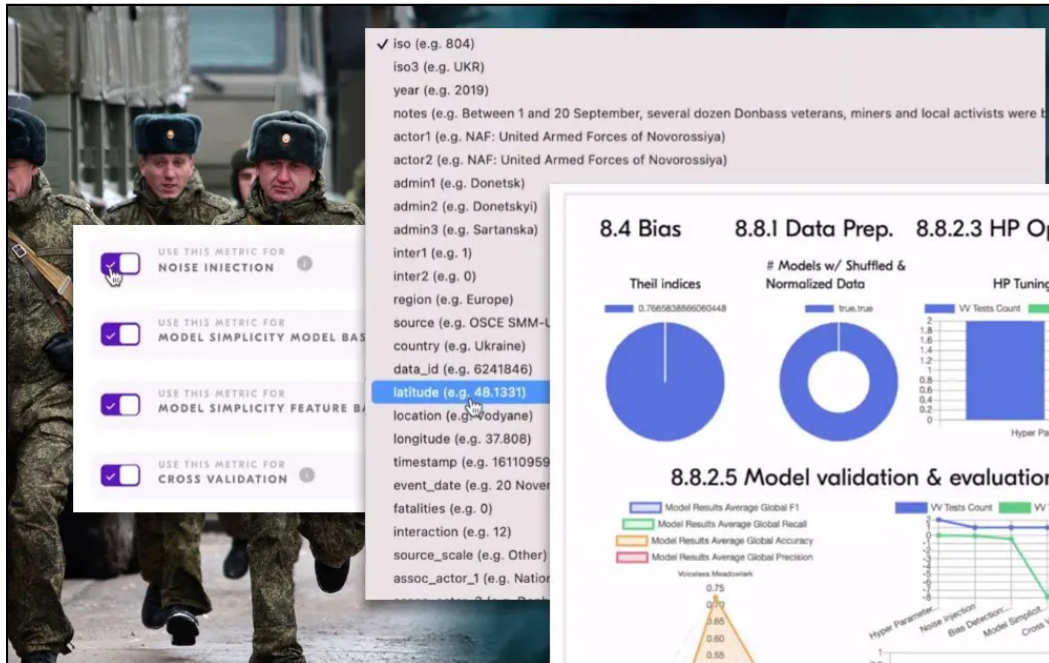


US Federal and Military AI: New Platform Offers Algorithm Validation and Accreditation


Originally published at <https://www.unite.ai/usa-federal-and-military-ai-vespr/>



A startup that features an advisory board of former government military luminaries has released a new platform designed to evaluate the security and deployability of AI applications. Early adopters of the system are said to include the U.S. Air Force and the Department of Homeland Security.

The platform is called VESPR, from CalypsoAI, founded in 2018, with headquarters in Silicon Valley, Dublin, and an officially 'undisclosed' location in Virginia – terra firma for the CIA at Langley.

VESPR is a model risk management (MRM) system designed to facilitate a federally-compliant system of accreditation for deployed algorithms. It offers both a user-friendly dashboard-style GUI environment, and a CLI interface for more advanced usage.



Click to enlarge. Source: <https://www.youtube.com/watch?v=IMhS6j7t2pl>

VESPR is founded on CalypsoAI's machine learning validation, verification and accreditation standards, and features hand-crafted adversarial machine learning libraries. It also offers automated stress-testing routines for potential deployable algorithms.

National Artificial Intelligence Research Resource Task Force

The timing of the release may relate to yesterday's [launch](#) by the Biden administration of a new National Artificial Intelligence Research Resource Task Force, a body designed to serve as a federal advisory committee pursuant to Congress's [National AI Initiative Act of 2020](#).

Pressure has been mounting on the USA, and around the world, for meaningful regulatory standards for machine learning systems, not least in mission-critical areas such as key infrastructure and military usage. Since ML systems are still enduring a formative stage and a rapid rate of progress, they represent a relatively unstable and frequently controversial recourse from which it's now essential to identify replicable and reliable analytical algorithms – if this proves to be possible.


In April CalypsoAI posted its support of the [Endless Frontiers Act](#), a congressional bill designed to reform science funding in the face of China's growing eminence as an AI power, though the act was ultimately [watered down](#) at the Senate stage.

Validation For Federal AI

According to the VESPR press release, areas covered by the framework include computer vision and natural language processing (NLP).

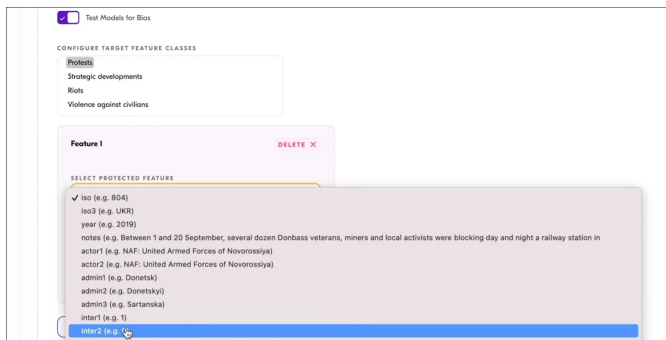
CalypsoAI claims that VESPR was created 'with critical input from existing national security customers and born out of years of independently funded research into adversarial machine learning'.

Images of the system seen in a promotional video (see end of article) appear to include detection and/or simulation routines for data poisoning and noise injection, providing simulations for the actions of potential attackers to deployed systems.




Click to enlarge.

The system seems to utilize historic data of national as well as foreign action. The target classes featured include 'Protests' and 'Riots', in addition to the less clear 'Strategic developments'. National terrorism incidents also seem to be included in the system's reference databases, with 'Violence against civilians' another available target class. Other available target classes include 'Battles' and 'Explosions/Remote Violence'.



Click to enlarge.

The system appears to allow for the protection of features in a “BIAS management” section of configuration, apparently designed to combat overfitting or to avoid the unwanted elimination of minor outlier events that may be of interest in an analytical routine. In the video, VESPR is processing tabular historical data about ‘Ukraine’.



Beyond this initial promotional blitz, it’s not likely (perhaps by design) that we’re going to hear much more about this government-facing SaaS product, however it fares; it shares its name with a coffee bar franchise, a social dating app, and a streaming album, and is relentlessly beaten back in ranked results by the VSEPR chemistry model.

CalypsoAI received [\\$13 million](#) in Series A fundraising from Paladin Capital Venture Group in July Of 2020. Other investors [included](#) 8VC, Lockheed Martin  LMT -0.36% Ventures, Manta Ray Ventures, Frontline Ventures, Lightspeed Venture Partners and Pallas Ventures.

In a blog post on the company site, CalypsoAI’s founder Neil Serebryany, who undertook unspecified research work at the Department of Defense in 2018, states that the company was founded as a possible solution to governments’ fear of deploying advanced algorithmic systems in an unregulated climate:

‘The primary reason for this fear of AI projects, leading to them being abandoned inside the Government, sounds prosaic, but is actually quite complex. They were being abandoned due to a lack of quality assurance [...] AI models cannot be assessed the same way traditional software models are. This is due to the underlying nature of model structure and the highly-complex ways in which they can fail. Lacking a mechanism to assess these non-deterministic systems in a deterministic, auditable way, organizations

inside the government were unable to assess the so-called 'quality' of AI models against a benchmark. This led to fear that they could fail, could malfunction, or could be hacked by an adversary at the time they are needed most for example in combat, in flight, or during a complex medical procedure. '

Advisory Board

A month prior to the investment round, the company created a National Security Advisory Board including **Tony DeMartino**, a former aide to Defense Secretary Jim Mattis, and now a founding partner of Washington-based strategic advisory firm Pallas Advisors; former Principal Deputy Under Secretary of Defense for Intelligence (under President Trump) **Kari Bingen**; former Associate Deputy Director for CIA Digital Innovation **Sean Roche**, an erstwhile cyber-intelligence specialist in that organization; and **Michael Molino**, former Executive Vice President of corporate development at ASRC Federal, which provides advisory, research and technical migration capabilities across a range of critical federal services.

According to the release:

'VESPR provides advanced AI testing capabilities with a streamlined workflow to ensure that every machine learning algorithm put into production has been verified secure. VESPR provides unparalleled security and assurance to a variety of AI systems, from computer vision to natural language processing. The VESPR process ensures testing, evaluation, verification, and validation (TEVV) throughout the secure machine learning lifecycle (SMLC), from the research and development phase through model deployment. The end result is AI systems that provide accurate and comprehensive monitoring and reporting on model capabilities, vulnerabilities, and performance.'

Updated 11:07 AM EST to reflect that Michael Molino no longer works at ASRC Federal, an errata in original article.

Updated 1 January 28, 2024 to remove broken YouTube video.

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More