Hewlett Packard Enterprise

HPE's Response to NIST RFI on Developing a Federal AI Standards Engagement Plan

The hyper-connected world makes new things possible, now…

June 10, 2019

Contacts:
David Showalter
Account Manager
Tel.: 410-610-2960
David.Showalter@hpe.com

Shawn M. Vasell
Senior Director, Corporate Affairs
Tel: 202-772-4464
Shawn.Vasell@hpe.com

Executive Summary

Hewlett Packard Enterprise knows artificial intelligence (AI) can amplify human capabilities and turn exponentially growing data into insight, action, and value for governments and businesses around the world. That's why we appreciate the opportunity to respond to the Request for Information (RFI) on AI Standards issued by the National Institute of Standards and Technology (NIST). We are excited to share with you HPE's approach and our current understanding of the:

- current status and plans regarding the availability, use, and development of AI technical standards and tools in support of reliable, robust, and trustworthy systems that use AI technologies;
- needs and challenges regarding the existence, availability, use, and development of AI standards and tools; and
- current and potential future role of Federal agencies regarding the existence, availability, use, and development of AI technical standards and tools in order to meet the nation's needs.

Our experience and portfolio of solutions developed in collaboration with our industry partners advances our ability to help make artificial intelligence work for many of our existing customers, regardless of where they are in the adoption of AI. Furthermore, our long-standing partnership with the United States government has been a driving force in developing product standardizations that have been the cornerstone of our industry.

Hewlett Packard Enterprise has been in the innovation business for more than 75 years. Our vast intellectual property portfolio and global research and development capabilities are part of an innovation roadmap designed to help organizations of all sizes—from global enterprises to local startups—transition from traditional technology platforms to the IT systems of the future.

Innovation is a key element of our culture and critical to our success. Our research and development efforts are focused on designing and developing products, services, and solutions that anticipate customers' changing needs and desires and emerging technological trends.

Hewlett Packard Labs, together with the various research and development groups within our business segments, is responsible for our research and development efforts. Hewlett Packard Labs is part of our Corporate Investments segment.

Expenditures for research and development were $1.7 billion in fiscal 2018, $1.5 billion in fiscal 2017, and $1.7 billion in fiscal 2016. We anticipate that we will continue to have significant research and development expenditures in the future to provide a continuing flow of innovative, high-quality products and services to maintain and enhance our competitive position.

We applaud the efforts of NIST in developing this AI Standardization Engagement Plan. We encourage NIST to collaborate with industry, participate in existing international standardization efforts, and maintain a technology neutral policy. We believe this approach will accelerate the tooling, testing, and deployment of AI technology and certainly help in achieving U.S. leadership. HPE is pleased to offer any assistance deemed necessary to promote this effort.

HPE Response

The development of AI technical standards and tools in support of reliable, robust, and trustworthy systems that use AI technologies is becoming increasingly urgent due to a market that is changing faster than ever before. Digital transformation is happening everywhere – as everyone and everything is connected and sharing data: a manufacturing floor, a building, a campus, a city, your house, a crop field, a wind farm, a power plant, an oil rig, a sports arena, your running shoe, a coffee machine, your car, and much more. This explosion of connected things is happening all around us, at the edge of the digital world, and AI technology will play a critical role in the successful deployment of these solutions.

As governments and industries contemplate the options in facilitating the development of these technologies, policies, regulations, and laws, the role of standardization is paramount in the dialogue to ensure meaningful and practical implementations. In responding to the RFI, HPE will present our recommendations based on the three areas of focus you identified:

1. AI Standards and Tools
2. Defining and Achieving U.S. AI Technical Standards Leadership
3. Prioritizing Federal Government Engagement in AI Standardization

AI Standards and Tools

AI should be reproducible and bias-free. HPE sees a need for the community to develop and evangelize voluntary, industry-led standards around the creation, maintenance, and security of open/public datasets used for machine training. In theory everyone should be able to download a public dataset and train an AI that will demonstrate similar/identical behavior, allowing reproducibility, regardless of the hardware or conditions of training.

However, you cannot improve what you cannot measure. HPE, in conjunction with Hewlett Packard Labs, created an open source performance benchmark suite for comparing DL frameworks, models, and compute platforms. The code for the Deep Learning Benchmark Suite is available on GitHub: https://github.com/HewlettPackard/dlcookbook-dlbs.

This suite is vendor agnostic and supports user-defined workloads, models, and frameworks. It is extensible and supports accelerators and measures units of work per wall clock time so the impact or server, network, and storage variations are reflected in the results.

Additionally, there are a number of tools being developed industry-wide as enabling technology to advance recent revival of artificial intelligence. Already adopted by a handful of tech giants and industry leaders, deep learning has the potential to change the way we work, across industries. At HPE, we see a lot of interest in this technology from our customers, but many of them don't have the expertise and resources of tech giants to make informed decisions on optimal hardware and software configurations to run deep learning workloads efficiently. To remedy that, we have created a new tool set enabler, the Deep Learning Cookbook.

The Deep Learning Cookbook is based on a massive collection of performance results for different deep learning workloads on different hardware/software (HW/SW) stacks, and

analytical performance models. This combination enables us to estimate the performance of a given workload and to recommend an optimal HW/SW stack for that workload.

Currently, our Cookbook is based on the extensive benchmarking of eleven workloads with eight deep learning frameworks and six hardware systems. These include the most popular open source frameworks such as TensorFlow, Caffe, and Caffe2, and hardware systems from different HPE product lines, including HPE Apollo 6500, HPE Apollo 6000, and HPE Edgeline. And this list keeps on growing.

The results of the benchmarks are incorporated into our models to predict performance for all possible system configurations, even those which we have not yet had the chance to evaluate.

We believe the Deep Learning Cookbook will make it easier to employ deep learning in real world vertical applications. Additionally, we use the Cookbook to detect bottlenecks in deep learning workloads and to guide the design of future HPE systems for artificial intelligence and deep learning.

Benchmarks for machine learning are necessary extensions to the tools and guidelines developed to ensure robust AI technologies. As mentioned earlier, if you cannot measure it, you cannot improve it. HPE would welcome engagement with the community on the definition and promotion of key AI benchmarks that can help the industry/community assess behavior and performance of disparate solutions. Alignment with DOE labs would be productive and HPE could help engage. Our long standing partnership with Federal agencies such has DOE and our own Hewlett Packard Labs has spawned many innovative developments.

Sub-industry Focus Area

AI for health and public safety would benefit from a standard set so vendors can be free to innovate without fear of frivolous lawsuit as long as they meet or exceed the standard. A possible standard could be if the accuracy of the AI results exceed that of a well-trained medical professional (i.e. if the norm for a radiologist is 90% accuracy, then any AI used to supplement this process must demonstrate 90% or better accuracy in the same task). As long as the AI meets this standard then the vendor should be held blameless.

Otherwise, AI in healthcare will continue to lag because vendors will fear lawsuits and medical practitioners will avoid the perceived risk.

Defining and Achieving U.S. AI Technical Standards Leadership

In the United States, a number of companies are currently leveraging open source software provisions to collaborate on research and innovation in the development of AI solutions. These efforts have spawned projects that will enable standards to establish foundational concepts and other technical aspects to enable interoperable AI. Today SC42 (Artificial Intelligence) under the Joint Technical Committee 1 (JTC 1) is the primary standards development technical committee in which AI standards are being developed. The United States is represented on the committee by the American National Standards Institute (ANSI) and we believe NIST should continue to participate and support these efforts and establish standards to achieve/maintain leadership.

The financial services industry (FSI) is leading the adoption of AI because it realized early on the competitive advantage and ROI of making AI aided financial decisions. FSI manages their own risk/reward balance and standards are of little value here.

While healthcare shows the greatest potential benefit, it is the farthest behind in the potential adoption of AI. Healthcare AI R&D is robust, but it has not crossed over into patient care. Similar to how fear of malpractice suits have resulted in “defensive medicine”, this fear has slowed the transition of AI from R&D to clinical use. (Note comment in section 1, Sub-industry focus)

The U.S. maintains a lead in high performance computing thanks to government and NSF grants to US Labs, research organizations, and universities, but there is no similar push for large scale AI infrastructure. NIST should actively advocate similar research grants to foster the growth of AI research within academia and other public forums.

The Executive Order on Maintaining American Leadership in Artificial Intelligence states all government agencies must adopt AI whenever possible. Policies to measure this progress and incentivize increased compliance are essential next steps to attain a worldwide leadership position.

AI as a Service (AIaaS) needs to be made available to all agencies for their AI development to speed up adoption. DoD related agencies are ahead and have begun building AIaaS. Civilian agencies are behind and could benefit most from the AIaaS model. Pittsburgh Supercomputing Center (PSC) “Bridges” cluster is a great example of such a pan-organizational model that could jumpstart many Civilian Agencies while they develop their AI strategy.

Prioritizing Federal Government Engagement in AI Standardization

The unique needs of the Federal government are such that all data used must be openly available for others to determine its:

1. Lack of bias
2. Transparency
3. Anonymity/Privacy
4. Security

The DoD Joint Artificial Intelligence Center (JAIC) has established AI guidelines and was the first to create a focal point for AI research and practices with the DoD. The civilian agencies could benefit from a similar JAIC-like entity with the added benefit of providing cross agency resources, including AIaaS.

It is our strong assertion the full and open access to big data will go a long way in facilitating more rapid development implementation of AI solutions across a broader spectrum of users and citizens. Without this access to data for the purpose of R&D those in small startups and academia have very little recourse. NIST and other government agencies should use their leverage to encourage leaders of industry in all sectors to share the wealth in making data more broadly available.

HPE would encourage NIST to participate in international standards development organizations and their efforts to develop standards that define safe, robust and trustworthy AI technologies. Industry growth will be enabled by the ability to clearly communicate and validate these key principles. These are examples of standards work we are following today:

I. The IEEE Global Initiative on Ethics of Autonomous and Intelligent System
II. The European Commission's High-Level Expert Group on Artificial Intelligence – current "Draft Ethics Guidelines for Trustworthy AI"

There are clearly limitations on the reach of a single government entity to effect global technological development, without the collaboration of the industry at large. The following questions are salient to the effectiveness of the ongoing effort of NIST:

- Is there an opportunity for NIST to encourage technical research into methods of transparency in AI decision making?
- Are there measurement techniques that could be lead to better auditability of AI systems?

NIST should consider engaging in research and/or partnering with industry to explore techniques for new/improved risk measurement of AI technologies. There are both government applications and commercial applications where this will be critical for adoption/acceptance.

HPE is poised to advocate on behalf of our worldwide customers, for standards that will assist the development of AI solutions, and the consumption of AI solutions that support significant business outcomes.