

EXPERT EDITION

How desired outcomes help agencies plan for operational AI

Insights from

- DHS
- NSF
- VA





Put people at the core of your mission with AI

In government, you're always striving to meet the evolving needs of your constituents – seamlessly and efficiently. That's why agencies today embrace AI as a tool to enhance their services, increase citizen engagement, and maximize employee satisfaction.

[Read more](#)



TABLE OF CONTENTS

As AI goes mainstream, NSF programs try to 'respond to the moment' ... **4**

National AI Institute strives for uniformity across VHA's artificial intelligence approach **7**

New DHS task force looks to the future of AI use at the agency **9**

For success in AI, start with a clear idea of the outcomes **12**



Preparing for AI prime time

"Ain't no stopping us now. We're on the move."

That pithy lyric (and '70s one-hit wonder from McFadden & Whitehead) well describes where many agencies are in understanding the possibilities and delivering on the promises of artificial intelligence.

Agencies across government have pilots and use cases aplenty underway. But now the focus seems squarely on setting up appropriate frameworks and governance policies and identifying AI candidates for broad use.

As Gil Alterovitz, chief AI officer at the Department of Veterans Affairs, put it: "It's really interesting that sometimes it may be an idea that comes from a single person in one site that then informs us and allows us to learn about what potentially could scale." (*Read the full story about harmonizing VHA AI priorities.*)

Figuring that out often comes down to the outcomes the agency or organization is trying to achieve and then prioritizing the value of changing those outcomes against one another.

"What are you trying to do? What are the limitations of the current practice? What's new in your approach? Who cares?" said Amy Henninger, a senior advisor for advanced computing in the Science and Technology Directorate at the Homeland Security Department. And then she added, is it "better, faster, cheaper? And finally, is it compatible with the longer-term strategic vision and higher-order guidance?" (*Read the full story about DHS' AI efforts.*)

In the pages ahead, you'll learn how agencies are tackling just such questions and, we hope, pick up some worthwhile lessons learned and tips to inform your own AI efforts.

Vanessa Roberts
Editor, Custom Content
Federal News Network

As AI goes mainstream, NSF programs try to 'respond to the moment'

BY DAISY THORNTON



Federal agencies are increasingly excited by the results they're getting from training artificial intelligence systems using government data. But that's mostly what it is so far: training, with a carefully curated and limited dataset.

Now, a new program at the National Science Foundation is seeking answers to two questions:

- What happens when that same algorithm is fed real-world data that continues to expand and grow?
- Will the results remain as accurate?

NSF's new Safe Learning Enabled Systems program wants to apply principles of safety derived from building physical things like bridges and airplanes to software.

"How do we build safe software that isn't just safe when we first deploy it, but remains safe even as it continues to adapt to the data that it's given?" Michael Littman, director of the Division of Information and Intelligent Systems at NSF, said on [Federal Monthly Insights – Artificial Intelligence](#).

"This is an outstanding problem. This is not a small problem. But we've put that challenge out there to the research community, and we've gotten in what appear to be some really fantastic proposals."

In a way, these challenges look to go even further than the standard principles of safety. For physical construction, engineers must consider things like metal fatigue, where materials will degrade over time and eventually give out. The goal in that case is to maintain and preserve the materials as long as possible.

AI question: Can data keep improving the outcomes?

The difference with the AI program is that NSF wants to do more than maintain, Littman said. It wants to improve results over time through the introduction of more data. The idea is that the algorithm will hopefully learn to be better, rather than degrade. NSF wants to see more data yielding better results, rather than drifting biases and outcomes.

And that has a lot to do with the data itself. For example, if you train a facial recognition algorithm on 120 faces, with varied skin tones evenly distributed throughout, you'll get much better results than if you introduce 10,000 similarly hued faces.

"This is something that we understand both from the mathematical standpoint but also from the social science standpoint," Littman said on the [Federal Drive with Tom Temin](#).

"They all kind of agree that this is the phenomenon that you get: If you train with biased data that doesn't really cover the space particularly well, then these systems are going to fall into traps. They're going to take shortcuts because they're lazy. Fundamentally, these are lazy systems. They're just trying to do what the data tells them they need to do. So if you give them limited data, they're going to be limited systems."

Current efforts build on NSF's long-standing work in AI

That's not the only new program centered on AI currently in development at NSF. The agency has been involved in AI research for decades, but Littman said it's striving to "respond to the moment" now that AI has hit the mainstream and sparked imaginations, particularly in the area of generative AI.

"There's a tremendous amount of attention. There's a tremendous amount of opportunity. There's some new vistas to explore," he said. "And so we want to make sure that the academic community has the resources that they need to pursue these questions and really answer what society needs them to answer."

Littman said in his own division, there's the Information, Integration and Informatics program, which has been focused for years on the problem of bringing real-world data to bear on solving problems. That's foundational, because without a clear path to actually integrating real world data into AI, the question of safety and how to use that data to improve



If you train with biased data that doesn't really cover the space particularly well, then these systems are going to fall into traps. They're going to take shortcuts because they're lazy. Fundamentally, these are lazy systems. They're just trying to do what the data tells them they need to do.

– Michael Littman, Director of Information and Intelligent Systems Division, NSF



rather than degrade the algorithms would be merely academic.

Then there's the Human-Centered Computing program, focused on building systems that interact well with people. Eventually, these systems need to be able to explain themselves and how they arrived at their conclusions to the people who are using them, who likely won't have the technical expertise to figure that out themselves.

Right now, even the most technically proficient people are struggling with figuring out the

processes being used by things like neural networks or deep networks, he said.

Finally, there's the Robust Intelligence program, which revolves around understanding the limitations and power in the core algorithms that make AI and machine learning possible.

"And so these are just the core programs," Littman said. "Then we have a lot of programs that we've spun up to address particular problems that have come up in AI and computer security." 🤖

There's a tremendous amount of attention. There's a tremendous amount of opportunity. There's some new vistas to explore. And so we want to make sure that the academic community has the resources that they need to pursue these questions and really answer what society needs them to answer.

— NSF's Michael Littman



Listen to the complete discussion between the Federal Drive's Tom Temin and [NSF's Michael Littman on the agency's continue research into AI.](#)

Through its National AI Institute, VA intends to develop uniformity across AI use cases at VHA

BY DERACE LAUDERDALE

The National Artificial Intelligence Institute at the Veterans Affairs Department is developing an approach to bring uniformity to AI projects across the Veterans Health Administration and its medical centers.

The VHA plans align with broader VA efforts to integrate big data principles with a trustworthy AI framework at the department, which VA Secretary Denis McDonough announced during NAAI's International Summit for AI in Health Care. Given the high rates of burnout for health care workers generally, the agency also has an AI tech sprint focused on developing AI capabilities and testing use cases to improve worker experience. For instance, at the summit, VA displayed an AI touchscreen to help better support frontline care workers.

In addition to the touchscreen, the AI tech sprint will drive the development of additional AI capabilities to relieve worker burnout. One will be spoken language converted into text through AI and natural language processing when a patient is with a provider. The second will integrate AI with optical character recognition so that care providers can find information more quickly.



AI, like any other technology, really depends on essentially what is your use case and what you want to do.

— Gil Alterovitz, Chief AI Officer, VA

AI goal at VHA: Provide faster access to critical care information

"So all of these are really ways to make the systems work faster. They augment what the provider is doing. And of course, the provider then has the ability to then leverage that to decide what to do moving forward," Gil Alterovitz, VA chief AI officer and director of NAAI, said on [Federal Monthly Insights - Artificial Intelligence](#). "AI, like any other technology, really depends on essentially what is your use case and what you want to do."

Through iterative development and pilots, NAI hopes to expand the department's use of AI, particularly in ways that can help improve its health care and medical services, Alterovitz said.

"There's a lot of really exciting work happening right now, and we are focused on a number of different priorities and developing AI capabilities to allow for both translating from research and development and moving forward to enable operational use cases of artificial intelligence," he said on the [Federal Drive with Tom Temin](#). "In doing so, we have created a network of National Artificial Intelligence Institute centers across the United States."

Looking to veterans and employees for feedback


A key component for moving from R&D to operational technology will be getting input from veterans and its medical center staffs, Alterovitz said. Feedback and recommendations will help VA prioritize its AI projects and decide which it should scale up first.

That's needed, he said, because VHA medical centers now have many use cases underway. The positive side of all of this early work is that the centers have fostered a community of people who are interested in AI and continue to develop new ideas, he added.

It's really interesting that sometimes it may be an idea that comes from a single person in one site that then informs us and allows us to learn about what potentially could scale.

— VA's Gil Alterovitz

"There's an approach being developed around governance within the Veterans Health Administration as well, to ensure and bring uniformity across the different parts of the different medical centers across the country and different regions," Alterovitz said. "And so work is ongoing there. It's really interesting that sometimes it may be an idea that comes from a single person in one site that then informs us and allows us to learn about what potentially could scale." 🔄

 Listen to the complete discussion between the Federal Drive's Tom Temin and [VA's Gil Alterovitz on prioritizing AI initiatives across VHA and its medical centers](#).

New DHS task force looks to the future of AI use at the agency

BY MICHELE SANDIFORD

Like the majority of industry, the federal government is grappling with how to apply artificial intelligence to its mission. For the Homeland Security Department, data collection and AI impacts not just how the agency manages national security issues, but has public-facing impact on American citizens, their personal data and potential implications regarding everything from travel to criminal investigations.

"AI is somewhat nascent at DHS. We're really just starting to get into it in a big way. And I think that's demonstrated by the recent stand-up of the AI Task Force," Amy Henninger, a senior advisor for advanced computing in DHS' Science and Technology Directorate, said on [Federal Monthly Insights - Artificial Intelligence](#).

The new task force will define the goals and strategy for the use of AI across the department and its component agencies, Henninger said. She added that DHS is looking for applications of the technology that "have the most bang for the buck."



Setting DHS priorities for AI adoption

DHS is working to prioritize the use of AI for functions that would be best served by the technology in speed, data management and staying ahead of adversaries.

Henninger noted on the [Federal Drive with Tom Temin](#) that the priorities coming from the DHS secretary through the AI Task Force are "looking at applications related to border security, supply chain security, disruption of fentanyl or synthetic opioid production."

Additionally, the department is "very interested in digital forensics and combating online child exploitation," she said. "We're very interested in adversarial AI, the types of AI security above and beyond cybersecurity that AI enables new attack services for. And then we're looking at AI applied to critical infrastructure."

Speaking about the approach the agency is using to prioritize AI applications, Henninger said she asks herself several questions that



derive from an approach developed by the former chief of the Defense Advanced Research Projects Agency, George Heilmeier: “What are you trying to do? What are the limitations of the current practice? What’s new in your approach? Who cares?”

The agency also looks at whether a proposed AI capability will provide sufficient value. Is it “better, faster, cheaper?” she said. “And finally, is it compatible with the longer-term strategic vision and higher-order guidance?”

Creating AI data use best practices

Regarding how DHS is evaluating the use of AI related to data collection, storage, use and regard to its future implications, Henninger said, “The other thing you have to make sure you have access to is appropriate data in terms of quality, quantity ... sufficiency. If you don’t have that for your AI project, it’s pretty



AI is somewhat nascent at DHS. We’re really just starting to get into it in a big way.

— Amy Henninger, Senior Advisor for Advanced Computing, Science and Technology Directorate, DHS

much dead on arrival. And you have to make sure it’s consistent with ethical and regulatory considerations.”

She pointed out that the department has extensive regulations and policies in place to protect citizen data.

“Coming into DHS from the Defense Department, that was really sort of a whole new thing for me to get used to and to understand,”

Henninger said. “I can report as a citizen, I am thrilled and so grateful that DHS is as careful with our data, our private data, as we are. So I was a little bit shocked at all the oversight and regulations and policies making sure that our citizen data is kept private.”

Another issue faced by the DHS task force is how to monitor AI use for bias in datasets. “You have to be very, very careful about watching for unstable trajectories because the model and execution starts to diverge from the initial training datasets in ways that are hard for the model to correct for,” she said.

Regarding the speed of automation and risks brought on by large scale data collection,


The other thing you have to make sure you have access to is appropriate data in terms of quality, quantity ... sufficiency. If you don’t have that for your AI project, it’s pretty much dead on arrival.

— DHS’ Amy Henninger

Henninger said, “all of these AI-based systems especially can be tricked and spoofed or manipulated in nefarious ways.”

With the likelihood that most adversarial AI will be automated, the government will need to be able to use comparable automation to counter attacks and keep ahead of its adversaries, she said.

Henninger said that there is the risk that at some point the amount of data collected can become slightly overwhelming, but the technology itself will make the process of monitoring information more manageable and effective. 🤖



Listen to the complete discussion between the Federal Drive’s Tom Temin and [DHS’ Amy Henninger on the department’s plans to plot an AI roadmap.](#)

For success in AI, start with a clear idea of the outcomes

In some ways, getting started presents the most difficult part of deploying artificial intelligence. One way to get started though: Think less about algorithms and training data and more about the business problems you'd most like to solve.

Peter van der Putten, director of the AI Lab at [Pegasystems](#), offered these questions: "What is it that you're trying to achieve? Is it about providing customer-centric, citizen-centric service? Or is it more focused on doing more with less and improving your productivity?"

As for AI itself, van der Putten added, "It's not just technology. It's also a way of working, becoming more evidence-based as an organization." He said AI is also a way of helping employees focus on higher-value work by removing repetitive or routine tasks. In such cases, think of AI as augmented intelligence. If used effectively and ethically, an agency can think of AI as "accepted intelligence," he suggested.

In an interview on [Federal Monthly Insights – Artificial Intelligence](#), van der Putten named a third application

domain for AI that's neither improving direct online experiences for constituents nor quite improving internal operations. Namely, showing people how the government is improving access to information and services in the first place. He called that "nudging, similar to how commercial companies are nudging their customers to their products and services."

AI-powered Q&As, for instance, can give people personalized views of the programs and services most relevant to them, van der Putten offered.

This all means that while the technology staff must be ready with the data and computing



It's not just technology. It's also a way of working, becoming more evidence-based as an organization.

— Peter van der Putten, Director of AI Lab, Pegasystems

resources to handle AI, decisions about where and how to use AI require a team approach.

"Ultimately, the entire organization needs to think about strategic priorities," van der Putten said.

Deploy AI based on where it delivers most value

AI projects generally don't require armies of data architects and AI specialists, van der Putten suggested, but rather some careful thinking.

"It is much better to start with outcomes, the various parts of the organization where we want to improve," he said. Then, work backward toward the processes that contribute to the outcomes. He used the example of case backlogs, where AI might improve the routing of cases in a manner more closely aligned to the skills of the people making decisions.

Van der Putten said it's important to continuously measure outcomes and check in with the people whose work AI is aimed at supporting, as they are the experts. He recommended an incremental approach, applying the technology to one process at a time, constantly

Maybe we shouldn't start with thinking about data. We should flip it around and think about the business priority. What kind of automation do support agents and citizens need to improve their respective journeys?

— Pegasystems' Peter van der Putten

feeding data back to the algorithm for continuous improvement.

Such an approach can make modernization dollars go far, van der Putten said. Often AI can help integrate legacy applications and improve process outcomes.

"If you want to modernize, you maybe shouldn't rip and replace but think more in terms of putting an agility layer on top" of existing applications, he said. That lets the agency improve the experience or solve internal logjams, while buying time to rework backend applications in a more measured way.

Applications, of course, don't work without data. When thinking about the data that will inform and then work operationally with AI, van der Putten again advised starting with

the desired business outcomes and decision processes that support them.

Data matters but may not be the first priority when planning AI use

Too often, people start thinking about data immediately.


“Maybe we shouldn’t start with thinking about data. We should flip it around and think about the business priority,” he said. “What kind of automation do support agents and citizens need to improve their respective journeys?”

That line of thinking prompts a sharp focus on precisely what data the AI deployment will need, such as from natural language processing or inbound emails.

“You work your way back from the top,” van der Putten said. An agency may discover that a small, carefully curated dataset will achieve the most effective and bias-free algorithm performance. Avoid thinking that the more data is thrown at an algorithm, the better it is, he said.

It’s also important to periodically

retest algorithms to make sure results don’t drift over time.

“We need to continuously monitor all the decisions [for] whether there’s bias in our systems. Rules change quite often,” van der Putten said. “So every time we make changes to rules or changes to models, we need to retest.” He suggested keeping synthetic datasets on hand to recalibrate AI systems, adding that agencies might want to consider generative AI services to build reference datasets. 

Watch and listen to the complete discussion between the Federal Drive’s Tom Temin and [Pegasystems’ Peter van der Putten on why desired outcomes should define an organization’s AI priorities.](#)