

AI RFI Responses, October 26, 2018

Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan RFI Responses

DISCLAIMER: The [RFI public responses](#) received and posted do not represent the views and/or opinions of the U.S. Government, National Science and Technology Council (NSTC) Select Committee on Artificial Intelligence (AI), NSTC Subcommittee on Machine Learning and AI, NSTC Subcommittee on Networking and Information Technology Research and Development (NITRD), NITRD National Coordination Office, and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality or content of all external links included in this document.



RFI Response: National AI Research and Development Plan

Networking and Information Technology Research and Development (NITRD);
National Coordination Office (NCO); National Science Foundation

ATTN: Faisal D'Souza, NCO, 2415 Eisenhower Ave., Alexandria, VA 22314
USA

CONTACT

Trae Stephens
Chairman of the Board

Babak Siavoshy
Legal Counsel

(949) 891-1607

contact@anduril.com

ANDURIL INDUSTRIES

1340 Reynolds Ave

Suite 116

Irvine, CA 92614

www.anduril.com

Anduril Industries’s Comment on the National AI Strategic Plan: Leveraging private industry for the national interest

Company Overview

Founded in June 2017, Anduril Industries is a venture-capital backed startup headquartered in Orange County, California focused on combining the latest in applied artificial intelligence and autonomous systems with commercial-off-the-shelf sensor technology to solve complex national security challenges for America and its allies. Anduril is comprised of an elite team of industry experts in optics, sensor fusion, computer vision, integrated systems, and data analysis, including senior engineers from Oculus VR (acquired by Facebook in 2014), Palantir, General Atomics, SpaceX, Tesla, and Google.

Anduril is building the next generation of technology that will aid and protect those who serve on the front lines defending the nation and its interests. Our mission is to develop cutting edge technology that enables America and its allies to maintain global leadership now and into the future. After spending the last two years learning about the defense technology ecosystem and identifying technology gaps we can work to address, we have become increasingly concerned that the US is falling behind in the global competition for AI leadership. We believe a strong AI strategic plan is critical to maintaining the high ground that will be required to preserve Western values in a world in which our adversaries have taken a strong interest in out-investing and out-developing us in key technology areas, including AI.

Introduction

The AI Strategic Plan was developed to maintain the US’s technological leadership in the field of AI by outlining a “clear set of R&D priorities” to ensure that federal “investments can achieve their full potential.” The Plan focuses on “areas of strong societal importance that are not aimed at consumer markets—areas such as AI for public health, urban systems and smart communities, social welfare, criminal justice, environmental sustainability, and national security, as well as long-term research that accelerates the production of AI knowledge and technologies.”

While we agree with the Plan’s laudable goals, we believe the current draft falls short of reaching them. The Plan outlines several strategies—centered on investment in Government R&D, standards setting, and policymaking—which, in our view, are insufficient to ensure the US’s continued technological leadership in the field of AI, especially in light of significant investments made by China and others in the field.

In particular, we believe an effective national AI strategy should envision a more prominent role for the commercial tech sector, which is home to the highest concentration of AI engineering talent in the world. In order to effectively leverage this talent advantage, the US must use its strength in procurement to reward promising businesses working on the next generation of AI technologies with meaningful government contracts to meet national priorities. By focusing, instead, on government R&D, policymaking, and standards setting, the Plan under-leverages the US’s most powerful strategic asset: it’s dynamic commercial technology industry.

Traditional R&D is unlikely to lead to the next breakthrough in AI

While government-funded R&D has been effective in driving innovation in a variety of industries, it has a poor track record in the field of software engineering, which is the focal point for advancements in AI, machine learning, and autonomy.¹ That is in large part because software engineering talent is concentrated in the commercial tech sector, rather than in government labs, academic institutions, and traditional government contractors—the most common beneficiaries of government R&D funding. By focusing on R&D over commercial procurement, the plan rewards these traditional players at the expense of engineering talent-centers which, if history is a guide, are most likely to fuel the next generation of breakthroughs in Artificial Intelligence.

In addition to under-leveraging private industry talent, the Plan makes a number of assumptions about technological progress that, in our experience, are incorrect. First, the Plan presumes that important breakthroughs in AI and autonomy (or at least, those deserving of coordinated government funding) will stem from the creation of new technologies through long development cycles of the type common to large R&D projects. In fact, recent history demonstrates that paradigm-shifting technical breakthroughs in software and electronics often result from elegant integrations of existing technologies, rather than through long-term R&D.

This lesson is reflected in Anduril’s experience building autonomous security solutions, as well as in our founding team’s decades of experience building companies like Oculus VR, Palantir Technologies, and SpaceX. For example, Oculus VR (founded by Anduril co-founder Palmer Luckey, and sold to Facebook in 2014) revolutionized the commercial VR industry utilizing technology that had been in existence for half a decade. The same principle holds for many of the pathmaking technologies of the past two decades, such as the Apple iPhone and Amazon’s AWS, which were the result of improvements to and integration of existing technologies, rather than breakthroughs in fundamental research.

The Plan also takes the wrong lesson from the under-investment thesis—“that some important areas of research are unlikely to receive sufficient investment by industry”—by recommending government R&D as the sole antidote to the problem.² While under-investment in public goods is a serious challenge, the solution is not to further isolate commercial tech companies from public works through the creation of insular Government R&D programs, but rather to create new incentives for industry to participate in such projects.

It is no surprise, then, that the report cited in the Plan in support of the under-investment thesis³ also endorses the use of research funds to incentivize industry: “Granting government research funds directly to companies makes it more feasible for them to invest in high-risk, high-payoff

¹ See <https://www.politico.com/story/2011/06/computer-bugs-hurt-army-ops-058051>. To be sure, government R&D was instrumental to the development of the Internet and modern computing, particularly in the post-cold war period. Its track record over the past thirty year in those fields has been less impressive, however.

² See AI Strategic Plan, p. 6.

³ See AI Strategic Plan, p. 17.

innovations. Industry is competitive, not cooperative, and government funding can encourage otherwise risky development that can lead to economic growth.”⁴ Leveraging industry for the public good is all the more important in the field of AI, where (as already noted) engineering talent is disproportionately concentrated in the commercial sector.

Using commercial procurement to accelerate breakthroughs in AI

In order to effectively leverage its massive talent-advantage in AI, the US must use its strength in procurement to reward promising commercial AI companies working in the national interest. A well-planned procurement strategy would, in our view, do more to spur breakthroughs in AI than any other R&D funding initiatives to date.

This approach—focused on directing funding to proven technology talent-centers—is familiar to Venture Capital (VC) firms, which routinely use concentrated investment to jump-start and accelerate technological breakthroughs in emerging technologies. What VCs do through direct investment in startups, the US can and should do through government procurement. As the largest purchaser of goods and services in the world, the US has the unmatched ability to spur the development of AI in the national interest by creating meaningful incentives for businesses working on the next generation of AI to work with the Government.

This approach, it is worth noting, has also been adopted by other nations as part of their respective national AI strategies.⁵ China’s AI Strategic Plan has perhaps been the most ambitious in this regard, including a “multibillion-dollar national investment initiative to support “moonshot” projects, start-ups and academic research in A.I.”⁶ China’s Sense Time, which is now the world’s most valuable non-public AI company, has the Chinese government both as one of its largest investors and one of its largest customers.⁷ China has also entered into public-private partnerships with Baidu, Tencent, and Alibaba, three of the largest software companies in the country.⁸ The Chinese government has put its stamp on these initiatives, reportedly leveraging its work with the commercial AI industry to create new—and troubling—tools for social control.⁹

⁴ See *Continuing Innovation in Information Technology* (Washington D.C.: The National Academies Press, 2012).

⁵ See <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd> for a running list of other national AI strategic plans.

⁶ See <https://www.nytimes.com/2017/07/20/business/china-artificial-intelligence.html>.

⁷ See

<https://www.reuters.com/article/us-sensetime-funding/chinas-sensetime-raises-620-million-its-second-funding-round-in-two-months-idUSKCN1IW07I>.

⁸ See

<https://www.scmp.com/tech/china-tech/article/2120913/china-recruits-baidu-alibaba-and-tencent-ai-national-team>.

⁹ See <https://www.forbes.com/sites/arthurherman/2018/08/30/chinas-brave-new-world-of-ai/#10a01bfb28e9>.

For the United States to retain its leadership in AI and set ethical norms on its use and development, it must first hold the technological high ground.¹⁰ The US can achieve that goal by leveraging its significant procurement power to jump-start and accelerate the commercial development of AI in the national interest.

To be sure, the US has recently engaged in a number of funding efforts targeting the commercial AI sector.¹¹ Those efforts, however, have fallen short of drawing meaningful interest from private industry. While the overall funding commitments are large—\$2 billion dollars from DARPA, for example—those commitments have resulted in few, if any, contracts for private companies other than traditional defense contractors. They have therefore failed to create significant incentives for the commercial tech sector to invest in government applications of AI. At this juncture, the US has arguably spent more time and money debating the future of AI than in building it.

Accelerating innovation through concentrated investment: Lessons from Venture Capital

The US's current AI funding approach is in need of an overarching strategy—one that, in our view, must leverage the dynamic talent base in the US technology sector to be successful. We believe several strategies used by Venture Capital firms to accelerate innovation in emerging technologies may also be relevant to federal investment efforts in AI. VCs use three strategies to spur breakthroughs in emerging technologies, which are instructive here: First, they make big bets, instead of spreading funds across many small investments. Second, they target talent centers, rewarding nimble and dynamic engineering teams rather than traditional players with fewer incentives to innovate. Third, and finally, they anchor their investment strategy to an overarching thesis, rather than to an overspecified vision of a specific solution, product, or service.

The Government should incorporate elements of these strategies into its AI Strategic Plan. First, and most importantly, it should award meaningful, high dollar-value contracts to talented businesses innovating in the field of AI. The Government's current AI procurement approach consists in awarding many small contracts, rather than a small number of large ones. Indeed, a cursory search through recent awards data and news releases detailing unclassified contracts reveals just a few awards over \$5m specifically targeting on AI, autonomy, and machine learning.¹² Experience in the private sector has shown that this “spray and pray” approach is unlikely to lead to technological breakthroughs.

Next, Government procurement in AI should target engineering talent centers, rather than merely focusing on traditional government contractors and government-funded labs. While the US

¹⁰ See

https://www.washingtonpost.com/opinions/silicon-valley-should-stop-ostracizing-the-military/2018/08/08/7a7e0658-974f-11e8-80e1-00e80e1fdf43_story.html?utm_term=.aa6087e16277.

¹¹ See

https://www.washingtonpost.com/technology/2018/09/07/defense-department-pledges-billions-toward-artificial-intelligence-research/?noredirect=on&utm_term=.3abaa0f9f76f.

¹² Based on keyword searches through USASpending.gov.

commercial software industry is the largest and deepest AI talent base in the world, in our search the largest government contracts relating to AI were awarded to traditional defense contractors and systems integrators,¹³ rather than commercial companies. Only by aggressively pursuing its engineering talent-centers can the US achieve its goal of maintaining its technological leadership in AI.

Finally, government procurement in AI should be guided by an overarching thesis, but avoid overspecified and requirements-driven funding vehicles that are common in typical Government R&D programs. There has been a great deal of progress on this front, with federal agencies using innovative contracting vehicles--such as OT agreements and commercial-item contracting--to encourage greater participation by industry in government programs. Whatever the specific mechanism, procurement efforts should encourage, not impede, the type of product-focused development and innovation that has powered the commercial technology market for the 30 years.

¹³ See <https://www.fedscoop.com/booz-allen-lands-885m-pentagon-ai-contract/>, detailing the Pentagon's \$885m deal with Booz Hamilton to help the DoD rapidly deploy AI, neural, and deep networks across its infrastructure.