruptions and opportunities.

AI has the potential to enhance various roles across the entire workforce through human-machine teaming, requiring a foundation of AI literacy.

Federal agencies can invest in a culture of adaptability, establish strategies for broad access to AI, and harness partnerships to augment their workforce and navigate the power of AI for inclusive growth and progress.

The Age of Principled AI

EXPLORING RISK, RESPONSIBILITY, AND POSSIBILITY

John Larson and Geoff Schaefer
Contributions from Megan Smith-Branch

Recent news cycles have fluctuated between two extremes: the transformative opportunities of AI across industries from scientific discovery to education; or the possibility for existential risk to society and our way of life.

In either case, industry analysts project AI investments to grow at near light speed—with global spending projected to increase at a compound annual growth rate of 27% from 2021 to 2026, according to IDC. China is expected to more than double its annual investments by 2026, while European AI expenditures, roughly 20% of the current world market, are expected to reach $70 billion in the next three years.

The train has left the proverbial station, and with AI's potential for such sweeping, positive impacts, calls for pausing or ceasing its development seem not only untenable but undesirable. Critical industries are seeing the potential to unravel long-puzzling mysteries of physics and predicting potentially far-reaching developments in national defense, infrastructure, precision medicine, and beyond. Ultimately, the future of AI as a transformative partner demands a return to an ancient question from classical philosophy: **What can we do to flourish?**

Beyond Ethics: Defining Responsible AI

Aristotle famously focused on a principle he called eudaimonia, the fusion of human happiness and flourishing that signified a life well lived. Centuries later, utilitarianism held that morality stems from producing the greatest good for the greatest number of people. As Aristotle might say if he were around: The future of AI is bright *if* we make specific choices to harness its power for the benefit of civilization.

Society is already recognizing that AI is poised to become the most impactful technology of our lifetime, in ways not yet fully imagined. Glimpses of what's possible can be found in stories such as DeepMind's AlphaFold solution to the "gene-folding" problem that had been impenetrable to scientists for more than 50 years. Using AI, AlphaFold was able to "fold" and predict the structure of nearly all proteins known to science—a landmark accomplishment set to fundamentally innovate disease detection and novel drug development.

As we look to the future of AI innovation, decision-makers and developers who are asking, "Is AI right or wrong?" should instead use their resources to advance tools that reinforce Responsible AI and embrace a mindset that focuses on harnessing the net good that AI systems will generate over their lifespan.

This field of Responsible AI is still in its earliest, formative days, but it has been in discussion for some time, from NIST's Responsible AI Framework to the Department of Defense's Ethical Principles for AI. The Federal Government is a steadfast leader in the pursuit of policies and protocols, with already established high-profile directives for agency implementation and White House initiatives to bolster investments in Responsible AI to protect the American public. Lawmakers are moving forward cautiously with new legislation—focused in the short term on understanding the complexities of the technology before proposing regulations. Meanwhile, there has been a profusion of guidance from industry and academic institutions to help evaluate the ethical dimensions of AI systems.

But beyond matters of ethics, what is Responsible AI?

A comprehensive, pragmatic definition of Responsible AI includes three distinct domains, each with its own body of research and technical considerations:

SAFETY:
The ability to adequately control an AI system's behavior, use, and functionality in the field—and remain safe against nefarious actors.

GOVERNANCE:
The alignment of relevant policies, laws, executive orders, or other regulatory guidance in a programmatic fashion.

ETHICS:
The assessment of risks and values of an AI system's deployment that considers net positive gains versus potential harms—and the decision to use that system or not.

Many of the principles and frameworks for using AI responsibly are left to platitudes and lofty language, lacking tangible tools or clear-cut, actionable approaches. The above domains, taken together, provide a roadmap and broad set of practices and operational parameters for Responsible AI, but necessitate that leaders start to test tangible solutions.

A Practice of Solutions—Not Barriers

From transforming the provision and costs of healthcare, to securing our nation's most critical infrastructure and augmenting our warfighters' capabilities, the risk is too great for missions of national priority to get by without AI. In the future, organizations that don't effectively deploy these emerging capabilities will be vastly constrained.

THEIR **WORLD**
WILL GET SMALLER.

THEIR **IMPACT**
WILL BE RESTRICTED.

THEIR **DECISIONS**
WILL LACK AN EDGE.

Why? Their world will get smaller because they'll no longer be able to keep pace with the exponential growth of information expanding around them. Their impact will be restricted because modern global challenges will not be a fit for old solutions. And their decisions will lack an edge because speed, precision, and information superiority are the markers of competitive advantage.

Because of this reality, the practice of Responsible AI is not (and cannot be) about establishing a risk police as a barrier to innovation, or a mechanism to tell teams "No!" or "Start over!" Rather, it needs to be a solution-oriented practice that enables and empowers AI deployment in as many critical missions as possible. Its core purpose is to help navigate risks and challenges before AI systems are set free in the wild, so that organizations can, ultimately, extend use cases and accelerate the technology boldly and with confidence.

In short, Responsible AI is a discipline that weighs risk against possibility. It positions leaders to consider real-world opportunity costs and put decision-making tools into the hands of those who can evaluate fundamental questions:

- What is the **maximal good** that AI could do for a particular program and mission?
- What are the **potential risks** of using AI in this way?
- How can we **mitigate these risks** so that the maximal good can be achieved?

The goal: deploying AI systems into the field without crossing any fingers, but with high confidence that the systems are safe, transparent, ethically sound, and conform to our societal and democratic values.

Of course, efforts to evaluate the opportunity cost of an AI system must calibrate to the mission context, specific objectives, and potential impacts—good and bad. In many cases, the situational context is straightforward, such as scenarios in which using AI can result in life-or-death outcomes (e.g., combat scenarios) versus scenarios in which AI helps drive more effective operations and management (e.g., customer service). In other cases, the risk is dependent on sociotechnical factors, such as bias in the AI system's training data. But even bias itself is complicated, with multiple nuanced dimensions. For instance, imagine a physician uses an AI system to help diagnose a veteran's medical condition. The patient is from an underrepresented demographic and has a complicated condition. In this imagined scenario, historical discrimination impacts the quality of the data the AI was trained on, *and* the AI simply hasn't seen enough cases of this particular condition and its variations to accurately evaluate it. The combined result is a "biased" diagnosis of the patient—for classically prejudicial reasons and because of the rarity of the disease itself. Both constitute bias, but only one is recognized as such.

Zooming out, these issues of bias will manifest in all types of AI systems and applications. But the healthcare scenarios above demonstrate how using AI in certain daily situations can pose more risk and personal harm to a much bigger cross-section of society than an AI system powering satellites and drones.

The power of Responsible AI in practice is that it codifies the language through which teams are discussing and contemplating this kind of risk calculation. Standardizing the lens through which AI is viewed makes it possible to measure that risk within a broader AI portfolio while calibrating it to the specific mission context. It opens the door to practical actions and creative solutions to mitigate anticipated risks and let AI get to work.

While going from principles to operations is hard, federal agencies and other organizations don't need an AI ethicist on staff or an extensive suite of capabilities in place to put a Responsible AI approach into practice. And they certainly don't have to hold back AI development. However, they do need a well-designed framework and a purposeful process to navigate trade-offs and opportunity costs. Robust technical guardrails and engineering practices empower leaders to evaluate the use of AI across increasing lines of work and expanding mission sets—easing significant pressure and risk that would otherwise prevent mission-critical innovation.

"As we look to the future of the field of AI ethics, an ideal outcome would be that the phrase *Responsible AI* becomes redundant."

Navigating Risk in Emerging, Uncharted Spaces

Without a doubt, focusing on the risks inherent in AI technology is necessary and urgent. Questions surrounding issues such as data privacy protection, bias, disinformation delivery, lack of model transparency, and unintended behaviors within AI must be acknowledged and addressed.

But let's ensure that risk is balanced appropriately within the context of the possibilities and power of AI as it continues to mature and systems have access to more data. Already, some large language models are demonstrating the capacity to self-correct issues related to disinformation delivery and bias. In short order, these risks—and others like them—may become self-contained as AI systems advance.

Moreover, drawing a page from classical ethics, one could argue that society has a moral responsibility to shepherd AI to its fullest potential. To do that, organizations and AI engineers will need to navigate uncharted territory on the path to Responsible AI development. Ongoing research is critical to address areas that include:

COMPLEX EMERGENT BEHAVIOR.
How to design controls and assess risk for unanticipated AI actions or outputs when it's impossible to predict what the AI system will do.

REINFORCEMENT LEARNING.
How to specify parameters and avoid AI misalignment when AI systems pursue the intended goal but employ unanticipated or unintended methods to do so.

HUMAN-MACHINE TEAMING.
How to design controls to mitigate unpredictable human behavior when augmented by AI and maintain human trust in their own operating skills and expertise.

As these capabilities emerge and become more sophisticated, there will be more questions than answers along the AI journey. Therefore, a commitment to Responsible AI requires continuous optimization, with humans in the loop. For instance, when using an AI system to inform military pilots about intended targets, the humans involved need proper training and experience to know when to trust themselves over the machine, should a situation demand it.

Of course, there are scenarios and critical missions today that are too high-risk to employ AI responsibly, particularly with reinforcement learning, where the machine may act in limitless ways to achieve a goal. Ultimately, complex emergent behavior has the capacity for both risk *and* reward.

However, stopping or slowing AI capability development across the board is not a shortcut to safety. Instead of admiring the problem, the best way to discover risks and solutions is by maximally engaging with the technology in sandbox environments. In those contained environments, teams can demonstrate and address unexpected behaviors—and then move the best designs forward into controlled, real-world applications to monitor deployments, learning and refining over time. At that point, things *will* go wrong. But what matters is how organizations mitigate the chances of material risk and how they minimize the impact, learn, and correct for the next time.

Engineering a Responsible AI Future

As we look to the future of the field of AI ethics, an ideal outcome would be that the phrase *Responsible AI* becomes redundant. The foundational principles of a responsible approach—encompassing safety, governance, and ethical considerations in AI deployment—will simply evolve as the standard for how *all* AI work is done. And in this future, the work will transcend compliance or risk, becoming an integral component of machine learning operations (MLOps) platforms and advanced, transformative AI tradecraft.

While today's AI systems can be overhyped, tomorrow's AI is almost certainly underestimated in terms of capability and scale. When deployed responsibly, AI stands as an enduring catalyst, opening new doors of possibility in medicine, public policy, and everyday life. AI will soon integrate data and domains in ways that have never before been integrated or correlated. New technologies will enable physical sensors and communications links in disconnected and denied environments. And as these systems, domains, and effects advance, AI may become the key to securing our nation against the predominant threats of this century.

Embracing a new calculus that computes the risks *and* rewards present in AI systems puts federal agencies and organizations in a better position to harness AI's potential for transformational good. Here's where you, the reader, come in. Whether you are an AI engineer, a policy maker, a change management expert, or a human that engages with AI in any capacity, you are a practitioner of Responsible AI. Your journey starts now.

Credo AI
INDUSTRY PERSPECTIVE

Credo AI *is a leader in embedding Responsible AI and regulatory requirements into operations and the technical development process. The company's CEO and founder, Navrina Singh, shared her perspective on this rapidly shifting priority for technology and mission leaders alike.*

There's so much hype today around AI. What's one of the biggest questions you get from organizational leadership? By far, the most frequent questions we come across are related to the pace at which organizations should be moving to adopt AI and AI governance: Should they pause, wait and see, or move ahead at full speed? In research we recently did with customers and industry, we found that organizations with a lack of expertise in generative AI combined with concerns over security, privacy and intellectual property are by and large taking a "wait, review, and test" approach toward generative AI governance. We've also seen very public examples of organizations in the private sector banning these tools for their employees altogether. I strongly believe that the emphasis should not be on halting AI development or innovation or halting the use of AI in one company or the other, but rather on catching up as a society to invest; research; and build governance, oversight, and alignment over the remarkable capabilities that have emerged.

What dynamic areas of research are you most excited about? On the technical research front, the industry is still in the early innings of understanding these complex AI systems, especially foundation models. There is an active body of technical research trying to make sense of the unexpected capabilities, deployment safety, and proliferation problems that come with these frontier AI systems. I have been critically thinking about these systems from an application perspective—exploring how humans and AI systems can best work together, the environmental impact of training large models and running data centers, the safety and long-term consequences of artificial general intelligence (AGI) systems, and how AI can be used to address pressing issues facing humanity like poverty, healthcare, climate change—all of which are still in the early research phase.

As you look to the future, what does the world of Responsible AI look like? A year from now my hope is that Responsible AI will emerge as the cornerstone of all AI. Every person that touches AI in some capacity will have a role in accountability, transparency, and fairness embedded in AI systems and organizations. With that, I hope to see continued public commitments to Responsible AI practices by organizations in and outside of the tech industry. What I mean by that is organizations publicly disclosing R&D and investments in AI safety and governance, disclosures around AI systems, and impact assessments of AI applications. As an industry, we can only be held accountable to outcomes if we know and are transparent about what those outcomes are.

John Larson *leads Booz Allen's AI practice, with a focus on ensuring leaders across federal missions achieve AI understanding, purpose-built solutions, and accelerated adoption.*

Geoff Schaefer *serves as chief AI ethics advisor at Booz Allen, working to make AI ethics and safety more practical, tangible, and measurable for clients and within the organization.*

SPEED READ

Industry's AI investments are surging, with particular emphasis on responsible development to unravel scientific mysteries and bolster national defense, infrastructure, and medicine.

Responsible AI is not about stifling innovation or imposing barriers but about enabling and empowering AI deployment for critical missions. It involves weighing the risks against the possibilities, considering the maximum good AI can achieve for a program or mission while mitigating potential risks.

The intent for Responsible AI is that it becomes a foundational norm rather than a catchphrase. Safety, governance, and ethical considerations in AI will evolve into standard practices across all AI endeavors. Tomorrow's AI, when deployed responsibly, will lead to innovation and become crucial to safeguarding against predominant threats.

CIO CORNER

Navigating IT Through AI Adoption

Q&A with Brad Stone, Chief Information Officer and Chief Data Officer, Booz Allen

Although large language models aren't new, the scale, wide availability, and imaginative use cases of tools like ChatGPT are unprecedented. Along with novel concerns, possibilities are emerging—in areas ranging from knowledge management to software development and ticketing—with organizations rapidly positioning to invest in and advance their capabilities. As IT leaders explore this new territory, we interviewed Booz Allen CIO and Chief Data Officer Brad Stone to learn how he's navigating this new era of AI-related uncertainty and optimism.

Q **Given the intense media coverage of large language models in recent months, there is a lot of anxiety—and excitement—around AI advancement. How are you approaching this landscape?**

A **Brad:** For some CIOs, the answer to that anxiety is to take a strategic "pause" before fully investing in generative AI, and that's understandable. We're taking a different approach. We've decided to embrace the opportunities that large language models present while, of course, being proactive in managing the high risks they carry. Beyond the hype, my focus is on using generative AI with purpose, in alignment with business objectives, and with a full understanding of the security requirements to safeguard sensitive information. In our case, it's not just our internal information that we're protecting—it's entrusted information that allows us to support many of the most important client missions across the nation.

Generative AI gives us access to new productivity and innovative capabilities, and I expect it to provide additional value as it unlocks our collective ingenuity. But it's important to consider the specific business requirements for using it. We consider the **enterprise team view** (How are we using it for corporate needs?), the **broader user view** (How are our employees using it?), and the **client mission view** (How are we helping our clients move their missions forward with it?). When addressing any of these areas, I strongly believe that AI is about accelerating and improving the human role in the mission. We're committed to strengthening trust in AI by keeping human beings at the center of the strategy going forward.

Q **Enterprise IT organizations are being asked to integrate generative AI and large language models into business operations—but there are some foundational steps that come first. What building blocks for enterprise "readiness" do you prioritize?**

A **Brad:** AI models are obviously only as good as the data they are trained on. For that reason, a critical building block is to have a comprehensive data management strategy in place. An integrated and effective data strategy helps an enterprise understand what data exists and where it is—and provides the ability to track

***Brad Stone**, Chief Information Officer and Chief Data Officer, Booz Allen*

the owners, curation, and quality of that data so it can efficiently and successfully train and interact with the models. With a sound data management strategy, it becomes possible to navigate key security, privacy, and regulatory requirements and avoid data exposure and data loss while still embracing these large language models. Ultimately, it's important to keep in mind that the investment to build a large language model from the ground up is likely to be prohibitive. In many cases, organizations can rapidly adopt and build capabilities on top of open-source or software-as-a-service solutions. And again, that begins with having a strong data management strategy, with data protection built in.

We have lived this firsthand with our IT organization's migration to a business systems data lake, which was foundational to our journey to modern data management and data protection. As I look forward, it will be especially important to further mature these concepts as generative AI opportunities are introduced.

Q **The technology industry is quickly adopting generative AI as an integrated piece of as-a-service products, many of which are already in use across organizations. How should IT leaders continuously evaluate their portfolio of industry tools without curbing employee access to valuable innovation?**

A **Brad:** So many vendors now offer generative AI capabilities that it can be difficult to prioritize service delivery opportunities within an individual enterprise. For this reason, it's important for organizations to invest and innovate based on what their users truly need and what their larger data problems and essential use cases actually are. They can then sequence efforts with industry partners to collaborate on prototyping, testing, and piloting potential tools. This prioritized approach with multiple pilots allows organizations to develop and quickly refine governance processes for these third-party solutions and respond to evolving risks or changing regulations.

We're focused on giving our employees access to AI tools wherever they are needed to fulfill their mission with the right guardrails, guidance, and support. Initial areas of focus for us will be using large language models, natural language processing, and information retrieval to enhance IT support with chatbots. Our organization is working with key industry partners to pilot as-a-service capabilities that include AI-enabled knowledge management solutions and accelerated IT ticket resolution—all with the goal of improving employee self-service and automating employee-focused transactions. As a key design principle, we are pursuing goals regarding the standardization and modularity of these technology stacks and are intent on leveraging the power of generative AI tools while ensuring that restricted, proprietary, and other entrusted information stays compartmentalized within trusted environments.

> "We've decided to embrace the opportunities that large language models present while, of course, being proactive in managing the high risks they carry. Beyond the hype, my focus is on using generative AI with purpose, in alignment with business objectives, and with a full understanding of the security requirements to safeguard sensitive information."

Q **As IT organizations think about how to proceed with the latest AI technologies, what advice would your team share?**

A **Brad:** We are following the concept of "think big but start small," and we advise other enterprises to follow this approach as they take on emerging AI capabilities. Making use of sandbox environments to test new use cases in a safe, controlled setting and then pursuing pilot programs is a good example of this concept. Participating in industry information-sharing groups, starting vendor AI capability discussions, and building acceptable usage of available services into training are other examples.

In my experience, some examples of this approach include stacking models to optimize resource management, accelerating performance by using AI for vulnerability scanning and remediation, harnessing AI-powered paired programming capabilities to enhance employee productivity, implementing conversational AI capabilities that connect into our broader IT ecosystem, and accelerating task execution to significantly reduce time to completion across the enterprise. Ultimately, starting small with AI technology will enable enterprises to fail fast and understand the impact of technical decisions before they affect real-world processes.

Booz Allen team working at an engineering lab in Panama City, FL

PERFORMANCE

UNLOCKING HUMAN POTENTIAL

Personalized intelligence enhances cognitive and mission performance

Munjeet Singh, Cameron Mayer, and Dave Prakash
Contributions from Irik Johnson, Andie Rauta, and Sonya Rahmani

Wade Spann, a Marine corporal from Virginia, arrived at the 2023 Department of Defense Warrior Games Challenge to compete as an ultimate champion in eight individual events spanning categories from cycling to powerlifting. Like other elite athletes onsite, he donned a wearable device that was tracking physiological data throughout the week.

"Having eight events, it's really critical for me to relax, unwind, let my body recuperate," Spann said. "For me, I'm actually using [the wearable] not to make sure I hit my 8 hours of standing but to really try to make it purposeful, to rest while I can, so that way my recovery the next day is not as harsh."

In elite athletic environments like the Warrior Games, human-worn technologies are increasingly used to provide insights for individuals and teams to train and perform at their best. But as Spann emphasized, peak performance is not just about physical training; it hinges on cognitive readiness and purposeful recovery.

Although sports science has traditionally focused on enhancing performance through physical training and fortitude, many practitioners of the discipline today embrace a human-as-a-system perspective, which links the brain and the body to optimize individual physical and psychological readiness, well-being, performance, and rest.

From competitive games to real-world combat, the U.S. military is integrating this sort of human performance science and technology to keep service members in peak shape to succeed in their critical missions. While coaches continue to collect and analyze physical data—for example, information related to cardiorespiratory and musculoskeletal systems—the cognitive domain is one of the most impactful components of warfighter readiness and resiliency. Researchers are increasingly investigating the important effects that mental preparedness, stress, sleep, nutrition, hydration, and other factors have on training, performance, recovery, and injury prevention and management in the military context.

But now, leaders in the field of mission readiness are taking this holistic view a step further, combining cognitive and physical analytics *with mission analytics* to provide a fuller picture of a warfighter's performance and resilience. This data-driven insight into individuals can optimize performance and accelerate readiness of the entire unit, which in turn fuels mission success. And analysis of biometric and mission performance data supports leaders in making responsible, effective decisions about their formations. Let's explore the powerful effect of combining these disciplines.

"Having eight events, it's really critical for me to relax, unwind, let my body recuperate," Spann shared. "For me, I'm actually using [the wearable] not to make sure I hit my 8 hours of standing but to really try to make it purposeful, to rest while I can, so that way my recovery the next day is not as harsh."

Physiology Meets Mission Operations

Wearables, sensors, and AI make it possible for elite athletes to incorporate individual analytics on heart rate, respiration, sleep cycles, and more to predict outcomes and to personalize training for and recovery from high-intensity activity. But the physical and mental load experienced in the military has few parallels with competitive sports.

For example, while athletes can use periodization to fine-tune their training and peak for a specific function or event, warfighters—and especially unconventional forces—need to be ready to perform whenever and wherever the mission calls, often with little to no advanced notice or rest and recovery time between deployments. Similarly, athletes play on a known field and face known adversaries, while warfighters operate in unknown areas, in unfamiliar conditions, against anonymous enemies. What's more, warfighters face unique military-specific stressors, such as the possibility of having to take actions that may cause serious injury or death.

Over time, these stressors may exact a profound toll on the warfighter's cognitive, physical, and emotional well-being. The impact of this toll on the warfighter's career, life, and family cannot be understated. By quantifying these types of stressors through biometrics, the field of human performance applied to the unique conditions of the military impacts both personal and mission success—extending an individual's fitness for duty and for their return to civilian life.

The delivery of biometrics through human performance platforms creates a snapshot of the warfighter's complex physical and psychological system. But for a complete picture, the military has another valuable tool: *mission analytics,* which are specific to the management and execution of the mission. Applied to warfighters, mission analytics evaluate the effectiveness and precision of operational activity. For example, in modern military training, warfighters and their weapons can be outfitted with sensors that provide data about what soldiers are doing at any given moment, such as where they are looking, whether they are pointing their weapons in the right direction at the right time, or how they follow doctrine in activities such as room clearing.

Together, integrated data from mission analytics and human performance offers the military a powerful new way to enhance training and accelerate readiness (see Figure 1). And it empowers individuals to make changes to improve their own readiness with data-driven feedback. Here's a basic example: A marksmanship trainer can evaluate an individual's shooting technique through observation but cannot see what hidden factors might be influencing poor performance. But integrated data may show that although an individual's technique is good, their pulse, blood pressure, and breathing rise when they squeeze the trigger, leading to underperformance. With this data, the trainer can work with the individual on controlling their heart rate and respiration, paired with proper marksmanship, to hit the target.

This integrated analysis can also be used in environments that are adjacent to the battlespace. For example, a special forces unit could use integrated analytics to support the performance and well-being of operators who disseminate critical, sensitive analytics to the field in real time. Imagine that these operators sit behind a desk for 8 to 10 hours each shift and are under extreme pressure to perform consistently and rapidly, especially during periods of high-intensity operations and cognitive overload. Data from wearables can capture a range of biometrics from these operators, including heart rate variability (HRV; see sidebar below) and sleep quality. By combining these metrics with mission analytics, leaders can track not only the accuracy of the decisions and work itself, but also understand how operators are functioning based on reaction time, working memory, and overall resilience. With this information in hand, units can better support operators with training and education for more consistent, controlled, and resilient performance.

However, the benefits of combining mission analytics and human performance go well beyond individual readiness and recovery. Elite tactical units can be thought of as complex adaptive systems that work together to carry out the mission. In such a system, outcomes can be nonlinear with respect to the number of people on the team. When faced with unexpected adversity or the loss of teammates, members of high-performing teams improvise and adapt to overcome. Their collective output can exceed all expectations. Such outsized team performance requires everyone to perform at the peak of their potential. Advanced analytics is the key to unlocking this force multiplier effect.

Figure 1: Human performance and mission analytics together are a force multiplier

Heart Rate Variability: An Indicator of Resiliency and Readiness

When it comes to strengthening resilience and accelerating readiness, HRV is an especially important metric that provides insight into how a person's stress response affects their performance on and off the battlefield.

HRV is the variation among a set of temporally ordered inter-beat intervals from a continuous measure of heart rate. A high HRV is often associated with high resiliency and the ability to reach peak levels of physical and cognitive ability and quickly return to baseline, whereas low HRV is often associated with low resiliency and can be related to poor physical conditioning, high cognitive stress, illness, pain, fatigue, and overall lower health. Tracking their HRV with wearable technology, warfighters can learn self-regulatory techniques and habits to activate their parasympathetic nervous system and increase their HRV, such as managing their energy, incorporating tactical pauses in operations, and using breathing to refocus on their assignment. When HRV is correlated with other information such as nutrition, physical fitness, and sleep data and analyzed over time, the combined measures provide insight into the quality of sustained operations and can indicate when someone might be on the verge of overtraining or burning out.

Advanced Analytics for Accelerated Readiness

The power of coupling biometric and mission performance data lies in how it is analyzed and applied so leaders can make responsible, effective decisions about their formations. For example, instead of making decisions about personnel based on intuition and observation, leaders can use psychophysiological data to track subtle indicators and anomalies that predict compromised performance. That's where predictive data analytics, AI, and machine learning (ML) come into play.

Automated AI algorithms empower leaders with a current and future-state snapshot of individual and team readiness. With predictive analytics, they can anticipate when members may not perform at their peak capacity because they need to recover from injury, illness, or other stressors. Detecting potential threats in advance creates opportunities to deliver preventive care individuals need before compromising their health or the mission.

In practice, however, how does this intelligence come together for individuals and their leaders to digest and use for decision-making? Dashboards and visualizations (see Figure 2) must be set up with end users in mind to deliver an analysis that is accessible, actionable, and relevant so clear actions can be taken.

These dashboards cater to various user personas, including individual warfighters, unit leaders, and commanders. Individuals can use the information to improve readiness on their own by identifying and addressing things like sleep quality, cardio conditioning, and stress levels. Unit commanders can use information about individuals to assemble the best-performing team for a particular mission. And senior commanders can focus on overall readiness for the battlespace by projecting the performance of U.S. forces.

Glimpsing the Future

Technological breakthroughs will continue to expand the potential for integrated and personalized human performance and mission analytics. The research community is also advancing the field, applying emerging methods, such as the physics-derived Lyapunov model of how chaotic systems work together, with sensor data to derive new insights about human performance. The opportunity to accelerate force readiness and effectiveness by harnessing advances in biometric and mission-sensing technologies and cutting-edge data analytics is limitless. And the landscape is changing quickly, with future mission applications coming into view.

Figure 2: Performance Dashboard Illustration

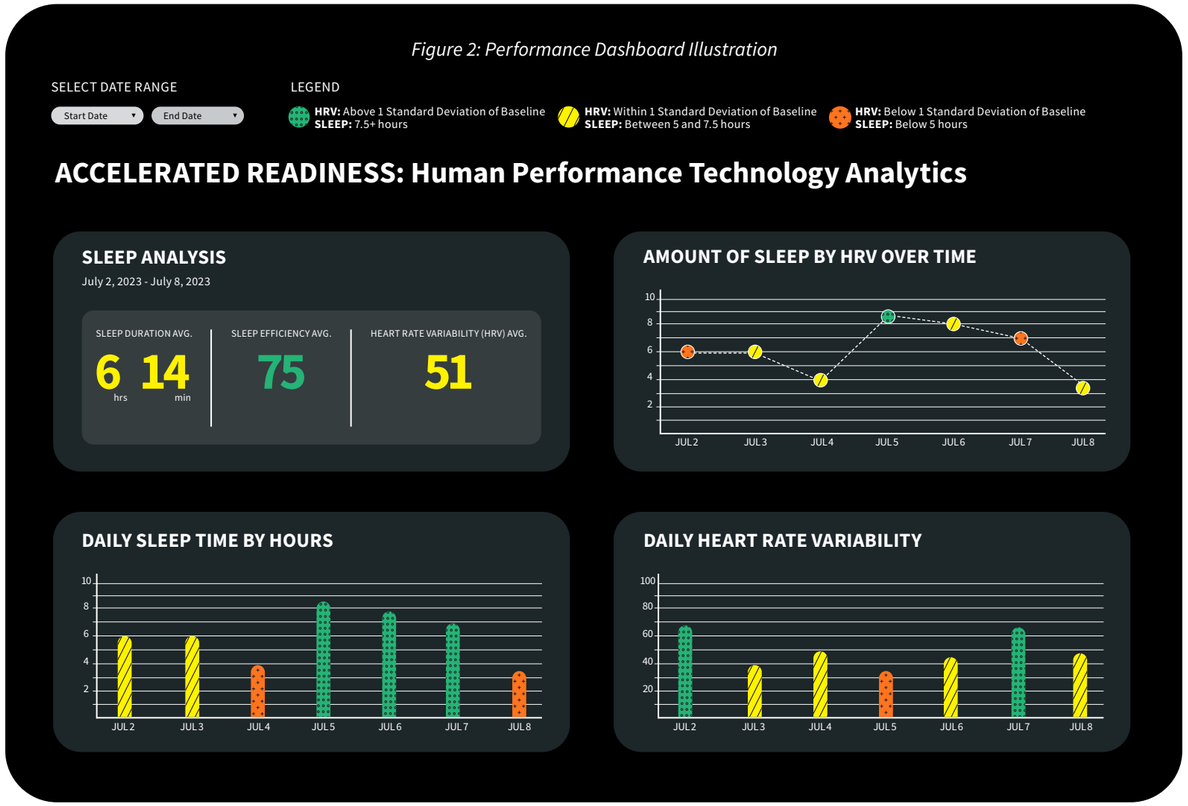

As we look to the future of human performance, areas of high-impact innovation include:

- **Human-machine teaming.** Advances in autonomous sensors can one day allow machines to recognize the strengths and weaknesses of their human partner, then make up the deficit to improve performance of the team as a whole.
- **Digital twins.** Next-generation simulations—including digital twins of humans—will continue informing how defense organizations design, integrate, and field new systems into combat ensembles.
- **Noninvasive neurosensing.** Neuroplasticity means human performance changes over time, so the neuroscience element of warfighter training is essential. Technological innovations in noninvasive sensing devices—through electroencephalography (EEG)-based brain-computer interface (BCI) systems—are further expanding ways to get reliable, precise data about brain activity during high-intensity training periods.
- **Wearables.** Investments in advanced levels of customization will help wearable technology meet the specialized needs of individuals in the military community. "Innovation on this front will allow for more unique, rich biometrics," Jeff Collins, a human performance lead at ŌURA, maker of Oura Ring, said recently. "It will also allow for new mechanisms, such as measuring data optically through the skin."

Each of these advancements should be centered around empowering warfighters to understand and act on their own data, as Spann did at the Warrior Games. "Information is great, but it's what you do with it that matters," Collins added. "Metrics should give insight into how a person is exerting and recovering, offering suggestions and assistance along the way."

As individuals swiftly identify and tackle cognitive and physical issues, they can enhance their performance, boost their recovery, and remain healthy and resilient for longer, leading to extended and more sustainable military careers. This career extension is especially critical today, given the convergence of factors that is making it extremely challenging for the military to meet its recruiting needs. The benefits of improved warfighter well-being to the military overall compound quickly, supporting a more effective force.

> "Information is great, but it's what you do with it that matters. Metrics should give insight into how a person is exerting and recovering, offering suggestions and assistance along the way."
>
> —Jeff Collins, Human Performance, Oura Ring

But these breakthroughs are not ultimately about the technology or the insights that can be derived through AI. They are about the military's greatest asset—its force, from warfighters to operators—and optimizing their individual physical and psychological readiness and performance on and off the battlefield.

Munjeet Singh *is a senior vice president and the leader of Booz Allen's BrightLabs incubator, an experimentation organization designed to develop, test, and incubate mission-centric solutions rooted in emerging technology.*

Cameron Mayer*, a senior vice president at Booz Allen, leads the firm's Defense Accelerated Readiness business with a focus on data driven solutions.*

Dave Prakash, MD, *is an executive advisor at Booz Allen, an Air Force veteran pilot, and a physician—focusing his expertise on AI for defense and health.*

SPEED READ

The integration of wearable technology in elite sports emphasizes both physical training and cognitive readiness, promoting peak performance through holistic approaches.

In the military, the fusion of human performance science with mission analytics ensures optimal readiness by addressing mental preparedness, stress, and individual factors, reshaping training outcomes.

The convergence of biometric technology, mission sensing, and data analytics propels improved decision-making, effective training, and extended careers, supporting a resilient and effective military force.

ENVIRONMENT

THE PHYSICAL IMPACT OF DATA

Achieving Both Computational Power and Sustainability

Brianna Hogan
Contributions from Chris Christou, Brad Beaulieu, and Alex Hanson

As organizations rapidly move to the cloud and embed AI into every aspect of business operations, rarely does the connection between these innovative technologies and their impact on the environment come to mind. Even in the technology sector we tend to think of data—the new "oil" of the Information Age—as having no physical form and therefore no tangible impact on the environment. This could not be further from the truth.

If data is the new oil, then data centers are the refineries that sustain the entire cloud computing industry and fuel carbon emissions from the telecom sector.

Cloud services make it easy to click a button and deploy a virtual server to run an application, start cloud-native service for business analytics, check email, and, notably, provide the processor-intense computational power necessary to train AI models. What's less obvious is the corresponding carbon emissions of each button-click, leading to potentially inefficient cloud resource usage. A study of IT professionals at 100 companies spending about $1 million on cloud computing found that more than half of these companies use only 20%–40% of available central processing units (CPUs), due to idle servers, easily contributing to tens if not hundreds of millions of tons of carbon emissions (CO_2e). While migrating to the cloud offers 1.4–2 times greater energy efficiency in comparison to on-premises data centers, the technology sector's race for green IT has begun, and AI offers the starting point to our growing problems and scalable solutions to achieving cloud and AI sustainability.

Enter AI, Both the Challenge and the Solution

The dramatic proliferation of large language models is generating massive amounts of carbon emissions from the computational power required to train these models, which goes virtually unnoticed by the public. A single AI algorithm, such as the transformer (big) model (see Figure 1), can generate upward of 300 metric tons of CO_2e, the equivalent of 470 people taking flights from New York to San Francisco. However, training the same model using algorithmic techniques to improve energy efficiency reduced carbon emissions by over 1,000%. Training Generative Pre-trained Transformer 3 (GPT-3) with 175 billion parameters required 1,287 megawatt hours of electricity and 502 metric tons of carbon, with one ChatGPT search consuming 100 times more energy than one Google search. The only thing greater than AI's effect on the environment is our opportunity to design AI algorithms in ways that increase our chances of a greener future.

Curbing the Proliferation of Energy Emissions

While the technology sector soars to unprecedented levels of consumption, the Biden Administration established a Federal Sustainability Plan that can serve as a guide to agencies to curb emissions and decarbonize supply chains increasingly powered by the cloud.

Data Centers Across Industries

- **The worldwide consumption of electricity by data centers** was estimated to be between 220 and 320 terawatt hours (TWh) in 2021, more than the electrical consumption of entire countries such as Sweden or Egypt.
- **The telecom sector,** supported by constantly humming servers, produces 5% of global carbon emissions, more than double those of the entire airline industry, and is set to increase to 14% by 2040.
- **The infrastructure-as-a-service market,** which drives data center growth and electrical power consumption in the technology and telecom sector, is projected to grow from $130.9 billion in 2023 to $325.9 billion by 2028, with a compound annual growth rate (CAGR) of 20% (Yahoo Finance).

Figure 1: Carbon Emissions Consumption (CO_2e lbs)

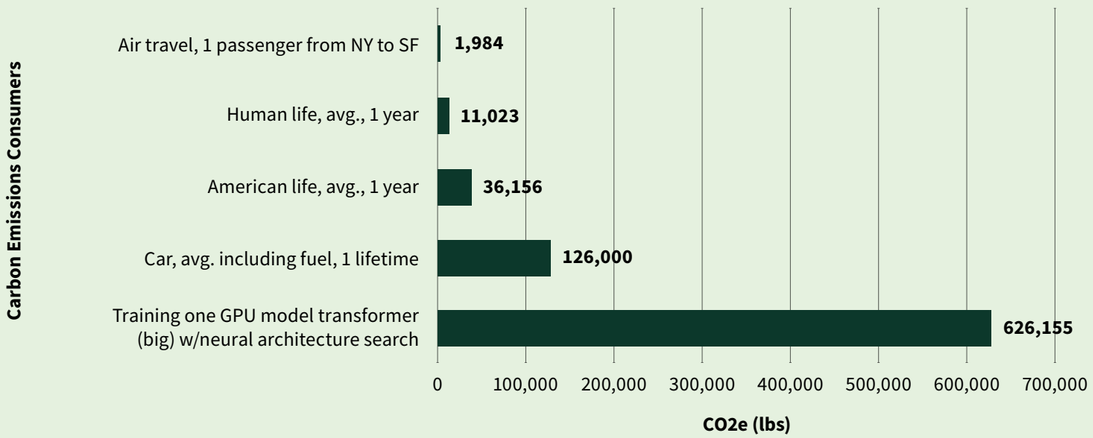

Estimated carbon dioxide equivalent (CO_2e) from everyday activities compared with training a common natural language processing model, the transformer (big) model, with 251 million parameters trained on 8 NVIDIA P100 graphics processing units (GPUs) for 3.5 days (84 hours; 300,000 steps). Source: Energy and Policy Considerations for Deep Learning in NLP

While AI fuels demand for cloud computing to train and execute the models, AI offers solutions for curbing carbon emissions through standard machine learning (ML) techniques that improve energy efficiency. Even before designing an algorithm, AI sustainability begins with the engineer's decision to deploy cloud resources for training and running AI models in renewable energy-powered data centers, such as those listed in Google's human- and machine-readable table. Algorithmic techniques for improving energy efficiency involve achieving the same accuracy of the model with less overall energy consumption and reduced computational cost. Some of these techniques include distillation, the transfer of knowledge from larger models into smaller ones; fine-tuning, which involves refining already-trained models for efficiency; and pruning to remove redundant parameters that do not affect the accuracy of the model but increase computational needs and carbon emissions. Right-sizing ML models for a given purpose also saves considerable amounts of energy by reserving the computationally intense models for consequential efforts, such as cancer research, and less computationally intense models for mundane tasks like summarizing meeting notes.

The advent of TinyML and the BabyLM Challenge throws the gauntlet to ML engineers and researchers to design ML algorithms on low-cost, low-power microcontroller systems that can reduce carbon emissions by 5 to 38 times, according to initial testing. In addition, designing TinyML "at the edge," or closest to where the data exists on a network, significantly reduces the cost and energy required to transfer the data to a central location and train massive models with potentially unnecessary parameters. Government missions at the tactical edge, ranging from the warfighter on the battlefield to the cyber analyst defending satellites in space, can greatly benefit from algorithm optimization and TinyML algorithms on edge compute devices to rapidly develop data and mission insights in areas with intermittent or zero connectivity.

With these methods, smaller and optimized AI algorithms can more effectively enable government missions at the tactical edge and more easily offset their own carbon emissions by reducing emissions in cloud computing and across industry. A sustainability-aware AI algorithm (and engineer) can rapidly identify idle servers, underused compute resources, and more efficient options for cloud data storage with the insights and scalability from AI, while mitigating climate change by increasing the efficiency of power grids, batteries, manufacturing, and supply chains.

Doing more with less and at the edge is the name of the AI sustainability game. Gaining awareness, visibility, and transparency into the connection between cloud and AI carbon emissions is a crucial step toward making data-driven decisions to advance sustainable cloud and AI.

> "Edge AI will be key in helping to build sustainable development practices that reduce carbon emissions while also lowering cloud costs. The amount of data being collected by devices like drones and cameras means cloud processing and energy demands have nowhere to go but up. Moving processing to the edge with models optimized for memory, power, and compute not only lowers cloud carbon emissions and costs, but can also extend missions and their impact."
>
> —***Jags Kandasamy,*** *CEO and co-founder of Latent AI*

Reporting Carbon Emissions

We can't change what we can't measure, and current cloud and AI carbon emissions reporting is still advancing to empower people to make carbon-conscious decisions. For most organizations, cloud and AI carbon emissions fall into Scope 3 reporting. This reporting includes carbon emissions that an organization does not produce directly from operations but indirectly from their supply chains, such as with purchased goods and services. Measuring Scope 3 emissions presents notorious challenges, due to imprecise data and scope for what falls within an organization's supply chain. This problem is especially acute for cloud and AI carbon reporting for which hardly any emissions data exists. Federal agencies can begin to tackle carbon emissions reporting requirements by starting with setting goals to reduce greenhouse gas (GHG) emissions, as described in Biden's "Executive Order 14057 on catalyzing American clean

energy industries and jobs through Federal sustainability" and partnering with carbon-conscious organizations to accelerate progress.

While major cloud service providers (CSPs) offer some form of carbon footprint tracking tools for cloud-hosted workloads, the reporting does not currently include near real-time data with granularity at the project or workload level that could inform the C-suite, project manager, engineers, and regulators of how best to use and report this data. Some CSPs have roadmaps that include new features to approach more granular reporting of carbon emissions. However, this only includes the compute resources that power AI and not the software and algorithms themselves. In research and academia, most scholarly publications include information about an AI model's accuracy, number of parameters, and duration of time to run, but few include efficiency metrics such as carbon emissions, training costs, or model accelerators. We have reached the frontier of carbon emissions tracking, and further progress requires a coalition of the willing and a commitment to environmental, social, and governance (ESG) causes to create the measures by which regulators and the public should assess our carbon footprint.

Technology titans, technologists, and trainees have access to tools that can begin to close the chasm between what metrics we see and what metrics we should be tracking for cloud and AI carbon emissions reporting. The Green Software Foundation, an organization committed to carbon reduction in software development, created the Software Carbon Intensity (SCI) Specification, which describes how to calculate the carbon intensity of a software application. The SCI GitHub page provides a methodology for "calculating the SCI score for any software application, from a large, distributed cloud system to a small, monolithic open-source library, any on-premises application, or even a serverless function." Climatiq's REST API supports organizations' GHG data in technology products with real-time emissions calculations. In addition to climate-focused tools, most CSPs offer cloud optimization services that use automation to identify ways to improve cloud cost and usage efficiency and well-architected frameworks that include sustainability and/or performance optimization in their design principles. While no one tool singlehandedly resolves the transparency problem of measuring AI carbon emissions, education about these tools and methodologies can empower people to reduce carbon emissions during the design, development, and deployment of AI algorithms.

Imminent changes to federal regulations for carbon emissions reporting have also upped the ante for action. In November 2022, the Department of Defense, General Services Administration, and NASA submitted a proposal to amend the Federal Acquisition Regulation to ensure major federal contractors disclose their GHG emissions and climate-related financial risks and set science-based targets to reduce GHG emissions. Earlier the same year, the Securities and Exchange Commission released a draft rule requiring public organizations to disclose Scope 1, Scope 2, and Scope 3 GHG emissions. While disclosure requirements are still evolving, the message is clear: Upcoming changes to policy aim to ensure federal suppliers make required disclosures and set targets to reduce GHG emissions.

The IT systems we design today that use cloud and AI technology will be the subject of regulator assessments tomorrow. Organizations can prepare for the requirement to report on Scope 3 carbon emissions by gaining visibility into cloud carbon metrics reporting from the data generators. Due to the shared responsibility model between cloud providers and cloud consumers, both sides need to work together to tackle the cloud carbon data visibility challenges. Cloud providers can evolve existing carbon footprint tools to support real-time reporting services with APIs; publicly available and repeatable emissions calculations methodologies; and granular tracking of carbon emissions per service, region, and workload with the ability to tag resources based on customer-defined groupings across projects and business units. Consumers of cloud computing, especially systems integrators for federal entities, can relay carbon reporting requirements to CSPs for feature enhancements and embed sustainability metrics, thresholds, and goals into cloud and AI algorithm design.

Any investment in education, strategy, and implementation of sustainable cloud and AI practices can lay the foundation for a greener future and avoid costly rework of applications, cloud infrastructure, and AI models to comply with federal GHG reporting requirements.

What Can Organizations Do Today?

Curbing cloud and AI carbon emissions begins with awareness of the connection between them and understanding every person's role in the solution—from enterprise leadership to the end user. Based on our on-the-ground experience navigating enterprise sustainability, here are specific areas of focus for key roles that can start to make an impact today:

- **Senior Executive and Agency Leader.** Identify sustainable cloud and AI as an organizational priority, incorporate objectives into the organization's strategic roadmap, and design science-based targets looking toward federal regulation.
- **Procurement Official.** Include strategies, methods, and targets for carbon emissions reductions in requests for information (RFIs) and requirements for corporate experience in solicitations.
- **Auditor.** Require clear, repeatable, and publicly available methodologies for carbon emissions calculations from federal and commercial organizations.
- **Project Manager.** Identify areas of waste and excess in cloud spend due to inefficient resource usage (e.g., leaving lower environment servers running 24/7) and work with engineering/development teams to optimize resource usage to reduce cost and carbon emissions.
- **Developer and Engineer.** Design carbon-aware systems with published thresholds, metrics, and dashboards; write energy-efficient code using software carbon intensity measuring tools; and optimize cloud and AI compute resources to reduce energy consumption.

Ultimately, we all have a role in achieving cloud and AI sustainability in the actions we take, in big and small ways. For example, individuals can reduce the amount of data stored in the cloud by deleting emails, pictures, and applications and avoiding the increase of additional cloud storage space and costs. The difference between achieving the promise and realizing the downstream impacts of cloud computing and AI lies in our action or inaction for advancing green IT. Awareness of cloud carbon emissions and AI's potential to be both the greatest cause and most efficient solution for the climate crisis offers a starting place for embedding sustainability practices within every organization. As the effects of climate change intensify, so too does our individual and collective responsibility to curb cloud carbon emissions in order to build a greener future at home, in cyberspace, and at the tactical edge.

Brianna Hogan *focuses on emerging cloud capabilities and development in Booz Allen's BrightLabs incubator, an experimentation organization designed to develop, test, and incubate mission-centric solutions rooted in emerging technology.*

SPEED READ

While AI fuels demand for cloud computing to train and execute the models, AI offers solutions for curbing carbon emissions through standard machine learning (ML) techniques that improve energy efficiency.

Right-sizing ML models for a given purpose saves considerable amounts of energy by reserving computationally intense models for consequential efforts and less computationally intense models for mundane tasks.

Navigating enterprise sustainability begins with awareness of the connection between the cloud and carbon emissions and understanding every person's role in the solution—from enterprise leadership to end user.

CRITICAL SERVICES

THE EMERGING CITIZEN EXPERIENCE

Advantages and Challenges of Human-Machine Service Delivery

Will Healy, Ernest Sohn, and Santiago Milian
Contributions from Jeff Leon, Melanie Laffin, and Dia Metropolis

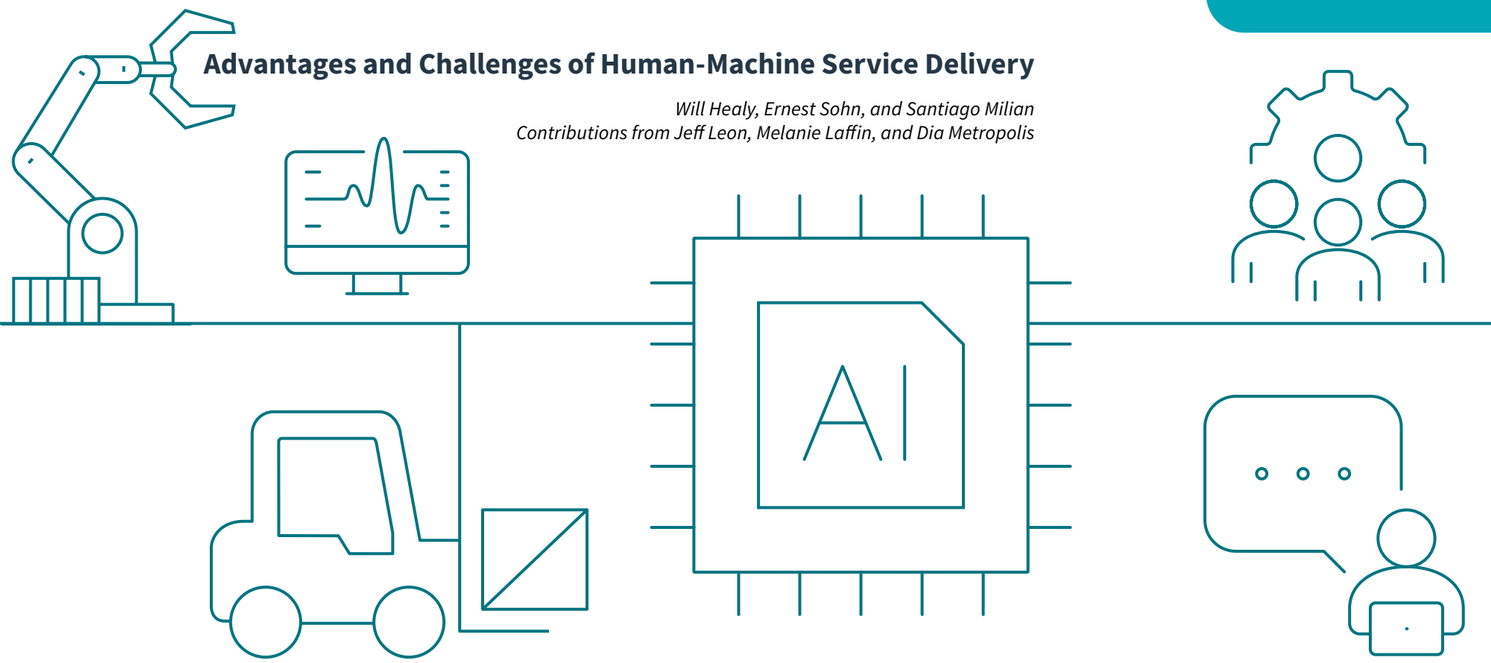

As our nation's most critical service provider, the Federal Government is navigating fast-changing customer expectations. Along with commercial organizations, agencies are rapidly investing in AI-powered tools—from user interfaces to software platforms on the back end—to augment and streamline critical citizen services.

Unlike in consumer culture, however, customer experience in the Federal Government is not a matter of competitive advantage or profit optimization, but it can make or break the mission of delivering vital services and benefits. And for citizens, the value of service delivery can significantly impact a person's life—whether that involves seeking health benefits, finding shelter and food, or looking for employment.

Here we explore the opportunities and nuances inherent in AI innovation for the citizen experience.

AI for Citizen Experience

AI capabilities, such as natural language processing, have been part of government services for years, providing transformative support to citizens—from the Department of Veterans Affairs analyzing free-text data to identify veterans at risk of suicide or homelessness to Recreation.gov using voice recognition to analyze customers' speech and triage their questions (see "Spotlight on Recreation.gov"). Emerging advances in AI and large language models are rapidly enhancing tools, such as citizen-facing chatbots, and finding new use cases—from helping with data interoperability and improving search functions to providing personalized recommendations.

As agencies advance their use of AI to modernize and streamline the citizen experience, they will need to weigh new challenges and considerations, given the unique and critical services they deliver to the American public.

Four important issues to navigate:

1 **Reconciling access and equity with AI solutions.** While the commercial sector can choose to cater to only the most profitable customer segments, the Federal Government takes pride in serving *all* of the population, delivering vital services equitably. AI can help increase the reach, accessibility, and impact of federal programs. An AI-enabled government of the future would be better able to serve citizens in ways that fit their situations and needs, on various devices, in nearly any language, at any time of day—not requiring people to use high-speed internet or to visit federal offices during business hours (such requirements inadvertently harm working families that don't have paid time off or those without computers or broadband).

A well-known shortcoming of AI technology is its potential to propagate biases, but AI tools can also be used to reach populations that may have been historically underserved. In healthcare, for example, race, gender, and socioeconomic biases in AI models have the potential to harm already marginalized populations and reinforce or magnify inequities in the system. AI can also easily exclude people who are not tech savvy. For instance, a chatbot powered by generative AI should be fine-tuned to accommodate the digital literacy of the entire citizen population. Agencies that orient responsible product development around broad and diverse customer needs, experiences, and levels of digital literacy can use AI to advance, rather than to impede, equity and access.

Since federal agencies have a mission to support every citizen, they have a front-row seat to advances in responsible AI and can be a model for how other industries can experiment with inclusive innovation (read more about this topic on page 36). However, AI and technology are still unable to replace human-to-human exchanges for the most critical citizen needs.

2 **Using AI to augment rather than replace humans.** There are many areas where government should not put an AI agent between a person and their needs—seeking health benefits, natural disaster relief, or unemployment compensation is not like buying a pair of shoes.

While AI cannot replace human agents, it can increase their capacity to do uniquely human tasks that are more complex and specialized. For instance, AI agents can answer routine citizen questions, provide personalized assistance, triage inquiries, and help human agents rapidly filter through large data sets. Enabling workforce efficiencies allows employees more time to handle complicated needs and questions, where they can add more value.

Meanwhile, citizens should be kept fully in the loop with AI, given the high-stakes nature of federal services. While this may be counterintuitive, AI developers in government should consider what areas of an AI-powered system should actually be *harder* to use to reduce user error and encourage human interaction. If an AI system is too easy, a user might overlook a small error or inconsistency on a tax form or housing application—leading to real-world, personal consequences.

Spotlight on Recreation.gov

Recreation.gov is the government's central travel planning platform and reservation system for nine federal agencies, where the public can reserve camping sites, buy park passes, find recreation information, and much more. Its customer service operation handles high call volumes covering a wide variety of issues that can take significant time and resources to address.

Similarly, developers should be aware of how the design of automated systems can exacerbate or mitigate automation bias, which is the propensity of people to believe and act on suggestions from automated decision-making systems instead of relying on their own analysis and expertise. This issue applies to anything dealing with safety, from medical diagnosis to aviation, but also to lower-risk situations, such as digitizing paper forms that contain critical information to process taxes and benefits.

3 **Streamlining experiences through customer data while overcoming mistrust.** Personalized, AI-driven services hinge on whether citizens trust the government enough to provide relevant information that can power the algorithms. However, according to 2023 data from the Partnership for Public Service, only about one-third of Americans say they trust the Federal Government, compared with 46% who do not.

While the challenges of trust in digital society are deeper than the topic of customer experience (last year's Velocity publication focused extensively on this issue), agencies deploying AI in citizen applications must navigate the "value exchange" inherent in digital society: People are more willing to share data with companies if they perceive there will be immediate and tangible value in return. This same value exchange underlies customer interactions in the federal space.

Providing value in the moment through an AI output creates a continuous value exchange:

- Agencies use AI to improve the immediacy and tangibility of their services.
- Citizens perceive the value from exchanging data for services and become more willing to share data on a continued basis.

An AI solution, deployed across channels including web chat, voice, and text messaging, is streamlining the customer experience and improving the quality of services. The technology uses voice recognition and natural language processing to analyze callers' speech, and machine learning to handle requests for information and triage calls before they are transferred to live agents. Automated screening and routing of simple requests provides faster and more accurate responses to many questions, which allows live agents to spend the majority of their time on resolving more complex customer interactions.

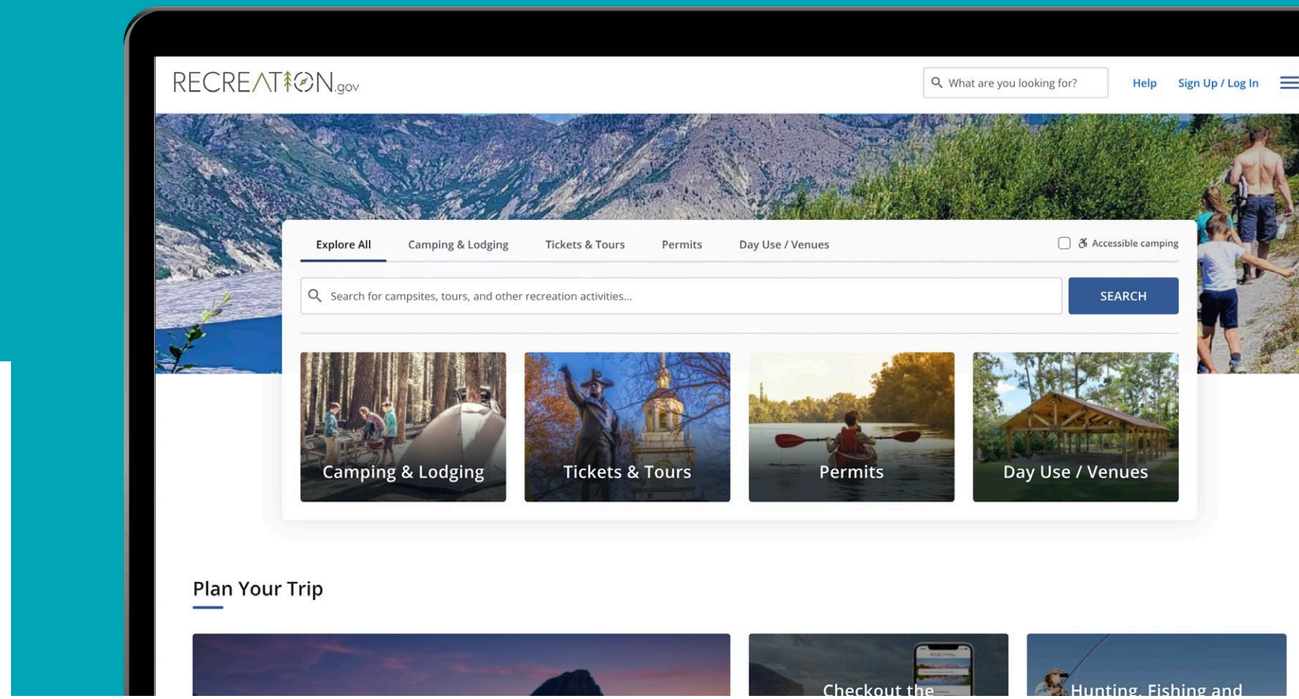

- The cycle continues as they receive immediate impact and ultimately reduce the time and effort needed to search and apply for services.

Of course, the significant benefits to the customer from this critical data exchange need to come with increased investments in more sophisticated data protection.

4 **Mobilizing for an increased attack surface and new threat actors employing AI for fraud.** Private-and public-sector organizations, even those not on the cutting edge of using AI for customer service, face growing threats from bad actors using tools like ChatGPT to commit fraud. Ever-evolving schemes include impersonating people through advances in deepfake visuals and bypassing filters and monitoring systems to steal information. Cyber criminals are growing increasingly skilled at getting through defenses such as continuous monitoring. Their successful exploits disrupt operations, generate negative publicity and political backlash, and degrade customer experience. Consider a stolen tax refund—the targeted individual suffers financially and has to spend hours correcting the mistake, which wasn't of their making.

While AI is the source of new fraud schemes against customers, AI can also be employed to counter these criminal activities, as part of a multidisciplinary approach (for more on cybersecurity, see page 64). For example, because predictive models are largely trained on known categories of fraud schemes, they routinely fail to detect novel techniques. More sophisticated analytics and the latest AI and machine learning techniques can detect new fraud schemes that are still needles in a haystack. Oftentimes bad actors test a few new schemes to determine which ones will work and are driven to other fraud schemes once they realize a particular avenue is monitored. AI can detect faint signals and suspicious patterns to identify these schemes so they can be stopped before they have a broad impact that affects customers.

This changing digital ecosystem will require ongoing dialogue and educational effort among federal employees and the public to create broader awareness around advances in fraud and to help citizens spot tell-tale signs of bad actors using AI against them.

Table 1: Enhanced Customer Experience Operations

Basic Practices	Examples of AI Enhancements
Customer Experience Management & Measurement Quantitative and qualitative customer experience data is collected through multiple mediums before it is analyzed and shared across the organization, to identify opportunities for short-term service recovery and long-term systems improvement.	Using natural language processing techniques and large language models (LLMs) to parse large amounts of customer experience data, summarize that data, and provide customer trends.
Customer Research & Synthesis Research provides actionable insight into who customers are, what they value, and how they feel when interacting with the organization.	Using generative AI to create personas, based on this research, to interact with customers or even test new features before they go live.
Product & Service Design Solutions are designed to improve interactions with customers and the underlying operations, systems, and policies shaping service delivery.	Designing AI agents that seamlessly integrate to aid humans in end-to-end service delivery to address key pain points and enhance customer experience.
Agile Implementation Products, services, and systems are developed, deployed, and seamlessly integrated into portfolios to create unified customer and employee experiences.	Integrating AI tools across the software development lifecycle (e.g., testing, generative software) to increase productivity of agile development teams.
Service Delivery & Support Organizations manage the change inherent in the introduction of new solutions to customers, employees, and stakeholders by providing the necessary supporting infrastructure, training, communications, and governance.	Using LLMs to integrate advanced self-service and on-demand resources to help users navigate change and help the enterprise rapidly disseminate information.
Customer Service & Engagement Individual customer feedback is received and issues are identified and resolved in a multichannel environment, to build and maintain positive long-term customer relationships.	Using AI agents to route and prioritize customer requests and using natural language processing for sentiment analysis, to decipher language used in customer interactions.

> “This changing digital ecosystem will require ongoing dialogue and educational effort among federal employees and the public to create broader awareness around advances in fraud and to help citizens spot tell-tale signs of bad actors using AI against them.”

Keeping the Customer at the Center

Ultimately, citizen services through the Federal Government will benefit from AI in targeted areas rather than carte blanche (see Table 1: “Enhanced Customer Experience Operations”). But whether AI solutions are powering user interfaces, streamlining back-end systems, or detecting fraud, the technology in and of itself does not improve experience. Therefore, as agencies embrace new AI-powered systems for service delivery, multidisciplinary disciplines—from human-centered design to data engineering to cybersecurity—should be brought together to address the nuances of mission requirements and deploy solutions effectively and responsibly for public use.

Will Healy *is a senior leader in Booz Allen’s technology organization with a focus on the recreation management space.*

Ernest Sohn *is a data scientist in Booz Allen’s AI practice, focused on the integration of AI-enabled solutions for federal civilian clients.*

Santiago Milian *leads Booz Allen’s CX capabilities, delivering human-centered digital strategies for advanced citizen services.*

SPEED READ

The commercial sector’s increasing investment in AI for customer experience is shaping expectations for individualized and integrated services across industries, including the Federal Government. The government is leveraging AI to augment and streamline critical citizen services that can significantly impact citizens’ lives, such as health benefits, shelter, food, and employment.

The government must ensure equitable service delivery to all citizens and consider potential biases in AI. AI tools should cater to the entire citizen population and varying levels of digital literacy. Moreover, the trust deficit in the government means that agencies must provide tangible value in return for the citizen-provided data, ensuring data protection and addressing concerns about AI-driven fraud.

For AI to truly benefit citizen services, it must be integrated thoughtfully and responsibly. This will involve multidisciplinary approaches, combining human-centered design, data engineering, and cybersecurity. The ultimate goal is to keep the citizen at the center of all AI enhancements, ensuring technology genuinely improves the experience.

A Digital Tapestry

WEAVING AI INTO UNIFIED CYBER DEFENSE

Patrick Myers and Aaron Sant-Miller
Contributions from Michael Dahlberg and Colin Friedman

Every day, federal agencies are locked in an intense struggle to defend, monitor, and remediate vulnerabilities across their ever-expanding attack surface. It's a mission that gets more difficult each day as cybercriminals, insider threats, and nation-state adversaries adopt increasingly creative and dynamic tactics to breach networks. Many of these tactics are powered by AI and machine learning (ML), which allow adversaries to generate new attacks with increasing scale and efficiency while providing methods to tailor attacks to the characteristics of the cyber defenses that are in place.

Traditional defensive strategies—from blocklists to signature-based detection—aren't enough to stop sophisticated adversaries: Novel, advanced detection techniques are needed. These new defenses must be able to identify emerging threats as they evolve, uncover previously undetectable behaviors, differentiate the malicious from the anomalous, and harden defenses before adversaries can deliver their effects.

To protect against today's most advanced adversaries, cyber defenders must develop and deliver tradecraft that is equally advanced and dynamic. This means operationalizing AI/ML models of our own that can learn patterns in higher dimensional space, across a wider range of data sources, and be correlated with adversarial actions or known threats, tactics, and procedures (TTPs). While adversaries can dynamically adjust their execution of attacks, their methods are often consistent and generalizable. AI/ML models are designed to detect these methods, enabling more proactive defenses that are harder to circumvent. When applied correctly and correlated with reputable threat intelligence sources, AI/ML can harden networks against previously undetectable signatures faster and with more precision.

While individual federal agencies may invest in the development of AI/ML for advanced detection, these investments struggle to scale across the enterprise. As a result, a unified, government-wide approach to cyber defense is increasingly difficult to sustain.

Challenges with an "Acquisition-First" Mentality

To stand up defensive cyber suites, federal agencies often buy commercial tools to collect network and endpoint data, procure products to monitor that data, and purchase licenses for software to store and analyze it. In effect, acquisition trends dictate the tools and technologies agencies use to defend and monitor networks.

In this product-rich market, a comprehensive cyber solution stack all too often amounts to a jumble of vendor tools that don't integrate well. What's more, establishing these defensive suites consumes significant time and resources, which limits the Government's ability to adjust tooling over time. Now, as adversaries grow more inventive, the Government's defenses evolve only as quickly as commercial tools advance. Without the flexibility to rapidly field advanced capabilities like AI/ML, federal cyber defenses may be constrained by enhancements to vendor capabilities that have already been procured.

To fight the adversaries of tomorrow, defensive minds must focus on the development of innovative detection and prevention capabilities. Improving fundamental architectures, establishing common standards, and maturing analytic approaches can make integrating tools and software less cumbersome and free up time for more original thinking.

More broadly, federal agencies are faced with several infrastructure challenges that limit effective fielding of AI/ML, including:

- **Data visibility limitations**, where data is siloed to specific locations, tools, or vendors
- **Procuring duplicative capability**, such that an analogous defensive and monitoring suite exists within each platform

When the Federal Government chooses to invest directly in advanced detection tools (e.g., AI/ML), these tools often become deeply intertwined with the commercial platforms used by the agency. This integration limits the Government's ability to retain intellectual property (IP) and consequently hampers the sharing of tradecraft across the federal community. In practice, adversaries attack the Federal Government in diverse ways, using consolidated campaigns against multiple federal networks simultaneously, increasing the need for defenders to unify their defenses. This requires the free exchange of advanced detection capability, AI/ML models, and intelligence across systems and organizations. As tradecraft is often tightly coupled to commercial tools, distributing these capabilities either requires a common tool procurement for all agencies or the reengineering of tradecraft across products. The Government is, therefore, confronted with either overtly expensive tradecraft-sharing costs or a vendor-lock scenario that disincentivizes commercial innovation.

To fight the adversaries of tomorrow, defensive minds must focus on the development of innovative detection and prevention capabilities. Improving fundamental architectures, establishing common standards, and maturing analytic approaches can make integrating tools and software less cumbersome and free up time for more original thinking. This will improve the solutions of today and help futureproof existing platforms to deliver more advanced capability and tradecraft. As adversaries develop new and nuanced attacks, defenders can train and develop AI/ML to detect the methods used in those attacks, sharing these models across the community to harden all federal networks simultaneously.

A Balanced Approach to Manage Advanced Capabilities

Federal agencies can extract maximum value from commercial offerings while retaining the agility to implement advanced detection capabilities (e.g., AI/ML) without being solely reliant on commercial entities. The key to this equilibrium, where both commercial and noncommercial tools can harmoniously coexist, is adopting open, modular, and purpose-driven architectures. As illustrated in Figure 1, this architecture centers around two core components:

1. A data broker that seamlessly channels data flow across multiple systems
2. A versatile framework for deploying and managing advanced detection tools, like AI/ML

TRADITIONAL APPROACH

Multiple and manual data-to-platform connections, where AI and advanced detection capabilities are coupled with the platforms in which they are executed.

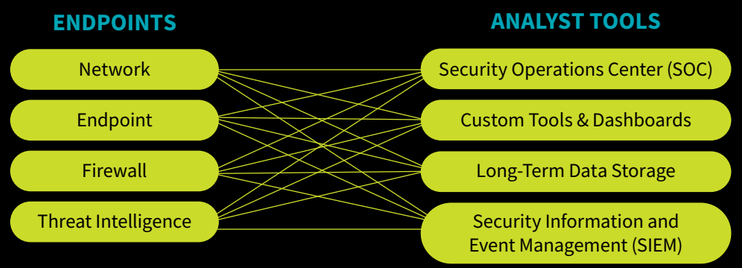

MODERN APPROACH

Central data broker connects sensor data and platforms, where modular AI and advanced detection capabilities run upstream to downstream platforms.

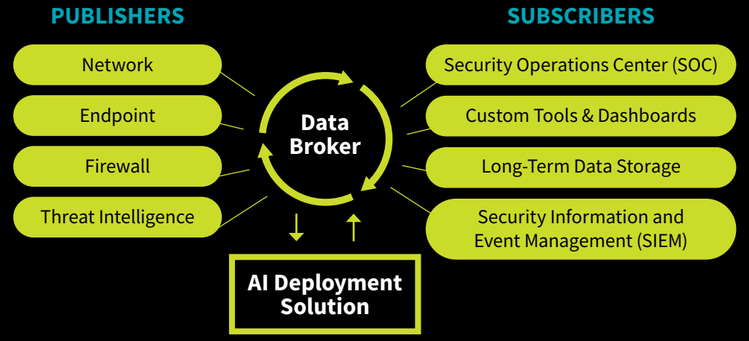

Figure 1: Comparison of Traditional vs. Modern Data Broker Approach

MISSION HIGHLIGHT

The CDAO

The Department of Defense (DoD) Chief Digital and Artificial Intelligence Office (CDAO) is "the senior official responsible for the acceleration of the DoD's adoption of data, analytics, and AI to generate decision advantage across the boardroom to the battlefield." In support of that mission, the CDAO is developing a suite of infrastructure, tools, services, and best practices to unify and scale AI-enabled solutions for national defense. To this aim, the CDAO offers Perceptor, an AI/ML deployment and monitoring platform owned and operated by the CDAO and intended for joint community use. With Perceptor, cyber AI/ML or advanced detection capabilities can be decoupled from the infrastructure in which they are executed and instead aligned to a government standard and shared across agencies. The CDAO has successfully partnered with agencies across DoD to integrate Perceptor with existing defensive suites, accelerating integration of advanced detection tools and AI/ML capabilities, enhancing the defenses of DoD information networks, and allowing government IP retention and cross-agency capability sharing. In practice, this has allowed DoD to fight as one, exchanging and reusing AI/ML capabilities to work together against common adversaries.

In this architecture, data brokers serve as the connective tissue between data feeds and data lakes. Data brokers employ a "pub/sub" (publish/subscribe) model in which data feeds publish data to the broker and data lakes subscribe to the data needed. This allows organizations to treat data as a living artifact, where modular applications consume data streams off the data broker, process that data, and publish it back so other applications can access data products with ease. This breaks down data silos, allowing new solutions to be deployed with full access to data, and addresses fundamental data limitations that exist in commercially diverse defensive suites. Consequently, federal agencies can enhance data visibility across platforms without rigid, time-consuming data-integration activities.

In parallel, a dedicated solution for advanced detection deployment reduces reliance on commercial vendors to execute new tradecraft. This deployment solution works in concert with the data broker, where it subscribes to streaming input data, enriches input data with advanced detection outputs, and publishes all data back to the broker for consumption by operational platforms and data lakes. Additionally, this simplifies mission integration of advanced detection tools (e.g., AI/ML), where the solution handles the connection of data to n number of detection capability tools.

This open architecture allows federal agencies to use the commercial infrastructure they are comfortable with while increasing their flexibility to deploy new defenses. Advanced detection capabilities become increasingly shareable between mission partners as tradecraft is deployed as modular applications. This empowers the Federal Government to fight together, scaling investments on top of those made by other agencies. By reducing cumbersome data and capability integration activities, agencies can reallocate resources to invest in the development of more advanced detection methods required to defend the nation against increasingly creative and dynamic adversaries. By decoupling these capabilities with the infrastructure in which they are executed, new investments can truly have an enterprise-wide impact.

Enabling the Shift

Winning tomorrow's fight requires a cultural shift *and* an architectural shift to keep pace with sophisticated adversaries and drive the development of advanced tradecraft. Rather than fighting in silos, federal agencies and organizations must work together to harden shared defenses. This requires the adoption of open architectures, intelligent protection of federal IP, and decreased dependence on out-of-the-box commercial offerings. By continuing to address the infrastructure challenges of today that limit scalable, enterprise-wide use of AI/ML for cybersecurity, the federal community can shift resources to defend our nation against the cyber adversaries of tomorrow.

Patrick Myers *and* ***Aaron Sant-Miller*** *focus on deploying AI and ML solutions that secure the nation's IT, respond to threats, and support critical missions across federal agencies.*

SPEED READ

As cyber adversaries become more inventive, traditional defense strategies are proving inadequate. There is a need for advanced detection techniques, particularly AI/ML, to proactively identify and counter emerging threats, differentiating malicious activities from anomalies.

The current trend of acquiring tools often leads to a fragmented defensive landscape. This approach not only limits the Government's agility in adapting to new threats but also ties them to vendors' motivations, potentially stunting innovation and adaptability.

A unified vision for the future requires a balanced approach, blending commercial and noncommercial tools through open, modular architectures. This would enable federal agencies to share advanced detection capabilities, breaking down data silos and fostering a collaborative defensive strategy against sophisticated cyber adversaries.

SNAP SHOT

*2023 Opening of Booz Allen's Pax River Mission Systems Integration Facility (Speaker: **Judi Dotson**, Booz Allen Global Defense Sector President). At 20,000 square feet, this facility is a first-of-its-kind space in the region to rapidly design, develop, prototype, integrate, test, and evaluate innovative solutions that address the warfighter's evolving and dynamic needs.*

From Science to Practical Climate Resilience

THE VISION, BARRIERS, AND NOVEL AI APPROACHES TO OPERATIONALIZE CLIMATE SCIENCE

Prachi Sukhatankar
Contributions from Christopher Holder and Marie Nguyen

In 1985, Carl Sagan—one of the most prominent scientists of our times—testified in front of Congress about increased greenhouse gas emissions and their potential global and intergenerational impact. He urged progress in climate science and research and indicated how that could be done by leveraging historical data.

Climate science has come a long way since its inception in the late 1950s. The coupled Earth system models, primarily developed by federal agencies and research centers such as NOAA, NASA, the National Center for Atmospheric Research (NCAR), Intergovernmental Panel on Climate Change (IPCC) and others rely on sophisticated algorithms and advanced high performance computing to process large-scale datasets. This combination of science and technology ushered in a new era that significantly enhanced the understanding of our dynamic planet and the different forces at play. Several decades later, as we look to build a climate-resilient nation and globe, we are uniquely positioned to leverage climate models and other diverse datasets and to capitalize on the rise of AI, harnessing its power of pattern detection across massive amounts of structured and unstructured data to sort signal from noise, better extract insights, and inform climate adaptation and impact mitigation.

The future era of climate resilience requires far more than sophisticated climate science and scientific computing. It demands a commitment to translating advanced climate science and purposefully embedding that intelligence into operational use where it is most relevant: at the hyperlocal level.

Climate resilience, as the term suggests, means being prepared in the face of climate adversities and recovering to the same state of well-being or better.

Adversities include more frequent and extreme weather events, such as extended droughts, heat waves, and ocean warming, and the direct and indirect impact of those events on natural ecosystems, built infrastructure, and daily life.

The impact of these events and our aspirations toward resiliency are closely related to conditions across various domains of civilization, such as society, economics, the environment, government, and technology (see Figure 1 on page 70).

Moreover, new mechanisms can fuel powerful scientific collaboration and operationalization, so that small interventions done at scale yield big impacts. This opportunity to drive increased collaboration and operationalization is where novel AI approaches can help in powerful ways.

The rise of ubiquitous AI—increasingly embedded in an ever-growing array of applications and devices—offers a unique chance to bridge the gap between advanced research and on-the-ground operations. Federal organizations and the climate community can harness the power of AI tools to better bring the science to localized operations. From infrastructure improvements to proactive disaster planning, bringing climate insights to decision-making can achieve the level of scale and speed required to respond to rapidly evolving climate threats. But what barriers to the operationalization of climate science lie in front of us, and how do we flip the narrative through novel approaches to integrate data, algorithms, and compute power, so that these challenges become opportunities to build real-world resilience?

Barriers and Opportunities for Operationalizing the Science for Climate Resilience

Accessing relevant climate information is difficult for many reasons. While Earth-observing satellites capture hundreds of terabytes of data daily, the downlink stations can process only a fraction of that data, which is further transformed into analysis-ready data and stored online before being made available for broader analytical use. Simulations of climate scenarios are often stored in large nonintuitive repositories that use unique file-naming structures and file formats. For example, basic access to data may require that the systems ingest files with names like *r1i1p1f1*[1], with a NetCDF, BUFR, or GRIB2 file structure, where latency is often as critical as accuracy. The files stored on repositories can be enormous—on the order of petabytes—and may cover only a few years of a single global climate variable. Downloading and working with multiple variables over a time frame of several

[1]*The naming convention of* r1i1p1f1 *specifies the "variant-ID." The letter "r" represents the ensemble (aka realization) number; "i" is for the initial conditions used in the model; "p" stands for the physics parameterization; and "f" represents the forcing conditions.*

decades typically requires using remote servers and computational resources, which may not be available to the average researcher. Plus, researchers and scientists may not always have a deep background in information technology (IT) and can find it challenging to work with advanced IT systems. Breaking down these barriers and allowing subject matter experts to harness the power of advanced technologies will be key to unlocking large-scale scientific collaboration.

Agencies can take several steps to address these issues. For example, with edge computing—including deployable algorithms; powerful, distributed graphics processing units (GPUs); and advanced compression and pruning techniques to streamline models for greater efficiency—we can limit the need to backhaul Earth observation satellite data through the internet into cloud-based repositories. "For AI to reach its full potential, processing needs to occur where the data is being collected. However, most light form factor hardware like drones can't run standard AI models that tend to be large and power intensive," shared Jags Kandasamy, chief executive officer and co-founder of Latent AI. "By compressing models while simultaneously improving inference, processing can take place on the edge and deliver real-time reliable results and extended mission capabilities." Performing data processing and analytics at the edge, whether in space or at ground stations, and combining it with federated learning approaches reduces latency and unlocks options for critical decision-making needed at the speed of relevance, such as in response to disaster threats and recovery.

Additionally, open-source data platforms can be made available and leveraged by scientists and citizen scientists alike. These platforms provide access to various data sources that are translated into more universally recognized formats on authoritative data repositories. AI can also be used to conduct topical analyses that return novel data associations for intuitive search and discoverability.

Ultimately, the pace of scientific discovery and its application to the climate mission hinge on the level at which collaboration can be fostered and the depth at which the relevance of scientific products can be measured. This is no easy feat, owing to scientific information and research residing in large-scale, unstructured data, such as journals and web repositories. Applying AI to that data can identify patterns within scientific information, ranging from large-scale analysis of citations and mining web and social data to identify interrelationship, to trends around usage and the relevance of scientific analyses. Generative AI systems, for instance, can be trained on a vast corpus of scientific climate journals and academic papers to aid with summarization and intuitive queries. Additionally, natural language processing (NLP), a subfield of AI, can help with extraction of specific patterns for decision-making around key issues. For example, an AI solution deployed for the Department of Defense (DoD) helps search and uncover critical policy documents, within massive troves of context-rich text, that discuss the policies of interest and highlight groups tackling similar problems. This solution is changing the policy game for DoD and is open sourced for broader access and benefit.

This is the kind of democratization of AI that will provide pathways to improve scientific collaboration and empower the shift to open data and open science paradigms, a first step toward the broader operationalization we envision.

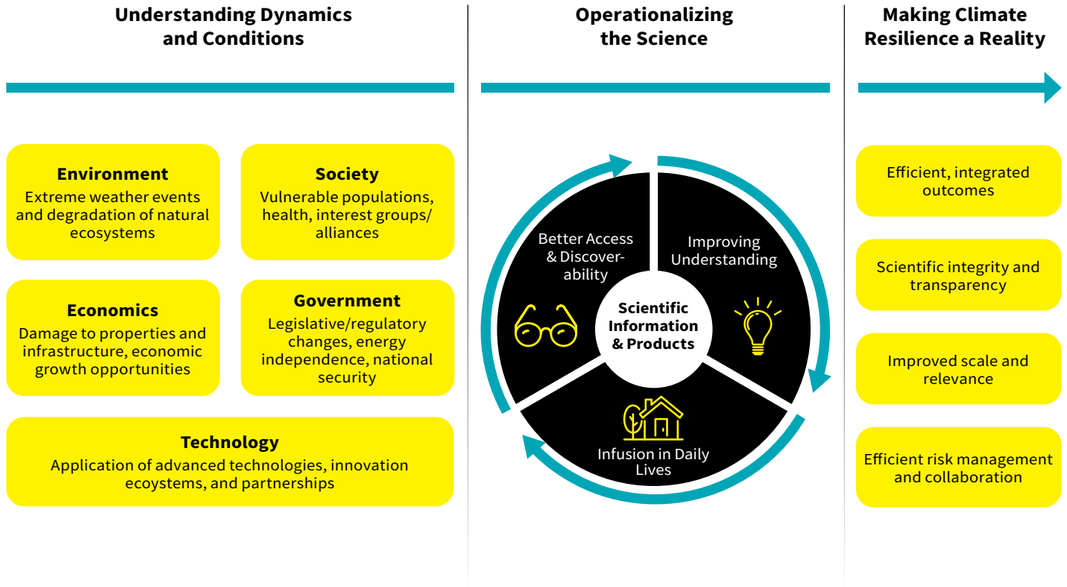

Figure 1: Climate resilience requires an understanding of interconnected dynamic conditions along with broader application and access to information.

DELIVERING AI-READY DATA

A project within a space agency gathers Earth-observation data and parameters to serve to the public in several free, easy-to-access, easy-to-use methods. It helps communities become resilient amid observed climate variability by improving data accessibility, aiding research in renewable energy development, building energy efficiency, and supporting agriculture projects.

All parameters and outputs are provided in commonly employed community formats, naming conventions, and units. The project also provides thoroughly documented application programming interfaces (APIs), a catalog of geospatially enabled analysis-ready data (ARD), and the capability to visually explore data via geospatial services, serving over 100,000 unique users with more than 50 million data requests in a year.

Infusion of Scientific Products into Daily Lives

Climate science deals with information that is presented across large time frames (decades and centuries) and spatial scales (multi-kilometer boxes), which can make it difficult to grapple with the relevance of it. Resilience planners and disaster response personnel may labor to translate climate information from computer models into actionable planning for climate resilience or to deploy it for disaster prediction, mitigation, and recovery. For example, if climate models predict that a city will see a 50% yearly increase in precipitation by 2100, what changes should be made to allow the city to become more resilient to this increase? How does this environmental threat interrelate with economic, societal, and technological conditions?

As a result of this difficulty in translating information into action, there is a limited infusion of valuable scientific information and products into daily operations, which further hinders crucial long-term planning and management of localities. Agencies can take steps to address this disconnect, including using advanced technologies to generate analyses and outputs that are relevant to specific industries or communities. For example, AI, specifically predictive analytics, is increasingly being used to identify demand-supply equations and characterize trends related to natural resources, such as water, food, energy, and critical minerals. AI could also be used to report livability metrics for cities now and to predict them in the future by integrating data sources such as local weather patterns, documented hazard sites, locations of critical infrastructure, city walkability, and reliability of public transport.

Reducing or eliminating barriers to operationalizing climate science could lead to astonishing advances in various sectors. Significant opportunities include infusing climate AI into our daily lives and multi-domain insights for decision-making, rapid preparedness, scenario planning, and recovery.

Novel Data-Driven Science and AI Scenarios Where the Science Is Operationalized

While many of us may not be conscious of the technological changes around us, AI has already penetrated many moments in our day-to-day lives—from avoiding highway collisions to communicating with our home devices through speech recognition. Leveraging these already pervasive technologies and infusing important environmental and climatological information would be game-changing. Envision global positioning system (GPS) navigation with an AI application for a long road trip that recommends an optimal route to the driver of an electric vehicle (EV) based on weather conditions, overall traffic-flow management, locations of charging stations, and personal preferences. We are seeing the beginning of these types of innovations already.

Imagine an AI application that dynamically aggregates and controls thousands of devices that are plugged into the electricity grid (including EVs) so they can supply power when demand is greatest while minimizing inconvenience to vehicle owners. This would decrease costs by minimizing the need to build new power plants and increase the reliability of electricity supply by adding grid resources at scale. This aggregation of distributed energy resources—called Virtual Power Plants (VPPs)—could reduce peak demand in the U.S. by 60 gigawatts, almost 5% of current total electric grid capacity, with EV batteries potentially playing a significant role. Embedding AI algorithms with even small impact, when multiplied across hundreds of millions of devices, can compound the effect, facilitating optimization of electricity generation and load, creating a more intelligent electricity grid, and revolutionizing the energy sector.

Delivering nuanced information in near-real time is possible with the combination of AI and climate science. Here's an example: A storm is approaching a region and the region's grid operators are trying to predict where damages and outages are most likely to occur. The operators pull up a dashboard showing historical data about the area, such as average and maximum rainfall. AI also displays predictions within the dashboard for the affected area, such as downed powerlines, wind speeds, and hail. A digital twin of the locality visualizes the likely course of action, continuously updates the model as new information from the field is streamed in and allows planning according to different scenarios and outcomes. Post-event, the feedback "bots" or intelligent automations work in the background to refine the recommendations, based on the data and learnings from this event, to build resilience toward potential future extreme weather events.

Extreme weather events have been causing billions of dollars in damages to infrastructure. Consider the conventional approaches to managing a large-scale infrastructure rebuilding project. From building to security and IT systems to operations, the process is often siloed and manual, forcing decisions based on incomplete or obsolete data. However, for critical infrastructure—particularly in regions that are prone to climate adversities—new AI-based capabilities are allowing for insight-driven planning and sustainment that dramatically improves operational efficiency, safety, and overall resilience. For example, AI and digital twins are being used to provide a digital thread that links authoritative data sources, such as enterprise data systems, geographic information system layers, utility assets, and building information model (BIM) updates to enable integrated analysis; real-time feeds; and human-centered, contextualized visualizations for diverse stakeholders to make critical design and engineering decisions for improved climate resilience.

Agencies are already implementing climate resilience solutions that assess risk across key climate threats. For example, DoD Climate Assessment Tool (DCAT) allows facility managers of DoD installations to assess risks from environmental threats under various climate scenarios. The U.S. Army Corps of Engineers utilizes relevant climate information to enhance infrastructure climate resilience. The Army's Climate Hydrology Assessment Tool (CHAT) allows users to visualize and calculate trends for historical and future projections of annual streamflow, precipitation, and temperature.

Similar constructs and solutions could be repurposed for localities where trained AI systems provide climate assessments and recommendations based on livability and workability parameters that are within the acceptable thresholds for a community. Communities could access environmental information and predictions, ranging from seasonal, extreme heat and precipitation forecasts to drought conditions. Consolidating this data with information from other domains, such as demographics, infrastructure costs, and cultural context, would provide comprehensive, actionable insights; improve understanding; and, ultimately, help communities ask the right questions and provide feedback on issues that matter to them.

A Catalyst for Climate Resilience

The data-driven science and AI scenarios envisioned here are only the beginning. By fully operationalizing the science, we will achieve a level of scale and agility where agencies, the private sector, academic and community partners, citizen scientists, and the public at large will be able to play a key role in building climate resilience.

The pathways described here have the potential to solve the challenges already at our doorstep and those that Carl Sagan alluded to in his 1985 remarks where he urged for international amity. Almost 40 years later, the need for operationalization of science couldn't be more urgent. Using novel AI approaches to address this need can be the catalyst that makes climate resilience a reality.

Prachi Sukhatankar *leads climate and infrastructure engagements for clients, with a focus on using technology solutions such as climate intelligence, smart transport infrastructure, and advanced energy technologies to meet mission objectives.*

"Not only is reliable, accessible, compelling data necessary to build long-term resilience and sustainability planning, it also provides context and timely warnings for those living in areas vulnerable to high-impact events in the here and now. Existential matters related to water and energy are certainly in this mix as are immediate threats including flooding, fire, and excessive heat. None of these issues respect state boundaries, thus demanding high levels of collaboration and a thoughtful, coordinated exchange of information and analytics."

—***Wellington "Duke" Reiter,*** *executive director of Ten Across, a regional resilience initiative that focuses on proactive decision making around social, economic, and climate change.*

SPEED READ

The future era of climate resilience requires a commitment to translating advanced climate science and purposefully embedding it into operational use; this is where novel AI approaches can help in powerful ways.

Barriers exist to operationalizing climate science, including limited access to scientific data and discoverability hindering collaboration and the limited infusion of scientific products into daily lives at scale.

Democratization of AI can help overcome these barriers and bridge the gap between science and operations. Examples include application of edge AI, NLP, predictive analytics, and generative AI toward advances that range from Earth-observation data acquisition, scientific collaboration, climate risk assessments, and community resilience recommendations to grid optimization and large-scale electrification. These are the catalysts for making climate resilience a reality—they are here and available now.

Code Writing Code

THE NEXT MAJOR DISRUPTION IN SOFTWARE DEVELOPMENT

Jarid Cottrell, Josh Gordon, Michael Martoccia, and Sahil Sanghvi

There are some technologies that have a before and after moment: the internet, hosting applications in the cloud, containerization with Docker technology, and automation with continuous integration and continuous deployment (CI/CD). All these innovations changed the landscape of how software developers interact with code and data, using new technologies as the medium. And now it's happening again with AI, and software development is officially disrupted.

Rapid advancements in generative AI have impacted every stage across the entire software development lifecycle. New AI-powered tools are being developed in functional verticals, with a concentrated focus on the development phase—and momentum is building around the review and testing phase. Massive innovation has brought about a seismic shift in a short period of time: from a world where humans solely wrote code to one where AI is an indispensable partner in coding. Tools like GitHub Copilot and ChatGPT have become the new normal, revolutionizing the way teams approach every aspect of software development.

The next disruption will be the application of large language models (LLMs) in the form of AI agents. With AI as a partner in software development, are we on the cusp of a revolution that could render traditional coding obsolete?

AI Agents: The Future of Coding

AI agents are self-governed, goal-driven entities focused on a role that aims to tackle a challenge. The AI agent uses a series of instructions or prompts against an LLM to develop a list of tasks, determine the reasons for completing them, and initiate action. Unlike traditional language models, such as ChatGPT, which provide single responses to user queries, AI agents engage in complex workflows with minimal human intervention. In these workflows, the model converses with itself to perform complex decision making and can take actions using tools given to them. This allows the agent to work through the different phases of the software development lifecycle, including planning, design, implementation, and testing (see Figure 1).

In the not-so-distant future, a software developer will be able to describe tasks in natural language and have an AI agent build the step-by-step plan, complete all the tasks, check for errors and make revisions, and provide an output, whether that be in the form of code, text, audio, or video. And it will do so at the speed of light, awaiting the developer's feedback and enabling faster software development cycles than ever before. AI agents will allow development teams to scale their productivity, with humans providing oversight.

The engine powering AI agents are the LLMs, which have been trained on a massive amount of data, including code libraries and technical examples. This enables AI agents to be used within development workflows and across the software development lifecycle as the new medium for interaction between developers and code. With AI as a partner in coding, software engineers can halve their coding time. There are predictions that AI agents could boost developer productivity more than 4-fold by 2030, and there are those who predict it will take half that amount of time.

To illustrate this further, here is an example from the development phase of the software development lifecycle. A software engineer could task an AI developer agent with any task, from front-end website development to machine learning model creation. With only a natural language prompt as an input, the AI agent would autonomously output code files to a local directory. The developer can instruct the agent to revise its work throughout the process. This allows the developer to increase their productivity through AI code creation while remaining in the loop through transparent communication from the agent as it writes, tests, and finalizes the code. Booz Allen is currently building AI developer agents to review GitHub pull requests, which take significant time to thoroughly assess. By using an agent, pull requests can be automatically reviewed, with the agent suggesting code for the developer to incorporate before merging the code back to the main repository.

> "Pre-trained AI models represent the most important architectural change in software since the internet."
>
> —***Matt Bornstein*** *and* ***Rajko Radovanovic****, in "Emerging Architectures for LLM Applications" (Andreessen Horowitz)*

Figure 1: Capabilities of AI agents across every phase of the software development lifecycle

LIFECYCLE PHASE	EXAMPLE AGENT CAPABILITIES
Planning	An **AI agent** can help with solution planning and requirements gathering by functioning as a chatbot that interacts with stakeholders to gather requirements through natural language conversations. The chatbot can help elicit detailed project requirements and specifications from nontechnical users and turn them into various deliverables, such as product roadmaps, architectural documents, or agile stories to scope out.
Design	A **generative design AI agent** can automatically create design concepts based on input parameters and constraints. This can assist in generating multiple design options quickly, allowing designers to explore creative possibilities.
Development	An **AI developer agent** can perform many functions, from code creation to pull request reviews and suggestions. Coding agents can operate as plug-ins and autonomous actions within the tools developers already use, such as VS Code and GitHub.
Review and Testing	An **AI testing agent** can autonomously review code and create unit tests within developer repositories. This allows developers to spend more time on innovative, mission-focused work and less time testing code for errors.
Deployment	An **AI deployment agent** can monitor the deployment process and automatically roll back code changes if anomalies or performance issues are detected. This ensures smoother deployments and reduces downtime.
Maintenance and Operations	An **anomaly detection AI agent** can continuously monitor system performance metrics and user interactions. It can identify abnormal patterns, trigger alerts for proactive maintenance, and make code suggestions to improve performance based on its findings.

The value these agents bring is multi-faceted:

- **Collaborative Efficiency.** AI agents streamline task execution by managing multiple tasks simultaneously, fostering a collaborative environment where human intervention is minimized and long-term efficiency gains are realized.
- **Autonomous and Collaborative Problem Solving.** AI agents independently devise solutions to complex problems, working in tandem with human counterparts to achieve comprehensive and accurate results.
- **Empowered Learning and Workforce Development.** AI agents facilitate interactive and personalized learning experiences, automating routine tasks and allowing the workforce to focus on strategic initiatives, thereby enhancing job satisfaction and productivity.
- **Adaptive Evolution.** AI agents continuously evolve to solve increasingly complex tasks, enhancing their value and adaptability. They leverage advancements in AI models and tools, ensuring safety, security, and the generation of data-driven insights.

As AI agents become more prevalent and task-specific, AI systems will autonomously generate and optimize code in collaboration with other AI systems—a new era of AI helping AI to code. This will increase the iteration speed of development, yielding improvements from faster feedback across the various phases of the software development lifecycle. Human trust of AI will also increase, resulting in development teams expanding the responsibilities given to their AI developer agents. Eventually, these models will evolve into multimodal systems, combining text, speech, and vision to unlock a new level of conversational generative AI. However, this advancement is not without cost. Organizations will need to develop frameworks to manage AI agents and monitor token usage to effectively control the AI-powered development lifecycle.

How Will Agents Be Managed? With a Familiar, Yet New, Platform

AI agent capabilities extend to advanced reasoning, planning, tool usage, maintaining memory, and self-reflection. These attributes equip AI applications with a fundamentally new set of capabilities, reshaping the way we construct and interact with AI systems. Establishing the AI Agents Management Platform within the solution architecture empowers software developers to become managers of the agents they control—selecting the best LLM for the task at hand, managing token usage and budget among deployable agents, choosing personas and fine-tuning prompts, and managing the data pipelines that provide the agents with contextual memory to better solve the problem they're presented with.

The AI Agents Management Platform (see Figure 2) is becoming an integral part of the LLM application architecture. As these frameworks mature and reliability is enhanced, there will be a significant shift in the AI landscape that allows the developer to be a manager instead of a worker.

The AI Agents Management Platform serves as a control center, providing a unified interface for businesses to monitor and manage their AI agents. It encompasses several key areas:

- **Agent Specialization.** AI agents can be specialized to handle specific tasks. The management platform allows businesses to assign and manage these specializations, ensuring that each agent is used in areas where it can deliver the most value.
- **Regression and Drift Management.** AI agents, like any machine learning models, can suffer from model drift over time. The management platform provides tools to monitor for drift and implement corrective measures when necessary. It also allows for regression testing to ensure that updates or changes to the AI agents do not negatively impact their performance.
- **Resource Management.** AI agents consume resources, such as tokens, during their operations. The

Figure 2: Approach to managing and developing with AI agents

This framework can be used when architecting an AI Agent application and provides the foundation for generative AI implementation and a holistic approach across people, processes, and technology. The approach outlines how the Management Platform connects with and controls AI Agents, the Memory Layer, and the Thought Engine that powers the application.

ACTIONS AND STEPS, EXPLAINED

1. **User Input (Developer)**: User makes a request through the AI agent application to generate an output (e.g., code).

   Through the management dashboard, the **business user** manages all agent functionality to control the permissions within the application, LLMs being used, tools that are integrated, and data that can be accessed. A business leader's actions happen in parallel to the work being done by the AI agents.

2. **Prompt Management**: Improve contextual understanding by analyzing the conversation, user's chat history, role, and stored data.
3. **API Call to LLM**: Send user's input and contextual information to the LLM.
4. **LLM Orchestration**: The agent communicates with itself, other agents, applications that it has access to, and long-term memory which adds contextual information that the agent uses to plan and implement a solution.
5. **Refinement**: If needed, refine the response by sending it back to the LLM with specific instructions for solution improvement.
6. **Data Storage**: Save the final response and LLM output in the appropriate data resources.
7. **Agent Response**: The AI agent sends the output to the user, which can be in the form of text, code, or an action made within an integrated tool. The AI agent also outputs information to the management platform by logging relevant data and metrics.

Management Platform

At the core of AI agent operations lies the management layer, empowering both business and technical leaders with comprehensive control. This layer equips users with a robust set of tools to effectively oversee AI agent activities and user interactions. Key features encompass data management, integrations management, and model administration. In addition, it offers user permission controls, resource cost analysis, and API key management to ensure optimal utilization.

AI Agents

The AI agents layer serves as the bridge between users' requests and meaningful outputs. Users input their requirements, such as tasks like developing Python scripts, resolving bugs, or translating code, or even submitting items like GitHub pull requests. These inputs are dynamically transformed into prompts that direct the AI agents' activities. The agents collaborate not only among themselves but also with integrated tools, generating valuable outputs for users. This layer also records comprehensive data, facilitating subsequent evaluation through the intuitive management dashboard.

Memory

Enhancing the AI agents' proficiency is the memory layer, a repository of contextual information specific to tasks and enterprises. This layer grants AI agents access to an array of data stores ranging from structured to unstructured and vectorized formats. By augmenting the agents' knowledge with mission-critical insights, the memory layer significantly improves their overall performance. This contextual supplementation enables AI agents to deliver outputs that align precisely with the user's requirements.

Thought Engine

The hosting of the AI agents application happens within the thought engine. This is a combination of cloud resources (multi-cloud compatible), local compute, and in-application processing.

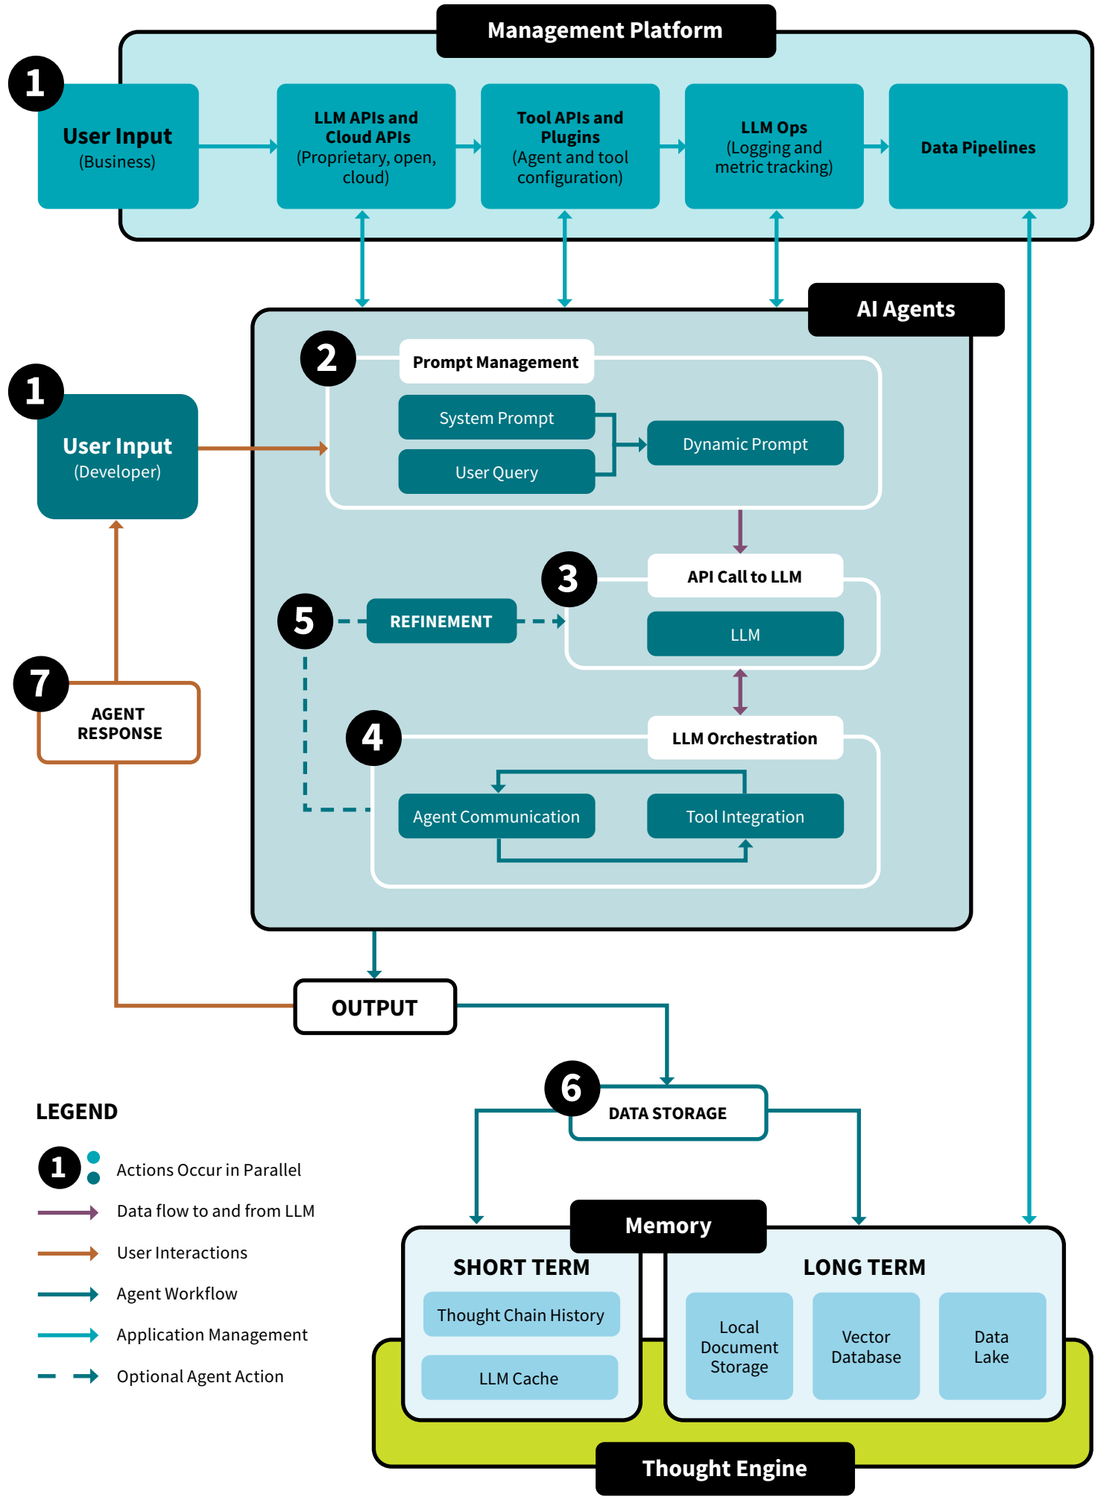

management platform provides tools to monitor and manage resource consumption, ensuring that the AI agents operate within their allocated resources.

- **Ethics and Compliance.** As AI agents take on more complex tasks, ensuring ethical use and compliance with regulations becomes crucial. The management platform provides mechanisms to enforce ethical guidelines and ensure compliance with regulations such as those related to data privacy.

- **Continuous Learning and Improvement.** AI agents can learn and improve over time. The management platform provides a framework for continuous learning, allowing the AI agents to adapt to new tasks, improve their performance, and stay updated with the latest advancements in AI.

The AI Agents Management Platform is a critical component in the successful deployment of AI agents. It provides the necessary tools and controls to manage AI agents effectively, ensuring they deliver maximum value while operating within ethical and regulatory boundaries. As AI continues to evolve, this management platform will play a pivotal role in shaping the future of AI in software development. It will not only boost developer productivity but will also fundamentally transform the way developers interact with code.

AI agents and AI agent management are already revolutionizing the way developers code and how software development teams scale their work across the entire software development lifecycle. Today, we are on the cusp of witnessing not just an exponential growth in productivity but also a fundamental shift in human-AI relationships across software development teams.

The Path Forward

Over the coming year, we anticipate AI agents to evolve from simple task executors to advanced problem solvers. With the integration of LLMs, memory solutions, and tools, AI agents will be able to plan, execute, and reprioritize tasks more effectively, thereby overcoming current bottlenecks in task execution and reprioritization. This period will also see a surge in the development of specialized tools designed to enhance the capabilities of AI agents. These tools will address the current limitations and will be instrumental in improving the quality of task execution and the ability of AI agents to adapt to changing priorities. Major cloud providers are also expected to evolve their offerings to include trained AI agents callable via an API that can be deployed for specific functions at various stages of the development lifecycle.

As AI agents take on more complex tasks, ensuring ethical use and compliance with regulations will become crucial. We foresee the emergence of robust frameworks for managing AI agents, including features for monitoring ethics and bias and mechanisms for dealing with potential threats from bad actors.

We are on the brink of a new era where AI agents will become an integral part of the software development process, transforming it from a labor-intensive task to a highly efficient and automated process. This transformation will redefine the role of developers, turning them into managers of AI agents, and usher in a new age of software development where code is capable of writing code.

Jarid Cottrell, Michael Martoccia, *and* ***Josh Gordon*** *manage software delivery and digital strategy projects for federal clients. They also lead strategic initiatives and solution development at the intersection of AI and software innovation.*

Sahil Sanghvi *is a distinguished technologist and head of strategy for Booz Allen's Technology business, with a focus on emerging technologies and their applications in government and industry.*

SPEED READ

Rapid advancements in generative AI have impacted every stage across the entire software development lifecycle. Massive innovation has brought about a seismic shift in a short period of time: from a world where humans solely wrote code to one where AI is an indispensable partner in coding.

AI agents will allow development teams to scale their productivity, with humans providing oversight. LLMs which have been trained on a massive amount of data, including code libraries and technical examples, enable AI agents to be used within development workflows and across the software development lifecycle as the new medium for interaction between developers and code.

The AI Agents Management Platform is a critical component in the successful deployment of AI agents. It provides the necessary tools and controls to manage AI agents effectively, ensuring they deliver maximum value while operating within ethical and regulatory boundaries.

SNAP SHOT

Booz Allen employees at lab in Laurel, MD

AI& EVERYTHING

A FUTURE OF LIMITLESS POSSIBILITIES

Horacio Rozanski, *President and Chief Executive Officer, Booz Allen*

We have entered a new technological era. Pushing the boundaries of what we once thought possible, generative AI and large language models are ushering in unprecedented change, excitement, and, indeed, trepidation. The fascination with this technology may be akin to the awe and wonder of the first person to peer through a microscope or hear a faraway voice over the telephone. AI is reshaping human life and work, from education to professional tasks to scientific breakthroughs. This era is marked by reimagined potential, with AI's true value extending well beyond mere algorithms.

In this year's issue of *Velocity*, Booz Allen looked at how AI impacts nearly every facet of life—from the workforce and citizen services to climate and human performance. Through our work in all these areas, we see how AI, when used responsibly, is a remarkable accelerant across almost every variable and challenge. The dynamic utility of AI is rooted in the algorithm *and* its application, and AI's promise will be fully realized when it ceases to be used on its own. That's why, as we approach the next wave of technological evolution, the focus will shift to "AI&."

THE POWER OF AI&

AI is becoming a part of everything, but it's not the end of anything. Its true potential is unleashed when paired with other technologies to amplify capabilities, identify anomalies, and accelerate transformational outcomes.

In every corner of our technological landscape, the emerging implications of AI are vast and groundbreaking. For instance, AI is primed to boost the resilience and adaptability of cyber, space, and communications capabilities. Zero-day exploits could be thwarted by coupling cyber and AI to create dynamic defenses against previously unknown threats. In the borderless and congested realm of space, AI offers enhanced awareness and collision avoidance capabilities for better protection and navigation of our space assets. Wireless and satellite communications networks can leverage AI to adapt to real-time conditions, optimizing resources to avoid interference or congestion and improve clarity, even in the most remote locations.

Over time, these types of AI& applications deliver increased value as they learn and evolve. Yet AI's potential doesn't stop at augmenting our current technologies. When converged with quantum mechanics, AI stands to usher in quantum machine learning to enhance processing power and speed, as we learned about in "Can Quantum Supercharge AI?" on page 11.

Ultimately, though, AI will need to be paired with *people* to achieve its greatest outputs. As you read in "A Workforce Disrupted" on page 30, AI is already redefining the workplace, automating routine tasks, and empowering people to engage in higher cognitive functions. The partnership between humans and machines does not merely improve efficiency; it elevates human potential. AI will liberate our brainpower to focus on more strategic and impactful endeavors, fostering an environment where innovation thrives and breakthroughs come to fruition.

For example, in the age of AI, software developers can move beyond coding and debugging to imagine the vast possibilities that lie within the data. Their personal programming abilities will no longer constrict their vision of the attainable. They can explore and experiment with an AI coding partner like the one we described in "Code Writing Code" on page 74—with the human focusing on the "what," while guiding and training their AI partner to handle the "how." Uniting these powerful forces will deliver solutions more quickly and surmount previously intractable problems.

"AI is becoming a part of everything, but it's not the end of anything. Its true potential is unleashed when paired with other technologies to amplify capabilities, identify anomalies, and accelerate transformational outcomes."

Historically, the fusion of technologies and human intellect has driven societal progress. At our current inflection point, AI& people's impact on society promises to be revolutionary as it transforms economies, drives unprecedented growth, and strengthens national security. Despite all the good that can come from the power of AI&, it is not immune to a downside.

AI& THREAT ACTORS

The gap between the introduction of new technology and its malicious exploitation is narrowing. As we embrace AI's capabilities, we must also be prepared for amplified technological risks, spanning from biases to security threats. The digital age brought heightened cybersecurity concerns, and AI presents new vulnerabilities.

Adversarial AI introduces threats like data manipulation—by subtly altering input data, malicious actors can deceive AI systems, leading to erroneous outputs. Such data manipulation can lead to significant misclassifications and ultimately harm. Furthermore, AI-powered deepfakes—hyper-realistic but entirely fake content generated by algorithms—pose a significant risk to information integrity. In the hands of the "bad guys," these can be used to spread misinformation, manipulate public opinion, or impersonate public officials.

As we navigate these challenges, we must not lose sight of the horizon. The good and the bad of AI will continue to rapidly evolve, and so we must evolve with it. We need to simultaneously drive AI adoption, defend against its risks, and ensure its responsible use. Only then can we fully realize AI's potential for groundbreaking advancements without halting our progress. As we look ahead, the future of AI& beckons.

AI& THE FUTURE

Now is not the time to let fear outweigh opportunity. We are poised to write the next transformative chapter of our history. Embracing and investing in ethical AI use and human capital will be crucial as we envision and build a future where humans and machines join forces to create extraordinary outcomes. AI literacy, redefined roles, and strategic partnerships are essential components for maximizing AI adoption and its benefits.

True transformation with AI is a collective endeavor, transcending individual companies and sectors. Partnerships across industries, governments, and societies are required to achieve AI's full potential. The promise of AI is within reach, when the technology ceases to be a tool on its own, but *the* critical ingredient in solving the greatest challenges of our time.

06. Linking Large Language Models for Space Domain Awareness

"Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST, January 2023, https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

Erik Gregersen, "megaconstellation." Encyclopedia Britannica, May 12, 2023, https://www.britannica.com/technology/megaconstellation.

Kari A. Bingen, Kaitlyn Johnson, Makena Young, and John Raymond, "Space Threat Assessment 2023," CSIS, April 14, 2023, https://www.csis.org/analysis/space-threat-assessment-2023.

"Large Language Models Explained," NVIDIA, accessed August 2023, https://www.nvidia.com/en-us/glossary/data-science/large-language-models/.

Laura Duffy and James Lake, "Cislunar Spacepower: The New Frontier," *Space Force Journal,* no. 2, (December 31, 2021), https://spaceforcejournal.org/3859-2/.

"Overview," USSPACECOM Magazine, 2023, https://media.defense.gov/2023/Apr/14/2003199875/-1/-1/1/2023%20USSPACECOM%20MAGAZINE.PDF.

Sandra Erwin, "U.S. Sharpens Plan for Military Space Race: Space Force Challenged to Compete with China Amid 'Lack of Communication, Zero Transparency,'" July 11, 2023, https://spacenews.com/u-s-sharpens-plan-for-military-space-race/.

"Space Debris and Human Spacecraft," NASA, May 26, 2021, https://www.nasa.gov/mission_pages/station/news/orbital_debris.html.

Tereza Pultarova and Elizabeth Howell, "Starlink Satellites: Everything You Need to Know About the Controversial Internet Megaconstellation," Space.com, last updated August 2, 2023.

University of Plymouth, "Scientists Call for Global Push to Eliminate Space Junk," Phys.org, March 9, 2023, https://phys.org/news/2023-03-scientists-global-space-junk.html.

11. Can Quantum Supercharge AI?

Amine Zeguendry, Zahi Jarir, and Mohamed Quafafou, "Quantum Machine Learning: A Review and Case Studies," Entropy 25, no. 2 (February 2023), https://www.mdpi.com/1099-4300/25/2/287.

Antonio Acín et al., "The Quantum Technologies Roadmap: A European Community View," *New Journal of Physics 20,* (August 2018), https://iopscience.iop.org/article/10.1088/1367-2630/aad1ea.

B. Jaderberg et al., "Quantum Self-Supervised Learning," *Quantum Science and Technology 7*, no. 3 (May 2022), https://iopscience.iop.org/article/10.1088/2058-9565/ac6825.

Dan O'Shea, "Honeywell Sees $2 Billion in Sales for Quantinuum by 2026," Inside Quantum Technology, February 15, 2022, https://www.insidequantumtechnology.com/news-archive/honeywell-sees-2-billion-in-sales-for-quantinuum-by-2026/.

Ian Hellström, "Roads to Quantum Advantage," Ian Hellström, September 8, 2022, https://ianhellstrom.org/roads-to-quantum-advantage/.

Kinshuk Sengupta and Praveen Ranjan Srivastava, "Quantum Algorithm for Quicker Clinical Prognostic Analysis: An Application and Experimental Study Using CT Scan Images of COVID-19 Patients," *BMC Medical Informatics and Decision Making* 21 (2021), https://bmcmedinformdecismak.biomedcentral.com/articles/10.1186/s12911-021-01588-6.

Kushal Batra et al., "Quantum Machine Learning Algorithms for Drug Discovery Applications," *Journal of Chemical Information and Modeling* 61, no. 6 (2021): 2641-2647, https://pubs.acs.org/doi/10.1021/acs.jcim.1c00166.

Lei Zhang, Andriy Miranskyy, and Walid Rjaibi, "Quantum Advantage and Y2K Bug: Comparison," Cornell University, last modified March 31, 2020, https://arxiv.org/abs/1907.10454.

M. Cerezo et al., *Challenges and Opportunities in Quantum Machine Learning,* Cornell University, March 16, 2023, https://arxiv.org/pdf/2303.09491.pdf.

Michele Mosca and Marco Piani, *The Quantum Threat Timeline Report 2022*, Global Risk Institute, December 2022, https://quantum-safe.ca/wp-content/uploads/2023/01/2022-quantum-threat-timeline-report-dec.pdf.

Paul Smith-Goodson, "IBM On Track to Achieve Quantum Advantage by 2026 Using Error Mitigation," Forbes, December 7, 2022, https://www.forbes.com/sites/moorinsights/2022/12/07/ibm-on-track-to-achieve-quantum-advantage-by-2026-using-error-mitigation/?sh=1851139c12b5.

15. Open-Source Overload

Bob Ashley and Neil Wiley, "How the Intelligence Community Can Get Better at Open Source Intel," Defense One, July 16, 2021, https://www.defenseone.com/ideas/2021/07/intelligence-community-open-source/183789/.

*Hearing to Receive Testimony on Worldwide Threats, Before the United States Senate Committee on Armed Service*s, May 4, 2023, https://www.armed-services.senate.gov/imo/media/doc/23-44_05-04-2023.pdf.

Office of the Director of National Intelligence Senior Advisory Group Panel on Commercially Available Information, *Report to the Director of National Intelligence,* January 27, 2022, https://www.dni.gov/files/ODNI/documents/assessments/ODNI-Declassified-Report-on-CAI-January2022.pdf.

Ritu Gill, "What Is Open-Source Intelligence?" SANS, February 23, 2023, https://www.sans.org/blog/what-is-open-source-intelligence/.

"Volume of Data/Information Created, Captured, Copied, and Consumed Worldwide from 2010 to 2020, with Forecasts from 2021 to 2025," Statista, accessed August 2023, https://www.statista.com/statistics/871513/worldwide-data-created/.

19. Generative AI in the Wild

Ashish Vaswani et al., "Attention Is All You Need," arXiv preprint, June 2017, https://arxiv.org/abs/1706.03762.

Cecily Mauran, "Whoops, Samsung Workers Accidentally Leaked Trade Secrets via ChatGPT," Mashable, April 6, 2023, https://mashable.com/article/samsung-chatgpt-leak-details.

Daniel Levi, "ChatGPT Crosses 1 Million Users Five Days After Launch," Tech Startups, December 5, 2022, https://techstartups.com/2022/12/05/chatgpt-crosses-1-million-users-five-days-launch/.

Elizabeth Reid, "Supercharging Search with Generative AI," *The Keyword* (blog), May 10, 2023, https://blog.google/products/search/generative-ai-search/.

Frances Karamouzis, Erick Brethenoux, and Bern Elliot, "Beyond the Hype: Enterprise Impact of ChatGPT and Generative AI," (webinar, Gartner, on-demand), https://www.gartner.com/en/webinar/464445/1096048.

"Gartner Experts Answer the Top Generative AI Questions for Your Enterprise," Gartner, n.d., https://www.gartner.com/en/topics/generative-ai.

Marina Temkin, "In the World of Startup Valuations, There's Generative AI—and Everything Else," PitchBook, May 22, 2023, https://pitchbook.com/news/articles/early-stage-valuations-generative-AI-compare-VC.

Rowan Curran et al., *Predictions 2023: Artificial Intelligence,* Forrester, October 27th, 2022, https://www.forrester.com/report/predictions-2023-artificial-intelligence/RES178186.

Yusuf Mehdi, "Reinventing Search with a New AI-Powered Microsoft Bing and Edge, Your Copilot for the Web," *Official Microsoft Blog*, February 7, 2023, https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web/.

28. AI for Everyone

Katia Savchuk, "Do Investors Really Care About Gender Diversity? Shareholders Penalize Tech and Finance Companies for Not Hiring Women, New Research Shows," Insights by Stanford Business, September 17, 2019, https://stanford.io/2mng5EX.

Larry Hardesty, "Study Finds Gender and Skin-Type Bias in Commercial Artificial-Intelligence Systems: Examination of Facial-Analysis Software Shows Error Rate of 0.8 Percent for Light-Skinned Men, 34.7 Percent for Dark-Skinned Women," MIT News, February 11, 2018, https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212.

Remco Zwetsloot, James Dunham, Zachary Arnold, and Tina Huang, *Keeping Top AI Talent in the United States: Findings and Policy Options for International Graduate Student Retention,* Center for Security and Emerging Technology (CSET), December 2019, https://cset.georgetown.edu/wp-content/uploads/Keeping-Top-AI-Talent-in-the-United-States.pdf.

30. A Workforce Disrupted?

Chips and Science Act, an act that provides funds to support the domestic production of semiconductors and authorizes various programs and activities of the federal science agencies, 117th Congress, Public Law No: 117-167 (became law August 9, 2022), https://www.congress.gov/bill/117th-congress/house-bill/4346.

Final Report, National Security Commission on Artificial Intelligence, March 1, 2021, https://assets.foleon.com/eu-west-2/uploads-7e3kk3/48187/nscai_full_report_digital.04d6b124173c.pdf.

Future of Jobs Report, World Economic Forum, May 2023, https://www.weforum.org/reports/the-future-of-jobs-report-2023/.

Marc Andreessen, "Why AI Will Save the World," Andreessen Horowitz, June 6, 2023, https://a16z.com/2023/06/06/ai-will-save-the-world/.

Select Committee on Artificial Intelligence of the National Science and Technology Council, *National Artificial Intelligence Research and Development Strategic Plan 2023 Update*, Executive Office of the President of the United States, May 2023, https://www.whitehouse.gov/wp-content/uploads/2023/05/National-Artificial-Intelligence-Research-and-Development-Strategic-Plan-2023-Update.pdf.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," Cornell University, August 22, 2023, https://arxiv.org/pdf/2303.10130.pdf?.

36. The Age of Principled AI

"AlphaFold," Google DeepMind, accessed August 2023, https://www.deepmind.com/research/highlighted-research/alphafold.

"DoD Adopts Ethical Principles for Artificial Intelligence," U.S. Department of Defense, February 24, 2020, https://www.defense.gov/News/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/.

"European Artificial Intelligence Market Driven by the Need for More Efficient Operations and Enhanced Security, Says IDC," IDC, March 17, 2023, https://www.idc.com/getdoc.jsp?containerId=prEUR250505223.

Robert C. O'Brien, "Investments Are Needed to Stop the U.S. from Falling Behind China in AI," The Hill, June 7, 2023, https://thehill.com/opinion/national-security/4034909-investments-are-needed-to-stop-the-us-from-falling-behind-china-in-ai/.

Ryan Heath, "Schumer, Humbled by AI, Crafts Crash Course for Senate," July 18, 2023, https://www.axios.com/2023/07/18/schumer-ai-forums-senate.

"The Capacity for Moral Self-Correction in Large Language Models," Anthropic, February 15, 2023, https://www.anthropic.com/index/the-capacity-for-moral-self-correction-in-large-language-models.

"Worldwide Spending on AI-Centric Systems Forecast to Reach $154 Billion in 2023, According to IDC," IDC press release, March 7, 2023, https://www.idc.com/getdoc.jsp?containerId=prUS50454123.

46. Unlocking Human Potential

Andrew Flood and Richard J. Keegan, "Cognitive Resilience to Psychological Stress in Military Personnel," *Frontiers in Psychology*, 13 (March 16, 2022), https://doi.org/10.3389/fpsyg.2022.809003.

George E. Billman, "Heart Rate Variability — A Historical Perspective," *Frontiers in Physiology,* November 2011, https://doi.org/10.3389/fphys.2011.00086.

Hayley A. Young and David Benton, "Heart-rate Variability: A Biomarker to Study the Influence of Nutrition on Physiological and Psychological Health?" *Behavioural Pharmacology,* March 15, 2018, doi: 10.1097/FBP.0000000000000383.

Nathalie Pattyn et al., "Lessons from Special Forces Operators for Elite Team Sports Training: How to Make the Whole Greater Than the Sum of the Parts," *Frontiers in Sports and Active Living* 4 (March 21, 2022), https://doi.org/10.3389/fspor.2022.780767.

52. The Physical Impact of Data

Alan Weissberger, “Future Market Insights: Telecom Cloud Market CAGR at 15.2% from 2022–2032,” Telecom Cloud Market, IEEE Communications Society, October 7, 2022, https://techblog.comsoc.org/category/telecom-cloud-market/.

“BabyLM Challenge,” GitHub, accessed August 2023, https://babylm.github.io/.

“Carbon Free Energy for Google Cloud Regions,” Google Cloud, last updated August 29, 2023, https://cloud.google.com/sustainability/region-carbon.

David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean, *Carbon Emissions and Large Neural Network Training,* Cornell University, April 23, 2021, https://arxiv.org/ftp/arxiv/papers/2104/2104.10350.pdf.

Emma Strubell, Ananya Ganesh, and Andrew McCallum, *Energy and Policy Considerations for Deep Learning in NLP,* University of Massachusetts Amherst, June 2019, https://arxiv.org/pdf/1906.02243.pdf.

“Infrastructure as a Service Market Size & Share Analysis – Growth Trends & Forecasts (2023–2028),” Yahoo Finance, August 16, 2023.

“Net-Zero Emissions Operations by 2050, Including a 65% Reduction by 2030,” Office of the Chief Sustainability Officer, accessed August 2023, https://www.sustainability.gov/federalsustainabilityplan/emissions.html.

Nick Bild, “Going Green with TinyML,” hackster.io, accessed August 2023, https://www.hackster.io/news/going-green-with-tinyml-a44390ec6169.

Renee Cho, “AI’s Growing Carbon Footprint,” State of the Planet, June 9, 2023, https://news.climate.columbia.edu/2023/06/09/ais-growing-carbon-footprint.

“Software Carbon Intensity (SCI) Specification,” GitHub, accessed August 2023, https://github.com/Green-Software-Foundation/sci/blob/main/Software_Carbon_Intensity/Software_Carbon_Intensity_Specification.md.

Sylvain Duranton, Green IT Is No Longer an Option for the Tech Sector,” Forbes, February 27, 2023, https://www-forbes-com.cdn.ampproject.org/c/s/www.forbes.com/sites/sylvainduranton/2023/02/27/green-it-is-no-longer-an-option-for-the-tech-sector/amp/.

“What Is Green Software?,” Green Software Foundation, accessed August 2023, https://greensoftware.foundation/articles/what-is-green-software.

58. The Emerging Citizen Experience

David Sulek and Sahil Sanghvi, “Equity as a Platform: Applying a New Mindset to Scale Innovation,” MeriTalk, February 15, 2022, accessed July 19, 2023, https://www.meritalk.com/equity-as-a-platform-applying-a-new-mindset-to-scale-innovation/.

Partnership for Public Service, *The State of Public Trust in Government 2023*, May 30, 2023, https://ourpublicservice.org/publications/the-state-of-public-trust-in-government-2023/.

Sarah Perez, “FTC Warns That AI Technology Like ChatGPT Could ‘Turbocharge’ Fraud,” TechCrunch, April 18, 2023, https://techcrunch.com/2023/04/18/ftc-warns-congress-that-ai-technology-like-chatgpt-could-turbocharge-fraud-and-scams/.

“Transforming the Veteran’s Experience: Customer Experience Tool Powered by Advanced Analytics,” Booz Allen Hamilton, accessed July 13, 2023, https://www.boozallen.com/s/insight/thought-leadership/transforming-the-veterans-experience.html.

64. A Digital Tapestry

“CDAO,” Chief Digital and Artificial Intelligence Office, https://www.ai.mil/.

68. From Science to Practical Climate Resilience

“Carl Sagan Testifying before Congress in 1985 on Climate Change,” YouTube, December 10, 1985, https://www.youtube.com/watch?v=Wp-WiNXH6hI.

Kevin Brehm, Avery McEvoy, Connor Usry, and Mark Dyson, *Virtual Power Plants, Real Benefits: How Aggregating Distributed Energy Resources Can Benefit Communities, Society, and the Grid*, RMI, 2023, https://rmi.org/insight/virtual-power-plants-real-benefits.

74. Code Writing Code

ARK Investment Management, *Big Ideas 2023*, January 21, 2023, https://ark-invest.com/big-ideas-2023/artificial-intelligence/.

Matt Bornstein and Rajko Radovanovic, “Emerging Architectures for LLM Applications,” Andreessen Horowitz, n.d., https://a16z.com/2023/06/20/emerging-architectures-for-llm-applications/.

BOOZ ALLEN AUTHOR INFORMATION

Special thanks to all of the leaders from within Booz Allen who contributed to this collection and provided technical and mission expertise along the way.

Rebecca Allegar
Allegar_Rebecca@bah.com
Brad Beaulieu
Beaulieu_Brad@bah.com
Ryan Caulfield
Caulfield_Ryan@bah.com
Chris Christou
Christou_Chris@bah.com
Jarid Cottrell
Cottrell_Jarid@bah.com
Ron Craig
Craig_Ronald@bah.com
Michael Dahlberg
Dahlberg_Michael@bah.com
Ted Edwards
Edwards_William@bah.com
Colin Friedman
Friedman_Colin@bah.com
Logan Gibson
Gibson_Logan@bah.com
Josh Gordon
Gordon_Joshua@bah.com
Alex Hanson
Hanson_Alexander@bah.com
Michelle Harper
Harper_Michelle@bah.com
Will Healy
Healy_Will@bah.com
Ramon Hill
Hill_ThaddeusRamon@bah.com
Brianna Hogan
Hogan_Brianna@bah.com
Christopher Holder
Holder_Christopher@bah.com
Irik Johnson
Johnson_Irik@bah.com
Greg Kacprzynski
Kacprzynski_Gregory@bah.com
Melanie Laffin
Laffin_Melanie@bah.com
John Larson
Larson_John@bah.com
Jeff Leon
Leon_Jeffrey@bah.com
Isabella Bello Martinez
Martinez_Isabella@bah.com
Michael Martoccia
Martoccia_Michael@bah.com
Cameron Mayer
Mayer_Cameron@bah.com
Julie McPherson
McPherson_Julie@bah.com
Dia Metropolis
Metropolis_Dia@bah.com
Santiago Milian
Milian_Santiago@bah.com
Patrick Myers
Myers_Patrick@bah.com
Marie Nguyen
Nguyen_Marie@bah.com
Beau Oliver
Oliver_Beau@bah.com
Susan Penfield
Penfield_Susan@bah.com
Dave Prakash
Prakash_Dave@bah.com
Sonya Rahmani
Rahmani_Sonya@bah.com
Andie Rauta
Rauta_Andreia@bah.com
Joe Rohner
Rohner_Joseph@bah.com
Brian Rost
Rost_Brian@bah.com
Horacio Rozanski
Rozanski_Horacio@bah.com
Gabi Rubin
Rubin_Gabriela@bah.com
Sahil Sanghvi
Sanghvi_Sahil@bah.com
Aaron Sant-Miller
Sant-Miller_Aaron@bah.com
Geoff Schaefer
Schaefer_Geoffrey@bah.com
Munjeet Singh
Singh_Munjeet@bah.com
Alison Smith
Smith_Alison@bah.com
Megan Smith-Branch
Smith-Branch_Megan@bah.com
Ernest Sohn
Sohn_Ernest@bah.com
Brad Stone
Stone_Brad@bah.com
Josh Strosnider
Strosnider_Joshua@bah.com
Prachi Sukhatankar
Sukhatankar_Prachi@bah.com
Eric Syphard
Syphard_J_Eric@bah.com
Betty Thompson
Thompson_Betty@bah.com
Eric Zitz
Zitz_Eric@bah.com

VELOCITY, A BOOZ ALLEN PUBLICATION

CHIEF TECHNOLOGY OFFICER
Susan Penfield

OFFICE OF THE CTO
John Larson
Julie McPherson
Beau Oliver
Munjeet Singh
Brad Stone

EDITORIAL
Executive Editor: John Larson
Managing Editor: Elana Akman
Editorial Strategy: Abby Vaughan, Sahil Sanghvi
Editorial Consultant: Emily Primeaux, Dragonfly Editorial

DESIGN
Creative Direction: Chris Walker
Art Direction and Design: Brody Rose and David Everett

MARKETING AND COMMUNICATIONS
Head of Brand and Marketing: Sandhya Davis

For media inquiries please contact
Amanda Allison: Allison_Amanda@bah.com

COMMUNICATIONS AND DIGITAL
Special acknowledgements to individuals across content, digital, creative, brand, media, legal, website, accessibility, and project management for their support of this project.

NOTABLE CONTRIBUTIONS

ABOUT BOOZ ALLEN

Trusted to transform missions with the power of tomorrow’s technologies, Booz Allen Hamilton advances the nation’s most critical civil, defense, and national security priorities. We lead, invest, and invent where it’s needed most—at the forefront of complex missions, using innovation to define the future. We combine our in-depth expertise in AI and cybersecurity with leading-edge technology and engineering practices to deliver impactful solutions. Combining more than 100 years of strategic consulting expertise with the perspectives of diverse talent, we ensure results by integrating technology with an enduring focus on our clients. We’re first to the future—moving missions forward to realize our purpose: **Empower People to Change the World℠**

To learn more, visit boozallen.com.

Booz Allen®

Velocity, a Booz Allen publication, studies the complex issues that are emerging for mission and technology leaders on the front lines of government innovation.

BoozAllen.com/Velocity