STRATEGIC POLICIES IN THE AI RACE

AI Political Strategy in the USA

Peter Stone



Araya Peralta

Compared to other countries, the USA has been relatively slow to develop a national strategy pertaining specifically to Artificial Intelligence. However that has not slowed down the rate of progress in American academia and industry that has led to many noteworthy AI technical advances over the past several years, both in fundamental algorithms and in practical applications. This high rate of AI-related technological progress shows no sign of slowing down. Meanwhile, the federal government has recently become more proactive in its organization of a national strategy and providing guidance and possibly new resources for AI development.

*

Artificial Intelligence (AI) is changing rapidly in all ways. Exciting new technical results are being published regularly, thanks both to the high public interest drawing many top students to the field, and also to the huge investments in AI being made by industry. As a result, anything written about the “state of the art” in AI, whether focussing on its capabilities, its limitations, or the technical details of how “AI” works, risks being obsolete even before it is published.

The political landscape pertaining to AI is no exception. Governments around the world are actively seeking ways to encourage AI innovation within their boundaries, while simultaneously wrestling with how best to regulate AI-based technologies. As a result, the political landscape is ever-shifting.

As a case in point, when I was first asked to write this article, in March of 2019, the

narrative would have been relatively pessimistic. The federal government had paid lip service to AI being important to the nation. But no concrete plans or funding had been announced. Seven months later, as I write this article in October, I'm pleased to be able to give a much more positive and optimistic report.

> The rate of progress in American academia and industry that has led to many noteworthy AI technical advances over the past several years shows no sign of slowing down

In the remainder of this article, I summarize the documents, events, and reports that have come to my attention pertaining to the AI political strategy in the USA. Note that I do so mainly as an external observer -an American academic researcher living in the US and employeed by an American university. I do not have any governmental role, nor any inside knowledge of future plans, which in all likelihood will change the landscape even further after this writing. I also make no claim to covering the landscape in any comprehensive way. That being said, I focus in on some notable activities at the end of the Obama administration in 2016, some non-Governmental activities from around the same time, and finally the relatively recent follow-up from the Trump administration in 2019.

2016 White House Reports and National Strategic Plan

In 2016, during the final months of the Obama administration, there was a flury of activity that laid the groundwork for concrete action pertaining to AI policy. The Office of Science Technology and Policy (OSTP) – specifically the Subcommittee on Machine Learning and Artificial Intelligence – sponsored a series of workshops between May and July under the heading "Preparing for the Future of Artificial Intelligence." These workshops led directly to a National AI research and Development Strategic Plan that was published in October 2016 by the National Science and Technology Council (NSTC). While these events and document were valuable steps, they unfortunately did not lead to concrete action in the near term.

The OSTP workshops were spearheaded by Ed Felten, a Professor at Princeton who was serving at the time as Deputy U.S. Chief Technology Officer. The plans for the workshops were announced in May 2016. The workshop topics and dates were as follows:

- Legal and Governance Implications of Artificial Intelligence: May 2016
- Artificial Intelligence for Social Good: June 2016
- Safety and Control for Artificial Intelligence: June 2016
- The Social and Economic Implications of Artificial Intelligence Technologies in the Near-Term: July 2016

The output from these workshops, along with one additional workshop on "Future of AI:

Emerging Topics and Societal Benefit" at the Global Entrepreneurship Summit (June 2016) is summarized in a document entitled "Preparing for The Future of Artificial Intelligence" posted in October 2016.

Meanwhile, in a related effort also directed from the White House, an AI Task force comprising members from multiple agencies and led by the NSTC put out a National AI Research and Development Strategic Plan, also in October 2016.

In that report, the following seven strategies were identified:

1. Make long-term investments in AI research.
2. Develop effective methods for human-AI.
3. Understand and address the ethical, legal, and societal implications of AI.
4. Ensure the safety and security of AI systems.
5. Develop shared public datasets and environments for AI training and testing.
6. Measure and evaluate AI technologies through standards and benchmarks.
7. Better understand the national AI Research and Development (R&D) work force needs.

It further made two recommendations as follows:

- Develop an AI R&D implementation framework to identify Science and Technology opportunities and support effective coordination of AI R&D investments, consistent with Strategies 1-6.
- Study the national landscape for creating and sustaining a healthy AI R&D workforce, consistent with Strategy number 7.

The task force was co-chaired by Prof. Lynne Parker from the University of Tennessee who was serving at the time as Division Director of Information and Intelligent Systems at the National Science Foundation. She described the process of creating the strategic plan in detail in an article published by AI Magazine.

2016 non-Governmental Efforts

Also in 2016, there were two notable efforts from outside the government: the Partnership on AI and the 100 Year Study on AI.

Partnership on AI was founded in late 2016 by six large companies, based mainly in the USA. Specifically, Apple, Amazon, DeepMind and Google, Facebook, IBM, and Microsoft banded together for the purpose of the following four stated goals:

- Develop and share best practices.
- Advance public understanding.
- Provide an open and inclusive platform for discussion & engagement.

- Identify and foster aspirational efforts in AI for socially beneficial purposes.

There are now more than 90 members from 13 countries [1]. In addition to the stated goals, the organization could be seen as an effort to self-regulate and thus forestall and/or moderate eventual government regulation.

On a related note, the inaugural report from the 100 Year Study on AI strongly cautioned against over-regulation::

“The Study Panel’s consensus is that attempts to regulate ‘AI’ in general would be misguided, since there is no clear definition of AI (it isn’t any one thing), and the risks and considerations are very different in different domains. Instead, policymakers should recognize that to varying degrees and over time, various industries will need distinct, appropriate, regulations that touch on software built using AI or incorporating AI in some way.”

The 100 Year Study on AI is a longitudinal study designed “to anticipate how the effects of artificial intelligence will ripple through every aspect of how people work, live, and play.” Reports are planned every five years [2].

Along with Prof. Barbara Grosz from Harvard, I summarized the process of writing this report in an article in the Communications of the ACM.

Concrete Progress in 2019

After the flurry of governmental activity in 2016, there wasn’t much visible progress from the federal government in the following two years. Few concrete steps that I am aware of were taken to enact the recommendations from the strategic plan. As a result, it seemed that the USA was falling markedly behind other countries and regions in this regard. However, the situation has changed markedly over the past several months with an update of the strategic plan in June 2019 and especially a call for National AI Institutes in October of this same year.

In March 2018, the White House did host a Summit on “Artificial Intelligence for American Industry” as a way to bring together “more than 100 senior government officials, technical experts from top academic institutions, heads of industrial research labs, and American business leaders who are adopting AI technologies to benefit their customers, workers, and shareholders”.

Then in February 2019, the President issued an executive order on “Maintaining American Leadership in Artificial Intelligence“. As described in a companion posting on “Accelerating America’s Leadership in Artificial Intelligence” this executive order, which launched the American AI Initiative, directed Federal agencies to prioritize AI investments in their R&D missions, while making data, models, and computing resources relevant to AI more available to America’s AI R&D experts. It also called for efforts to establish governance

standards that foster public trust, and for a focus on building up the nation's "AI workforce." As was noted by many experts at the time, however, the executive order did not designate any new resources for these purposes.

A few months later, in June, the NSTC updated its 2016 strategic plan. Juntament amb les set recomanacions estratègiques de la versió original, enumerades anteriorment, va afegir una vuitena estratègia:

- Expand public-private partnerships to accelerate advances in AI.

Shortly thereafter, in August, the Computing Community Consortium (CCC) and the Association for the Advancement of AI (AAAI), published "A 20-Year Community Roadmap for Artificial Intelligence Research in the US". Co-chaired by Prof. Yolanda Gil from the University of Southern California and Prof. Bart Selman from Cornell University (the AAAI President and President Elect respectively), the roadmap's first recommendation was to create and operate a national AI infrastructure.

To this end, on October 8th, 2019, roughly a week before the initial writing of this article, the National Science Foundation (NSF), in partnership with four other agencies (the National Institute of Food and Agriculture, the Department of Homeland Security, the Department of Transportation, and the Department of Veterans Affairs), issued a call for National AI Research Institutes. The call projected funding of one to six institutes for up to $20 million each on the following designated topics:

- Trustworthy AI
- Foundations of Machine Learning
- AI-Driven Innovation in Agriculture and the Food System
- AI-Augmented Learning
- AI for Accelerating Molecular Synthesis and Manufacturing
- AI for Discovery in Physics.

While it is not known to the author whether the funding behind this call is from newly designated resources or the agencies' existing budgets, the call seems to be being taken as a very positive step by the American AI research community.

Update as of February 2020

As anticipated in the article, since it was written in October 2019, the U.S. AI policy landscape has continued to take concrete shape. Most significantly, in January of 2020, the government published a Memorandum on "Guidance for Regulation of Artificial Intelligence Applications" that provides specific guidance to government agencies on how to balance the risks and benefits of regulation. It specifically calls for regulation only as needed to maintain public trust in AI applications, while keeping in mind the need to foster high-paced innovation. It is another very encouraging step in the right direction.

REFERENCES

1 — Full disclosure: my start-up company, Cogitai, is a member.
2 — Full disclosure: I was Chair of the first Study Panel, and am currently Chair of the Standing Committee.



Peter Stone

Peter Stone is the founder and director of the *Learning Agents Research Group* (LARG) that is part of the *Artificial Intelligence Laboratory* in the Department of Information Sciences at Austin University, Texas. He is also associate department chair of Texas Robotics. He is currently the Executive Director of Sony AI America, and he is also an associate member of the *American Association of Artificial Intelligence* (AAAI). He us also the president of the *Stanford University AI100 Standing Committee*, which has published the *One Hundred Year Study on AI.* His research focuses on understanding how we can best create complete intelligent agents. More specifically, he is interested in *machine learning*, multiagent systems and robotics_ robot soccer, autonomous bidding agents, autonomous vehicles and human-interactive agents. He was awarded wit the CAREER prize from the *National Science Foundation* (2003), the *Fullbright Award* (2008) and the *IJCAI Computers and Thought Award* (2007).